


Population Games and Evolutionary Dynamics

William H. Sandholm

March 21, 2010



Chaos under the replicator dynamic

CONTENTS IN BRIEF

Preface	xv
1 Introduction	1
I Population Games	19
2 Population Games	21
3 Potential Games, Stable Games, and Supermodular Games	53
II Deterministic Evolutionary Dynamics	123
4 Revision Protocols and Deterministic Evolutionary Dynamics	125
5 Deterministic Dynamics: Families and Properties	147
6 Best Response and Projection Dynamics	187
III Convergence and Nonconvergence of Deterministic Dynamics	233
7 Global Convergence of Evolutionary Dynamics	235
8 Local Stability under Evolutionary Dynamics	289
9 Nonconvergence of Evolutionary Dynamics	339

IV Stochastic Evolutionary Models	387
10 Stochastic Evolution and Deterministic Approximation	389
11 Stationary Distributions and Infinite Horizon Behavior	421
12 Limiting Stationary Distributions and Stochastic Stability	477
Bibliography	573

CONTENTS

Preface	xv
1 Introduction	1
1.1 Population Games	2
1.1.1 Modeling Interactions in Large Populations	2
1.1.2 Definitions and Classes of Population Games	5
1.2 Evolutionary Dynamics	6
1.2.1 Knowledge, Rationality, and Large Games	7
1.2.2 Foundations for Evolutionary Dynamics	8
1.2.3 Deterministic Evolutionary Dynamics	10
1.2.4 Orders of Limits for Stochastic Evolutionary Models	12
1.2.5 Stationary Distributions and Stochastic Stability	12
1.3 Remarks on History, Motivation, and Interpretation	13
1.N Notes	16
I Population Games	19
2 Population Games	21
2.0 Introduction	21
2.1 Population Games	22
2.1.1 Populations, Strategies, and States	22
2.1.2 Payoffs	23
2.1.3 Best Responses and Nash Equilibria	24
2.1.4 Prelude to Evolutionary Dynamics	25
2.2 Examples	25
2.2.1 Matching in Normal Form Games	25

2.2.2	Congestion Games	27
2.2.3	Two Simple Externality Models	29
2.3	The Geometry of Population Games and Nash Equilibria	29
2.3.1	Drawing Two-Strategy Games	30
2.3.2	Displacement Vectors and Tangent Spaces	31
2.3.3	Orthogonal Projections	34
2.3.4	Drawing Three-Strategy Games	36
2.3.5	Tangent Cones and Normal Cones	38
2.3.6	Normal Cones and Nash Equilibria	41
2.A	Affine Spaces, Tangent Spaces, and Projections	44
2.A.1	Affine Spaces	44
2.A.2	Affine Hulls of Convex Sets	45
2.A.3	Orthogonal Projections	46
2.A.4	The Moreau Decomposition Theorem	50
2.N	Notes	50
3	Potential Games, Stable Games, and Supermodular Games	53
3.0	Introduction	53
3.1	Full Potential Games	54
3.1.1	Full Population Games	54
3.1.2	Definition and Characterization	55
3.1.3	Examples	56
3.1.4	Nash Equilibria of Full Potential Games	59
3.1.5	The Geometry of Nash Equilibrium in Full Potential Games	65
3.1.6	Efficiency in Homogeneous Full Potential Games	67
3.1.7	Inefficiency Bounds for Congestion Games	69
3.2	Potential Games	72
3.2.1	Motivating Examples	73
3.2.2	Definition, Characterizations, and Examples	74
3.2.3	Potential Games and Full Potential Games	77
3.2.4	Passive Games and Constant Games	79
3.3	Stable Games	81
3.3.1	Definition	81
3.3.2	Examples	83
3.3.3	Invasion	87
3.3.4	Global Neutral Stability and Global Evolutionary Stability	88
3.3.5	Nash Equilibrium and Global Neutral Stability in Stable Games	92

3.4	Supermodular Games	96
3.4.1	Definition	96
3.4.2	Examples	99
3.4.3	Best Response Monotonicity in Supermodular Games	101
3.4.4	Nash Equilibria of Supermodular Games	102
3.A	Multivariate Calculus	103
3.A.1	Univariate Calculus	103
3.A.2	The Derivative as a Linear Map	103
3.A.3	Differentiation as a Linear Operation	106
3.A.4	The Product Rule and the Chain Rule	106
3.A.5	Homogeneity and Euler’s Theorem	108
3.A.6	Higher Order Derivatives	109
3.A.7	The Whitney Extension Theorem	110
3.A.8	Vector Integration and the Fundamental Theorem of Calculus	111
3.A.9	Potential Functions and Integrability	111
3.B	Affine Calculus	112
3.B.1	Linear Forms and the Riesz Representation Theorem	112
3.B.2	Dual Characterizations of Multiples of Linear Forms	113
3.B.3	Derivatives of Functions on Affine Spaces	115
3.B.4	Affine Integrability	117
3.N	Notes	119

II Deterministic Evolutionary Dynamics 123

4	Revision Protocols and Deterministic Evolutionary Dynamics	125
4.0	Introduction	125
4.1	The Basic Stochastic Evolutionary Model	126
4.1.1	Inertia and Myopia	126
4.1.2	Revision Protocols	127
4.2	Mean Dynamics	129
4.2.1	Derivation	129
4.2.2	Target Protocols and Target Dynamics	130
4.3	Examples	132
4.3.1	Imitative Protocols and Dynamics	133
4.3.2	Direct Protocols and Dynamics	135
4.4	Deterministic Evolutionary Dynamics	136

4.A	Ordinary Differential Equations	138
4.A.1	Basic Definitions	138
4.A.2	Existence, Uniqueness, and Continuity of Solutions	140
4.A.3	Ordinary Differential Equations on Compact Convex Sets	142
4.N	Notes	144
5	Deterministic Dynamics: Families and Properties	147
5.0	Introduction	147
5.1	Information Requirements for Revision Protocols	149
5.2	Incentives and Aggregate Behavior	152
5.3	Families of Evolutionary Dynamics	156
5.4	Imitative Dynamics	161
5.4.1	Definition	161
5.4.2	Examples	162
5.4.3	Biological Derivations of the Replicator Dynamic	168
5.4.4	Extinction and Invariance	169
5.4.5	Monotone Percentage Growth Rates and Positive Correlation	170
5.4.6	Rest Points and Restricted Equilibria	173
5.5	Excess Payoff Dynamics	174
5.5.1	Definition and Interpretation	174
5.5.2	Incentives and Aggregate Behavior	176
5.6	Pairwise Comparison Dynamics	178
5.6.1	Definition	178
5.6.2	Incentives and Aggregate Behavior	179
5.6.3	Desiderata Revisited	181
5.7	Multiple Revision Protocols and Combined Dynamics	182
5.N	Notes	183
6	Best Response and Projection Dynamics	187
6.0	Introduction	187
6.1	The Best Response Dynamic	188
6.1.1	Definition and Examples	188
6.1.2	Construction and Properties of Solution Trajectories	190
6.1.3	Incentive Properties	197
6.2	Perturbed Best Response Dynamics	198
6.2.1	Revision Protocols and Mean Dynamics	198
6.2.2	Perturbed Optimization: A Representation Theorem	200

6.2.3	Logit Choice and the Logit Dynamic	202
6.2.4	Perturbed Incentive Properties via Virtual Payoffs	207
6.3	The Projection Dynamic	210
6.3.1	Definition	210
6.3.2	Solution Trajectories	211
6.3.3	Incentive Properties	213
6.3.4	Revision Protocols and Connections with the Replicator Dynamic . .	214
6.A	Differential Inclusions	216
6.A.1	Basic Theory	216
6.A.2	Differential Equations Defined by Projections	219
6.B	The Legendre Transform	220
6.B.1	Legendre Transforms of Functions on Open Intervals	220
6.B.2	Legendre Transforms of Functions on Multidimensional Domains . .	223
6.C	Perturbed Optimization	225
6.C.1	Proof of the Representation Theorem	225
6.C.2	Additional Results	227
6.N	Notes	229

III Convergence and Nonconvergence of Deterministic Dynamics 233

7	Global Convergence of Evolutionary Dynamics	235
7.0	Introduction	235
7.1	Potential Games	237
7.1.1	Potential Functions as Lyapunov Functions	237
7.1.2	Gradient Systems for Potential Games	243
7.2	Stable Games	247
7.2.1	The Projection and Replicator Dynamics in Strictly Stable Games . .	247
7.2.2	Integrable Target Dynamics	250
7.2.3	Impartial Pairwise Comparison Dynamics	261
7.2.4	Summary	263
7.3	Supermodular Games	264
7.3.1	The Best Response Dynamic in Two-Player Normal Form Games . .	266
7.3.2	Stochastically Perturbed Best Response Dynamics	269
7.4	Dominance Solvable Games	273
7.4.1	Dominated and Iteratively Dominated Strategies	274
7.4.2	The Best Response Dynamic	274

7.4.3	Imitative Dynamics	275
7.A	Limit and Stability Notions for Deterministic Dynamics	277
7.A.1	ω -Limits and Notions of Recurrence	277
7.A.2	Stability of Sets of States	278
7.B	Stability Analysis via Lyapunov Functions	279
7.B.1	Lyapunov Stable Sets	279
7.B.2	ω -Limits and Attracting Sets	280
7.B.3	Asymptotically Stable and Globally Asymptotically Stable Sets . . .	282
7.C	Cooperative Differential Equations	282
7.N	Notes	285
8	Local Stability under Evolutionary Dynamics	289
8.0	Introduction	289
8.1	Non-Nash Rest Points of Imitative Dynamics	290
8.2	Local Stability in Potential Games	291
8.3	Evolutionarily Stable States	293
8.3.1	Single-Population Games	294
8.3.2	Multipopulation Games	299
8.3.3	Regular Taylor ESS	300
8.4	Local Stability via Lyapunov Functions	301
8.4.1	The Replicator and Projection Dynamics	301
8.4.2	Target and Pairwise Comparison Dynamics: Interior ESS	302
8.4.3	Target and Pairwise Comparison Dynamics: Boundary ESS	304
8.5	Linearization of Imitative Dynamics	309
8.5.1	The Replicator Dynamic	310
8.5.2	General Imitative Dynamics	314
8.6	Linearization of Perturbed Best Response Dynamics	316
8.6.1	Deterministically Perturbed Best Response Dynamics	317
8.6.2	The Logit Dynamic	318
8.A	Matrix Analysis	319
8.A.1	Rank and Invertibility	319
8.A.2	Eigenvectors and Eigenvalues	320
8.A.3	Similarity, (Block) Diagonalization, and the Spectral Theorem	322
8.A.4	Symmetric Matrices	324
8.A.5	The Real Jordan Canonical Form	325
8.A.6	The Spectral Norm and Singular Values	326
8.A.7	Hines's Lemma	327

8.B	Linear Differential Equations	328
8.B.1	Examples	329
8.B.2	Solutions	331
8.B.3	Stability and Hyperbolicity	332
8.C	Linearization of Nonlinear Differential Equations	334
8.N	Notes	336
9	Nonconvergence of Evolutionary Dynamics	339
9.0	Introduction	339
9.1	Conservative Properties of Evolutionary Dynamics	340
9.1.1	Constants of Motion in Null Stable Games	340
9.1.2	Preservation of Volume	345
9.2	Games with Nonconvergent Evolutionary Dynamics	347
9.2.1	Circulant Games	348
9.2.2	Continuation of Attractors for Parameterized Games	351
9.2.3	Mismatching Pennies	354
9.2.4	The Hypnodisk Game	358
9.3	Chaotic Evolutionary Dynamics	362
9.4	Survival of Dominated Strategies	365
9.4.1	A General Survival Theorem	368
9.4.2	Examples and Discussion	372
9.A	Three Classical Theorems on Nonconvergent Dynamics	377
9.A.1	Liouville's Theorem	377
9.A.2	The Poincaré-Bendixson and Bendixson-Dulac Theorems	380
9.B	Attractors and Continuation	381
9.B.1	Attractors and Repellors	381
9.B.2	Continuation of Attractors	383
9.N	Notes	384
IV	Stochastic Evolutionary Models	387
10	Stochastic Evolution and Deterministic Approximation	389
10.0	Introduction	389
10.1	The Stochastic Evolutionary Process	390
10.2	Finite Horizon Deterministic Approximation	392
10.2.1	Kurtz's Theorem	392

10.2.2	Deterministic Approximation of the Stochastic Evolutionary Process	395
10.3	Extensions	398
10.3.1	Discrete-Time Models	398
10.3.2	Finite-Population Adjustments	399
10.A	The Exponential and Poisson Distributions	400
10.A.1	Basic Properties	400
10.A.2	The Poisson Limit Theorem	403
10.B	Probability Models and their Interpretation	405
10.B.1	Countable Probability Models	405
10.B.2	Uncountable Probability Models and Measure Theory	406
10.B.3	Distributional Properties and Sample Path Properties	408
10.C	Countable State Markov Chains and Processes	412
10.C.1	Countable State Markov Chains	412
10.C.2	Countable State Markov Processes: Definition and Construction	413
10.C.3	Countable State Markov Processes: Transition Probabilities	416
10.N	Notes	418
11	Stationary Distributions and Infinite Horizon Behavior	421
11.0	Introduction	421
11.1	Irreducible Evolutionary Processes	423
11.1.1	Full Support Revision Protocols	423
11.1.2	Stationary Distributions and Infinite Horizon Behavior	426
11.1.3	Reversibility	427
11.2	Stationary Distributions for Two-Strategy Games	427
11.2.1	Birth and Death Processes	428
11.2.2	The Stationary Distribution of the Evolutionary Process	431
11.2.3	Examples	432
11.3	Waiting Times and Infinite Horizon Prediction	437
11.3.1	Examples	438
11.3.2	Discussion	441
11.4	Model Adjustments for Finite Populations	443
11.4.1	Finite-Population Games	443
11.4.2	Clever Payoff Evaluation	444
11.4.3	Committed Agents and Imitative Protocols	446
11.5	Exponential Protocols and Potential Games	449
11.5.1	Finite-Population Potential Games	449
11.5.2	Exponential Revision Protocols	453

11.5.3	Reversibility and Stationary Distributions	455
11.A	Long Run Behavior of Markov Chains and Processes	459
11.A.1	Communication, Recurrence, and Irreducibility	459
11.A.2	Periodicity	462
11.A.3	Hitting Times and Hitting Probabilities	463
11.A.4	The Perron-Frobenius Theorem	466
11.A.5	Stationary Distributions for Markov Chains	467
11.A.6	Reversible Markov Chains	468
11.A.7	Stationary Distributions and Reversibility for Markov Processes . . .	469
11.A.8	Convergence in Distribution	471
11.A.9	Ergodicity	474
11.N	Notes	475
12	Limiting Stationary Distributions and Stochastic Stability	477
12.0	Introduction	477
12.1	Definitions of Stochastic Stability	480
12.1.1	Small Noise Limits	480
12.1.2	Large Population Limits	482
12.1.3	Double Limits	484
12.1.4	Double Limits: A Counterexample	486
12.2	Exponential Protocols and Potential Games	490
12.2.1	Direct Exponential Protocols: The Small Noise Limit	491
12.2.2	Direct Exponential Protocols: The Large Population Limit	493
12.2.3	Direct Exponential Protocols: Double Limits	499
12.2.4	Imitative Exponential Protocols with Committed Agents	500
12.3	Noisy Best Response Protocols in Two-Strategy Games	502
12.3.1	Noisy Best Response Protocols and their Cost Functions	502
12.3.2	The Small Noise Limit	505
12.3.3	The Large Population Limit	507
12.3.4	Double Limits	508
12.3.5	Stochastic Stability: Examples	510
12.3.6	Risk Dominance, Stochastic Dominance, and Stochastic Stability . .	515
12.4	Imitative Protocols in Two-Strategy Games	519
12.4.1	Imitative Protocols with Mutations	520
12.4.2	Imitative Protocols with Committed Agents	523
12.4.3	Imitative Protocols, Mean Dynamics, and Stochastic Stability	524
12.5	Small Noise Limits	528

12.5.1	Noisy Best Response Protocols and Cost Functions	528
12.5.2	Limiting Stationary Distributions via Trees	531
12.5.3	Two-Strategy Games and Risk Dominance	533
12.5.4	The Radius-Coradius Theorem	539
12.5.5	Half Dominance	541
12.6	Large Population Limits	544
12.6.1	Convergence to Recurrent States of the Mean Dynamic	545
12.6.2	Convergence to Stable Rest Points of the Mean Dynamic	547
12.A	Trees, Escape Paths, and Stochastic Stability	549
12.A.1	The Markov Chain Tree Theorem	550
12.A.2	Limiting Stationary Distributions via Trees	551
12.A.3	Limiting Stationary Distributions via Trees on Recurrent Classes	554
12.A.4	Radius-Coradius Theorems	557
12.A.5	Lenient Transition Costs and Weak Stochastic Stability	559
12.B	Stochastic Approximation Theory	562
12.B.1	Convergence to the Birkhoff Center	562
12.B.2	Sufficient Conditions for Convergence to Stable Rest Points	564
12.N	Notes	568

Bibliography		573
---------------------	--	------------

Preface

Why I wrote this book

This book is my attempt to provide a systematic, rigorous, and unified presentation of evolutionary game theory.

Evolutionary game theory concerns the behavior of large populations of strategically interacting agents who occasionally reevaluate their choices in light of current payoff opportunities. (This is an economist's one-sentence description; a biologist's might refer instead to populations of interacting animals whose different genetic programs lead to varying levels of reproductive success—see Section 1.3.) To the outside observer, it might seem that the theory offers a collection of loosely related models, with the deterministic models (e.g., the replicator dynamic) and the stochastic models (e.g., stochastic stability theory) being nearly completely isolated from one another. This point of view is reinforced by the fact that basic techniques used in each case—deterministic dynamical systems and Markov processes—are traditionally viewed as belonging to different branches of mathematics, despite their obvious complementarity from an applied point of view. One might notice further that many (though certainly not all) researchers in evolutionary game theory work either on deterministic models or on stochastic models, but not on both.

With these thoughts in mind, I aimed to write book that presents evolutionary game theory as a coherent whole. I do so by deriving models of aggregate behavior from explicit microfoundations. By starting with descriptions of how individual agents update their strategies over time, one can derive both the deterministic and the stochastic parts of the theory within a single framework, with the two analyses arising by way of different assumptions about population sizes and the time horizon of interest. The idea that microfoundations can provide a unified presentation of evolutionary game theory has been articulated a number of researchers over the past decade, but this book may be the first to offer this point of view to a broad audience.

Over the years, I have had many conversations with other economists end with ex-

pressions of regret that gaps in their mathematical background preclude their becoming acquainted with evolutionary game theory in more than a superficial way. While I realize that you were all just being nice, it is still the case that the standard graduate training of economists and game theorists includes very little about ordinary differential equations and Markov processes. To address this gap, and to reduce costs of entry, I supplement the main text with appendices that present the mathematical tools needed to work (or even to read deeply) in evolutionary game theory. These appendices, which account for about one quarter of the book, are written so as to be readable independently of the main text, and may be useful for anyone seeking a quick introduction to the methods of dynamic modeling.

A number of books providing overviews of evolutionary game theory appeared between 1988 and 1998, and there have been important surveys and monographs since that time (see the Notes to Chapter 1). But it has been more than ten years since a picture of the field in full has been attempted. New ideas introduced during this interval have made it easier to view evolutionary game theory as a unified whole, and make it natural to present classic results alongside the most recent developments.

What is in this book?

This book contains twelve chapters. The first is an introductory chapter devoted to stage-setting, and the remaining eleven chapters are each assigned to one of four parts:

Part I: Population games (Chapters 2–3)

Part II: Deterministic evolutionary dynamics (Chapters 4–6)

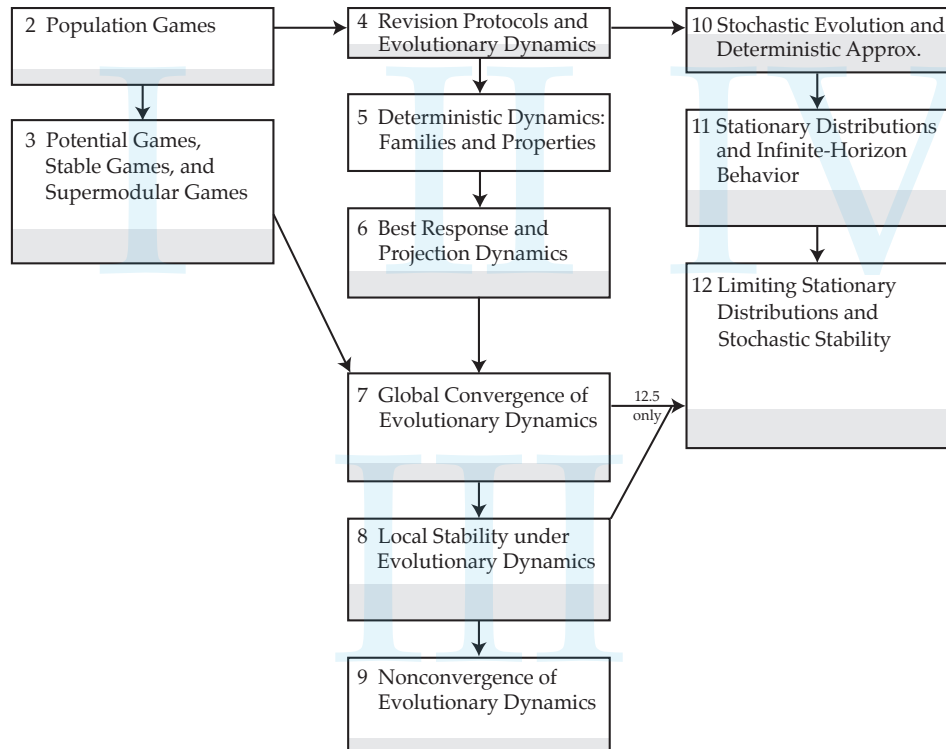
Part III: Convergence and nonconvergence of deterministic dynamics (Chapters 7–9)

Part IV: Stochastic evolutionary models (Chapters 10–12)

The chart below describes the logical flow of the chapters. The size of each chapter's box corresponds to the chapter's length, and the size of the box's shaded area to the length of the chapter's math appendices. A detailed overview of the contents of the book is provided in Chapter 1, but I offer an exceedingly brief one here.

Part I introduces population games, which offer a simple, general model for studying strategic interactions among large numbers of agents. Basic definitions and examples are presented in Chapter 2. Chapter 3 studies three classes of especially well-behaved games: potential games, stable games, and supermodular games.

Part II introduces deterministic evolutionary dynamics. Basic definitions and examples are presented in Chapter 4. The notion of a revision protocol, which underlies the



microfoundations uniting all of the subsequent analyses, first appears here. Chapters 5 and 6 introduce a variety of families of deterministic dynamics and the revision protocols that generate them, and investigate their basic properties.

Part III studies convergence and nonconvergence under deterministic evolutionary dynamics. Chapter 7 establishes global convergence results, mainly for combinations of the classes of games from Chapter 2 and the families of dynamics from Chapters 5 and 6. Chapter 8 considers local stability of rest points, and connects local stability with the classical notion of ESS. Chapter 9 considers examples of cycling and chaotic behavior, as well as the implications of these examples for predicting behavior in games.

Part IV considers models of stochastic evolution. Chapter 10 describes the theory of finite-horizon deterministic approximation, providing a rigorous foundation for the deterministic dynamics studied in Parts II and III. Chapters 11 and 12 consider infinite-horizon behavior of stochastic evolutionary processes, with analyses of stationary distributions, limiting stationary distributions, and stochastic stability.

It is possible to cover all of the chapters of the book in a single semester, provided that not every topic is covered in detail, although this path is not for the faint of heart. One alternative is a course that focuses on deterministic dynamics, which emphasizes Chapters 2–9. Another is a course focusing on stochastic evolution, which would start

with Chapters 2, 4, and 10–12, leaving time to return to some topics that initially were passed over.

The reader will discover that many of the analyses and discussions in this book have a decidedly geometric flavor. Each chapter contains a large number of figures, which I believe are necessary to develop a real feeling for many of the topics I cover.

The majority of the phase diagrams for deterministic evolutionary dynamics were created using the Dynamo software suite (Sandholm and Dokumacı (2007)). This software, which runs within Mathematica (for now, versions 4.0–5.2), automates the process of drawing phase diagrams for evolutionary dynamics. Dynamo is free and open-source, and is available for download from my website. I find the software very helpful for running quick numerical experiments, and I encourage you to try it out as you work through the text. Please feel welcome to use Dynamo to create publishable figures for your own research.

What is not in this book?

While this book provides a unified treatment of a large swath of evolutionary game theory, there are many important topics that I have omitted entirely. Among them are

- evolution in extensive form games (deterministic models, stochastic models)
- evolution in local interaction games (on lattices, on general graphs)
- evolution in games with continuous strategy sets
- stochastic evolution via diffusion processes

These are important topics, but each of them requires new ideas and techniques, and each could easily fill one or more additional chapters. I omitted them only to leave myself some hope of finishing the book in finite time.

An area of research with very close connections to evolutionary game theory is heuristic learning. Like models in evolutionary game theory, models of heuristic learning, starting with the classical model of fictitious play, consider the repeated play of a game by players who use simple rules to decide how to act. But these models consider small groups of players—typically, one for each role in a normal form game—and the rules the players employ condition on statistics summarizing the time-average of past play, rather than on the current play of a large population. The two approaches to disequilibrium game dynamics overlap, with a number of results having formulations in both contexts, but this book maintains the evolutionary point of view.

A basic test of the worth of any theory is its ability to offer insights into applications. Except for certain fundamental examples (for instance, on flows and routing in transportation networks), this book says little about applied topics. Nevertheless, a wide range

of applications of evolutionary game theory have been developed in economic and other social science contexts, including

- markets
- bargaining and hold-up problems
- signaling and cheap talk
- public good provision
- implementation and decentralized control
- residential segregation
- preference evolution and cultural evolution

This list does not mention applications in biology, where evolutionary game theory originated, or more recent applications in engineering, control theory, and computer science. References to applications in all of these areas, as well as to the theoretical topics mentioned above, can be found in the Notes to Chapter 1.

What about the math?

This book is written at a level appropriate for second-year economics graduate students and senior undergraduate mathematics majors. All but the most basic mathematical tools are reviewed in the appendices, which address the following topics:

Calculus: 3.A

Convex analysis: 2.A, 3.B, 6.B, 6.C

Matrix analysis: 8.A

Dynamical systems: 4.A, 6.A, 7.A, 7.B, 7.C, 8.B, 8.C, 9.A, 9.B

Probability theory and stochastic processes: 10.A, 10.B, 10.C, 11.A, 12.A, 12.B

While Appendix 3.A provides a review of the necessary tools from multivariate calculus, I expect most readers to have the background provided by an undergraduate real analysis sequence. Since most of the action takes place on the simplex or on products of simplices, certain basic tools from convex analysis turn out to be quite useful, and are explained as needed. The review of matrix analysis contains what is necessary to understand thoroughly how to linearize a differential equation, and so summarizes a number of results that might be left out of an undergraduate linear algebra course.

The treatments of dynamical systems and stochastic processes assume no prior knowledge in these areas. In both cases I start with basic material and proceed to topics that typically would not be covered until graduate-level courses. A few of the topics treated here (for instance, differential inclusions (Appendix 6.A), attractors and continuation (Appendix 9.B), stochastic stability analysis (Appendix 12.A), and stochastic approximation

theory (Appendix 12.B)) are not treated in typical courses at any level, but will be very important here.

The appendices are written so that they can be read independently of the main text. A reader with limited exposure to, say, ordinary differential equations might be better off reading the relevant appendices before attempting the chapter texts; others might peek at the appendices as necessary. In general, the appendices introduce definitions and notation, state results, and provide intuitions and examples. Results are only proved when the proof aids intuition, or when I could not find a proof in the literature.

While I think these overviews are useful, they are pale substitutes for full treatments of their subject matter. For readers looking for mathematics references, I offer a few recommendations here. Lang (1997) provides a clear treatment of calculus and real analysis at about the level of the present book. Hiriart-Urruty and Lemaréchal (2001) is a wonderful reference on convex analysis, and Horn and Johnson (1985) is an excellent book on matrix analysis. Hirsch and Smale (1974) offer a solid undergraduate-level course on dynamical systems, while Robinson (1995) provides a fine graduate-level treatment. My favorite undergraduate-level introduction to Markov chains and processes is the concise account of Norris (1997). Many additional references on a variety of specialized topics are provided in the notes to the chapters where they become relevant.

Thanks

The beginning of my collaboration with Josef Hofbauer in the spring of 2000 was the decisive moment of my research career, and I will always be in his debt.

I had the idea to write this book during a series of lectures that Josef and I gave at the University of Tokyo in June of 2002. Portions of the book were written in the collegial atmosphere of the Stockholm School of Economics during May 2005 and May 2008. I thank both institutions for their warm hospitality. I also express my gratitude to the National Science Foundation for supporting my research since 2001.

I was extremely fortunate to have Larry Samuelson as a colleague at Wisconsin from 1998 through 2007; his influence pervades my work. Jörgen Weibull first invited me to speak in Stockholm in the fall of 2000, and has been a constant source of encouragement and inspiration ever since.

Emin Dokumacı is the main programmer of the Dynamo project, and he created many of the figures in this book; our collaboration is as close as I am likely to come to producing visual art. Joint work with Michel Benaïm, Jeff Ely, and Ratul Lahkar shaped my thinking on various topics presented here. Ignacio Monzón and Danqing Hu gamely undertook the laborious task of file conversion in 2006, when I became the last person on the planet

to switch from Microsoft Word to LaTeX. Katsuhiko Aiba, Fernando Louge, and Ryoji Sawa each read large portions of the manuscript, and their suggestions and corrections have made for a clearer and cleaner presentation.

Many other people provided valuable feedback on this book and on related research projects. I thank Ken Binmore, Ross Cressman, Eddie Dekel, Drew Fudenberg, Ed Hopkins, Bart Lipman, Aki Matsui, Nolan Miller, Jeroen Swinkels, Peyton Young, and students in many Economics 806 classes at the University of Wisconsin and in a minicourse at the Stockholm School of Economics for stimulating conversations and helpful comments. Many anonymous reviewers provided useful suggestions for improvements, and my editors at MIT Press, John Covell and Elizabeth Murry, were patient guides during the seemingly endless process of creating this book.

I dedicate this book to my wife Leah and my daughter Lily. In the words of Lina Lamont, you bring a little joy into my humdrum life, and it makes me feel as though my hard work ain't been in vain for nothin'.

Introduction

This book describes an approach to modeling recurring strategic interactions in large populations of small, anonymous agents. The approach is built upon two basic elements. The first, called a *population game*, describes the strategic interaction that is to occur repeatedly. The second, called a *revision protocol*, specifies the myopic procedure that agents employ to decide when and how to choose new strategies. Starting with a population game and a revision protocol, one can derive dynamic processes, both deterministic and stochastic, that describe how the agents' aggregate behavior changes over time. These processes are known as *evolutionary game dynamics*.

This introductory chapter begins the work of adding substance to this austere account of evolutionary game theory, providing motivations for and overviews of the analyses to come. For the most part, the chapter is written with an eye toward modeling in economics and other social sciences, though it also discusses the biological origins of the field. But these perspectives should not be viewed as constraints, as the methods presented in this book have ready applications in other disciplines that require models of interacting populations of humans, animals, or machines.

Section 1.1 introduces the notion of a population game by presenting applications, offering informal definitions, and discussing connections with normal form games. It then previews our treatment of population games in Chapters 2–3. Section 1.2 describes dynamic models of behavior in recurrent play of population games, and contrasts this dynamic approach with the equilibrium approach traditionally used in game theory. This section also offers an overview of our presentation of evolutionary dynamics in Chapters 4–12. Section 1.3 concludes with some remarks on motivations for and interpretations of evolutionary game theory. References relevant to the discussions in the text can be found in the Notes at the end of the chapter, where the references from the Preface can also be found.

1.1 Population Games

[O]nly after the theory for moderate numbers of participants has been satisfactorily developed will it be possible to decide whether extremely great numbers of participants simplify the situation. . . We share the hope. . . that such simplifications will indeed occur. . .

von Neumann and Morgenstern (1944, p. 14)

We shall now take up the “mass-action” interpretation of equilibrium points. . . It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.

To be more detailed, we assume that there is a population (in the sense of statistics) of participants for each position of the game. Let us also assume that the “average playing” of the game involves n participants selected at random from the n populations, and that there is a stable average frequency with which each pure strategy is employed by the “average member” of the appropriate population.

Nash (1950b, p. 21)

There are many situations, however, in which an individual is, in effect, competing not against an individual opponent but against the population as a whole, or some section of it. Such cases can loosely be described as “playing the field”. . . [S]uch contests against the field are probably more widespread and important than pairwise contests.

Maynard Smith (1982, p. 23)

1.1.1 Modeling Interactions in Large Populations

One can imagine many economic, social, and technological environments in which large collections of small agents make strategically interdependent decisions:

- Network congestion. Drivers commute over a highway network. The delay each driver experiences depends not only on the route he selects, but also on the congestion created by other agents along this route.
- Public goods and externalities. A local government maintains a collection of public recreation facilities. The benefit that a family obtains from using a facility depends on the quality of the facility and the number of other families that use it.
- Industrial organization. Software developers choose whether to copyright their products or make them freely available under a public license; they can also choose to work within an existing open-source framework. The latter options entail a loss of control, but allow products to improve through the accumulation of uncoordinated individual efforts.

- The emergence of conventions, norms, and institutions. Firms in a developing economy choose among various business practices: whether to accept credit or require cash, whether to fight or acquiesce to corrupt officials, whether to reward merit or engage in nepotism. Through historical precedent and individual decisions, conventions about business conduct are formed. Whether or not these conventions are efficient, they enable firms to form accurate expectations about how trading partners will act.
- Cultural integration and assimilation. The behavior of immigrants settling in a new country is influenced by traditions imported from their home country, and, to the extent that interactions with the incumbent population require coordination, by the practices of these incumbents as well. At the same time, the choices of the incumbents are influenced by the need to coordinate with the immigrants. The interplay of these forces determines how behavior in the society as a whole evolves.
- Language and communication. Agents without a common language interact repeatedly, attempting to communicate their intentions to the others on each occasion. Whether these attempts are ultimately understood, and whether they enable the agents to coordinate on mutually beneficial behavior, depends on meanings determined by the aggregate communication pattern.
- Task allocation and decentralized control. The employees of a large firm provide services to customers in a number of distinct locations. To take advantage of the ground-level information possessed by the employees, the firm has them allocate themselves among the customers requiring service, providing incentives to ensure that the employees' choices further the firm's objectives.
- Markets and bargaining. Large numbers of buyers and sellers participate in a centralized exchange. Each individual specifies acceptable terms of trade in the hopes of obtaining the greatest benefit from his initial endowment.

While the environments listed above are quite varied, they have certain basic features in common. First, each environment contains a large number of agents capable of making decisions independently. Second, each agent is small, in that his choices have only a minor impact on other agents' outcomes. Third, agents are anonymous: an agent's outcome from the interaction depends on his own strategy and the distribution of others' strategies; further individuation of the opponents is not required.

Simultaneous interactions exhibiting these three properties can be modeled using *population games*. The participants in a population game form a *society* consisting of one or more *populations* of agents. Agents in a given population are assumed to be identical: an agent's population determines his role in the game, the strategies available to him,

and his preferences. These preferences are described by a payoff function that conditions on the agent's own strategy and the distribution of strategies in each population. The populations may have either finite numbers of agents or continua of agents, whichever is more convenient. When populations are continuous, the payoffs to each strategy are assumed to depend on the society's aggregate behavior in a continuous fashion, reflecting the idea that very small changes in aggregate behavior do not lead to large changes in the consequences of playing any given strategy.

Aggregate behavior in a population game is described by a *social state*, which specifies the empirical distribution of strategy choices in each population. For simplicity, we assume throughout this book that there are a finite number of populations, and that members of each population choose from a finite strategy set. Doing so ensures that the social state is finite-dimensional, expressible as a vector with a finite number of components; if populations are continuous the set of social states is a polytope.

Except when they are revising (see the next section), agents in population games are assumed to play pure strategies. One of the main reasons for introducing randomized strategies is moot here: when the populations are continuous, and payoffs are continuous in the social state, pure strategy Nash equilibria always exist. This guarantee may be one of the "simplifications" that von Neumann and Morgenstern had in mind when they looked ahead to the study of games with large numbers of players. But this fact is not as essential to our theory as it is to traditional approaches to analyzing games: while traditional approaches are grounded on the assumption of equilibrium play, we emphasize the process through which agents adjust their behavior in response to their current strategic environment.

From our point of view, a more important simplification provided by the population games framework is the description of behavior using social states—that is, by distributions of agents' choices. To understand the advantages this approach brings, let us contrast it with the standard framework for modeling simultaneous-move interactions: normal form games. To define a normal form game, we must specify a (finite or infinite) set of players, and to describe behavior in such a game, we must stipulate each player's strategy choice. If there are many players, these tasks can be laborious.

One can view a population game as a normal form game that satisfies certain restrictions on the diversity and anonymity of the players. But it is nevertheless preferable to work directly with population games, and so to avoid the extra work involved in individuating the players. Moreover, describing behavior using empirical strategy distributions shifts our attention from questions like "who chose strategy i ?" to ones like "how many chose strategy i ?" and "what happens to the payoff of strategy j players if some agents

switch to strategy i ?" This is just where our attention should be: as we will see in Chapter 3, the answers to the last question determine the incentive structure of a population game.

Much research in evolutionary game theory has focused on population games of a particularly simple sort: those generated when populations of agents are matched to play a normal form game. Indeed, Nash informally introduced population games of this sort in 1949 in proposing the "mass action" interpretation of his equilibrium concept, a development that seems to have gone unnoticed for the next 45 years (see the Notes). This population-matching interpretation of normal form games has a clear appeal, and we will use normal form games as a ready source of simple examples throughout the book.

At the same time, if our main interest is in large population interactions rather than in normal form games, then focusing only on matching in normal form games is quite restrictive. Maynard Smith observes that matching is a rather special sort of interaction in large populations. Instead, interactions in which each agent's payoffs are determined directly from all agents' behavior—what Maynard Smith terms "playing the field"—seem to us to be the rule rather than the exception. Only a few of the applications listed at the onset are most naturally modeled using matching; some, like congestion in highway networks, require payoffs to depend nonlinearly on the population state, and so are mathematically inconsistent with a random matching approach. While one might expect that moving from linear to nonlinear payoffs might lead to intractable models, we will see that it does not: the dynamics we study are nonlinear even when payoffs in the underlying game are not, so allowing nonlinear payoffs does not lead to a qualitative increase in the complexity of the analysis.

1.1.2 Definitions and Classes of Population Games

Our formal introduction to population games begins in Chapter 2, which offers definitions of population states, payoff functions, and various other items mentioned and unmentioned above. To reduce the degree of abstraction, and to illustrate the range of possible applications, the chapter presents a number of basic examples of population games; among these are *congestion games*, which provide a general but tractable model of network congestion and of related sorts of multilateral externalities. Finally, by showing how low-dimensional population games can be represented in pictures, Chapter 2 ushers in the geometric methods of analysis that are emphasized throughout the book.

From a purely formal point of view, a population game is defined by an arbitrary collection of continuous, real-valued functions on an appropriate domain. While some basic results, including existence of Nash equilibrium, can be proved at this level of generality, obtaining more specific conclusions requires us to focus on classes of games

defined by certain structural properties. In Chapter 3, we introduce three important classes of population games: *potential games*, *stable games*, and *supermodular games*. Each of these classes of games includes a number of important examples, and each is characterized by restrictions on the nature of the externalities that users of each strategy impose on one another. The structure imposed by these restrictions makes analyses of games in these classes relatively simple. For instance, in games from all three classes, one can prove existence of Nash equilibrium without recourse to fixed point theorems. But more important to us is the impact of this structure on disequilibrium behavior: we will see in Chapter 7 that in potential games, stable games, and supermodular games, broad classes of evolutionary dynamics are assured of converging to Nash equilibrium from arbitrary initial conditions.

1.2 Evolutionary Dynamics

The state of equilibrium. . . is therefore *stable*; i.e., if either of the producers, misled as to his true interest, leaves it temporarily, he will be brought back to it by a series of reactions, constantly declining in amplitude, and of which the dotted lines of the figure give a representation by their arrangement in steps.

Cournot (1838)

We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and preferable.

von Neumann and Morgenstern (1944, p. 44–45)

An equilibrium would be just an extreme state of rare occurrence if it were not stable—that is, if there were no forces which tended to restore equilibrium as soon as small deviations from it occurred.

Besides this stability “in the small”, one may consider stability “in the large”—that is, the ability of the system to reach an equilibrium from any initial position.

Beckmann, McGuire, and Winsten (1956, p. 70)

An obvious weakness of the game-theoretic approach to evolution is that it places great emphasis on equilibrium states, whereas evolution is a process of continuous, or at least periodic, change. It is, of course, mathematically easier to analyse equilibria than trajectories of change. There are, however, two situations in which game theory models force us to think about change as well as constancy. The first is that a game may not have an ESS, and hence the population cycles indefinitely. . .

The second situation. . . is when, as is often the case, a game has more than one ESS. Then, in order to account for the present state of the population, one has to allow for initial conditions—that is, for the state of the ancestral population.

Maynard Smith (1982, p. 8)

After modeling a strategic interaction using a population game, one would like to use the game as the basis for predicting how agents in the interaction will behave. Traditionally, most predictions in game theory have been based on equilibrium analysis: one introduces some notion of equilibrium play, and then finds all behaviors in the game that agree with the equilibrium notion. While the equilibrium approach is standard practice in applications of game theory, the quotations above, drawn from seminal works in economic theory, game theory, transportation science, and theoretical biology, all emphasize that this approach is incomplete, and should be complemented by an analysis of dynamics. The latter quotes go further: they point out that local stability analysis, which checks whether equilibrium play will be restored after small disturbances in behavior, is only a first step, as it begs the question of how equilibrium is established in the first place. These concerns are most pronounced in settings with large numbers of players, where the interplayer coordination of beliefs and actions associated with equilibrium seems most difficult to achieve.

The majority of this book studies dynamic models of behavior in large population games, and so is an attempt to provide some answers to the questions raised above.

1.2.1 Knowledge, Rationality, and Large Games

The fundamental solution concept of noncooperative game theory is *Nash equilibrium*: the requirement that each agent choose a strategy that is optimal given the choices of the others. There are many other solution concepts for games, but most of them are refinements of Nash's definition, and are called upon to reduce the set of predictions in games with multiple Nash equilibria.

Despite the central role of the Nash equilibrium concept, the traditional, rationalistic justification for applying this concept is not especially convincing. This justification is based on three assumptions about the players in the game. First, each player is assumed to be *rational*, acting to maximize his payoffs given what he knows. Second, players have *knowledge of the game* they are playing: they know what strategies are available, and what payoffs result from every strategy profile. Third, the players have *equilibrium knowledge*: they are able to anticipate correctly what their opponents will do. If all players expect a certain strategy profile to be played, and if each player is rational and understands the payoff consequences of switching strategies, then each player is content to play his part in the strategy profile if and only if that profile is a Nash equilibrium.

Of the three assumptions listed above, the equilibrium knowledge assumption is the hardest to accept. Certainly, shared expectations about play can be put into place by a disinterested moderator who guides the players to a particular strategy profile. But

without such guidance, it is hard to explain how players can introspectively anticipate how others will act, particularly in games with large numbers of participants.

In fact, when we consider games with many players or strategies, even apparently innocuous conditions for equilibrium play may be called into question. Under the traditional interpretation of equilibrium play in a traffic network, a driver choosing a route to work has a complete mental account of all of the routes he could take, and he is able to anticipate the delay that would arise on each route for any possible profile of choices by his fellow drivers. Evidently, the assumption of knowledge of the game, while seemingly innocent, may actually be quite bold when the game is large.

This discussion suggests that in large games, even the force of seemingly weak solution concepts, ones that do not require equilibrium knowledge, should not be taken for granted. For instance, a basic tenet of traditional game-theoretic analysis holds that a strictly dominated strategy—a strategy that performs worse than some single alternative strategy regardless of how opponents behave—should not be chosen. This requirement is uncontroversial when players have full knowledge of the game. But if players are unable or unwilling to keep the entire game in mind, they may well not notice that one strategy is dominated by another. While one might expect that an accumulation of experience might ensure that players eventually avoid dominated strategies, we will see that this is not necessarily so: in Section 9.4, we will present a set of seemingly mild conditions on players' updating rules that are enough to ensure that strictly dominated strategies must survive in perpetuity in some games.

Obtaining a convincing rationalistic justification for equilibrium play in games seems an impossible task. But in settings where the same game is played many times, the possibilities become brighter, as one can replace introspection with repetition as the basis for coordination of behavior. In large population settings, repetition may be enough to coordinate behavior even when agents' information and abilities are quite limited. But while dynamic approaches can support and even refine traditional game-theoretic predictions, they also can lead to predictions of cyclical or more complex nonstationary behavior, possibilities that are ignored by traditional analyses.

1.2.2 Foundations for Evolutionary Dynamics

There are a variety of approaches one could take to studying disequilibrium dynamics in games, depending on the number of players involved, the information the players are expected to possess, and the importance of the interaction to the players. *Evolutionary game theory*, the approach studied in this book, considers the dynamics of behavior in large, strategically interacting populations. This approach posits that agents only occasionally

switch strategies, and then use simple myopic rules to decide how to act. While these assumptions are certainly not appropriate for every application, they seem natural when the interaction in question is just one among many the agent faces, so that the sporadic application of a rule of thumb is a reasonable way for the agent to proceed.

While it is possible to proceed directly with a description of aggregate behavior dynamics, we find it preferable to begin by specifying when and how individual agents make decisions. We accomplish these tasks using a modeling device called a *revision protocol*. A revision protocol is a function that takes the strategies' payoffs and utilization levels as inputs; it returns as outputs the overall rate of switching strategies, and the probabilities with which each alternative strategy will be chosen.

In defining a revision protocol, we implicitly specify the informational burden that the agents must bear. Starting in Chapter 4, we show that revision protocols come in many different varieties, from ones that embody exact myopic optimization, to others that require each agent to know nothing more than his own current payoff. In all cases, though, the protocol only relies on information about current strategic conditions: historical information, as well as counterfactual information about strategies' performances under other conditions, are not considered.

Also implicit in the definition of a revision protocol is the method to be used to identify alternative strategies. One can place protocols into two broad categories according to this criterion. Under *imitative protocols*, an agent obtains a candidate strategy by observing the strategy of a randomly chosen member of his population. Under *direct protocols*, agents are assumed to choose candidate strategies directly; a strategy's popularity does not directly influence the probability with which it is considered. (Agents may also meander among different protocols as time passes, in which case they are said to employ a *hybrid protocol*.) After obtaining a candidate strategy, an agent can evaluate its current payoff by briefly experimenting with it. If the strategy is currently in use, its current payoff can also be determined by observing the outcomes of an opponent who employs it.

We will see in Chapter 5 that the aggregate behavior dynamics generated by imitative protocols are very similar to those studied in mathematical biology. Thus, the *replicator dynamic*, introduced in the mathematical biology literature to model natural selection, also describes the aggregate behavior of agents who use certain imitative protocols. In contrast, dynamics based on direct selection—for instance, the *best response dynamic*, which is based on optimal myopic choices—behave rather differently than those studied in biology. Direct selection allows unused strategies to be introduced to the population, which is impossible under pure imitation (or under biological reproduction without mutations). For its part, imitation generates dynamics with relatively simple functional forms and behavior. But

the simple forms and special properties of imitative dynamics are rather special: we will find that when agents use hybrid protocols, their aggregate behavior agrees in broad qualitative terms with what one would see under direct selection of alternative strategies.

1.2.3 Deterministic Evolutionary Dynamics

Suppose that one or more large populations of agents recurrently play a population game, with each agent occasionally updating his choice of strategies using a fixed revision protocol. Since there are many agents, and since the stochastic elements of the agents' updating procedures are idiosyncratic, one expects these stochastic influences to be averaged away, leaving aggregate behavior to evolve in an essentially deterministic fashion. In Chapter 4, we explain how such a deterministic evolutionary process can be described by an ordinary differential equation, and how this equation can be derived from the population game and revision protocol. A formal justification for our study of this differential equation, which we call the *mean dynamic*, is deferred until Chapter 10; see Section 1.2.4 below.

The mean dynamic specifies the rate of change in the use of each strategy i . It is the difference of two terms: an inflow term, which captures agents' switches from other strategies to strategy i , and an outflow term, which captures agents' switches from strategy i to other strategies. Of course, the exact specification of the dynamic depends on the primitives of the model: namely, the protocol the agents employ and the game they play.

Taking a slightly different point of view, we note that any fixed revision protocol defines a map that takes population games as inputs, and returns specific instances of mean dynamics as outputs. We call this map from population games to ordinary differential equations an *evolutionary dynamic*.

Chapters 5 and 6 introduce a variety of families of evolutionary dynamics and investigate their properties. Each family of dynamics is generated from a set of revision protocols sharing certain qualitative features: for instance, being based on imitation, or on myopic optimization. The main results in Chapters 5 and 6 establish properties of evolutionary dynamics that hold regardless of the population game at hand.

One of the basic issues considered in these chapters is the relationship between the rest points of the dynamics and the payoff structure of the underlying game. We show in Chapter 5 that the rest points of imitative dynamics include all Nash equilibria of the game being played. Further analyses reveal that under many direct and hybrid dynamics, the sets of rest points and Nash equilibria are identical. This latter property, which we call *Nash stationarity*, provides a first link between the population dynamics and traditional equilibrium analyses. A second issue we address is the connection between out-of-

equilibrium dynamics and incentives in the underlying game. Many of the dynamics we study satisfy a monotonicity property called *positive correlation*, which requires strategies' growth rates to be positively correlated with their payoffs. This property and its relatives are a basic ingredients in our later analyses of dynamic stability.

Properties like Nash stationarity only provide a weak justification for the prediction of Nash equilibrium play. To obtain a more convincing defense of this prediction, one must address the questions of stability raised at the start of this section.

Chapter 8 considers local stability of equilibrium: whether equilibrium will be reached if play begins at a nearby social state. It is easy to see that some Nash equilibria are unlikely to be locally stable. For instance, if play begins near the mixed equilibrium of a coordination game, then myopic adjustment will lead the population away from this equilibrium. Additional restrictions beyond Nash's condition are thus needed to obtain general stability results. A main finding in Chapter 8 is that the notion of an *evolutionarily stable state (ESS)*, introduced by Maynard Smith and Price for model of evolution in populations of mixed strategists, provides a general sufficient condition for local stability under the pure-strategist dynamics studied here.

Local stability results only offer a partial justification of equilibrium predictions. To be convinced that equilibrium will be reached, one must look instead for global convergence results, which establish that equilibrium is attained from any initial state. We offer such results in Chapter 7, where we prove that in certain classes of games—namely, the classes of potential games, stable games, and supermodular games, introduced in Chapter 3—there are classes of dynamics that converge to equilibrium from all or almost all initial conditions.

These stability results provide strong support for the Nash prediction in some settings, but they say little about behavior in games outside of the classes the results cover. While one might hope that these results are nevertheless representative of behavior in most games, this seems not to be so. In Chapter 9, we present a variety of examples in which deterministic dynamics enter cycles far from Nash equilibrium, and others in which the dynamics display chaotic behavior. Thus, if we take the model of myopic adjustment dynamics seriously, we must sometimes accept these more complicated limit behaviors as more credible predictions than equilibrium play. Moreover, as we hinted earlier, the possibility of nonconvergent behavior has some counterintuitive consequences: we show in Section 9.4 that under "typical" evolutionary dynamics, we can always find simple games in which strictly dominated strategies survive.

1.2.4 Orders of Limits for Stochastic Evolutionary Models

The deterministic dynamics studied in Chapters 4–9 form the largest portion of the literature on evolutionary game dynamics. But there is another major branch of the literature that focuses on the infinite horizon behavior of stochastic evolutionary dynamics, and whose central aim is to obtain unique predictions of play in games with multiple equilibria. In Chapter 10, we bring together these two literatures by showing how deterministic and stochastic evolutionary game dynamics can be derived from a single foundation.

To accomplish this, we consider population games played by large but finite populations of agents. Together, a finite-population game and a revision protocol define a stochastic evolutionary process—in particular, a Markov process—on the now finite set of population states. What methods are most useful for studying this process depends on the parameter values—the population size N and the time horizon T —that are relevant to the application at hand.

This point is made most clearly by taking limits. To begin, suppose that we fix the time horizon T and take the population size N to infinity. The main result of Chapter 10 shows that once N is large enough, the stochastic evolutionary process is very likely to behave in a nearly deterministic way, mirroring a solution trajectory of the relevant mean dynamic through time T . For intuition, observe that the stochastic aspects of the evolutionary process—the random arrivals of revision opportunities, and the randomizations among candidate strategies during these opportunities—are idiosyncratic. When the population size becomes very large, the idiosyncratic noise is averaged away. Thus, the behavior of the process is driven by the expected changes in the state, which are precisely what the mean dynamic captures.

Alternatively, suppose we fix the population size N and take the time horizon T to infinity. In this case we are studying the infinite horizon behavior of a finite-state Markov process. If this process is irreducible—that is, if all states are mutually accessible, as is true, for instance, under any revision protocol that always places positive probability on each alternative strategy—then results from probability theory tell us that every social state will be visited infinitely often, and that the proportion of time spent in each state over the infinite horizon does not depend on the initial state. This invariance property opens the door to obtaining unique predictions of play.

1.2.5 Stationary Distributions and Stochastic Stability

The proportion of time that an irreducible Markov process spends in each state is described by the process's *stationary distribution*. By computing this distribution, we can

obtain an exact description of the infinite horizon behavior of our evolutionary process. Chapter 11 presents some basic results on stationary distributions of stochastic evolutionary models, and studies in detail those cases—namely, the models that generate *reversible* Markov processes—in which the stationary distribution can be computed exactly.

These analyses provide a first illustration of a fundamental idea. Even if the underlying game has multiple equilibria, it is nevertheless the case that in the infinite horizon, a stochastic evolutionary dynamic may spend the vast majority of periods in the vicinity of a single state. But we also argue that the amount of time necessary for this analysis to become relevant can be exceedingly long, limiting the range of possible applications.

Once one moves beyond reversible cases, obtaining an exact expression for the stationary distribution is typically impossible. One way of circumventing this difficulty is to take certain parameters of the evolutionary process to their limiting values. If one takes the noise level that parameterizes the agents' revision protocol to zero, or the population size to infinity, one can sometimes describe the limiting stationary distribution, even though the stationary distributions for fixed parameter values cannot be computed explicitly. There are many interesting cases in which the limiting stationary distribution places all of its mass on a single state, providing an especially tidy prediction of infinite horizon play.

States that retain mass in the limiting stationary distribution are said to be *stochastically stable*. In Chapter 12 we offer a complete presentation of stochastic stability theory. We describe the main analytical techniques used to determine the stochastically stable states, and we present key selection results. We also establish formal connections between the states that are stochastically stable in the large population limit and the recurrent states of the relevant mean dynamic. In doing so, we provide another link between stochastic stability analysis and the deterministic dynamics studied in earlier chapters.

1.3 Remarks on History, Motivation, and Interpretation

Evolutionary game theory dates from the early 1970s, when John Maynard Smith introduced the ESS concept as a way of understanding ritualized animal conflict (see the Notes). In pursuing this line of research, Maynard Smith and other biologists borrowed the standard game-theoretic definitions of mathematicians and economists, but reinterpreted these definitions to suit their own purposes. In a biological model, a strategy is not something that an animal chooses, but a behavior that an animal is hardwired to perform by virtue of its genetic endowment. Similarly, payoffs are not a numerical representation of a preference relation, but instead represent “Darwinian fitnesses”: they describe how the use of the strategy improves an animal's prospects for survival and reproduction.

Thus, the development of a biological approach to game theory was marked not only by the invention of new solution concepts and methods of analysis, but also by a recasting of existing definitions.

In the late 1980s, economists surmised that this biological paradigm could be used to address problems with the foundations of traditional game theory. During the heyday of the equilibrium refinements literature, game theorists introduced solution concepts that seemed to impose ever larger demands on the reasoning abilities of the players whose behavior they purported to describe. This development raised the criticism that the solution concepts could only describe the behavior of hyper-rational players. If, though, it could be shown that the behavior of populations of preprogrammed creatures could be described by means of these same solution concepts, then this would seem to provide a way of dispelling the hyper-rationality critique.

In view of this motivation, it is not surprising that many early efforts by economists in evolutionary game theory adopted the biological paradigm wholesale, with populations of agents hardwired to choose certain strategies, and payoffs interpreted as Darwinian fitnesses. These analyses often found that behavior in these populations, as described by ESS and other biological solution concepts, agreed in a formal sense with the most demanding of the equilibrium refinements introduced by economists.

While these links between the biological and economic approaches to game theory might provide some sense of reassurance, it is important to be mindful of the limitations of this line of reasoning. In many economic applications of game theory, the game in question is played by, say, two or three players. Nash equilibrium and its refinements offer predictions about the mixed strategies these players will employ. For their part, the evolutionary models introduced by biologists concern interactions among large populations of animals, and their solution concepts describe the effects of natural selection on population shares of different genetic types. The fact that economists' predictions about rational agents' choices of mixed strategies formally agree with biologists' predictions about population shares under natural selection is certainly intriguing. But there is no obvious logic by which this agreement enables either analysis to justify the other.

Economists working on population dynamics soon realized that to obtain a modeling tool useful for typical social science environments, the Darwinian assumptions and interpretations of evolutionary game theory would need to be replaced with ones reflecting human decision making. This new vision of evolutionary game theory brought back the original interpretation of the game itself: strategies returned to being objects of choice, and payoffs resumed being descriptions of individual preferences. But other parts of the biologists' approach—in particular, the very idea of population dynamics, with changes

in strategies' population shares driven by differences in current payoffs—were retained. Of course, some reinterpretation was necessary here, as the “evolution” of population shares would no longer be driven by relative rates of reproduction, but rather by conscious decisions to switch from one strategy to another. More intriguing agreements were discovered between the biological and economic approaches. For instance, Taylor and Jonker introduced the replicator dynamic in 1978 to provide a dynamic foundation for Maynard Smith's static approach to natural selection. But in the 1990s, economists found that this dynamic could also be derived from models of imitation in populations of economic agents—see Section 5.4.

This history has created some understandable confusion about the meaning of the term “evolutionary game theory”. From the origins of the field through the earliest contributions of economists, this phrase referred exclusively to game-theoretic models of natural selection. But since the mid 1990s, it has also encompassed models of myopic behavior in games played by large populations of active decision makers. In this interpretation, which predominates in this book, the word “evolution” should not be understood as a direct reference to Darwinian survival of the fittest, but instead should be taken in its broader sense, referring to a process of gradual change.

As the interpretation of evolutionary game theory has changed, so too has its role within economic modeling. As we noted earlier, the initial flush of interest in evolutionary game theory in the economics community stemmed from the theory's potential to provide low-rationality foundations for high-rationality solution concepts. This motivation helps explain why economists working in evolutionary game theory have devoted so much effort to the study of random matching in normal form games. If the point of the theory is to motivate solution concepts for such games, then there is not much reason to move beyond the random matching setting.

As we have emphasized here, predictions about the aggregate behavior of large populations do not directly address how two players would play a two-player normal form game. For such predictions to play more than a metaphorical role, they must do so in applications that specifically concern the behavior of populations. But once such applications are brought to the fore, restricting attention to matching in normal form games begins to feel artificial: if one is interested in modeling behavior in large populations, the games one should write down are those one believes populations actually play.

1.N Notes

General references on evolutionary game theory. Books: Maynard Smith (1982), Hofbauer and Sigmund (1988, 1998), Bomze and Pötscher (1989), Cressman (1992, 2003), Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Fudenberg and Levine (1998), Young (1998b). Surveys: van Damme (1991, Chapter 9), Hines (1987), Kandori (1997), Mailath (1998), Weibull (2002), Hofbauer and Sigmund (2003), Sandholm (2009a).

References on omitted topics. *Evolution in extensive form games:* Cressman (1996a, 2000, 2003), Cressman and Schlag (1998), Chamberland and Cressman (2000), Binmore et al. (1995a), Binmore and Samuelson (1999), Ponti (2000) (deterministic models); Samuelson (1994, 1997), Nöldeke and Samuelson (1993), Hart (2002), Kuzmics (2004) (stochastic models). *Local interaction models:* Blume (1993, 1995, 1997), Ellison (1993, 2000), Kosfeld (2002), Miękisz (2004); Herz (1994), Ely (2002), Eshel et al. (1998), Anderlini and Ianni (1996), Goyal and Janssen (1997), Alós-Ferrer and Weidenholzer (2006, 2007, 2008), Goyal (2007), Vega-Redondo (2007), Jackson (2008); Nowak and May (1992, 1993), Nowak (2006), Szabó and Fáth (2007), Hauert (2007). *Continuous strategy sets:* Eshel (1983), Bomze (1990, 1991), Friedman and Yellin (1997), Oechssler and Riedel (2001, 2002), Eshel and Sansone (2003), Cressman (2005, 2009), Cressman and Hofbauer (2005), Cressman et al. (2006), Friedman and Ostrov (2008), Norman (2008), Hofbauer et al. (2009). *Stochastic evolution via diffusion processes:* Foster and Young (1990), Fudenberg and Harris (1992), Cabrales (2000), Imhof (2005), Benaïm et al. (2008), Hofbauer and Imhof (2009).

References on heuristic learning. Books: Fudenberg and Levine (1998), Young (2004), Cesa-Bianchi and Lugosi (2006). Survey: Hart (2005). Brief overview: Sandholm (2008c).

References on applications. *Markets:* Hopkins and Seymour (2002), Lahkar (2007), Droste et al. (2002), Sethi (1999), Ben-Shoham et al. (2004), Agastya (2004), Vega-Redondo (1997), Alós-Ferrer et al. (2000), Ania et al. (2002), Alós-Ferrer and Ania (2005), Alós-Ferrer et al. (2006), Kandori et al. (2008), Friedman and Ostrov (2008). *Bargaining and hold-up problems:* Young (1993b, 1998a,b), Burke and Young (2001), Ellingsen and Robles (2002), Tröger (2002), Binmore et al. (2003), Abreu and Sethi (2003), Dawid and MacLeod (2008), Robles (2008). *Signaling and cheap talk:* Nöldeke and Samuelson (1993, 1997), Jacobsen et al. (2001), Robson (1990), Matsui (1991), Wärneryd (1993), Blume et al. (1993), Kim and Sobel (1995), Bhaskar (1998), Banerjee and Weibull (2000), Trapa and Nowak (2000), Skyrms (2002), Pawlowitsch (2008), Demichelis and Weibull (2008). *Public good provision:* Sethi and Somanathan (1996), Myatt and Wallace (2008a,b, 2009). *Implementation and decentralized control:* Cabrales (1999), Cabrales and Ponti (2000), Sandholm (2002, 2005b, 2007b), Cabrales and Serrano (2007), Mathevet (2007), Arslan et al. (2007), Marden et al. (2009a). *Residential segregation:* Young (1998b, 2001), Möbius (2000), Zhang (2004a,b), Dokumacı

and Sandholm (2007a). *Preference evolution*: Güth and Yaari (1992), Güth (1995), Huck and Oechssler (1999), Koçkesen et al. (2000), Sethi and Somanathan (2001), Ok and Vega-Redondo (2001), Ely and Yilankaya (2001), Dekel et al. (2007), Heifetz et al. (2007), Herold and Kuzmics (2008). *Cultural evolution*: Bisin and Verdier (2001), Sandholm (2001b), Kuran and Sandholm (2008), Montgomery (2009). *Applications in biology*: Maynard Smith (1982), Hofbauer and Sigmund (1988) (books); Hammerstein and Selten (1994), Dugatkin and Reeve (1998) (surveys).

Section 1.1. Noncooperative games with a continuum of agents were introduced by Schmeidler (1973), who proves that pure strategy Nash equilibria exist when each player's payoffs only depend on his own action and on the aggregate behavior of all players. Schmeidler (1973) uses a normal form approach, in that he names and describes the behavior of each agent separately. Mas-Colell (1984) introduces a simpler formulation based on empirical distributions of strategies rather than on strategy profiles; his formulation can be understood as a population game in which both the set of populations and the set of strategies may be continuous. Mas-Colell (1984) also points out the overlap between his model and the distributional approach to Bayesian games with a continuum of types (Milgrom and Weber (1985)). This connection can be used as the basis for studying evolution in Bayesian games; see Ely and Sandholm (2005). The literature on noncooperative games with large numbers of players is surveyed in Khan and Sun (2002); see Balder (2002) and Carmona and Podczeck (2009) for recent results.

Nash's "mass action" interpretation of his equilibrium concept appears in his dissertation (Nash (1950b)), but not in either of its published versions (Nash (1950a, 1951)). It was later rediscovered by Leonard (1994), Weibull (1994, 1996), and Björnerstedt and Weibull (1996); see also Hofbauer (2000).

Section 1.2. There are many good general references on traditional game-theoretic analysis: see Osborne (2004) at the undergraduate level, and Fudenberg and Tirole (1991), Myerson (1991), Ritzberger (2002), and van Damme (1991) at the graduate level. For an overview of the use of rationality and knowledge assumptions in game theory, see Dekel and Gul (1997).

While evolutionary game theory models assume that agents are myopic, one can imagine large population interactions whose participants are better modeled as forward-looking dynamic optimizers, who attempt to maximize discounted flows of future payoffs. The *perfect foresight dynamics* of Matsui and Matsuyama (1995) introduce forward-looking behavior to large population models while retaining the assumption of inertia from the evolutionary approach. For additional references on this modeling approach, see the Notes to Chapter 5.

Section 1.3. A series of papers by John Maynard Smith in the early 1970s (Maynard Smith (1972, 1974), Maynard Smith and Price (1973)) are traditionally viewed as heralding the birth of evolutionary game theory. But Maynard Smith (1982, p. 2) himself notes the appearance of game-theoretic ideas in earlier studies of the evolution of sex ratios: they appeared implicitly in the work of Fisher (1930), and explicitly in the work of Hamilton (1967). But it was only with Maynard Smith's contributions that game theory began to take on its role as a standard framework for biological modeling.

Part I

Population Games

Population Games

2.0 Introduction

As we saw in Chapter 1, there are many important strategic interactions that exhibit the following properties:

- (i) *The number of agents is large.*
- (ii) *Individual agents are small:* Any one agent's behavior has little or no effect on other agents' payoffs.
- (iii) *Agents interact anonymously:* Each agent's payoffs only depend on opponents' behavior through the distribution of their choices.

While these three properties are basic, two additional restrictions of a more technical nature will help us focus on the environments that interest us most. For the "distribution of opponents' choices" mentioned in property (iii) to exist, there must be collections of agents who share the same set of strategies. This structure and more are ensured by

- (iv) *The number of roles is finite:* Each agent is a member of one of a finite number of populations. Members of a population choose from the same set of strategies, and their payoffs are identical functions of own behavior and the distribution of opponents' behavior.

Finally, to ensure that very small changes in aggregate behavior do not lead to large changes in strategies' payoffs, we introduce

- (v) *Payoffs are continuous:* The dependence of each agent's payoffs on the distribution of opponents' choices is continuous.

Population games, the subject of this chapter, provide a unified framework for studying strategic interactions with the properties listed above. Applications range from economics

(externalities, macroeconomic spillovers, centralized markets) to biology (animal conflict, genetic natural selection), to transportation science (highway network congestion, mode choice) to computer science (selfish routing of Internet traffic). Section 1.1 provides further applications and discussions.

This chapter formally defines population games and offers a variety of explicit examples. It also introduces some basic tools for the geometric analysis of population games, an approach we develop further in subsequent chapters of the book. The mathematics that underlies this approach—in particular, the notions of affine spaces, tangent spaces, and orthogonal projections—is reviewed in Appendix 2.A.

Until the final chapters of the book, we find it convenient to study games with continuous populations of agents. This framework offers a few distinct advantages: Chapter 3 will demonstrate the power of tools from analysis in studying continuous-population games, and these games are the natural setting for the differential equation models that are the focus of Chapters 4–9. However, one can equally well define large games with finite populations of players. We will do so in Chapters 11 and 12, where these games provide a tractable environment for defining stochastic evolutionary dynamics. Since continuous-population models can be obtained by taking limits of finite-population models (see Section 11.4), the choice between the two settings should be viewed less as a matter of substance than as a matter of convenience.

2.1 Population Games

This section formally defines population games and their equilibria. To reduce the level of abstraction, it may be helpful to refer ahead to the presentation of congestion games in Example 2.2.5 to see how the definitions introduced below are employed in an application.

2.1.1 Populations, Strategies, and States

Let $\mathcal{P} = \{1, \dots, p\}$, be a *society* consisting of $p \geq 1$ *populations of agents*. Agents in population p form a continuum of *mass* $m^p > 0$. (Thus, p is the number of populations, while p is an arbitrary population.)

The set of *strategies* available to agents in population p is denoted $S^p = \{1, \dots, n^p\}$, and has typical elements i, j , and (in the context of normal form games) s^p . We let $n = \sum_{p \in \mathcal{P}} n^p$ equal the total number of pure strategies in all populations.

During game play, each agent in population p selects a (pure) strategy from S^p . The set

of *population states* (or *strategy distributions*) for population p is $X^p = \{x^p \in \mathbf{R}_+^{m^p} : \sum_{i \in S^p} x_i^p = m^p\}$. The scalar $x_i^p \in \mathbf{R}_+$ represents the mass of players in population p choosing strategy $i \in S^p$. Elements of X_v^p , the set of vertices of X^p , are called *pure population states*, since at these states all agents choose the same strategy.

Elements of $X = \prod_{p \in \mathcal{P}} X^p = \{x = (x^1, \dots, x^p) \in \mathbf{R}_+^n : x^p \in X^p\}$, the set of *social states*, describe behavior in all p populations at once. The elements of $X_v = \prod_{p \in \mathcal{P}} X_v^p$ are the vertices of X , and are called the *pure social states*.

When there is just one population ($p = 1$), we assume that its mass is 1, and we omit the superscript p from all of our notation: thus, the strategy set is $S = \{1, \dots, n\}$, the state space is $X = \{x \in \mathbf{R}_+^n : \sum_{i \in S} x_i = 1\}$, the simplex in \mathbf{R}^n , and the set of pure states $X_v = \{e_i : i \in S\}$ is the set of standard basis vectors in \mathbf{R}^n .

2.1.2 Payoffs

We generally take the sets of populations and strategies as fixed and identify a game with its payoff function. A *payoff function* $F : X \rightarrow \mathbf{R}^n$ is a continuous map that assigns each social state a vector of payoffs, one for each strategy in each population. $F_i^p : X \rightarrow \mathbf{R}$ denotes the payoff function for strategy $i \in S^p$, while $F^p : X \rightarrow \mathbf{R}^{m^p}$ denotes the payoff functions for all strategies in S^p .

While our standing assumption is that F is continuous, we often impose the stronger requirements that F be Lipschitz continuous or continuously differentiable (C^1). These additional assumptions will be made explicit whenever we use them.

We define

$$\bar{F}^p(x) = \frac{1}{m^p} \sum_{i \in S^p} x_i^p F_i^p(x)$$

to be the (*weighted*) *average payoff* obtained by members of population p at social state x . Similarly, we let

$$\bar{F}(x) = \sum_{p \in \mathcal{P}} \sum_{i \in S^p} x_i^p F_i^p(x) = \sum_{p \in \mathcal{P}} m^p \bar{F}^p(x)$$

denote the *aggregate payoff* achieved by the society as a whole.

2.1.3 Best Responses and Nash Equilibria

To describe optimal behavior, we define population p 's *pure best response correspondence*, $b^p : X \Rightarrow S^p$, which specifies the strategies in S^p that are optimal at each social state x :

$$b^p(x) = \operatorname{argmax}_{i \in S^p} F_i^p(x).$$

Let $\Delta^p = \{y^p \in \mathbf{R}_+^{n^p} : \sum_{i \in S^p} y_i^p = 1\}$ denote the simplex in \mathbf{R}^{n^p} . The *mixed best response correspondence* for population p , $B^p : X \Rightarrow \Delta^p$ is given by

$$B^p(x) = \{y^p \in \Delta^p : y_i^p > 0 \Rightarrow i \in b^p(x)\}.$$

In words, $B^p(x)$ is the set of probability distributions in Δ^p whose supports only contain pure strategies that are optimal at x . Geometrically, $B^p(x)$ is the convex hull of the vertices of Δ^p corresponding to elements of $b^p(x)$.

Social state $x \in X$ is a *Nash equilibrium* of the game F if each agent in every population chooses a best response to x :

$$NE(F) = \{x \in X : x^p \in m^p B^p(x) \text{ for all } p \in \mathcal{P}\}.$$

Equivalently, x is a Nash equilibrium if in each population, every strategy in use earns the maximal payoff:

$$NE(F) = \{x \in X : [x_i^p > 0 \Rightarrow F_i^p(x) \geq F_j^p(x)] \text{ for all } i, j \in S^p, p \in \mathcal{P}\}.$$

In Section 2.3.6, we will use this second definition to characterize the set of Nash equilibria in a purely geometric way.

Nash equilibria always exist:

Theorem 2.1.1. *Every population game admits at least one Nash equilibrium.*

Theorem 2.1.1 can be proved by applying Kakutani's Theorem to the profile of best response correspondences. But we will see that in each of the three classes of games we focus on in Chapter 3—potential games (Sections 3.1 and 3.2), stable games (Section 3.3), and supermodular games (Section 3.4)—existence of Nash equilibrium can be established without recourse to fixed point theorems.

2.1.4 Prelude to Evolutionary Dynamics

In traditional game-theoretic analyses, it is usual to assume that players follow some Nash equilibrium of the game at hand. But because population games involve large numbers of agents, the equilibrium assumption is quite strong, making it more appealing to rely on less demanding assumptions. Therefore, rather than assume equilibrium play, we suppose that individual agents gradually adjust their choices to their current strategic environment. We then ask whether or not the induced behavior trajectories converge to Nash equilibrium. When they do, the Nash prediction can be justified; when they do not, the Nash prediction may be unwarranted.

The question of convergence to equilibrium is a central issue in this book. We will see later on that in the three classes of games studied in Chapter 3, convergence results can be established in some generality—that is, without being overly specific about the exact nature of the agents’ revision protocols. But in this chapter and the next, we confine ourselves to introducing population games and studying their equilibria.

2.2 Examples

To fix ideas, we offer four examples of population games. These examples and many more will be developed and analyzed through the remainder of the book.

2.2.1 Matching in Normal Form Games

Let us begin with the canonical example of evolutionary game theory.

Example 2.2.1. Matching in a single population to play a symmetric game. A symmetric two player normal form game is defined by a strategy set $S = \{1, \dots, n\}$ and a payoff matrix $A \in \mathbf{R}^{n \times n}$. A_{ij} is the payoff a player obtains when he chooses strategy i and his opponent chooses strategy j ; this payoff does not depend on whether the player in question is called player I or player II. Below is the bimatrix corresponding to A when $n = 3$.

		Player II		
		1	2	3
Player I	1	A_{11}, A_{11}	A_{12}, A_{21}	A_{13}, A_{31}
	2	A_{21}, A_{12}	A_{22}, A_{22}	A_{23}, A_{32}
	3	A_{31}, A_{13}	A_{32}, A_{23}	A_{33}, A_{33}

To obtain a population game from this normal form game, suppose that agents in a single (unit mass) population are matched to play A , with each pair of agents meeting

exactly once. Aggregating payoffs over all matches, we find that the payoff to strategy i when the population state is x is $F_i(x) = \sum_{j \in S} A_{ij}x_j$. It follows that the population game associated with A is described by the linear map $F(x) = Ax$.

Instead of assuming a complete, deterministic matching, one can assume alternatively that each agent is randomly matched against a single opponent. Then if agents evaluate probability distributions over payoffs by taking expectations (i.e., if the entries of the matrix A are von Neumann–Morgenstern utilities), we can again represent the payoffs to the interaction as $F(x) = Ax$. §

When we consider foundations for evolutionary dynamics beginning in Chapter 4, some of the choice rules we consider are easiest to interpret when $F_i(x)$ represents a realized payoff rather than an expected payoff. For this reason, we prefer to use complete matchings rather than random matchings to derive population games from normal form games.

Of course, one can also think of the population game $F(x) = Ax$ as describing an interaction that is not based on matching in a normal form game at all. Such interactions are sometimes referred to “playing the field” models; congestion games, defined below, offer one basic example.

Example 2.2.2. Matching in two populations. A (possibly asymmetric) two player game is defined by two strategy sets, $S^1 = \{1, \dots, n^1\}$ and $S^2 = \{1, \dots, n^2\}$, and two payoff matrices, $U^1 \in \mathbf{R}^{n^1 \times n^2}$ and $U^2 \in \mathbf{R}^{n^2 \times n^1}$. The corresponding bimatrix when $n^1 = 2$ and $n^2 = 3$ is as follows.

		Player II		
		1	2	3
Player I	1	U_{11}^1, U_{11}^2	U_{12}^1, U_{12}^2	U_{13}^1, U_{13}^2
	2	U_{21}^1, U_{21}^2	U_{22}^1, U_{22}^2	U_{23}^1, U_{23}^2

To define the corresponding population game, we suppose that there are two unit mass populations, one corresponding to each player role. Each agent is matched with every member of the other population to play the game (U^1, U^2) . Alternatively, one can assume that each agent is randomly matched with a single member of the other population—see the previous example. Either way, the payoff functions for populations 1 and 2 are given by $F^1(x) = U^1x^2$ and $F^2(x) = (U^2)'x^1$, so the entire population game is described by the linear map

$$F(x) = \begin{pmatrix} F^1(x) \\ F^2(x) \end{pmatrix} = \begin{pmatrix} 0 & U^1 \\ (U^2)' & 0 \end{pmatrix} \begin{pmatrix} x^1 \\ x^2 \end{pmatrix} = \begin{pmatrix} U^1x^2 \\ (U^2)'x^1 \end{pmatrix}. \text{ §}$$

Example 2.2.3. Matching in p populations. To generalize the previous example, we define a p player normal form game. Let $S^p = \{1, \dots, n^p\}$ denote player p 's strategy set and $S = \prod_{q \in \mathcal{P}} S^q$ the set of pure strategy profiles; player p 's payoff function U^p is a map from S to \mathbf{R} .

In the population game, agents in p unit mass populations are matched to play the normal form game $U = (U^1, \dots, U^p)$, with each matching of p agents, one from each population, occurring exactly once. This procedure yields a population game with the multilinear (i.e., linear in each x^p) payoff function

$$F_{s^p}^p(x) = \sum_{s^{-p} \in S^{-p}} U^p(s^1, \dots, s^p) \prod_{r \neq p} x_{s^r}^r, \text{ where } S^{-p} = \prod_{q \neq p} S^q. \S$$

We conclude with an observation relating the Nash equilibria of population games generated by matching to those of the underlying normal form games.

- Observation 2.2.4.** (i) *In the single population case (Example 2.2.1), the Nash equilibria of F are the symmetric Nash equilibria of the symmetric normal form game $U = (A, A')$.*
(ii) *In the multipopulation cases (Examples 2.2.2 and 2.2.3), the Nash equilibria of F are the Nash equilibria of the normal form game $U = (U^1, \dots, U^p)$.*

2.2.2 Congestion Games

Because of the linearity of the expectation operator, matching in normal form games generates population games with linear or multilinear payoffs. Moreover, when $p \geq 2$, each agent's payoffs are independent of the behavior of other members of his population. Outside of the matching context, neither of these properties need hold. Our next class of example provides a case in point.

Example 2.2.5. Congestion games. Consider the following model of highway congestion. A collection of towns is connected by a network of *links* (Figure 2.2.1). For each ordered pair of towns there is a population of agents, each of whom needs to commute from the first town in the pair (where he lives) to the second (where he works). To accomplish this, the agent must choose a *path* connecting the two towns. The payoff the agent obtains is the negation of the delay on the path he takes. The delay on the path is the sum of the delays on its constituent links, while the delay on a link is a function of the number of agents who use that link.

Congestion games are used to study not only highway congestion, but also more general settings involving "symmetric" externalities. To define a congestion game, we begin with a finite collection of *facilities* (e.g., links in a highway network), denoted Φ .




Figure 2.2.1: A highway network.

Every strategy $i \in S^p$ requires the use of some collection of facilities $\Phi_i^p \subseteq \Phi$ (e.g., the links in route i). The set $\rho^p(\phi) = \{i \in S^p : \phi \in \Phi_i^p\}$ contains those strategies in S^p that require facility ϕ .

Each facility ϕ has a *cost function* $c_\phi : \mathbf{R}_+ \rightarrow \mathbf{R}$ whose argument is the facility's *utilization level* u_ϕ , the total mass of agents using the facility:

$$u_\phi(x) = \sum_{p \in \mathcal{P}} \sum_{i \in \rho^p(\phi)} x_i^p.$$

Payoffs in the congestion game are obtained by summing the appropriate facility costs and multiplying by -1 .

$$F_i^p(x) = - \sum_{\phi \in \Phi_i^p} c_\phi(u_\phi(x)).$$

Since driving on a link increases the delays experienced by other drivers on that link, cost functions in models of highway congestion are increasing; they are typically convex as well. On the other hand, when congestion games are used to model settings with positive externalities (e.g., consumer technology choice), cost functions are decreasing. Evidently, payoffs in congestion games depend on own-population behavior, and need only be linear if the underlying cost functions are linear themselves.

Congestion games are the leading examples of potential games (Sections 3.1 and 3.2; congestion games with increasing cost functions are also stable games (Section 3.3). §

2.2.3 Two Simple Externality Models

We conclude this section with two simpler models of externalities.

Example 2.2.6. Asymmetric negative externalities. Agents from a single population choose from a set of n activities. There are externalities both within and across activities; the increasing C^1 function $c_{ij} : [0, 1] \rightarrow \mathbf{R}$ represents the cost imposed on agents who choose activity i by agents who choose activity j . Payoffs in this game are described by

$$F_i(x) = - \sum_{j \in S} c_{ij}(x_j).$$

If “own activity” externalities are strong, in the sense that the derivatives of the cost functions satisfy

$$2c'_{ii}(x_i) \geq \sum_{j \neq i} (c'_{ij}(x_j) + c'_{ji}(x_i)),$$

then F is a stable game (Section 3.3 of Chapter 3). §

Example 2.2.7. Search with positive externalities. Consider this simple model of macroeconomic spillovers. Members of a single population choose levels of search effort from the set $S = \{1, \dots, n\}$. Stronger efforts increase the likelihood of finding trading partners, so that payoffs are increasing both in own search effort and in aggregate search effort. In particular, payoffs are given by

$$F_i(x) = m(i) b(a(x)) - c(i),$$

where $a(x) = \sum_{k=1}^n kx_k$ represents aggregate search effort, the increasing function $b : \mathbf{R}_+ \rightarrow \mathbf{R}$ represents the benefits of search as a function of aggregate effort, the increasing function $m : S \rightarrow \mathbf{R}$ is the benefit multiplier, and the arbitrary function $c : S \rightarrow \mathbf{R}$ captures search costs. In Section 3.4, we will show that F is a supermodular game. §

2.3 The Geometry of Population Games and Nash Equilibria

In low-dimensional cases, we can present the payoff vectors generated by a population game in pictures. Doing so provides a way of visualizing the strategic forces at work; moreover, the geometric insights we obtain can be extended to games that we cannot draw.




Figure 2.3.1: Payoffs in 12 Coordination.

2.3.1 Drawing Two-Strategy Games

The population games that are easiest to draw are *two-strategy games*: i.e., games played by a single population of agents who choose between a pair of strategies. When drawing a two-strategy game, we represent the simplex as a subset of \mathbf{R}^2 . We synchronize the drawing with the layout of the payoff matrix by using the *vertical* coordinate to represent the mass on the *first* strategy and the *horizontal* coordinate to represent the mass on the *second* strategy. We then select a group of states spaced evenly through the simplex; from each state x , we draw an arrow representing the payoff vector $F(x)$ that corresponds to x . (Actually, we draw scaled-down versions of the payoff vectors in order to make the diagrams easier to read.)

In Figures 2.3.1 and 2.3.2, we present the payoff vectors generated by the two-strategy coordination game F^{C^2} and the Hawk-Dove game F^{HD} :

$$F^{C^2}(x) = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \end{pmatrix}; \quad F^{HD}(x) = \begin{pmatrix} -1 & 2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_H \\ x_D \end{pmatrix} = \begin{pmatrix} 2x_D - x_H \\ x_D \end{pmatrix}$$

Let us focus on the coordination game F^{C^2} . At the pure state $e_1 = (1,0)$ at which all agents play strategy 1, the payoffs to the two strategies are $F_1^{C^2}(e_1) = 1$ and $F_2^{C^2}(e_1) = 0$; hence, the arrow representing $F^{C^2}(e_1)$ points directly *upward* from state e_1 . At the interior Nash equilibrium $x^* = (x_1^*, x_2^*) = (\frac{2}{3}, \frac{1}{3})$, each strategy earns a payoff of $\frac{2}{3}$, so the arrow representing payoff vector $F^{C^2}(x^*) = (\frac{2}{3}, \frac{2}{3})$ is drawn at a right angle to the simplex at x^* .




Figure 2.3.2: Payoffs in the Hawk-Dove game.

Similar logic explains how the payoff vectors are drawn at other states, and how the Hawk-Dove figure is constructed as well.

The diagrams of F^{C2} and F^{HD} help us visualize the incentives faced by agents playing these games. In the coordination game, the payoff vectors “push outward” toward the two axes, reflecting an incentive structure that drives the population toward the two pure Nash equilibria. In contrast, payoff vectors in the Hawk-Dove game “push inward”, away from the axes, reflecting forces leading the population toward the interior Nash equilibrium $x^* = (\frac{1}{2}, \frac{1}{2})$.

2.3.2 Displacement Vectors and Tangent Spaces

To draw games with more than two strategies we need to introduce two new objects: TX , the tangent space of the state space X ; and Φ , the orthogonal projection of \mathbf{R}^n onto TX . We summarize the relevant concepts in this subsection and the next; for a fuller treatment, see the Appendix.

To start, let us focus on a single-population game F . Imagine that the population is initially at state x , and that a group of agents of mass ε switch from strategy i to strategy j . These revisions move the state from x to $x + \varepsilon(e_j - e_i)$: the mass of agents playing strategy i goes down by ε , while the mass of agents playing strategy j goes up by ε . Vectors like $\varepsilon(e_j - e_i)$, which represent the effects of such strategy revisions on the population state, are called *displacement vectors*. (Since these vectors are tangent to the state space X , we also




Figure 2.3.3: Displacement vectors for two-strategy games.

call them *tangent vectors*—more on this below.)

In Figure 2.3.3, we illustrate displacement vectors for two-strategy games. In this setting, displacement vectors can only point in two directions: when agents switch from strategy 1 to strategy 2, the state moves in direction $e_2 - e_1$, represented by an arrow pointing southeast; when agents switch from strategy 2 to strategy 1, the state moves in direction $e_1 - e_2$, represented by an arrow pointing northwest. Both of these vectors are tangent to the state space X .

(Two clarifications are in order here. First, remember that a vector is characterized by its direction and its length, not where we position its base. When we draw an arrow representing the vector z , we use the context to determine an appropriate position x for the arrow's base; the arrow takes the form of a directed line segment from x to $x + z$. Second, since we are mainly interested in displacement vectors' *relative* sizes, we rescale them before drawing them, just as we did with payoff vectors in Figures 2.3.1 and 2.3.2.)

Now consider a *three-strategy game*: a game with one population and three strategies, whose state space X is thus the simplex in \mathbf{R}^3 . A “three-dimensional” picture of X is provided in Figure 2.3.4, where X is situated within the plane in \mathbf{R}^3 that contains it. This plane is called the *affine hull* of X , and is denoted by $\text{aff}(X)$ (see Appendix 2.A.2). For future reference, note that displacement vectors drawn from states in X are situated in the plane $\text{aff}(X)$.

Instead of representing the state space X explicitly in \mathbf{R}^3 , it is more common to present it as a two-dimensional equilateral triangle (Figure 2.3.5). When we follow this approach,




Figure 2.3.4: The simplex in \mathbf{R}^3 .




Figure 2.3.5: Displacement vectors for three-strategy games.

our sheet of paper itself represents the affine hull $\text{aff}(X)$, and so arrows drawn on the paper represent displacement vectors. Figure 2.3.5 presents arrows describing the $3 \times 2 = 6$ displacement vectors of the form $e_j - e_i$, which correspond to switches between distinct ordered pairs of strategies. Each of these arrows is parallel to some edge of the simplex. For purposes of orientation, note that if we resituate the simplex from Figure 2.3.5 in three-dimensional space (i.e., in Figure 2.3.4), then each of these six arrows is obtained by subtracting one standard basis vector from another.

Switches between pairs of strategies are not the only ways of generating displacement vectors—they can also come from switches involving three or more strategies, and, in multipopulation settings, from switches occurring within more than one population. The set of all displacement vectors from states in X forms a subspace of \mathbf{R}^n ; this subspace is called the *tangent space* TX .

To formally define TX , let us first consider population $p \in \mathcal{P}$ in isolation. The state space for population p is $X^p = \{x^p \in \mathbf{R}_+^{n^p} : \sum_{i \in S^p} x_i^p = m^p\}$. The *tangent space* of X^p , denoted TX^p , is the smallest subspace of \mathbf{R}^{n^p} that contains all vectors describing motions between points in X^p . In other words, if $x^p, y^p \in X^p$, then $y^p - x^p \in TX^p$, and TX^p is the span of all vectors of this form. It is not hard to see that $TX^p = \mathbf{R}_0^{n^p} \equiv \{z^p \in \mathbf{R}^{n^p} : \sum_{i \in S^p} z_i^p = 0\}$: that is, TX^p contains exactly those vectors in \mathbf{R}^{n^p} whose components sum to zero. The restriction on the sum embodies the fact that changes in the population state leave the population's mass constant.

The definition above is sufficient for studying single population games. What if there are multiple populations? In this case, any change in the social state $x \in X = \prod_{p \in \mathcal{P}} X^p$ is a combination of changes occurring within the individual populations. Therefore, the grand tangent space TX is just the product of the tangent spaces for each set X^p : in other words, $TX = \prod_{p \in \mathcal{P}} TX^p$.

2.3.3 Orthogonal Projections

Suppose we would like to draw a diagram representing a three-strategy game F . One possibility is to draw a “three-dimensional” representation of F in the fashion of Figure 2.3.4. We would place more modest demands on our drafting skills if we instead represented F in just two dimensions. But this simplification comes at a cost: since three-dimensional payoff vectors $F(x) \in \mathbf{R}^3$ will be presented as two-dimensional objects, some of the information contained in these vectors will be lost.

From a geometric point of view, the most natural way to proceed is pictured in Figure 2.3.6: instead of drawing an arrow from state x corresponding to the vector $F(x)$ itself, we instead draw the arrow closest to $F(x)$ among those that lie in the plane $\text{aff}(X)$. This arrow




Figure 2.3.6: Projected payoff vectors for three-strategy games.

represents a vector in the tangent space TX : namely, the orthogonal projection of $F(x)$ onto TX .

Let Z be a linear subspace of \mathbf{R}^n . The *orthogonal projection* of \mathbf{R}^n onto Z is a linear map that sends each $\pi \in \mathbf{R}^n$ to the closest point to π in Z . Each orthogonal projection can be represented by a matrix $P_Z \in \mathbf{R}^{n \times n}$ via the map $\pi \mapsto P_Z \pi$, and it is common to identify the projection with its matrix representation. We treat orthogonal projections in some detail in Appendix 2.A.3; here we focus only on the orthogonal projections we need.

First consider population $p \in \mathcal{P}$ in isolation. The orthogonal projection of \mathbf{R}^{n^p} onto the tangent space TX^p , denoted $\Phi \in \mathbf{R}^{n^p \times n^p}$, is defined by $\Phi = I - \frac{1}{n^p} \mathbf{1}\mathbf{1}'$, where $\mathbf{1} = (1, \dots, 1)'$ is the vector of ones; thus $\frac{1}{n^p} \mathbf{1}\mathbf{1}'$ is the matrix whose entries are all $\frac{1}{n^p}$.

If π^p is a payoff vector in \mathbf{R}^{n^p} , the projection of π^p onto TX^p is

$$\Phi \pi^p = \pi^p - \frac{1}{n^p} \mathbf{1}\mathbf{1}' \pi^p = \pi^p - \mathbf{1} \left(\frac{1}{n^p} \sum_{k \in S^p} \pi_k^p \right).$$

The i th component of $\Phi \pi^p$ is the difference between the actual payoff to strategy i and the *unweighted* average payoff of all strategies in S^p . Thus, $\Phi \pi^p$ discards information about average payoffs while retaining information about *relative* payoffs of different strategies in S^p . This interpretation is important from a game-theoretic point of view, since incentives, and hence Nash equilibria, only depend on payoff differences. Therefore, when incentives (as opposed to, e.g., efficiency) are our main concern, we do not need to know the actual payoff vectors π^p ; looking at the projected payoff vectors $\Phi \pi^p$ is enough.

In multipopulation settings, the tangent space $TX = \prod_{p \in \mathcal{P}} TX^p$ has a product structure; hence, the orthogonal projection onto TX , denoted $\Phi \in \mathbf{R}^{n \times n}$, has a block diagonal structure: $\Phi = \text{diag}(\Phi, \dots, \Phi)$. (Note that the blocks on the diagonal of Φ are generally not identical: the p th block is an element of $\mathbf{R}^{n^p \times n^p}$.) If we apply Φ to the society's payoff vector $\pi = (\pi^1, \dots, \pi^p)$, the resulting vector $\Phi\pi = (\Phi\pi^1, \dots, \Phi\pi^p)$ lists the relative payoffs in each population.

2.3.4 Drawing Three-Strategy Games

Before using orthogonal projections to draw three-strategy games, let us see how this device affects our pictures of two-strategy games. Applying the projection $\Phi = I - \frac{1}{2}\mathbf{1}\mathbf{1}'$ to the payoff vectors from the coordination game F^{C2} and the Hawk-Dove game F^{HD} yields

$$\begin{aligned}\Phi F^{C2}(x) &= \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} x_1 \\ 2x_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{2}x_1 - x_2 \\ -\frac{1}{2}x_1 + x_2 \end{pmatrix} \text{ and} \\ \Phi F^{HD}(x) &= \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} 2x_D - x_H \\ x_D \end{pmatrix} = \begin{pmatrix} \frac{1}{2}(x_D - x_H) \\ \frac{1}{2}(x_H - x_D) \end{pmatrix}.\end{aligned}$$

We draw the projected payoffs along with the original payoffs in Figures 2.3.7 and 2.3.8.

The projected payoff vectors $\Phi F(x)$ lie in the tangent space TX , and so are represented by arrows running parallel to the simplex X . Projecting away the orthogonal component of payoffs makes the “outward force” in the coordination game and the “inward force” in the Hawk-Dove game more transparent. Indeed, Figures 2.3.7 and 2.3.8 are suggestive of evolutionary dynamics for these two games—a topic we take up starting in Chapter 4.

Now, let us consider the three-strategy coordination game F^{C3} and the Rock-Paper-Scissors game F^{RPS} .

$$F^{C3}(x) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix}; \quad F^{RPS}(x) = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \end{pmatrix} = \begin{pmatrix} x_S - x_P \\ x_R - x_S \\ x_P - x_R \end{pmatrix}.$$

These games are pictured in Figures 2.3.9 and 2.3.10. The arrows in Figure 2.3.9 represent the projected payoff vectors $\Phi F^{C3}(x)$, defined by

$$\Phi F^{C3}(x) = \left(I - \frac{1}{3}\mathbf{1}\mathbf{1}'\right) \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix} = \begin{pmatrix} \frac{1}{3}(2x_1 - 2x_2 - 3x_3) \\ \frac{1}{3}(-x_1 + 4x_2 - 3x_3) \\ \frac{1}{3}(-x_1 - 2x_2 + 6x_3) \end{pmatrix}.$$




Figure 2.3.7: Payoffs and projected payoffs in 12 Coordination.




Figure 2.3.8: Payoffs and projected payoffs in the Hawk-Dove game.




Figure 2.3.9: Projected payoffs in 123 Coordination.

But in the Rock-Paper-Scissors game, the column sums of the payoff matrix all equal 0, implying that the maps F^{RPS} and ΦF^{RPS} are identical; by drawing one, we draw the other.

As with that of F^{C2} , the diagram of the coordination game F^{C3} shows forces pushing outward toward the extreme points of the simplex. In contrast, Figure 2.3.10 displays a property that cannot occur with just two strategies: instead of driving toward some Nash equilibrium, the arrows in Figure 2.3.10 cycle around the simplex. Thus, the figure suggests that in the Rock-Paper-Scissors game, evolutionary dynamics need not converge to Nash equilibrium, but instead may avoid equilibrium in perpetuity. We return to questions of convergence and nonconvergence of evolutionary dynamics beginning in later chapters.

2.3.5 Tangent Cones and Normal Cones

To complete our introduction to the geometric approach to population games, we explain how one can find a game's Nash equilibria by examining a picture of the game.

To begin this discussion, note that the constraint that defines vectors z as lying in the tangent space TX —the constraint that keeps population masses constant—is not always enough to ensure that motion in direction z is feasible. Motions in every direction in TX are feasible if we begin at a state x in the interior of the state space X . But if $x_i^p = 0$ for some strategy $i \in S^p$, then motion in any direction z with $z_i^p < 0$ would cause the mass of




Figure 2.3.10: Payoffs (= projected payoffs) in Rock-Paper-Scissors.

agents playing strategy i to become negative, taking the state out of X .

To describe the *feasible* displacement directions from an arbitrary state $x \in X$, we introduce the notion of a tangent cone. To begin, recall that the set $K \subseteq \mathbf{R}^n$ is a *cone* if whenever it contains the vector z , it also contains the vector αz for every $\alpha > 0$. Most often one is interested in convex cones (i.e., cones that are convex sets). The *polar* of the convex cone K is a new convex cone

$$K^\circ = \{y \in \mathbf{R}^n : y'z \leq 0 \text{ for all } z \in K\}.$$

In words, the polar cone of K contains all vectors that form a weakly obtuse angle with each vector in K (Figure 2.3.11).

Exercise 2.3.1. Let K be a convex cone. Show that

- (i) K° is a closed convex cone, and $K^\circ = (\text{cl}(K))^\circ$. (Hence, K° contains the origin.)
- (ii) K is a subspace of \mathbf{R}^n if and only if K is symmetric, in the sense that $K = -K$. Moreover, in this case, $K^\circ = K^\perp$.
- (iii) $(K^\circ)^\circ = \text{cl}(K)$. (Hint: To show that $(K^\circ)^\circ \subseteq \text{cl}(K)$, use the separating hyperplane theorem.)

The last result above tells us that $(K^\circ)^\circ = K$ for any closed convex cone K ; thus, polarity defines an *involution* on the set of closed convex cones.




Figure 2.3.11: A convex cone and its polar cone.

Another fundamental result about closed convex cones and their polar cones, the *Moreau Decomposition Theorem*, is not needed until later chapters. But as the preceding discussion provides the proper context to present this result, we do so in Appendix 2.A.4.

If $C \subseteq \mathbf{R}^n$ is a closed convex set, then the *tangent cone* of C at state $x \in C$, denoted $TC(x)$, is the closed convex cone

$$TC(x) = \text{cl}(\{z \in \mathbf{R}^n : z = \alpha(y - x) \text{ for some } y \in C \text{ and some } \alpha \geq 0\}).$$

If $C \subset \mathbf{R}^n$ is a *polytope* (i.e., the convex hull of a finite number of points), then the closure operation is redundant. In this case, $TC(x)$ is the set of directions of motion from x that initially remain in C ; more generally, $TC(x)$ also contains the limits of such directions. (To see the difference, construct $TC(x)$ for $x \in \text{bd}(C)$ when C is a square and when C is a circle.) If x is in the relative interior of C (i.e., the interior of C relative to $\text{aff}(C)$), then $TC(x)$ is just TC , the tangent space of C ; otherwise, $TC(x)$ is a strict subset of TC .

Finally, define the *normal cone* of C at x to be the polar of the tangent cone of C at x : that is, $NC(x) = (TC(x))^\circ$. By definition, $NC(x)$ is a closed convex cone, and it contains every vector that forms a weakly obtuse angle with every feasible displacement vector at x .

In Figures 2.3.12 and 2.3.13, we sketch examples of tangent cones and normal cones when X is the state space for a two strategy game (i.e., the simplex in \mathbf{R}^2) and for a three

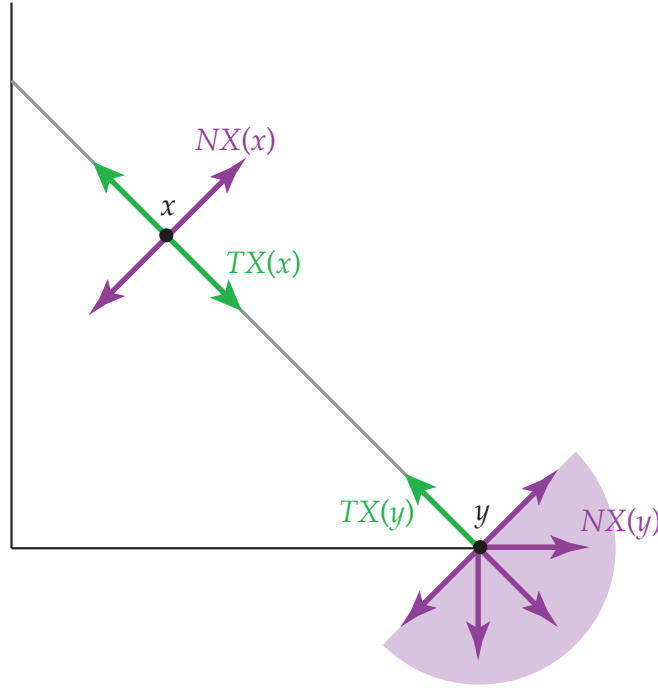


Figure 2.3.12: Tangent cones and normal cones for two-strategy games.

strategy game (the simplex in \mathbf{R}^3). Since the latter figure is two-dimensional, with the sheet of paper representing the affine hull of X , the figure actually displays the projected normal cones $\Phi(NX(x))$.

2.3.6 Normal Cones and Nash Equilibria

At first glance, normal cones might appear to be less relevant to game theory than tangent cones. Theorem 2.3.2 shows that this impression is false: normal cones and Nash equilibria are intimately linked.

Theorem 2.3.2. *Let F be a population game. Then $x \in NE(F)$ if and only if $F(x) \in NX(x)$.*

$$\begin{aligned}
 \text{Proof. } x \in NE(F) &\Leftrightarrow [x_i^p > 0 \Rightarrow F_i^p(x) \geq F_j^p(x)] \text{ for all } i, j \in S^p, p \in \mathcal{P} \\
 &\Leftrightarrow (x^p)'F^p(x) \geq (y^p)'F^p(x) \text{ for all } y^p \in X^p, p \in \mathcal{P} \\
 &\Leftrightarrow (y^p - x^p)'F^p(x) \leq 0 \text{ for all } y^p \in X^p, p \in \mathcal{P} \\
 &\Leftrightarrow (z^p)'F^p(x) \leq 0 \text{ for all } z^p \in TX^p(x^p), p \in \mathcal{P} \\
 &\Leftrightarrow F^p(x) \in NX^p(x^p) \text{ for all } p \in \mathcal{P} \\
 &\Leftrightarrow F(x) \in NX(x). \blacksquare
 \end{aligned}$$




Figure 2.3.13: Tangent cones and normal cones for three-strategy games.

Exercise 2.3.3. Justify the last equivalence above.

Theorem 2.3.2 tells us that state x is a Nash equilibrium if and only if the payoff vector $F(x)$ lies in the normal cone of the state space X at x . This result provides us with a simple, purely geometric description of Nash equilibria of population games. Its proof is very simple: some algebra shows that x is a Nash equilibrium if and only if it solves a *variational inequality problem*—that is, if it satisfies

$$(2.1) \quad (y - x)'F(x) \leq 0 \text{ for all } y \in X.$$

Applying the definitions of tangent and normal cones then yields the result.

In many cases, it is more convenient to speak in terms of projected payoff vectors and projected normal cones. Corollary 2.3.4 restates Theorem 2.3.2 in these terms.

Corollary 2.3.4. $x \in NE(F)$ if and only if $\Phi F(x) \in \Phi(NX(x))$.

Proof. Clearly, $F(x) \in NX(x)$ implies that $\Phi F(x) \in \Phi(NX(x))$. The reverse implication follows from the facts that $NX(x) = \Phi(NX(x)) + (TX)^\perp$ (see Exercise 2.3.5) and that $\Phi((TX)^\perp) = \{0\}$ (which is the equation that defines Φ as the *orthogonal* projection of \mathbf{R}^n onto TX). ■

Exercise 2.3.5. (i) Using the notions of relative and average payoffs discussed in Section 2.3.3, explain the intuition behind Corollary 2.3.4 in the single population case.

(ii) Prove that $NX(x) = \Phi(NX(x)) + (TX)^\perp$.

(iii) Only one of the two statements to follow is equivalent to $x \in NE(F) : F(x) \in \Phi(NX(x))$, or $\Phi F(x) \in NX(x)$. Which is it?

In Figures 2.3.7 through 2.3.10, we mark the Nash equilibria of our four population games with dots. In the two-strategy games F^{C2} and F^{HD} , the Nash equilibria are those states x at which the payoff vector $F(x)$ lies in the normal cone $NX(x)$, as Theorem 2.3.2 requires. In both these games and in the three-strategy games F^{C3} and F^{RPS} , the Nash equilibria are those states x at which the projected payoff vector $\Phi F(x)$ lies in the projected normal cone $\Phi(NX(x))$, as Corollary 2.3.4 demands. Even if the dots were not drawn, we could locate the Nash equilibria of all four games by examining the arrows alone.

Exercise 2.3.6. Compute the Nash equilibria of the four games studied above, and verify that the equilibria appear in the correct positions in Figures 2.3.7 through 2.3.10.

Exercise 2.3.7. Two-population two-strategy games. Let F be a game played by two unit mass populations ($p = 2$) with two strategies for each ($n^1 = n^2 = 2$).

(i) Describe the state space X , tangent space TX , and orthogonal projection Φ for this setting.

(ii) Show that the state space X can be represented on a sheet of paper by a unit square, with the upper left vertex representing the state at which all agents in both populations play strategy 1, and with the upper right vertex representing the state at which all agents in population 1 play strategy 1 and all agents in population 2 play strategy 2. Explain how the projected payoff vectors $\Phi F(x)$ can be represented as arrows in this diagram.

(iii) At (a) a point in the interior of the square, (b) a non-vertex boundary point, and (c) a vertex, draw the tangent cone $TX(x)$ and the projected normal cone $\Phi(NX(x))$, and give algebraic descriptions of each.

(iv) Suppose we draw projected payoff vectors $\Phi F(x)$ in the manner you described in part (ii) and projected normal cones in the manner you described in part (iii). Verify that in each of the cases considered in part (iii), the arrow representing $\Phi F(x)$ is contained in the sketch of $\Phi(NX(x))$ if and only if x is a Nash equilibrium of F .

Appendix

2.A Affine Spaces, Tangent Spaces, and Projections

The simplex in \mathbf{R}^n , the state space for single population games, is an $n - 1$ dimensional subset of \mathbf{R}^n ; state spaces for multipopulation games are Cartesian products of scalar multiples of simplices. For this reason, linear subspaces, affine spaces, and orthogonal projections all play important roles in the study of population games.

2.A.1 Affine Spaces

The set $Z \subseteq \mathbf{R}^n$ is a (*linear*) *subspace* of \mathbf{R}^n if it is closed under linear combination: if $z, \hat{z} \in Z$ and $a, b \in \mathbf{R}$, then $az + b\hat{z} \in Z$ as well. Suppose that Z is a subspace of \mathbf{R}^n of dimension $\dim(Z) < n$, and that the set A is a translation of Z by some vector $v \in \mathbf{R}^n$:

$$A = Z + \{v\} = \{x \in \mathbf{R}^n : x = z + v \text{ for some } z \in Z\}.$$

Then we say that A is an *affine space* of dimension $\dim(A) = \dim(Z)$.

Observe that any vector representing a direction of motion through A is itself an element of Z : if $x, y \in A$, then $y - x = (z^y + v) - (z^x + v) = z^y - z^x$ for some z^x and z^y in Z ; since Z is closed under linear combinations, $z^y - z^x \in Z$. For this reason, the set Z is called the *tangent space* of A , and we often write TA in place of Z .

Since the origin is an element of Z , the translation vector v in the definition $A = Z + \{v\}$ can be any element of A . But is there a “natural” choice of v ? Recall that the *orthogonal complement* of Z , denoted by Z^\perp , contains the vectors in \mathbf{R}^n orthogonal to all elements of Z : that is, $Z^\perp = \{v \in \mathbf{R}^n : v'z = 0 \text{ for all } z \in Z\}$. It is easy to show that the set $A \cap Z^\perp$ contains a single element, which we denote by z_A^\perp , and that this *orthogonal translation vector* is the closest point in Z^\perp to every point in A (in the language of Section 2.A.3 below, $P_{Z^\perp}x = z_A^\perp$ for all $x \in A$). We will see that for many purposes, this translation vector is the most convenient choice.

Example 2.A.1. Consider the subspace $\mathbf{R}_0^n = \{z \in \mathbf{R}^n : \mathbf{1}'z = 0\}$ and the affine space $A = \mathbf{R}_0^n + \{e_1\} = \{z \in \mathbf{R}^n : \mathbf{1}'z = 1\}$, where $\mathbf{1} = (1, \dots, 1)'$. Since $(\mathbf{R}_0^n)^\perp = \text{span}(\{\mathbf{1}\})$ and $A \cap \text{span}(\{\mathbf{1}\}) = \{\frac{1}{n}\mathbf{1}\}$, the vector $\frac{1}{n}\mathbf{1}$ is the orthogonal translation vector that generates A . In particular, $A = \mathbf{R}_0^n + \{\frac{1}{n}\mathbf{1}\}$, and $\frac{1}{n}\mathbf{1}$ is the closest point in $\text{span}(\{\mathbf{1}\})$ to every $x \in A$. We illustrate the case in which $n = 2$ in Figure 2.A.1; note again our convention of using the vertical axis to represent the first component of $x = (x_1, x_2)$. §




Figure 2.A.1: The state space and its affine hull for two-strategy games.

2.A.2 Affine Hulls of Convex Sets

Let $Y \subseteq \mathbf{R}^n$. The *affine hull* of Y , denoted $\text{aff}(Y)$, is the smallest affine space that contains Y . This set can be described as

$$(2.2) \quad \text{aff}(Y) = \left\{ x \in \mathbf{R}^n : x = \sum_{i=1}^k \lambda^i y^i \text{ for some } \{y^i\}_{i=1}^k \subset Y \text{ and } \{\lambda^i\}_{i=1}^k \subset \mathbf{R} \text{ with } \sum_{i=1}^k \lambda^i = 1 \right\}.$$

The vector x is called an *affine combination* of the vectors y^i . If we also required the λ^i to be nonnegative, x would instead be a *convex combination* of the y^i , and (2.2) would become $\text{conv}(Y)$, the *convex hull* of Y .

Now suppose that Y is itself convex, let $A = \text{aff}(Y)$ be its affine hull, and let $Z = TA$ be the tangent space of A ; then we also call $Z = TY$ the *tangent space* of Y , as Z contains directions of motion from points in the (relative) interior of Y that stay in Y . We also call $\dim(Y) = \dim(Z)$ the *dimension* of Y .

In constructing the affine hull of a convex set as in (2.2), it is enough to take affine combinations of a fixed set of $\dim(Y) + 1$ points in Y . To accomplish this, let $d = \dim(Y)$, fix $y^0 \in Y$ arbitrarily, and choose y^1, \dots, y^d so that $\{y^1 - y^0, \dots, y^d - y^0\}$ is a basis for Z . Then letting $\lambda^0 = 1 - \sum_{i=1}^d \lambda^i$, we see that

$$Z + \{y^0\} = \text{span}(\{y^1 - y^0, \dots, y^d - y^0\}) + \{y^0\}$$

$$\begin{aligned}
&= \left\{ x \in \mathbf{R}^n : x = \sum_{i=1}^d \lambda^i (y^i - y^0) + y^0 \text{ for some } \{\lambda^i\}_{i=1}^d \subset \mathbf{R} \right\}. \\
&= \left\{ x \in \mathbf{R}^n : x = \sum_{i=0}^d \lambda^i y^i \text{ for some } \{\lambda^i\}_{i=0}^d \subset \mathbf{R} \text{ with } \sum_{i=0}^d \lambda^i = 1 \right\} \\
&= \text{aff}(Y).
\end{aligned}$$

Example 2.A.2. Population states. Let $X^p = \{x^p \in \mathbf{R}_+^{n^p} : \mathbf{1}'x^p = m^p\}$ be the set of population states for a population of mass m^p . This convex set has affine hull $\text{aff}(X^p) = \{x^p \in \mathbf{R}^{n^p} : \mathbf{1}'x^p = m^p\}$ and tangent space $TX^p = \{z^p \in \mathbf{R}^{n^p} : \mathbf{1}'z^p = 0\} = \mathbf{R}_0^{n^p}$ (cf Example 2.A.1). §

Example 2.A.3. Social states. Let $X = \prod_{p \in \mathcal{P}} X^p$ be the set of social states for a collection of populations $\mathcal{P} = \{1, \dots, p\}$ with masses m^1, \dots, m^p . This convex set has affine hull $\text{aff}(X) = \prod_{p \in \mathcal{P}} \text{aff}(X^p)$ and tangent space $TX = \prod_{p \in \mathcal{P}} \mathbf{R}_0^{n^p}$. Thus, if $z = (z^1, \dots, z^p) \in TX$, then each z^p has components that sum to zero. §

2.A.3 Orthogonal Projections

If V and W are subspaces of \mathbf{R}^n , their sum is $V + W = \text{span}(V \cup W)$, the set of linear combinations of elements of V and W . If $V \cap W = \{\mathbf{0}\}$, every $x \in V + W$ has a unique decomposition $x = v + w$ with $v \in V$ and $w \in W$. In this case, we write $V + W$ as $V \oplus W$, and call it the *direct sum* of V and W . For instance, $V \oplus V^\perp = \mathbf{R}^n$ for any subspace $V \subseteq \mathbf{R}^n$.

Every matrix $A \in \mathbf{R}^{n \times n}$ defines a linear operator from \mathbf{R}^n to itself via $x \mapsto Ax$. To understand the action of this operator, remember that the i th column of A is the image of the standard basis vector e_i , and, more generally, that Ax is a linear combination of the columns of A .

We call the linear operator $P \in \mathbf{R}^{n \times n}$ a *projection* onto the subspace $V \subseteq \mathbf{R}^n$ if there is a second subspace $W \subseteq \mathbf{R}^n$ satisfying $V \cap W = \{\mathbf{0}\}$ and $V \oplus W = \mathbf{R}^n$ such that

- (i) $Px = x$ for all $x \in V$, and
- (ii) $Px = \mathbf{0}$ for all $x \in W$.

If $W = V^\perp$, we call P the *orthogonal projection* onto V , and write P_V in place of P .

Every projection onto V maps all points in \mathbf{R}^n to points in V . While for any given subspace V there are many projections onto V , the orthogonal projection onto V is unique. For example,

$$P_1 = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} \text{ and } P_2 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

both define projections of \mathbf{R}^2 onto the horizontal axis $\{x \in \mathbf{R}^2 : x_1 = 0\}$. (Recall again our convention of representing x_1 on the vertical axis.) However, since P_2 maps the vertical

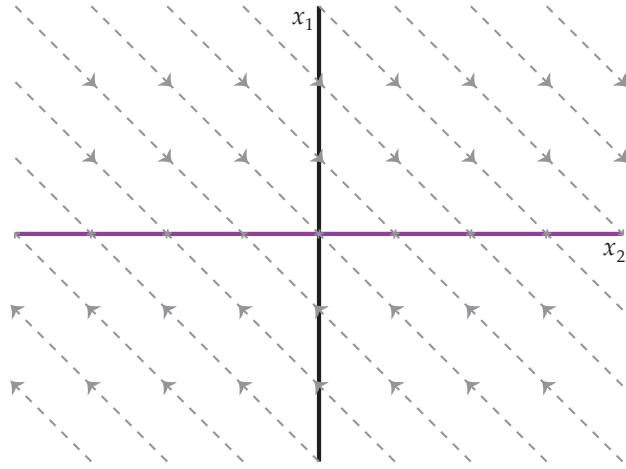


Figure 2.A.2: A projection.




Figure 2.A.3: An orthogonal projection.

axis $\{x \in \mathbf{R}^2 : x_2 = 0\}$ to the origin, it is the orthogonal projection. The action of the two projections is illustrated in Figures 2.A.2 and 2.A.3 below. The latter figure illustrates a geometrically obvious property of orthogonal projections: the orthogonal projection of \mathbf{R}^n onto V maps each point $y \in \mathbf{R}^n$ to the closest point to y in V :

$$P_V y = \operatorname{argmin}_{v \in V} |y - v|^2.$$

Projections admit simple algebraic characterizations. Recall that the matrix $A \in \mathbf{R}^{n \times n}$ is *idempotent* if $A^2 = A$. It is easy to see that projections are represented by idempotent matrices: once the first application of P projects \mathbf{R}^n onto the subspace V , the second application of P does nothing more. In fact, we have

Theorem 2.A.4. (i) P is a projection if and only if P is idempotent.




Figure 2.A.4: The orthogonal projection Φ in \mathbf{R}^2 .

(ii) P is an orthogonal projection if and only if P is symmetric idempotent.

Example 2.A.5. The orthogonal projection onto $\mathbf{R}_0^{n^p}$. In Example 2.A.2, we saw that the set of population states $X^p = \{x^p \in \mathbf{R}_+^{n^p} : \mathbf{1}'x^p = m^p\}$ has tangent space $TX^p = \mathbf{R}_0^{n^p} = \{z^p \in \mathbf{R}^{n^p} : \mathbf{1}'z^p = 0\}$. We can decompose the space \mathbf{R}^{n^p} into the direct sum $\mathbf{R}_0^{n^p} \oplus \mathbf{R}_1^{n^p}$, where $\mathbf{R}_1^{n^p} = (\mathbf{R}_0^{n^p})^\perp = \text{span}(\{\mathbf{1}\})$. The orthogonal projection of \mathbf{R}^{n^p} onto $\mathbf{R}_1^{n^p}$ is $\Xi = \frac{1}{n^p}\mathbf{1}\mathbf{1}'$, the matrix whose entries all equal $\frac{1}{n^p}$; to verify this, note that $\Xi z^p = \mathbf{0}$ for $z^p \in \mathbf{R}_0^{n^p}$ and $\Xi\mathbf{1} = \mathbf{1}$. The orthogonal projection of \mathbf{R}^{n^p} onto $\mathbf{R}_0^{n^p}$ is $\Phi = I - \Xi$, since $\Phi z^p = z^p$ for $z^p \in \mathbf{R}_0^{n^p}$ and $\Phi\mathbf{1} = \mathbf{0}$ (see Figure 2.A.4 for the case of $n^p = 2$). Both Ξ and Φ are clearly symmetric, and since $\Xi^2 = (\frac{1}{n^p}\mathbf{1}\mathbf{1}')(\frac{1}{n^p}\mathbf{1}\mathbf{1}') = \frac{1}{n^p}\mathbf{1}\mathbf{1}' = \Xi$ and $\Phi^2 = (I - \Xi)(I - \Xi) = I - 2\Xi + \Xi^2 = I - 2\Xi + \Xi = I - \Xi = \Phi$, both are idempotent as well. §

More generally, it is easy to show that if P is the orthogonal projection of \mathbf{R}^n onto V , then $I - P$ is the orthogonal projection of \mathbf{R}^n onto V^\perp . Or, in the notation introduced above, $P_{V^\perp} = I - P_V$.

Example 2.A.6. The orthogonal projection onto TX . Recall from Example 2.A.3 that the set of social states $X = \prod_{p \in \mathcal{P}} X^p$ has tangent space $TX = \prod_{p \in \mathcal{P}} \mathbf{R}_0^{n^p}$. We can decompose \mathbf{R}^n into the direct sum $\prod_{p \in \mathcal{P}} \mathbf{R}_0^{n^p} \oplus \prod_{p \in \mathcal{P}} \mathbf{R}_1^{n^p} = TX \oplus \prod_{p \in \mathcal{P}} \text{span}(\{\mathbf{1}\})$. The orthogonal projection of \mathbf{R}^n onto $\prod_{p \in \mathcal{P}} \text{span}(\{\mathbf{1}\})$ is the block diagonal matrix $\Xi = \text{diag}(\Xi, \dots, \Xi)$, while the orthogonal projection of \mathbf{R}^n onto TX is $\Phi = I - \Xi = \text{diag}(\Phi, \dots, \Phi)$. Of course, Ξ and Φ are both symmetric idempotent. §

Example 2.A.7. Ordinary least squares. Suppose we have a collection of $n > k$ data points, $\{(x^i, y^i)\}_{i=1}^n$, where each $x^i \in \mathbf{R}^k$ contains k components of “explanatory” data and each $y^i \in \mathbf{R}$ is the corresponding component of “explainable” data. We write these data as

$$X = \begin{pmatrix} (x^1)' \\ \vdots \\ (x^n)' \end{pmatrix} \in \mathbf{R}^{n \times k} \text{ and } y = \begin{pmatrix} y^1 \\ \vdots \\ y^n \end{pmatrix} \in \mathbf{R}^n$$

and assume that X is of full rank. We seek the *best linear predictor*: the map $x \mapsto x'\beta$ that minimizes the sum of squared prediction errors $\sum_{i=1}^n (y^i - (x^i)'\beta)^2 = |y - X\beta|^2$.

(The prediction function $x^i \mapsto (x^i)'\beta$ is a truly linear function of x , in the sense that the input vector $\mathbf{0}$ generates a prediction of 0. Typically, one seeks an affine prediction function—that is, one that allows for a nonzero constant term. To accomplish this, one sets $x_1^i = 1$ for all i , leaving only $k - 1$ components of true explanatory data. In this case, the component β_1 serves as a constant term in the affine prediction function $(x_2, \dots, x_k) \mapsto \beta_1 + \sum_{i=2}^k x_i \beta_i$.)

Let $\text{span}(X) = \{Xb : b \in \mathbf{R}^k\}$ be the column span of X . That $\beta \in \mathbf{R}^k$ minimizes $|y - X\beta|^2$ is equivalent to the requirement that $X\beta$ be the closest point to y in the column span of X :

$$X\beta = \operatorname{argmin}_{v \in \text{span}(X)} |y - v|^2.$$

Both calculus and geometry tell us that for this to be true, the vector of prediction errors $y - X\beta$ must be orthogonal to $\text{span}(X)$, and hence to each column of X .

$$X'(y - X\beta) = \mathbf{0}.$$

One can verify that $X \in \mathbf{R}^{n \times k}$ and $X'X \in \mathbf{R}^{k \times k}$ have the same null space, and hence the same (full) rank. Therefore, $(X'X)^{-1}$ exists, and we can solve the previous equation for β :

$$\beta = (X'X)^{-1}X'y.$$

To this point, we have taken $X \in \mathbf{R}^{n \times k}$ and $y \in \mathbf{R}^n$ as given and used them to find the vector $\beta \in \mathbf{R}^k$, which we have viewed as defining a map from vectors of explanatory data $x \in \mathbf{R}^k$ to predictions $x'\beta \in \mathbf{R}$. Now, let us take X alone as given and consider the map from vectors of “explainable” data $y \in \mathbf{R}^n$ to vectors of predictions $X\beta = X(X'X)^{-1}X'y \in \mathbf{R}^n$. By construction, this linear map $P = X(X'X)^{-1}X' \in \mathbf{R}^{n \times n}$ is the orthogonal projection of \mathbf{R}^n onto $\text{span}(X)$. P is clearly symmetric (since the inverse of a symmetric matrix is symmetric),

and since $P^2 = X(X'X)^{-1}X'X(X'X)^{-1}X' = X(X'X)^{-1}X' = P$ it is idempotent as well. §

2.A.4 The Moreau Decomposition Theorem

A basic fact about projection onto subspaces holds that for any vector $v \in \mathbf{R}^n$ and any subspace $Z \subseteq \mathbf{R}^n$, the sum $v = P_Z(v) + P_{Z^\perp}(v)$ is the unique decomposition of v into the sum of elements of Z and Z^\perp . The *Moreau Decomposition Theorem* is a generalization of this result that replaces the subspace Z and its orthogonal complement with a closed convex cone and its polar cone. We use this theorem repeatedly in Chapter 6 in our analysis of the projection dynamic.

To state this result, we need an appropriate analogue of orthogonal projection for the context of closed, convex sets. To this end, we define $\Pi_C : \mathbf{R}^n \rightarrow C$, the (*closest point*) *projection* of \mathbf{R}^n onto the closed convex set C by

$$\Pi_C(y) = \operatorname{argmin}_{x \in C} |y - x|.$$

This definition generalizes that of the projection P_Z onto the subspace $Z \subseteq \mathbf{R}^n$ to cases in which the target set is not linear, but merely closed and convex. With this definition in hand, we can state our new decomposition theorem; an illustration is provided in Figure 2.A.5.

Theorem 2.A.8 (The Moreau Decomposition Theorem). *Let $K \subseteq \mathbf{R}^n$ and $K^\circ \subseteq \mathbf{R}^n$ be a closed convex cone and its polar cone, and let $v \in \mathbf{R}^n$. Then the following are equivalent:*

- (i) $v_K = \Pi_K(v)$ and $v_{K^\circ} = \Pi_{K^\circ}(v)$.
- (ii) $v_K \in K$, $v_{K^\circ} \in K^\circ$, $v = v_K + v_{K^\circ}$, and $v_K'v_{K^\circ} = 0$.

2.N Notes

Congestion games are introduced in Beckmann et al. (1956); see the notes to Chapter 3 for further references. For the biological motivation for the Hawk-Dove game, see Maynard Smith (1982, Chapter 2).

Portions of Section 2.3 follow Lahkar and Sandholm (2008). The link between normal cones and Nash equilibria is known from the literature on variational inequalities; see Harker and Pang (1990) and Nagurney (1999). For more on affine spaces, tangent cones, normal cones, the Moreau Decomposition Theorem, and related notions, see Hiriart-Urruty and Lemaréchal (2001). The algebra of orthogonal projections is explained, e.g., in Friedberg et al. (1989, Section 6.6).




Figure 2.A.5: The Moreau Decomposition Theorem.

Potential Games, Stable Games, and Supermodular Games

3.0 Introduction

In the previous chapter, we offered a general definition of population games and characterized their Nash equilibria geometrically. Still, since any continuous map F from the state space X to \mathbf{R}^n defines a population game, population games with even a moderate number of strategies can be difficult to analyze. In this chapter, we define three important classes of population games: *potential games*, *stable games*, and *supermodular games*. From an economic point of view, each definition places constraints on the sorts of externalities agents impose on one another through their choices in the game. From a mathematical point of view, each definition imposes a structure on payoff functions that renders their analysis relatively simple.

We show through examples that potential games, stable games, and supermodular games each encompass a variety of interesting applications. We also establish the basic properties of each class of games. Among other things, we show that for games in each class, existence of Nash equilibrium can be proved using elementary methods. Beginning in Chapter 7, we investigate the behavior of evolutionary dynamics in the three classes of games; there, our assumptions on the structure of externalities will allow us to establish a range of global convergence results.

The definitions of our three classes of games only require continuity of the payoff functions. If we instead make the stronger assumption that payoffs are smooth (in particular, continuously differentiable), we can avail ourselves of the tools of calculus. Doing so not only simplifies computations, but also allows us to express our definitions and results in simple, useful, and intuitively appealing ways. The techniques from calculus that we

require are reviewed in the Appendices 3.A and 3.B.

3.1 Full Potential Games

In potential games, all information about payoffs that is relevant to agents' incentives can be captured in a single scalar-valued function. The existence of this function—the game's potential function—underlies potential games' many attractive properties. In this section, we consider full potential games, which can be analyzed using standard multivariate calculus techniques (Appendix 3.A), but at the expense of requiring an extension of the payoff functions' domain. In Section 3.2, we introduce a definition of potential games that does not use this device, but that instead requires analyses that rely on affine calculus (Appendix 3.B).

Potential games were first defined in the context of normal form games—see the Notes and Exercise 3.2.10. It is possible combine the normal form and population approaches by defining potential games for large but finite populations—see Section 11.5.1.

3.1.1 Full Population Games

To understand the issues alluded to above, consider a game F played by a single population of agents. Since population states for this game are elements of $X = \{x \in \mathbf{R}_+^n : \sum_{k \in S} x_k = 1\}$, the simplex in \mathbf{R}^n , the payoff F_i to strategy i is a real-valued function with domain X .

In looking for useful properties of population games, a seemingly natural characteristic to consider is the marginal effect of adding new agents playing strategy j on the payoffs of agents currently choosing strategy i . This effect is captured by the partial derivative $\frac{\partial F_i}{\partial x_j}$. But herein lies the difficulty: if F_i is only defined on the simplex, then even if the function F is differentiable, the partial derivative $\frac{\partial F_i}{\partial x_j}$ does not exist.

To ensure that partial derivatives exist, we extend the domain of the game F from the state space $X = \{x \in \mathbf{R}^n : \sum_{k \in S} x_k = 1\}$ to the entire positive orthant \mathbf{R}_+^n . In multipopulation settings, the analogous extension is from the original set of social states $X = \{x = (x^1, \dots, x^p) \in \mathbf{R}_+^n : \sum_{i \in S^p} x_i^p = m^p\}$ to \mathbf{R}_+^n . In either setting, we call the game with payoffs defined on the positive orthant a *full population game*. In many interesting cases, one can interpret the extensions of payoffs as specifying the values that payoffs would take were the population sizes to change—see Section 3.1.3.

3.1.2 Definition and Characterization

With these preliminaries addressed, we are now prepared to define full potential games.

Definition. Let $F : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ be a full population game. We call F a full potential game if there exists a continuously differentiable function $f : \mathbf{R}_+^n \rightarrow \mathbf{R}$ satisfying

$$(3.1) \quad \nabla f(x) = F(x) \text{ for all } x \in \mathbf{R}_+^n.$$

Property (3.1) can be stated more explicitly as

$$\frac{\partial f}{\partial x_i^p}(x) = F_i^p(x) \text{ for all } p \in \mathcal{P}, i \in S^p \text{ and } x \in \mathbf{R}_+^n.$$

The function f , which is unique up to the addition of a constant, is called the *full potential function* for the game F . It represents the game's payoffs in an integrated form.

To explain the potential function's role, suppose that $x \in X$ is a population state at which $F_j^p(x) > F_i^p(x)$, so that an agent choosing strategy $i \in S^p$ would be better off choosing strategy $j \in S^p$. Now suppose some small group of agents switch from strategy i to strategy j . These switches are represented by the displacement vector $z = e_j^p - e_i^p$, where e_i^p is the (i, p) th standard basis vector in \mathbf{R}^n . The marginal impact that these switches have on the value of potential is therefore

$$\frac{\partial f}{\partial z}(x) = \nabla f(x)'z = \frac{\partial f}{\partial x_j^p}(x) - \frac{\partial f}{\partial x_i^p}(x) = F_j^p(x) - F_i^p(x) > 0.$$

In other words, profitable strategy revisions increase potential. More generally, we will see in later chapters that the "uphill" directions of the potential function include all directions in which reasonable adjustment processes might lead. This fact underlies the many attractive properties that potential games possess.

If the map $F : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ is C^1 (continuously differentiable), it is well known that F admits a potential function if and only if its derivative matrices $DF(x)$ are symmetric (see Appendix 3.A.9). In the current game-theoretic context, we call this condition *full externality symmetry*.

Observation 3.1.1. Suppose the population game F is C^1 . Then F is a full potential game if and only if it satisfies full externality symmetry:

$$(3.2) \quad DF(x) \text{ is symmetric for all } x \in \mathbf{R}_+^n$$

More explicitly, F is a potential game if and only if

$$(3.3) \quad \frac{\partial F_i^p}{\partial x_j^q}(x) = \frac{\partial F_j^q}{\partial x_i^p}(x) \text{ for all } i \in S^p, j \in S^q, p, q \in \mathcal{P}, \text{ and } x \in \mathbf{R}_+^n.$$

Observation 3.1.1 characterizes smooth full potential games in terms of a simple, economically meaningful property: condition (3.2) requires that the effect on the payoffs to strategy $i \in S^p$ of introducing new agents choosing strategy $j \in S^q$ always equals the effect on the payoffs to strategy j of introducing new agents choosing strategy i .

3.1.3 Examples

Our first two examples build on ones studied in Chapter 2.

Example 3.1.2. Matching in normal form games with common interests. Suppose that agents in a single population are matched to play symmetric the two player normal form game $A \in \mathbf{R}^{n \times n}$, generating the population game $F(x) = Ax$. While earlier we used this formula to define F on the state space X , here we will use it to define F on all of \mathbf{R}_+^n . (While this choice works very well in the present example, it is not always innocuous, as will see in Section 3.2.)

The symmetric normal form game A has *common interests* if both players always receive the same payoff. This means that $A_{ij} = A_{ji}$ for all i and j , or, equivalently, that the matrix A is symmetric. Since $DF(x) = A$, this is precisely what we need for F to be a full potential game. The full potential function for F is

$$f(x) = \frac{1}{2}x'Ax,$$

which is one-half of $x'Ax = \sum_{i \in S} x_i F_i(x) = \bar{F}(x)$, the aggregate payoff function for F .

To cover the multipopulation case, call the normal form game $U = (U^1, \dots, U^p)$ a *common interest game* if there is a function $V : S \rightarrow \mathbf{R}$ such that $U^p(s) = V(s)$ for all $s \in S$ and $p \in \mathcal{P}$. As before, this means that under any pure strategy profile, all p players earn the same payoff. This normal form game generates the full population game

$$F_{S^p}^p(x) = \sum_{s^p \in S^p} V(s) \prod_{r \neq p} x_{s^r}^r$$

on \mathbf{R}_+^n . Aggregate payoffs in F are given by

$$\bar{F}(x) = \sum_{p \in \mathcal{P}} \sum_{s^p \in S^p} x_{s^p}^p F_{s^p}^p(x) = p \sum_{s \in S} V(s) \prod_{r \in \mathcal{P}} x_{s^r}^r.$$

Hence, if we let $f(x) = \sum_{s \in S} V(s) \prod_{r \in \mathcal{P}} x_{s^r}^r = \frac{1}{p} \bar{F}(x)$, we obtain

$$\frac{\partial f}{\partial x_{s^p}^p}(x) = \sum_{s^{-p} \in S^{-p}} V(s) \prod_{r \neq p} x_{s^r}^r = F_{s^p}^p(x).$$

So once again, matching in a common interest game generates a full potential game in which potential is proportional to aggregate payoffs. §

Exercise 3.1.3. In the multipopulation case, check directly that condition (3.2) holds.

Example 3.1.4. Congestion games. For ease of exposition, suppose that the congestion game F models behavior in a traffic network. In this environment, an agent taking path $j \in S^q$ affects the payoffs of agents choosing path $i \in S^p$ through the marginal increases in congestion on the links $\phi \in \Phi_i^p \cap \Phi_j^q$ that the two paths have in common. But the marginal effect of an agent taking path i on the payoffs of agents choosing path j is identical:

$$\frac{\partial F_i^p}{\partial x_j^q}(x) = - \sum_{\phi \in \Phi_i^p \cap \Phi_j^q} c'_\phi(u_\phi(x)) = \frac{\partial F_j^q}{\partial x_i^p}(x).$$

In other words, congestion games satisfy condition (3.2), and so are full potential games.

The full potential function for the congestion game F can be written explicitly as

$$f(x) = - \sum_{\phi \in \Phi} \int_0^{u_\phi(x)} c_\phi(z) dz.$$

Hence, potential is typically unrelated to aggregate payoffs, which are given by

$$\bar{F}(x) = \sum_{p \in \mathcal{P}} \sum_{i \in S^p} x_i^p F_i^p(x) = - \sum_{\phi \in \Phi} u_\phi(x) c_\phi(u_\phi(x)).$$

In Section 3.1.6, we offer conditions under which potential and aggregate payoffs are directly linked. §

Example 3.1.5. Cournot competition. Consider a unit mass population of firms who choose production quantities from the set $S = \{1, \dots, n\}$. The firms' aggregate production is given

by $a(x) = \sum_{i \in S} i x_i$. Let $p : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ denote inverse demand, a decreasing function of aggregate production. Let the firms' production cost function $c : S \rightarrow \mathbf{R}$ be arbitrary. Then the payoff to a firm producing quantity $i \in S$ at population state $x \in X$ is $F_i(x) = i p(a(x)) - c(i)$.

It is easy to check that F is a full potential game with full potential function

$$f(x) = \int_0^{a(x)} p(z) dz - \sum_{i \in S} x_i c(i).$$

In contrast, aggregate payoffs in F are

$$\bar{F}(x) = \sum_{i \in S} x_i F_i(x) = a(x)p(a(x)) - \sum_{i \in S} x_i c(i).$$

The difference between the two is

$$f(x) - \bar{F}(x) = \int_0^{a(x)} (p(z) - p(a(x))) dz,$$

which is simply consumers' surplus. Thus, the full potential function $f = \bar{F} + (f - \bar{F})$ measures the total surplus received by firms and consumers. (Total surplus differs from aggregate payoffs because the latter ignores consumers, who are not modeled as active agents.) §

Example 3.1.6. Games generated by variable externality pricing schemes. Population games can be viewed as models of externalities for environments with many agents. One way to force agents to internalize the externalities they impose upon others is to introduce pricing schemes. Given an arbitrary full population game F with aggregate payoff function \bar{F} , define an augmented game \tilde{F} as follows:

$$\tilde{F}_i^p(x) = F_i^p(x) + \sum_{q \in \mathcal{P}} \sum_{j \in S^q} x_j^q \frac{\partial F_j^q}{\partial x_i^p}(x).$$

The double sum represents the marginal effect that an agent choosing strategy i has on other agents' payoffs.

Suppose that when the game F is played, a social planner charges each agent choosing strategy i a tax equal to this double sum, and that each agent's payoff function is separable in this tax. The population game generated by this intervention is \tilde{F} .

Now observe that

$$(3.4) \quad \frac{\partial \bar{F}}{\partial x_i^p}(x) = \frac{\partial}{\partial x_i^p} \sum_{q \in \mathcal{P}} \sum_{j \in S^q} x_j^q F_j^q(x) = F_i^p(x) + \sum_{q \in \mathcal{P}} \sum_{j \in S^q} x_j^q \frac{\partial F_j^q}{\partial x_i^p}(x) = \tilde{F}_i^p(x).$$

Equation (3.4) tells us that the augmented game \tilde{F} is a full potential game, and that its full potential function is the aggregate payoff function of the original game F . Hence, changes in strategy which are profitable in the *augmented* game increase efficiency with respect to the payoffs of the *original* game. §

3.1.4 Nash Equilibria of Full Potential Games

We saw in Section 3.1.2 that in full potential games, profitable strategy revisions increase potential. It is therefore natural to expect that Nash equilibria of full potential games are related to local maximizers of potential. To investigate this idea, consider the nonlinear program

$$\begin{aligned} \max f(x) \quad \text{subject to} \quad & \sum_{i \in S^p} x_i^p = m^p \quad \text{for all } p \in \mathcal{P}, \text{ and} \\ & x_i^p \geq 0 \quad \text{for all } i \in S^p \text{ and } p \in \mathcal{P}. \end{aligned}$$

The Lagrangian for this maximization problem is

$$L(x, \mu, \lambda) = f(x) + \sum_{p \in \mathcal{P}} \mu^p \left(m^p - \sum_{i \in S^p} x_i^p \right) + \sum_{p \in \mathcal{P}} \sum_{i \in S^p} \lambda_i^p x_i^p,$$

so the Kuhn-Tucker first order necessary conditions for maximization are

$$(3.5) \quad \frac{\partial f}{\partial x_i^p}(x) = \mu^p - \lambda_i^p \quad \text{for all } i \in S^p \text{ and } p \in \mathcal{P},$$

$$(3.6) \quad \lambda_i^p x_i^p = 0, \quad \text{for all } i \in S^p \text{ and } p \in \mathcal{P}, \text{ and}$$

$$(3.7) \quad \lambda_i^p \geq 0 \quad \text{for all } i \in S^p \text{ and } p \in \mathcal{P}.$$

Let

$$KT(f) = \{x \in X : (x, \mu, \lambda) \text{ satisfies (3.5)-(3.7) for some } \lambda \in \mathbf{R}^n \text{ and } \mu \in \mathbf{R}^p\}.$$

Theorem 3.1.7 shows that the Kuhn-Tucker first order conditions for maximizing f on X

characterize the Nash equilibria of F .

Theorem 3.1.7. *If F is a full potential game with full potential function f , then $NE(F) = KT(f)$.*

Proof. If x is a Nash equilibrium of F , then since $F = \nabla f$, the Kuhn-Tucker conditions are satisfied by x , $\mu^p = \max_{j \in S^p} F_j^p(x)$, and $\lambda_i^p = \mu^p - F_i^p(x)$. Conversely, if (x, μ, λ) satisfies the Kuhn-Tucker conditions, then for every $p \in \mathcal{P}$, (3.5) and (3.6) imply that $F_i^p(x) = \frac{\partial f}{\partial x_i^p}(x) = \mu^p$ for all i in the support of x^p . Furthermore, (3.5) and (3.7) imply that $F_j^p(x) = \mu^p - \lambda_j^p \leq \mu^p$ for all $j \in S^p$. Hence, the support of x^p is a subset of $\operatorname{argmax}_{j \in S^p} F_j^p(x)$, and so x is a Nash equilibrium of F . ■

Note that the multiplier μ^p represents the equilibrium payoff in population p , and that the multiplier λ_i^p represents the “payoff slack” of strategy $i \in S^p$.

Since the set X satisfies constraint qualification, satisfaction of the Kuhn-Tucker conditions is necessary for local maximization of the full potential function. Thus, Theorem 3.1.7, along with the fact that a continuous function on a compact set achieves its maximum, gives us a simple proof of existence of Nash equilibrium in full potential games.

On the other hand, the Kuhn-Tucker conditions are not sufficient for maximizing potential. Therefore, while all local maximizers of potential are Nash equilibria, not all Nash equilibria locally maximize potential.

Example 3.1.8. Consider again the 123 Coordination game introduced in Chapter 2:

$$F(x) = \begin{pmatrix} F_1(x) \\ F_2(x) \\ F_3(x) \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix}.$$

The full potential function for this game is the convex function $f(x) = \frac{1}{2}(x_1)^2 + (x_2)^2 + \frac{3}{2}(x_3)^2$. The three pure states, $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$, and $e_3 = (0, 0, 1)$, all locally maximize potential, and so are Nash equilibria. To focus on one instance, note that the Kuhn-Tucker conditions are satisfied at state e_1 by the multipliers $\mu = 1$, $\lambda_1 = 0$, and $\lambda_2 = \lambda_3 = 1$. The global minimizer of potential, $(\frac{6}{11}, \frac{3}{11}, \frac{2}{11})$, is a state at which payoffs to all three strategies are equal, and is therefore a Nash equilibrium as well; the Kuhn-Tucker conditions are satisfied here with multipliers $\mu = \frac{6}{11}$ and $\lambda_1 = \lambda_2 = \lambda_3 = 0$. Finally, at each of the boundary states $(\frac{2}{3}, \frac{1}{3}, 0)$, $(\frac{3}{4}, 0, \frac{1}{4})$, and $(0, \frac{3}{5}, \frac{2}{5})$, the strategies which are played receive equal payoffs, which exceed the payoff accruing to the unused strategy; thus, these states are Nash equilibria as well. These states, coupled with the appropriate multipliers, also satisfy the Kuhn-Tucker conditions: for example, $x = (\frac{2}{3}, \frac{1}{3}, 0)$ satisfies the conditions with $\mu = \frac{2}{3}$, $\lambda_1 = \lambda_2 = 0$ and $\lambda_3 = \frac{2}{3}$. This exhausts the set of Nash equilibria of F .





Figure 3.1.1: Graph (i) and contour plot (ii) of the potential function of 123 Coordination.

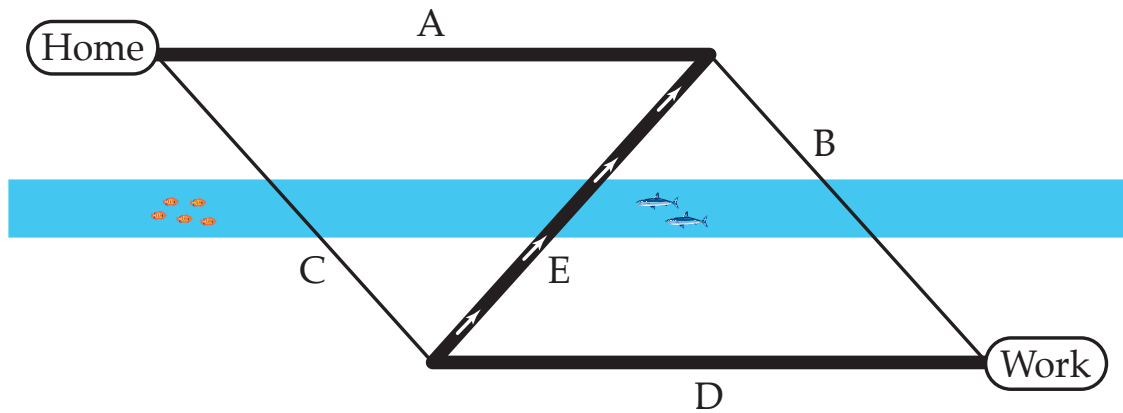


Figure 3.1.2: A highway network.

Figure 3.1.1 contains a graph and a contour plot of the full potential function f , and show the connection between this function and the Nash equilibria of F . §

The previous example demonstrates that in general, potential games can possess Nash equilibria that do not maximize potential. But if the full potential function f is concave, the Kuhn-Tucker conditions are not only necessary for maximizing f ; they are also sufficient. This fact gives us the following corollary to Theorem 3.1.7.

- Corollary 3.1.9.** (i) *If f is concave on X , then $NE(F)$ is the convex set of maximizers of f on X .*
(ii) *If f is strictly concave on X , then $NE(F)$ is a singleton containing the unique maximizer of f on X .*

Example 3.1.10. A network of highways connects Home and Work. The two towns are separated by a river. Highways A and D are expressways that go around bends in the river, and that do not become congested easily: $c_A(u) = c_D(u) = 4 + 20u$. Highways B and C cross the river over two short but easily congested bridges: $c_B(u) = c_C(u) = 2 + 30u^2$. In order to create a direct path between the towns, a city planner considers building a new expressway E that includes a third bridge over the river. Delays on this new expressway are described by $c_E(u) = 1 + 20u$. The highway network as a whole is pictured in Figure 3.1.2.

Before link E is constructed, there are two paths from Home to Work: path 1 traverses links A and B, while path 2 traverses links C and D. The equilibrium driving pattern splits the drivers equally over the two paths, yielding an equilibrium driving time (= –equilibrium payoff) of 23.5 on each.

After link E is constructed, drivers may also take path 3, which uses links C , E , and B . (We assume that traffic on link E only flows to the right.) The resulting population game has payoff functions

$$F(x) = \begin{pmatrix} F_1(x) \\ F_2(x) \\ F_3(x) \end{pmatrix} = \begin{pmatrix} -(6 + 20x_1 + 30(x_1 + x_3)^2) \\ -(6 + 20x_2 + 30(x_2 + x_3)^2) \\ -(5 + 20x_3 + 30(x_1 + x_3)^2 + 30(x_2 + x_3)^2) \end{pmatrix}$$

and full potential function

$$f(x) = -\left(6x_1 + 6x_2 + 5x_3 + 10((x_1)^2 + (x_2)^2 + (x_3)^2 + (x_1 + x_3)^3 + (x_2 + x_3)^3)\right).$$

Figure 3.1.3 contains a graph and a contour plot of the full potential function. Note that the full potential function for the two-path game is the restriction of f to the states at which $x_3 = 0$.

Evidently, the full potential function f is concave. (This is no coincidence—see Exercise 3.1.11 below.) The unique maximizer of potential on X , the state $x \approx (.4616, .4616, .0768)$, is the unique Nash equilibrium of the game. In this equilibrium, the driving time on each path is approximately 23.93, which exceeds the original equilibrium time of 23.5. In other words, adding an additional link to the network actually *increases* equilibrium driving times—a phenomenon known as *Braess' paradox*.

The intuition behind this phenomenon is easy to see. By opening up the new link E , we make it possible for a single driver on path 3 to use both of the easily congested bridges, B and C . But while using path 3 is bad for the population as a whole, it is appealing to individual drivers, as drivers do not account for the negative externalities their use of the bridges imposes on others. §





Figure 3.1.3: Graph (i) and contour plot (ii) of the potential function of a congestion game.

- Exercise 3.1.11. Uniqueness of equilibrium in congestion games.* (i) Let F be a congestion game with cost functions c_ϕ and full potential function f . Show that if each c_ϕ is nondecreasing, then f is concave, which implies that $NE(F)$ is the convex set of maximizers of f on X . (Hint: Fix $y, z \in X$, let $x(t) = (1 - t)y + tz$, and show that $g(t) = f(x(t))$ is concave.)
- (ii) Construct a congestion game in which each c_ϕ is increasing but in which $NE(F)$ is not a singleton.
- (iii) Show that in case (ii), the equilibrium link utilization levels u_ϕ are unique. (Hint: Since $f(x)$ only depends on the state x through the utilization levels $u_\phi(x)$, we can define a function $g : U \rightarrow \mathbf{R}$ on $U = \{v_\phi\}_{\phi \in \Phi} : v_\phi = u_\phi(x) \text{ for some } x \in \mathbf{R}_+^n\}$ by $g(u_\phi(x)) = f(x)$. Show that x maximizes f on X if and only if $u_\phi(x)$ maximizes g on U .)

Exercise 3.1.12. Example 3.1.6 shows that by adding state-dependent congestion charges to a congestion game, a planner can ensure that drivers use the network efficiently, in the sense of minimizing average travel times. Show that these congestion charges can be imposed on a link-by-link basis, and that the price on each link need only depend on the number of drivers on that link.

Exercise 3.1.13. Show that Cournot competition (Example 3.1.5) with a decreasing inverse demand function generates a potential game with a strictly concave potential function, and hence admits a unique Nash equilibrium.

Exercise 3.1.14. Entry and exit. When we define a full population game $F : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$, we specify the payoffs of each of the n strategies for all possible vectors of population masses. It is only a small additional step to allow agents to enter and leave the game. Fixing a vector of population masses (m^1, \dots, m^p) , we define a *population game with entry and exit* by assuming that the set of feasible social states is $\bar{X} = \{x = (x^1, \dots, x^p) \in \mathbf{R}_+^n : \sum_{i \in S^p} x_i^p \leq m^p\}$, and that an agent who exits the game receives a payoff of 0.

- (i) State an appropriate definition of Nash equilibrium for population games with entry and exit.
- (ii) A population game with entry and exit is a *potential game* if it satisfies full externality symmetry (3.2). Prove an analogue of Theorem 3.1.7 for such games.

3.1.5 The Geometry of Nash Equilibrium in Full Potential Games

Theorem 3.1.7 shows that if F is a potential game with potential function f , then the set of states satisfying the Kuhn-Tucker first order conditions for maximizing f are precisely

the Nash equilibria of F . We now offer a geometric proof of this result, and discuss its implications.

The nonlinear program from Section 3.1.4 seeks to maximize the function f on the polytope X . What do the Kuhn-Tucker conditions for this program mean?

The Kuhn-Tucker conditions adapt the classical approach to optimization based on linearization to settings with both equality and inequality constraints. In the current context, these conditions embody the following construction: To begin, one linearizes the objective function f at the state $x \in X$ of interest, replacing it with the function $l_{f,x}(y) = f(x) + \nabla f(x)'(y - x)$. Then, one determines whether the linearized function reaches its maximum on X at state x . Of course, this method can accept states that are not maximizers: for instance, if x is an interior local *minimizer* of f , then the linearization $l_{f,x}$ is a constant function, and so is maximized everywhere in X . But because X is a polytope (in particular, since constraint qualification holds), x must maximize $l_{f,x}$ on X if it is to maximize f on X .

With this interpretation of the Kuhn-Tucker conditions in hand, we can offer a simple geometric proof that $NE(F) = KT(f)$. The analysis employs our normal cone characterization of Nash equilibrium from Chapter 2.

Theorem 3.1.7: *If F is a full potential game with full potential function f , then $NE(F) = KT(f)$.*

$$\begin{aligned}
 \text{Second proof.} \quad x \in KT(f) &\Leftrightarrow x \text{ maximizes } l_{f,x} \text{ on } X. \\
 &\Leftrightarrow [z \in TX(x) \Rightarrow \nabla f(x)'z \leq 0] \\
 &\Leftrightarrow \nabla f(x) \in NX(x) \\
 &\Leftrightarrow F(x) \in NX(x) \\
 &\Leftrightarrow x \in NE(F). \blacksquare
 \end{aligned}$$

This proof is easy to explain in words. As we have argued, satisfying the Kuhn-Tucker conditions for f on X is equivalent to maximizing the linearized version of f on X . This in turn is equivalent to the requirement that if z is in the tangent cone of X at x —that is, if z is a feasible displacement direction from x —then z forms a weakly obtuse angle with the gradient vector $\nabla f(x)$, representing the direction in which f increases fastest. But this is precisely what it means for $\nabla f(x)$ to lie in the normal cone $NX(x)$. The definition of potential tells us that we can replace $\nabla f(x)$ with $F(x)$; and we know from Chapter 2 that $F(x) \in NX(x)$ means that x is a Nash equilibrium of F .

This argument sheds new light on Theorem 3.1.7. The Kuhn-Tucker conditions, which provide a way of finding the maximizers of the function f , are stated in terms of the

gradient vectors $\nabla f(x)$. At first glance, it seems rather odd to replace $\nabla f(x)$ with some non-integrable map F : after all, what is the point of the Kuhn-Tucker conditions when there is no function to maximize? But from the geometric point of view, replacing $\nabla f(x)$ with F makes perfect sense. When the Kuhn-Tucker conditions are viewed in geometric terms—namely, in the form $\nabla f(x) \in NX(x)$ —they become a restatement of the Nash equilibrium condition; the fact that $\nabla f(x)$ is a gradient vector plays no role. So to summarize, the Nash equilibrium condition $F(x) \in NX(x)$ is identical to the Kuhn-Tucker conditions, but applies whether or not the map F is integrable.

Exercise 3.1.15. Let F be a full potential game with full potential function f . Let $C \subseteq NE(F)$ be *smoothly connected*, in the sense that if $x, y \in C$, then there exists a piecewise C^1 path $\alpha : [0, 1] \rightarrow C$ with $\alpha(0) = x$ and $\alpha(1) = y$. Show that f is constant on C . (Hint: Use the Fundamental Theorem of Calculus and the fact that $F(x) \in NX(x)$ for all $x \in NE(F)$, along with the fact that when $\alpha(t) = x$ and α is differentiable at x , both $\alpha'(t)$ and $-\alpha'(t)$ are in $TX(x)$.)

3.1.6 Efficiency in Homogeneous Full Potential Games

We saw in Section 3.1.3 that when agents are matched to play normal form games with common interests, the full potential function of the resulting population game is proportional to the game's aggregate payoff function. How far can we push this connection?

Definition. We call a full potential game F homogeneous of degree k if each of its payoff functions $F_i^p : \mathbf{R}_+^n \rightarrow \mathbf{R}$ is a homogeneous function of degree k (that is, if $F_i^p(tx) = t^k F_i^p(x)$ for all $x \in \mathbf{R}_+^n$ and $t > 0$), where $k \neq -1$.

Example 3.1.16. Matching in normal form games with common interests. In the single population setting, each payoff function $F(x) = Ax$ is linear, so the full potential game F is homogeneous of degree 1. With $p \geq 2$ populations, the payoffs F^p to population p 's strategies are multilinear in $(x^1, \dots, x^{p-1}, x^{p+1}, \dots, x^p)$, so the full potential game F is homogeneous of degree $p - 1$. §

Example 3.1.17. Isoelastic congestion games. Let F be a congestion game with cost functions c_ϕ . For each facility $\phi \in \Phi$, let

$$\eta_\phi(u) = \frac{uc'_\phi(u)}{c_\phi(u)}$$

denote ϕ 's cost elasticity, which is well defined whenever $c_\phi(u) \neq 0$. We call a congestion

game *isoelastic with elasticity* $\eta \in \mathbf{R}$ if $\eta_\phi = \eta$ for all $\phi \in \Phi$. Thus, a congestion game is isoelastic if all facilities in Φ are equally sensitive to congestion at all levels of use.

Isoelasticity implies that all cost functions are of the form $c_\phi(u) = a_\phi u^\eta$, where the a_ϕ are arbitrary (i.e., positive or negative) scalar constants. (Notice that η cannot be negative, as this would force facility costs to become infinite at $u = 0$.) Since each u_ϕ is linear in x , each payoff function F_i^p is a sum of functions that are homogeneous of degree η in x , and so is itself homogeneous of degree η . Therefore, any isoelastic congestion game with elasticity η is a homogeneous potential game of degree η . §

The efficiency properties of homogeneous potential games are consequences of the following theorem.

Theorem 3.1.18. *The full potential game F is homogeneous of degree $k \neq -1$ if and only if the normalized aggregate payoff function $\frac{1}{k+1}\bar{F}(x)$ is a full potential function for F and is homogeneous of degree $k+1 \neq 0$.*

Proof. If the potential game F is homogeneous of degree $k \neq -1$, then $\frac{1}{k+1}\bar{F}(x) = \frac{1}{k+1} \sum_{p \in \mathcal{P}} \sum_{j \in S^p} x_j^p F_j^p(x)$ is clearly homogeneous of degree $k+1$. Therefore, condition (3.2) and Euler's Theorem (see Appendix 3.A.5) imply that

$$\begin{aligned} \frac{\partial}{\partial x_i^p} \left(\frac{1}{k+1} \bar{F}(x) \right) &= \frac{1}{k+1} \left(\sum_{q \in \mathcal{P}} \sum_{j \in S^q} x_j^q \frac{\partial F_j^q}{\partial x_i^p}(x) + F_i^p(x) \right) \\ &= \frac{1}{k+1} \left(\sum_{q \in \mathcal{P}} \sum_{j \in S^q} x_j^q \frac{\partial F_i^p}{\partial x_j^q}(x) + F_i^p(x) \right) \\ &= \frac{1}{k+1} (kF_i^p(x) + F_i^p(x)) \\ &= F_i^p(x), \end{aligned}$$

so $\frac{1}{k+1}\bar{F}$ is a full potential function for F . On the other hand, if $\frac{1}{k+1}\bar{F}$ is homogeneous of degree $k+1 \neq 0$ and is a full potential function for F , then each payoff function $F_i^p = \frac{\partial}{\partial x_i^p} \left(\frac{1}{k+1} \bar{F} \right)$ is homogeneous of degree k , so the converse statement follows. ■

To understand the connection between homogeneity and efficiency, consider the expression $\frac{\partial}{\partial x_i^p} \bar{F}(x)$, which represents the impact of an agent who chooses strategy i on aggregate payoffs. Recalling Example 3.1.6, we split this impact into two terms. The first term, $\sum_q \sum_j x_j^q \frac{\partial F_j^q}{\partial x_i^p}(x)$, represents the impact of this agent's behavior on his opponents' payoffs. The second term, $F_i^p(x)$, represents the agent's own payoffs. In homogeneous potential games, these two effects are precisely balanced: the payoff an agent receives

from choosing a strategy is directly proportional to the social impact of his choice. For this reason, self-interested behavior leads to desirable social outcomes.

Observe that if a potential game is homogeneous of degree less than -1 , its full potential function is *negatively* proportional to aggregate payoffs. In this case, self-interested behavior leads to undesirable social outcomes. To remove this case from consideration, we call a potential game *positively homogeneous* if its full potential function is homogeneous of positive degree, so that the game itself is homogeneous of degree $k > -1$.

With this definition in hand, we can present a result on the efficiency of Nash equilibria. We call the social state x *locally efficient* in game F ($x \in LE(F)$) if there exists an $\varepsilon > 0$ such that $\bar{F}(x) \geq \bar{F}(y)$ for all $y \in X$ within ε of x . If this inequality holds for all $y \in X$, we call x *globally efficient* ($x \in GE(F)$).

Corollary 3.1.19.

- (i) If the full potential game F is positively homogeneous, then $LE(F) \subseteq NE(F)$.
- (ii) If in addition its full potential function f is concave, then $GE(F) = LE(F) = NE(F)$.

Exercise 3.1.20. Establish these claims.

Exercise 3.1.21. Let F be a congestion game with nondecreasing affine cost functions: $c_\phi(u) = a_\phi u + b_\phi$. Suppose that within each population, the fixed cost of each route is equal:

$$\sum_{\phi \in \Phi_i^p} b_\phi = b^p \text{ for all } i \in S^p \text{ and } p \in \mathcal{P}.$$

Show that $NE(F) = GE(F)$.

3.1.7 Inefficiency Bounds for Congestion Games

The results from the previous section provide stringent conditions under which Nash equilibria of congestion games are efficient. Since exact efficiency rarely obtains, it is natural to ask just how inefficient equilibrium behavior can be. We address this question in the context of congestion games with nondecreasing cost functions—in other words, congestion games in which congestion is a bad.

It will be convenient to use notation tailored to the questions at hand. Given the facilities ϕ and the nondecreasing cost functions c_ϕ , we let

$$C_i^p(x) = -F_i^p(x) = \sum_{\phi \in \Phi_i^p} c_\phi(u_\phi(x))$$

denote the *cost* of strategy $i \in S^p$ at state x and let

$$\bar{C}(x) = -\bar{F}(x) = \sum_{p \in \mathcal{P}} \sum_{i \in S^p} x_i^p C_i^p(x) = \sum_{\phi \in \Phi} u_\phi(x) c_\phi(u_\phi(x))$$

denote *social cost* at state x . We refer to the resulting congestion game either as C or as (C, m) (to emphasize the population masses m). When we introduce alternative cost functions γ_ϕ , we replace C with Γ in the notation above.

One approach to bounding the inefficiency of equilibria is to compare the equilibrium social cost to the minimal social cost in a game with additional agents.

Proposition 3.1.22. *Let C be a congestion game with nondecreasing cost functions. Let x^* be a Nash equilibrium of (C, m) , and let y be a feasible state in $(C, 2m)$. Then $\bar{C}(x^*) \leq \bar{C}(y)$.*

Exercise 3.1.23. This exercise outlines a proof of Proposition 3.1.22.

- (i) Define the cost functions γ_ϕ by $\gamma_\phi(u) = \max\{c_\phi(u_\phi(x^*)), c_\phi(u)\}$. Show that $u(\gamma_\phi(u) - c_\phi(u)) \leq c_\phi(u_\phi(x^*)) u_\phi(x^*)$.
- (ii) Show that $\Gamma_i^p(y) \geq \min_{j \in S^p} C_j^p(x^*)$.
- (iii) Use parts (i) and (ii) to show that $\bar{\Gamma}(y) - \bar{C}(y) \leq \bar{C}(x^*)$ and that $\bar{\Gamma}(y) \geq 2\bar{C}(x^*)$, and conclude that $\bar{C}(x^*) \leq \bar{C}(y)$.

Exercise 3.1.24. This exercise applies Proposition 3.1.22 to settings with fixed population masses but varying cost functions.

- (i) Show that the equilibrium social cost under cost functions $\tilde{c}_\phi(u) = \frac{1}{2}c_\phi(\frac{u}{2})$ is bounded above by the minimal social cost under cost functions c_ϕ .
- (ii) Let C be a congestion game with cost functions $c_\phi(u) = (k_\phi - u)^{-1}$ for some *capacities* $k_\phi > 0$. (We assume that population masses are small enough that no edge can reach its capacity.) Using part (i), show that the equilibrium social cost when capacities are $2k$ is bounded above by the minimal social cost when capacities are k . In other words, doubling the capacities of the edges reduces costs at least as much as enforcing efficient behavior under the original capacities.

A more direct way of understanding inefficiency is to bound a game's *inefficiency ratio*: the ratio between the game's equilibrium social cost and its minimal feasible social cost.

Example 3.1.25. A highway network consisting of two parallel links is to be traversed by a unit mass of drivers. The links' cost functions are $c_1(u) = 1$ and $c_2(u) = u$. In the unique Nash equilibrium of this game, all drivers travel on route 2, creating a social cost of 1. The efficient state, which minimizes $\bar{C}(x) = x_1 + (x_2)^2$, is $x_{\min} = (\frac{1}{2}, \frac{1}{2})$; it generates a social cost of $\bar{C}(x_{\min}) = \frac{3}{4}$. Thus, the inefficiency ratio in this game is $\frac{4}{3}$.

The next result describes an easily established upper bound on inefficiency ratios.

Proposition 3.1.26. *Suppose that the cost functions c_ϕ are nondecreasing and satisfy $uc_\phi(u) \leq \alpha \int_0^u c_\phi(z) dz$ for all $u \geq 0$. If $x^* \in NE(C)$ and $x \in X$, then $\bar{C}(x^*) \leq \alpha \bar{C}(x)$.*

Exercise 3.1.27. (i) Prove Proposition 3.1.26. (Hint: Use a potential function argument.)

(ii) Show that if cost functions in C are polynomials of degree at most k with nonnegative coefficients, then the inefficiency ratio in C is at most $k + 1$.

Exercise 3.1.27 tells us that the inefficiency ratio of a congestion game with affine cost functions cannot exceed 2. Is it possible to establish a smaller upper bound? We saw in Example 3.1.25 that inefficiency ratios as high as $\frac{4}{3}$ can arise in very simple games with nonnegative affine cost functions. Amazingly, $\frac{4}{3}$ is the highest possible inefficiency ratio for congestion games with cost functions of this form.

Theorem 3.1.28. *Let C be a congestion game whose cost functions c_ϕ are nonnegative, nondecreasing, and affine: $c_\phi(u) = a_\phi + b_\phi u$ with $a_\phi, b_\phi \geq 0$. If $x^* \in NE(C)$ and $x \in X$, then $\bar{C}(x^*) \leq \frac{4}{3} \bar{C}(x)$.*

Proof. Fix $x^* \in NE(C)$ and $x \in X$, and write $v_\phi^* = u_\phi(x^*)$ and $v_\phi = u_\phi(x)$. Let $\Phi_L = \{\phi \in \Phi : v_\phi < v_\phi^*\}$ be the set of facilities that are underutilized at x relative to x^* .

Every $r \in \mathbf{R}$ satisfies $(1 - r)r \leq \frac{1}{4}$. Multiplying both sides of this inequality by $b_\phi (v_\phi^*)^2$ and setting $r = v_\phi/v_\phi^*$ yields

$$(3.8) \quad b_\phi (v_\phi^* - v_\phi) \cdot v_\phi \leq \frac{1}{4} b_\phi (v_\phi^*)^2 \text{ whenever } \phi \in \Phi_L.$$

Thus, since x^* is a Nash equilibrium, and by our assumptions on c_ϕ , we have that

$$\begin{aligned} \bar{C}(x^*) &= \sum_{\phi \in \Phi} c_\phi(v_\phi^*) v_\phi^* \\ &\leq \sum_{\phi \in \Phi} c_\phi(v_\phi^*) v_\phi \\ &= \sum_{\phi \in \Phi} c_\phi(v_\phi) v_\phi + \sum_{\phi \in \Phi} (c_\phi(v_\phi^*) - c_\phi(v_\phi)) v_\phi \\ &\leq \bar{C}(x) + \sum_{\phi \in \Phi_L} b_\phi (v_\phi^* - v_\phi) \cdot v_\phi \\ &\leq \bar{C}(x) + \frac{1}{4} \sum_{\phi \in \Phi_L} b_\phi (v_\phi^*)^2 \end{aligned}$$

$$\begin{aligned}
&\leq \bar{C}(x) + \frac{1}{4} \sum_{\phi \in \Phi_L} v_\phi^* \cdot c_\phi(v_\phi^*) \\
&= \bar{C}(x) + \frac{1}{4} \bar{C}(x^*).
\end{aligned}$$

Rearranging yields $\bar{C}(x^*) \leq \frac{4}{3} \bar{C}(x)$. ■

That the highest inefficiency ratio for a given class of cost functions can be realized in a very simple network is true quite generally. Consider a two-link network with link cost functions $c_1(u) = 1$ and $c_2(u) = u^k$, where $k \geq 1$. With a unit mass population, the Nash equilibrium for this network is $x^* = (0, 1)$, and has social cost $\bar{C}(x^*) = 1$; the efficient state is $x_{\min} = (1 - (k + 1)^{-1/k}, (k + 1)^{-1/k})$, and has social cost $\bar{C}(x_{\min}) = 1 - k(k + 1)^{-(k+1)/k}$. Remarkably, it is possible to show that the resulting inefficiency ratio of $(1 - k(k + 1)^{-(k+1)/k})^{-1}$ is the highest possible in any network whose cost functions are polynomials of degree at most k . See the Notes for further details.

A minor modification of the proof of Theorem 3.1.28 allows us to improve the bound presented in Proposition 3.1.22. The earlier result showed that equilibrium social cost in a congestion game cannot exceed the efficient level of social cost in the game with identical cost functions but with population sizes that are twice as large. Corollary 3.1.29 shows that in the affine case, this multiplicative factor can be reduced from 2 to $\frac{5}{4}$.

Corollary 3.1.29. *Let C be a congestion game whose cost functions c_ϕ are nonnegative, nondecreasing, and affine. Let x^* be a Nash equilibrium of (C, m) , and let y be a feasible state in $(C, \frac{5}{4}m)$. Then $\bar{C}(x^*) \leq \bar{C}(y)$.*

Exercise 3.1.30. Prove Corollary 3.1.29. (Hint: The proof of Theorem 3.1.28 establishes that $\sum_{\phi \in \Phi} c_\phi(u_\phi(x^*)) u_\phi(y) \leq \bar{C}(y) + \frac{1}{4} \bar{C}(x^*)$ for any $y \in \mathbf{R}_+^n$. Combine this inequality with the fact that x^* is a Nash equilibrium of (C, m) .)

While our treatment of inefficiency bounds has focused on congestion games with affine and polynomial cost functions, it is possible to establish such bounds while placing much less structure on the costs. In fact, one can generalize the arguments presented above to obtain inefficiency bounds for general population games. See the Notes for references to the relevant literature.

3.2 Potential Games

To define full potential games, we first defined full population games by extending the domain of payoffs from the state space X to the positive orthant \mathbf{R}_+^n . While this device for

introducing potential functions is simple, it is often artificial. By using ideas from affine calculus (Appendix 3.B), we can define potential functions for population games without recourse to changes in domain.

3.2.1 Motivating Examples

We can motivate the developments to come not only by parsimony, but also by generality, as the following two examples show.

Example 3.2.1. Matching in symmetric normal form potential games. Recall that the symmetric normal form game $C \in \mathbf{R}^{n \times n}$ is a common interest game if C is a symmetric matrix, so that both players always receive the same payoff. We call the symmetric normal form game $A \in \mathbf{R}^{n \times n}$ a *potential game* if $A = C + \mathbf{1}r'$ for some common interest game C and some arbitrary vector $r \in \mathbf{R}^n$. Thus, each player's payoff is the sum of a common interest term and a term that only depends on his opponent's choice of strategy. (For the latter point, note that $A_{ij} = C_{ij} + r_j$.)

Suppose a population of agents is matched to play game A . Since the second payoff term has no effect on agents' incentives, it is natural to expect our characterization of equilibrium from the previous section to carry over to the current setting. But this does not follow directly from our previous definitions. Suppose we define the full population game $F : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ as in Example 3.1.2: $F(x) = Ax$. Then the resulting derivative matrix is $DF(x) = A = C + \mathbf{1}r'$, and so

$$\frac{\partial F_i}{\partial x_j}(x) = C_{ij} + r_j, \quad \text{but} \quad \frac{\partial F_j}{\partial x_i}(x) = C_{ji} + r_i.$$

Therefore, unless r is a constant vector (in which case A itself is symmetric), the full population game F defined above is not a full potential game. §

Example 3.2.2. Two-strategy games. Recall that the population game $F : X \rightarrow \mathbf{R}^n$ is a two-strategy game if $p = 1$ and $n = 2$. In this setting, the state space X is the simplex in \mathbf{R}^2 , which can be viewed as a relabelling of the unit interval. Because all functions defined on the unit interval are integrable, it seems natural to expect two-strategy games to admit potential functions. If we wanted to show that F defines a full potential game, we would first need to extend its domain to \mathbf{R}_+^2 . Once we do this, the domain is no longer one-dimensional, so our intuition about the existence of a potential function is lost. §

3.2.2 Definition, Characterizations, and Examples

Example 3.2.2 suggests that the source of our difficulties is the extension of payoffs from the original state space X to the full-dimensional set \mathbf{R}_+^n . As the definition of full potential games relied on this extension, our new notion of potential games will require some additional ideas. The key concepts are the tangent spaces and orthogonal projections introduced in Chapter 2, which we briefly review here.

Recall that the state space for population p is given by $X^p = \{x^p \in \mathbf{R}_+^{m^p} : \sum_{i \in S^p} x_i^p = m^p\}$. The tangent space of X^p , denoted TX^p , is the smallest subspace of \mathbf{R}^{m^p} that contains all directions of motion through X^p ; it is defined by $TX^p = \mathbf{R}_0^{m^p} \equiv \{z^p \in \mathbf{R}^{m^p} : \sum_{i \in S^p} z_i^p = 0\}$. The matrix $\Phi \in \mathbf{R}^{m^p \times m^p}$, representing the orthogonal projection of \mathbf{R}^{m^p} onto TX^p , is defined by $\Phi = I - \frac{1}{m^p} \mathbf{1}\mathbf{1}'$. If $\pi^p \in \mathbf{R}^{m^p}$ is a payoff vector, then the projected payoff vector $\Phi\pi^p$ represents relative payoffs under π^p : it preserves the differences between components of π^p while normalizing their sum to zero. Changes in the social state $x \in X = \prod_{p \in \mathcal{P}} X^p$ are represented by elements of $TX = \prod_{p \in \mathcal{P}} TX^p$, the tangent space of X . The matrix $\Phi \in \mathbf{R}^{n \times n}$, representing the orthogonal projection of \mathbf{R}^n onto TX , is the block diagonal matrix $\text{diag}(\Phi, \dots, \Phi)$. If $\pi = (\pi^1, \dots, \pi^p) \in \mathbf{R}^n$ is a payoff vector for the society, then $\Phi\pi = (\Phi\pi^1, \dots, \Phi\pi^p)$ normalizes each of the p pieces of the vector π separately.

With these preliminaries in hand, we are ready for our new definition.

Definition. Let $F : X \rightarrow \mathbf{R}^n$ be a population game. We call F a potential game if it admits a potential function: a C^1 function $f : X \rightarrow \mathbf{R}$ that satisfies

$$(3.9) \quad \nabla f(x) = \Phi F(x) \text{ for all } x \in X.$$

Since the potential function f has domain X , the gradient vector $\nabla f(x)$ is by definition an element of the tangent space TX (see Appendix 3.B.3). Our definition of potential games requires that this gradient vector always equal $\Phi F(x)$, the projection of the payoff vector $F(x)$ onto the subspace TX .

At the cost of sacrificing parsimony, one can define potential games without affine calculus by using a function defined throughout \mathbf{R}_+^n to play the role of the potential function f . To do so, one simply includes the projection Φ on both sides of the analogue of equation (3.9).

Observation 3.2.3. If F is a potential game with potential function $f : X \rightarrow \mathbf{R}$, then any C^1 extension $\tilde{f} : \mathbf{R}_+^n \rightarrow \mathbf{R}$ of f satisfies

$$(3.10) \quad \Phi \nabla \tilde{f}(x) = \Phi F(x) \text{ for all } x \in X.$$

Conversely, if the population game F admits a function \tilde{f} satisfying condition (3.10), then F is a potential game, and the restriction $f = \tilde{f}|_X$ is a potential function for F .

This observation is immediate from the relevant definitions. In particular, if \tilde{f} and f agree on X , then for all $x \in X$ the gradient vectors $\nabla\tilde{f}(x)$ and $\nabla f(x)$ define identical linear operators on TX , implying that $\Phi\nabla\tilde{f}(x) = \Phi\nabla f(x)$. But since $\Phi\nabla f(x) = \nabla f(x)$ by definition, it follows that $\Phi\nabla\tilde{f}(x) = \nabla f(x)$; this equality and definition (3.9) yield the result.

Like full potential games, potential games can be characterized by a symmetry condition on the payoff derivatives $DF(x)$. Since potential games generalize full potential games, the new symmetry condition is less restrictive than the old one.

Theorem 3.2.4. *Suppose the population game $F : X \rightarrow \mathbf{R}^n$ is C^1 . Then F is a potential game if and only if it satisfies externality symmetry:*

$$(3.11) \quad DF(x) \text{ is symmetric with respect to } TX \times TX \text{ for all } x \in X.$$

Proof. Immediate from Theorem 3.B.8 in Appendix 3.B. ■

Condition (3.11) demands that at each state $x \in X$, the derivative $DF(x)$ define a symmetric bilinear form on $TX \times TX$:

$$z'DF(x)\hat{z} = \hat{z}'DF(x)z \text{ for all } z, \hat{z} \in TX \text{ and } x \in X.$$

Observation 3.2.5 offers a version of this condition that does not require affine calculus, just as Observation 3.2.3 did for definition (3.9).

Observation 3.2.5. *Suppose that the population game $F : X \rightarrow \mathbf{R}^n$ is C^1 , and let $\tilde{F} : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ be any C^1 extension of F . Then F satisfies externality symmetry (and so is a potential game) if and only if*

$$\Phi D\tilde{F}(x)\Phi \text{ is symmetric for all } x \in X.$$

The next exercise characterizes externality symmetry in a more intuitive way.

Exercise 3.2.6. Show that externality symmetry (3.11) holds if and only if the previous equality holds whenever $z = e_j^p - e_i^p$ and $\hat{z} = e_l^q - e_k^q$. In other words, show that (3.11) is equivalent to

$$(3.12) \quad \frac{\partial(F_j^p - F_i^p)}{\partial(e_l^q - e_k^q)}(x) = \frac{\partial(F_l^q - F_k^q)}{\partial(e_j^p - e_i^p)}(x) \text{ for all } i, j \in S^p, k, l \in S^q, p, q \in \mathcal{P}, \text{ and } x \in X.$$

The left hand side of equation (3.12) captures the change in the payoff to strategy $j \in S^p$ relative to strategy $i \in S^p$ as agents switch from strategy $k \in S^q$ to strategy $l \in S^q$. This effect must equal the change in the payoff of l relative to k as agents switch from i to j , as expressed on the right hand side of (3.12). This description is akin to that of full externality symmetry (3.2) (see the discussion after equation (3.3)), but it only refers to *relative* payoffs and to *feasible* changes in the social state.

Exercise 3.2.7. Let F be a C^1 single population game. Show that F is a potential game if and only if it satisfies *triangular integrability*:

$$\frac{\partial F_i}{\partial(e_j - e_k)}(x) + \frac{\partial F_j}{\partial(e_k - e_i)}(x) + \frac{\partial F_k}{\partial(e_i - e_j)}(x) = 0 \text{ for all } i, j, k \in S \text{ and } x \in X.$$

We now return to the examples that led off the section.

Example 3.2.8. Two-strategy games revisited. If $F : X \rightarrow \mathbf{R}^2$ is a smooth two-strategy game, its state space X is the simplex in \mathbf{R}^2 , whose tangent space TX is spanned by the vector $d = e_1 - e_2$. If z and \hat{z} are vectors in TX , then $z = kd$ and $\hat{z} = \hat{k}d$ for some real numbers k and \hat{k} ; thus, however F is defined, we have that $z'DF(x)\hat{z} = k\hat{k}d'DF(x)d = z'DF(x)z$ for all $x \in X$. In other words, F is a potential game. Even if F is merely continuous, the function $f : X \rightarrow \mathbf{R}$ defined by

$$(3.13) \quad f(x_1, 1 - x_1) = \int_0^{x_1} (F_1(t, 1 - t) - F_2(t, 1 - t)) dt$$

is a potential function for F , so F is still a potential game. (If you think that a $\frac{1}{2}$ is needed on the right hand side of equation (3.13), convince yourself that it is not.) §

Exercise 3.2.9. Matching in symmetric normal form potential games. Let $A = C + \mathbf{1}r'$ be a symmetric normal form potential game: $C \in \mathbf{R}^{n \times n}$ is symmetric, and $r \in \mathbf{R}^n$ is arbitrary. Define the population game $F : X \rightarrow \mathbf{R}^n$ by $F(x) = Ax$. Use one of the derivative conditions above to verify that F is a potential game, and find a potential function $f : X \rightarrow \mathbf{R}$ for F .

Exercise 3.2.10. Matching in normal form potential games. The normal form game $U = (U^1, \dots, U^p)$ is a *potential game* if there is a *potential function* $V : S \rightarrow \mathbf{R}$ such that

$$U^p(\hat{s}^p, s^{-p}) - U^p(s) = V(\hat{s}^p, s^{-p}) - V(s) \text{ for all } s \in S, \hat{s}^p \in S^p, \text{ and } p \in \mathcal{P}.$$

That is, after any unilateral deviation, the change in the deviator's payoffs is equal to the change in potential. It is easy to see that pure strategy profile $s \in S$ is a Nash equilibrium of U if and only if it is a local maximizer of the potential function V .

- (i) Show that U is a potential game with potential function V if and only if there are auxiliary functions $W^p : S^{-p} \rightarrow \mathbf{R}$ such that

$$U^p(s) = V(s) + W^p(s^{-p}) \quad \text{for all } s \in S \text{ and } p \in \mathcal{P}.$$

In words: each player's payoff function is the sum of a of a common payoff term given by the value of potential, and a term that only depends on opponents' behavior. This characterization accords with the definition of symmetric normal form potential games from the previous exercise.

- (ii) Define the full population game $\tilde{F} : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ by

$$\tilde{F}_{s^p}^p(x) = \sum_{s^{-p} \in S^{-p}} U^p(s) \prod_{r \neq p} x_{s^r}^r = \sum_{s^{-p} \in S^{-p}} (V(s) + W^p(s^{-p})) \prod_{r \neq p} x_{s^r}^r.$$

Show that \tilde{F} is not a full potential game.

- (iii) Define the population game $F : X \rightarrow \mathbf{R}^n$ using the equation from part (ii). By verifying condition (3.11), show that F is a potential game.
 (iv) Construct a potential function for F .

Exercise 3.2.11. This exercise provides the converse to Exercise 3.2.10(iii).

Let the population game $F : X \rightarrow \mathbf{R}^n$ be generated by matching in a p player normal form game U . Show that if F is potential game with potential function $f : X \rightarrow \mathbf{R}$, then U is a potential game with potential function $V(s) = f(\xi(s))$, where $\xi(s) \in X$ is the pure population state with $\xi(s)_{s^p}^p = 1$ for all $p \in \mathcal{P}$. (Hint: To evaluate $f(\xi(s^p, s^{-p})) - f(\xi(s))$, use the Fundamental Theorem of Calculus, along with the fact that $F^p(x)$ is independent of x^p .)

3.2.3 Potential Games and Full Potential Games

What is the relationship between full potential games and potential games? In the former case, condition (3.1) requires that payoffs be completely determined by the potential function, which is defined on \mathbf{R}_+^n ; in the latter, condition (3.9) asks only that relative payoffs be determined by the potential function, now defined just on X .

To understand the relationship between the two definitions, take a potential game $F : X \rightarrow \mathbf{R}^n$ with potential function $f : X \rightarrow \mathbf{R}$ as given, and extend f to a full potential function $\tilde{f} : \mathbf{R}_+^n \rightarrow \mathbf{R}$. Theorem 3.2.12 shows that the link between the full potential game $\tilde{F} \equiv \nabla \tilde{f}$ and the original game F depends on how the extension \tilde{f} is chosen.

Theorem 3.2.12. *Let $F : X \rightarrow \mathbf{R}^n$ be a potential game with potential function $f : X \rightarrow \mathbf{R}$. Let $\tilde{f} : \mathbf{R}_+^n \rightarrow \mathbf{R}$ be any C^1 extension of f , and define the full potential game $\tilde{F} : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ by*

$\tilde{F}(x) = \nabla \tilde{f}(x)$. Then

- (i) The population games F and $\tilde{F}|_X$ have the same relative payoffs: $\Phi F(x) = \Phi \tilde{F}(x)$ for all $x \in X$.
- (ii) One can choose the extension \tilde{f} in such a way that F and $\tilde{F}|_X$ are identical.

Part (i) of the theorem shows that the full potential game \tilde{F} generated from an arbitrary extension of the potential function f exhibits the same relative payoffs as F on their common domain X . It follows that F and \tilde{F} have the same best response correspondences and Nash equilibria, but may exhibit different average payoff levels. Part (ii) of the theorem shows that by choosing the extension \tilde{f} appropriately, we can make \tilde{F} and F identical on X . To accomplish this, we construct the extension \tilde{f} in such a way (equation (3.14) below) that its derivatives at states in X evaluated in directions orthogonal to TX encode information about average payoffs from the original game F .

In conclusion, Theorem 3.2.12(ii) demonstrates that if population masses are fixed, so that the relevant set of social states is X , then definition (3.1), while more difficult to check, does not entail a loss of generality relative to definition (3.9).

Proof of Theorem 3.2.12: Part (i) follows from the fact that $\Phi \tilde{F}(x) = \Phi \nabla \tilde{f}(x) = \nabla f(x) = \Phi F(x)$ for all $x \in X$; compare the discussion following Observation 3.2.3.

To prove part (ii), we first extend f and F from the state space X to its affine hull $\text{aff}(X)$. Let $\hat{f} : \text{aff}(X) \rightarrow \mathbf{R}$ be a C^1 extension of $f : X \rightarrow \mathbf{R}$, and let $\hat{g}^p : \text{aff}(X) \rightarrow \mathbf{R}$ be a continuous extension of population p 's average payoff function, $\frac{1}{n^p} \mathbf{1}' F^p : X \rightarrow \mathbf{R}$. (The existence of these extensions follows from the Whitney Extension Theorem.) Then define $\hat{G} : \text{aff}(X) \rightarrow \mathbf{R}^n$ by $\hat{G}^p(x) = \mathbf{1} \hat{g}^p(x)$, so that $F(x) = \Phi F(x) + (I - \Phi)F(x) = \nabla \hat{f}(x) + \hat{G}(x)$ for all $x \in X$. If after this we define $\hat{F} : \text{aff}(X) \rightarrow \mathbf{R}^n$ by $\hat{F}(x) = \nabla \hat{f}(x) + \hat{G}(x)$, then \hat{F} is a continuous extension of F , and $\nabla \hat{f}(x) = \Phi \hat{F}(x)$ for all $x \in \text{aff}(X)$.

With this groundwork complete, we can extend f to all of \mathbf{R}_+^n via

$$(3.14) \quad \tilde{f}(y) = f(\xi(y)) + (y - \xi(y))' F(\xi(y)),$$

where $\xi(y) = \Phi y + z_{TX}^\perp$ is the closest point to y in $\text{aff}(X)$. (Here, z_{TX}^\perp is the orthogonal translation vector that sends TX to $\text{aff}(X)$: namely, $(z_{TX}^\perp)^p = \frac{m^p}{n^p} \mathbf{1}$.) Theorem 3.B.10 shows that $\nabla \tilde{f}|_X = \tilde{F}|_X$ is identical to F . ■

Theorem 3.2.12 implies that all of our results from Sections 3.1.4 and 3.1.5 on Nash equilibria of full potential games apply unchanged to potential games. On the other hand, the efficiency results from Section 3.1.6 do not. In particular, the proof of Theorem 3.1.18 depends on the game F being a full population game, as the application of Euler's Theorem makes explicit use of the partial derivatives of F . In fact, to establish that a potential game

F has efficiency properties of the sorts described in Section 3.1.6, one must show that F can be extended to a homogeneous full potential game. This should come as no surprise: since the potential function $f : X \rightarrow \mathbf{R}$ only captures relative payoffs, it cannot be used to prove efficiency results, which depend on both relative and average payoffs.

Exercise 3.2.13. Consider population games with entry and exit (Exercise 3.1.14). Which derivative condition is the right one for defining potential games in this context, (3.2) or (3.11)? Why?

Exercise 3.2.14. Prove this simple “converse” to Theorem 3.2.12: Suppose $\tilde{F} : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ is a full potential game with full potential function $\tilde{f} : \mathbf{R}_+^n \rightarrow \mathbf{R}$. Let $F = \tilde{F}|_X$ and $f = \tilde{f}|_X$. Then F is a potential game with potential function f .

3.2.4 Passive Games and Constant Games

We conclude this section by introducing two simple classes of population games.

Definition. The population game $H : X \rightarrow \mathbf{R}^n$ is a passive game if for each state $x \in X$ and each population $p \in \mathcal{P}$, the payoffs to all of population p 's strategies are equal:

$$H_i^p(x) = H_j^p(x) \text{ for all } i, j \in S^p, p \in \mathcal{P}, \text{ and } x \in X.$$

Definition. The population game $K : X \rightarrow \mathbf{R}^n$ is a constant game if all strategies' payoffs are independent of the state: that is, if $K(x) = \pi$ for all $x \in X$, or, more explicitly, if

$$K_i^p(x) = \pi_i^p \text{ for all } i \in S^p, p \in \mathcal{P}, \text{ and } x \in X.$$

In a passive game, an agent's own behavior has no bearing on his payoffs; in a constant game, each agent's behavior is the sole determinant of his payoffs.

The following two propositions provide some alternate characterizations of these games.

Proposition 3.2.15. The following statements are equivalent:

- (i) H is a passive game.
- (ii) There are functions $c^p : X \rightarrow \mathbf{R}$ such that $H^p(x) = c^p(x)\mathbf{1}$ for all $p \in \mathcal{P}$ and $x \in X$.
- (iii) $H(x) \in (TX)^\perp$ for all $x \in X$.
- (iv) $\Phi H(x) = \mathbf{0}$ for all $x \in X$.
- (v) $z'H(x) = \mathbf{0}$ for all $z \in TX$ and $x \in X$.
- (vi) H is a potential game whose potential function is constant.

Proposition 3.2.16. *The following statements are equivalent:*

- (i) K is a constant game.
- (ii) $DK(x) = \mathbf{0}$ for all $x \in X$.
- (iii) K is a potential game that admits a linear potential function.

In particular, if $K(x) = \pi$ is a constant game, then $k(x) = \pi'x$ is a potential function for K .

One reason that passive and constant games are interesting is that adding them to a population game from a certain class (the potential games, the stable games, the supermodular games) results in a new game from the same class. For instance, suppose that F is a potential game with potential function f , let H be a passive game, and let K be a constant game with potential function k . Evidently, $F + H$ is also a potential game with potential function f ; thus, adding H to F leaves the Nash equilibria of F unchanged. $F + K$ is also a potential game, but its potential function is not f , but $f + k$; thus, $NE(F)$ and $NE(F + K)$ generally differ. Similar observations are true for stable games and for supermodular games: adding a passive game or a constant game to a game from either of these classes keeps us in the class, but only adding passive games leaves incentives unchanged.

When payoffs are smooth, the invariances just described can be represented in terms of payoff derivatives. As an illustration, recall that the C^1 population game $F : X \rightarrow \mathbf{R}^n$ is a potential game if and only if it satisfies externality symmetry:

$$(3.11) \quad DF(x) \text{ is symmetric with respect to } TX \times TX \text{ for all } x \in X.$$

The first TX tells us that condition (3.11) constrains the effects of left multiplication of $DF(x)$ by elements of TX ; this restricts the purview of (3.11) to *changes in relative payoffs*. The second TX tells us that (3.11) constrains the effects of right multiplication of $DF(x)$ by elements of TX ; this reflects that we can only evaluate how payoffs change in response to *feasible changes in the state*. In summary, the action of the derivative matrices $DF(x)$ on $TX \times TX$ captures changes in relative payoffs due to feasible changes in the state. We have seen that this action is enough to characterize potential games, and we will soon find that it is enough to characterize stable and supermodular games as well.

It follows from this discussion that the additions to F that do not affect the action of its derivative matrices on $TX \times TX$ are the ones that do not alter F 's class. These additions are characterized by the following proposition.

Proposition 3.2.17. *Let G be a C^1 population game. Then $DG(x)$ is the null bilinear form on $TX \times TX$ for all $x \in X$ if and only if $G = H + K$, where H is a passive game and K is a constant game.*

Exercise 3.2.18. Prove Propositions 3.2.15, 3.2.16, and 3.2.17. (Hints: For Proposition 3.2.16, prove the equivalence of (i) and (iii) using the Fundamental Theorem of Calculus. For 3.2.17, use the previous propositions, along with the fact that $DG(x)$ is the null bilinear form on $TX \times TX$ if and only if $\Phi DG(x) = \mathbf{0}$.)

Exercise 3.2.19.

- (i) Suppose $H(x) = Ax$ is a single population passive game. Describe A .
- (ii) Suppose $K(x) = Ax$ is a single population constant game. Describe A .

3.3 Stable Games

There are a variety of well-known classes of games whose Nash equilibria lie in a single convex component: for instance, two player zero-sum games, wars of attrition, games with an interior ESS or NSS, and potential games with concave potential functions. This shared property of these seemingly disparate examples springs from a common source: all of these examples are stable games.

3.3.1 Definition

The common structure in the examples above is captured by the following definition.

Definition. *The population game $F : X \rightarrow \mathbf{R}^n$ is a stable game if*

$$(3.15) \quad (y - x)'(F(y) - F(x)) \leq 0 \text{ for all } x, y \in X.$$

If the inequality in condition (3.15) holds strictly whenever $x \neq y$, we call F a strictly stable game, while if this inequality always binds, we call F a null stable game.

For a first intuition, imagine for the moment that $F \equiv \nabla f(x)$ is also a full potential game. In this case, condition (3.15) is simply the requirement that the potential function f be concave. Our definition of stable games thus extends the defining property of concave potential games to games whose payoffs are not integrable.

Stable games whose payoffs are differentiable can be characterized in terms of the action of their derivative matrices $DF(x)$ on $TX \times TX$.

Theorem 3.3.1. *Suppose the population game F is C^1 . Then F is a stable game if and only if it satisfies self-defeating externalities:*

$$(3.16) \quad DF(x) \text{ is negative semidefinite with respect to } TX \text{ for all } x \in X.$$

Before proving Theorem 3.3.1, let us provide some intuition for condition (3.16). This condition asks that

$$z'DF(x)z \leq 0 \text{ for all } z \in TX \text{ and } x \in X.$$

This requirement is in turn equivalent to

$$\sum_{p \in \mathcal{P}} \sum_{i \in S^p} z_i^p \frac{\partial F_i^p}{\partial z}(x) \leq 0 \text{ for all } z \in TX \text{ and all } x \in X.$$

To interpret this expression, recall that the displacement vector $z \in TX$ describes the aggregate effect on the population state of strategy revisions by a small group of agents. The derivative $\frac{\partial F_i^p}{\partial z}(x)$ represents the marginal effect that these revisions have on the payoffs of agents currently choosing strategy $i \in S^p$. Condition (3.16) considers a weighted sum of these effects, with weights given by the changes in the use of each strategy, and requires that this weighted sum be negative.

Intuitively, a game exhibits self-defeating externalities if the improvements in the payoffs of strategies to which revising agents are switching are always exceeded by the improvements in the payoffs of strategies which revising agents are abandoning. For example, suppose the tangent vector z takes the form $z = e_j^p - e_i^p$, representing switches by some members of population p from strategy i to strategy j . In this case, the requirement in condition (3.16) reduces to $\frac{\partial F_j^p}{\partial z}(x) \leq \frac{\partial F_i^p}{\partial z}(x)$: that is, any performance gains that the switches create for the newly chosen strategy j are dominated by the performance gains created for the abandoned strategy i .

Exercise 3.3.2. (i) Characterize the C^1 two-strategy stable games using a derivative condition.

(ii) Recall the Hawk-Dove game introduced in Chapter 2:

$$F^{HD}(x) = \begin{pmatrix} -1 & 2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_H \\ x_D \end{pmatrix} = \begin{pmatrix} 2x_D - x_H \\ x_D \end{pmatrix}.$$

Verify that F is a stable game. Also, fill in the numerical details of the argument from the previous paragraph for this specific choice of payoff function.

Proof of Theorem 3.3.1: To begin, suppose that F is a stable game. Fix $x \in X$ and $z \in TX$; we want to show that $z'DF(x)z \leq 0$. Since F is C^1 , it is enough to consider x in the interior

of X . In this case, $y_\varepsilon = x + \varepsilon z$ lies in X whenever $|\varepsilon|$ is sufficiently small, and so

$$F(y_\varepsilon) = F(x) + DF(x)(y_\varepsilon - x) + o(|y_\varepsilon - x|).$$

by the definition of $DF(x)$. Premultiplying by $y_\varepsilon - x$ and rearranging yields

$$(y_\varepsilon - x)'(F(y_\varepsilon) - F(x)) = (y_\varepsilon - x)'DF(x)(y_\varepsilon - x) + o(|y_\varepsilon - x|^2).$$

Since the left hand side is nonpositive and since $y_\varepsilon - x = \varepsilon z$, it follows that $\varepsilon^2 z' DF(x) z + o(\varepsilon^2) \leq 0$, and hence that $z' DF(x) z \leq 0$.

Next, suppose that condition (3.16) holds. Then if we let $\alpha(t) = ty + (1 - t)x$, the Fundamental Theorem of Calculus implies that

$$\begin{aligned} (y - x)'(F(y) - F(x)) &= (y - x)' \left(\int_0^1 DF(\alpha(t))(y - x) dt \right) \\ &= \int_0^1 (y - x)' DF(\alpha(t))(y - x) dt \leq 0. \blacksquare \end{aligned}$$

Exercise 3.3.3. The derivative condition that characterizes potential games, externality symmetry (3.11), requires that $z' DF(x) \hat{z} = \hat{z}' DF(x) z$. That z and \hat{z} are chosen separately means that $DF(x)$ is treated as a *bilinear form*. Exercise 3.2.6 shows that in order to check that (3.11) holds for all z and \hat{z} in TX , it is enough to show that it holds for all z and \hat{z} in a basis for TX —for example, the set of vectors of the form $e_j^p - e_i^p$.

In contrast, self-defeating externalities (3.16), which requires that $z' DF(x) z \leq 0$, places the same vector z on both sides of $DF(x)$, thus viewing $DF(x)$ as a *quadratic form*. Explain why the conclusion of Exercise 3.2.6 does not extend to the present setting. Also, construct a 3×3 symmetric game A such that $z' Az \leq 0$ whenever z is of the form $e_j^p - e_i^p$ but such that $F(x) = Ax$ is not a stable game.

3.3.2 Examples

Our first example introduces the notion of an evolutionarily stable state (ESSs) in the context of symmetric normal form games. A full treatment of ESSs of general population games is offered in Section 8.3.

Example 3.3.4. *Matching in symmetric normal form games with an interior evolutionarily or neutrally stable state.* Let $A \in \mathbf{R}^{n \times n}$ be a symmetric normal form game. State $x \in X$ is an

evolutionarily stable state (or an *evolutionarily stable strategy*, or simply an *ESS*) of A if

$$(3.17) \quad x'Ax \geq y'Ax \text{ for all } y \in X; \text{ and}$$

$$(3.18) \quad x'Ax = y'Ax \text{ implies that } x'Ay > y'Ay.$$

Condition (3.17) says that x is a symmetric Nash equilibrium of A . Condition (3.18) says that x performs better against any alternative best reply y than y performs against itself. (Alternatively, (3.17) says that no $y \in X$ can strictly invade x , and (3.17) and (3.18) together say that if y can weakly invade x , then x can strictly invade y —see Section 3.3.3 below.) If we weaken condition (3.18) to

$$(3.19) \quad \text{If } x'Ax = y'Ax, \text{ then } x'Ay \geq y'Ay,$$

then a state satisfying conditions (3.17) and (3.19) is called a *neutrally stable state* (NSS).

Suppose that the ESS x lies in the interior of X . Then as x is an interior Nash equilibrium, all pure and mixed strategies are best responses to it: for all $y \in X$, we have that $x'Ax = y'Ax$, or, equivalently, that $(x - y)'Ax = 0$. Next, we can rewrite the inequality in condition (3.18) as $(x - y)'Ay > 0$. Subtracting this last expression from the previous one yields $(x - y)'A(x - y) < 0$. But since x is in the interior of X , all tangent vectors $z \in TX$ are proportional to $x - y$ for some choice of $y \in X$. Therefore, $z'DF(x)z = z'Az < 0$ for all $z \in TX$, and so F is a strictly stable game. Similar reasoning shows that if F admits an interior NSS, then F is a stable game. §

Example 3.3.5. Matching in Rock-Paper-Scissors. In Rock-Paper-Scissors, Paper covers Rock, Scissors cut Paper, and Rock smashes Scissors. If a win in a match is worth $w > 0$, a loss $-l < 0$, and a draw 0, we obtain the symmetric normal form game

$$A = \begin{pmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{pmatrix}, \text{ where } w, l > 0.$$

When $w = l$, we refer to A as (*standard*) *RPS*; when $w > l$, we refer to A as *good RPS*, and when $w < l$, we refer to A as *bad RPS*. In all cases, the unique symmetric Nash equilibrium of A is $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

To determine the parameter values for which this game generates a stable population game, define $d = w - l$. Since $y'Ay = \frac{1}{2}y'(A + A')y$, it is enough to see when the symmetric

matrix

$$\hat{A} = A + A' = \begin{pmatrix} 0 & d & d \\ d & 0 & d \\ d & d & 0 \end{pmatrix}$$

is negative semidefinite with respect to TX . Now \hat{A} has one eigenvalue of $2d$ corresponding to the eigenvector $\mathbf{1}$, and two eigenvalues of $-d$ corresponding to the orthogonal eigenspace TX . Thus, $z'\hat{A}z = -dz'z$ for each $z \in TX$. Since $z'z > 0$ whenever $z \neq \mathbf{0}$, we conclude that F is stable if and only if $d \geq 0$. Thus, good RPS is strictly stable, standard RPS is stable, and bad RPS is not stable. §

Exercise 3.3.6. Matching in wars of attrition. A war of attrition is a two player symmetric normal form game. Strategies represent amounts of time committed to waiting for a scarce resource. If the two players choose times i and $j > i$, then the j player obtains the resource, worth v , while both players pay a cost of c_i : once the first player leaves, the other seizes the resource immediately. If both players choose time i , the resource is split, so payoffs are $\frac{v}{2} - c_i$ each. Show that for any resource value $v \in \mathbf{R}$ and any cost vector $c \in \mathbf{R}^n$ satisfying $c_1 \leq c_2 \leq \dots \leq c_n$, matching in a war of attrition generates a stable game. §

Example 3.3.7. Matching in symmetric zero-sum games. A symmetric two player normal form game A is *symmetric zero-sum* if A is skew-symmetric: that is, if $A_{ji} = -A_{ij}$ for all $i, j \in S$. This condition ensures that under single population matching, the total utility generated in any match is zero. Since payoffs in the resulting single population game are $F(x) = Ax$, we find that $z'DF(x)z = z'Az = 0$ for all vectors $z \in \mathbf{R}^n$, and so F is a null stable game. §

Example 3.3.8. Matching in standard zero-sum games. A two player normal form game $U = (U^1, U^2)$ is *zero-sum* if $U^2 = -U^1$, so that the two players' payoffs always add up to zero. Matching of two populations to play U generates the population game

$$F(x^1, x^2) = \begin{bmatrix} 0 & U^1 \\ (U^2)' & 0 \end{bmatrix} \begin{pmatrix} x^1 \\ x^2 \end{pmatrix} = \begin{bmatrix} 0 & U^1 \\ -(U^1)' & 0 \end{bmatrix} \begin{pmatrix} x^1 \\ x^2 \end{pmatrix}.$$

If z is a vector in $\mathbf{R}^n = \mathbf{R}^{n^1+n^2}$, then

$$z'DF(x)z = \begin{pmatrix} (z^1)' & (z^2)' \end{pmatrix} \begin{bmatrix} 0 & U^1 \\ -(U^1)' & 0 \end{bmatrix} \begin{pmatrix} z^1 \\ z^2 \end{pmatrix} = (z^1)'U^1z^2 - (z^2)'(U^1)'z^1 = 0,$$

so F is a null stable game. §

Exercise 3.3.9. Matching in multi-zero-sum games. Let U be a p player normal form game in which each player $p \in \mathcal{P}$ chooses a single strategy from S^p to simultaneously play a distinct zero-sum contest with each of his $p - 1$ opponents. We call such a U a *multi-zero-sum game*.

- (i) When $p < q$, let $Z^{pq} \in \mathbf{R}^{n^p \times n^q}$ denote player p 's payoff matrix for his zero-sum contest against player q . Define the normal form game U in terms of the Z^{pq} matrices.
- (ii) Let F be the p population game generated by matching in U . Show that $z'DF(x)z = 0$ for all $x \in X$ and $z \in \mathbf{R}^n$, and hence that F is a null stable game.

The previous example and exercise show that matching across multiple populations can generate a null stable game. Proposition 3.3.10 reveals that null stable games are the only stable games that can be generated in this way.

Proposition 3.3.10. *Suppose F is a C^1 stable game without own-population interactions: $F^p(x)$ is independent of x^p for all $p \in \mathcal{P}$. Then F is a null stable game.*

Proof. By Theorem 3.3.1, F is stable if and only if for all $x \in X$, $DF(x)$ is negative semidefinite with respect to TX . This requirement on $DF(x)$ can be restated as (i) $\Phi DF(x)\Phi$ is negative semidefinite (with respect to \mathbf{R}^n); or as (ii) $\Phi(DF(x) + DF(x)')\Phi$ is negative semidefinite, or (since the previous matrix is symmetric) as (iii) $\Phi(DF(x) + DF(x)')\Phi$ has all eigenvalues nonpositive. By similar logic, F is null stable if and only if for all $x \in X$, $\Phi(DF(x) + DF(x)')\Phi$ has all eigenvalues zero (and so is the null matrix).

Let $D^q F^p(x)$ be the (p, q) th block of the derivative matrix $DF(x)$. Since F^p is independent of x^p , it follows that $D^p F^p(x) = \mathbf{0}$, and hence that $\Phi(D^p F^p(x) + D^p F^p(x)')\Phi = \mathbf{0}$. Since this product is the (p, p) th block of the symmetric matrix $\Phi(DF(x) + DF(x)')\Phi$, the latter has zero trace, and so its eigenvalues sum to zero. Therefore, the only way $\Phi(DF(x) + DF(x)')\Phi$ can be negative semidefinite is if all of its eigenvalues are zero. In other words, if F is stable, it is null stable. ■

Proposition 3.3.10 tells us that within-population interactions are required to obtain a strictly stable game. Thus, strictly stable games can arise when there is matching within a single population to play a symmetric normal form game, but not when there is matching in multiple populations to play a standard normal form game.

On the other hand, strictly stable games can arise in multipopulation matching settings that allow matches both across and within populations (see the Notes). Moreover, in general population games—for instance, in congestion games—within-population interactions are the norm, and strictly stable games are not uncommon. Our remaining examples illustrate this point.

Example 3.3.11. (Perturbed) concave potential games. We call $F : X \rightarrow \mathbf{R}^n$ a *concave potential game* if it is a potential game whose potential function $f : X \rightarrow \mathbf{R}$ is concave. Then since $y - x \in TX$, since the orthogonal projection matrix Φ is symmetric, and since $\nabla f \equiv \Phi F$, we find that

$$\begin{aligned} (y - x)'(F(y) - F(x)) &= (\Phi(y - x))'(F(y) - F(x)) \\ &= (y - x)'(\Phi F(y) - \Phi F(x)) \\ &= (y - x)'(\nabla f(y) - \nabla f(x)) \\ &\leq 0, \end{aligned}$$

so F is a stable game. If the inequalities above are satisfied strictly, then they will continue to be satisfied if the payoff functions are slightly perturbed. In other words, perturbations of strictly concave potential games remain strictly stable games. §

Example 3.3.12. Negative dominant diagonal games. We call the full population game F a *negative dominant diagonal game* if it satisfies

$$\frac{\partial F_i^p}{\partial x_i^p}(x) < 0 \text{ and } \left| \frac{\partial F_i^p}{\partial x_i^p}(x) \right| \geq \frac{1}{2} \sum_{(j,q) \neq (i,p)} \left(\left| \frac{\partial F_j^q}{\partial x_i^p}(x) \right| + \left| \frac{\partial F_i^p}{\partial x_j^q}(x) \right| \right)$$

for all $i \in S^p, p \in \mathcal{P}$, and $x \in X$. The first condition says that choosing strategy $i \in S^p$ imposes a negative externality on other users of this strategy. The second condition requires that this externality exceeds the average of (i) the total externalities that strategy i imposes on other strategies and (ii) the total externalities that other strategies impose on strategy i . These conditions are precisely what is required for the matrix $DF(x) + DF(x)'$ to have a negative dominant diagonal. The dominant diagonal condition implies that all of the eigenvalues of $DF(x) + DF(x)'$ are negative; since $DF(x) + DF(x)'$ is also symmetric, it is negative definite. Therefore, $DF(x)$ is negative definite too, and so F is a strictly stable game. §

3.3.3 Invasion

In Section 3.3.4, we introduce new equilibrium concepts that are of basic importance for stable games: global neutral stability and global evolutionary stability. These concepts are best understood in terms of the notion of invasion to be presented now.

Let $F : X \rightarrow \mathbf{R}^n$ be a population game, and let $x, y \in X$ be two social states. We say that y can *weakly invade* x ($y \in \bar{I}_F(x)$) if $(y - x)'F(x) \geq 0$. Similarly, y can *strictly invade* x ($y \in I_F(x)$)

if $(y - x)'F(x) > 0$.

The intuition behind these definitions is simple. Consider a single population of agents who play the game F , and whose initial behavior is described by the state $x \in X$. Now imagine that a very small group of agents decide to switch strategies. After these agents select their new strategies, the distribution of choices within their group is described by some $y \in X$, but since the group is so small the impact of its behavior on the overall population state is negligible. Thus, the average payoff in the invading group is at least as high as that in the incumbent population if $y'F(x) \geq x'F(x)$, or equivalently, if $y \in \bar{I}_F(x)$. Similarly, the average payoff in the invading group exceeds that in the incumbent population if $y \in I_F(x)$.

The interpretation of invasion does not change much when there are multiple populations. If we write $(y - x)'F(x)$ as $\sum_p (y^p - x^p)'F^p(x)$, we see that if $y \in I_F(x)$, there must be some population p for which the small group switching to y^p outperforms the incumbent population playing x^p at social state x .

These stories suggest a link with evolutionary dynamics. If y is any state in X , then the vector $y - x$ is a feasible displacement direction from state x . If in addition $y \in I_F(x)$, then the direction $y - x$ is not only feasible, but also respects the incentives provided by the underlying game.

The invasion conditions also have simple geometric interpretations. That $y \in \bar{I}_F(x)$ means that the angle between the displacement vector $y - x$ and the payoff vector $F(x)$ is weakly acute; if $y \in I_F(x)$, this angle is strictly acute. Figure 3.3.1 sketches the set $I_F(x)$ at various states x in a two strategy game. Figure 3.3.2 does the same for a three-strategy game. To draw the latter case, we need the observation that

$$\begin{aligned} y \in I_F(x) &\Leftrightarrow (y - x)'F(x) > 0 \\ &\Leftrightarrow (\Phi(y - x))'F(x) > 0 \\ &\Leftrightarrow (y - x)'\Phi F(x) > 0. \end{aligned}$$

In other words, $y \in I_F(x)$ if and only if the angle between the displacement vector $y - x$ and the *projected* payoff vector $\Phi F(x)$ is strictly acute.

3.3.4 Global Neutral Stability and Global Evolutionary Stability

Before introducing our new solution concepts, we first characterize Nash equilibrium in terms of invasion: a Nash equilibrium is a state that no other state can strictly invade.

Proposition 3.3.13. $x \in NE(F)$ if and only if $I_F(x) = \emptyset$.




Figure 3.3.1: Invasion in a two strategy game.




Figure 3.3.2: Invasion in a three strategy game.

Proof. $x \in NE(F) \Leftrightarrow (y - x)'F(x) \leq 0$ for all $y \in X \Leftrightarrow I_F(x) = \emptyset$. ■

With this background at hand, we call $x \in X$ a *globally neutrally stable state* (GNSS) if

$$(y - x)'F(y) \leq 0 \text{ for all } y \in X.$$

Similarly, we call x a *globally evolutionarily stable state* (GESS) if

$$(y - x)'F(y) < 0 \text{ for all } y \in X - \{x\}.$$

We let $GNSS(F)$ and $GESS(F)$ denote the sets of globally neutrally stable strategies and globally evolutionarily stable strategies, respectively.

To see the reason for our nomenclature, note that the inequalities used to define GNSS and GESS are the same ones used to define NSS and ESS in symmetric normal form games (Example 3.3.4), but that they are now required to hold not just at those states y that are optimal against x , but at all $y \in X$. NSS and ESS also require a state to be a Nash equilibrium, but our new solution concepts implicitly require this as well—see Proposition 3.3.15 below.

It is easy to describe both of these concepts in terms of the notion of invasion.

Observation 3.3.14. (i) $GNSS(F) = \bigcap_{y \in X} \bar{I}_F(y)$, and so is convex.

(ii) $x \in GESS(F)$ if and only if $x \in \bigcap_{y \in X - \{x\}} I_F(y)$.

In words: a GNSS is a state that can weakly invade every state (or, equivalently, every other state), while a GESS is a state that can strictly invade every other state.

Our new solution concepts can also be described in geometric terms. For example, x is a GESS if a small motion from any state $y \neq x$ in the direction $F(y)$ (or $\Phi F(y)$) moves the state closer to x (see Figure 3.3.3). If we allow not only these acute motions, but also orthogonal motions, we obtain the weaker notion of GNSS.

We conclude this section by relating our new solution concepts to Nash equilibrium.

Proposition 3.3.15. (i) If $x \in GNSS(F)$, then $x \in NE(F)$.

(ii) If $x \in GESS(F)$, then $NE(F) = \{x\}$. Hence, if a GESS exists, it is unique.

Proof. To prove part (i), let $x \in GNSS(F)$ and let $y \neq x$. Define $x_\varepsilon = \varepsilon y + (1 - \varepsilon)x$. Since x is a GNSS, $(x - x_\varepsilon)'F(x_\varepsilon) \geq 0$ for all $\varepsilon \in (0, 1]$. Simplifying and dividing by ε yields $(x - y)'F(x_\varepsilon) \geq 0$ for all $\varepsilon \in (0, 1]$, so taking ε to zero yields $(y - x)'F(x) \leq 0$. In other words, $x \in NE(F)$.




Figure 3.3.3: The geometric definition of GNSS.




Figure 3.3.4: Why every GNSS is a Nash equilibrium.

To prove part (ii), it is enough to show that if x is a GESS, then no $y \neq x$ is Nash. But if $x \in \text{GESS}(F)$, then $x \in I_F(y)$; since $I_F(y)$ is nonempty, $y \notin \text{NE}(F)$. ■

Evidently, this proposition implies that every GNSS is an NSS, and that every GESS is an ESS.

The proof that every GNSS is Nash is easy to explain in pictures. In Figure 3.3.4, we draw the GNSS x and an arbitrary state y , and place the state x_ε on the segment between y and x . Since x is a GNSS, the angle between $F(x_\varepsilon)$ and $x - x_\varepsilon$, and hence between $\Phi F(x_\varepsilon)$ and $x - x_\varepsilon$, is weakly acute. Taking ε to zero, it is apparent that the angle between $\Phi F(x)$ and $y - x$, and hence between $y - x$ and $\Phi F(x)$, must be weakly obtuse. Since y was arbitrary, x is a Nash equilibrium.

3.3.5 Nash Equilibrium and Global Neutral Stability in Stable Games

Proposition 3.3.15 tells us that every GNSS of an arbitrary game F is a Nash equilibrium. Theorem 3.3.16 shows that much more can be said if F is stable: in these case, the sets of globally neutrally stable states and Nash equilibria coincide. Together, this fact and Observation 3.3.14 imply that the Nash equilibria of any stable game form a convex set. In fact, if we can replace certain of the weak inequalities that define stable games with strict ones, then the Nash equilibrium is actually unique.

Theorem 3.3.16. (i) *If F is a stable game, then $\text{NE}(F) = \text{GNSS}(F)$, and so is convex.*

(ii) *If in addition F is strictly stable at some $x \in \text{NE}(F)$ (that is, if $(y - x)'(F(y) - F(x)) < 0$ for all $y \neq x$), then $\text{NE}(F) = \text{GESS}(F) = \{x\}$.*

Proof. Suppose that F is stable, and let $x \in \text{NE}(F)$. To establish part (i), it is enough to show that $x \in \text{GNSS}(F)$. So fix an arbitrary $y \neq x$. Since F is stable,

$$(3.20) \quad (y - x)'(F(y) - F(x)) \leq 0.$$

And since $x \in \text{NE}(F)$, $(y - x)'F(x) \leq 0$. Adding these inequalities yields

$$(3.21) \quad (y - x)'F(y) \leq 0,$$

As y was arbitrary, x is a GNSS.

Turning to part (ii), suppose that F is strictly stable at x . Then inequality (3.20) holds strictly, so inequality (3.21) holds strictly as well. This means that x is a GESS of F , and hence the unique Nash equilibrium of F . ■




Figure 3.3.5: The GESS of good RPS.

Example 3.3.17. Rock-Paper-Scissors revisited. Recall from Example 3.3.5 that good RPS is a (strictly) stable game; standard RPS is a zero-sum game, and hence a (weakly) stable game. The unique Nash equilibrium of both of games is $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. In Figure 3.3.5, for a selection of states x , we draw the projected payoff vectors $\Phi F(x)$ generated by good RPS (with $w = 3$ and $l = 1$), as well as the vector from x to x^* . For each x , the angle between this pair of vectors is acute, reflecting the fact that the Nash equilibrium x^* is a GESS. In Figure 3.3.6, we perform the same exercise for standard RPS. In this case, the vectors $\Phi F(x)$ and $x^* - x$ always form a right angle, so x^* is a GNSS but not a GESS. §

Exercise 3.3.18. Let F be a stable game. Show that if x^* is a Nash equilibrium of F such that $DF(x^*)$ is negative definite with respect to $TX \times TX$, then x^* is a GESS, and hence the unique Nash equilibrium of F .

Exercise 3.3.19. Pseudostable games. We call the population game F *pseudostable* if for all $x, y \in X$, $(y - x)'F(x) \leq 0$ implies that $(x - y)'F(y) \geq 0$. In other words, if y cannot strictly invade x , then x can weakly invade y .

- (i) Show that every stable game is pseudostable.
- (ii) Show that if F is pseudostable, then $NE(F) = GNSS(F)$, and so is convex.




Figure 3.3.6: The GNSS of standard RPS.

(A smooth real-valued function f is *pseudoconcave* if its gradient ∇f is pseudostable. Given facts (i) and (ii) above and the discussion in Section 3.1.5, it should be no surprise that many results from concave programming (e.g., the convexity of the set of maximizers) remain true when the objective function is only pseudoconcave.)

In addition to its role in establishing that the set of Nash equilibria of a stable game is convex, the concept of global neutral stability enables us to carry out an important theoretical exercise: that of devising an elementary proof of existence of Nash equilibrium in stable games—in other words, one that does not rely on an appeal to a fixed point theorem. The heart of the proof, Proposition 3.3.20, is a finite analogue of the result we seek.

Proposition 3.3.20. *Let F be a stable game, and let Y be a finite subset of X . Then there exists a state $x^* \in \text{conv}(Y)$ such that $(y - x^*)'F(y) \leq 0$ for all $y \in Y$.*

In words: if F is a stable game, then given any finite set of states Y , we can always find a state in the convex hull of Y that can weakly invade every element of Y . The proof of this result uses the Minmax Theorem.

Proof. Suppose that Y has m elements. Define a two player zero-sum game $U =$

$(U^1, U^2) = (Z, -Z)$ with $n^1 = n^2 = m$ as follows:

$$Z_{xy} = (x - y)'F(y).$$

In this game, player 2 chooses a “status quo” state $y \in Y$, player 1 chooses an “invader” $x \in Y$, and the payoff Z_{xy} is the invader’s “relative payoff” in F . Split Z into its symmetric and skew-symmetric parts:

$$Z^S = \frac{1}{2}(Z + Z') \text{ and } Z^{SS} = \frac{1}{2}(Z - Z').$$

Since F is stable, equation (3.20) from the previous proof shows that

$$Z_{xy}^S = \frac{1}{2}((x - y)'F(y) + (y - x)'F(x)) = \frac{1}{2}(x - y)'(F(y) - F(x)) \geq 0$$

for all $x, y \in Y$.

The Minmax Theorem tells us that in any zero sum game, player 1 has a strategy that guarantees him the value of the game. In the skew-symmetric game $U^{SS} = (Z^{SS}, -Z^{SS}) = (Z^{SS}, (Z^{SS})')$, the player roles are interchangeable, so the game’s value must be zero. Since $Z = Z^{SS} + Z^S$ and $Z^S \geq 0$, the value of $U = (Z, -Z)$ must be *at least* zero. In other words, if $\lambda \in \mathbf{R}^m$ is a maxmin strategy for player 1, then

$$\sum_{x \in Y} \sum_{y \in Y} \lambda_x Z_{xy} \mu_y \geq 0$$

for all mixed strategies μ of player 2. If we let

$$x^* = \sum_{x \in Y} \lambda_x x \in \text{conv}(Y)$$

and fix an arbitrary pure strategy $y \in Y$ for player 2, we find that

$$0 \leq \sum_{x \in Y} \lambda_x Z_{xy} = \sum_{x \in Y} \lambda_x (x - y)'F(y) = (x^* - y)'F(y). \blacksquare$$

With this result in hand, existence of Nash equilibrium in stable games follows from a simple compactness argument. Theorem 3.3.16 and Observation 3.3.14 tell us that

$$NE(F) = GNSS(F) = \bigcap_{y \in X} \{x \in X : (y - x)'F(y) \leq 0\}.$$

Proposition 3.3.20 shows that if we take the intersection above over an arbitrary finite set

$Y \subset X$ instead of over X itself, then the intersection is nonempty. Since X is compact, the finite intersection property (see the Notes) allows us to conclude that $GNSS(F)$ is nonempty itself.

Exercise 3.3.21. In Exercise 3.1.14, we defined population games with entry and exit. If $F : \mathbf{R}_+^n \rightarrow \mathbf{R}$ is C^1 and defines such a game, what condition on the derivative matrices $DF(x)$ is the appropriate definition of stable games for this context? Argue that all of the results in this section continue to hold when entry and exit are permitted.

3.4 Supermodular Games

Of the classes of games we study in this chapter, supermodular games, a class that includes models of coordination, search, and Bertrand competition, are the most familiar to economists. By definition, supermodularity requires that higher choices by one's opponents make one's own higher strategies look relatively more desirable. This complementarity condition imposes a monotone structure on the agents' best response correspondences, which in turn imposes structure on the set of Nash equilibria.

3.4.1 Definition

Each strategy set $S^p = \{1, \dots, n^p\}$ is naturally endowed with a linear order. To define supermodular games, we introduce a corresponding partial order on the set of population states X^p (and, implicitly, on the set of mixed strategies for population p). Define the matrix $\Sigma \in \mathbf{R}^{(n^p-1) \times n^p}$ by

$$\Sigma = \begin{pmatrix} 0 & 1 & \cdots & 1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 \end{pmatrix}.$$

Then

$$(\Sigma x^p)_i = \sum_{j=i+1}^{n^p} x_j^p$$

equals the total mass on strategies greater than i at population state x^p . If we view x^p as a discrete density function on S^p with total mass m^p , then Σx^p defines the corresponding "decumulative distribution function" for x^p . In particular, $\Sigma y^p \geq \Sigma x^p$ if and only if y^p stochastically dominates x^p .

We extend this partial order to all of X using the matrix $\Sigma \in \mathbf{R}^{(n-p) \times n}$, which we define as the block diagonal matrix $\Sigma = \text{diag}(\Sigma, \dots, \Sigma)$. Note that $\Sigma y \geq \Sigma x$ if and only if y^p stochastically dominates x^p for all $p \in \mathcal{P}$.

With these preliminaries in hand, we are ready to define our class of games.

Definition. We call the population game $F : X \rightarrow \mathbf{R}^n$ a supermodular game if it exhibits strategic complementarities:

$$(3.22) \quad \text{If } \Sigma y \geq \Sigma x, \text{ then } F_{i+1}^p(y) - F_i^p(y) \geq F_{i+1}^p(x) - F_i^p(x) \text{ for all } i < n^p, p \in \mathcal{P}, x \in X.$$

In words: if y stochastically dominates x , then for any strategy $i < n^p$, the payoff advantage of $i + 1$ over i is greater at y than at x .

By introducing a bit more notation, we can express condition (3.22) in a more concise way. Define the matrices $\tilde{\Sigma} \in \mathbf{R}^{n^p \times (n^p - 1)}$ and $\tilde{\Sigma} \in \mathbf{R}^{n \times (n - p)}$ by

$$\tilde{\Sigma} = \begin{pmatrix} -1 & 0 & \cdots & 0 \\ 1 & -1 & \ddots & \vdots \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & -1 \\ 0 & \ddots & 0 & 1 \end{pmatrix} \quad \text{and} \quad \tilde{\Sigma} = \text{diag}(\tilde{\Sigma}, \dots, \tilde{\Sigma}).$$

Observation 3.4.1. F is a supermodular game if and only if the following condition holds:

$$(3.23) \quad \Sigma y \geq \Sigma x \text{ implies that } \tilde{\Sigma}' F(y) \geq \tilde{\Sigma}' F(x).$$

As with potential games and stable games, we can characterize smooth supermodular games in terms of conditions on the derivatives $DF(x)$.

Theorem 3.4.2. Suppose the population game F is C^1 . Then F is supermodular if and only if either of the following equivalent conditions holds.

$$(3.24) \quad \frac{\partial(F_{i+1}^p - F_i^p)}{\partial(e_{j+1}^q - e_j^q)}(x) \geq 0 \text{ for all } i < n^p, j < n^q, p, q \in \mathcal{P}, \text{ and } x \in X.$$

$$(3.25) \quad \tilde{\Sigma}' DF(x) \tilde{\Sigma} \geq \mathbf{0} \text{ for all } x \in X.$$

Condition (3.24) is the most transparent of the four conditions. It requires that if some players in population q switch from strategy j to strategy $j + 1$, the performance of strategy $i + 1 \in S^p$ improves relative to that of strategy i . On the other hand, condition

(3.25) provides the most concise characterization of supermodular games. Moreover, since the range of $\tilde{\Sigma}$ is TX (i.e., since each column of $\tilde{\Sigma}$ lies in TX), condition (3.25) is a restriction of the action of $DF(x)$ on $TX \times TX$ —just like our earlier conditions (3.11) and (3.16) characterizing potential games and stable games.

Proof. The equivalence of conditions (3.24) and (3.25) is easily verified. Given Observation 3.4.1, it is enough to show that (3.22) implies (3.24) and that (3.25) implies (3.23).

So suppose condition (3.22) holds, and fix $x \in X$; since F is C^1 it is enough to consider x in the interior of X . Let $y_\varepsilon = x + \varepsilon(e_{j+1}^q - e_j^q)$, which lies in X whenever $|\varepsilon|$ is sufficiently small, and which satisfies $\Sigma y_\varepsilon \geq \Sigma x$. By the definition of $DF(x)$, we have that

$$F_{i+1}^p(y_\varepsilon) - F_i^p(y_\varepsilon) = F_{i+1}^p(x) - F_i^p(x) + \varepsilon \frac{\partial(F_{i+1}^p - F_i^p)}{\partial(e_{j+1}^q - e_j^q)}(x) + o(|y_\varepsilon - x|).$$

Thus, condition (3.22) implies that

$$\varepsilon \frac{\partial(F_{i+1}^p - F_i^p)}{\partial(e_{j+1}^q - e_j^q)}(x) + o(|\varepsilon|) \geq 0,$$

which implies (3.24).

We now show that (3.25) implies (3.23). We consider only the single population case, leaving the general case as an exercise. The idea behind the proof is simple. If state y stochastically dominates state x , then we can transit from state x to state y by shifting mass from strategy 1 to strategy 2, from strategy 2 to strategy 3, ..., and finally from strategy $n - 1$ to strategy n . Condition (3.24) \equiv (3.25) says that each such shift improves the payoff of each strategy $k + 1$ relative to that of strategy k . Since transiting from x to y means executing all of the shifts, this transition too must improve the performance of $k + 1$ relative to k , which is exactly what condition (3.22) \equiv (3.23) requires.

Our matrix notation makes it possible to formalize this argument in a streamlined way. Recall the definitions of $\tilde{\Sigma} \in \mathbf{R}^{n \times (n-1)}$ and $\Sigma \in \mathbf{R}^{(n-1) \times n}$, and define $\Omega \in \mathbf{R}^{n \times n}$ as follows:

$$\tilde{\Sigma} = \begin{pmatrix} -1 & 0 & \cdots & 0 \\ 1 & -1 & \ddots & \vdots \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & -1 \\ 0 & \ddots & 0 & 1 \end{pmatrix}, \Sigma = \begin{pmatrix} 0 & 1 & \cdots & \cdots & 1 \\ 0 & 0 & 1 & & \vdots \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & 0 & 1 \end{pmatrix}, \text{ and } \Omega = \begin{pmatrix} 1 & 1 & \cdots & \cdots & 1 \\ 0 & 0 & \cdots & \cdots & 0 \\ 0 & 0 & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \cdots & 0 \end{pmatrix}$$

Then it is easy to verify this next observation.

Observation 3.4.3. $\tilde{\Sigma}\Sigma = I - \Omega \in \mathbf{R}^{n \times n}$.

In words, Observation 3.4.3 says that the stochastic dominance operator Σ is “inverted” by the difference operator $\tilde{\Sigma}$, except for a remainder Ω that is a null operator on TX (i.e., that satisfies $\Omega z = \mathbf{0}$ for all $z \in TX$). (For completeness, we also note that $\Sigma\tilde{\Sigma} = I \in \mathbf{R}^{(n-1) \times (n-1)}$.)

Now suppose that $\Sigma x \leq \Sigma y$, and let $\alpha(t) = ty + (1-t)x$, so that $\alpha(0) = x$, $\alpha(1) = y$, and $\alpha'(t) = y - x \in TX$. Then using the Fundamental Theorem of Calculus, Observation 3.4.3, condition (3.25), and the fact that $\Sigma(y - x) \geq 0$, we find that

$$\begin{aligned} \tilde{\Sigma}'(F(y) - F(x)) &= \tilde{\Sigma}' \int_0^1 DF(\alpha(t))(y - x) dt \\ &= \int_0^1 \tilde{\Sigma}' DF(\alpha(t)) (\tilde{\Sigma}\Sigma + \Omega)(y - x) dt \\ &= \int_0^1 (\tilde{\Sigma}' DF(\alpha(t)) \tilde{\Sigma}) \Sigma(y - x) dt \\ &\geq 0. \blacksquare \end{aligned}$$

3.4.2 Examples

Exercise 3.4.4. Matching in supermodular normal form games. The normal form game $U = (U^1, \dots, U^p)$ is *supermodular* if the difference $U^p(s^p + 1, s^q, s^{-\{p,q\}}) - U^p(s^p, s^q, s^{-\{p,q\}})$ is nondecreasing in s^q for all $s^p < n^p$, $s^{-\{p,q\}} \in \prod_{r \notin \{p,q\}} S^r$ and distinct $p, q \in \mathcal{P}$. Show that matching of p populations to play U generates a supermodular game.

Exercise 3.4.5. Which symmetric normal form games generate supermodular population games?

Example 3.4.6. Bertrand oligopoly with differentiated products. A population of firms produce output at zero marginal cost and compete in prices $S = \{1, \dots, n\}$. Suppose that the demand faced by a firm increases when competitors raise their prices, and that this effect does not diminish when the firm itself charges higher prices. More precisely, let $q_i(x)$, the demand faced by a firm that charges price i when the price distribution is x , satisfy

$$\frac{\partial q_i}{\partial(e_{j+1} - e_j)}(x) \geq 0 \quad \text{and} \quad \frac{\partial(q_{k+1} - q_k)}{\partial(e_{j+1} - e_j)}(x) \geq 0 \quad \text{for all } i \leq n \text{ and all } j, k < n.$$

The payoff to a firm that charges price i is $F_i(x) = i q_i(x)$, and so

$$\frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) = (i+1) \frac{\partial q_{i+1}}{\partial(e_{j+1} - e_j)}(x) - i \frac{\partial q_i}{\partial(e_{j+1} - e_j)}(x)$$

$$= i \frac{\partial(q_{i+1} - q_i)}{\partial(e_{j+1} - e_j)}(x) + \frac{\partial q_{i+1}}{\partial(e_{j+1} - e_j)}(x) \geq 0.$$

Therefore, F is a supermodular game. §

Example 3.4.7. Search with positive externalities. A population of agents choose levels of search effort in $S = \{1, \dots, n\}$. The payoff to choosing effort i is

$$F_i(x) = m(i) b(a(x)) - c(i),$$

where $a(x) = \sum_{k \leq n} kx_k$ is the aggregate search effort, b is some increasing benefit function, m is an increasing multiplier function, and c is an arbitrary cost function. Notice that the benefits from searching are increasing in both own search effort and in the aggregate search effort. Since

$$\begin{aligned} \frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) &= m(i+1) b'(a(x)) ((j+1) - j) - m(i) b'(a(x)) ((j+1) - j) \\ &= (m(i+1) - m(i)) b'(a(x)) \geq 0, \end{aligned}$$

F is a supermodular game. §

Example 3.4.8. Relative consumption effects/Arms races. Agents from a single population choose consumption levels (or armament levels) in $S = \{1, \dots, n\}$. Payoffs take the form

$$F_i(x) = r(i - a(x)) + u(i) - c(i).$$

Here, r is a concave function of the difference between the agent's consumption level and the average consumption level in the population, while u and c are arbitrary functions of the consumption level. (One would typically assume that r is increasing, but this property is not needed for supermodularity.) Since

$$\begin{aligned} \frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) &= r'((i+1) - a(x)) (-(j+1) + j) - r'(i - a(x)) (-(j+1) + j) \\ &= r'(i - a(x)) - r'((i+1) - a(x)) \geq 0, \end{aligned}$$

F is a supermodular game. §

Exercise 3.4.9. Characterize the C^1 two-strategy supermodular games using a derivative condition. Compare them with the C^1 two-strategy stable games (Exercise 3.3.2(i)). Are all C^1 two-strategy games in one class or the other?

3.4.3 Best Response Monotonicity in Supermodular Games

Recall the definition of the pure best response correspondence for population p :

$$b^p(x) = \operatorname{argmax}_{i \in S^p} F_i^p(x).$$

Theorem 3.4.10 establishes a fundamental property of supermodular games: their pure best response correspondences are increasing.

Theorem 3.4.10. *Let F be a supermodular game with pure best response correspondences b^p . If $\Sigma x \leq \Sigma y$, then $\min b^p(x) \leq \min b^p(y)$ and $\max b^p(x) \leq \max b^p(y)$ for all $p \in \mathcal{P}$.*

This property is intuitively obvious: when opponents choose higher strategies, an agent's own higher strategies look relatively better, so his best strategies must be (weakly) higher as well.

Proof. We consider the case in which $p = 1$, focusing on the first inequality; we leave the remaining cases as exercises.

Let $\Sigma x \leq \Sigma y$ and $i < j$. Then condition (3.22) implies that

$$(F_j(y) - F_i(y)) - (F_j(x) - F_i(x)) = \sum_{k=i}^{j-1} ((F_{k+1}(y) - F_k(y)) - (F_{k+1}(x) - F_k(x))) \geq 0.$$

Thus, if $j = \min b(x) > i$, then $F_j(y) - F_i(y) \geq F_j(x) - F_i(x) > 0$, so i is not a best response to y . As $i < \min b(x)$ was arbitrary, we conclude that $\min b(x) \leq \min b(y)$. ■

To state a version of Theorem 3.4.10 for mixed best responses, we need some additional notation. Let $v_i^p \in \mathbf{R}^n$ denote the i th vertex of the simplex Δ^p : that is, $(v_i^p)_j$ equals 1 if $j = i$ and equals 0 otherwise. (To summarize our notation to date: $x_i^p \in \mathbf{R}$, $v_i^p \in \mathbf{R}^n$, and $e_i^p \in \mathbf{R}^n$. Of course, the notation v_i^p is unnecessary in the single population case.) We can describe population p 's mixed best response correspondence in the following equivalent ways:

$$\begin{aligned} B^p(x) &= \{x^p \in \Delta^p : x_i^p > 0 \Rightarrow i \in b^p(x)\} \\ &= \operatorname{conv}(\{v_i^p : i \in b^p(x)\}), \end{aligned}$$

We can also define the minimal and maximal elements of $B^p(x)$ as follows:

$$\underline{B}^p(x) = v_{\min b^p(x)}^p \quad \text{and} \quad \bar{B}^p(x) = v_{\max b^p(x)}^p.$$

To extend this notation to the multipopulation environment, define

$$\underline{B}(x) = (\underline{B}^1(x), \dots, \underline{B}^p(x)) \quad \text{and} \quad \bar{B}(x) = (\bar{B}^1(x), \dots, \bar{B}^p(x)).$$

Then the following corollary follows immediately from Theorem 3.4.10.

Corollary 3.4.11. *If F is supermodular and $\Sigma x \leq \Sigma y$, then $\Sigma \underline{B}(x) \leq \Sigma \underline{B}(y)$ and $\Sigma \bar{B}(x) \leq \Sigma \bar{B}(y)$.*

3.4.4 Nash Equilibria of Supermodular Games

We now use the monotonicity of the best response correspondence to show that every supermodular game has a minimal and a maximal Nash equilibrium. The derivation of this result includes a finite iterative method for computing the minimal and maximal equilibria, and so provides a simple proof of the existence of equilibrium. We focus attention on the case where each population has mass one, so that each set of population states X^p is just the simplex in \mathbf{R}^{n^p} ; the extension to the general case is a simple but notationally cumbersome exercise.

Let \underline{x} and \bar{x} be the minimal and maximal states in $X : \underline{x}^p = v_1^p$ and $\bar{x}^p = v_{n^p}^p$ for all $p \in \mathcal{P}$. Recall that X_v denotes the set of vertices of X , and let $n^* = \#X_v = \prod_{p \in \mathcal{P}} n^p$. Finally, for states $y, z \in X$, define the interval $[y, z] \subseteq X$ by $[y, z] = \{x \in X : \Sigma y \leq \Sigma x \leq \Sigma z\}$.

Theorem 3.4.12. *Suppose F is a supermodular game. Then*

- (i) *The sequences $\{\underline{B}^k(\underline{x})\}_{k \geq 0}$ and $\{\bar{B}^k(\bar{x})\}_{k \geq 0}$ are monotone sequences in X_v , and so converge within n^* steps to their limits, \underline{x}^* and \bar{x}^* .*
- (ii) *$\underline{x}^* = \underline{B}(\underline{x}^*)$ and $\bar{x}^* = \bar{B}(\bar{x}^*)$, so \underline{x}^* and \bar{x}^* are pure Nash equilibria of F .*
- (iii) *$NE(F) \subseteq [\underline{x}^*, \bar{x}^*]$. Thus, if $\underline{x}^* = \bar{x}^*$, then this state is the Nash equilibrium of F .*

In short, iterating \underline{B} and \bar{B} from the minimal and maximal states in X yields Nash equilibria of F , and all other Nash equilibria of F lie between the two so obtained.

Proof. Part (i) follows immediately from Corollary 3.4.11. To prove part (ii), note that since $\underline{x}^* = \underline{B}^{n^*}(\underline{x})$ and $\underline{B}^{n^*+1}(\underline{x}) = \underline{B}^{n^*}(\underline{x})$ by part (i), it follows that

$$\underline{B}(\underline{x}^*) = \underline{B}(\underline{B}^{n^*}(\underline{x})) = \underline{B}^{n^*+1}(\underline{x}) = \underline{B}^{n^*}(\underline{x}) = \underline{x}^*.$$

An analogous argument shows that $\bar{B}(\bar{x}^*) = \bar{x}^*$.

We finish with the proof of part (iii). If $Y \subseteq X$ and $\min Y$ and $\max Y$ exist, then the monotonicity of B implies that $B(Y) \subseteq [\underline{B}(\min Y), \bar{B}(\max Y)]$. Iteratively applying B to the set X therefore yields $B^{n^*}(X) \subseteq [\underline{B}^{n^*}(\underline{x}), \bar{B}^{n^*}(\bar{x})] = [\underline{x}^*, \bar{x}^*]$. Also, if $x \in NE(F)$, then $x \in B(x)$,

and so $B^{k-1}(x) \subseteq B^{k-1}(B(x)) = B^k(x)$, implying that $x \in B^k(x)$ for all $k \geq 1$. We therefore conclude that $x \in B^{n^*}(x) \subseteq B^{n^*}(X) \subseteq [\underline{x}^*, \bar{x}^*]$. ■

Appendix

3.A Multivariate Calculus

3.A.1 Univariate Calculus

Before discussing multivariate calculus we review some ideas from univariate calculus. A function f from the real line to itself is *differentiable* at the point x if

$$f'(x) = \lim_{y \rightarrow x} \frac{f(y) - f(x)}{y - x}$$

exists; this limit is called the *derivative* of f at x . Three useful facts about derivatives are

The Product Rule:

$$(fg)'(x) = f(x)g'(x) + g(x)f'(x);$$

The Chain Rule:

$$(g \circ f)'(x) = g'(f(x))f'(x);$$

The Fundamental Theorem of Calculus:

$$f(y) - f(x) = \int_x^y f'(z) dz.$$

The definition of $f'(x)$ above is equivalent to the requirement that

$$(3.26) \quad f(y) = f(x) + f'(x)(y - x) + o(y - x),$$

where $o(z)$ represents a remainder function $r : \mathbf{R} \rightarrow \mathbf{R}$ satisfying

$$\lim_{z \rightarrow 0} \frac{r(z)}{z} = 0.$$

(In words: $r(z)$ approaches zero faster than z approaches zero.) In the approximation (3.26), $f'(x)$ acts as a linear map from \mathbf{R} to itself; it sends the displacement of the input, $y - x$, to the displacement of the output, $f'(x)(y - x)$.

3.A.2 The Derivative as a Linear Map

Let $L(\mathbf{R}^n, \mathbf{R}^m)$ denote the space of linear maps from \mathbf{R}^n to \mathbf{R}^m :

$$L(\mathbf{R}^n, \mathbf{R}^m) = \{\lambda : \mathbf{R}^n \rightarrow \mathbf{R}^m \mid \lambda(az + b\hat{z}) = a\lambda(z) + b\lambda(\hat{z}) \text{ for all } a, b \in \mathbf{R} \text{ and } z, \hat{z} \in \mathbf{R}^n\}.$$

Each matrix $A \in \mathbf{R}^{m \times n}$ defines a linear map in $L(\mathbf{R}^n, \mathbf{R}^m)$ via $\lambda(z) = Az$, and such a matrix can be found for every map λ in $L(\mathbf{R}^n, \mathbf{R}^m)$ (see Appendix 3.B.1). It is common to identify a linear map with its matrix representation. But it is important to be aware of the distinction between these two objects: if we replace the domain \mathbf{R}^n with a proper subspace of \mathbf{R}^n , matrix representations of linear maps are no longer unique—see Appendix 3.B.

Let F be a function from \mathbf{R}^n to \mathbf{R}^m . (Actually, we can replace the domain \mathbf{R}^n with any open set in \mathbf{R}^n , or even with a closed set in \mathbf{R}^n , as discussed in Appendix 3.A.7.) We say that F is *differentiable* at x if there is a linear map $DF(x) \in L(\mathbf{R}^n, \mathbf{R}^m)$ satisfying

$$(3.27) \quad F(y) = F(x) + DF(x)(y - x) + o(y - x)$$

Here, $o(z)$ represents a remainder function $r : \mathbf{R}^n \rightarrow \mathbf{R}^m$ that satisfies

$$\lim_{z \rightarrow \mathbf{0}} \frac{r(z)}{|z|} = \mathbf{0}.$$

If the function $DF : \mathbf{R}^n \rightarrow L(\mathbf{R}^n, \mathbf{R}^m)$ is continuous, we say that F is *continuously differentiable* or of class C^1 .

When we view $DF(x)$ as a matrix in $\mathbf{R}^{m \times n}$, we call it the *Jacobian matrix* or *derivative matrix* of F at x . To express this matrix explicitly, define the *partial derivatives* of F at x by

$$\frac{\partial F_i}{\partial x_j}(x) = \lim_{y_j \rightarrow x_j} \frac{F_i(y_j, x_{-j}) - F_i(x)}{y_j - x_j}.$$

Then the derivative matrix $DF(x)$ can be expressed as

$$DF(x) = \begin{pmatrix} \frac{\partial F_1}{\partial x_1}(x) & \cdots & \frac{\partial F_1}{\partial x_n}(x) \\ \vdots & \vdots & \vdots \\ \frac{\partial F_m}{\partial x_1}(x) & \cdots & \frac{\partial F_m}{\partial x_n}(x) \end{pmatrix}.$$

If f is a function from \mathbf{R}^n to \mathbf{R} (i.e., if $m = 1$), then its derivative at x can be represented by a vector. We call this vector the *gradient* of f at x , and define it by

$$\nabla f(x) = \begin{pmatrix} \frac{\partial f}{\partial x_1}(x) \\ \vdots \\ \frac{\partial f}{\partial x_n}(x) \end{pmatrix}.$$

Our notations for derivatives are related by $Df(x) = \nabla f(x)'$, where the prime represents

transposition, and also by

$$DF(x) = \begin{pmatrix} \nabla F_1(x)' \\ \vdots \\ \nabla F_m(x)' \end{pmatrix}.$$

Suppose we are interested in how quickly the value of f changes as we move from the point $x \in \mathbf{R}^n$ in the direction $z \in \mathbf{R}^n - \{\mathbf{0}\}$. This rate is described by the *directional derivative* of f at x in direction z , defined by

$$(3.28) \quad \frac{\partial f}{\partial z}(x) = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon z) - f(x)}{\varepsilon}.$$

It is easy to verify that

$$\frac{\partial f}{\partial z}(x) = \nabla f(x)'z.$$

More generally, the rate of change of the vector-valued function F at x in direction z can be expressed as $DF(x)z$.

It is worth noting that a function can admit directional derivatives at x in every direction $z \neq \mathbf{0}$ without being differentiable at x (i.e., without satisfying definition (3.27)). Amazingly, such a function need not even be continuous at x , as the following example shows.

Example 3.A.1. Define the function $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ by

$$f(x_1, x_2) = \begin{cases} \frac{x_1(x_2)^2}{(x_1)^2 + (x_2)^4} & \text{if } x_1 \neq 0, \\ 0 & \text{if } x_1 = 0. \end{cases}$$

Using definition (3.28), it is easy to verify that the directional derivatives of f at the origin in every direction $z \neq \mathbf{0}$ exist:

$$\frac{\partial f}{\partial z}(\mathbf{0}) = \begin{cases} \frac{(z_2)^2}{z_1} & \text{if } z_1 \neq 0, \\ 0 & \text{if } z_1 = 0. \end{cases}$$

But while $f(\mathbf{0}) = 0$, $f(x) = \frac{1}{2}$ at all other x that satisfy $x_1 = (x_2)^2$, and so f is discontinuous at $\mathbf{0}$. ■

On the other hand, if all (or even all but one) of the partial derivatives f exist and are continuous in a neighborhood of x , then f is differentiable at x .

3.A.3 Differentiation as a Linear Operation

We can view differentiation as an operation that takes functions as inputs and returns functions as outputs. From this point of view, differentiation is a linear operation between spaces of functions. As an example, suppose that f and g are functions from \mathbf{R} to itself, and that a and b are real numbers. Then the scalar product af is a function from \mathbf{R} to itself, as is the linear combination $af + bg$. (In other words, the set of functions from \mathbf{R} to itself is a *vector space*.) The fact that differentiation is linear means that the derivative of the linear combination, $(af + bg)'$, is equal to the linear combination of the derivatives, $af' + bg'$.

We can express this idea in a multivariate setting using a simple formula. Suppose that $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a differentiable function and that A is a matrix in $\mathbf{R}^{l \times m}$. Then AF is the function from \mathbf{R}^n to \mathbf{R}^l defined by $(AF)_k(x) = \sum_{j=1}^m A_{kj}F_j(x)$ for $k \in \{1, \dots, l\}$. Linearity of differentiation says that $D(AF) = A(DF)$, or, more explicitly, that

$$\text{Linearity of differentiation: } D(AF)(x) = A(DF)(x) \text{ for all } x \in \mathbf{R}^n.$$

Put differently, the differential operator D and the linear map A commute.

3.A.4 The Product Rule and the Chain Rule

Suppose f and g are differentiable functions from \mathbf{R} to itself. Then the product rule tells us that $(fg)'(x) = f(x)g'(x) + g(x)f'(x)$. In other words, to find the effect of changing x on the value $(fg)(x)$ of the product function, first find the effect of changing x on $g(x)$, and scale this effect by $f(x)$; then, find the effect of changing x on $f(x)$, and scale this effect by $g(x)$; and finally, add the two terms.

This same idea can be applied in multidimensional cases as well. Let $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and $G : \mathbf{R}^m \rightarrow \mathbf{R}^l$ be differentiable vector-valued functions. Then $F'G : \mathbf{R}^n \rightarrow \mathbf{R}^l$, defined by $(F'G)(x) = F'(x)G(x)$, is a scalar-valued function. The derivative $D(F'G)(x) \in \mathbf{R}^{1 \times n}$ of our new function is described by the following product rule:

$$\text{Product Rule 1: } D(F'G)(x) = (\nabla(F'G)(x))' = F'(x)DG(x) + G(x)'DF(x).$$

(Notice that in the previous paragraph, a prime (') denoted the derivative of a scalar-valued function, while here it denotes matrix transposition. So long as we keep these scalar and matrix usages separate, no confusion should arise.)

If $a : \mathbf{R}^n \rightarrow \mathbf{R}$ is a differentiable scalar-valued function, then $aF : \mathbf{R}^n \rightarrow \mathbf{R}^m$, defined by $(aF)(x) = a(x)F(x)$, is a vector-valued function. Its derivative $D(aF)(x) \in \mathbf{R}^{m \times n}$ is described our next product rule:

$$\text{Product Rule 2: } D(aF)(x) = a(x)DF(x) + F(x)\nabla a(x)' = a(x)DF(x) + F(x)Da(x).$$

Finally, we can create a vector-valued function from $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and $G : \mathbf{R}^n \rightarrow \mathbf{R}^m$ by introducing the componentwise product $F \bullet G : \mathbf{R}^n \rightarrow \mathbf{R}^m$. This function is defined by $(F \bullet G)_i(x) = F_i(x)G_i(x)$, or, in matrix notation, by $(F \bullet G)(x) = \text{diag}(F(x))G(x) = \text{diag}(G(x))F(x)$, where $\text{diag}(v)$ denotes the diagonal matrix whose diagonal entries are the components of the vector v . The derivative of the componentwise product, $D(F \bullet G)(x) \in \mathbf{R}^{m \times n}$, is described by our last product rule:

$$\text{Product Rule 3: } D(F \bullet G)(x) = \text{diag}(F(x))DG(x) + \text{diag}(G(x))DF(x).$$

One can verify each of the formulas above by expanding them and then applying the univariate product rule term by term. To remember the product rules, bear in mind that the end result must be a sum of two terms of the same dimensions, and that each of the terms must end with a derivative, so as to operate on a displacement vector $z \in \mathbf{R}^n$ to be placed on the right hand side.

In the one dimensional setting, the chain rule tells us that $(g \circ f)'(x) = g'(f(x))f'(x)$. In words, the formula says that we can decompose the effect of changing x on $(g \circ f)(x)$ into two pieces: the effect of changing x on the value of $f(x)$, and the effect of this change in $f(x)$ on the value of $g(f(x))$.

This same idea carries through to multivariate functions. Let $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and $G : \mathbf{R}^m \rightarrow \mathbf{R}^l$ be differentiable, and let $G \circ F : \mathbf{R}^n \rightarrow \mathbf{R}^l$ be their composition. The chain rule says that the derivative of this composition at $x \in \mathbf{R}^n$, $D(G \circ F)(x) \in \mathbf{R}^{l \times n}$, is obtained as the product of the derivative matrices $DG(F(x)) \in \mathbf{R}^{l \times m}$ and $DF(x) \in \mathbf{R}^{m \times n}$.

$$\text{The Chain Rule: } D(G \circ F)(x) = DG(F(x))DF(x).$$

This equation can be stated more explicitly as

$$\frac{\partial (G \circ F)_k}{\partial x_i}(x) = \sum_{j=1}^m \frac{\partial G_k}{\partial y_j}(F(x)) \frac{\partial F_j}{\partial x_i}(x).$$

The chain rule can be viewed as a generalization of the earlier formula on linearity of differentiation, with the linear map A replaced by the nonlinear function G .

3.A.5 Homogeneity and Euler's Theorem

Let f be a differentiable function from \mathbf{R}^n to \mathbf{R} . (We can replace the domain \mathbf{R}^n with an open (or even a closed) *convex cone*: a convex set which, if it contains $x \in \mathbf{R}^n$, also contains tx for all $t > 0$.) We say that f is *homogeneous of degree k* if

$$(3.29) \quad f(tx) = t^k f(x)$$

for all $x \in \mathbf{R}^n$ and $t > 0$.

By definition, homogeneous functions are monomials along each ray from the origin. Indeed, when $n = 1$ the homogeneous functions are precisely the monomials: if $x \in \mathbf{R}$, $g(tx) = t^k g(x)$, and $g(1) = a$, then $g(x) = ax^k$. But when $n > 1$ more complicated homogeneous functions can be found.

Nevertheless, the basic properties of homogeneous functions are generalizations of properties of monomials. If we take the derivative of each side of equation (3.29) with respect to x_i , applying the chain rule on the left hand side, we obtain

$$\nabla f(tx)'(te_i) = t^k \frac{\partial f}{\partial x_i}(x).$$

Dividing both sides of this equation by t and simplifying yields

$$\frac{\partial f}{\partial x_i}(tx) = t^{k-1} \frac{\partial f}{\partial x_i}(x).$$

In other words, the partial derivatives of a homogeneous function of degree k are themselves homogeneous of degree $k - 1$.

If we instead take the derivative of each side of (3.29) with respect to t , again using the chain rule on the left hand side, we obtain

$$\nabla f(tx)'x = \begin{cases} 0 & \text{if } k = 0, \\ kt^{k-1} f(x) & \text{otherwise.} \end{cases}$$

Setting $t = 1$ yields *Euler's Theorem*: if f is homogeneous of degree k , then

$$\nabla f(x)'x = kf(x) \text{ for all } x \in \mathbf{R}^n.$$

In fact, the converse of Euler's Theorem is also true: one can show that if f satisfies the previous identity, it is homogeneous of degree k .

3.A.6 Higher Order Derivatives

As we have seen, the derivative of a function $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a new function

$$(3.30) \quad DF : \mathbf{R}^n \rightarrow L(\mathbf{R}^n, \mathbf{R}^m).$$

For each $x \in \mathbf{R}^n$, $DF(x)$ describes how the value of F in \mathbf{R}^m changes as we move away from x in any direction $z \in \mathbf{R}^n$. Notice that in expression (3.30), the point x around which we evaluate the function F inhabits the first \mathbf{R}^n , while the displacement vector z inhabits the second \mathbf{R}^n .

The second derivative of F at x , $D^2F(x) = D(DF(x))$, describes how the value of the first derivative $DF(x) \in L(\mathbf{R}^n, \mathbf{R}^m)$ changes as we move away from x in direction $\hat{z} \in \mathbf{R}^n$. Thus, $D^2F(x)$ is an element of the set of maps $L(\mathbf{R}^n, L(\mathbf{R}^n, \mathbf{R}^m))$, which we denote by $L^2(\mathbf{R}^n, \mathbf{R}^m)$. Elements of $L^2(\mathbf{R}^n, \mathbf{R}^m)$ are called *bilinear maps* from $\mathbf{R}^n \times \mathbf{R}^n$ to \mathbf{R}^m : they take two vectors in \mathbf{R}^n as inputs, are linear in each of these vectors, and return elements of \mathbf{R}^m as outputs.

If F is twice continuously differentiable (i.e., if DF and D^2F are both continuous in x), then it can be shown that $D^2F(x)$ is *symmetric*, in the sense that $D^2F(x)(z, \hat{z}) = D^2F(x)(\hat{z}, z)$ for all $z, \hat{z} \in \mathbf{R}^n$. We therefore say that $D^2F(x)$ is an element of $L_s^2(\mathbf{R}^n, \mathbf{R}^m)$, the set of *symmetric bilinear maps* from $\mathbf{R}^n \times \mathbf{R}^n$ to \mathbf{R}^m .

More generally, the k th derivative of F is a map $D^kF : \mathbf{R}^n \rightarrow L_s^k(\mathbf{R}^n, \mathbf{R}^m)$. For each $x \in \mathbf{R}^n$, $D^kF(x)$ is a symmetric multilinear map; it takes k displacement vectors in \mathbf{R}^n as inputs, is linear in each, and returns an output in \mathbf{R}^m ; this output does not depend on the order of the inputs. If F has continuous derivatives of orders zero through K , we say that it is in class C^K .

We can use higher order derivatives to write the K th order version of Taylor's Formula, which provides a polynomial approximation of a C^K function F around the point x .

$$\text{Taylor's Formula:} \quad F(y) = F(x) + \sum_{k=1}^K \frac{1}{k!} D^kF(x)(y-x, \dots, y-x) + o(|y-x|^K).$$

Here, $D^kF(x)(y-x, \dots, y-x) \in \mathbf{R}^m$ is the output generated when the multilinear map $D^kF(x) \in L_s^k(\mathbf{R}^n, \mathbf{R}^m)$ acts on k copies of the displacement vector $(y-x) \in \mathbf{R}^n$. (To see where the factorial terms come from, try expressing the coefficients of a K th order polynomial in terms of the polynomial's derivatives.)

The higher order derivative that occurs most frequently in applications is the second derivative of a scalar valued function $f : \mathbf{R}^n \rightarrow \mathbf{R}$. This second derivative, D^2f , sends each $x \in \mathbf{R}^n$ to a symmetric bilinear map $D^2f(x) \in L_s^2(\mathbf{R}^n, \mathbf{R})$. We can represent this map using a *Hessian matrix* $\nabla^2 f(x) \in \mathbf{R}^{n \times n}$, the elements of which are the second order partial

derivatives of f :

$$\nabla^2 f(x) = \begin{pmatrix} \frac{\partial^2 f}{(\partial x_1)^2}(x) & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n}(x) \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1}(x) & \cdots & \frac{\partial^2 f}{(\partial x_n)^2}(x) \end{pmatrix}$$

When f is C^2 , the symmetry of the map $D^2 f(x)$ is reflected in the fact that the Hessian matrix is symmetric: corresponding pairs of mixed partial derivatives are equal.

The value $D^2 f(x)(z, \hat{z})$ is expressed in terms of the Hessian matrix in this way:

$$D^2 f(x)(z, \hat{z}) = z' \nabla^2 f(x) \hat{z}.$$

Using the gradient vector and Hessian matrix, we can express the second-order Taylor approximation of a C^2 scalar-valued function as follows:

$$f(y) = f(x) + \nabla f(x)'(y - x) + \frac{1}{2}(y - x)' \nabla^2 f(x)(y - x) + o\left(\|y - x\|^2\right).$$

3.A.7 The Whitney Extension Theorem

While we have defined our K times continuously differentiable functions to have domain \mathbf{R}^n , nothing we have discussed so far would change were our functions only defined on open subsets of \mathbf{R}^n . In fact, it is also possible to define C^K functions on *closed* sets $X \subset \mathbf{R}^n$. However, doing so requires some care, as the following example illustrates.

Example 3.A.2. Let $f : [0, \infty) \rightarrow \mathbf{R}$ be defined by

$$f(x) = \begin{cases} x^2 \sin \frac{1}{x} & \text{if } x > 0; \\ 0 & \text{if } x = 0. \end{cases}$$

Then f is continuous on $[0, \infty)$, and f is continuously differentiable on $(0, \infty)$ with $f'(x) = 2x \sin(\frac{1}{x}) - \cos(\frac{1}{x})$. It is possible to extend f to a differentiable function defined throughout \mathbf{R} : for instance, one can set $f(x) = 0$ for $x < 0$, in which case $f'(x) = 0$ for all $x \leq 0$. But since f' takes values -1 and 1 at points arbitrarily close to 0 , f cannot be extended to a continuously differentiable function defined throughout \mathbf{R} . §

The *Whitney Extension Theorem* gives under which a function $F : X \rightarrow \mathbf{R}^m$ can be extended to a C^K function defined throughout \mathbf{R}^n . Example 3.A.2 shows that it is not

enough for F to be continuous on X and C^K on $\text{int}(X)$: in addition, F must admit “local uniform Taylor expansions” at each x on $\text{bd}(X)$. In effect, the Whitney Extension Theorem provides a definition of (K times) continuously differentiability for functions defined on closed sets in \mathbf{R}^n .

3.A.8 Vector Integration and the Fundamental Theorem of Calculus

Let $\alpha : \mathbf{R} \rightarrow \mathbf{R}^n$ be a vector-valued function defined on the real line. Integrals of α are computed componentwise: in other words,

$$(3.31) \quad \left(\int_a^b \alpha(t) dt \right)_i = \int_a^b \alpha_i(t) dt.$$

It is easy to verify that integration, like differentiation, is linear: if $A \in \mathbf{R}^{m \times n}$ then

$$\int_a^b A \alpha(t) dt = A \int_a^b \alpha(t) dt.$$

With definition (3.31) in hand, we can state a multivariate version of the Fundamental Theorem of Calculus. Suppose that $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a C^1 function. Let $\alpha : [0, 1] \rightarrow \mathbf{R}^n$ be a C^1 function satisfying $\alpha(0) = x$ and $\alpha(1) = y$, and call its derivative $\alpha' : \mathbf{R} \rightarrow \mathbf{R}^n$. Then we have

$$\text{The Fundamental Theorem of Calculus:} \quad F(y) - F(x) = \int_0^1 DF(\alpha(t)) \alpha'(t) dt.$$

3.A.9 Potential Functions and Integrability

When can a continuous vector field $F : \mathbf{R}^n \rightarrow \mathbf{R}^n$ be expressed as the gradient of some scalar valued function f ? In other words, when does $F = \nabla f$ for some *potential function* $f : \mathbf{R}^n \rightarrow \mathbf{R}$? One can characterize the vector fields that admit potential functions in terms of their integrals over closed curves: if $F : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is continuous, it admits a potential function if and only if

$$(3.32) \quad \int_0^1 F(\alpha(t))' \left(\frac{d}{dt} \alpha(t) \right) dt = 0$$

for every piecewise C^1 function $\alpha : [0, 1] \rightarrow \mathbf{R}^n$ with $\alpha(0) = \alpha(1)$. If we use C to denote the closed curve through \mathbf{R}^n traced by α , then (3.32) can be expressed more concisely as

$$\oint_C F(x) \cdot dx = 0.$$

When F is not only continuous, but also C^1 , the question of the integrability of F can be answered by examining cross-partial derivatives. Note first that if F admits a C^2 potential function f , then the symmetry of the Hessian matrices of f implies that

$$(3.33) \quad \frac{\partial F_i}{\partial x_j}(x) = \frac{\partial^2 f}{\partial x_i \partial x_j}(x) = \frac{\partial^2 f}{\partial x_j \partial x_i}(x) = \frac{\partial F_j}{\partial x_i}(x),$$

and hence that the derivative matrix $DF(x)$ is symmetric for all $x \in \mathbf{R}^n$. The converse statement is also true, and provides the characterization of integrability we seek: if F is C^1 , with $DF(x)$ symmetric for all $x \in \mathbf{R}^n$ (i.e., whenever the *integrability condition* (3.33) holds), there is a function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ such that $\nabla f = F$. This sufficient condition for integrability remains valid whenever the domain of F is an open (or closed) convex subset of \mathbf{R}^n . However, condition (3.33) does not ensure the existence of a potential function for vector fields defined on more general domains.

3.B Affine Calculus

The simplex in \mathbf{R}^n , which serves as our state space in single population games, is an $n - 1$ dimensional set. As a consequence, derivatives of functions defined on the simplex can not be computed in the manner described in Appendix 3.A, as partial derivatives of such functions do not exist. To understand differential calculus in this context, and in the more general context of multipopulation games, we must develop the tools of calculus for functions defined on affine spaces.

3.B.1 Linear Forms and the Riesz Representation Theorem

Let Z be a subspace of \mathbf{R}^n , and let $L(Z, \mathbf{R})$ be the set of linear maps from Z to \mathbf{R} . $L(Z, \mathbf{R})$ is also known as the *dual space* of Z , and elements of $L(Z, \mathbf{R})$, namely, maps $\lambda : Z \rightarrow \mathbf{R}$ that satisfy $\lambda(az + b\hat{z}) = a\lambda(z) + b\lambda(\hat{z})$, are also known as *linear forms*.

Each vector $y \in Z$ defines a linear form $\lambda \in L(Z, \mathbf{R})$ via $\lambda(z) = y'z$. In fact, the converse statement is also true: every linear form can be uniquely represented in this way.

Theorem 3.B.1 (The Riesz Representation Theorem). *For each linear form $\lambda \in L(Z, \mathbf{R})$, there is a unique $y \in Z$, the Riesz representation of λ , such that $\lambda(z) = y'z$ for all $z \in Z$.*

Another way of describing the Riesz representation theorem is to say that Z and $L(Z, \mathbf{R})$ are *linearly isomorphic*: the map from Z to $L(Z, \mathbf{R})$ described above is linear, one-to-one, and onto.

It is crucial to note that when Z is a proper subspace of \mathbf{R}^n , the linear form λ can be represented by many vectors in \mathbf{R}^n . What Theorem 3.B.1 tells us is that λ can be represented by a unique vector *in Z itself*.

Example 3.B.2. Let $Z = \mathbf{R}_0^2 = \{z \in \mathbf{R}^2 : z_1 + z_2 = 0\}$, and define the linear form $\lambda \in L(Z, \mathbf{R})$ by $\lambda(z) = z_1 - z_2$. Then not only

$$y = \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \text{ but also } \hat{y} = \begin{pmatrix} 3 \\ 1 \end{pmatrix}$$

represents λ : if $z \in Z$, then $\hat{y}'z = 3z_1 + z_2 = 3z_1 + (-z_1) = 2z_1 = z_1 - z_2 = y'z = \lambda(z)$. But since y is an element of Z , it is the Riesz representation of λ . §

In this example, the reason that both y and \hat{y} can represent λ is that their difference,

$$\hat{y} - y = \begin{pmatrix} 2 \\ 2 \end{pmatrix}$$

is orthogonal to Z . This suggests a simple way of recovering the Riesz representation of a linear form from an arbitrary vector representation: eliminate the portion orthogonal to Z by applying the orthogonal projection P_Z .

Theorem 3.B.3. *Let $\lambda \in L(Z, \mathbf{R})$ be a linear form. If $\hat{y} \in \mathbf{R}^n$ represents λ , in the sense that $\lambda(z) = \hat{y}'z$ for all $z \in Z$, then $y = P_Z \hat{y}$ is the Riesz representation of λ .*

Example 3.B.4. Recall that the orthogonal projection onto \mathbf{R}_0^2 is $\Phi = I - \frac{1}{2}\mathbf{1}\mathbf{1}'$. Thus, in the previous example, we can recover y from \hat{y} in the following way:

$$y = \Phi \hat{y} = (I - \frac{1}{2}\mathbf{1}\mathbf{1}') \begin{pmatrix} 3 \\ 1 \end{pmatrix} = \begin{pmatrix} 3 \\ 1 \end{pmatrix} - \begin{pmatrix} 2 \\ 2 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}. \text{ §}$$

3.B.2 Dual Characterizations of Multiples of Linear Forms

Before turning our attention to calculus, we present some results that characterize when two linear forms are scalar multiples of one another. We will use these results

when studying imitative dynamics in Chapters 5 and 8; see especially Exercise 5.4.18 and Theorem 8.5.9.

If the vectors $v \in \mathbf{R}^n$ and $w \in \mathbf{R}^n$ are non-zero multiples of one another, then v and w clearly are orthogonal to the same set of vectors in \mathbf{R}^n . Conversely, if $\{v\}^\perp = \{y \in \mathbf{R}^n : v'y = 0\}$ equals $\{w\}^\perp$, then v and w must be (non-zero) multiples of one another, as they are both normal vectors of the same hyperplane.

When are v and w positive multiples of one another? This is the case if and only if the set $\mathcal{H}(v) = \{y \in \mathbf{R}^n : v'y \geq 0\}$, the closed half-space consisting of those vectors with which v forms an acute or right angle, is equal to the corresponding set $\mathcal{H}(w)$. Clearly, $\mathcal{H}(v) = \mathcal{H}(w)$ implies that $\{v\}^\perp = \{w\}^\perp$, and so that $v = cw$; since $v \in \mathcal{H}(v) = \mathcal{H}(w)$, it must be that $c > 0$.

In summary, we have

Observation 3.B.5.

- (i) $\{x \in \mathbf{R}^n : v'x = 0\} = \{x \in \mathbf{R}^n : w'x = 0\}$ if and only if $v = cw$ for some $c \neq 0$.
- (ii) $\{x \in \mathbf{R}^n : v'x \geq 0\} = \{x \in \mathbf{R}^n : w'x \geq 0\}$ if and only if $v = cw$ for some $c > 0$.

Proposition 3.B.6 provides analogues of the characterizations above for settings in which one can only compare how v and w act on vectors in some subspace $Z \subseteq \mathbf{R}^n$. Since these comparisons relate v and w as linear forms on Z , Theorem 3.B.3 suggests that the characterizations should be expressed in terms of the orthogonal projections of v and w onto Z .

Proposition 3.B.6.

- (i) $\{z \in Z : v'z = 0\} = \{z \in Z : w'z = 0\}$ if and only if $P_Z v = c P_Z w$ for some $c \neq 0$.
- (ii) $\{z \in Z : v'z \geq 0\} = \{z \in Z : w'z \geq 0\}$ if and only if $P_Z v = c P_Z w$ for some $c > 0$.

Proof. The “if” direction of part (i) is immediate. For the “only if” direction, observe that $v'z = 0$ for all $z \in Z$ if and only if $v'P_Z x = 0$ for all $x \in \mathbf{R}^n$. Since the matrix P_Z is symmetric, we can rewrite the equality above as $(P_Z v)'x = 0$; thus, the conclusion that $P_Z v = c P_Z w$ with $c \neq 0$ follows from Observation 3.B.5(i). The proof of part (ii) follows similarly from Observation 3.B.5(ii). ■

To cap this discussion, we note that both parts of Observation 3.B.5 are the simplest cases of more general duality results that link a linear map $A \in L(\mathbf{R}^m, \mathbf{R}^n) \equiv \mathbf{R}^{n \times m}$ with its transpose $A' \in L(\mathbf{R}^n, \mathbf{R}^m) \equiv \mathbf{R}^{m \times n}$. Part (i) is essentially the $m = 1$ case of the *Fundamental Theorem of Linear Algebra*:

$$(3.34) \quad \text{range}(A) = (\text{nullspace}(A'))^\perp.$$

In equation (3.34), the set $\text{range}(A) = \{w \in \mathbf{R}^n : w = Ax \text{ for some } x \in \mathbf{R}^m\}$ is the span of the columns of A . The set $\text{nullspace}(A') = \{y \in \mathbf{R}^n : A'y = \mathbf{0}\}$ consists of the vectors that A' maps to the origin; equivalently, it is the set of vectors that are orthogonal to every column of A . Viewed in this light, equation (3.34) says that w is a linear combination of the columns of A if and only if any y that is orthogonal to each column of A is also orthogonal to w . While (3.34) is of basic importance, it is quite easy to derive after taking orthogonal complements:

$$(\text{range}(A))^\perp = \{y \in \mathbf{R}^n : y'Ax = 0 \text{ for all } x \in \mathbf{R}^m\} = \{y \in \mathbf{R}^n : y'A = \mathbf{0}'\} = \text{nullspace}(A').$$

Part (ii) of Observation 3.B.5 is essentially the $m = 1$ case of *Farkas's Lemma*:

$$(3.35) \quad [w = Ax \text{ for some } x \in \mathbf{R}_+^m] \text{ if and only if } [[A'y \geq \mathbf{0} \Rightarrow w'y \geq 0] \text{ for all } y \in \mathbf{R}^n].$$

In words: w is a *nonnegative* linear combination of the columns of A if and only if any y that forms a *weakly acute* angle with each column of A also forms a *weakly acute* angle with w . Despite their analogous interpretations, statement (3.35) is considerably more difficult to prove than statement (3.34)—see the Notes.

3.B.3 Derivatives of Functions on Affine Spaces

Before considering calculus on affine spaces, let us briefly review differentiation of scalar-valued functions on \mathbf{R}^n . If f is a C^1 function from \mathbf{R}^n to \mathbf{R} , then its derivative at x , denoted $Df(x)$, is an element of $L(\mathbf{R}^n, \mathbf{R})$, the set of linear maps from \mathbf{R}^n to \mathbf{R} . For each $x \in \mathbf{R}^n$, the map $Df(x)$ takes vectors $z \in \mathbf{R}^n$ as inputs and returns scalars $Df(x)z \in \mathbf{R}$ as outputs. The latter expression appears in the first order Taylor expansion

$$f(x + z) = f(x) + Df(x)z + o(z) \text{ for all } z \in \mathbf{R}^n.$$

By the Riesz Representation Theorem, there is a unique vector $\nabla f(x) \in \mathbf{R}^n$ satisfying $Df(x)z = \nabla f(x)'z$ for all $z \in \mathbf{R}^n$. We call $\nabla f(x)$ the gradient of f at x . In the present full-dimensional case, $\nabla f(x)$ is the vector of partial derivatives $\frac{\partial f}{\partial x_i}(x)$ of f at x .

Now, let $A \subseteq \mathbf{R}^n$ be an affine space with tangent space TA , and consider a function $f : A \rightarrow \mathbf{R}$. (As in Appendix 3.A, the ideas to follow can also be applied to functions whose domain is a set that is open (or closed) relative to A .) We say that f is *differentiable* at $x \in A$ if there is a linear map $Df(x) \in L(TA, \mathbf{R})$ satisfying

$$f(x + z) = f(x) + Df(x)z + o(z) \text{ for all } z \in TA.$$

The *gradient* of f at x is the Riesz representation of $Df(x)$. In other words, it is the unique vector $\nabla f(x) \in TA$ such that $Df(x)z = \nabla f(x)'z$ for all $z \in TA$. If the function $\nabla f : A \rightarrow TA$ is continuous, then f is *continuously differentiable*, or of class C^1 .

When $A = \mathbf{R}^n$, this definition of the gradient is simply the one presented earlier, and $\nabla f(x)$ is the only vector in \mathbf{R}^n that represents $Df(x)$. But in lower dimensional cases, there are many vectors in \mathbf{R}^n that can represent $Df(x)$. The gradient vector $\nabla f(x)$ is the only one lying in TA ; all others are obtained by summing $\nabla f(x)$ and an element of $(TA)^\perp$.

When $A = \mathbf{R}^n$, the gradient of f at x is just the vector of partial derivatives of f at x . But in other cases, the partial derivatives of f may not even exist. How does one compute $\nabla f(x)$ then? Usually, it is easiest to extend the function f to all of \mathbf{R}^n in some smooth way, and then to compute the gradient by way of this extension. In some cases (e.g., when f is a polynomial), obtaining the extension is just a matter of declaring that the domain is \mathbf{R}^n . But even in this situation, there is an alternative extension that is often handy.

Proposition 3.B.7. *Let $f : A \rightarrow \mathbf{R}$ be a C^1 function on the affine set A , and let $Z = TA$.*

- (i) *Let $\tilde{f} : \mathbf{R}^n \rightarrow \mathbf{R}$ be any C^1 extension of f . Then $\nabla f(x) = P_Z \nabla \tilde{f}(x)$ for all $x \in A$.*
- (ii) *Define $\bar{f} : \mathbf{R}^n \rightarrow \mathbf{R}$ by*

$$\bar{f}(y) = f(P_Z y + z_A^\perp),$$

where z_A^\perp is the unique element of $A \cap Z^\perp$. Then $\nabla \bar{f}(x) = \nabla f(x)$ for all $x \in A$.

In words, \bar{f} assigns the value $f(x)$ to each point in \mathbf{R}^n whose orthogonal projection onto $TA = Z$ is the same as that of $x \in A$; the gradient of \bar{f} is identical to the gradient of f on the set A .

Proof. Part (i) follows immediately from the relevant definitions. To prove part (ii), suppose that $x \in A$. Then by the chain and product rules,

$$D\bar{f}(x) = D(f(P_Z x + z_A^\perp)) = Df(x)P_Z.$$

This linear form on \mathbf{R}^n is represented by the (column) vector $\nabla \bar{f}(x) = (\nabla f(x)'P_Z)' \in \mathbf{R}^n$. But since the orthogonal projection matrix P_Z is symmetric, and since $\nabla f(x) \in Z$, we conclude that

$$\nabla \bar{f}(x) = (\nabla f(x)'P_Z)' = P_Z' \nabla f(x) = P_Z \nabla f(x) = \nabla f(x). \blacksquare$$

The fact that P_Z is an *orthogonal* projection makes this proof simple: since P_Z is symmetric, we are able to transfer its action from the displacement direction $z \in Z$ to the vector $\nabla f(x)$ itself.

Similar considerations arise for vector-valued functions defined on affine spaces, and also for higher order derivatives. If $F : A \rightarrow \mathbf{R}^m$ is C^1 , its derivative at $x \in A$ is a linear map $DF(x) \in L(Z, \mathbf{R}^m)$, where we once again write Z for TA . While there are many matrices in $\mathbf{R}^{m \times n}$ that represent this derivative, applying the logic above to each component of F shows that there is a unique such matrix, called the *Jacobian matrix* or *derivative matrix*, whose rows are elements of Z . As before, we abuse notation by denoting this matrix $DF(x)$. But unlike before, this abuse can create some confusion: if F is “automatically” defined on all of \mathbf{R}^n , one must be careful to distinguish between the derivative matrix of $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ at x and the derivative matrix of its restriction $F|_A : A \rightarrow \mathbf{R}^m$ at x ; they are related by $DF|_A(x) = DF(x)P_Z$.

If the function $f : A \rightarrow \mathbf{R}$ is C^2 , then its second derivative at $x \in A$ is a symmetric bilinear map $D^2f(x) \in L_s^2(Z, \mathbf{R})$. There are many symmetric matrices in $\mathbf{R}^{n \times n}$ that represent $D^2f(x)$, but there is a unique such matrix whose rows and columns are in Z . We call this matrix the *Hessian* of f at x , and denote it $\nabla^2 f(x)$. If $\tilde{f} : \mathbf{R}^n \rightarrow \mathbf{R}$ is any C^2 extension of f , then we can compute the Hessian of f as $\nabla^2 f(x) = P_Z \nabla^2 \tilde{f}(x) P_Z$; if $\bar{f}(y) = f(P_Z y + z_A^\perp)$ is the constant orthogonal extension of f to \mathbf{R}^n , then $\nabla^2 f(x) = \nabla^2 \bar{f}(x)$.

3.B.4 Affine Integrability

A necessary and sufficient condition for a C^1 vector field $F : \mathbf{R}^n \rightarrow \mathbf{R}^n$ to admit a potential function—that is, a scalar valued function f satisfying $\nabla f(x) = F(x)$ for all $x \in \mathbf{R}^n$ —is that its derivative matrix $DF(x)$ be symmetric for all $x \in \mathbf{R}^n$. We now state a definition of potential functions for cases in which the map F is only defined on an affine space, and show that an appropriate symmetry condition on $DF(x)$ is necessary and sufficient for a potential function to exist. We also relate these notions to their full-dimensional analogues.

Let $A \subseteq \mathbf{R}^n$ be an affine space with tangent space $Z = TA$, and let z_A^\perp be the unique element of $A \cap Z^\perp$. Suppose that the map $F : A \rightarrow \mathbf{R}^n$ is continuous. We call the function $f : A \rightarrow \mathbf{R}$ a *potential function* for F if

$$(3.36) \quad \nabla f(x) = P_Z F(x) \text{ for all } x \in A.$$

What does this definition require? Since $\nabla f(x) \in Z$, the action of $\nabla f(x)$ on Z^\perp is null (that is, $(z^\perp)' \nabla f(x) = 0$ whenever $z^\perp \in Z^\perp$). But since $F(x) \in \mathbf{R}^n$, the action of $F(x)$ on Z^\perp is not restricted in this way. Condition (3.36) requires that $F(x)$ have the same action as $\nabla f(x)$ on Z , but places no restriction on how $F(x)$ acts on the complementary set Z^\perp .

Theorem 3.B.8 characterizes the smooth maps on A that admit potential functions. The

characterization is stated in terms of a symmetry condition on the derivatives $DF(x)$.

Theorem 3.B.8. *The C^1 map $F : A \rightarrow \mathbf{R}^n$ admits a potential function if and only if $DF(x)$ is symmetric with respect to $Z \times Z$ for all $x \in A$ (i.e., if and only if $z'DF(x)\hat{z} = \hat{z}'DF(x)z$ for all $z, \hat{z} \in Z$ and $x \in A$).*

Proof. To prove the “only if” direction, suppose that F admits a potential function f satisfying condition (3.36). This means that for all $x \in A$, $F(x)$ and $\nabla f(x)$ define identical linear forms in $L(Z, \mathbf{R})$. By taking the derivative of each side of this identity, we find that $DF(x) = \nabla^2 f(x)$ as bilinear forms in $L^2(Z, \mathbf{R})$. But since $\nabla^2 f(x)$ is a symmetric bilinear form on $Z \times Z$ (by virtue of being a second derivative), $DF(x)$ is as well.

The “if” direction is a consequence of the following proposition.

Proposition 3.B.9. *Define the map $\bar{F} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ by*

$$\bar{F}(y) = P_Z F(P_Z y + z_A^\perp).$$

Then \bar{F} admits a potential function $\bar{f} : \mathbf{R}^n \rightarrow \mathbf{R}$ if and only if $DF(x)$ is symmetric with respect to $Z \times Z$ for all $x \in A$. In this case, $f = \bar{f}|_A$ is a potential function for F .

Proof. Define the function $\xi : \mathbf{R}^n \rightarrow A$ by $\xi(y) = P_Z y + z_A^\perp$. Then

$$(3.37) \quad D\bar{F}(y) = D(P_Z F(\xi(y))) = P_Z (DF(\xi(y))) P_Z.$$

Now \bar{F} admits a potential function if and only if $D\bar{F}(y)$ is symmetric for all $y \in \mathbf{R}^n$. Equation (3.37) tells us that the latter statement is true if and only if $DF(x)$ is symmetric with respect to $Z \times Z$ for all $x \in A$, proving the first statement in the proposition.

To prove the second statement, suppose that \bar{f} is a potential function for \bar{F} , and let $f = \bar{f}|_A$. Then since $\xi(x) = x$ for all $x \in A$, we find that

$$\nabla f(x) = P_Z \nabla \bar{f}(x) = P_Z \bar{F}(x) = P_Z (P_Z F(\xi(x))) = P_Z F(x). \blacksquare$$

This completes the proof of Theorem 3.B.8. \blacksquare

If the C^1 map $F : A \rightarrow \mathbf{R}^n$ is integrable (i.e., if it admits a potential function $f : A \rightarrow \mathbf{R}$), can we extend F to all of \mathbf{R}^n in such a way that the extension is integrable too? One natural way to proceed is to extend the potential function f to all of \mathbf{R}^n . If one does so in an arbitrary way, then the *projected* maps $P_Z F$ and $P_Z \tilde{F}$ will agree regardless of how the extended potential function \tilde{f} is chosen (cf Observation 3.2.3 and the subsequent

discussion). But is it always possible to choose \tilde{f} in such a way that that F and \tilde{F} are identical on A , so that the function \tilde{F} is a genuine extension of the function F ? Theorem 3.B.10 shows one way that this can be done.

Theorem 3.B.10. *Suppose $F : A \rightarrow \mathbf{R}^n$ is continuous with potential function $f : A \rightarrow \mathbf{R}$. Define $\tilde{f} : \mathbf{R}^n \rightarrow \mathbf{R}$ by*

$$\tilde{f}(y) = f(\xi(y)) + (y - \xi(y))'F(\xi(y)), \text{ where } \xi(y) = P_Z y + z_A^\perp,$$

and define $\tilde{F} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ by $\tilde{F}(y) = \nabla \tilde{f}(y)$. Then $\tilde{F}|_A = F$. Thus, any integrable map/potential function pair defined on A can be extended to a vector field/potential function pair defined on all of \mathbf{R}^n .

Proof. We can compute \tilde{F} from \tilde{f} using the chain and product rules:

$$\begin{aligned} \tilde{F}(y)' &= \nabla \tilde{f}(y)' \\ &= \nabla f(\xi(y))'P_Z + (y - \xi(y))'DF(\xi(y))P_Z + F(\xi(y))'(I - P_Z) \\ &= \left(P_Z F(\xi(y)) \right)' P_Z + (y - \xi(y))' DF(\xi(y)) P_Z + F(\xi(y))' - F(\xi(y))' P_Z \\ &= F(\xi(y))' P_Z P_Z + (y - \xi(y))' DF(\xi(y)) P_Z + F(\xi(y))' - F(\xi(y))' P_Z \\ &= F(\xi(y))' + (y - \xi(y))' DF(\xi(y)) P_Z \end{aligned}$$

If $x \in A$, then $\xi(x) = x$, allowing us to conclude that $\tilde{F}(x) = F(x)$. ■

If F takes values in Z , so that $F(x) = P_Z F(x)$ for all $x \in A$, then $\tilde{f}(y)$ is simply $f(\xi(y))$, and so $\tilde{F}(y) = P_Z \nabla f(\xi(y)) = P_Z F(\xi(y))$; in this case, the construction in Theorem 3.B.10 is identical to the one introduced in Proposition 3.B.9. The novelty in Theorem 3.B.10 is that it lets us extend the domain of F to all of \mathbf{R}^n in an integrable fashion even when F takes values throughout \mathbf{R}^n .

3.N Notes

Section 3.1. Sections 3.1.1 through 3.1.6 follow Sandholm (2001a), while Section 3.1.7 follows Roughgarden and Tardos (2002, 2004) and Correa et al. (2004, 2008).

Matching in two-player games with common interests defines a fundamental model from population genetics; the common interest assumption reflects the shared fate of two genes that inhabit the same organism. See Hofbauer and Sigmund (1988, 1998) for further discussion.

Congestion games first appear in the seminal book of Beckmann et al. (1956), who define a general model of traffic flow with inelastic demand, and use a potential function argument to establish the existence and uniqueness of Nash equilibrium. Boyce et al. (2005) describe this book's influence on later work in transportation science. In this literature, Nash equilibrium of congestion games are sometimes referred to as *Wardrop equilibria*, after Wardrop (1952). The textbook of Sheffi (1985) treats congestion games from a transportation science perspective at an undergraduate level; the more recent monograph of Patriksson (1994) provides a comprehensive treatment of the topic from this point of view.

Important examples of finite player potential games are introduced by Rosenthal (1973) and Slade (1994), and characterizations of this class of normal form games are provided by Monderer and Shapley (1996b), Ui (2000), and Sandholm (2008a). Example 3.1.6 and Exercise 3.1.12 are due to Sandholm (2005b). Braess's paradox (Example 3.1.10) was first reported in Braess (1968). Exercise 3.1.11 is well known in the transportation science literature; it also corrects a mistake (!) in Corollary 5.6 of Sandholm (2001a). Versions of the efficiency results in Section 3.1.6 are established by Dafermos and Sparrow (1969) for a model of traffic congestion model and by Hofbauer and Sigmund (1988) for single population games. For further discussion of constraint qualification and of the interpretation of the Kuhn-Tucker first order conditions, see Avriel (1976, Section 3.1) and Harker and Pang (1990).

Inefficiency bounds were introduced in the computer science literature by Koutsoupias and Papadimitriou (1999) and by Papadimitriou (2001), who introduced the term "price of anarchy" to refer to these bounds. Most of the results in presented in Section 3.1.7 are due to Roughgarden and Tardos (2002, 2004) and are presented in the book of Roughgarden (2005). The simple proof of Theorem 3.1.28 presented here is due to Correa et al. (2004, 2008), as is Corollary 3.1.29. These references present a variety of additional inefficiency bounds for congestion games with general classes of cost functions, as well as inefficiency bounds for general population games.

Section 3.2. This section follows Sandholm (2009c).

The general definition and basic properties of normal form potential games are established by Monderer and Shapley (1996b). The triangular integrability condition from Exercise 3.2.7 is due to Hofbauer (1985). The fact that constant games are potential games in which potential equals aggregate payoffs is important in models of evolutionary implementation; see Sandholm (2002, 2005b, 2007b).

Section 3.3. This section follows Hofbauer and Sandholm (2009).

Evolutionarily stable strategies and neutrally stable strategies are introduced in the

single population random matching context by Maynard Smith and Price (1973) and Maynard Smith (1982), respectively. The connection between interior ESS and negative definiteness of the payoff matrix was first noted by Haigh (1975). See Hines (1987) for a survey of early work on these and related concepts. A version of the GESS concept is used by Hamilton (1967) in his pioneering analysis of sex-ratio selection under the name “unbeatable strategy”. See Hamilton (1996, p. 373–374) for an intriguing discussion of the links between the notions of unbeatable strategy and ESS, and see Kojima (2006b) for a recent treatment. Further discussion of ESS can be found in the Notes to Chapter 8.

For more on Rock-Paper-Scissors, see Gaunersdorfer and Hofbauer (1995). The War of Attrition is introduced in Bishop and Cannings (1978); for economic applications, see Bulow and Klemperer (1999) and the references therein. Imhof (2005) derives a closed-form expression for the Nash equilibrium of the war of attrition in terms of Chebyshev polynomials of the second kind. The dominant diagonal condition used in Example 3.3.12 is a consequence of the Geršgorin Disk Theorem; see Horn and Johnson (1985). This reference also presents the trace condition used in proving Proposition 3.3.10.

In the convex analysis literature, functions that satisfy our definition of stability (though typically with the inequality reversed) are called “monotone”—see Rockafellar (1970) or Hiriart-Urruty and Lemaréchal (2001). For more on pseudomonotonicity and pseudoconvexity, see Avriel (1976, Chapter 6) and Crouzeix (1998). The elementary proof of existence of Nash equilibrium in stable games presented in Section 3.3.5 is a translation to the present context of work on monotone operators on vector spaces due to Minty (1967). Good references on the Minmax Theorem and its connection with the Separating Hyperplane Theorem are Kuhn (2003) and Luce and Raiffa (1957). For the finite intersection property, see Folland (1999).

Section 3.4. The definition of supermodular population games here comes from Hofbauer and Sandholm (2007). Finite player analogues of the results presented here are established by Topkis (1979), Vives (1990), and Milgrom and Roberts (1990). Accounts of these results can be found in Fudenberg and Tirole (1991, Sec. 12.3) and Vives (2005); Topkis (1998) and Vives (2000) are book-length studies. For macroeconomic applications, see Cooper (1999).

Appendix 3.A. For a textbook treatment of multivariate calculus that emphasizes the notion of the derivative as a linear map, see Lang (1997, Chapter 17). For the Whitney Extension Theorem, see Abraham and Robbin (1967) or Krantz and Parks (1999).

Appendix 3.B. The version of the Riesz Representation Theorem presented here, along with further discussion of calculus on affine spaces, can be found in Akin (1990). For further discussion of the dual characterizations described at the end of Section 3.B.2, see

Lax (2007, Chapter 13) or Hiriart-Urruty and Lemaréchal (2001, Section A.4.3).

Part II

Deterministic Evolutionary Dynamics

Revision Protocols and Deterministic Evolutionary Dynamics

4.0 Introduction

The theory of population games developed in the previous chapters provides a simple framework for describing strategic interactions among large numbers of agents. Having explored these games' basic properties, we now turn to modeling the behavior of the agents who play them.

Traditionally, predictions of behavior in games are based on some notion of equilibrium, typically Nash equilibrium or some refinement thereof. These notions are founded on the assumption of equilibrium knowledge, which posits that each player correctly anticipates how his opponents will act. The equilibrium knowledge assumption is difficult to justify, and in contexts with large numbers of agents it is particularly strong.

As an alternative to the equilibrium approach, we introduce an explicitly dynamic model in which each agent occasionally reconsiders his choice of strategy, using myopic rules to adjust his action in response to the current strategic environment. This dynamic model does not assume the automatic coordination of agents' beliefs, and it accommodates a variety of specifications of agents' myopic updating procedures.

We define these updating procedures formally in Section 4.1, where we introduce the notion of a *revision protocol*. A revision protocol ρ takes current payoffs and aggregate behavior as inputs; its outputs are *conditional switch rates* $\rho_{ij}^p(\pi^p, x^p)$, which describe how frequently agents playing strategy $i \in S^p$ who are considering switching strategies switch to strategy $j \in S^p$, given that the current payoff vector and population state are π^p and x^p . Revision protocols are flexible enough to accommodate a wide range of choice paradigms, including ones based on imitation and optimization, and other possibilities.

A population game F describes a strategic environment; a revision protocol ρ describes the procedures agents follow in adapting their behavior to that environment. Together, F and ρ define a stochastic evolutionary process, a process whose stochastic elements are independent conditional on the current population state. Since the number of agents are large, intuition from the law of large numbers suggests that the idiosyncratic noise in this process will be averaged away, leaving aggregate behavior to evolve according to an essentially deterministic process. We derive this process, which we call the *mean dynamic* generated by F and ρ , in Section 4.2. In Section 4.3, we present examples illustrating how well-known dynamics from the literature can be provided with microfoundations by means of this approach.

In the story above, we began with a game and a revision protocol and derived a differential equation on the state space X . But if our goal is to investigate the consequences of a particular choice procedure, it is preferable to fix this revision protocol and let the game F vary. By doing so, we generate a map from population games to differential equations that we call a *deterministic evolutionary dynamic*. We develop this notion in detail in Section 4.4.

Our derivation of deterministic evolutionary dynamics in this chapter is informal, based solely on an appeal to the idea that idiosyncratic noise should be averaged away when populations are large. We will formalize this logic in Chapter 10, where we introduce a Markov process to describe stochastic evolution in a large but finite population. We prove there that over finite time spans, this Markov process converges to a deterministic limit—namely, a solution trajectory of the mean dynamic—as the population size becomes arbitrarily large. While this deterministic approximation result does not hold over infinite time spans, we show in Chapter 12 that the possible infinite-horizon behaviors of the Markov process can still be ascertained through an analysis of the mean dynamic.

Until then, we spend Chapters 4–9 working directly with the deterministic limit. As background, we introduce the rudiments of the theory of ordinary differential equations in Appendix 4.A, leaving further developments for the appendices of the chapters to come.

4.1 The Basic Stochastic Evolutionary Model

4.1.1 Inertia and Myopia

There many possible definitions of behavior dynamics for games played by large populations of agents. The evolutionary approach to population game dynamics is based on two assumptions—inertia and myopia—that dictate the form that the dynamics take.

By *inertia*, we mean that individual agents do not continually reevaluate their choices in the game, but instead reconsider their strategies only sporadically. It is natural to expect inertia when the strategic environment being modeled is just one of many with which the agents simultaneously contend. Such a multiplicity of concerns, rather than a single-minded focus on one problem over days, weeks, months, or years, seems more usual than exceptional.

By *myopia*, we mean that revising agents condition their choices on current behavior and payoff opportunities; they do not attempt to incorporate beliefs about the future course of play into their decisions. Large population interactions make myopia relatively appealing. When populations are large enough that individual members become anonymous, contingent punishments, reputations, and other possibilities that are central to the theory of repeated game can safely be ignored. Moreover, myopia and inertia are mutually reinforcing, as myopic behavior is most sensible when opponents' behavior adjusts slowly, so that strategies that are appealing now are likely to continue to be so.

4.1.2 Revision Protocols

We now introduce a simple, general model of choice in population games. It is based on the notion of a *revision protocol*, which describes when and how individual agents update their strategies as time passes. Inertia and myopia are built into the fabric of this model: the model has agents wait a random amount of time before they consider switching strategies, and their decisions at such moments only condition on current payoffs and the current social state.

Let $F : X \rightarrow \mathbf{R}^n$ be a population game with pure strategy sets (S^1, \dots, S^p) and integer-valued population masses (m^1, \dots, m^p) . We suppose for now that each population is large but finite: population $p \in \mathcal{P}$ has Nm^p members, where N is a positive integer. The set of feasible social states is therefore $X^N = X \cap \frac{1}{N}\mathbf{Z}^n = \{x \in X : Nx \in \mathbf{Z}^n\}$, a discrete grid embedded in the original state space X . We refer to the parameter N somewhat loosely as the *population size*.

The procedures agents follow in deciding when to switch strategies and which strategies to switch to are called *revision protocols*.

Definition. A revision protocol ρ^p is a map $\rho^p : \mathbf{R}^{np} \times X^p \rightarrow \mathbf{R}_+^{np \times np}$. The scalar $\rho_{ij}^p(\pi^p, x^p)$ is called the conditional switch rate from strategy $i \in S^p$ to strategy $j \in S^p$ given payoff vector π^p and population state x^p .

For convenience, we will write $\rho_{ij}^p(\pi^p)$ in place of $\rho_{ij}^p(\pi^p, x^p)$ when conditional switch rates do not depend directly on x^p , and we will refer to the collection $\rho = (\rho^1, \dots, \rho^p)$ as a revision

protocol when no confusion will arise. Analogous simplifications are used whenever it is helpful to do so.

A population game F , a population size N , and a revision protocol ρ define a continuous time evolutionary process on X^N . A one-size-fits-all interpretation of this process is as follows. Each agent in the society is equipped with a “stochastic alarm clock”. The times between rings of an agent’s clock are independent, each with a rate R exponential distribution. (This modeling device is often called a “Poisson alarm clock”; see Appendix 10.A.1 and Example 10.C.1 for an explanation.) We assume that the rate R satisfies

$$R \geq \max_{x,i,p} \sum_{j \in S^p - \{i\}} \rho_{ij}^p(F^p(x), x^p),$$

and that the ring times of different agents’ clocks are independent of one another.

The ringing of a clock signals the arrival of a revision opportunity for the clock’s owner. If an agent playing strategy $i \in S^p$ receives a revision opportunity, he switches to strategy $j \neq i$ with probability ρ_{ij}^p/R , and he continues to play strategy i with probability $1 - \sum_{j \neq i} \rho_{ij}^p/R$; this decision is made independently of the timing of the clocks’ rings. If a switch occurs, the population state changes accordingly, from the old state x to a new state y that accounts for the agent’s choice. As the evolutionary process proceeds, the alarm clocks and the revising agents are only influenced by the prior history of the process by way of the current values of payoffs and the social state.

This interpretation of the evolutionary process can be applied to any revision protocol. Still, simpler interpretations are often available for protocols with additional structure. To motivate one such case, observe that in the interpretation provided above, the diagonal components ρ_{ii}^p of the revision protocol play no role whatsoever. But if the protocol is *exact*—that is, if there is a constant $R > 0$ such that

$$(4.1) \quad \sum_{j \in S^p} \rho_{ij}^p(\pi^p, x^p) = R \text{ for all } \pi^p \in \mathbf{R}^{m^p}, x^p \in X^p, i \in S^p, \text{ and } p \in \mathcal{P},$$

then the values of these diagonal components become meaningful: in this case, $\rho_{ii}^p/R = 1 - \sum_{j \neq i} \rho_{ij}^p/R$ is the probability that a strategy i player who receives a revision opportunity does not switch strategies.

Exact protocols are particularly easy to interpret when $R = 1$: in this case, agents’ clocks ring at rate 1, and for every strategy $j \in S^p$, ρ_{ij}^p itself represents the probability that an i player whose clock rings proceeds by playing strategy j . We will henceforth assume that protocols described as exact have clock rate $R = 1$ unless a different clock rate is specified explicitly. This focus on unit clock rates is not very restrictive: the only effect

of replacing a protocol ρ with its scalar multiple $\frac{1}{R}\rho$ is to change the speed at which the evolutionary process runs by a constant factor.

Other examples of protocols that allow alternative interpretations of the evolutionary process can be found in Section 4.3.

4.2 Mean Dynamics

4.2.1 Derivation

The model above defines a stochastic process $\{X_t^N\}$ on the state space \mathcal{X}^N . We now derive a deterministic process that describes the expected motion of $\{X_t^N\}$. In Chapter 10, we will prove that this deterministic process provides a very good approximation of the behavior of the stochastic process $\{X_t^N\}$ so long as the time horizon of interest is finite and the population size is sufficiently large. But having noted this result, we will focus in the intervening chapters on the deterministic process itself.

The times between rings of each agent's stochastic alarm clock are independent and follow a rate R exponential distribution. How many times will this agent's clock ring during the next t time units? A basic result from probability theory shows that the number of rings during time interval $[0, t]$ follows a Poisson distribution with mean Rt . This fact is all we need to perform the analysis below; a detailed account of the exponential and Poisson distributions can be found in Appendix 10.A.

Let us now compute the expected motion of the stochastic process $\{X_t^N\}$ over the next dt time units, where dt is small. To rein in our notation we focus on the single population case.

Each agent in the population receives revision opportunities according to an exponential distribution with rate R , and so each expects to receive $R dt$ opportunities during the next dt time units. Thus, if the current state is x , the expected number of revision opportunities received by agents currently playing strategy i is approximately

$$Nx_i R dt.$$

We say "approximately" because the value of x_i may change during time interval $[0, dt]$, but this change is very likely to be small if dt is small.

Since an i player who receives a revision opportunity switches to strategy j with probability ρ_{ij}/R , the expected number of such switches during the next dt time units is

approximately

$$Nx_i \rho_{ij} dt.$$

It follows that the expected change in the use in strategy i during the next dt time units is approximately

$$(4.2) \quad N \left(\sum_{j \in S} x_j \rho_{ji} - x_i \sum_{j \in S} \rho_{ij} \right) dt.$$

The first term in expression (4.2) captures switches to strategy i from other strategies, while the second captures switches to other strategies from strategy i . Dividing expression (4.2) by N yields the expected change in the *proportion* of agents choosing strategy i : that is, in component x_i of the social state. We obtain a differential equation for the social state by eliminating the time differential dt :

$$\dot{x}_i = \sum_{j \in S} x_j \rho_{ji} - x_i \sum_{j \in S} \rho_{ij}.$$

This ordinary differential equation is the *mean dynamic* corresponding to revision protocol ρ .

We now describe the mean dynamic for the general multipopulation case.

Definition. Let F be a population game, and let ρ be a revision protocol. The mean dynamic corresponding to F and ρ is

$$(M) \quad \dot{x}_i^p = \sum_{j \in S^p} x_j^p \rho_{ji}^p(F^p(x), x^p) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(F^p(x), x^p).$$

4.2.2 Target Protocols and Target Dynamics

To conclude this section, we introduce a condition on revision protocols that is satisfied in many interesting examples, and that generates mean dynamics that are easy to describe in geometric terms.

We say that ρ is a *target protocol* if conditional switch rates under ρ do not depend on agents' current strategies: in other words, ρ_{ij}^p may depend on the candidate strategy j , but not on the incumbent strategy i . We can represent target protocols using maps of the form $\tau^p : \mathbf{R}^{n^p} \times X^p \rightarrow \mathbf{R}_+^{n^p}$, where $\rho_{ij}^p \equiv \tau_j^p$ for all $i \in S^p$. This restriction yields mean dynamics of

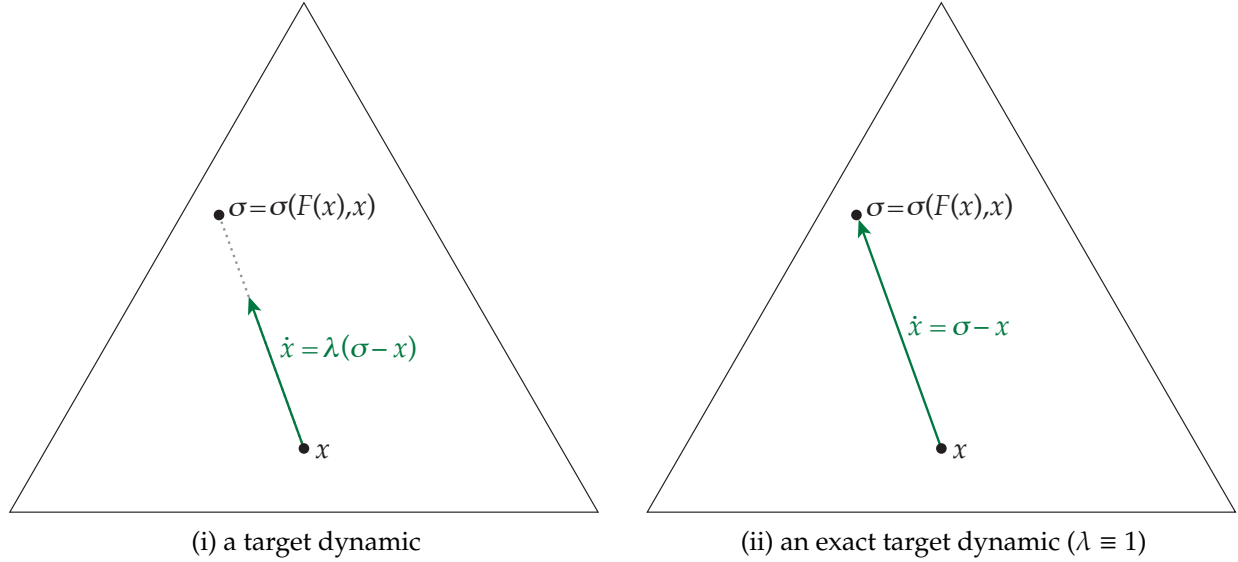


Figure 4.2.1: Target dynamics in a single population.

the form

$$(4.3) \quad \dot{x}_i^p = m^p \tau_i^p(F^p(x), x^p) - x_i^p \sum_{j \in S^p} \tau_j^p(F^p(x), x^p),$$

which we call *target dynamics*.

What is the geometric interpretation of these dynamics? If $\tau^p(\pi^p, x^p) \in \mathbf{R}_+^{m^p}$ is not the zero vector, we can define

$$\lambda^p(\pi^p, x^p) = \sum_{i \in S^p} \tau_i^p(\pi^p, x^p) \quad \text{and} \quad \sigma_i^p(\pi^p, x^p) = \frac{\tau_i^p(\pi^p, x^p)}{\lambda^p(\pi^p, x^p)}.$$

Then $\sigma^p(\pi^p, x^p) \in \Delta^p$ is a probability vector, and we can rewrite equation (4.3) as

$$(4.4) \quad \dot{x}^p = \begin{cases} \lambda^p(F^p(x), x^p) (m^p \sigma^p(F^p(x), x^p) - x^p) & \text{if } \tau^p(F^p(x), x^p) \neq \mathbf{0}, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

The first case of equation (4.4) tells us that the population state $x^p \in X^p$ moves in the direction of the target state $m^p \sigma^p \in X^p$, the representative of the probability vector $\sigma^p \in \Delta^p$ in the state space $X^p = m^p \Delta^p$; moreover, motion toward the target state proceeds at rate λ^p . Figure 4.2.1(i) illustrates this idea in the single population case; since here the population's mass is 1, the target state is just the probability vector $\sigma^p \in X^p = \Delta^p$.

Now suppose that protocol τ is an *exact target protocol*: a target protocol that is exact

with clock rate $R = 1$ (see equation (4.1) and the subsequent discussion). In this case, we call the resulting mean dynamic an *exact target dynamic*. Since exactness implies that $\lambda^p \equiv 1$, we often denote exact target protocols by σ rather than τ , emphasizing that the values of $\sigma^p : \mathbf{R}^{n^p} \times X^p \rightarrow \Delta^p$ are probability vectors. Exact target dynamics take the especially simple form

$$(4.5) \quad \dot{x}^p = m^p \sigma^p(F^p(x), x^p) - x^p.$$

The vector of motion in (4.5) can be drawn with its tail at the current state x^p and its head at the target state $m^p \sigma^p$, as illustrated in Figure 4.2.1(ii) in the single population case.

4.3 Examples

We now offer a number of examples of revision protocols and their mean dynamics that we will revisit throughout the remainder of the book. Recall that

$$\bar{F}^p(x) = \frac{1}{m^p} \sum_{i \in S^p} x_i^p F_i^p(x)$$

represents the average payoff obtained by members of population p . It is useful to define the *excess payoff* to strategy $i \in S^p$,

$$\hat{F}_i^p(x) = F_i^p(x) - \bar{F}^p(x),$$

as the difference between strategy i 's payoff and the average payoff in population p . The *excess payoff vector* for population p is written as

$$\hat{F}^p(x) = F^p(x) - \mathbf{1}\bar{F}^p(x).$$

To conserve on notation, the examples to follow are stated for the single population setting. When introducing revision protocols, we let $\pi \in \mathbf{R}^n$ denote an arbitrary payoff vector; when the population state $x \in X$ is also given, we let $\hat{\pi} = \pi - \mathbf{1}x'\pi$ denote the resulting excess payoff vector.

It is useful to place revision protocols in one of two categories: imitative protocols and direct protocols.

4.3.1 Imitative Protocols and Dynamics

We begin with examples of *imitative protocols*. These protocols are of the form

$$(4.6) \quad \rho_{ij}(\pi, x) = x_j r_{ij}(\pi, x),$$

with additional restrictions placed on the form of the function r_{ij} (see Section 5.4.1). Under an imitative protocol, an agent who receives a revision opportunity chooses an opponent at random and observes her strategy. If the agent is playing strategy i and the opponent strategy j , the agent switches from i to j with probability proportional to r_{ij} . Note that the value of x_j need not be observed; instead, the initial term on the right hand side of (4.6) reflects the agent's observation of a randomly chosen opponent.

Example 4.3.1. Pairwise proportional imitation. Suppose that after selecting an opponent, the agent imitates the opponent only if the opponent's payoff is higher than his own, doing so in with probability proportional to the payoff difference:

$$\rho_{ij}(\pi, x) = x_j [\pi_j - \pi_i]_+.$$

The mean dynamic generated by this revision protocol is

$$\begin{aligned} \dot{x}_i &= \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x) \\ &= \sum_{j \in S} x_j x_i [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} x_j [F_j(x) - F_i(x)]_+ \\ &= x_i \sum_{j \in S} x_j (F_i(x) - F_j(x)) \\ &= x_i (F_i(x) - \bar{F}(x)). \end{aligned}$$

This equation, which we can rewrite as

$$(R) \quad \dot{x}_i = x_i \hat{F}_i(x).$$

defines the *replicator dynamic*, the best known dynamic in evolutionary game theory. Under this dynamic, the percentage growth rate \dot{x}_i/x_i of each strategy currently in use is equal to that strategy's current excess payoff; unused strategies always remain so. §

Example 4.3.2. Pure imitation driven by dissatisfaction. Suppose that when a strategy i player receives a revision opportunity, he opts to switch strategies with a probability that is linearly decreasing in his current payoff. (For example, agents might revise when their

payoffs do not meet a uniformly distributed random aspiration level.) In the event that the agent decides to switch, he imitates a randomly selected opponent. This leads to the revision protocol

$$\rho_{ij}(\pi, x) = (K - \pi_i)x_j,$$

where the constant K is sufficiently large that conditional switch rates are always positive.

The mean dynamic generated by this revision protocol is

$$\begin{aligned} \dot{x}_i &= \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x) \\ &= \sum_{j \in S} x_j (K - F_j(x)) x_i - x_i (K - F_i(x)) \\ &= x_i \left(K - \sum_{j \in S} x_j F_j(x) - K + F_i(x) \right) \\ &= x_i \hat{F}_i(x). \end{aligned}$$

Thus, this protocol's mean dynamic is the replicator dynamic as well. §

Exercise 4.3.3. Imitation of success. Consider the revision protocol

$$\rho_{ij}(\pi, x) = \tau_j(\pi, x) = x_j(\pi_j - K),$$

where the constant K is smaller than any feasible payoff.

- (i) Offer an interpretation of this protocol.
- (ii) Show that this protocol generates the replicator dynamic as its mean dynamic.
- (iii) Part (ii) implies that the replicator dynamic is a target dynamic. Compute the rate $\lambda(F(x), x)$ and target state $\sigma(F(x), x)$ corresponding to population state x . Describe how these vary as one changes the value of K .

While the three protocols above all generate the replicator dynamic as their mean dynamics, the class of imitative dynamics contains many other members: see Section 5.4.2 for a variety of additional examples.

Exercise 4.3.4. In the single population setting, we call a mean dynamic an *anti-target dynamic* if it can be expressed as

$$\dot{x} = \tilde{\lambda}(F(x), x) \left(x - \tilde{\sigma}(F(x), x) \right),$$

where $\tilde{\lambda}(\pi, x) \in \mathbf{R}_+$ and $\tilde{\sigma}(\pi, x) \in \Delta$.

- (i) Give a geometric interpretation of anti-target dynamics.
- (ii) Show that the replicator dynamic is an anti-target dynamic.

4.3.2 Direct Protocols and Dynamics

Under imitative protocols, agents obtain candidate strategies by observing the behavior of randomly chosen opponents. In settings where agents are aware of the full set S of available strategies, we can assume instead that they choose candidate strategies directly, without having to see the choices of others. The analyses in the coming chapters will show that these *direct protocols* generate dynamics that differ in a number of important ways from those generated by imitative protocols, some fairly obvious, and others less so.

Example 4.3.5. Logit choice. Suppose that choices are made according to the *logit choice protocol*, the exact target protocol defined by

$$\rho_{ij}(\pi) = \sigma_j(\pi) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}.$$

The parameter $\eta > 0$ is called the *noise level*. If η is large, choice probabilities under the logit rule are nearly uniform. But if η is near zero, choices are optimal with probability close to one, at least when the difference between the best and second best payoff is not too small. By equation (4.5), the exact target dynamic generated by protocol σ is

$$\begin{aligned} \text{(L)} \quad \dot{x}_i &= \sigma_i(F(x)) - x_i \\ &= \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} - x_i. \end{aligned}$$

This is the *logit dynamic* with noise level η . §

Example 4.3.6. Comparison to the average payoff. Consider the target protocol

$$\rho_{ij}(\pi, x) = \tau_j(\pi, x) = [\hat{\pi}_j]_+ = \left[\pi_j - \sum_{k \in S} x_k \pi_k \right]_+.$$

When an agent's clock rings, he chooses a strategy at random. If that strategy's payoff is above average, the agent switches to it with probability proportional to its excess payoff. By equation (4.3), the induced target dynamic is

$$\text{(BNN)} \quad \dot{x}_i = \tau_i(F(x), x) - x_i \sum_{j \in S} \tau_j(F(x), x)$$

$$= [\hat{F}_i(x)]_+ - x_i \sum_{j \in S} [\hat{F}_j(x)]_+.$$

This is the *Brown-von Neumann-Nash (BNN) dynamic*. §

Example 4.3.7. Pairwise comparisons. Suppose that

$$\rho_{ij}(\pi) = [\pi_j - \pi_i]_+.$$

When an agent's clock rings, he selects a strategy at random. If the new strategy's payoff is higher than his current strategy's payoff, he switches strategies with probability proportional to the difference between the two payoffs. The resulting mean dynamic,

$$\begin{aligned} \text{(S)} \quad \dot{x}_i &= \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x) \\ &= \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+, \end{aligned}$$

is called the *Smith dynamic*. §

4.4 Deterministic Evolutionary Dynamics

With this background established, we now provide a formal definition of deterministic evolutionary dynamics. Let $\mathcal{P} = \{1, \dots, p\}$ be a set of populations with masses m^p and strategy sets S^p . Let X be the corresponding set of social states:

$$X = \{x \in \mathbf{R}_+^n : x = (x^1, \dots, x^p), \text{ where } \sum_{i \in S^p} x_i^p = m^p\}.$$

Define the sets \mathcal{F} and \mathcal{T} as follows:

$$\begin{aligned} \mathcal{F} &= \{F : X \rightarrow \mathbf{R}^n : F \text{ is Lipschitz continuous}\}; \\ \mathcal{T} &= \{\{x\}_{t \geq 0} \subseteq X : x_{(\cdot)} \text{ is continuous}\} \end{aligned}$$

\mathcal{F} is the set of population games with Lipschitz continuous payoffs; \mathcal{T} is the set of continuous forward-time trajectories through the state space X .

Definition. A deterministic evolutionary dynamic is a set-valued map $\mathbf{D} : \mathcal{F} \Rightarrow \mathcal{T}$. It assigns each population game $F \in \mathcal{F}$ a set of trajectories $\mathbf{D}(F) \subset \mathcal{T}$ satisfying

Existence and forward invariance: For each $\xi \in X$, there is a $\{x_t\}_{t \geq 0} \in \mathbf{D}(F)$ with $x_0 = \xi$.

Thus, for each game F and each initial condition $\xi \in X$, an evolutionary dynamic must specify at least one solution trajectory that begins at ξ and then remains in X at all positive times.

This definition of a deterministic evolutionary dynamic is rather general, in that it does not impose a uniqueness requirement (i.e., since it allows multiple trajectories in $\mathbf{D}(F)$ to emanate from a single initial condition). This generality allows us to handle dynamics defined by discontinuous differential equations and by differential inclusions—see Chapter 6. But for dynamics defined by Lipschitz continuous differential equations, this level of generality is unnecessary: in this case, standard results allow us to ensure not only existence of solutions, but also:

Uniqueness: For each $\xi \in X$, there is exactly one $\{x_t\}_{t \geq 0} \in \mathbf{D}(F)$ with $x_0 = \xi$.

Lipschitz continuity: For each $t, x_t = x_t(\xi)$ is a Lipschitz continuous function of ξ .

The basic results on existence and uniqueness of solutions to ordinary differential equations concern equations defined on open sets. To contend with the fact that our mean dynamics are defined on the compact, convex set X , we need conditions ensuring that solution trajectories do not leave this set. The required conditions are provided by Theorem 4.4.1: if the vector field $V_F : X \rightarrow \mathbf{R}^n$ is Lipschitz continuous, and if at each state $x \in X$, the growth rate vector $V_F(x)$ is contained in the tangent cone $TX(x)$, the set of directions of motion from x that do not point out of X , then all of our desiderata for solution trajectories are satisfied.

Theorem 4.4.1. *Suppose $V_F : X \rightarrow \mathbf{R}^n$ is Lipschitz continuous, and let $S(V_F) \subset \mathcal{T}$ be the set of solutions $\{x_t\}_{t \geq 0}$ to $\dot{x} = V_F(x)$. If $V_F(x) \in TX(x)$ for all $x \in X$, then $S(V_F)$ satisfies existence and forward invariance, uniqueness, and Lipschitz continuity.*

Theorem 4.4.1 follows directly from Theorems 4.A.2, 4.A.6, and 4.A.8 in Appendix 4.A.

In Section 4.2, we showed how a population game F and revision protocol ρ together define a mean dynamic $\dot{x} = V_F(x)$ via equation (M). Viewing this development in the present light, we can think of the revision protocol ρ by itself as defining a map from population games to differential equations, or, going one step further, to sets of forward-time trajectories through the state space X . To conclude this chapter, we show that as long as ρ is Lipschitz continuous, the map so defined is a deterministic evolutionary dynamic in the sense specified above. Indeed, by proposing a Lipschitz continuous revision protocol, one implicitly identifies a unique aggregate behavior trajectory from every initial social state and for every population game.

Theorem 4.4.2. Fix a Lipschitz continuous revision protocol ρ . For each population game $F \in \mathcal{F}$, let

$$(M) \quad \dot{x}_i^p = (V_F)_i^p(x) = \sum_{j \in S^p} x_j^p \rho_{ji}^p(F^p(x), x^p) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(F^p(x), x^p).$$

be the mean dynamic defined by F and ρ , and let

$$\mathbf{D}(F) = \{\{x_t\} \in \mathcal{T} : \{x_t\} \text{ solves } \dot{x} = V_F(x)\}$$

be the set of solutions of this dynamic. Then the map $\mathbf{D}: \mathcal{F} \Rightarrow \mathcal{T}$ so defined is a deterministic evolutionary dynamic. Indeed, for each $F \in \mathcal{F}$, the set $\mathbf{D}(F) \subset \mathcal{T}$ satisfies not only existence and forward invariance, but also uniqueness and Lipschitz continuity.

Proof. We begin with a simple observation about the form of the tangent cones $TX(x)$.

Observation 4.4.3. Let $z \in \mathbf{R}^n$. Then $z \in TX(x)$ if and only if

- (i) For all $p \in \mathcal{P}$, $\sum_{i \in S^p} z_i^p = 0$, and
- (ii) For all $i \in S^p$ and $p \in \mathcal{P}$, $x_i^p = 0$ implies that $z_i^p \geq 0$.

In words: for z to be a feasible vector of strategy growth rates when the social state is x , the growth rates for each population must sum to zero (so that population masses stay constant over time), and the growth rates of unused strategies must be nonnegative.

Since the population shares x_i^p and conditional switch rates $\rho_{ij}^p(\pi^p, x^p)$ are nonnegative, it is clear from equation (M) that the vector $V_F(x)$ satisfies the conditions from the observation, and so lies in $TX(x)$. Moreover, since F and ρ are Lipschitz continuous, V_F is Lipschitz continuous as well. The conclusions of the theorem thus follow from Theorem 4.4.1. ■

Appendix

4.A Ordinary Differential Equations

4.A.1 Basic Definitions

Every continuous vector field $V: \mathbf{R}^n \rightarrow \mathbf{R}^n$ defines an *ordinary differential equation* (ODE) on \mathbf{R}^n , namely

$$\frac{d}{dt}x_t = V(x_t).$$




Figure 4.A.1: A solution of an ordinary differential equation.

Often we write \dot{x}_t for $\frac{d}{dt}x_t$; we also express the previous equation as

$$(D) \quad \dot{x} = V(x).$$

Equation (D) describes the evolution of a state variable x_t over time. When the current state is x_t , the current velocity of state—in other words, the speed and direction of the change in the state—is $V(x_t)$. The trajectory $\{x_t\}_{t \in I}$ is a *solution* to (D) if $\dot{x}_t = V(x_t)$ at all times t in the interval I , so that at each moment, the time derivative of the trajectory is described by the vector field V (see Figure 4.A.1).

In many applications, one is interested in solving an *initial value problem*: that is, in characterizing the behavior of solution(s) to (D) that start at a given initial condition $\xi \in \mathbf{R}^n$.

Example 4.A.1. Exponential growth and decay. The simplest differential equation is the linear equation $\dot{x} = ax$ on the real line. What are the solutions to this equation starting from initial condition $\xi \in \mathbf{R}$? It is easy to verify that $x_t = \xi \exp(at)$ is a solution to this equation on the full time interval $(-\infty, \infty)$, since

$$\frac{d}{dt}x_t = \frac{d}{dt}(\xi \exp(at)) = a(\xi \exp(at)) = ax_t,$$

as required. This solution describes a process of exponential growth or decay according to whether a is positive or negative.

In fact, $x_t = \xi \exp(at)$ is the only solution to $\dot{x} = ax$ from initial condition ξ . If $\{y_t\}$ is a solution to $\dot{x} = ax$ from any initial condition, then

$$\frac{d}{dt} (y_t \exp(-at)) = \dot{y}_t \exp(-at) - ay_t \exp(-at) = 0.$$

Hence, $y_t \exp(-at)$ is constant, and so $y_t = \psi \exp(at)$ for some $\psi \in \mathbf{R}$. Since $y_0 = \xi$, it must be that $\psi = \xi$. §

4.A.2 Existence, Uniqueness, and Continuity of Solutions

Except in cases where the state variable x is one dimensional or the vector field V is linear, explicit solutions to ODEs are usually impossible to obtain. To investigate dynamics for which explicit solutions are unavailable, one begins by verifying that a solution exists and is unique, and then uses various indirect methods to determine its properties.

The main tool for ensuring existence and uniqueness of solutions to ODEs is the *Picard-Lindelöf Theorem*. To state this result, fix an open set $O \subseteq \mathbf{R}^n$. We call the function $f : O \rightarrow \mathbf{R}^m$ *Lipschitz continuous* if there exists a scalar K such that

$$|f(x) - f(y)| \leq K|x - y| \text{ for all } x, y \in O.$$

More generally, we say that f is *locally Lipschitz continuous* if for all $x \in O$, there exists an open neighborhood $O_x \subseteq O$ containing x such that the restriction of f to O_x is Lipschitz continuous. It is easy to verify that every C^1 function is locally Lipschitz.

Theorem 4.A.2 (The Picard-Lindelöf Theorem). *Let $V : O \rightarrow \mathbf{R}^n$ be locally Lipschitz continuous. Then for each $\xi \in O$, there exists a scalar $T > 0$ and a unique trajectory $x : (-T, T) \rightarrow O$ with $x_0 = \xi$ such that $\{x_t\}$ is a solution to (D).*

The Picard-Lindelöf Theorem is proved using the method of successive approximations. Given an approximate solution $x^k : (-T, T) \rightarrow O$ with $x_0^k = \xi$, one constructs a new trajectory $x^{k+1} : (-T, T) \rightarrow O$ using the map C that is defined as follows:

$$x_t^{k+1} = C(x^k)_t \equiv \xi + \int_0^t V(x_s^k) ds.$$

It is easy to see the trajectories $\{x_t\}$ that are fixed points of C are the solutions to (D) with $x_0 = \xi$. Thus, if C has a unique fixed point, the theorem is proved. But it is possible to show that if T is sufficiently small, then C is a contraction in the supremum norm; therefore, the desired conclusion follows from the *Banach* (or *Contraction Mapping*) *Fixed Point Theorem*.

If V is continuous but not Lipschitz, *Peano's Theorem* tells us that solutions to (D) exist, but in this case solutions need not be unique. The following example shows that when V does not satisfy a Lipschitz condition, so that small changes in x can lead to arbitrarily large changes in $V(x)$, it is possible for solution trajectories to escape from states at which the velocity under V is zero.

Example 4.A.3. Consider the ODE $\dot{x} = 3x^{2/3}$ on \mathbf{R} . The right hand side of this equation is continuous, but it fails to be Lipschitz at $x = 0$. One solution to this equation from initial condition $\xi = 0$ is the stationary solution $x_t \equiv 0$. Another solution is given by $x_t = t^3$. In fact, for each $t_0 \in [0, \infty)$, the trajectory that equals 0 until time t_0 and satisfies $x_t = (t - t_0)^3$ thereafter is also a solution. §

The Picard-Lindelöf Theorem guarantees the existence of a solution to (D) over some open interval of times. This open interval need not be the full time interval $(-\infty, \infty)$, as the following example illustrates.

Example 4.A.4. Consider the C^1 ODE $\dot{x} = x^2$ on \mathbf{R} . The unique solution with initial condition $\xi = 1$ is $x_t = \frac{1}{1-t}$. This solution exists for all negative times, but it “explodes” in forward time at $t = 1$. §

When V is locally Lipschitz, one can always find a maximal open time interval over which the solution to (D) from initial condition ξ exists in the domain O . If V is defined throughout \mathbf{R}^n and is bounded, then the speed of all solution trajectories is bounded as well, which implies that solutions exist for all time.

Theorem 4.A.5. *If $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is locally Lipschitz continuous and bounded, then for each $\xi \in \mathbf{R}^n$, then $\{x_t\}$, the unique solution to (D) with $x_0 = \xi$, exists for all $t \in (-\infty, \infty)$.*

We will often find it convenient to discuss solutions to (D) from more than one initial condition at the same time. To accomplish this most easily, we introduce the flow of differential equation (D).

Suppose that $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is Lipschitz continuous, and let $A \subseteq \mathbf{R}^n$ be an *invariant* set under (D): that is, solutions to (D) with initial conditions in A exist and remain in A at all times $t \in (-\infty, \infty)$. Then the *flow* $\phi : (-\infty, \infty) \times A \rightarrow A$ generated by (D) is defined by $\phi_t(\xi) = x_t$, where $\{x_t\}_{t \in (-\infty, \infty)}$ is the solution to (D) with initial condition $x_0 = \xi$. If we fix $\xi \in A$ and vary t , then $\{\phi_t(\xi)\}_{t \in (-\infty, \infty)}$ is the solution orbit of (D) through initial condition ξ ; note also that ϕ satisfies the group property $\phi_t(\phi_s(\xi)) = \phi_{s+t}(\xi)$. If we instead fix t and vary ξ , then $\{\phi_t(\xi)\}_{\xi \in A'}$ describes the positions at time t of solutions to (D) with initial conditions in $A' \subseteq A$.

Using this last notational device, we can describe the continuous variation of solutions to (D) in their initial conditions.

Theorem 4.A.6. *Suppose that $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is Lipschitz continuous with Lipschitz constant K , and that $A \subseteq \mathbf{R}^n$ is invariant under (D). Let ϕ be the flow of (D), and fix $t \in (-\infty, \infty)$. Then $\phi_t(\cdot)$ is Lipschitz continuous with Lipschitz constant $e^{K|t|}$: for all $\xi, \chi \in A$, we have that $|\phi_t(\xi) - \phi_t(\chi)| \leq |\xi - \chi| e^{K|t|}$.*

The assumption that A is invariant is only made for notational convenience; the theorem is valid as long as solutions to (D) from ξ and χ exist throughout the time interval from 0 to t .

The proof of Theorem 4.A.6 is a direct consequence of the following inequality, which is of importance in its own right.

Lemma 4.A.7 (Grönwall's Inequality). *Let $z : [0, T] \rightarrow \mathbf{R}_+$ be continuous. Suppose $C \geq 0$ and $K \geq 0$ are such that $z_t \leq C + \int_0^t Kz_s ds$ for all $t \in [0, T]$. Then $z_t \leq Ce^{Kt}$ for all $t \in [0, T]$.*

If we set $z_t = |\phi_t(\xi) - \phi_t(\chi)|$, then the antecedent inequality in the lemma is satisfied when $C = |\xi - \chi|$ and K is the Lipschitz constant for V , so Theorem 4.A.6 immediately follows. Also note that setting $\xi = \chi$ establishes the uniqueness of solutions to (D) from each initial condition.

4.A.3 Ordinary Differential Equations on Compact Convex Sets

The Picard-Lindelöf Theorem concerns ODEs defined on open subsets of \mathbf{R}^n . In contrast, evolutionary dynamics for population games are defined on the set of population states X , which is compact and convex. Fortunately, existence and uniqueness of *forward time* solutions can still be established in this setting.

To begin, we introduce the notion of forward invariance. The set $C \subseteq \mathbf{R}^n$ is *forward invariant* under the Lipschitz ODE (D) if every solution to (D) that starts in C at time 0 remains in C at all positive times: if $\{x_t\}$ is the solution to (D) from $\xi \in C$, then x_t exists and lies in C at all $t \in [0, \infty)$.

When C is forward invariant but not necessarily invariant under (D), we can speak of the *semiflow* $\phi : [0, \infty) \times A \rightarrow A$ generated by (D). While semiflows are not defined for negative times, they resemble flows in many other respects: by definition, $\phi_t(\xi) = x_t$, where $\{x_t\}_{t \geq 0}$ is the solution to (D) with initial condition $x_0 = \xi$; also, ϕ is continuous in t and ξ , and ϕ satisfies the group property $\phi_t(\phi_s(\xi)) = \phi_{s+t}(\xi)$.

Now suppose that the domain of the vector field V is a compact, convex set C . Intuition suggests that as long as V never points outward from C , solutions to (D) should be well defined and remain in C for all positive times.

Theorem 4.A.8 tells us that if we are given a Lipschitz continuous vector field V that is defined on a compact convex set C and that never points outward from the boundary

of C , then the ODE $\dot{x} = V(x)$ leaves C forward invariant. If in addition the negation of V never points outward from the boundary of C , then C is both forward and backward invariant under the ODE.

Theorem 4.A.8. *Let $C \subset \mathbf{R}^n$ be a compact convex set, and let $V : C \rightarrow \mathbf{R}^n$ be Lipschitz continuous.*

- (i) *Suppose that $V(\hat{x}) \in TC(\hat{x})$ for all $\hat{x} \in C$. Then for each $\xi \in C$, there exists a unique $x : [0, \infty) \rightarrow C$ with $x_0 = \xi$ that solves (D).*
- (ii) *Suppose that $V(\hat{x}) \in TC(\hat{x}) \cap (-TC(\hat{x}))$ for all $\hat{x} \in C$. Then for each $\xi \in C$, there exists a unique $x : (-\infty, \infty) \rightarrow C$ with $x_0 = \xi$ that solves (D).*

Proof. (i) Let $\bar{V} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ be the extension of $V : C \rightarrow \mathbf{R}^n$ defined by $\bar{V}(y) = V(\Pi_C(y))$, where $\Pi_C : \mathbf{R}^n \rightarrow C$ is the closest point projection onto C (see Section 2.A.4). Then \bar{V} is Lipschitz continuous and bounded. Thus, Theorem 4.A.5 tells us that the ODE

$$(4.7) \quad \dot{y} = \bar{V}(y)$$

admits unique solutions from all initial conditions in \mathbf{R}^n , and that these solutions exist for all (forward and backward) time. Now let $\xi \in C$, let $\{x_t\}_{t \in (-\infty, \infty)}$ be the unique solution to (4.7) with $x_0 \in C$, and suppose that $x_t \in C$ for all positive t ; then since V and \bar{V} agree on C , $\{x_t\}_{t \geq 0}$ must be the unique forward solution to (D) with $x_0 = \xi$. Thus, to prove our result, it is enough to show that the set C is forward invariant under the dynamic (4.7).

Define the squared distance function $\delta_C : \mathbf{R}^n \rightarrow \mathbf{R}$ by

$$\delta_C(y) = \min_{x \in C} |y - x|^2.$$

One can verify that δ_C is differentiable with gradient

$$\nabla \delta_C(y) = 2(y - \Pi_C(y)).$$

Hence, if $\{y_t\}$ is a solution to (4.7), then the chain rule tells us that

$$(4.8) \quad \frac{d}{dt} \delta_C(y_t) = \nabla \delta_C(y_t)' \dot{y}_t = 2(y_t - \Pi_C(y_t))' \bar{V}(y_t) = 2(y_t - \Pi_C(y_t))' V(\Pi_C(y_t)).$$

Suppose we could show that this quantity is bounded above by zero (i.e., that when $y_t - \Pi_C(y_t)$ and $V(\Pi_C(y_t))$ are nonzero, the angle between them is weakly obtuse.) This would imply that the distance between y_t and C is nonincreasing over time, which would in turn imply that C is forward invariant under (4.7). (The function δ_C is thus a *Lyapunov function* for the set C under the dynamic (4.7); see Appendix 7.B.)

We divide the analysis into two cases. If $y_t \in C$, then $y_t = \Pi_C(y_t)$, so expression (4.8)




Figure 4.A.2: The proof of Theorem 4.A.8.

evaluates to zero. On the other hand, if $y_t \notin C$, then the difference $y_t - \Pi_C(y_t)$ is in the normal cone $NC(\Pi_C(y_t))$ (see Figure 4.A.2). Since $V(\Pi_C(y_t)) \in TC(\Pi_C(y_t))$, it follows that $(y_t - \Pi_C(y_t))' V(\Pi_C(y_t)) \leq 0$, so the proof is complete.

(ii) If $V(\hat{x}) \in TC(\hat{x}) \cap (-TC(\hat{x}))$, then a slight modification of the argument above shows that $\frac{d}{dt} \delta_C(y_t) = 2(y_t - \Pi_C(y_t))' V(\Pi_C(y_t)) = 0$, and so that the distance between y_t and C is constant over time under the dynamic (4.7). Therefore, C is both forward and backward invariant under (4.7), and hence under (D) as well. ■

4.N Notes

Sections 4.1 and 4.2: Björnerstedt and Weibull (1996) introduce a version of the revision protocol model and derive the mean dynamics associated with certain imitative decision rules; see Weibull (1995, Sections 4.4 and 5.3) for a summary. The model studied here builds on Benaïm and Weibull (2003) and Sandholm (2003, 2008b). Versions of target dynamics are considered in Sandholm (2005a) and Hofbauer and Sandholm (2009).

The evolutionary approach is not the only alternative for modeling recurring strategic interactions in large populations. The perfect foresight dynamics of Matsui and Matsuyama (1995) retain the assumption of inertia from evolutionary approach, but replace the assumptions of myopia and limited information with those of dynamic optimization and equilibrium knowledge. For recent work on perfect foresight dynamics, see Hofbauer

and Sorger (1999), Oyama (2002), Matsui and Oyama (2006), Kojima (2006a), Kojima and Takahashi (2008), Takahashi (2008), Oyama et al. (2008), and Oyama and Tercieux (2009).

The traditional approach to studying repeated strategic interactions is the theory of repeated games, which considers equilibrium behavior among small groups of forward-looking, non-anonymous agents. Mailath and Samuelson (2006) provide an excellent and comprehensive presentation of this subject.

Section 4.3: The replicator dynamic was introduced by Taylor and Jonker (1978), but is closely related to a number of older models from mathematical biology—see Schuster and Sigmund (1983). The latter authors coined the term “replicator dynamic”, borrowing the term “replicator” from Dawkins (1976, 1982). Example 4.3.1, Example 4.3.2, and Exercise 4.3.3 are due to Schlag (1998), Björnerstedt and Weibull (1996), and Hofbauer (1995a), respectively.

The logit dynamic is studied by Fudenberg and Levine (1998) and Hofbauer and Sandholm (2002, 2007). The BNN dynamic was introduced in the context of symmetric zero sum games by Brown and von Neumann (1950). Nash (1951) uses a discrete version of this dynamic to devise a proof of existence of Nash equilibrium in normal form games based on Brouwer’s fixed point theorem. The Smith dynamic was introduced by Smith (1984) to study the dynamics of traffic flow. Generalizations of all of the dynamics from this section are studied in the next two chapters, where additional references can be found.

Appendix 4.A: Hirsch and Smale (1974) and Robinson (1995) are fine introductions to ordinary differential equations at the undergraduate and graduate levels, respectively. Theorem 4.A.8 is adapted from Smirnov (2002, Theorem 5.7).

Deterministic Dynamics: Families and Properties

5.0 Introduction

In the model of stochastic evolution introduced in Chapter 4, a large society of agents recurrently play a population game F by applying a revision protocol ρ . Through an informal appeal to the law of large numbers (one that will be made formal in Chapter 10), we argued that aggregate behavior in the society can be described by a differential equation on the set of social states X . This differential equation, which we call the *mean dynamic*, describes the inflows and outflows of agents to and from each available strategy:

$$(M) \quad \dot{x}_i^p = \sum_{j \in S^p} x_j^p \rho_{ji}^p(F^p(x), x^p) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(F^p(x), x^p).$$

Each revision protocol ρ can be viewed as defining a map from population games F to differential equations (M). It is this map from games to differential equations that we call a *deterministic evolutionary dynamic*.

In this chapter, we begin to investigate revision protocols and evolutionary dynamics in a systematic fashion. One of our primary goals in doing so is to understand the connections between the predictions of evolutionary dynamics and those of the traditional solution concepts of game theory, particularly Nash equilibrium. As we discussed in Sections 1.2 and 4.0, the prediction of Nash equilibrium play is commonly justified by strong assumptions about agents' knowledge, both about other agents' behavior and about the game itself. In this chapter, we aim to determine whether and when Nash equilibria correspond to stationary patterns of behavior among agents with limited information about their strategic environment.

We begin in Section 5.1 by considering the informational demands that different revision protocols impose. We argue first that revision protocols should be continuous functions of their inputs: in large-population environments, very precise information about aggregate behavior may be difficult to obtain, prohibiting agents from implementing revision protocols that are extremely sensitive to the exact values of payoffs or the population state. Next, we categorize revision protocols by the pieces of data they require. One key distinction made here is between protocols which only condition on payoffs, and those that require information about the social state itself. Protocols that respect our continuity and data requirements do not make unrealistic demands on the amount of information that agents in an evolutionary model must possess, and so can provide a convincing basis for predicting Nash equilibrium play.

Section 5.2 takes a further step toward formalizing this goal, introducing conditions that relate aggregate behavior under evolutionary dynamics to incentives in the underlying games. The first such condition, *positive correlation* (PC), requires that whenever a population is not at rest, its strategies' growth rates are positively correlated with their payoffs. This condition demands a weak but fundamental connection between individual incentives and disequilibrium aggregate dynamics, and is a key building block in our analyses of dynamic stability in Chapters 7–9. The second condition, *Nash stationarity* (NS), asks that the rest points of the mean dynamic be precisely the Nash equilibria of the game being played. Dynamics satisfying (NS) evince a basic agreement between the evolutionary dynamic and game-theoretic notions of equilibrium play.

Each revision protocol specifies a particular functional relationship between its inputs (payoffs and population states) and its outputs (conditional switch rates). To the extent that the precise form of this functional relationship affects our results, our evolutionary justification of the Nash prediction is weakened. Fortunately, many of the results in this book depend only on the qualitative properties of protocols and dynamics, not on details about functional forms. We begin to make this idea precise in Section 5.3, where we introduce *families of evolutionary dynamics*: these are collections of evolutionary dynamics derived from qualitatively similar revision protocols.

Our study of these families begins in Section 5.4, which introduces *imitative dynamics*, the most thoroughly studied dynamics in the evolutionary literature. Imitative dynamics, exemplified by the replicator dynamic, have many appealing properties. They can be defined using protocols that require very little information, and they satisfy the disequilibrium condition (PC); the latter fact implies that all Nash equilibria are rest points of these dynamic (see Proposition 5.2.1). However, because pure imitation precludes the introduction of unused strategies, imitative dynamics admit rest points that are not Nash

equilibria, and so fail Nash stationarity (NS).

In Section 5.5, we introduce *excess payoff dynamics*. These dynamics are derived from direct protocols (see Section 4.3.2) under which agents compare the payoffs of candidate strategies with the average payoff obtained in their population. While excess payoff dynamics satisfy both positive correlation (PC) and Nash stationarity (NS), we argue that their strong informational requirements leave them unsuited to provide a broadly applicable justification of the Nash prediction.

We provide such a justification in Section 5.6, where we define *pairwise comparison dynamics*. The direct revision protocols that generate these dynamics only require agents to compare the payoffs of their current and candidate strategies. Pairwise comparison dynamics are not only informationally undemanding, but also satisfy positive correlation (PC) and Nash stationarity (NS), and so yield an appealing justification of the Nash prediction.

Because imitation is a basic feature of human decision process, the fact that imitative dynamics fail one of our basic desiderata, Nash stationarity (NS), may seem a cause for concern. Fortunately, this conclusion is somewhat misleading: we show in Section 5.7 that the failure of Nash stationarity is not due to imitation per se, but rather to its exclusive use. Under hybrid protocols that temper imitation with occasional direct selection of alternative strategies, Nash stationarity (NS) is achieved.

While this chapter focuses on the connection between stationary states and Nash equilibria, a more compelling justification of the Nash prediction would go further, showing that evolutionary dynamics lead to Nash equilibrium from disequilibrium states. This fundamental question requires sophisticated analyses, and is the focus of Chapters 7–9.

5.1 Information Requirements for Revision Protocols

In the model of evolution introduced in Chapter 4, we used revision protocols to describe the process through which individual agents update their choices over time. By design, the behavior of agents who employ revision protocols is characterized by inertia and myopia. We argued in Chapter 4 that inertia and myopia are natural assumptions to make in the context of population games, particularly when the game in question is one of many interactions that agents concurrently face.

In environments of this sort, one might also expect agents to possess limited information about current strategic conditions. In interactions involving large numbers of agents, exact and comprehensive information about payoffs and (especially) opponents' behavior may be difficult or costly to obtain. Agents who are content to revise their choices

sporadically and myopically seem unlikely to make the efforts necessary to gain precise knowledge about the current strategic environment. With this concern in mind, we now introduce conditions that help us to assess the informational demands that a revision protocol imposes.

It seems contrary to the evolutionary paradigm to posit revision protocols that are extremely sensitive to the exact values of payoffs or of the population state. This concern is reflected in the following condition, which requires that agents' revision protocols be continuous functions of payoffs and the state.

(C) *Continuity:* ρ is Lipschitz continuous.

Less formally, condition (C) asks that small changes in aggregate behavior not lead to large changes in players' responses. (We require Lipschitz continuity rather than just continuity because it implies that solutions to the mean dynamic are unique: see Section 4.4 and Appendix 4.A.)

Revision protocols also vary in terms of the specific pieces of data needed to implement them. A particularly simple protocol might only condition on the payoff to the agent's current strategy. Others might require agents to gather information about the payoffs to other strategies, whether by briefly experimenting with these strategies, or by asking others about their experiences with them. Still other protocols might require data beyond that provided by payoffs alone.

Before introducing our data requirements for revision protocols, we must single out an important case that requires special handling. Section 4.3 (and in Section 5.4 below) consider imitative protocols, which take the form

$$\rho_{ij}^p(\pi^p, x^p) = \frac{x_j^p}{m^p} r_{ij}^p(\pi^p, x^p).$$

An imitative protocol can be interpreted as a description of a two-stage revision procedure. When an agent using an imitative protocol receives a revision opportunity, he first selects a member of his population at random and observes this member's strategy; thus, the probability that a strategy j player is observed is x_j^p/m^p . Strategy j becomes the revising agent's candidate strategy. The agent switches to this candidate strategy with probability proportional to $r_{ij}^p(\pi^p, x^p)$. We do not view the act of observing the strategy of a randomly chosen opponent as imposing an informational burden. For this reason, we express our data requirements for imitative protocols in terms of the function r^p rather than the full protocol ρ^p .

Having noted this special case, we proceed by introducing our collection of data requirements.

- (D1) ρ_{ij}^p (or r_{ij}^p) depends only on π_i^p .
- (D1') ρ_{ij}^p (or r_{ij}^p) depends only on π_j^p .
- (D2) ρ_{ij}^p (or r_{ij}^p) depends only on π_i^p and π_j^p .
- (Dn) ρ_{ij}^p (or r_{ij}^p) depends on $\pi_1^p, \dots, \pi_{n^p}^p$, but not on $x_1^p, \dots, x_{n^p}^p$.
- (D+) ρ_{ij}^p (or r_{ij}^p) depends on $\pi_1^p, \dots, \pi_{n^p}^p$ and on $x_1^p, \dots, x_{n^p}^p$.

Protocols in classes (D1) and (D1') require only a single piece of payoff data: either the payoff of the agent's current strategy (under (D1)), or the payoff of the agent's candidate strategy (under (D1')). Protocols in class (D2) are slightly more demanding, as they require agents to know both of these strategies' payoffs. Protocols in class (Dn) require agents to know the payoffs of additional strategies.

Finally, protocols in class (D+) require not only information about the strategies' payoffs, but also information about the strategies' utilization levels. Unless information about these levels is provided by a central planner, we do not expect it to be readily available to agents in typical large-population settings.

To illustrate these conditions, we recall some examples of revision protocols from Section 4.3. As all of these examples satisfy continuity (C), we focus our attention on data requirements.

Example 5.1.1. The following three imitative protocols generate the replicator dynamic as their mean dynamics:

$$(5.1) \quad \rho_{ij}^p(\pi^p, x^p) = (K^p - \pi_i^p) \frac{x_j^p}{m^p},$$

$$(5.2) \quad \rho_{ij}^p(\pi^p, x^p) = \frac{x_j^p}{m^p} (\pi_j^p - K^p),$$

$$(5.3) \quad \rho_{ij}^p(\pi^p, x^p) = \frac{x_j^p}{m^p} [\pi_j^p - \pi_i^p]_+.$$

(In equations (5.1) and (5.2), we assume the constant K^p is chosen so that $\rho_{ij}^p(\pi^p, x^p) \geq 0$.) Protocol (5.1), *pure imitation driven by dissatisfaction*, is in class (D1). Protocol (5.2), *imitation of success*, is in class (D1'). Protocol (5.3), *pairwise proportional imitation*, is in class (D2). Evidently, each of these imitative protocols makes limited informational demands. §

Example 5.1.2. The logit dynamic is derived from the exact target protocol

$$\rho_{ij}^p(\pi^p) = \sigma_j^p(\pi^p) = \frac{\exp(\eta^{-1}\pi_j^p)}{\sum_{k \in S^p} \exp(\eta^{-1}\pi_k^p)}.$$

This protocol conditions on all strategies' payoffs, and so is in class (Dn). §

Example 5.1.3. The target protocol

$$\rho_{ij}^p(\pi^p, x^p) = \tau_j^p(\pi^p, x^p) = [\hat{\pi}_j^p]_+ = \left[\pi_j^p - \frac{1}{m^p} \sum_{k \in S^p} x_k^p \pi_k^p \right]_+$$

induces the BNN dynamic as its mean dynamic. This protocol conditions on strategy j 's excess payoff $\hat{\pi}_j^p$, and hence on the population average payoff $\frac{1}{m^p} (x^p)' \pi^p$. Since computing this average payoff requires knowledge of the payoffs and utilization levels of all strategies, this protocol is in class (D+). §

5.2 Incentives and Aggregate Behavior

The conditions introduced in the previous section provide ways of gauging the informational demands of revision protocols, and so can help us determine whether a protocol is appropriate for use in a given application. The properties defined in this section are of a different nature: rather than appraising the plausibility of a revision protocol for modeling boundedly rational behavior, they assess the extent to which the aggregate dynamics induced by the protocol respect the incentives provided by the underlying game.

To begin, we introduce our basic restriction on disequilibrium dynamics.

(PC) *Positive correlation:* $V_F^p(x) \neq \mathbf{0}$ implies that $V_F^p(x)' F^p(x) > 0$.

As we argue next, positive correlation requires that whenever a population is not at rest, the covariance between its strategies' growth rates and payoffs is positive; here, the covariance is defined with respect to a uniform random draw from the strategies in S^p .

To make this interpretation precise, we view the strategy set $S^p = \{1, \dots, n^p\}$ as a probability space endowed with the uniform probability measure: $\mathbb{P}(\{i\}) = \frac{1}{n^p}$ for all $i \in S^p$. Then for any fixed $x \in X$, we can think of the vectors $V_F^p(x) \in \mathbf{R}^{n^p}$ and $F^p(x) \in \mathbf{R}^{n^p}$ as functions from S^p to \mathbf{R} , or in other words, as scalar-valued random variables for the probability model (S^p, \mathbb{P}) . (For background on probability models, see Appendix 10.B.) It is therefore meaningful to speak of the covariance between $V_F^p(x)$ and $F^p(x)$.

With this background in hand, we now verify that

$$(5.4) \quad \text{Cov}(V_F^p(x), F^p(x)) = \frac{1}{n^p} V_F^p(x)' F^p(x),$$

which implies that condition (PC) admits the interpretation claimed above. To establish (5.4), observe first that when the vector $V_F^p(x)$ is viewed as a random variable on (S^p, \mathbb{P}) , its expectation is 0. This follows from the fact that the dynamic V_F keeps population p 's

mass constant (i.e., that $V_F^p(x) \in TX^p$): indeed,

$$\mathbb{E}(V_F^p(x)) \equiv \sum_{k \in S^p} \mathbb{P}(\{k\}) V_{F,k}^p(x) = \sum_{k \in S^p} \frac{1}{n^p} V_{F,k}^p(x) = 0.$$

A standard formula for the covariance then yields

$$\begin{aligned} \text{Cov}(V_F^p(x), F^p(x)) &= \mathbb{E}(V_F^p(x) F^p(x)) - \mathbb{E}(V_F^p(x)) \mathbb{E}(F^p(x)) \\ &= \sum_{k \in S^p} \mathbb{P}(\{k\}) V_{F,k}^p(x) F_k^p(x) - 0 \\ &= \frac{1}{n^p} V_F^p(x)' F^p(x), \end{aligned}$$

which is equation (5.4).

While the description of condition (PC) in terms of covariances is the most natural one from a game-theoretic point of view, the condition also has a very useful geometric interpretation. Recall that a pair of vectors emanating from the same point form an acute angle if and only if their inner product is positive. (This follows from the *Law of Cosines*: if θ is the angle between vectors u and v , then $u'v = |u||v| \cos(\theta)$.) Viewed in this light, condition (PC) requires that whenever the growth rate vector $V_F^p(x)$ is nonzero, the angle it forms with the payoff vector $F^p(x)$ is acute (see Examples 5.2.2 and 5.2.4 below). Ignoring considerations of feasibility (which are discussed next), condition (PC) allows considerable latitude in choosing a direction of motion for population p : the entire half-space of directions forming an acute angle with $F^p(x)$ are permitted.

Of course, this constraint on the direction of motion is in addition to those imposed by forward invariance. In other words, the growth rate vector $V_F^p(x)$ must be an element of the tangent cone $TX^p(x^p)$: the vector's components must sum to zero, and it must not assign negative growth rates to unused strategies (cf Observation 4.4.3). Because of these constraints, the direction of motion $V_F^p(x)$ often must distort the direction of the payoff vector $F^p(x)$ in order to remain feasible. (We will define the dynamic that minimizes this distortion, the *projection dynamic*, in Section 6.3.)

Our second condition on evolutionary dynamics asks for exact agreement between stationary behavior and the basic game-theoretic notion of equilibrium play.

(NS) *Nash stationarity*: $V_F(x) = \mathbf{0}$ if and only if $x \in NE(F)$.

Nash stationarity requires that the Nash equilibria of the game F and the rest points of the dynamic V_F coincide. It can be split into two distinct restrictions. First, (NS) asks that every Nash equilibrium of F be a rest point of V_F . If state x is a Nash equilibrium, then no

agent benefits from switching strategies; (NS) demands that in this situation, aggregate behavior is at rest under V_F . (This does not imply that individual agents' behavior is also at rest: see the discussion two paragraphs below.)

Second, condition (NS) asks that every rest point of V_F be a Nash equilibrium of F . If the current population state is not a Nash equilibrium, then by definition there are agents who would benefit from switching strategies. Condition (NS) requires that some of these agents eventually avail themselves of this opportunity.

In interpreting condition (NS), it is worth emphasizing that rest points of the mean dynamic V_F need not be states at which the underlying stochastic process is itself at rest. When we derived the mean dynamic in Section 4.2, we saw that $V_F(x)$ captures the *expected* changes in the use of each strategy under the stochastic evolutionary process. It is thus entirely possible that a rest point of $V_F(x)$ represents a balancing of stochastic forces rather than the absence of such forces. Indeed, some of the revision protocols we consider (like logit choice (Example 4.3.5)) always assign positive probabilities to all feasible transitions; under such protocols, all rest points of the mean dynamics must be balance points rather than dead spots of the stochastic evolutionary process.

There is an important link between our two restrictions on aggregate behavior: the out-of-equilibrium condition (PC) implies that all Nash equilibria of F are rest points under V_F . This is easiest to see in the single population setting. If x is a Nash equilibrium of F , then $F(x)$ is in the normal cone of X at x . Since $V_F(x)$ is a feasible direction of motion from x , it is in the tangent cone of X at x ; thus, the angle between $F(x)$ and $V_F(x)$ cannot be acute. Positive correlation therefore implies that x is a rest point of V_F .

More generally, we have the following result.

Proposition 5.2.1. *If V_F satisfies (PC), then $x \in NE(F)$ implies that $V_F(x) = \mathbf{0}$.*

Proof. Suppose that V_F satisfies (PC) and that $x \in NE(F)$. Recall that

$$x \in NE(F) \Leftrightarrow F(x) \in NX(x) \Leftrightarrow [v'F(x) \leq 0 \text{ for all } v \in TX(x)].$$

Now fix $p \in \mathcal{P}$, and define the vector $v \in \mathbf{R}^n$ by $v^p = V_F^p(x)$ and $v^q = \mathbf{0}$ for $q \neq p$. Then $v \in TX(x)$ by construction, and so $V_F^p(x)'F^p(x) = v'F(x) \leq 0$. Condition (PC) then implies that $V_F^p(x) = \mathbf{0}$. Since p was arbitrary, we conclude that $V_F(x) = \mathbf{0}$. ■

Example 5.2.2. Consider the two-strategy coordination game

$$F(x) = \begin{pmatrix} F_1(x) \\ F_2(x) \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \end{pmatrix},$$




Figure 5.2.1: Condition (PC) in 12 Coordination.

and the replicator dynamic for this game,

$$V(x) = \begin{pmatrix} V_1(x) \\ V_2(x) \end{pmatrix} = \begin{pmatrix} x_1 \hat{F}_1(x) \\ x_2 \hat{F}_2(x) \end{pmatrix} = \begin{pmatrix} x_1 (x_1 - ((x_1)^2 + 2(x_2)^2)) \\ x_2 (2x_2 - ((x_1)^2 + 2(x_2)^2)) \end{pmatrix},$$

both of which are graphed in Figure 5.2.1. At each state that is not a rest point, the angle between $F(x)$ and $V(x)$ is acute. At each Nash equilibrium, no vector that forms an acute angle with the payoff vector is a feasible direction of motion; thus, all Nash equilibria must be rest points under V . §

Exercise 5.2.3. Suppose that F is a two-strategy game, and let V_F and \hat{V}_F be Lipschitz continuous dynamics that satisfy condition (PC). Show that if neither dynamic is at rest at state $x \in X$, then $\hat{V}_F(x)$ is a positive multiple of $V_F(x)$. Conclude that if V_F and \hat{V}_F also satisfy condition (NS), then $\hat{V}_F(x) = k(x)V_F(x)$ for some positive function $k : X \rightarrow (0, \infty)$. In this case, the phase diagrams of V_F and \hat{V}_F are identical, and solutions to V_F and \hat{V}_F differ only by a change in speed (cf Exercise 5.4.10 below). §

Example 5.2.4. Consider the three-strategy coordination game

$$F(x) = \begin{pmatrix} F_1(x) \\ F_2(x) \\ F_2(x) \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_2 \end{pmatrix}.$$

Since payoffs are now vectors in \mathbf{R}^3 , they can no longer be drawn in a two-dimensional picture, so we draw the projected payoff vectors

$$\Phi F(x) = \left(I - \frac{1}{3} \mathbf{1}\mathbf{1}' \right) F(x) = \begin{pmatrix} x_1 - \frac{1}{3}(x_1 + 2x_2 + 3x_3) \\ 2x_2 - \frac{1}{3}(x_1 + 2x_2 + 3x_3) \\ 3x_3 - \frac{1}{3}(x_1 + 2x_2 + 3x_3) \end{pmatrix}$$

instead. Since dynamic V_F also takes values in TX , drawing the growth rate vectors $V_F(x)$ and the projected payoff vectors $\Phi F(x)$ is enough to evaluate property (PC) (cf Exercise 5.2.5). In Figure 5.2.2(i), we plot the projected payoffs ΦF and the replicator dynamic; in Figure 5.2.2(ii) we plot the projected payoffs ΦF and the BNN dynamic. In both cases, except when $V_F(x) = \mathbf{0}$, the angles between $V_F(x)$ and $\Phi F(x)$ are always acute. At each Nash equilibrium x , all directions of motion from x that form an acute angle with $\Phi F(x)$ are infeasible, and so both dynamics are at rest. §


Exercise 5.2.5. Let V_F be an evolutionary dynamic for the single population game F . Show that $\text{sgn}(V_F(x)'F(x)) = \text{sgn}(V_F(x)'\Phi F(x))$. Thus, to check that (PC) holds, it is enough to verify that it holds with respect to projected payoffs.

5.3 Families of Evolutionary Dynamics


In moving from equilibrium to evolutionary analysis of games, we are faced with a vast array of plausible assumptions about how agents make decisions. Indeed, an old critique of the use of bounded rationality in economic modeling holds that in departing from the traditional assumptions of rationality and equilibrium knowledge, one is left with too many possible alternatives. Since it is difficult to know which alternative is most apt in the application at hand, the predictions one obtains must therefore be viewed with suspicion.

This critique can be dispelled by proving results that are robust to a range of alternative specifications of agents' revision processes. When one's conclusions hold under any of number of protocols sharing a certain "family resemblance", one can argue that these conclusions are not artifacts of a particular choice of functional forms, but of more fundamental assumptions about how agents make decisions.

Following this principle, we proceed in this chapter and the next to define a variety of families of evolutionary dynamics. Each family is generated from a collection of revision protocols whose members all embody a particular style of decision making. We will find that many of our basic results hold for families of dynamics as a whole, suggesting that



(i) The replicator dynamic



(ii) The BNN dynamic

Figure 5.2.2: Condition (PC) in 123 Coordination.

Family	Leading example	(C)	$\leq (Dn)$	(PC)	(NS)
imitation	replicator	yes	yes	yes	no
excess payoff	BNN	yes	no	yes	yes
pairwise comparison	Smith	yes	yes	yes	yes
	best response	no	yes	yes ^a	yes ^a
perturbed best response	logit	yes	yes	no	no
	projection	no	no	yes	yes

^aThe best response dynamics satisfy versions of conditions (NS) and (PC) defined for differential inclusions.

Table 5.1: Families of evolutionary dynamics and their properties.

there is range of environments in which our results may be of practical relevance.

In Table 5.1, we list various families and examples of evolutionary dynamics to be introduced below, and we evaluate them in terms of our two conditions on revision protocols (continuity (C) and data requirements ((Dn) or weaker)) and our two conditions on aggregate dynamics (positive correlation (PC) and Nash stationarity (NS)). Let us briefly mention a few of the main ideas from the analyses to come.

- Imitative dynamics, including the replicator dynamic, satisfy all of the desiderata except for Nash stationarity (NS): these dynamics admit rest points that are not Nash equilibria.
- Excess payoff dynamics, including the BNN dynamic, satisfy all of our desiderata except data requirement ((Dn)): the revision protocols that generate these dynamics involve comparisons between the individual strategies' payoffs and the population's average payoff.
- By introducing revision protocols that only require pairwise comparisons of payoffs, we obtain a family of evolutionary dynamics that satisfy all four desiderata.
- The best response dynamic satisfies versions of all of the desiderata except continuity: its revision protocol depends discontinuously on payoffs.
- We can eliminate the discontinuity of the best response dynamic by introducing perturbations, but at the cost of violating the incentive conditions. In fact, choosing the level of perturbations involves a tradeoff between condition (C) and conditions (NS) and (PC): smaller perturbations reduce the degree of smoothing, while larger perturbations make the failures of the incentive conditions more severe.

- The projection dynamic minimizes the discrepancy at each state between the vector of payoffs and the vector representing the directions of motion. It satisfies both of incentive conditions, but neither of the limited information conditions. There are a variety of close connections between the projection dynamic and the replicator dynamic.

Figure 5.3.1 presents phase diagrams for the six basic dynamics in the standard Rock-Paper-Scissors game

$$F(x) = \begin{pmatrix} F_R(x) \\ F_P(x) \\ F_S(x) \end{pmatrix} = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \end{pmatrix} = \begin{pmatrix} x_S - x_P \\ x_R - x_S \\ x_P - x_R \end{pmatrix}.$$

The unique Nash equilibrium of RPS places equal mass on each strategy: $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. In the phase diagrams, colors represent speed of motion: within each diagram, motion is fastest in the red regions and slowest in the blue ones. In this example, the maximum speed under the replicator dynamic is $\frac{\sqrt{2}}{4} \approx .3536$, while the maximum speed under the other five dynamics is $\sqrt{2} \approx 1.4142$. Some remarks on the phase diagrams:

- The replicator and projection dynamics exhibit closed orbits around the Nash equilibrium. Under the other four dynamics, the Nash equilibrium is globally asymptotically stable.
- The replicator dynamic has rest points at the Nash equilibrium and at each of the pure states. Under the other dynamics, the only rest point is the Nash equilibrium.
- The phase diagram for the BNN dynamic can be divided into six regions. In the “odd” regions, exactly one strategy has above average payoffs, so the dynamic moves directly toward a pure state, just as under the best response dynamic. In the “even” regions, two strategies have above average payoffs; as these regions are traversed, the “target point” of the dynamic passes from one pure state to the next.
- Compared to those of the BNN dynamic, solutions of the Smith dynamic approach the Nash equilibrium at closer angles and at higher speeds.
- Under the best response dynamic, solution trajectories always aim directly toward the state representing the current best response. The trajectories are kinked whenever best responses change.

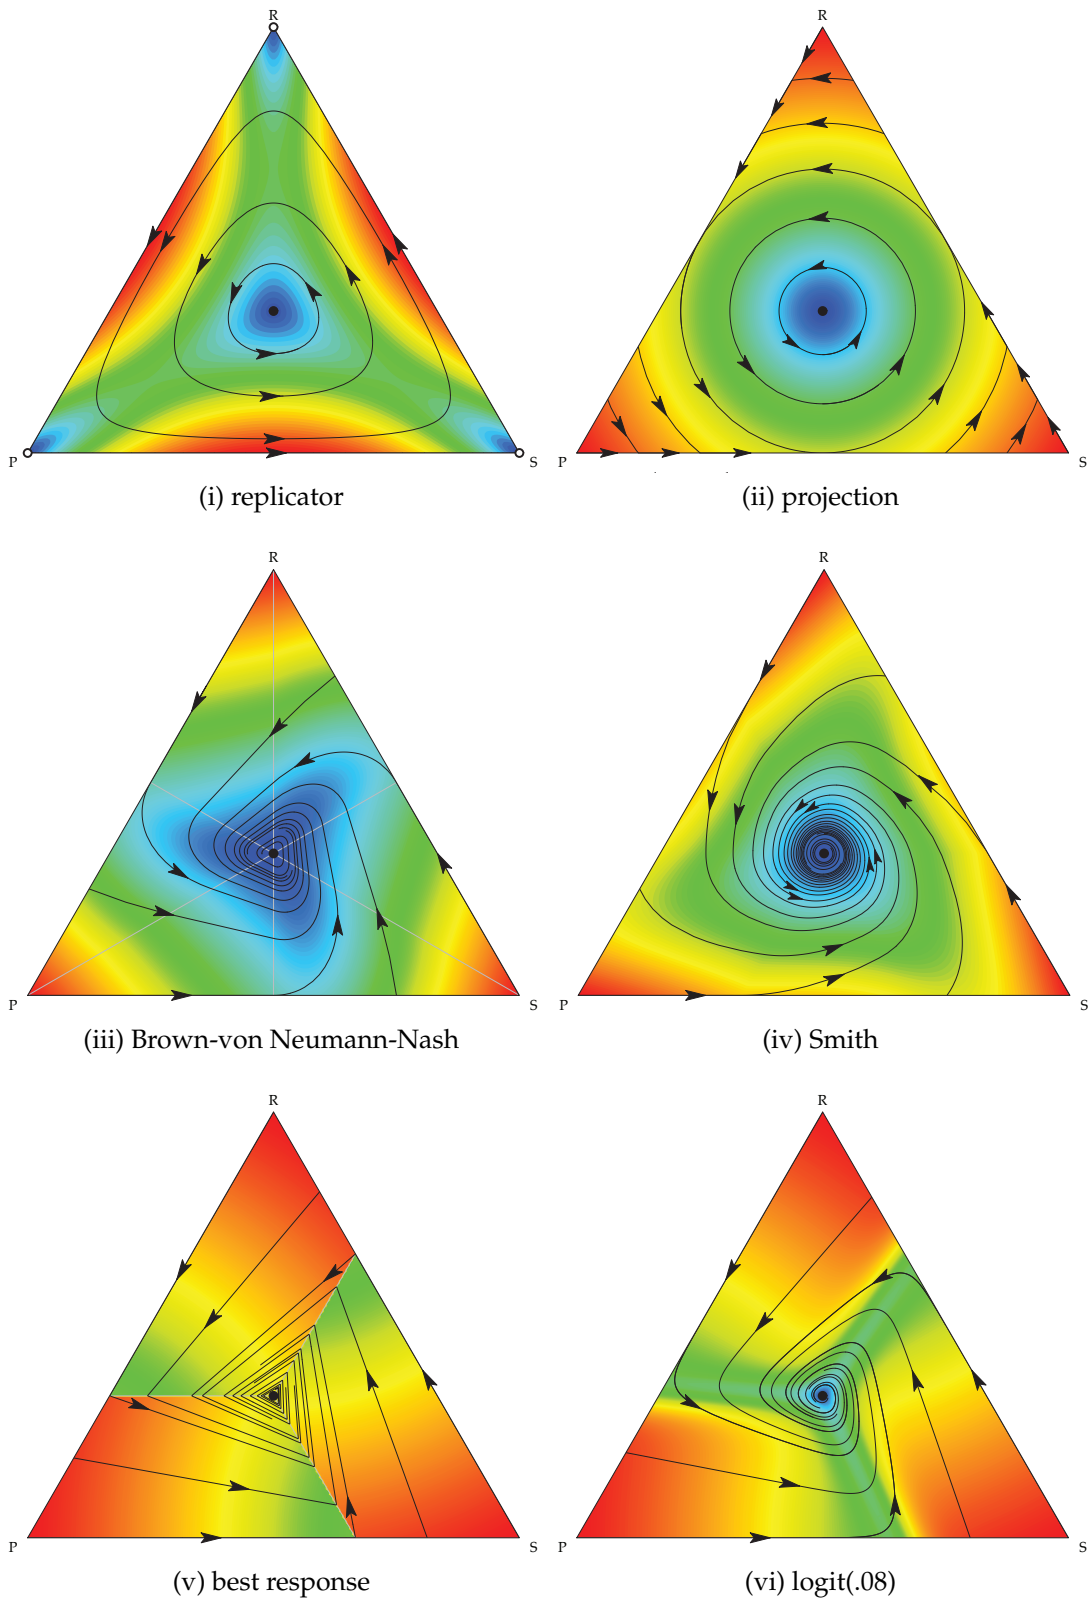


Figure 5.3.1: Six basic dynamics in the Rock-Paper-Scissors game.

- Unlike those of the best response dynamic, solutions trajectories of the logit dynamic are smooth. The directions of motion under the two dynamics are similar, except at states near the boundaries of the best response regions.
- Under the replicator dynamic, the boundary consists of three rest points and three *heteroclinic orbits* that connect distinct rest points. All told, the boundary forms what is known as a *heteroclinic cycle*.
- Under the projection dynamic, there is a unique forward solution from each initial condition, but backward solutions are not unique. For example, the outermost closed orbit (the inscribed circle) is reached in finite time by every solution trajectory that starts outside of it. In addition, there are solution trajectories that start in the interior of the state space and reach the boundary in finite time—an impossibility under any of the other dynamics

These and many further ideas will be developed below.

5.4 Imitative Dynamics

5.4.1 Definition

Imitative dynamics are based on revision protocols of the form

$$(5.5) \quad \rho_{ij}^p(\pi^p, x^p) = \hat{x}_j^p r_{ij}^p(\pi^p, x^p),$$

where $\hat{x}_j^p = x_j^p/m^p$ is the proportion of population p members playing strategy $j \in S^p$. We can interpret these protocols as follows: When an agent's clock rings, he randomly chooses an opponent from his population. If the agent is playing strategy $i \in S^p$ and the opponent strategy $j \in S^p$, then the agent imitates the opponent with probability proportional to the *conditional imitation rate* r_{ij}^p .

The revision protocol (5.5) generates a mean dynamic of the form

$$(5.6) \quad \begin{aligned} \dot{x}_i^p &= \sum_{k \in S^p} x_k^p \rho_{ki}^p(F^p(x), x^p) - x_i^p \sum_{k \in S^p} \rho_{ik}^p(F^p(x), x^p) \\ &= \sum_{k \in S^p} x_k^p \hat{x}_i^p r_{ki}^p(F^p(x), x^p) - x_i^p \sum_{k \in S^p} \hat{x}_k^p r_{ik}^p(F^p(x), x^p) \\ &= x_i^p \sum_{k \in S^p} \hat{x}_k^p (r_{ki}^p(F^p(x), x^p) - r_{ik}^p(F^p(x), x^p)). \end{aligned}$$

If the revision protocol satisfies the requirements below, the differential equation above defines an imitative dynamic.

Definition. Suppose that the conditional imitation rates r_{ij}^p are Lipschitz continuous, and that net conditional imitation rates are monotone:

$$(5.7) \quad \pi_j^p \geq \pi_i^p \Leftrightarrow r_{kj}^p(\pi^p, x^p) - r_{jk}^p(\pi^p, x^p) \geq r_{ki}^p(\pi^p, x^p) - r_{ik}^p(\pi^p, x^p) \text{ for all } i, j, k \in S^p \text{ and } p \in \mathcal{P}.$$

Then the map from population games $F \in \mathcal{F}$ to differential equations (5.6) is called an imitative dynamic.

Condition (5.7) says that whenever strategy $j \in S^p$ has a higher payoff than strategy $i \in S^p$, then the net rate of imitation from any strategy $k \in S^p$ to j exceeds the net rate of imitation from k to i . We illustrate this condition in the next subsection using a variety of examples; the condition's implications for aggregate behavior are developed thereafter.

Example 5.4.1. The replicator dynamic. The fundamental example of an imitative dynamic is the *replicator dynamic*, defined by

$$(R) \quad \dot{x}_i^p = x_i^p \hat{F}_i^p(x).$$

Under the replicator dynamic, the percentage growth rate of each strategy $i \in S^p$ currently in use equals its excess payoff $\hat{F}_i^p(x) = F_i^p(x) - \bar{F}^p(x)$; unused strategies remain so. We provide a variety of derivations of the replicator dynamic below. §

5.4.2 Examples

The examples to follow are expressed in the setting of a single, unit mass population, so that $\hat{x}_i = x_i$. They are easily recast for multipopulation settings.

Example 5.4.2. Imitation via pairwise comparisons. Suppose that $\rho_{ij}(\pi, x) = x_j \phi(\pi_j - \pi_i)$, where $\phi : \mathbf{R} \rightarrow \mathbf{R}_+$ equals 0 on $(-\infty, 0]$ and is increasing on $[0, \infty)$. In this case, an agent only imitates his randomly chosen opponent when the opponent's payoff is higher than the agent's own. Protocols of this form satisfy condition (5.7). If we write $\psi(d) = \phi(d) - \phi(-d)$, then we can express the corresponding mean dynamic as

$$\begin{aligned} \dot{x}_i &= x_i \sum_{k \in S} x_k \left(\phi(F_i(x) - F_k(x)) - \phi(F_k(x) - F_i(x)) \right) \\ &= x_i \sum_{k \in S} x_k \psi(F_i(x) - F_k(x)). \end{aligned}$$

Setting $\phi(d) = [d]_+$ gives us the *pairwise proportional imitation* protocol from Example 4.3.1. In this case $\psi(d) = d$, and the mean dynamic is the replicator dynamic (R). §

Exercise 5.4.3. Suppose we generalize the Example 5.4.2 by letting $\rho_{ij}(\pi, x) = x_j \phi_{ij}(\pi_j - \pi_i)$, where each function ϕ_{ij} equals 0 on $(-\infty, 0]$ and is increasing on $[0, \infty)$. Explain why the resulting mean dynamic need not satisfy condition (5.7), and so need not be an imitative dynamic. (For an interesting contrast, see Section 5.6.)

Example 5.4.4. Pure imitation driven by dissatisfaction. Suppose that $\rho_{ij}(\pi, x) = a(\pi_i) x_j$. Then when the clock of an i player rings, he abandons his current strategy with probability proportional to the *abandonment rate* $a(\pi_i)$; in such instances, he imitates a randomly chosen opponent. In this case, condition (5.7) requires that $a : \mathbf{R} \rightarrow \mathbf{R}_+$ be decreasing, and the mean dynamic becomes

$$(5.8) \quad \dot{x}_i = x_i \sum_{k \in S} x_k (a(F_k(x)) - a(F_i(x))) = x_i \left(\sum_{k \in S} x_k a(F_k(x)) - a(F_i(x)) \right).$$

If abandonment rates take the linear form $a(\pi_i) = K - \pi_i$ (where K is large enough), then (5.8) is again the replicator dynamic (R). §

Example 5.4.5. Imitation of success. Suppose $\rho_{ij}(\pi, x) = x_j c(\pi_j)$. Then when an agent's clock rings, he picks an opponent at random; if the opponent is playing strategy j , the player imitates him with probability proportional to the *copying rate* $c(\pi_j)$. In this case, condition (5.7) requires that $c : \mathbf{R} \rightarrow \mathbf{R}_+$ be increasing, and the mean dynamic becomes

$$(5.9) \quad \dot{x}_i = x_i \sum_{k \in S} x_k (c(F_i(x)) - c(F_k(x))) = x_i \left(c(F_i(x)) - \sum_{k \in S} x_k c(F_k(x)) \right).$$

Since ρ is a target protocol (i.e., since $\rho_{ij} \equiv \tau_j$), the mean dynamic (5.9) is actually a target dynamic:

$$\dot{x}_i = \begin{cases} \sum_{k \in S} x_k c(F_k(x)) \left(\frac{x_i c(F_i(x))}{\sum_{k \in S} x_k c(F_k(x))} - x_i \right) & \text{if } x_j c(F_j(x)) \neq 0 \text{ for some } j \in S, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

If copying rates are of the linear form $c(\pi_j) = \pi_j + K$ (for K large enough), then (5.9) is once again the replicator dynamic (R). If in addition payoffs are nonnegative and average

payoffs are positive, we can choose $c(\pi_j) = \pi_j$, so that (5.9) becomes

$$(5.10) \quad \dot{x}_i = \bar{F}(x) \left(\frac{x_i F_i(x)}{\bar{F}(x)} - x_i \right).$$

Here, the target state is proportional to the vector of popularity-weighted payoffs $x_i F_i(x)$, with the rate of motion toward this state governed by average payoffs $\bar{F}(x)$. §

Exercise 5.4.6. Why is the restriction on payoffs needed to obtain equation (5.10)?

Example 5.4.7. Imitation of success with repeated sampling. Suppose that

$$(5.11) \quad \rho_{ij}(\pi, x) = \frac{x_j w(\pi_j)}{\sum_{k \in S} x_k w(\pi_k)},$$

where $\sum_{k \in S} x_k w(\pi_k) > 0$. Here, when an agent's clock rings he chooses an opponent at random. If the opponent is playing strategy j , the agent imitates him with probability proportional to the *copying weight* $w(\pi_j)$. If the agent does not imitate this opponent, he draws a new opponent at random and repeats the procedure. In this case, condition (5.7) requires that $w : \mathbf{R} \rightarrow \mathbf{R}_+$ be increasing. Since ρ is an exact target protocol (i.e., since $\rho_{ij} \equiv \sigma_j$ and $\sum_{j \in S} \sigma_j \equiv 1$), it induces the exact target dynamic

$$(5.12) \quad \dot{x}_i = \frac{x_i w(F_i(x))}{\sum_{k \in S} x_k w(F_k(x))} - x_i. \quad \S$$

We conclude with two important instances of repeated sampling.

Example 5.4.8. The Maynard Smith replicator dynamic. If payoffs are nonnegative and average payoffs are positive, we can let copying weights equal payoffs: $w(\pi_j) = \pi_j$. The resulting exact target dynamic,

$$(5.13) \quad \dot{x}_i = \frac{x_i F_i(x)}{\bar{F}(x)} - x_i = \frac{x_i \hat{F}_i(x)}{\bar{F}(x)},$$

is known as the *Maynard Smith replicator dynamic*.

Example 5.4.5 showed that under the same assumptions on payoffs, the replicator dynamic takes the form (5.10). The Maynard Smith replicator dynamic (5.13) differs from (5.10) only in that the target state is approached at a unit rate rather than at a rate determined by average payoffs; thus, motion under (5.10) is relatively fast when average payoffs are relatively high. Comparing the protocol here to the one from Example 5.4.5 reveals the source of the difference in speeds: under repeated sampling, the overall payoff

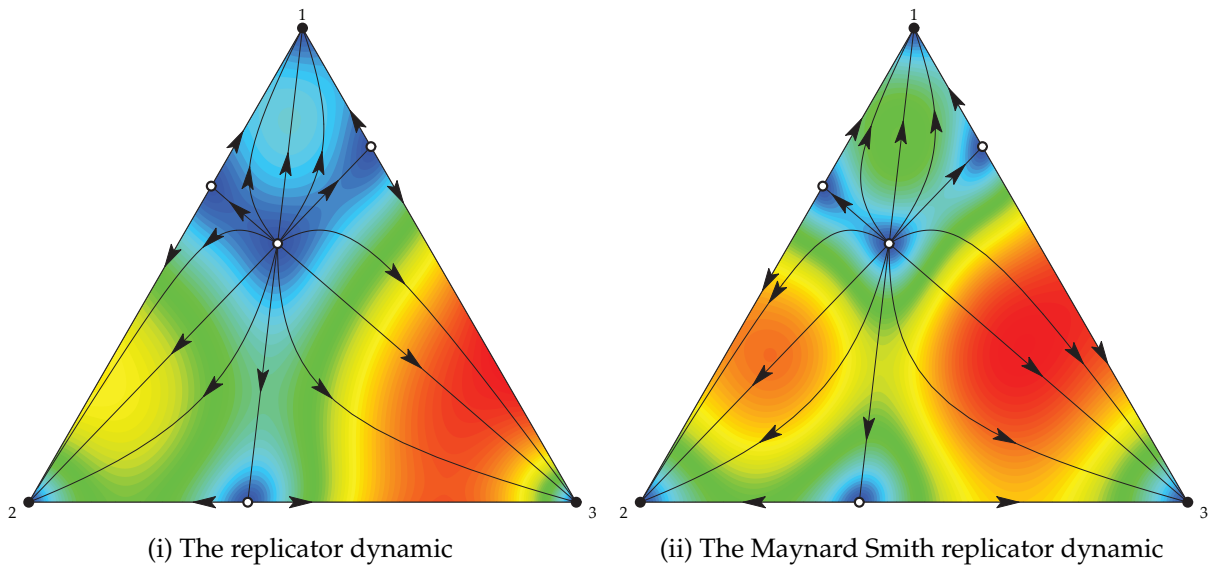


Figure 5.4.1: Two imitative dynamics in 123 Coordination.

level has little influence on the probability that a revising agent winds up switching strategies.

In the single population setting, the phase diagrams of (5.10) and (5.13) are identical, and the dynamics only differ in terms of the speed at which solution trajectories are traversed (cf Exercise 5.4.10). We illustrate this in Figure 5.4.1, which presents phase diagrams for the two dynamics in 123 Coordination.

When there are multiple populations, the fact that average payoffs differ across populations implies that the phase diagrams of (5.10) and (5.13) no longer coincide. This is illustrated in Figure 5.4.2, which presents phase diagrams for this Matching Pennies game:

	h	t
H	2, 1	1, 2
T	1, 2	2, 1

While interior solutions of (5.10) form closed orbits around the unique Nash equilibrium $x^* = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$, interior solutions of (5.13) converge to x^* . §

In the biology literature, the stochastic evolutionary process generated by the revision protocol in the previous example is known as a *frequency-dependent Moran process*—see the Notes for further discussion.

Example 5.4.9. The i-logit dynamic. If the copying weights $w(\pi_j) = \exp(\eta^{-1}\pi_j)$ are exponential functions of payoffs, the exact target dynamic (5.12) becomes the *i-logit dynamic* with

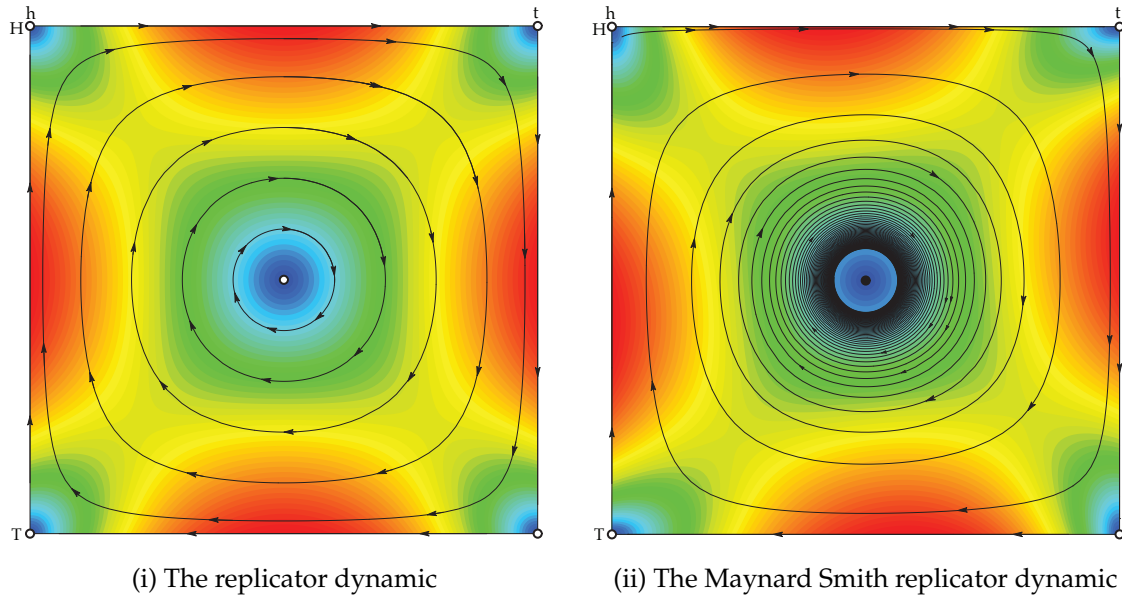


Figure 5.4.2: Two imitative dynamics in Matching Pennies.

noise level $\eta > 0$.

$$\dot{x}_i = \frac{x_i \exp(\eta^{-1} F_i(x))}{\sum_{k \in S} x_k \exp(\eta^{-1} F_k(x))} - x_i.$$

Here, the i th component of the target state is proportional both to the mass of agents playing strategy i and to an exponential function of strategy i 's payoff. If η is small, and x is not too close to the boundary of X or of any best response region, then the target state is close to $e_{b(x)}$, the vertex of X corresponding to the current best response. Therefore, in most games, the i-logit dynamic with small η approximates the best response dynamic $\dot{x} \in B(x) - x$ on much of the interior of the simplex. We illustrate this in Figure 5.4.3, which presents four i-logit dynamics (with $\eta = .5, .1, .05$, and $.01$) and the best response dynamic for the anticoordination game

$$F(x) = Ax = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} -x_1 \\ -x_2 \\ -x_3 \end{pmatrix}.$$

Examining the figure again, we see that regardless of the value of η , the i-logit dynamic does not resemble best response dynamic on the boundary of the simplex. Solutions of the best response dynamic starting on the boundary of the simplex immediately enter its interior and move rapidly toward the center of the simplex. For their part, solutions of the

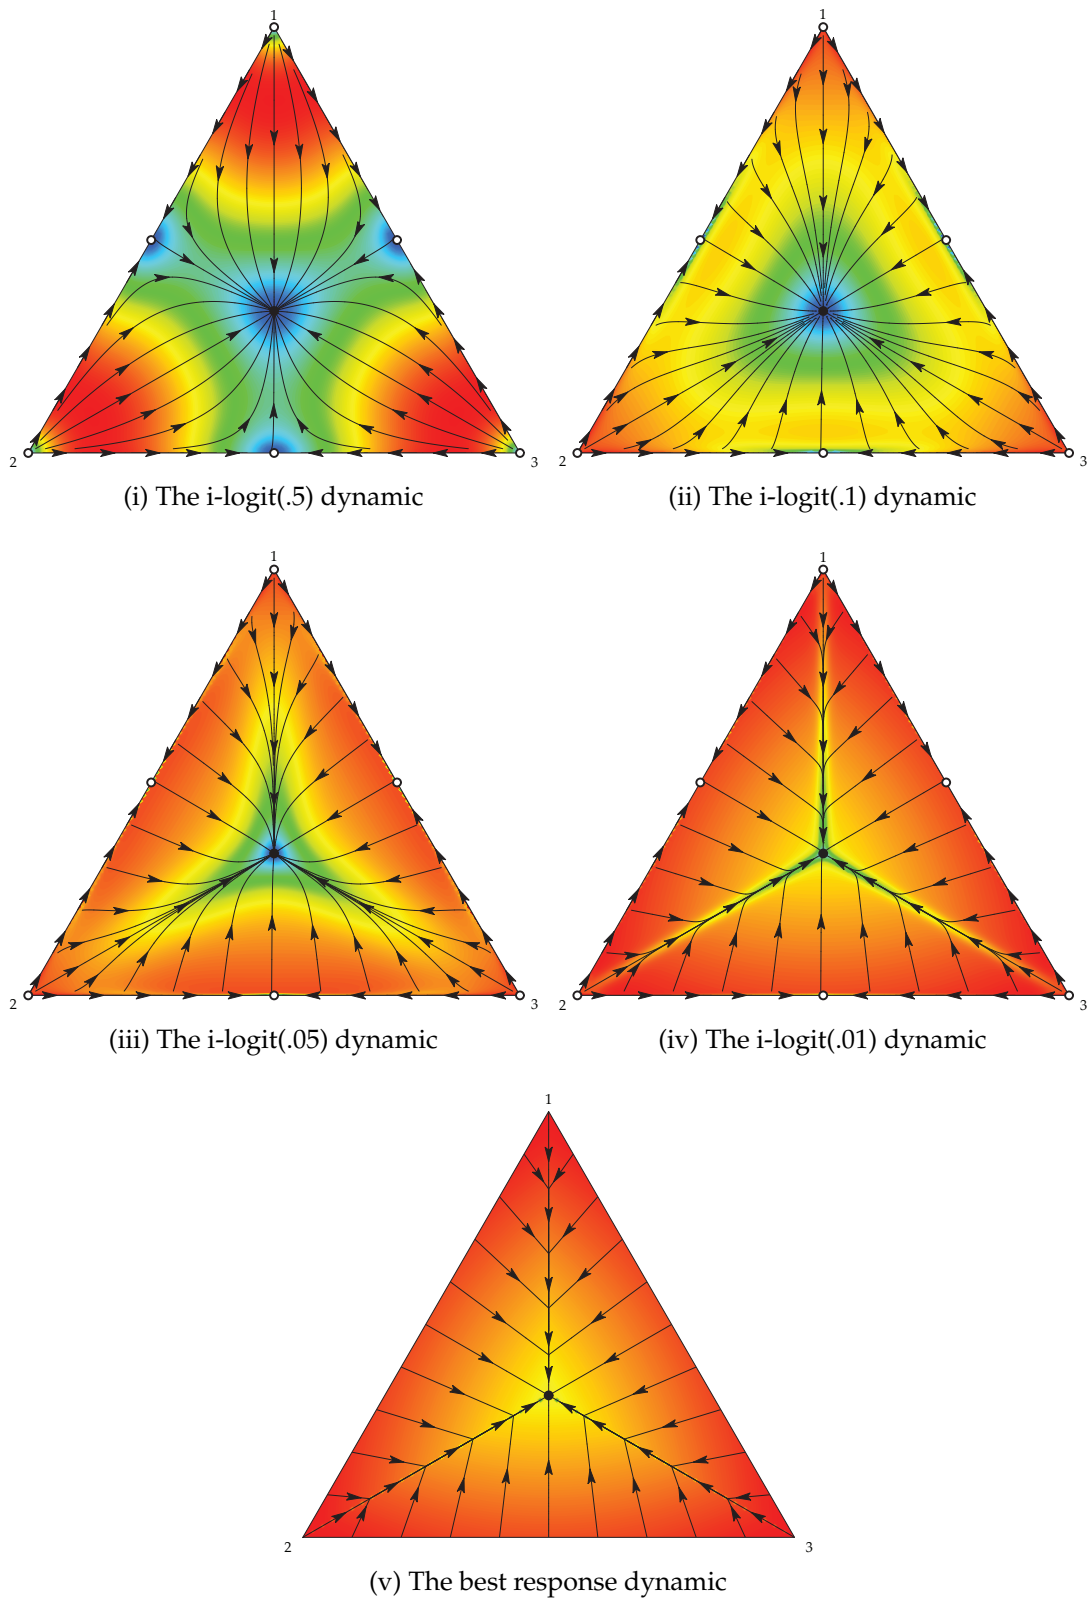


Figure 5.4.3: i-logit and best response dynamics in Anticoordination.

i-logit dynamic starting on the boundary remain there forever; by continuity, solutions starting close enough to the boundary initially move away from it quite slowly. These are general properties of imitative dynamics: see Section 5.4.4 below. §

Exercise 5.4.10. Changes of speed and reparameterizations of time. Let $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ be a Lipschitz continuous vector field and let $k : \mathbf{R}^n \rightarrow (0, \infty)$ be a positive Lipschitz continuous function. Let $\{x_t\}$ be a solution to $\dot{x} = V(x)$ with initial condition ξ , and let $\{y_t\}$ be a solution to $\dot{x} = k(x)V(x)$, also with initial condition ξ . Show that $y_t = x_{I(t)}$, where $I(t) = \int_0^t k(y_s) ds$.

5.4.3 Biological Derivations of the Replicator Dynamic

While we have derived the replicator dynamic from models of imitation, its origins lie in mathematical biology, where it arises from models of intra- and inter-species competition. The next two exercises, which are set in a single population, consider the replicator dynamic from this point of view.

Exercise 5.4.11. In the basic game theoretic model of natural selection within a single animal species, each strategy $i \in S$ represents a behavioral type. The value of $F_i(x)$ represents the (*reproductive*) *fitness* of type i when the current proportions of types are described by $x \in \text{int}(X)$. In particular, if we let $y_i \in (0, \infty)$ represent the (absolute) number animals of type i in the population, then the evolution of the population is described by

$$(5.14) \quad \dot{y}_i = y_i F_i(x), \quad \text{where } x_i = \frac{y_i}{\sum_{j \in S} y_j}.$$

Show that under equation (5.14), the vector x describing the proportions of animals of each of each type evolves according to the replicator equation (R).

Exercise 5.4.12. The Lotka-Volterra equation. The *Lotka-Volterra equation* is a fundamental model of biological competition among members of multiple species. When there are $n - 1$ species, the equation takes the form

$$(5.15) \quad \dot{y}_k = y_k (b_k + (My)_k), \quad k \in \{1, \dots, n - 1\},$$

where b_k is the baseline growth rate for species k , and the interaction matrix $M \in \mathbf{R}^{(n-1) \times (n-1)}$ governs cross-species effects. Show that after the change of variable

$$x_i = \frac{y_i}{1 + \sum_{l=1}^{n-1} y_l} \quad \text{and} \quad x_n = \frac{1}{1 + \sum_{l=1}^{n-1} y_l},$$

the $n - 1$ dimensional Lotka-Volterra equation (5.15) is equivalent up to a change of speed (cf Exercise 5.4.10) to the n strategy replicator dynamic

$$\dot{x}_i = x_i((Ax)_i - x'Ax), \quad i \in \{1, \dots, n\},$$

where the payoff matrix $A \in \mathbf{R}^{n \times n}$ is related to $M \in \mathbf{R}^{(n-1) \times (n-1)}$ and $b \in \mathbf{R}^{n-1}$ by the $\mathbf{R}^{(n-1) \times n}$ matrix equation

$$\begin{pmatrix} M & b \end{pmatrix} = \begin{pmatrix} I & (-\mathbf{1}) \end{pmatrix} A.$$

If M and b are given, this equation determines A up to an additive constant in each column. Thus, A can always be chosen so that either the elements of its last row or the elements of its diagonal are all 0.

5.4.4 Extinction and Invariance

We now derive properties shared by all imitative dynamics. First of all, it follows immediately from equation (5.6) that all imitative dynamics satisfy *extinction*: if a strategy is unused, its growth rate is zero.

$$(5.16) \quad \text{If } x_i^p = 0, \text{ then } V_i^p(x) = 0.$$

Extinction implies that the growth rate vectors $V(x)$ are always tangent to the boundaries of X : formally, $V(x)$ is not only in $TX(x)$, but also in $-TX(x)$ (cf Observation 4.4.3). Thus, since imitative dynamics are Lipschitz continuous, it follows from Theorem 4.A.8 in Chapter 4 that solutions to imitative dynamics exist for all positive *and negative* times.

Proposition 5.4.13 (Forward and backward invariance). *Let $\dot{x} = V_F(x)$ be an imitative dynamic. Then for each initial condition $\xi \in X$, this dynamic admits a unique solution trajectory in $\mathcal{T}_{(-\infty, \infty)} = \{x : (-\infty, \infty) \rightarrow X : x \text{ is continuous}\}$.*

Extinction also implies a second invariance property: if $\{x_t\}$ is a solution trajectory of an imitative dynamic, then the support of x_t is independent of t . Uniqueness of solution trajectories, which is implied by the Lipschitz continuity of the dynamic, is an essential ingredient of the proof of this result.

Theorem 5.4.14 (Support invariance). *If $\{x_t\}$ is a solution trajectory of an imitative dynamic, then the sign of component $(x_t)_i^p$ is independent of $t \in (-\infty, \infty)$.*

Proof. Let $\{x_t\}$ be a solution to the imitative dynamic $\dot{x} = V(x)$, and suppose that $x_0 = \xi$. Suppose that $\xi_i^p = 0$; we want to show that $(x_t)_i^p = 0$ for all $t \in (-\infty, \infty)$. To accomplish this, we define a new vector field $\hat{V} : X \rightarrow \mathbf{R}^n$ as follows:

$$\hat{V}_j^q(x) = \begin{cases} 0 & \text{if } j = i \text{ and } q = p, \\ V_j^q(x) & \text{otherwise.} \end{cases}$$

If $\{\hat{x}_t\} \subset X$ is the unique solution to $\dot{x} = \hat{V}(x)$ with $\hat{x}_0 = \xi$, then $(\hat{x}_t)_i^p = 0$ for all t . But V and \hat{V} are identical whenever $x_i^p = 0$ by extinction (5.16); therefore, $\{\hat{x}_t\}$ is also a solution to $\dot{x} = V(x)$. Since solutions to $\dot{x} = V(x)$ are unique, it must be that $\{\hat{x}_t\} = \{x_t\}$, and hence that $(x_t)_i^p = 0$ for all t .

Now suppose that $\xi_i^p > 0$. If $x_t = \chi$ satisfied $\chi_i^p = 0$, then the preceding analysis would imply that there are two distinct solutions to $\dot{x} = V(x)$ with $x_t = \chi$, one that is contained in the boundary of X and one that is not. As this would contradict uniqueness of solutions, we conclude $(x_t)_i^p > 0$ at all times t . ■

All of the phase diagrams presented in this section illustrate the face invariance property. The next example points out one of its more subtle consequences.

Example 5.4.15. Figure 5.4.4 presents the phase diagram of the replicator dynamic for a game with a strictly dominant strategy: for all $x \in X$, $F_1(x) = 1$ and $F_2(x) = F_3(x) = 0$. There are two connected components of rest points: one consisting solely of the unique Nash equilibrium e_1 , and the other containing those states at which strategy 1 is unused. Clearly, the latter component is unstable, as all nearby solution trajectories lead away from it and toward the Nash equilibrium. But as the coloring of the figure indicates, the speed of motion away from the unstable component is very slow: if a small behavior disturbance pushes the state off of the component, it may take a long time before the stable equilibrium is reached. §

5.4.5 Monotone Percentage Growth Rates and Positive Correlation

We now turn to monotonicity properties of imitative dynamics. All dynamics of form (5.6) can be expressed as

$$(5.17) \quad \dot{x}_i^p = V_i^p(x) = x_i^p G_i^p(x), \text{ where } G_i^p(x) = \sum_{k \in S^p} \hat{x}_k^p \left(r_{ki}^p(F^p(x), x^p) - r_{ik}^p(F^p(x), x^p) \right).$$




Figure 5.4.4: The replicator dynamic in a game with a strictly dominant strategy.

If strategy $i \in S^p$ is in use, then $G_i^p(x) = V_i^p(x)/x_i^p$ represents the *percentage growth rate* of the number of agents using this strategy.

Observation 5.4.16 notes that under every imitative dynamic (as defined in Section 5.4.1), strategies' percentage growth rates are ordered by their payoffs.

Observation 5.4.16. *All imitative dynamics exhibit monotone percentage growth rates:*

$$(5.18) \quad G_i^p(x) \geq G_j^p(x) \text{ if and only if } F_i^p(x) \geq F_j^p(x).$$

This observation is immediate from condition (5.7), which defines imitative dynamics.

Condition (5.18) is a strong restriction on strategies' *percentage* growth rates. We now show that it implies our basic payoff monotonicity condition, which imposes a weak restriction on strategies' *absolute* growth rates.

Theorem 5.4.17. *All imitative dynamics satisfy positive correlation (PC).*

Proof. Let x be a social state at which $V^p(x) \neq \mathbf{0}$; we need to show that $V^p(x)'F^p(x) > 0$. To do so, we define

$$S_+^p(x) = \{i \in S^p : V_i^p(x) > 0\} \quad \text{and} \quad S_-^p(x) = \{j \in S^p : V_j^p(x) < 0\}$$

to be the sets of population p strategies with positive and negative absolute growth rates,

respectively. By extinction (5.16), these sets are contained in the support of x^p . It follows that

$$S_+^p(x) = \{i \in S^p : x_i^p > 0 \text{ and } \frac{V_i^p(x)}{x_i^p} > 0\} \text{ and } S_-^p(x) = \{j \in S^p : x_j^p > 0 \text{ and } \frac{V_j^p(x)}{x_j^p} < 0\}.$$

Since $V(x) \in TX$, we know from Observation 4.4.3 that

$$\sum_{k \in S_+^p(x)} V_k^p(x) = - \sum_{k \in S_-^p(x)} V_k^p(x),$$

and since $V^p(x) \neq \mathbf{0}$, these expressions are positive. Therefore, Observation 5.4.16 enables us to conclude that

$$\begin{aligned} V^p(x)'F^p(x) &= \sum_{k \in S_+^p(x)} V_k^p(x) F_k^p(x) + \sum_{k \in S_-^p(x)} V_k^p(x) F_k^p(x) \\ &\geq \min_{i \in S_+^p(x)} F_i^p(x) \sum_{k \in S_+^p(x)} V_k^p(x) + \max_{j \in S_-^p(x)} F_j^p(x) \sum_{k \in S_-^p(x)} V_k^p(x) \\ &= \left(\min_{i \in S_+^p(x)} F_i^p(x) - \max_{j \in S_-^p(x)} F_j^p(x) \right) \sum_{k \in S_+^p(x)} V_k^p(x) > 0. \blacksquare \end{aligned}$$

We conclude this section by considering two other monotonicity conditions that appear in the literature.

Exercise 5.4.18. In the single population setting, an imitative dynamic (5.17) has *aggregate monotone percentage growth rates* if

$$(5.19) \quad \hat{y}'G(x) \geq y'G(x) \text{ if and only if } \hat{y}'F(x) \geq y'F(x)$$

for all population states $x \in X$ and mixed strategies $\hat{y}, y \in \Delta$.

- (i) Show that any imitative dynamic satisfying condition (5.19) is equivalent to the replicator dynamic up to a reparameterization of time (see Exercise 5.4.10). (Hint: Use Proposition 3.B.6 to show that condition (5.19) implies that $\Phi G(x) = c(x) \Phi F(x)$ for some $c(x) > 0$. Then use the fact that $G(x)'x = 0$ (why?) to conclude that $\dot{x}_i = c(x) x_i \hat{F}_i(x)$.)
- (ii) If a multipopulation imitative dynamic satisfies the natural analogue of condition (5.19), what can we conclude about the dynamic?

Exercise 5.4.19. An dynamic of form (5.17) has *sign-preserving percentage growth rates* if

$$(5.20) \quad \text{sgn}(G_i^p(x)) = \text{sgn}(\hat{F}_i^p(x)).$$

Show that any such dynamic satisfies positive correlation (PC). (Note that dynamics satisfying condition (5.20) need not satisfy condition (5.7), and so need not be imitative dynamics as we have defined them here. We do not know of an intuitive restriction on revision protocols that leads to condition (5.20).)

5.4.6 Rest Points and Restricted Equilibria

Since all imitative dynamics satisfy positive correlation (PC), Proposition 5.2.1 tells us that their rest points include all Nash equilibria of the underlying game F . On the other hand, face invariance tells us that non-Nash rest points can exist—for instance, while pure states in X are not always Nash equilibria of F , they are necessarily rest points of V_F .

To characterize the set of rest points, we first recall the definition of Nash equilibrium:

$$NE(F) = \{x \in X : x_i^p > 0 \Rightarrow F_i^p(x) = \max_{j \in S^p} F_j^p(x)\}.$$

Bearing this definition in mind, we define the set of *restricted equilibria* of F by

$$RE(F) = \{x \in X : x_i^p > 0 \Rightarrow F_i^p(x) = \max_{j \in S^p : x_j > 0} F_j^p(x)\}.$$

In words, x is a restricted equilibrium of F if it is a Nash equilibrium of a restricted version of F in which only strategies in the support of x can be played.

Exercise 5.4.20. Alternate definitions of restricted equilibrium.

- (i) Show that $x \in RE(F)$ if and only if within each population p , all strategies in the support of x^p achieve the same payoff: $RE(F) = \{x \in X : x_i^p > 0 \Rightarrow F_i^p(x) = \pi^p\}$.
- (ii) We can also offer a geometric definition of restricted equilibrium. Let $X_{[\hat{x}]}$ be the set of social states whose supports are contained in the support of \hat{x} : $X_{[\hat{x}]} = \{x \in X : \hat{x}_i^p = 0 \Rightarrow x_i^p = 0\}$. Show that $x \in RE(F)$ if and only if the payoff vector $F(x)$ is contained in the normal cone of $X_{[x]}$ at x : $RE(F) = \{x \in X : F(x) \in NX_{[x]}(x)\}$.

Because imitative dynamics exhibit face invariance, strategies that are initially unused are never subsequently chosen. This suggests a link between rest points of imitative dynamics and the restricted equilibria of the underlying game that is established in the following theorem.

Theorem 5.4.21. *If $\dot{x} = V_F(x)$ is an imitative dynamic, then $RP(V_F) = RE(F)$.*

$$\begin{aligned}
 \text{Proof. } x \in RP(V) &\Leftrightarrow V_i^p(x) = 0 \text{ for all } i \in S^p, p \in \mathcal{P} \\
 &\Leftrightarrow \frac{V_i^p(x)}{x_i^p} = 0 \text{ whenever } x_i^p > 0, p \in \mathcal{P} && \text{(by (5.16))} \\
 &\Leftrightarrow F_i^p(x) = \pi^p \text{ whenever } x_i^p > 0, p \in \mathcal{P} && \text{(by (5.18))} \\
 &\Leftrightarrow x \in RE(F). \blacksquare
 \end{aligned}$$

While there are rest points of imitative dynamics that are not Nash equilibria, we will see that non-Nash rest points are locally unstable—see Chapter 8. On the other hand, as Example 5.4.15 illustrates, the speed of motion away from these unstable rest points is initially rather slow.

- Exercise 5.4.22.* (i) Suppose that the payoffs of one population game are the negation of the payoffs of another. What is the relationship between the replicator dynamics of the two games?
- (ii) Give an example of a three-strategy game whose Nash equilibrium is unique and whose replicator dynamic admits seven rest points.

5.5 Excess Payoff Dynamics

In the next two subsections we consider revision protocols that are not based on imitation, but on the direct selection of candidate strategies. Under such protocols, good unused strategies will be discovered and chosen, raising the possibility that the dynamics satisfy Nash stationarity (NS).

5.5.1 Definition and Interpretation

In some settings, particularly those in which information about population aggregates is provided by a central planner, agents may know their population's current average payoff. Suppose that each agent's choices are based on comparisons between the various strategies' current payoffs and the population's average payoff, and that these choices do not condition on the agent's current strategy. Then the agents' choice procedure can be described using a target protocol of the form

$$\rho_{ij}^p(\pi^p, x^p) = \tau_j^p(\hat{\pi}^p),$$

where $\hat{\pi}_i^p = \pi_i^p(x) - \frac{1}{m^p}(x^p)' \pi^p$ represents the excess payoff to strategy $i \in S^p$. Such a protocol generates the target dynamic

$$(5.21) \quad \begin{aligned} \dot{x}_i^p &= m^p \tau_i^p(\hat{F}^p(x)) - x_i^p \sum_{j \in S^p} \tau_j^p(\hat{F}^p(x)) \\ &= \begin{cases} \sum_{j \in S} \tau_j^p(\hat{F}^p(x)) \left(m^p \frac{\tau_i^p(\hat{F}^p(x))}{\sum_{j \in S} \tau_j^p(\hat{F}^p(x))} - x^p \right) & \text{if } \tau^p(\hat{F}^p(x)) \neq \mathbf{0}, \\ \mathbf{0} & \text{otherwise.} \end{cases} \end{aligned}$$

To obtain our new class of dynamics, we introduce a monotonicity condition for the protocol τ . To do so, let us first observe that the excess payoff vector $\hat{F}^p(x)$ cannot lie in the interior of the negative orthant $\mathbf{R}_-^{n^p}$: for this to happen, every strategy would have to earn a below average payoff. Bearing this in mind, we can let the domain of the function τ^p be the set $\mathbf{R}_*^{n^p} = \mathbf{R}^{n^p} - \text{int}(\mathbf{R}_-^{n^p})$. Note that $\text{int}(\mathbf{R}_*^{n^p}) = \mathbf{R}^{n^p} - \mathbf{R}_-^{n^p}$ is the set of excess payoff vectors under which at least one strategy earns an above average payoff, while $\text{bd}(\mathbf{R}_*^{n^p}) = \text{bd}(\mathbf{R}_-^{n^p})$ is the set of excess payoff vectors under which no strategy earns an above average payoff. The definition of $\mathbf{R}_*^{n^p}$ will prove quite convenient in the analyses to follow: see especially Lemma 5.5.4 below.

With this notation in hand, we can define our family of dynamics.

Definition. Suppose the protocols $\tau^p : \mathbf{R}_*^{n^p} \rightarrow \mathbf{R}_+^{n^p}$ are Lipschitz continuous and satisfy acuteness:

$$(5.22) \quad \text{If } \hat{\pi}^p \in \text{int}(\mathbf{R}_*^{n^p}), \text{ then } \tau^p(\hat{\pi}^p)' \hat{\pi}^p > 0.$$

Then the map from population games $F \in \mathcal{F}$ to differential equations (5.21) is called an excess payoff dynamic.

How should one interpret condition (5.22)? If the excess payoff vector $\hat{\pi}^p$ has a positive component, this condition implies that

$$\sigma^p(\hat{\pi}^p) = \frac{1}{\sum_{i \in S} \tau_i^p(\hat{\pi}^p)} \tau^p(\hat{\pi}^p) \in \Delta^p,$$

the probability vector that defines the target state, is well defined. Acuteness requires that if we pick a component of the excess payoff vector $\hat{\pi}^p$ at random according to this probability vector, then the expected value of this randomly chosen component is positive. Put differently, acuteness asks that *on average*, revising agents switch to strategies with above average payoffs.

There is at least a formal resemblance between acuteness (5.22), which constrains the target protocol τ , and positive correlation (PC), which constrains the dynamic V_F . We will use this resemblance to our best advantage in Lemma 5.5.6 below.

Example 5.5.1. The BNN dynamic. Suppose that conditional switch rate to strategy $i \in S^p$ is given by the positive part of strategy i 's excess payoffs: $\tau_i^p(\hat{\pi}^p) = [\hat{\pi}_i^p]_+$. The resulting mean dynamic,

$$(BNN) \quad \dot{x}_i^p = m^p [\hat{F}_i^p(x)]_+ - x_i^p \sum_{j \in S^p} [\hat{F}_j^p(x)]_+,$$

is called the *Brown-von Neumann-Nash (BNN) dynamic*. §

Exercise 5.5.2. k-BNN dynamics. The *k-BNN dynamic* is generated by the revision protocol $\tau_i^p(\hat{\pi}^p) = [\hat{\pi}_i^p]^k$, where $k \geq 1$. Argue informally that if k is large, then at “typical” states, the direction of motion under the k -BNN dynamic is close to that under the best response dynamic, $\dot{x}^p \in m^p B^p(x) - x^p$ (see Chapter 6), but that the speed of motion is not.

5.5.2 Incentives and Aggregate Behavior

Our goal in this section is to show that every excess payoff dynamic satisfies our two incentive properties.

Theorem 5.5.3. *Every excess payoff dynamic $\dot{x} = V_F(x)$ satisfies positive correlation (PC) and Nash stationarity (NS).*

We prove this result under the assumption that τ^p satisfies *sign preservation*:

$$(5.23) \quad \text{sgn}(\tau_i^p(\hat{\pi}^p)) = \text{sgn}([\hat{\pi}_i^p]_+).$$

A proof using only acuteness is outlined in Exercise 5.5.8 below. We also focus on the single population case; the proof of the multipopulation case is a simple extension of the argument below.

The proof follows immediately from the following three lemmas.

Lemma 5.5.4. $\hat{F}(x) \in \text{bd}(\mathbf{R}_*^n)$ if and only if $x \in NE(F)$.

Proof. $\hat{F}(x) \in \text{bd}(\mathbf{R}_*^n) \Leftrightarrow F_i(x) \leq \sum_{k \in S} x_k F_k(x)$ for all $i \in S$
 \Leftrightarrow there exists a $c \in \mathbf{R}$ such that $F_i(x) \leq c$ for all $i \in S$,
with $F_j(x) = c$ whenever $x_j > 0$

$$\begin{aligned} &\Leftrightarrow F_j(x) = \max_{k \in S} F_k(x) \text{ whenever } x_j > 0 \\ &\Leftrightarrow x \in NE(F). \blacksquare \end{aligned}$$

Lemma 5.5.5. *If $\hat{F}(x) \in \text{bd}(\mathbf{R}_*^n)$, then $V_F(x) = \mathbf{0}$.*

Proof. Immediate from sign preservation (5.23). \blacksquare

Lemma 5.5.6. *If $\hat{F}(x) \in \text{int}(\mathbf{R}_*^n)$, then $V_F(x)'F(x) > 0$.*

Proof. Recall that $\hat{F}(x) = F(x) - \mathbf{1}\bar{F}(x)$ and that $V_F(x) = \tau(\hat{F}(x)) - \mathbf{1}'\tau(\hat{F}(x))x$. The first definition implies that x and $\hat{F}(x)$ are always orthogonal:

$$x'\hat{F}(x) = x'(F(x) - \mathbf{1}\bar{F}(x)) = x'F(x) - \bar{F}(x) = 0.$$

Combining this with the second definition, we see that if $\hat{F}(x) \in \text{int}(\mathbf{R}_*^n)$, then

$$\begin{aligned} V_F(x)'F(x) &= V_F(x)'(\hat{F}(x) + \mathbf{1}\bar{F}(x)) \\ &= V_F(x)'\hat{F}(x) && \text{since } V_F(x) \in TX \\ &= (\tau(\hat{F}(x)) - \mathbf{1}'\tau(\hat{F}(x))x)'\hat{F}(x) \\ &= \tau(\hat{F}(x))'\hat{F}(x) && \text{since } x'\hat{F}(x) = 0 \\ &> 0 && \text{by acuteness (5.22). } \blacksquare \end{aligned}$$

Exercise 5.5.7. Suppose that revision protocol τ^p is Lipschitz continuous, acute, and *separable*:

$$\tau_i^p(\pi^p) \equiv \tau_i^p(\pi_i^p).$$

Show that τ^p also satisfies sign preservation (5.23).

Exercise 5.5.8. This exercise shows how to establish properties (PC) and (NS) using only continuity and acuteness (5.22)—that is, without requiring sign preservation (5.23). The proofs of Lemmas 5.5.4 and 5.5.6 go through unchanged, but Lemma 5.5.5 requires additional work. Using acuteness and continuity, show that

- (i) If $\hat{\pi} \in \text{bd}(\mathbf{R}_*^n)$ and $\hat{\pi}_i < 0$, then $\tau_i(\hat{\pi}) = 0$. (Hint: Consider $\hat{\pi}^\varepsilon = \hat{\pi} + \varepsilon e_j$, where $\hat{\pi}_j = 0$.)
- (ii) If $\hat{\pi} \in \text{bd}(\mathbf{R}_*^n)$ and $\hat{\pi}_i = \hat{\pi}_j = 0$, then $\tau(\hat{\pi}) = \mathbf{0}$. (Hint: To show that $\tau_i(\hat{\pi}) = 0$, consider $\hat{\pi}^\varepsilon = \hat{\pi} - \varepsilon e_i + \varepsilon^2 e_j$.)

Then use these two facts to prove Lemma 5.5.5.

Exercise 5.5.9. This exercise demonstrates that in general, one cannot “normalize” a target dynamic in order to create an exact target dynamic. This highlights a nontrivial sense in which the former class of dynamics is more general than the latter.

Recall that in the single population setting, the BNN dynamic is defined by the target protocol $\tau_i(\hat{\pi}) = [\hat{\pi}_i]_+$.

- (i) It is tempting to try to define an exact target protocol by normalizing τ in an appropriate way. Explain why such a protocol would not be well-defined.
- (ii) To attempt to circumvent this problem, one can construct a dynamic that is derived from the normalized protocol whenever the latter is well-defined. Show that such a dynamic must be discontinuous in some games. (Hint: It is enough to consider two-strategy games.)

5.6 Pairwise Comparison Dynamics

Excess payoff dynamics satisfy Nash stationarity (NS), positive correlation (PC), and continuity (C), but they fall in the quite demanding data requirement class (D+). The revision protocols that underlie these dynamics require agents to compare their current payoff with the average payoff obtained in their population. Without the assistance of a central planner, the latter information is unlikely to be known to the agents.

A natural way to reduce these informational demands is to replace the population’s average payoff with another reference payoff, one whose value agents can directly access. We accomplish this by considering revision protocols based on pairwise payoff comparisons, which fall in data requirement class (D2). In the remainder of this section, we show that the resulting evolutionary dynamics can be made to satisfy our other desiderata as well.

5.6.1 Definition

Suppose that the revision protocol ρ^p only directly conditions on payoffs, not the population state. The induced mean dynamic is then of the form

$$(5.24) \quad \dot{x}_i^p = \sum_{j \in S^p} x_j^p \rho_{ji}^p(F^p(x)) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(F^p(x)),$$

This equation and a mild monotonicity condition on ρ defines our next class of dynamics.

Definition. Suppose that the revision protocol ρ is Lipschitz continuous and sign preserving:

$$(5.25) \quad \text{sgn}(\rho_{ij}^p(\pi^p)) = \text{sgn}([\pi_j^p - \pi_i^p]_+) \quad \text{for all } i, j \in S^p \text{ and } p \in \mathcal{P}.$$

Then the map from population games $F \in \mathcal{F}$ to differential equations (5.24) is called a pairwise comparison dynamic.

Sign preservation (5.25) is a particularly natural property: it says that the conditional switch rate from $i \in S^p$ to $j \in S^p$ is positive if and only if the payoff to j exceeds the payoff to i .

Example 5.6.1. The Smith dynamic. The simplest sign preserving revision protocol,

$$\rho_{ij}^p(\pi^p) = [\pi_j^p - \pi_i^p]_+.$$

generates the *Smith dynamic*:

$$(S) \quad \dot{x}_i^p = \sum_{j \in S^p} x_j^p [F_i^p(x) - F_j^p(x)]_+ - x_i^p \sum_{j \in S^p} [F_j^p(x) - F_i^p(x)]_+.$$

Exercise 5.6.2. The k -Smith dynamic. Consider instead the protocol $\rho_{ij}^p(\pi^p) = [\pi_j^p - \pi_i^p]_+^k$, where $k \geq 1$. Argue informally that in the single population case, when k is large, the direction of motion from most states x is approximately parallel to an edge of the simplex. How is this edge determined from the payoff vector $F(x)$?

5.6.2 Incentives and Aggregate Behavior

Our main result in this section is

Theorem 5.6.3. *Every pairwise comparison dynamic satisfies positive correlation (PC) and Nash stationarity (NS).*

The proof of this theorem relies on three equivalences between properties of Nash equilibria and evolutionary dynamics on the one hand, and requirements that sums of terms of the form $\rho_{ij}^p [F_j^p - F_i^p]_+$, or $\rho_{ij}^p [F_j^p - F_i^p]_+$ equal zero on the other. Sign preservation ensures that sums of the three types are identical, allowing us to establish the result.

In what follows, $\dot{x} = V(x)$ is the pairwise comparison dynamic generated by the population game F and revision protocol ρ .

Lemma 5.6.4. $x \in NE(F) \Leftrightarrow$ For all $i \in S^p$ and $p \in \mathcal{P}$, $x_i^p = 0$ or $\sum_{j \in S^p} [F_j^p(x) - F_i^p(x)]_+ = 0$.

Proof. Both statements say that each strategy in use at x is optimal. ■

Lemma 5.6.5. $V^p(x) = \mathbf{0} \Leftrightarrow$ For all $i \in S^p$, $x_i^p = 0$ or $\sum_{j \in S^p} \rho_{ij}^p(F^p(x)) = 0$.

Proof. (\Leftarrow) Immediate.

(\Rightarrow) Fix a population $p \in \mathcal{P}$, and suppose that $V^p(x) = \mathbf{0}$. If j is an optimal strategy for population p at x , then sign preservation implies that $\rho_{jk}^p(F^p(x)) = 0$ for all $k \in S^p$, and so that there is no “outflow” from strategy j :

$$x_j^p \sum_{i \in S^p} \rho_{ji}^p(F^p(x)) = 0.$$

Since $V_j^p(x) = 0$, there can be no “inflow” into strategy j either:

$$\sum_{i \in S^p} x_i^p \rho_{ij}^p(F^p(x)) = 0.$$

We can express this condition equivalently as

$$\text{For all } i \in S^p, \text{ either } x_i^p = 0 \text{ or } \rho_{ij}^p(F^p(x)) = 0.$$

If all strategies in S^p earn the same payoff at state x , the proof is complete. Otherwise, let i be a “second best” strategy—that is, a strategy whose payoff $F_i^p(x)$ is second highest among the payoffs available from strategies in S^p at x . The last observation in the previous paragraph and sign preservation tell us that there is no outflow from i . But since $V_i^p(x) = 0$, there is also no inflow into i :

$$\text{For all } k \in S^p, \text{ either } x_k^p = 0 \text{ or } \rho_{ki}^p(F^p(x)) = 0.$$

Iterating this argument for strategies with lower payoffs establishes the result. ■

Lemma 5.6.6. Fix a population $p \in \mathcal{P}$. Then

(i) $V^p(x)'F^p(x) \geq 0$.

(ii) $V^p(x)'F^p(x) = 0 \Leftrightarrow$ For all $i \in S^p$, $x_i^p = 0$ or $\sum_{j \in S^p} \rho_{ij}^p(F^p(x))[F_j^p(x) - F_i^p(x)]_+ = 0$.

Proof. We compute the inner product as follows:

$$V^p(x)'F^p(x) = \sum_{j \in S^p} \left(\sum_{i \in S^p} x_i^p \rho_{ij}^p(F^p(x)) - x_j^p \sum_{i \in S^p} \rho_{ji}^p(F^p(x)) \right) F_j^p(x)$$

$$\begin{aligned}
&= \sum_{j \in SP} \sum_{i \in SP} (x_i^p \rho_{ij}^p(F^p(x)) F_j^p(x) - x_j^p \rho_{ji}^p(F^p(x)) F_i^p(x)) \\
&= \sum_{j \in SP} \sum_{i \in SP} x_i^p \rho_{ij}^p(F^p(x)) (F_j^p(x) - F_i^p(x)) \\
&= \sum_{i \in SP} \left(x_i^p \sum_{j \in SP} \rho_{ij}^p(F^p(x)) [F_j^p(x) - F_i^p(x)]_+ \right),
\end{aligned}$$

where the last equality follows from sign-preservation. Both claims directly follow. ■

Theorem 5.6.3 follows easily from these three lemmas and sign preservation (5.25). Sign preservation implies that the second conditions in Lemmas 5.6.4, 5.6.5, and 5.6.6(ii) are equivalent. This observation and Lemmas 5.6.4 and 5.6.5 imply that $x \in NE(F)$ if and only if $V^p(x) = \mathbf{0}$ for all $p \in \mathcal{P}$; this is condition (NS). In addition, the observation, Lemma 5.6.5, and Lemma 5.6.6(ii) imply that $V^p(x) = \mathbf{0}$ if and only if $V^p(x)'F^p(x) = 0$; this fact and Lemma 5.6.6(i) imply that $V^p(x)'F^p(x) > 0$ whenever $V^p(x) \neq \mathbf{0}$, which is condition (PC). This completes the proof of Theorem 5.6.3.

5.6.3 Desiderata Revisited

Pairwise comparison dynamics satisfy all four of the desiderata proposed at the beginning of the chapter: their revision protocols are continuous (C), and satisfy the mild data requirement (D2), while the dynamics themselves satisfy Nash stationarity (NS) and positive correlation (PC). To provide some insight into this result, we compare revision protocols that generate the three key dynamics from this chapter:

$$\begin{aligned}
\text{replicator:} & \quad \rho_{ij}^p(\pi^p, x^p) = \hat{x}_j^p [\pi_j^p - \pi_i^p]_+; \\
\text{BNN:} & \quad \rho_{ij}^p(\pi^p, x^p) = [\pi_j^p - \bar{\pi}^p]_+ = \left[\pi_j^p - \sum_{k \in SP} \hat{x}_k^p \pi_k^p \right]_+; \\
\text{Smith:} & \quad \rho_{ij}^p(\pi^p) = [\pi_j^p - \pi_i^p]_+.
\end{aligned}$$

From the point of view of our desiderata, the protocol that generates the Smith dynamic combines the best features of the other two. Like the protocol for the BNN dynamic, the Smith protocol is based on direct selection of candidate strategies rather than imitation, allowing it to satisfy Nash stationarity (NS). Like the protocol for the replicator dynamic, the Smith protocol is based on comparisons of individual strategies' payoffs rather than comparisons involving aggregate statistics, and so satisfies data requirement (D2). Thus, while the BNN and replicator dynamics each satisfy three of our desiderata, the Smith dynamic satisfies all four.

5.7 Multiple Revision Protocols and Combined Dynamics

The results above might seem to suggest that dynamics satisfying all four desiderata are rather special, in that they must be derived from a very specific sort of revision protocol. We now argue to the contrary that these desiderata are satisfied rather broadly.

To make this point, let us consider what happens if an agent uses multiple revision protocols at possibly different intensities. If an agent uses the revision protocol ρ^V at intensity a and the revision protocol ρ^W at intensity b , then his behavior is described by the new revision protocol $\rho^C = a\rho^V + b\rho^W$. Moreover, since mean dynamics are linear in conditional switch rates, the mean dynamic for the combined protocol is a linear combination of the two original mean dynamics: $C_F = aV_F + bW_F$.

Theorem 5.7.1 links the properties of the original and combined dynamics.

Theorem 5.7.1. *Suppose that the dynamic V_F satisfies (PC), that the dynamic W_F satisfies (PC) and (NS), and that $a, b > 0$. Then the combined dynamic $C_F = aV_F + bW_F$ also satisfies (PC) and (NS).*

Proof. To show that C_F satisfies (PC), suppose that $C_F^p(x) \neq \mathbf{0}$. Then either $V_F^p(x), W_F^p(x)$, or both are not $\mathbf{0}$. Since V_F and W_F satisfy (PC), it follows that $V_F^p(x)'F^p(x) \geq 0$, that $W_F^p(x)'F^p(x) \geq 0$, and that at least one of these inequalities is strict. Consequently, $C_F^p(x)'F^p(x) > 0$, and so C_F satisfies (PC).

Our proof that C_F satisfies (NS) is divided into three cases. First, if x is a Nash equilibrium of F , then it is a rest point of both V_F and W_F by Proposition 5.2.1, and hence a rest point of C_F as well. Second, if x is a non-Nash rest point of V_F , then it is not a rest point of W_F . Since $V_F(x) = \mathbf{0}$ and $W_F(x) \neq \mathbf{0}$, it follows that $C_F(x) = bW_F(x) \neq \mathbf{0}$, so x is not a rest point of C_F . Finally, suppose that x is not a rest point of V_F . Then by Proposition 5.2.1, x is not a Nash equilibrium, and so x is not a rest point of W_F either. Since V_F and W_F satisfy condition (PC), we know that $V_F(x)'F(x) = \sum_{p \in \mathcal{P}} V_F^p(x)'F^p(x) > 0$ and that $W_F(x)'F(x) > 0$. Consequently, $C_F(x)'F(x) > 0$, implying that x is not a rest point of C_F . Thus, C_F satisfies (NS). ■

A key implication of Theorem 5.7.1 is that imitation and Nash stationarity are not incompatible. If agents usually rely on imitative protocols but occasionally follow protocols that directly evaluate strategies' payoffs, then the rest points of the resulting mean dynamics are precisely the Nash equilibria of the underlying game. Indeed, if we combine an imitative dynamic V_F with any small amount of a pairwise comparison dynamic W_F , we obtain a combined dynamic C_F that satisfies all four of our desiderata.




Figure 5.7.1: The $\frac{9}{10}$ replicator + $\frac{1}{10}$ Smith dynamic in RPS.

Example 5.7.2. Figure 5.7.1 presents a phase diagram for the $\frac{9}{10}$ replicator + $\frac{1}{10}$ Smith dynamic in standard Rock-Paper-Scissors. Comparing this diagram to those for the replicator and Smith dynamics alone (Figure 5.3.1), we see that the diagram for the combined dynamic more closely resembles the Smith phase diagram than the replicator phase diagram, and in more than one respect: the combined dynamic has exactly one rest point, the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and all solutions to the combined dynamic converge to this state. We will revisit this fragility of imitative dynamics in Chapter 9, where it will appear in a much starker form. §

5.N Notes

Section 5.1: This section follows Sandholm (2008b).

Section 5.2: A wide variety of payoff monotonicity conditions have been considered in the literature; for examples, see Nachbar (1990), Friedman (1991), Samuelson and Zhang (1992), Swinkels (1993), Ritzberger and Weibull (1995), Hofbauer and Weibull (1996), and Sandholm (2001a). Positive correlation is essentially the weakest condition that has been proposed; see the notes to Section 4.4 below. Friedman's (1991) *weak compatibility* is positive correlation plus the additional restriction that unused strategies are never subsequently chosen. Swinkels (1993) calls a dynamic a *myopic adjustment dynamic* if it satisfies positive

correlation, but he allows $F^p(x)'V^p(x) = 0$ even when $V^p(x) \neq \mathbf{0}$.

Section 5.4: The approach to imitative revision protocols and dynamics in this section builds on the work of Björnerstedt and Weibull (1996), Weibull (1995), and Hofbauer (1995a).

Taylor and Jonker (1978) introduce the replicator dynamic to provide a dynamic analogue of Maynard Smith and Price's (1973) equilibrium (ESS) model of animal conflict. Exercise 5.4.12, which shows that the replicator dynamic is equivalent after a nonlinear (barycentric) change of variable to the Lotka-Volterra equation is due to Hofbauer (1981b). Schuster and Sigmund (1983) further observe that fundamental models of population genetics (e.g., Crow and Kimura (1970)) and of biochemical evolution (e.g., Eigen and Schuster (1979)) can be viewed as special cases of the replicator dynamic; they are also the first to refer to the dynamic by this name. For more on these biological models, see Hofbauer and Sigmund (2003). For a detailed analysis of the replicator dynamic from an economic point of view, see Weibull (1995, Chapter 3). The derivations of the replicator dynamic in Examples 5.4.2, 5.4.4, and 5.4.5 are due to Schlag (1998), Björnerstedt and Weibull (1996), and Hofbauer (1995a), respectively.

The Maynard Smith replicator dynamic can be found in Maynard Smith (1982, Appendices D and J). For a contrast between the standard and Maynard Smith replicator dynamics from a biological point of view, see Hofbauer and Sigmund (1988, Section 27.1). The *i*-logit dynamic is due to Björnerstedt and Weibull (1996) and Weibull (1995).

In the biology literature, the stochastic evolutionary process generated by the revision protocol in Example 5.4.8 (i.e., protocol (5.11) with $w(\pi) = \pi$) is called a *frequency-dependent Moran process*, after Moran (1962). In this context, each animal is programmed to play a particular pure strategy; one interprets the arrival of a revision opportunity as the death of one of the animals in the population, and the revision protocol (5.11) as determining which animal will reproduce asexually to fill the vacancy. This process is usually studied in finite populations, often after the addition of mutations, in order to focus on its stochastic aspects. See Nowak (2006) and the Notes to Chapter 11 for further references and discussion.

Most early work by economists on deterministic evolutionary dynamics focuses on generalizations of the replicator dynamic expressed in terms of percentage growth rates, as in equation (5.17). The condition we call monotone percentage growth rates (5.18) has appeared in many places under a variety of names: *relative monotonicity* (Nachbar (1990)), *order compatibility of predynamics* (Friedman (1991)), *monotonicity* (Samuelson and Zhang (1992)), and *payoff monotonicity* (Weibull (1995)). Aggregate monotone percentage growth rates (5.19) and Exercise 5.4.18 are introduced by Samuelson and Zhang (1992).

Sign-preserving percentage growth rates (5.20) is a condition due to Nachbar (1990); see also Ritzberger and Weibull (1995), who call this condition *payoff positivity*. For surveys of the literature referenced here, see Weibull (1995, Chapters 4 and 5) and Fudenberg and Levine (1998, Chapter 3).

Sections 5.5, 5.6, and 5.7: These sections follow Sandholm (2005a, 2008b).

The Brown-von Neumann-Nash dynamic was introduced in the context of symmetric zero-sum games by Brown and von Neumann (1950). Nash (1951) uses a discrete time analogue of this dynamic as the basis for his simple proof of existence of equilibrium based on Brouwer's Theorem. More recently, the BNN dynamic was reintroduced by Skyrms (1990), Swinkels (1993), and Weibull (1996), and by Hofbauer (2000), who gave the dynamic its name. The Smith dynamic was introduced in the transportation science literature by Smith (1984).

Best Response and Projection Dynamics

6.0 Introduction

This chapter continues the procession of deterministic evolutionary dynamics commenced in Chapter 5. In the first two sections, the step from payoff vector fields to evolutionary dynamics is traversed through a traditional game-theoretic approach, by employing best response correspondences and perturbed versions thereof. The third section follows a geometric approach, defining an evolutionary dynamic via closest point projections of payoff vectors.

The *best response dynamic* embodies the assumption that revising agents always switch to their current best response. Because the best response correspondence is discontinuous and multivalued, the basic properties of solution trajectories under the best response dynamic are quite different from those of our earlier dynamics: multiple solution trajectories can sprout from a single initial condition, and solution trajectories can cycle in and out of Nash equilibria. Despite these difficulties, we will see that analogues of incentive properties (NS) and (PC) still hold true.

While the discontinuity of the best response protocol stands in violation of a basic desideratum from Chapter 5, one can obtain a continuous protocol by working with perturbed payoffs. The resulting *perturbed best response dynamics* are continuous (and even differentiable), and so have well-behaved solution trajectories. While the payoff perturbations prevent our incentive conditions from holding exactly, we show that appropriately perturbed versions of these conditions, defined in terms of so-called “virtual payoffs”, can be proved.

Our final evolutionary dynamic, the *projection dynamic*, is motivated by geometric

considerations: we define the growth rate vector under the projection dynamic to be the closest approximation of the payoff vector by a feasible vector of motion. While the resulting dynamic is discontinuous, its solutions still exist, are unique, and are continuous in their initial conditions; moreover, both of our incentive conditions are easily verified. We show that the projection dynamic can be derived from protocols that reflect “revision driven by insecurity”. These protocols also reveal surprising connections between the projection dynamic and the replicator dynamic, connections that we develop further when studying the global behavior of evolutionary dynamics in Chapter 7.

The dynamics studied in this chapter require us to introduce new mathematical techniques. Determining the basic properties of the best response dynamic and the projection dynamic requires ideas from the theory of differential inclusions (i.e., of set valued differential equations), which we develop in Appendix 6.A. A key tool for analyzing perturbed best response dynamics is the Legendre transform, whose basic properties are explained in Appendix 6.B. These properties are central to our analysis of perturbed maximization, which is offered in Appendix 6.C.

6.1 The Best Response Dynamic

6.1.1 Definition and Examples

Traditional game theoretic analysis is based on the assumption of equilibrium play. This assumption can be split into two distinct parts: that agents have correct beliefs about their opponents’ behavior, and that they choose their strategies optimally given those beliefs. When all agents simultaneously have correct beliefs and play optimal responses, their joint behavior constitutes a Nash equilibrium.

It is natural to introduce an evolutionary dynamic based on similar principles. To accomplish this, we suppose that each agent’s revision opportunities arrive at a fixed rate, and that when an agent receives such an opportunity, he chooses a best response to the current population state. Thus, we assume that each agent responds optimally to correct beliefs whenever he is revising, but not necessarily at other points in time.

Before introducing the best response dynamics, let us review the notions of exact target protocols and dynamics introduced in Section 4.2.2. Under an exact target protocol, conditional switch rates $\rho_{ij}^p(\pi^p, x^p) \equiv \sigma_j^p(\pi^p, x^p)$ are independent of an agent’s current strategy. These rates also satisfy $\sum_{j \in S} \sigma_j^p(\pi^p, x^p) \equiv 1$, so that $\sigma^p(\pi^p, x^p) \in \Delta^p$ is a mixed strategy. Such

a protocol induces the exact target dynamic

$$(6.1) \quad \dot{x}^p = m^p \sigma^p(F^p(x), x^p) - x^p.$$

Under (6.1), the vector of motion \dot{x}^p for population p has its tail at the current state x^p and its head at $m^p \sigma^p$, the representative of the mixed strategy $\sigma^p \in \Delta^p$ in the state space $X^p = m^p \Delta^p$.

The best response protocol is given by the multivalued map

$$(6.2) \quad \sigma^p(\pi^p) = M^p(\pi^p) \equiv \underset{y^p \in \Delta^p}{\operatorname{argmax}} (y^p)' \pi^p.$$

$M^p : \mathbf{R}^{n^p} \Rightarrow \Delta^p$ is the *maximizer correspondence* for population p : the set $M^p(\pi^p)$ consists of those mixed strategies that only place mass on pure strategies optimal under payoff vector π^p . Inserting this protocol into equation (6.1) yields the *best response dynamic*:

$$(BR) \quad \dot{x}^p \in m^p M^p(F^p(x)) - x^p.$$

We can also write (BR) as

$$\dot{x}^p \in m^p B^p(x) - x^p.$$

where $B^p = M^p \circ F^p$ is the *best response correspondence* for population p .

Definition. The best response dynamic assigns each population game $F \in \mathcal{F}$ the set of solutions to the differential inclusion (BR).

All of our dynamics from Chapter 5 are Lipschitz continuous, so the existence and uniqueness of their solutions is ensured by the Picard-Lindelöf Theorem. Since the best response dynamic (BR) is a discontinuous differential inclusion, that theorem does not apply here. But while the map M^p is not a Lipschitz continuous function, it does exhibit other regularity properties: in particular, it is a convex-valued, upper hemicontinuous correspondence. These properties impose enough structure on the dynamic (BR) to establish an existence result.

To state this result, we say that the Lipschitz continuous trajectory $\{x_t\}_{t \geq 0}$ is a *Carathéodory solution* to the differential inclusion $\dot{x} \in V(x)$ if it satisfies $\dot{x}_t \in V(x_t)$ at all but a measure zero set of times in $[0, \infty)$.

Theorem 6.1.1. Fix a continuous population game F . Then for each $\xi \in X$, there exists a trajectory $\{x_t\}_{t \geq 0}$ with $x_0 = \xi$ that is a Carathéodory solution to the differential inclusion (BR).

It is important to note that while solutions to the best response dynamic exist, they need not be unique: as the examples to follow will illustrate, multiple solution trajectories can emanate from a single initial condition. For a brief introduction to the theory of differential inclusions, see Appendix 6.A.1.

In Chapter 4, we justified our focus on the deterministic dynamic generated by a revision protocol through an appeal to a finite horizon approximation theorem. This result, which we present in Chapter 10, tells us that under certain regularity conditions, the stochastic evolutionary process $\{X_t^N\}$ generated by a game F and revision protocol ρ is well approximated by a solution to the mean dynamic (M) over any finite time horizon, so long as the population size is large enough. But because the revision protocol that generates the best response dynamic is discontinuous and multivalued, the finite horizon approximation theorem from Chapter 10 does not apply here: indeed, since σ is multivalued, the Markov process $\{X_t^N\}$ is not even uniquely defined! Nevertheless, we conjecture that it is possible to prove a version of the finite horizon approximation theorem that applies in the present setting (see the Notes).

6.1.2 Construction and Properties of Solution Trajectories

Because solutions to the best response dynamic need not be unique, they can be distinctly more complicated than solutions to Lipschitz continuous dynamics, as we demonstrate in a series of examples below. But before doing this, we show another sense in which solutions to the best response dynamic are rather simple.

Let $\{x_t\}$ be a solution to (BR), and suppose that at all times $t \in [0, T]$, population p 's unique best response to state x_t is the pure strategy $i \in S^p$. Then during this time interval, evolution in population p is described by the affine differential equation

$$\dot{x}^p = m^p e_i^p - x^p.$$

In other words, the population state x^p moves directly towards vertex $v_i^p = m^p e_i^p$ of the set X^p , proceeding more slowly as time passes. It follows that throughout the interval $[0, T]$, the state $(x_t)^p$ lies on the line segment connecting $(x_0)^p$ and v_i^p ; indeed, we can solve the previous equation to obtain an explicit formula for $(x_t)^p$:

$$(x_t)^p = (1 - e^{-t}) v_i^p + e^{-t} (x_0)^p \quad \text{for all } t \in [0, T].$$

Matters are more complicated at states that admit multiple best responses, since at such states more than one future course of evolution is possible. Still, not every element of




Figure 6.1.1: The best response dynamic in RPS.

$B^p(x)$ need define a feasible direction of motion for population p : if $\{(x_t)^p\}$ is to head toward state \hat{x}^p during a time interval of positive length, all pure strategies in the support of \hat{x}^p must remain optimal throughout the interval.

Example 6.1.2. Standard Rock-Paper-Scissors. Suppose a population of agents is randomly matched to play standard Rock-Paper-Scissors:

$$A = \begin{pmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{pmatrix}$$

with $w = l$. The phase diagram for the best response dynamic in $F(x) = Ax$ is presented in Figure 6.1.1. The upper, lower left, and lower right regions of the figure contain the states at which Paper, Scissors, and Rock are the unique best responses; in each of these regions, all solution trajectories head directly toward the appropriate vertex. When the boundary of a best response region is reached, multiple directions of motion are possible, at least in principle. But at all states other than the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, the only direction of motion that can persist for a positive amount of time is the one heading toward the new best response. Therefore, the solution from each initial conditions besides x^* is unique. (The solution from x^* is also unique, as we will see below.)

Figure 6.1.1 appears to show that every solution trajectory converges to x^* . To prove that this is so, we first note that along every solution trajectory $\{x_t\}$, whenever the best response to x_t is unique, we have that

$$(6.3) \quad \frac{d}{dt} \left(\max_{k \in S} F_k(x_t) \right) = - \max_{k \in S} F_k(x_t).$$

To see this, let x_t be a state at which there is a unique optimal strategy—say, Paper. At this state, $\dot{x}_t = e_P - x_t$. Since $F_P(x) = w(x_R - x_S)$, we can compute that

$$\begin{aligned} \frac{d}{dt} F_P(x_t) &= \nabla F_P(x_t)' \dot{x}_t \\ &= w(e_R - e_S)'(e_P - x_t) \\ &= -w(e_R - e_S)'x_t \\ &= -F_P(x_t). \end{aligned}$$

It is clear from Figure 6.1.1 that solutions with initial conditions other than x^* pass through states with multiple best responses at at most a countable number of times. We can therefore integrate equation (6.3) with respect to time. If we write $y_t = \max_{k \in S} F_k(x_t)$, then equation (6.3) says that $\frac{d}{dt} y_t = -y_t$. This implies that $y_t = e^{-t} y_0$, and hence that

$$(6.4) \quad \max_{k \in S} F_k(x_t) = e^{-t} \max_{k \in S} F_k(x_0).$$

Now in standard RPS, the maximum payoff function $\max_{k \in S} F_k$ is nonnegative, equalling zero only at the Nash equilibrium x^* . This fact and equation (6.4) imply that if play begins away from x^* , then the maximal payoff falls over time, converging as t approaches infinity to its minimum value of 0; over this same time horizon, the state x_t converges to the Nash equilibrium x^* . It also follows that the unique solution starting from x^* is the stationary one, since a continuous trajectory leading from x^* to any other state would necessarily violate equation (6.4) during its initial segment. §

Example 6.1.3. Two-strategy coordination. Suppose that agents are matched to play the two strategy game with strategy set $S = \{U, D\}$ and payoff matrix

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

The resulting matching game $F(x) = Ax$ has three Nash equilibria, the two pure equilibria e_U and e_D , and the mixed equilibrium $(x_U^*, x_D^*) = (\frac{2}{3}, \frac{1}{3})$.

To reduce the amount of notation, we let $d = x_D$ represent the proportion of players

choosing strategy D , so that the mixed Nash equilibrium becomes $d^* = \frac{1}{3}$. The best response dynamic for this game is described in terms of the state d as follows:

$$\dot{d} = \begin{cases} \{-d\} & \text{if } d < d^*, \\ [-\frac{1}{3}, \frac{2}{3}] & \text{if } d = d^*, \\ \{1-d\} & \text{if } d > d^*. \end{cases}$$

From every initial condition other than d^* , the dynamic admits a unique solution trajectory that converges to a pure equilibrium:

$$(6.5) \quad d_0 < d^* \Rightarrow d_t = e^{-t}d_0,$$

$$(6.6) \quad d_0 > d^* \Rightarrow d_t = e^{-t}d_0 + (1 - e^{-t}) = 1 - e^{-t}(1 - d_0).$$

But there are many solution trajectories starting from d^* : one solution is stationary; another proceeds to $d = 0$ according to equation (6.5), a third proceeds to $d = 1$ according to equation (6.6), and yet others follows the trajectories in (6.5) and (6.6) after some initial delay.

Notice that solutions (6.5) and (6.6) quickly leave the vicinity of d^* . This is unlike the behavior of Lipschitz continuous dynamics, under which solutions from all initial conditions are unique, and solutions that start near a stationary point move very slowly. §

Exercise 6.1.4. Two-strategy anti-coordination. Suppose players are matched to play the antcoordination game

$$A = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Show that there is a unique solution to this dynamic from each initial condition d_0 . Also, show that each solution reaches the unique Nash equilibrium $d^* = \frac{1}{2}$ in finite time, and express this time as a function of the initial condition d_0 . This is unlike the behavior of Lipschitz continuous dynamics, under which solutions can only reach rest points in the limit as the time t approaches infinity.

Example 6.1.5. Three-strategy coordination. Figure 6.1.2 presents the phase diagram for the




Figure 6.1.2: The best response dynamic in Pure Coordination.

best response dynamic generated by matching in the pure coordination game

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The speed of motion is fastest near the mixed Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. As in Example 6.1.3, solution trajectories are not unique: this time, whenever the state is on the Y-shaped set of boundaries between best response regions, it can leave this set and head into any adjoining basin of attraction. §

Exercise 6.1.6. Good and bad RPS. (i) Using a similar argument to that provided in Example 6.1.2, show that in any good RPS game, the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is globally stable, and that it is reached in finite time from every initial condition.

(ii) Show that in any bad RPS game, solutions starting from almost all initial conditions converge to a limit cycle in the interior of the state space. In addition, argue that there are multiple solutions starting from the Nash equilibrium x^* : one is stationary, while others spiral outward toward the limit cycle. The latter solutions are not differentiable at $t = 0$. It is therefore possible for a solution to escape a Nash




Figure 6.1.3: The best response dynamic in Zeeman's game.

equilibrium without the solution beginning its motion in a well-defined direction. (Hint: Consider backward solution trajectories from initial conditions in the region bounded by the cycle.)

Example 6.1.7. Zeeman's game. Consider the population game $F(x) = Ax$ generated by matching in the symmetric normal form game

$$A = \begin{pmatrix} 0 & 6 & -4 \\ -3 & 0 & 5 \\ -1 & 3 & 0 \end{pmatrix}$$

with strategy set $S = \{U, M, D\}$. The Nash equilibria of F are e_U , $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and $y^* = (\frac{4}{5}, 0, \frac{1}{5})$. The best response dynamic for F is presented in Figure 6.1.3. Solution trajectories from a majority of initial conditions are unique and converge to the pure equilibrium e_U . However, some initial conditions generate multiple solutions. Consider, for example, solutions starting at the interior Nash equilibrium x^* . There is a stationary solution at x^* , as well as solutions that head toward the vertex e_U , possibly after some delay. Other solutions head toward the Nash equilibrium y^* . Some of these converge to y^* ; others leave segment x^*y^* before reaching y^* . Of those that leave, some head to e_U , while others head toward e_D and then return to x^* . If x^* is revisited, any of the behaviors

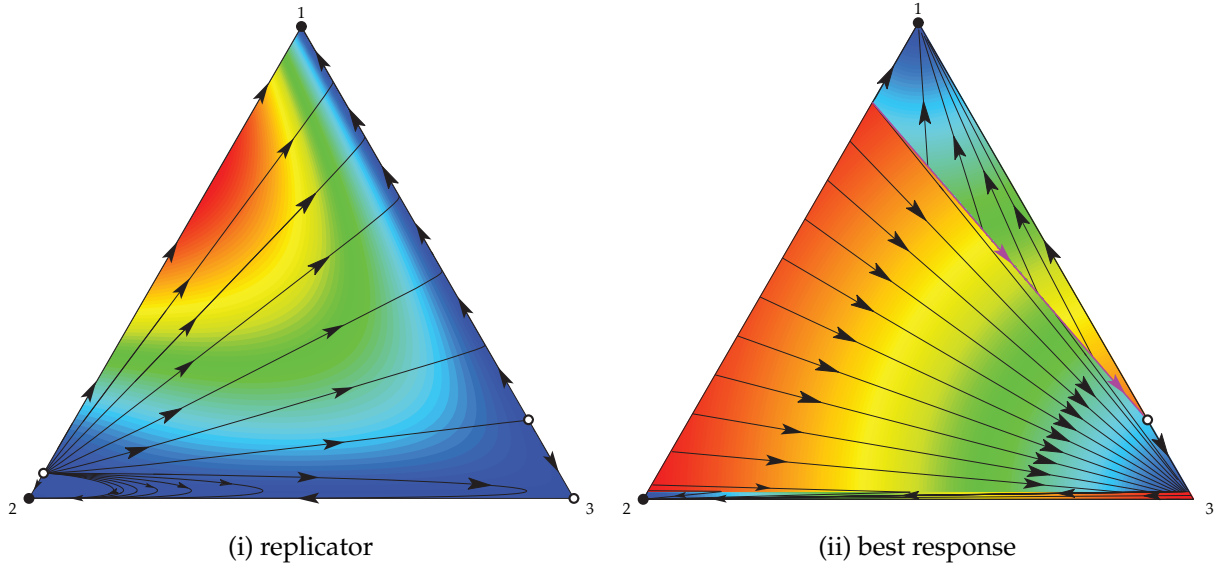


Figure 6.1.4: Unlike basins of attraction under imitative and optimizing dynamics.

just described can occur again. Therefore, there are solutions to (BR) that arrive at and depart x^* in perpetuity. §

Example 6.1.8. While a strict equilibrium is locally stable under the dynamics we have studied so far, the set of initial conditions from which the equilibrium is reached may vary dramatically across dynamics. As an illustration, consider the population game $F(x) = Ax$ generated by matching in the symmetric normal form game

$$A = \begin{pmatrix} 1 & -k & -\frac{1}{k} \\ 2 - k^3 & 2 & 2 \\ 0 & 0 & 0 \end{pmatrix}.$$

If $k > 0$, this game has three Nash equilibria: the strict equilibria e_1 and e_2 , and the mixed equilibrium, $x^* = (\frac{1}{k+1}, 0, \frac{k}{k+1})$. The game has additional restricted equilibria at states e_3 and $y = (\frac{k+2}{k^3+k+1}, \frac{k^3-1}{k^3+k+1}, 0)$; these are rest points under any imitative dynamic (see Section 5.4.6).

Suppose that $k = 5$, so that $x^* = (\frac{1}{6}, 0, \frac{5}{6})$ and $y = (\frac{7}{131}, \frac{124}{131}, 0) \approx (.0534, .9466, 0)$. Under the replicator dynamic (Figure 6.1.4(i)), solutions from most initial conditions converge to the strict equilibrium e_1 . Under the best response dynamic (Figure 6.1.4(ii)), solutions from most initial conditions head toward state e_3 before converging to the strict equilibrium e_2 . In fact, one can show that as k grows large, the basins of e_1 under the replicator dynamic and of e_2 under the best response dynamic include nearly all initial conditions in the simplex (see the Notes).

The sources of the distinct behaviors of the two dynamics are not difficult to explain. At most states in the simplex—in particular, states that are not too close to boundaries e_1e_2 or e_2e_3 —strategy 3 has a slightly higher payoff than strategy 1, while both of these strategies have much higher payoffs than strategy 2: $F_3(x) > F_1(x) \gg F_2(x)$. Under the replicator dynamic, this payoff configuration implies that agents switch from strategy 2 to strategies 1 and 3 in such a way that the ratio of strategy 1 players to strategy 3 players remains nearly constant (see equation (7.39)). If the value of x_1 exceeds $x_1^* = \frac{1}{k+1}$ when virtually no strategy 2 players remain, then $F_1(x)$ will exceed $F_3(x)$, and the state will converge to equilibrium e_1 .

Under the best response dynamic, the payoff ordering $F_3(x) > F_1(x) \gg F_2(x)$ implies that motion is directly toward vertex e_2 : the large payoff difference between strategies 1 and 2 is of no consequence. When the number of strategy 1 players becomes small enough, strategy 2 becomes and then remains optimal, and the state converges to equilibrium e_2 . §

6.1.3 Incentive Properties

In the previous chapter, we introduced two properties, positive correlation (PC) and Nash stationarity (NS), that link growth rates under evolutionary dynamics with payoffs in the underlying games.

(PC) $V_F^p(x) \neq \mathbf{0}$ implies that $V_F^p(x)'F^p(x) > 0$ for all $p \in \mathcal{P}$.

(NS) $V_F(x) = \mathbf{0}$ if and only if $x \in NE(F)$.

Both of these properties are designed for single-valued differential equations. We now establish that analogues of these two properties are satisfied by the differential inclusion (BR).

Theorem 6.1.9. *The best response dynamic satisfies*

$$(6.7) \quad (z^p)'F^p(x) = m^p \max_{j \in S^p} \hat{F}_j^p(x) \text{ for all } z^p \in V_F^p(x).$$

$$(6.8) \quad \mathbf{0} \in V_F(x) \text{ if and only if } x \in NE(F).$$

Condition (6.7) asks that the correspondence $x \mapsto V_F^p(x)'F^p(x)$ be single valued, always equaling the product of population p 's mass and its maximal excess payoff. It follows that this map is Lipschitz continuous and nonnegative, equaling zero if and only if all agents in population p are playing a best response (see Lemma 5.5.4). Summing over populations, we see that $V_F(x)'F(x) = \{\mathbf{0}\}$ if and only if x is a Nash equilibrium of F .

Condition (6.8) requires that the differential inclusion $\dot{x} \in V_F(x)$ have a stationary solution at every Nash equilibrium, but at no other states. As we have seen, this condition does not rule out the existence of additional solution trajectories that leave Nash equilibria.

Proof. Property (6.8) is immediate. To prove property (6.7), fix $x \in X$, and let $z^p \in V_F^p(x)$. Then $z^p = m^p y^p - x^p$ for some $y^p \in M^p(F^p(x))$. Therefore,

$$(z^p)' F^p(x) = (m^p y^p - x^p)' F^p(x) = m^p \max_{j \in S^p} F_j^p(x) - m^p \bar{F}^p(x) = m^p \max_{j \in S^p} \hat{F}_j^p(x). \blacksquare$$

6.2 Perturbed Best Response Dynamics

The best response dynamic is a fundamental model of evolution in games, as it provides an idealized description of the behavior of agents whose decisions condition on exact information about the current strategic environment. Of course, the flip side of exact information is discontinuity, a violation of our desideratum (C) for revision protocols (see Section 5.1).

We now introduce revision protocols under which agents choose best responses to payoffs that have been subjected to perturbations. While the perturbations can represent actual payoff noise, they can also represent errors in agents' perceptions of payoffs, or in the agents' implementations of the best response rule. Regardless of their interpretation, the perturbations lead to revision protocols that are smooth functions of payoffs, and so to dynamics that can be analyzed using standard techniques.

The use of perturbed best response functions is not unique to evolutionary game theory. To mention one prominent example, researchers in experimental economics employ perturbed best response functions when attempting to rationalize experimental data. Consequently, the ideas we develop in this section provide dynamic foundations for solution concepts in common use in experimental research (see the Notes).

6.2.1 Revision Protocols and Mean Dynamics

Perturbed best response protocols are exact target protocols defined in terms of *perturbed maximizer functions* $\tilde{M}^p : \mathbf{R}^{n^p} \rightarrow \text{int}(\Delta^p)$:

$$(6.9) \quad \sigma^p(\pi^p) = \tilde{M}^p(\pi^p).$$

Unlike the maximizer correspondence M^p , the function \tilde{M}^p is single-valued, continuous, and even differentiable. The mixed strategy $\tilde{M}^p(\pi^p) \in \text{int}(\Delta^p)$ places most of its mass on the




Figure 6.2.1: Logit choice probabilities $\tilde{M}_1(\pi_1, 0)$ for noise levels $\eta = .25$ (red), $\eta = .1$ (green), and $\eta = .02$ (blue), along with optimal choice probabilities $M_1(\pi_1, 0)$ (black).

optimal pure strategies, but places positive mass on all pure strategies. Precise definitions of \tilde{M}^p will be stated below.

Example 6.2.1. Logit choice. When $p = 1$, the logit choice function with noise level $\eta > 0$ is written as

$$\tilde{M}_i(\pi) = \frac{\exp(\eta^{-1}\pi_i)}{\sum_{j \in S} \exp(\eta^{-1}\pi_j)}.$$

For any value of $\eta > 0$, each strategy receives positive probability under \tilde{M} regardless of the payoff vector π . But if $\pi_i > \pi_j$ for all $j \neq i$, the probability with which strategy i is chosen approaches one as η approaches zero. Notice too that adding a constant vector to the payoff vector π has no effect on choice probabilities.

When there are just two strategies, the logit choice function reduces to

$$\tilde{M}_1(\pi) = \frac{\exp(\eta^{-1}(\pi_1 - \pi_2))}{\exp(\eta^{-1}(\pi_1 - \pi_2)) + 1} \quad \text{and} \quad \tilde{M}_1(\pi) + \tilde{M}_2(\pi) = 1.$$

In Figure 6.2.1, we fix π_2 at 0, and graph as a function of π_1 the logit(η) choice probabilities $\tilde{M}_1(\pi)$ for $\eta = .25, .1$, and $.02$, as well as the optimal choice probabilities $M_1(\pi)$. Evidently, \tilde{M}_1 provides a smooth approximation of the discontinuous map M_1 . While the function \tilde{M}_1 cannot converge uniformly to the correspondence M_1 as the noise level η goes to zero, one can show that the graph of \tilde{M}_1 converges uniformly (in the Hausdorff metric—see the Notes) to the graph of M_1 as η approaches zero. §

The protocol (6.9) induces the *perturbed best response dynamic*

$$(6.10) \quad \dot{x}^p = m^p \tilde{M}^p(F^p(x)) - x^p$$

as its mean dynamic. We can also write (6.10) as

$$\dot{x}^p = m^p \tilde{B}^p(x) - x^p,$$

where the function $\tilde{B}^p = \tilde{M}^p \circ F^p$, which maps social states to mixed strategies, is the *perturbed best response function* for population p ; it is a perturbed version of the best response correspondence $B^p = M^p \circ F^p$.

6.2.2 Perturbed Optimization: A Representation Theorem

We now consider two methods of defining perturbed maximizer functions. To avoid superscripts, we focus here on the single population case.

The traditional method of defining \tilde{M} , a method with a long history in the theory of discrete choice, is based on *stochastic perturbations* of the payoffs to each *pure* strategy. In this construction, an agent chooses the best response to the vector of payoffs $\pi \in \mathbf{R}^n$, but only after the payoffs to his alternatives have been perturbed by some random vector ε .

$$(6.11) \quad \tilde{M}_i(\pi) = \mathbb{P} \left(i = \operatorname{argmax}_{j \in S} \pi_j + \varepsilon_j \right).$$

We require the random vector ε to be an *admissible stochastic perturbation*: it must admit a positive density on \mathbf{R}^n , and this density must be smooth enough that the function \tilde{M} is continuously differentiable. For example, if the components ε_i are independent, standard results on convolutions imply that \tilde{M} is C^1 whenever the densities of the components ε_i are bounded. In the discrete choice literature, the definition of \tilde{M} via equation (6.11) is known as the *additive random utility model (ARUM)*.

We can also define \tilde{M} by introducing a *deterministic* perturbation of the payoffs to each *mixed* strategy. Call the function $v : \operatorname{int}(\Delta) \rightarrow \mathbf{R}$ an *admissible deterministic perturbation* if it is *differentiably strictly convex* and *steep near* $\operatorname{bd}(\Delta)$. That is, v is admissible if the second derivative at y , $D^2v(y) \in L_s^2(\mathbf{R}_0^n, \mathbf{R})$, is positive definite for all $y \in \operatorname{int}(\Delta)$, and if $|\nabla v(y)|$ approaches infinity whenever y approaches $\operatorname{bd}(\Delta)$. (Recall that \mathbf{R}_0^n is an alternate notation for $T\Delta$, the tangent space of the simplex.) With an admissible v in hand, we define the

function \tilde{M} by

$$(6.12) \quad \tilde{M}(\pi) = \operatorname{argmax}_{y \in \operatorname{int}(\Delta)} (y' \pi - v(y)).$$

One interpretation of the function v is that it represents a “control cost” that becomes large whenever an agent puts too little probability on any particular pure strategy. Because the base payoffs to each strategy are bounded, the steepness of v near $\operatorname{bd}(\Delta)$ implies that it is never optimal for an agent to choose probabilities too close to zero.

Note that under either definition, choice probabilities under \tilde{M} are unaffected by constant shifts in the payoff vector π . The projection of \mathbf{R}^n onto \mathbf{R}_0^n , $\Phi = I - \frac{1}{n}\mathbf{1}\mathbf{1}'$, employs just such a shift, so we can express this property of \tilde{M} as follows:

$$\tilde{M}(\pi) = \tilde{M}(\Phi\pi) \text{ for all } \pi \in \mathbf{R}^n.$$

With this motivation, we define $\bar{M} : \mathbf{R}_0^n \rightarrow \operatorname{int}(\Delta)$ to be the restriction of \tilde{M} to the subspace \mathbf{R}_0^n .

As we noted above, the stochastic construction (6.11) is the traditional way of defining perturbed maximizer functions, and this construction is more intuitively appealing than the deterministic construction (6.12). But the latter construction is clearly more convenient for analysis: while under (6.11) choice probabilities must be expressed as cumbersome multiple integrals, under (6.12) they are obtained as interior maximizers of a strictly concave function.

Happily, we need not trade off intuitive appeal for convenience: every \tilde{M} defined via equation (6.11) can be represented in form (6.12).

Theorem 6.2.2. *Let \tilde{M} be a perturbed maximizer function defined in terms of an admissible stochastic perturbation ε via equation (6.11). Then \tilde{M} satisfies equation (6.12) for some admissible deterministic perturbation v . In fact, $\bar{M} = \tilde{M}|_{\mathbf{R}_0^n}$ and ∇v are invertible, and $\bar{M} = (\nabla v)^{-1}$.*

Taking as given the initial statements in the theorem, it is easy to verify the last one. Indeed, suppose that \tilde{M} (and hence \bar{M}) can be derived from the admissible deterministic perturbation v , that the gradient $\nabla v : \operatorname{int}(\Delta) \rightarrow \mathbf{R}_0^n$ is invertible, and that the payoff vector π is in \mathbf{R}_0^n . Then $y^* = \bar{M}(\pi)$ satisfies

$$y^* = \operatorname{argmax}_{y \in \operatorname{int}(\Delta)} (y' \pi - v(y)).$$

This is a strictly concave maximization problem with an interior solution. Taking the first

order condition with respect to directions in \mathbf{R}_0^n yields

$$\Phi(\pi - \nabla v(y^*)) = \mathbf{0}.$$

Since π and $\nabla v(y^*)$ are already in \mathbf{R}_0^n , the projection Φ does nothing, so rearranging allows us to conclude that

$$\bar{M}(\pi) = y^* = (\nabla v)^{-1}(\pi).$$

In light of this argument, the main task in proving Theorem 6.2.2 is to show that a function v with the desired properties exists. Accomplishing this requires the use of the Legendre transform, a classical tool from convex analysis. We explain the basic properties of the Legendre transform in Appendix 6.B. This device is used to prove the representation theorem in Appendix 6.C, where some auxiliary results can also be found.

One such result is worth mentioning now. Theorem 6.2.2 tells us that every \tilde{M} defined in terms of stochastic perturbations can be represented in terms of deterministic perturbations. Exercise 6.2.3 shows that the converse statement is false, and thus that the deterministic definition of \tilde{M} is strictly more general than the stochastic one.

Exercise 6.2.3. Show that when $n \geq 4$, there is no stochastic perturbation of payoffs which yields the same choice probabilities as the admissible deterministic perturbation $v(y) = -\sum_{j \in S} \log y_j$. (Hint: Use Theorem 6.C.6 in the Appendix.)

6.2.3 Logit Choice and the Logit Dynamic

In Example 6.2.1, we introduced the best known example of a perturbed maximizer function: the *logit choice function* with *noise level* $\eta > 0$.

$$(6.13) \quad \tilde{M}_i(\pi) = \frac{\exp(\eta^{-1}\pi_i)}{\sum_{j \in S} \exp(\eta^{-1}\pi_j)}.$$

This function generates as its mean dynamic the *logit dynamic* with *noise level* η :

$$(L) \quad \dot{x}_i^p = m^p \frac{\exp(\eta^{-1}F_i^p(x))}{\sum_{j \in S^p} \exp(\eta^{-1}F_j^p(x))} - x_i^p.$$

Rest points of logit dynamics are called *logit equilibria*.

Example 6.2.4. In Figure 6.2.2, we present phase diagrams for the 123 Coordination game

$$F(x) = Ax = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix}$$

under logit dynamics with a range of noise levels. As η passes from .01 to 1, the dynamics pass through four distinct regimes. At the lowest noise levels, the dynamics admit seven rest points, three stable and four unstable, corresponding to the seven Nash equilibria of F . When η reaches $\approx .22$, two of the unstable rest points annihilate one another, leaving five rest points in total. At $\eta \approx .28$, the stable rest point corresponding to Nash equilibrium e_1 and an unstable rest point eliminate one another, so that three rest points remain. Finally, when $\eta \approx .68$, the stable rest point corresponding to Nash equilibrium e_2 and an unstable rest point annihilate each other, leaving just a single, stable rest point. If we continue to increase η , the last rest point ultimately converges to the central state $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

This example provides an illustration of a deep topological result called the *Poincaré-Hopf Theorem*. In the present two-dimensional context, this theorem ensures that generically, the number of sinks plus the number of sources equals the number of saddles plus one. §

Example 6.2.5. Stochastic derivation of logit choice. We can derive the logit choice function from stochastic perturbations that are i.i.d. with the *double exponential distribution*: $\mathbb{P}(\varepsilon_i \leq c) = \exp(-\exp(-\eta^{-1}c - \gamma))$, where $\gamma = \lim_{n \rightarrow \infty} (\sum_{k=1}^n \frac{1}{k} - \log n) \approx 0.5772$ is Euler's constant. For intuition, we mention without proof that $\mathbb{E}\varepsilon_i = 0$ and $\text{Var}(\varepsilon_i) = \frac{\eta^2 \pi^2}{6}$, so that $SD(\varepsilon_i) \approx 1.2826\eta$.

To see that these perturbations generate logit choice, note that the density of ε_i is $f(x) = \eta^{-1} \exp(-\eta^{-1}x - \gamma) \exp(-\exp(-\eta^{-1}x - \gamma))$. Using the substitutions $y = \exp(-\eta^{-1}x - \gamma)$ and $m_j = \exp(\eta^{-1}\pi_j)$, we compute as follows:

$$\begin{aligned} \mathbb{P}\left(i = \operatorname{argmax}_{j \in S} \pi_j + \varepsilon_j\right) &= \int_0^\infty f(x) \prod_{j \neq i} F(\pi_i + x - \pi_j) dx \\ &= - \int_0^\infty \eta^{-1} y \exp(-y) \prod_{j \neq i} \exp\left(-y \frac{m_j}{m_i}\right) \frac{\eta}{y} dy \\ &= - \int_0^\infty \exp\left(-y \sum_{j \in S} \frac{m_j}{m_i}\right) dy \\ &= \frac{m_i}{\sum_{j \in S} m_j} \end{aligned}$$

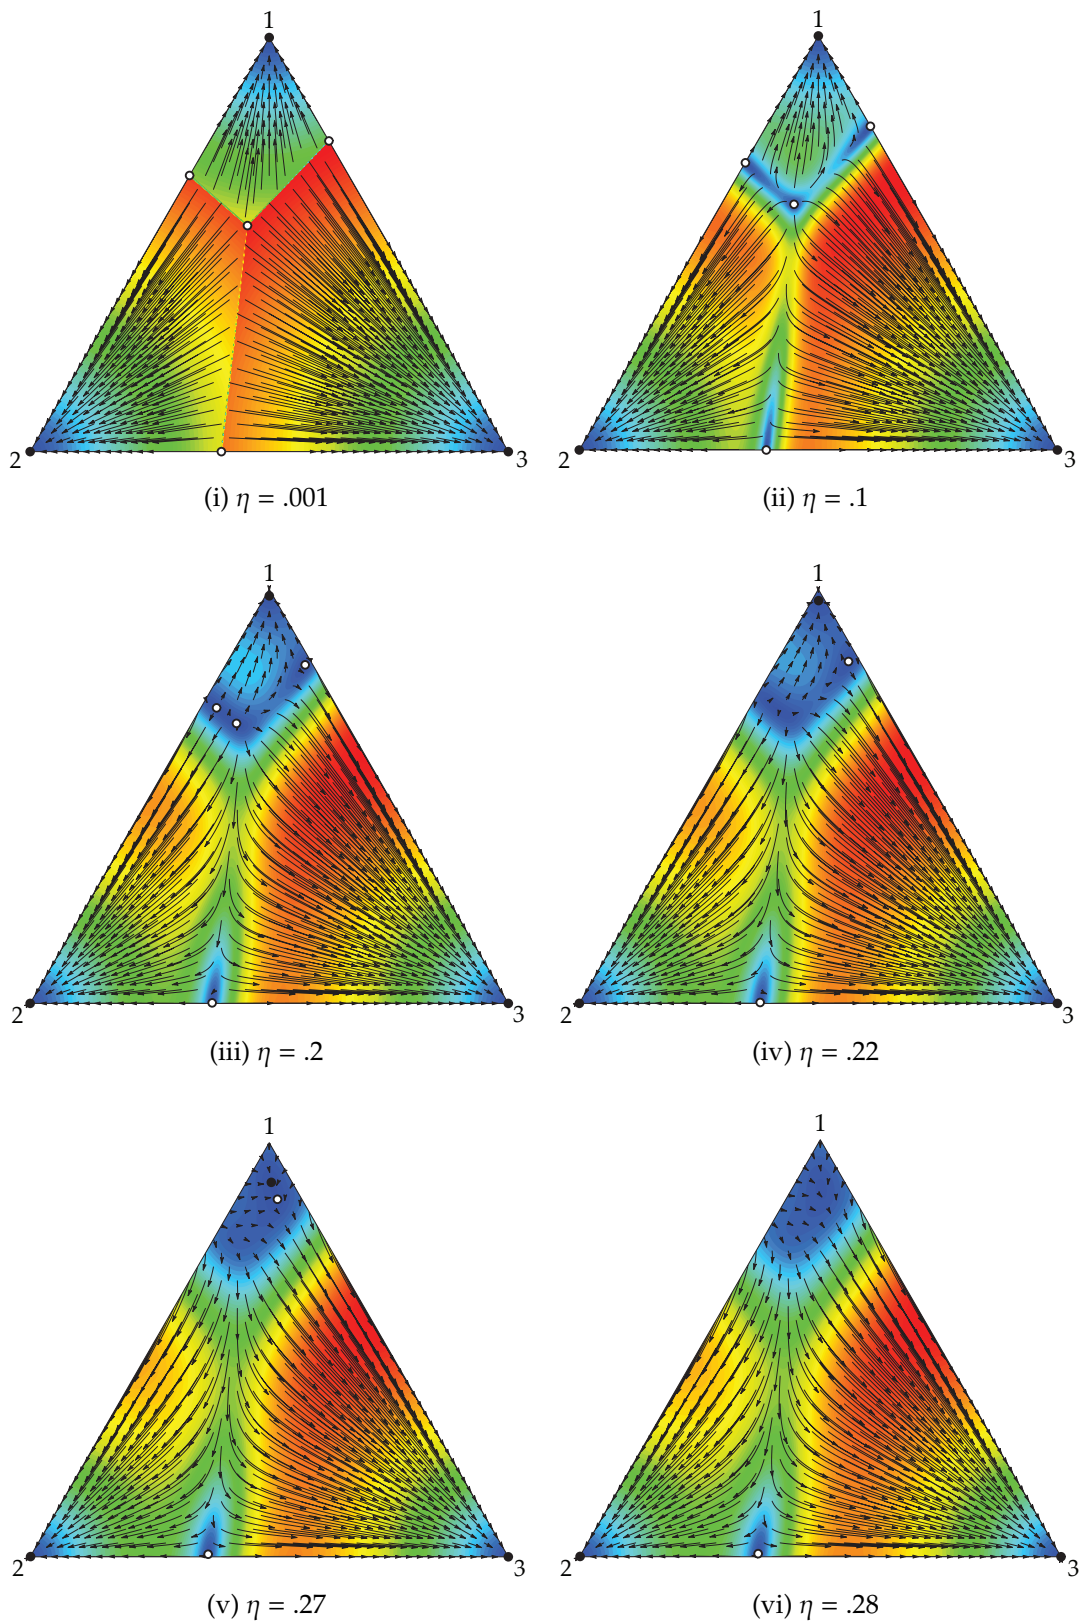


Figure 6.2.2: Logit dynamics in 123 Coordination.

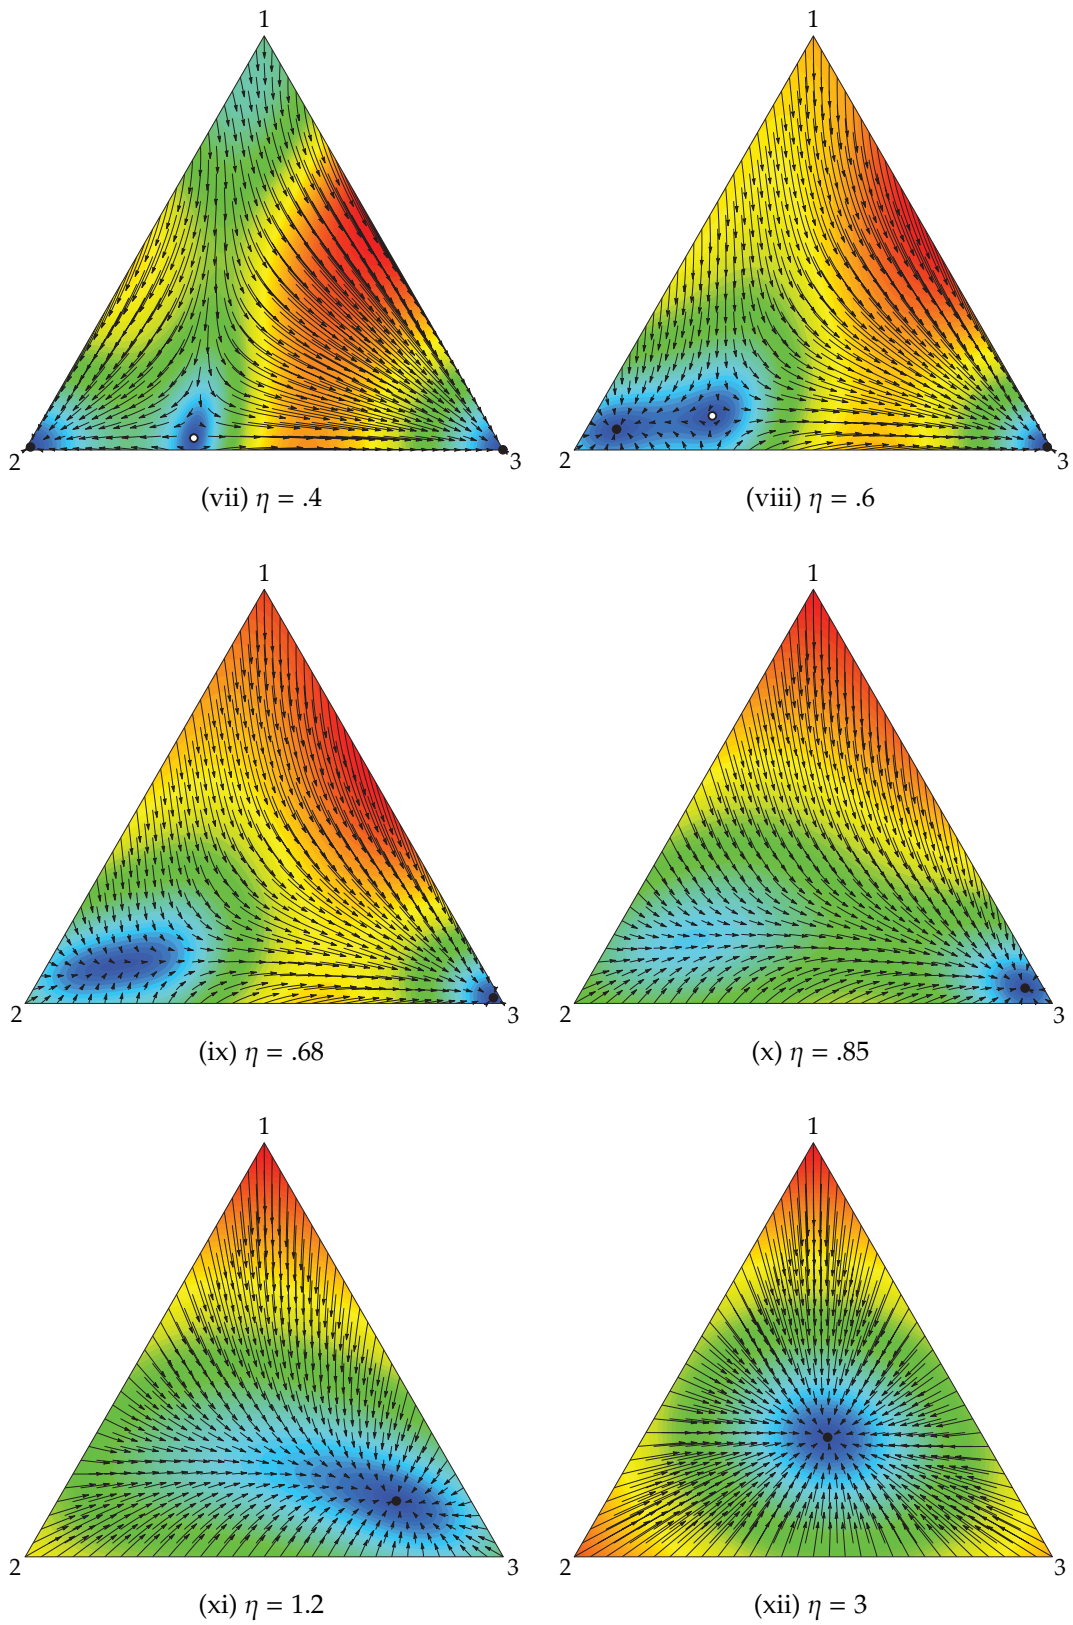


Figure 6.2.2: Logit dynamics in 123 Coordination.

$$= \frac{\exp(\eta^{-1}\pi_i)}{\sum_{j \in S} \exp(\eta^{-1}\pi_j)} \cdot \S$$

Exercise 6.2.6. Deterministic derivation of logit choice. According to the representation theorem, it must also be possible to derive the logit choice function from an admissible deterministic perturbation. Show that this is accomplished using the (negated) entropy function $v(y) = \eta \sum_{j \in S} y_j \log y_j$.

The next exercise gives explicit formulas for various functions from the proof of the representation theorem in the case of logit choice. Included is the derivative matrix $D\tilde{M}(\pi)$, a useful item in analyses of local stability (see Chapter 8.) The exercise also shows how the entropy function v can be derived from the function \tilde{M} .

Exercise 6.2.7. Additional results on logit choice.

- (i) Show that $\tilde{\mu}(\pi) = \eta \log(\sum_{j \in S} \exp(\eta^{-1}\pi_j))$ is a potential function for \tilde{M} . (For the interpretation of this function, see Observation 6.C.3 and Theorem 6.C.4 in the Appendix.)
- (ii) Let $\bar{\mu}$ be the restriction of $\tilde{\mu}$ to \mathbf{R}_0^n , so that $\nabla \bar{\mu}(\pi) = \Phi \tilde{M}(\pi) = \tilde{M}(\pi) - \frac{1}{n} \mathbf{1} = \bar{M}(\pi) - \frac{1}{n} \mathbf{1}$. For $y \in \text{int}(\Delta)$, let $\hat{y} \equiv y - \frac{1}{n} \mathbf{1}$. Show that

$$(\nabla \bar{\mu})^{-1}(\hat{y}) = \bar{M}^{-1}(y) = \eta \begin{pmatrix} \log y_1 - \frac{1}{n} \sum_{j \in S} \log y_j \\ \vdots \\ \log y_n - \frac{1}{n} \sum_{j \in S} \log y_j \end{pmatrix}.$$

- (iii) Let $(C^*, \bar{\mu}^*)$ be the Legendre transform of $(\mathbf{R}_0^n, \bar{\mu})$, and define $v : \text{int}(\Delta) \rightarrow \mathbf{R}$ by $v(y) = \bar{\mu}^*(\hat{y})$. Show by direct computation that $v(y) = \eta \sum_{j \in S} y_j \log y_j$.
- (iv) Show that $\nabla v(y) = \bar{M}^{-1}(y)$. (Hint: Let \tilde{v} be the natural extension of v to \mathbf{R}_+^n , and use the fact $\nabla v(y) = \Phi \nabla \tilde{v}(y)$.)
- (v) Show that $\nabla^2 v(y) = \eta \Phi \text{diag}([y^{-1}]) \Phi$, where $[y^{-1}]_j = y_j^{-1}$ for all $j \in S$.
- (vi) Show that if $\pi \in \mathbf{R}_0^n$, then

$$D\tilde{M}(\pi) = \nabla^2 \tilde{\mu}(\pi) = \eta^{-1} \left(\text{diag}(\bar{M}(\pi)) - \bar{M}(\pi) \bar{M}(\pi)' \right) = \nabla^2 \bar{\mu}(\pi) = D\bar{M}(\pi).$$

- (vii) Show that $\nabla^2 v(\bar{M}(\pi)) = (\nabla^2 \bar{\mu}(\pi))^{-1}$ when these matrices are viewed as linear maps from \mathbf{R}_0^n to \mathbf{R}_0^n . (Hint: Since both of these maps are of full rank on \mathbf{R}_0^n , it is enough to show that $\nabla^2 \bar{\mu}(\pi) \nabla^2 v(\bar{M}(\pi)) = \Phi$, the orthogonal projection onto \mathbf{R}_0^n .)

Exercise 6.2.8. Suppose that \tilde{M} is a perturbed maximizer function derived from an admissible deterministic perturbation as in equation (6.12) (or from an admissible stochastic

perturbation as in equation (6.11)). Show that if \tilde{M} can be expressed as

$$(6.14) \quad \tilde{M}_i(\pi) = \frac{\alpha(\pi_i)}{\sum_{j \in S} \alpha(\pi_j)}$$

for some increasing differentiable function $\alpha : \mathbf{R} \rightarrow (0, \infty)$, then \tilde{M} is a logit choice function with some noise level $\eta > 0$. (Hint: Combine equation (6.14) with the fact that the derivative matrix $D\tilde{M}(\pi)$ must be symmetric (see Corollary 6.C.5 and Theorem 6.C.6 in the Appendix).)

Exercise 6.2.9. The variable-rate logit dynamic. The *variable-rate logit dynamic* with noise level $\eta > 0$ is defined by

$$(6.15) \quad \dot{x}_i^p = m^p \exp(\eta^{-1} F_i^p(x)) - x_i^p \sum_{j \in S^p} \exp(\eta^{-1} F_j^p(x)).$$

The previous exercise shows that the logit dynamic is the only perturbed best response dynamic that admits a modification of this sort.

- (i) Describe a simple revision protocol that generates this dynamic, and provide an interpretation.
- (ii) Show that if $p = 1$, then (6.15) is equivalent to the logit dynamic (L) up to a change in the speed at which solution trajectories are traversed. Explain why this is not the case when $p \geq 2$.
- (iii) Compare this dynamic with the excess payoff dynamics from Chapter 5. Explain why those dynamics cannot be modified so as to resemble the logit dynamic (L).

6.2.4 Perturbed Incentive Properties via Virtual Payoffs

Because they incorporate payoff disturbances, perturbed best response dynamics cannot satisfy positive correlation (PC) or Nash stationarity (NS). We now show that these dynamics do satisfy suitably perturbed versions of the two incentive properties. In light of the representation theorem, there is no loss of generality in focusing on dynamics generated by admissible deterministic perturbations $v = (v^1, \dots, v^p)$.

We can describe the set of Nash equilibria of F in terms of the best response correspondences B^p :

$$NE(F) = \{x \in X : x^p \in m^p B^p(x) \text{ for all } p \in \mathcal{P}\}.$$

In similar fashion, we define the set of *perturbed equilibria* of the pair (F, v) in terms of the

perturbed best response functions \tilde{B}^p :

$$PE(F, v) = \{x \in X : x^p = m^p \tilde{B}^p(x) \text{ for all } p \in \mathcal{P}\}.$$

By definition, the rest points of the perturbed best response dynamic (6.10) are the perturbed equilibria of (F, v) .

Observation 6.2.10. *All perturbed best response dynamics satisfy perturbed stationarity:*

$$(6.16) \quad V(x) = \mathbf{0} \text{ if and only if } x \in PE(F, v).$$

We can derive an alternate characterization of perturbed equilibrium using the notion of virtual payoffs. Define the *virtual payoffs* $\tilde{F} : \text{int}(X) \rightarrow \mathbf{R}^n$ for the pair (F, v) by

$$\tilde{F}^p(x) = F^p(x) - \nabla v^p\left(\frac{1}{m^p} x^p\right).$$

Thus, the virtual payoff function for population p is the difference between the population's true payoff function and gradient of its deterministic perturbation.

For intuition, let us consider the single population case. When x is far from the boundary of the simplex X , the perturbation v is relatively flat, so the virtual payoffs $\tilde{F}(x)$ are close to the true payoffs $F(x)$. But near the boundary of X , true and virtual payoffs are quite different. For example, when x_i is the only component of x that is close to zero, then for each alternate strategy $j \neq i$, moving "inward" in direction $e_i - e_j$ sharply decreases the value of v ; thus, the directional derivative $\frac{\partial v}{\partial(e_i - e_j)}(x)$ is large in absolute value and negative. It follows that the difference $\tilde{F}_i(x) - \tilde{F}_j(x)$ between these strategies' virtual payoffs is large and positive. In other words, rare strategies are quite desirable in the "virtual game" \tilde{F} .

Individual agents do not use virtual payoffs to decide how to act: to obtain the maximized function in definition (6.12) from the virtual payoff function, we must replace the normalized population state $\frac{1}{m^p} x^p$ with the vector of choice probabilities y^p . But at perturbed equilibria, $\frac{1}{m^p} x^p$ and y^p agree. Therefore, perturbed equilibria of (F, v) correspond to "Nash equilibria" of the "virtual game" \tilde{F} .

Theorem 6.2.11. *Let $x \in X$ be a social state. Then $x \in PE(F, v)$ if and only if $\Phi \tilde{F}^p(x) = \mathbf{0}$ for all $p \in \mathcal{P}$.*

The equality $\Phi \tilde{F}^p(x) = \mathbf{0}$ means that $\tilde{F}^p(x)$ is a constant vector. Since uncommon strategies are quite desirable in the "virtual game" \tilde{F} , no state that includes an unused strategy can be a "Nash equilibrium" of \tilde{F} ; thus, equality of all virtual payoffs in each population is the right definition of "Nash equilibrium" in \tilde{F} .

Theorem 6.2.11 follows immediately from perturbed stationarity (6.16) and Lemma 6.2.12 below.

Lemma 6.2.12. *Let $x \in X$ be a social state. Then $V^p(x) = \mathbf{0}$ if and only if $\Phi\tilde{F}^p(x) = \mathbf{0}$.*

Proof. Using the facts that $\tilde{M}^p(\pi^p) = \bar{M}^p(\Phi\pi^p)$, that $\bar{M}^p = (\nabla v^p)^{-1}$, and that the range of ∇v^p is $\mathbf{R}_0^{n^p}$ (so that $\nabla v^p = \Phi \circ \nabla v^p$), we argue as follows:

$$\begin{aligned} V^p(x) = \mathbf{0} &\Leftrightarrow m^p \tilde{M}^p(F^p(x)) = x^p \\ &\Leftrightarrow \bar{M}^p(\Phi F^p(x)) = \frac{1}{m^p} x^p \\ &\Leftrightarrow \Phi F^p(x) = \nabla v^p\left(\frac{1}{m^p} x^p\right) \\ &\Leftrightarrow \Phi\tilde{F}^p(x) = \mathbf{0}. \blacksquare \end{aligned}$$

Turning now to disequilibrium behavior, recall that positive correlation is defined in terms of inner products of growth rate vectors and payoff vectors:

$$(PC) \quad V_F^p(x) \neq \mathbf{0} \text{ implies that } V_F^p(x)' F^p(x) > 0 \text{ for all } p \in \mathcal{P}.$$

In light of the discussion above, the natural analogue of property (PC) for perturbed best response dynamics replaces the true payoffs $F^p(x)$ with virtual payoffs $\tilde{F}^p(x)$. Doing so yields *virtual positive correlation*:

$$(6.17) \quad V^p(x) \neq \mathbf{0} \text{ implies that } V^p(x)' \tilde{F}^p(x) > 0 \text{ for all } p \in \mathcal{P}.$$

To conclude this section, we verify that all perturbed best response dynamics heed this property.

Theorem 6.2.13. *All perturbed best response dynamics satisfy virtual positive correlation (6.17).*

Proof. Let $x \in X$ be a social state at which $V^p(x) \neq \mathbf{0}$. Then by definition,

$$(6.18) \quad y^p \equiv \tilde{M}^p(F^p(x)) = \bar{M}^p(\Phi F^p(x)) \neq \frac{1}{m^p} x^p.$$

Since $\nabla v^p = (\bar{M}^p)^{-1}$, we can rewrite the equality in expression (6.18) as $\nabla v^p(y^p) = \Phi F^p(x)$. Therefore, since $V^p(x) \in TX^p$, we find that

$$\begin{aligned} V^p(x)' \tilde{F}^p(x) &= \left(m^p \tilde{M}^p(F^p(x)) - x^p \right)' \Phi \tilde{F}^p(x) \\ &= \left(m^p \bar{M}^p(\Phi F^p(x)) - x^p \right)' \left(\Phi F^p(x) - \nabla v^p\left(\frac{1}{m^p} x^p\right) \right) \\ &= m^p \left(y^p - \frac{1}{m^p} x^p \right)' \left(\nabla v^p(y^p) - \nabla v^p\left(\frac{1}{m^p} x^p\right) \right) > 0, \end{aligned}$$

where the final inequality follows from the fact that $y^p \neq \frac{1}{m^p}x^p$ and from the strict convexity of v^p . ■

6.3 The Projection Dynamic

6.3.1 Definition

Our main payoff monotonicity condition for evolutionary dynamics is positive correlation (PC). In geometric terms, (PC) requires that at each state where population p is not at rest, the growth rate vector $V^p(x)$ must form an acute angle with the payoff vector $F^p(x)$. Put differently, (PC) demands that growth rate vectors not distort payoff vectors to too great a degree. Is there an evolutionary dynamic that minimizes this distortion?

If the vector field V is to define an evolutionary dynamic, each growth rate vector $V(x)$ must represent a feasible direction of motion, in the sense of lying in the tangent cone $TX(x)$. Thus, the most direct approach to our question is to always take $V(x)$ to be the closest point in $TX(x)$ to the payoff vector $F(x)$.

Definition. *The projection dynamic associates each population game $F \in \mathcal{F}$ with a differential equation*

$$(P) \quad \dot{x} = \Pi_{TX(x)}(F(x)),$$

where $\Pi_{TX(x)}$ is the closest point projection of \mathbf{R}^n onto the tangent cone $TX(x)$.

It is easy to provide an explicit formula for (P) at social states in the interior of X . Since at such states $TX(x) = TX$, the closest point projection $\Pi_{TX(x)}$ is simply Φ , the orthogonal projection onto the subspace TX . In fact, whenever $x^p \in \text{int}(X^p)$, we have that

$$\dot{x}_i^p = (\Phi F^p(x))_i = F_i^p - \frac{1}{n} \sum_{k \in S} F_k^p(x).$$

Thus, when x^p is an interior population state, the growth rate of strategy $i \in S^p$ is the difference between its payoff and the unweighted average of the payoffs to population p 's strategies.

When x is a boundary state, then the projection $\Pi_{TX(x)}$ does not reduce to an orthogonal projection, so providing an explicit formula for (P) becomes more complicated. Exercise 6.3.1 describes the possibilities in a three-strategy game, while Exercise 6.3.2 provides an explicit formula for the general case.

Exercise 6.3.1. Let F be a three-strategy game. Give an explicit formula for $V(x) = \Pi_{TX(x)}F(x)$ when

- (i) $x \in \text{int}(X)$;
- (ii) $x_1 = 0$ but $x_2, x_3 > 0$;
- (iii) $x_1 = 1$.

Exercise 6.3.2. Let F be an arbitrary single population game. Show that the projection $\Pi_{TX(x)}(v)$ can be expressed as follows:

$$(\Pi_{TX(x)}(v))_i = \begin{cases} v_i - \frac{1}{\#\mathcal{S}(v,x)} \sum_{j \in \mathcal{S}(v,x)} v_j & \text{if } i \in \mathcal{S}(v,x). \\ 0 & \text{otherwise;} \end{cases}$$

Here, the set $\mathcal{S}(v,x) \subseteq S$ contains all strategies in $\text{support}(x)$, along with any subset of $S - \text{support}(x)$ that maximizes the average $\frac{1}{\#\mathcal{S}(v,x)} \sum_{j \in \mathcal{S}(v,x)} v_j$.

6.3.2 Solution Trajectories

The dynamic (P) is clearly discontinuous at the boundary of X , so the existence and uniqueness results for Lipschitz continuous differential equations do not apply. We nevertheless have the following result, which is an immediate consequence of Theorem 6.A.4 in the Appendix.

Theorem 6.3.3. *Fix a Lipschitz continuous population game F . Then for each $\xi \in X$, there exists a unique Carathéodory solution $\{x_t\}_{t \geq 0}$ to the projection dynamic (P) with $x_0 = \xi$. Moreover, solutions to (P) are Lipschitz continuous in their initial conditions: if $\{x_t\}_{t \geq 0}$ and $\{y_t\}_{t \geq 0}$ are solutions to (P), then $|y_t - x_t| \leq |y_0 - x_0| e^{Kt}$ for all $t \geq 0$, where K is the Lipschitz coefficient for F .*

Theorem 6.3.3 shows that the discontinuous differential equation (P) enjoys many of the properties of Lipschitz continuous differential equations. But there are important differences between the two types of dynamics. One difference is easy to spot: solutions to (P) are solutions in the Carathéodory sense, and so can have kinks at a measure zero set of times. Other differences are more subtle. For instance, while the theorem ensures the uniqueness of the forward solution trajectory from each state $\xi \in X$, backward solutions need not be unique. It is therefore possible for distinct solution trajectories of the projection dynamic to merge with one another.

Example 6.3.4. Figure 6.3.1 presents phase diagrams for the projection dynamic in good RPS ($w = 2, l = 1$), standard RPS ($w = l = 1$), and bad RPS ($w = 1, l = 2$). In all three games,




Figure 6.3.1: The projection dynamic in three Rock-Paper-Scissors games.

most solutions spiral around the Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ in a counterclockwise direction.

In good RPS (Figure 6.3.1(i)), all solutions converge to the Nash equilibrium. Solutions that begin close to a vertex hit and then travel along an edge of the simplex before heading into the interior of the simplex forever. Thus, there is a portion of each edge that is traversed by solutions starting from a positive measure set of initial conditions.

In standard RPS (Figure 6.3.1(ii)), all solutions enter closed orbits at a fixed distance from x^* . Solutions starting at distance $\frac{1}{\sqrt{6}}$ or greater from x^* (i.e., all solutions at least as far from x^* as the state $(0, \frac{1}{2}, \frac{1}{2})$) quickly enter the closed orbit at distance $\frac{1}{\sqrt{6}}$ from x^* ; other solutions maintain their initial distance from x^* forever.

In bad RPS (Figure 6.3.1(iii)), all solutions other than the one starting at x^* enter the same closed orbit. This orbit alternates between segments through the interior of X and segments along the boundaries.

Notice that in all three cases, solution trajectories starting in the interior of the state space can reach the boundary in finite time. This is impossible under any of our previous dynamics, including the best response dynamic. §

- Exercise 6.3.5.* (i) Under what conditions is the dynamic (P) described by $\dot{x} = \Phi F(x)$ at all states $x \in X$ (i.e., not just at interior states)?
- (ii) Suppose that $F(x) = Ax$ is generated by matching in the symmetric normal form game A . What do the conditions from part (i) reduce to in this case? (Note that under these conditions, $\dot{x} = \Phi Ax$ is a linear differential equation; it is therefore possible to write down explicit formulas for the solution trajectories (see Chapter 8).)

6.3.3 Incentive Properties

That solutions to the projection dynamic exist, are unique, and are continuous in their initial conditions is not obvious. But given this fact and the manner in which the dynamic is defined, it is not surprising that the dynamic satisfies both of our incentive properties. The proofs of these properties are simple applications of the Moreau Decomposition Theorem: given any closed convex cone $K \subseteq \mathbf{R}^n$ and any vector $\pi \in \mathbf{R}^n$, the projections $\Pi_K(\pi)$ and $\Pi_{K^\circ}(\pi)$ are the unique vectors satisfying $\Pi_K(\pi) \in K$, $\Pi_{K^\circ}(\pi) \in K^\circ$, and $\Pi_K(\pi) + \Pi_{K^\circ}(\pi) = \pi$ (see Appendix 2.A.4).

Theorem 6.3.6. *The projection dynamic satisfies positive correlation (PC) and Nash stationarity (NS).*

Proof. Using the Moreau Decomposition Theorem and the normal cone characterization of Nash equilibrium (see Theorem 2.3.2), we find that

$$\Pi_{TX(x)}(F(x)) = \mathbf{0} \Leftrightarrow F(x) \in NX(x) \Leftrightarrow x \in NE(F),$$

establishing (NS). To prove (PC), we again use the Moreau Decomposition Theorem:

$$\begin{aligned} V^p(x)'F^p(x) &= \Pi_{TX^p(x^p)}(F^p(x))' \left(\Pi_{TX^p(x^p)}(F^p(x)) + \Pi_{NX^p(x^p)}(F^p(x)) \right) \\ &= |\Pi_{TX^p(x^p)}(F^p(x))|^2 \\ &\geq 0. \end{aligned}$$

The inequality binds if and only if $\Pi_{TX^p(x^p)}(F^p(x)) = V^p(x) = \mathbf{0}$. ■

6.3.4 Revision Protocols and Connections with the Replicator Dynamic

To this point, we have motivated the projection dynamic entirely through geometric considerations. Can this dynamic be derived from a model of individual choice? In this section, we describe revision protocols that generate the projection dynamic as their mean dynamics, and use these protocols to argue that the projection dynamic models “revision driven by insecurity”. Our analysis reveals close connections between the projection dynamic and the replicator dynamic, connections that we will develop further in the next chapter.

In the remainder of this section, we focus on the single population setting; the extension to multipopulation settings is straightforward.

If we focus exclusively on interior states, the connections between the replicator and projection dynamics are especially strong. In Chapter 4, we introduced three revision protocols that generate the replicator dynamic as their mean dynamics:

$$(6.19) \quad \rho_{ij}(\pi, x) = x_j[\pi_j - \pi_i]_+;$$

$$(6.20) \quad \rho_{ij}(\pi, x) = x_j(K - \pi_i);$$

$$(6.21) \quad \rho_{ij}(\pi, x) = x_j(\pi_j + K).$$

The x_j term in each formula reflects the fact that these protocols are driven by imitation. For instance, to implement the first protocol, an agent whose clock rings picks an opponent from his population at random; he then imitates this opponent only if the opponents’ payoff is higher, doing so with probability proportional to the payoff difference. The

x_j term in these protocols endows their mean dynamic with a special functional form: the growth rate of each strategy is proportional to its prevalence in the population. For protocol (6.19), the derivation of the mean dynamic proceeds as follows:

$$\begin{aligned}
\dot{x}_i &= \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x) \\
&= \sum_{j \in S} x_j x_i [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} x_j [F_j(x) - F_i(x)]_+ \\
&= x_i \sum_{j \in S} x_j (F_i(x) - F_j(x)) \\
&= x_i \left(F_i(x) - \sum_j x_j F_j(x) \right).
\end{aligned}$$

To derive the projection dynamic on $\text{int}(X)$, we use analogues of the revision protocols above, replacing x_j with $\frac{1}{nx_i}$:

$$(6.22) \quad \rho_{ij}(\pi, x) = \frac{[\pi_j - \pi_i]_+}{nx_i};$$

$$(6.23) \quad \rho_{ij}(\pi, x) = \frac{K - \pi_i}{nx_i};$$

$$(6.24) \quad \rho_{ij}(\pi, x) = \frac{\pi_j + K}{nx_i}.$$

Thus, while in each of the imitative protocols, ρ_{ij} is *proportional* to the mass of agents playing the *candidate strategy* j , in the protocols just above, ρ_{ij} is *inversely proportional* to the mass of agents playing the *current strategy* i . One can therefore designate the projection dynamic as capturing “revision driven by insecurity”, as it describes the behavior of agents who are especially uncomfortable choosing strategies not used by many others.

It is easy to verify that protocols (6.22), (6.23), and (6.24) all induce the projection dynamic on the interior of the state space. In the case of protocol (6.22), the calculation proceeds as follows:

$$\begin{aligned}
\dot{x}_i &= \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x) \\
&= \sum_{j \in S} x_j \frac{[F_i(x) - F_j(x)]_+}{nx_j} - x_i \sum_{j \in S} \frac{[F_j(x) - F_i(x)]_+}{nx_i} \\
&= \frac{1}{n} \sum_{j \in S} (F_i(x) - F_j(x))
\end{aligned}$$

$$= F_i(x) - \frac{1}{n} \sum_{j \in S} F_j(x).$$

Because of the $\frac{1}{x_i}$ term in the revision protocol, the mean dynamic above does not depend directly on the value of x_i , allowing the disappearance rates of rare strategies to stay bounded away from zero. In other words, it is because unpopular strategies can be abandoned quite rapidly that solutions to the projection dynamic can travel from the interior to the boundary of the state space in a finite amount of time.

Except in cases where the projection dynamic is defined by $\dot{x} = \Phi F(x)$ at all states (cf Exercise 6.3.5), the revision protocols above do not generate the projection dynamic on the boundary of X . Exercise 6.3.7 presents a revision protocol that achieves this goal, even while maintaining connections with the replicator dynamic.

Exercise 6.3.7. Consider the following two revision protocols

$$(6.25) \quad \rho_{ij}(\pi, x) = \begin{cases} [\hat{\pi}_i]_- \cdot \frac{x_j [\hat{\pi}_j]_+}{\sum_{k \in S} x_k [\hat{\pi}_k]_+} & \text{if } \sum_{k \in S} x_k [\hat{\pi}_k]_+ > 0, \\ 0 & \text{otherwise.} \end{cases}$$

$$(6.26) \quad \rho_{ij}(\pi, x) = \begin{cases} \frac{[\tilde{\pi}_i^S]_-}{x_i} \cdot \frac{[\tilde{\pi}_j^S]_+}{\sum_{k \in \mathcal{S}(\pi, x)} [\tilde{\pi}_k^S]_+} & \text{if } \sum_{k \in \mathcal{S}(\pi, x)} x_k [\tilde{\pi}_k^S]_+ > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The set $\mathcal{S}(\pi, x)$ in equation (6.26) is defined in Exercise 6.3.2, and $\tilde{\pi}_i^S = \pi_i - \frac{1}{\#\mathcal{S}(\pi, x)} \sum_{k \in \mathcal{S}(\pi, x)} \pi_k$.

- (i) Provide an interpretation of protocol (6.25), and show that it generates the replicator dynamic as its mean dynamic.
- (ii) Provide an interpretation of protocol (6.26), and show that it generates the projection dynamic as its mean dynamic.

Appendix

6.A Differential Inclusions

6.A.1 Basic Theory

A correspondence (i.e., a set valued map) $V : \mathbf{R}^n \rightrightarrows \mathbf{R}^n$ defines a *differential inclusion* via

$$(DI) \quad \dot{x} \in V(x).$$

We call (DI) a *good upper hemicontinuous* (or *good UHC*) differential inclusion if V is:

- (i) Nonempty: $V(x) \neq \emptyset$ for all $x \in \mathbf{R}^n$;
- (ii) Convex valued: $V(x)$ is convex for all $x \in X$;
- (iii) Bounded: There exists a $K \in \mathbf{R}$ such that $\sup\{|y| : y \in V(x)\} \leq K$ for all $x \in \mathbf{R}^n$;
- (iv) Upper hemicontinuous: The *graph* of V , $\text{gr}(V) = \{(x, y) : y \in V(x)\}$, is closed.

While solutions to good UHC differential inclusions are neither as easily defined nor as well behaved as those of Lipschitz continuous differential equations, we will see that analogues of all the main properties of solutions to the latter can be established in the present setting.

The set of feasible directions of motion under (DI) changes abruptly at discontinuities of the correspondence V . Our solution notion for (DI) must therefore admit trajectories with kinks: rather than requiring the relation (DI) to hold at every instant in time, it asks only that (DI) hold at almost all times. To formalize this notion, recall that the set $Z \subseteq \mathbf{R}$ has *measure zero* if for every $\varepsilon > 0$, there is a countable collection of open intervals of total length less than ε that covers Z . A property is said to hold for *almost all* $t \in [0, T]$ if it holds on subset of $[0, T]$ whose complement has measure zero. Finally, we say that a trajectory $\{x_t\}_{t \in [0, T]}$ is a (*Carathéodory*) *solution* to (DI) if it is Lipschitz continuous and if $\dot{x}_t \in V(x_t)$ at almost all times $t \in [0, T]$. Since $\{x_t\}$ is Lipschitz continuous, its derivative \dot{x}_t exists for almost all $t \in [0, T]$, and the Fundamental Theorem of Calculus holds: $x_t - x_s = \int_s^t \dot{x}_u \, du$.

Observe that if $\{x_t\}$ is a Carathéodory solution to a continuous ODE $\dot{x} = V(x)$, it is also a solution to the ODE in the usual sense: $\dot{x}_t = V(x_t)$ at *all* times $t \in [0, T]$. While our new concept does not introduce new solutions to standard differential equations, it enables us to find solutions in settings where solutions of the old sort do not exist. In particular, we have the following existence result.

Theorem 6.A.1. *Let (DI) be a good UHC differential inclusion. Then for each $\xi \in \mathbf{R}^n$ there exists a (Carathéodory) solution $\{x_t\}_{t \in [0, T]}$ to (DI) with $x_0 = \xi$.*

Our forward invariance result for ODEs extends to the current setting as follows:

Theorem 6.A.2. *Let $C \subseteq \mathbf{R}^n$ be a closed convex set, and let $V : C \rightrightarrows \mathbf{R}^n$ satisfy conditions (i)-(iv) above. Suppose that $V(x) \subseteq TC(x)$ for all $x \in C$. Extend the domain of V to all of \mathbf{R}^n by letting $V(y) = V(\Pi_C(y))$ for all $y \in \mathbf{R}^n - C$, and let this extension define the differential inclusion (DI) on \mathbf{R}^n . Then*

- (i) (DI) is a good UHC differential inclusion.
- (ii) (DI) admits a forward solution $\{x_t\}_{t \in [0, T]}$ from each $x_0 \in \mathbf{R}^n$.
- (iii) C is forward invariant under (DI).

Our examples of best response dynamics in Section 6.1 show that differential inclusions can admit multiple solution trajectories from a single initial condition, and hence that solutions need not be continuous in their initial conditions. However, the set of solutions to a differential inclusion still possesses considerable structure. To formalize this claim, let $C_{[0,T]}$ denote the space of continuous trajectories through \mathbf{R}^n over the time interval $[0, T]$, equipped with the maximum norm:

$$C_{[0,T]} = \{x : [0, T] \rightarrow \mathbf{R}^n : x \text{ is continuous}\}, \text{ and}$$

$$\|x\| = \max_{t \in [0,T]} |x_t| \text{ for } x \in C_{[0,T]}.$$

Now recall two definitions from metric space topology. A set $\mathcal{A} \subseteq C_{[0,T]}$ is *connected* if it cannot be partitioned into two nonempty sets, each of which is disjoint from the closure of the other. The set \mathcal{A} is *compact* if every sequence of elements of \mathcal{A} admits a subsequence that converges to an element of \mathcal{A} .

Now let $S_{[0,T]}(V, \xi)$ be the set of solutions to (DI) with initial condition ξ :

$$S_{[0,T]}(V, \xi) = \{x \in C_{[0,T]} : x \text{ is a solution to (DI) with } x_0 = \xi\}.$$

Theorem 6.A.3. *Let (DI) be a good UHC differential inclusion. Then*

- (i) *For each $\xi \in \mathbf{R}^n$, $S_{[0,T]}(V, \xi)$ is connected and compact.*
- (ii) *The correspondence $S_{[0,T]}(V, \cdot) : \mathbf{R}^n \rightarrow C_{[0,T]}$ is upper hemicontinuous.*

Although an initial condition ξ may be the source of many solution trajectories of (DI), part (i) of the theorem shows that the set $S_{[0,T]}(V, \xi)$ of such trajectories has a simple structure: it is connected and compact. Given any continuous criterion $f : C_{[0,T]} \rightarrow \mathbf{R}$ (where continuity is defined with respect to the maximum norm on $C_{[0,T]}$) and any initial condition ξ , connectedness implies that the set of values $f(S_{[0,T]}(V, \xi))$ is an interval, while compactness implies that this set of values is compact; thus, there is a solution which is optimal according to criterion f among those that start at ξ . Part (ii) of the theorem provides an analogue of continuity in initial conditions. It tells us that if a sequence of solution trajectories $\{x^k\}_{k=1}^{\infty}$ to (DI) (with possibly differing initial conditions) converges to some trajectory $x \in C_{[0,T]}$, then x is also a solution to (DI).

6.A.2 Differential Equations Defined by Projections

Let $X \subseteq \mathbf{R}^n$ be a compact convex set, and let $F : X \rightarrow \mathbf{R}^n$ be Lipschitz continuous. We consider the differential equation

$$(P) \quad \dot{x} = \Pi_{TX(x)}(F(x)),$$

where $\Pi_{TX(x)}$ is the closest point projection onto the tangent cone $TX(x)$. This equation provides the closest approximation to the equation $\dot{x} = F(x)$ that is consistent with the forward invariance of X .

Since the right hand side of (P) changes discontinuously at the boundary of X , the Picard-Lindelöf Theorem does not apply here. Indeed, solutions to (P) have different properties than solutions of standard ODEs: for instance, solution trajectories from different initial conditions can merge after a finite amount of time has passed. But like solutions to standard ODEs, forward solutions to the dynamic (P) exist, are unique, and are Lipschitz continuous in their initial conditions.

Theorem 6.A.4. *Let F be Lipschitz continuous. Then for each $\xi \in X$, there exists a unique (Carathéodory) solution $\{x_t\}_{t \geq 0}$ to (P) with $x_0 = \xi$. Moreover, solutions are Lipschitz continuous in their initial conditions: $|y_t - x_t| \leq |y_0 - x_0| e^{Kt}$ for all $t \geq 0$, where K is the Lipschitz coefficient for F .*

We now sketch a proof of this result. Define the multivalued map $V : X \rightrightarrows \mathbf{R}^n$ by

$$V(x) = \bigcap_{\varepsilon > 0} \text{cl} \left(\text{conv} \left(\bigcup_{y \in X: |y-x| \leq \varepsilon} \Pi_{TX(y)}(F(y)) \right) \right).$$

In words, $V(x)$ is the closed convex hull of all values of $\Pi_{TX(y)}(F(y))$ that obtain at points y arbitrarily close to x . It is easy to check that V is upper hemicontinuous with closed convex values. Moreover, $V(x) \cap TX(x)$, the set of feasible directions of motion from x contained in $V(x)$, is always equal to $\{\Pi_{TX(x)}(F(x))\}$, and so in particular is nonempty. Because $V(x) \cap TX(x) \neq \emptyset$, an extension of Theorem 6.A.2 called the *Viability Theorem* implies that for each $\xi \in X$, a solution $\{x_t\}_{t \geq 0}$ to $\dot{x} \in V(x)$ exists. But since $V(x) \cap TX(x) = \{\Pi_{TX(x)}(F(x))\}$, this solution must also solve the original equation (P). This establishes the existence of solutions to (P).

To prove uniqueness and continuity, let $\{x_t\}$ and $\{y_t\}$ be solutions to (P). Using the chain rule, the Moreau Decomposition Theorem, and the Lipschitz continuity of F , we see that

$$\frac{d}{dt} |y_t - x_t|^2 = 2(y_t - x_t)' (\Pi_{TX(y_t)}(F(y_t)) - \Pi_{TX(x_t)}(F(x_t)))$$

$$\begin{aligned}
&= 2(y_t - x_t)'(F(y_t) - F(x_t)) - 2(y_t - x_t)'(\Pi_{NX(y_t)}(F(y_t)) - \Pi_{NX(x_t)}(F(x_t))) \\
&= 2(y_t - x_t)'(F(y_t) - F(x_t)) + 2(x_t - y_t)'\Pi_{NX(y_t)}(F(y_t)) \\
&\quad + 2(y_t - x_t)'\Pi_{NX(x_t)}(F(x_t)) \\
&\leq 2(y_t - x_t)'(F(y_t) - F(x_t)) \\
&\leq 2K|y_t - x_t|^2,
\end{aligned}$$

and hence that

$$|y_t - x_t|^2 \leq |y_0 - x_0|^2 + \int_0^t 2K|y_s - x_s| ds.$$

Gronwall's inequality then implies that

$$|y_t - x_t|^2 \leq |y_0 - x_0|^2 e^{2Kt}.$$

Taking square roots yields the inequality stated in the theorem.

6.B The Legendre Transform

The classical Legendre transform is the key tool for proving Theorem 6.2.2, the representation theorem for the additive random utility model. A generalization of this tool, the so-called Legendre-Fenchel transform, underlies the large deviations techniques we will introduce in Chapter 11. In this section, we introduce Legendre transforms of convex functions defined on open intervals and, more generally, on multidimensional convex domains.

6.B.1 Legendre Transforms of Functions on Open Intervals

Let $C = (a, b) \subseteq \mathbf{R}$ be an open interval, and let $f : C \rightarrow \mathbf{R}$ be a strictly convex, continuously differentiable function that becomes steep at the boundaries of C :

$$\lim_{x \downarrow a} f'(x) = -\infty \text{ if } a > -\infty, \text{ and } \lim_{x \uparrow b} f'(x) = \infty \text{ if } b < \infty.$$

The Legendre transform associates with the strictly convex function f a new strictly convex function f^* . Because $f : C \rightarrow \mathbf{R}$ is strictly convex, its derivative $f' : C \rightarrow \mathbf{R}$ is increasing, and thus invertible. We denote its inverse by $(f')^{-1} : C^* \rightarrow \mathbf{R}$, where the open interval C^* is the range of f' . Since $(f')^{-1}$ is itself increasing, its integral, which we denote $f^* : C^* \rightarrow \mathbf{R}$,

is strictly convex. With the right choice of the constant of integration K , the pair (C^*, f^*) is the Legendre transform of the pair (C, f) . In summary:

$$\begin{array}{ccc}
 f : C \rightarrow \mathbf{R} \text{ is strictly convex} & & f^* \equiv \int (f')^{-1} + K \text{ is strictly convex} \\
 \downarrow & & \uparrow \\
 f' : C \rightarrow C^* \text{ is increasing} & \longrightarrow & (f')^{-1} : C^* \rightarrow C \text{ is increasing}
 \end{array}$$

The cornerstone of the construction above is this observation: the derivative of f^* is the inverse of the derivative of f . That is,

$$(6.27) \quad (f^*)' = (f')^{-1}.$$

Or, in other words,

$$(6.28) \quad f^* \text{ has slope } x \text{ at } y \Leftrightarrow f \text{ has slope } y \text{ at } x.$$

Surprisingly enough, we can specify the function f^* described above in a simple, direct way. We define the *Legendre transform* (C^*, f^*) of the pair (C, f) by

$$C^* = \text{range}(f') \quad \text{and} \quad f^*(y) = \max_{x \in C} xy - f(x).$$

The first order condition of the program at right is $y = f'(x^*(y))$, or, equivalently, $(f')^{-1}(y) = x^*(y)$. On the other hand, if we differentiate f^* with respect to y , the envelope theorem yields $(f^*)'(y) = x^*(y)$. Putting these equations together, we see that $(f^*)'(y) = (f')^{-1}(y)$, which is property (6.27).

Suppose that f'' exists and is positive. Then by differentiating both sides of the identity $(f^*)'(y) = (f')^{-1}(y)$, we find this simple relationship between the second derivatives of f and f^* :

$$(f^*)''(y) = ((f')^{-1})'(y) = \frac{1}{f''(x)}, \quad \text{where } x = (f')^{-1}(y) = x^*(y).$$

In words: to find $(f^*)''(y)$, evaluate f'' at the point $x \in C$ corresponding to $y \in C^*$, and then take the reciprocal.

Our initial discussion of the Legendre transform suggests that it is a *duality relation*: in other words, that one can generate (C, f) from (C^*, f^*) using the same procedure through which (C^*, f^*) is generated from (C, f) . To prove this, we begin with the simple observations that C^* is itself an open interval, and that f^* is itself strictly convex and continuously

differentiable. It is also easy to check that $|(f^*)'(y)|$ diverges whenever y approaches $\text{bd}(C^*)$; in fact, this is just the contrapositive of the corresponding statement about f .

It is easy to verify that $(C^*)^* = C$:

$$(C^*)^* = \text{range}((f^*)') = \text{range}((f')^{-1}) = \text{domain}(f') = C.$$

To show that $(f^*)^* = f$, we begin with the definition of $(f^*)^*$:

$$(f^*)^*(x) = \max_{y \in C^*} xy - f^*(y)$$

Taking the first order condition yields $x = (f^*)'(y^*(x))$, and hence $y^*(x) = ((f^*)')^{-1}(x) = f'(x)$. Since $(f')^{-1}(y) = x^*(y)$, y^* and x^* are inverse functions. We therefore conclude that

$$(f^*)^*(x) = xy^*(x) - f^*(y^*(x)) = xy^*(x) - (x^*(y^*(x))y^*(x) - f(x^*(y^*(x)))) = f(x).$$

Putting this all together, we obtain our third characterization of the Legendre transform and of the implied bijection between C and C^* :

$$(6.29) \quad x \text{ maximizes } xy - f(x) \Leftrightarrow y \text{ maximizes } xy - f^*(y).$$

Example 6.B.1. If $C = \mathbf{R}$ and $f(x) = e^x$, then the Legendre transform of (C, f) is (C^*, f^*) , where $C^* = (0, \infty)$ and $f^*(y) = y \log y - y$. §

Example 6.B.2. Suppose that $c : \mathbf{R} \rightarrow \mathbf{R}$ is a strictly convex cost function. (For convenience, we allow negative levels of output; the next example shows that this is without loss of generality if $c'(0) = 0$.) If output can be sold at price $p \in C^* = \text{range}(c')$, then maximized profit equals

$$\pi(p) = \max_{x \in \mathbf{R}} xp - c(x).$$

Thus, by definition, (C^*, π) is the Legendre transform of (\mathbf{R}, c) . The duality relation tells us that if we started instead with the maximized profit function $\pi : C^* \rightarrow \mathbf{R}$, we could recover the cost function c via the dual program

$$c(x) = \max_{p \in C^*} xp - \pi(p). \quad \S$$

Example 6.B.3. To obtain the class of examples that are easiest to visualize, let the function




Figure 6.B.1: A Legendre transform.

$g : \mathbf{R} \rightarrow \mathbf{R}$ be continuous, increasing, and satisfy

$$\lim_{x \downarrow -\infty} g(x) = -\infty, \lim_{x \uparrow \infty} g(x) = \infty, \text{ and } g(0) = 0.$$

If we define $f(x) = \int_0^x g(s) ds$ on domain \mathbf{R} , then the Legendre transform of (\mathbf{R}, f) is (\mathbf{R}, f^*) , where $f^*(y) = \int_0^y g^{-1}(t) dt$. Evidently, $(f^*)' = g^{-1} = (f')^{-1}$. Indeed, Figure 6.B.1 illustrates that x maximizes $xy - f(x)$ if and only if y maximizes $xy - f^*(y)$, and that f^* has slope x at y if and only if f has slope y at x . §

6.B.2 Legendre Transforms of Functions on Multidimensional Domains

Analogues of all of the previous results can be established in settings with multidimensional domains. Let Z be a linear subspace of \mathbf{R}^n . We call (C, f) a *Legendre pair* if $C \subseteq Z$ is (relatively) open and convex, and f is C^1 , strictly convex, and steep near $\text{bd}(C)$, where f is *steep near* $\text{bd}(C)$ if $|\nabla f(x)| \rightarrow \infty$ whenever $x \rightarrow \text{bd}(C)$.

Our goal is to define a pair (C^*, f^*) that satisfies properties (6.30), (6.31), and (6.32):

$$(6.30) \quad \nabla f^* = (\nabla f)^{-1}.$$

$$(6.31) \quad f^* \text{ has slope } x \text{ at } y \Leftrightarrow f \text{ has slope } y \text{ at } x.$$

$$(6.32) \quad x \text{ maximizes } x'y - f(x) \Leftrightarrow y \text{ maximizes } x'y - f^*(y).$$

As before, we can imagine obtaining f^* from f by differentiating, inverting, and then integrating, as we illustrate in the diagram below:

$$\begin{array}{ccc}
 f : C \rightarrow \mathbf{R} \text{ is strictly convex} & & f^* \equiv \int (\nabla f)^{-1} + K \text{ is strictly convex} \\
 \downarrow & & \uparrow \\
 \nabla f : C \rightarrow C^* \text{ is invertible} & \longrightarrow & (\nabla f)^{-1} : C^* \rightarrow C \text{ is invertible}
 \end{array}$$

Since the domain of f is $C \subseteq Z$, the derivative of f , Df , is a map from C into $L(Z, \mathbf{R})$, the set of linear forms on Z . The gradient of f at x is the unique vector $\nabla f(x) \in Z$ that represents $Df(x)$; thus, ∇f is a map from C into Z .

We define the *Legendre transform* (C^*, f^*) of the pair (C, f) by

$$C^* = \text{range}(\nabla f) \quad \text{and} \quad f^*(y) = \max_{x \in C} x'y - f(x).$$

Theorem 6.B.4 summarizes the Legendre transform's basic properties.

Theorem 6.B.4. *Suppose that (C, f) is a Legendre pair. Then:*

- (i) (C^*, f^*) is a Legendre pair.
- (ii) $\nabla f : C \rightarrow C^*$ is bijective, and $(\nabla f)^{-1} = \nabla f^*$.
- (iii) $f(x) = \max_{y \in C^*} x'y - f^*(y)$.
- (iv) The maximizers x^* and y^* satisfy $x^*(y) = \nabla f^*(y) = (\nabla f)^{-1}(y)$ and $y^*(x) = \nabla f(x) = (\nabla f^*)^{-1}(x)$.

As in the one dimensional case, we can relate the second derivatives of f^* to the second derivatives of f . The second derivative D^2f is a map from C to $L_s^2(Z, \mathbf{R})$, the set of symmetric bilinear forms on $Z \times Z$. The Hessian of f at x , $\nabla^2 f(x) \in \mathbf{R}^{n \times n}$, is the unique representation of $D^2f(x)$ by a symmetric matrix whose rows and columns are in Z . In fact, since the map $z \mapsto \nabla^2 f(x)z$ has range Z , we can view the matrix $\nabla^2 f(x)$ as a linear map from Z to Z . We rely on this observation in the following result.

Corollary 6.B.5. *If $D^2f(x)$ exists and is positive definite for all $x \in C$, then $D^2f^*(y)$ exists and is positive definite for all $y \in C^*$. In fact, $\nabla^2 f^*(y) = (\nabla^2 f(x))^{-1}$ as linear maps from Z to Z , where $x = (\nabla f)^{-1}(y)$.*

In the one dimensional setting, the derivative f' is invertible because it is increasing. Both of these properties also follow from the stronger assumption that $f''(x) > 0$ for all $x \in C$. In the multidimensional setting, it makes no sense to ask whether ∇f is increasing. But there is an analogue of the second derivative condition: namely, that the Hessian $\nabla^2 f(x)$ is positive definite on $Z \times Z$ for all $x \in C$. According to the *Global Inverse Function*

Theorem, any function on a convex domain that is proper (i.e., preimages of compact sets are compact) and whose Jacobian determinant is everywhere nonvanishing is invertible; thus, the fact that $\nabla^2 f(x)$ is always positive definite implies that $(\nabla f)^{-1}$ exists. However, this deep result is not needed to prove Theorem 6.B.4 or Corollary 6.B.5.

6.C Perturbed Optimization

6.C.1 Proof of the Representation Theorem

We now use the results on Legendre transforms from Appendix 6.B to prove Theorem 6.2.2. We defined the perturbed maximizer function \tilde{M} using stochastic perturbations via

$$(6.11) \quad \tilde{M}_i(\pi) = \mathbb{P}\left(i = \operatorname{argmax}_{j \in S} \pi_j + \varepsilon_j\right).$$

Here, the random vector ε is an *admissible stochastic perturbation* if it has a positive density on \mathbf{R}^n , and if this density is sufficiently smooth that \tilde{M} is C^1 . We defined \tilde{M} using deterministic perturbations via

$$(6.12) \quad \tilde{M}(\pi) = \operatorname{argmax}_{y \in \operatorname{int}(\Delta)} (y' \pi - v(y)).$$

Here, the function $v : \operatorname{int}(\Delta) \rightarrow \mathbf{R}$ is an *admissible deterministic perturbation* if the Hessian matrix $\nabla^2 v(y)$ is positive definite on $\mathbf{R}_0^n \times \mathbf{R}_0^n$ for all $y \in \operatorname{int}(\Delta)$, and if $|\nabla v(y)|$ approaches infinity whenever y approaches $\operatorname{bd}(\Delta)$.

Theorem 6.2.2. *Let \tilde{M} be a perturbed maximizer function defined in terms of an admissible stochastic perturbation ε via equation (6.11). Then \tilde{M} satisfies equation (6.12) for some admissible deterministic perturbation v . In fact, $\bar{M} = \tilde{M}|_{\mathbf{R}_0^n}$ and ∇v are invertible, and $\bar{M} = (\nabla v)^{-1}$.*

Proof. The probability that alternative i is chosen when the payoff vector is π is

$$\begin{aligned} \tilde{M}_i(\pi) &= \mathbb{P}(\pi_i + \varepsilon_i \geq \pi_j + \varepsilon_j \text{ for all } j \in S) \\ &= \mathbb{P}(\varepsilon_j \leq \pi_i + \varepsilon_i - \pi_j \text{ for all } j) \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\pi_i + x_i - \pi_1} \cdots \int_{-\infty}^{\pi_i + x_i - \pi_{i-1}} \int_{-\infty}^{\pi_i + x_i - \pi_{i+1}} \cdots \int_{-\infty}^{\pi_i + x_i - \pi_n} f(x) dx_n \cdots dx_{i+1} dx_{i-1} \cdots dx_1 dx_i, \end{aligned}$$

where f is the joint density function of the random perturbations ε . The following lemma lists some properties of the derivative of \tilde{M} .

Lemma 6.C.1. For all $\pi \in \mathbf{R}^n$ we have

- (i) $D\tilde{M}(\pi)\mathbf{1} = \mathbf{0}$.
- (ii) $D\tilde{M}(\pi)$ is symmetric.
- (iii) $D\tilde{M}(\pi)$ has negative off-diagonal elements.
- (iv) $D\tilde{M}(\pi)$ is positive definite with respect to $\mathbf{R}_0^n \times \mathbf{R}_0^n$.

Proof. Part (i) follows from differentiating the identity $\tilde{M}(\pi) = \tilde{M}(\Phi\pi)$. To establish parts (ii) and (iii), let i and $j > i$ be two distinct strategies. Then using the change of variable $\hat{x}_j = \pi_i + x_i - \pi_j$, we find that

$$\begin{aligned} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) &= - \int_{-\infty}^{\infty} \int_{-\infty}^{\pi_i+x_i-\pi_1} \cdots \int_{-\infty}^{\pi_i+x_i-\pi_{i-1}} \int_{-\infty}^{\pi_i+x_i-\pi_{i+1}} \cdots \int_{-\infty}^{\pi_i+x_i-\pi_{j-1}} \int_{-\infty}^{\pi_i+x_i-\pi_{j+1}} \cdots \int_{-\infty}^{\pi_i+x_i-\pi_n} f(x_1, \dots, x_{j-1}, \\ &\quad \pi_i + x_i - \pi_j, x_{j+1}, \dots, x_n) dx_n \cdots dx_{j+1} dx_{j-1} \cdots dx_{i+1} dx_{i-1} \cdots dx_1 dx_i \\ &= - \int_{-\infty}^{\infty} \int_{-\infty}^{\pi_j+\hat{x}_j-\pi_1} \cdots \int_{-\infty}^{\pi_j+\hat{x}_j-\pi_{i-1}} \int_{-\infty}^{\pi_j+\hat{x}_j-\pi_{i+1}} \cdots \int_{-\infty}^{\pi_j+\hat{x}_j-\pi_{j-1}} \int_{-\infty}^{\pi_j+\hat{x}_j-\pi_{j+1}} \cdots \int_{-\infty}^{\pi_j+\hat{x}_j-\pi_n} f(x_1, \dots, x_{i-1}, \\ &\quad \pi_j + \hat{x}_j - \pi_i, x_{i+1}, \dots, x_n) dx_n \cdots dx_{j+1} dx_{j-1} \cdots dx_{i+1} dx_{i-1} \cdots dx_1 d\hat{x}_j \\ &= \frac{\partial \tilde{M}_j}{\partial \pi_i}(\pi), \end{aligned}$$

which implies claims (ii) and (iii). To establish claim (iv), let $z \in \mathbf{R}_0^n$. Then using claims (i), (ii), and (iii) in succession yields

$$\begin{aligned} z' D\tilde{M}(\pi)z &= \sum_{i \in S} \sum_{j \in S} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) z_i z_j = \sum_{i \in S} \sum_{j \neq i} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) z_i z_j + \sum_{i \in S} \left(- \sum_{j \neq i} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) \right) z_i^2 \\ &= \sum_i \sum_{j \neq i} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) (z_i z_j - z_i^2) = \sum_i \sum_{j < i} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) (2z_i z_j - z_i^2 - z_j^2) \\ &= - \sum_i \sum_{j < i} \frac{\partial \tilde{M}_i}{\partial \pi_j}(\pi) (z_i - z_j)^2 > 0. \quad \blacksquare \end{aligned}$$

Since the derivative matrix $D\tilde{M}(\pi)$ is symmetric, the vector field \tilde{M} admits a potential function $\tilde{\mu} : \mathbf{R}^n \rightarrow \mathbf{R}$ (that is, a function satisfying $\nabla \tilde{\mu}(\pi) = \tilde{M}(\pi)$ for all $\pi \in \mathbf{R}^n$). Let $\bar{\mu} = \tilde{\mu}|_{\mathbf{R}_0^n}$ be the restriction of $\tilde{\mu}$ to \mathbf{R}_0^n . Then for all $\pi \in \mathbf{R}_0^n$, $\nabla \bar{\mu}(\pi) \in \mathbf{R}_0^n$ is given by

$$\nabla \bar{\mu}(\pi) = \Phi \nabla \tilde{\mu}(\pi) = \Phi \tilde{M}(\pi) = \tilde{M}(\pi) - \frac{1}{n} \mathbf{1} = \bar{M}(\pi) - \frac{1}{n} \mathbf{1},$$

where the third equality uses the fact that $\tilde{M}(\pi) \in \Delta$.

Since $\nabla^2 \bar{\mu}(\pi) = D\bar{M}(\pi)$ is positive definite with respect to $\mathbf{R}_0^n \times \mathbf{R}_0^n$, $\bar{\mu}$ is strictly convex;

thus, since $\text{bd}(\mathbf{R}_0^n)$ is empty, $(\mathbf{R}_0^n, \bar{\mu})$ is a Legendre pair. Let the pair $(C^*, \bar{\mu}^*)$ be the Legendre transform of $(\mathbf{R}_0^n, \bar{\mu})$, and define the function $v : (C^* + \frac{1}{n}\mathbf{1}) \rightarrow \mathbf{R}$ by $v(y) = \bar{\mu}^*(y - \frac{1}{n}\mathbf{1})$. Theorem 6.2.2 then follows immediately from Lemma 6.C.2.

Lemma 6.C.2. (i) $C^* + \frac{1}{n}\mathbf{1} = \text{int}(\Delta)$.

(ii) $\nabla v : \text{int}(\Delta) \rightarrow \mathbf{R}_0^n$ is the inverse of $\bar{M} : \mathbf{R}_0^n \rightarrow \text{int}(\Delta)$.

(iii) v is an admissible deterministic perturbation.

(iv) $\tilde{M}(\pi) = \text{argmax}_{y \in \text{int}(\Delta)} y' \pi - v(y)$ for all $\pi \in \mathbf{R}^n$.

Proof. (i) The set $C^* = \text{range}(\nabla \bar{\mu}) = \text{range}(\bar{M}) - \frac{1}{n}\mathbf{1}$ is convex by Theorem 6.B.4(i). Moreover, if the components π_j , $j \in J \subset S$ stay bounded while the remaining components approach infinity, then $\tilde{M}_j(\pi) \rightarrow 0$ for all $j \in J$: that is, $\tilde{M}(\pi)$ converges to a subface of the simplex Δ . Thus, $\text{range}(\bar{M}) = \text{range}(\tilde{M}) \subseteq \text{int}(\Delta)$ contains points arbitrarily close to each corner of the simplex. Since $\text{range}(\bar{M})$ is convex, it must equal $\text{int}(\Delta)$.

(ii) Let $y \in \text{int}(\Delta)$. Using Theorem 6.B.4(ii), we find that

$$\nabla v(y) = \nabla \bar{\mu}^*(y - \frac{1}{n}\mathbf{1}) = (\nabla \bar{\mu})^{-1}(y - \frac{1}{n}\mathbf{1}) = \bar{M}^{-1}(y).$$

(iii) $(C^*, \bar{\mu}^*)$ is a Legendre pair by Theorem 6.B.4(i); thus, if $y \rightarrow \text{bd}(\Delta) = \text{bd}(C^*) + \frac{1}{n}\mathbf{1}$, then $|\nabla v(y)| = |\nabla \bar{\mu}^*(y - \frac{1}{n}\mathbf{1})|$ diverges. In addition, since $\nabla^2 \bar{\mu}(\pi) = D\bar{M}(\pi)$ is positive definite with respect to $\mathbf{R}_0^n \times \mathbf{R}_0^n$ for all $\pi \in \mathbf{R}_0^n$, Corollary 6.B.5 implies that $\nabla^2 v(y) = \nabla^2 \bar{\mu}^*(y - \frac{1}{n}\mathbf{1})$ is positive definite with respect to $\mathbf{R}_0^n \times \mathbf{R}_0^n$ for all $y \in \text{int}(\Delta)$.

(iv) Since $\tilde{M}(\cdot) = \tilde{M}(\Phi(\cdot))$, it is enough to consider $\pi \in \mathbf{R}_0^n$. For such π ,

$$\begin{aligned} \text{argmax}_{y \in \text{int}(\Delta)} y' \pi - v(y) &= \left(\text{argmax}_{\hat{y} \in \text{int}(\Delta) - \frac{1}{n}\mathbf{1}} \hat{y}' \pi - \bar{\mu}^*(\hat{y}) \right) + \frac{1}{n}\mathbf{1} \\ &= \nabla \bar{\mu}(\pi) + \frac{1}{n}\mathbf{1} \\ &= \tilde{M}(\pi), \end{aligned}$$

where the second equality follows from Theorem 6.B.4(iv). ■

This completes the proof of Theorem 6.2.2. ■

6.C.2 Additional Results

We conclude this section by stating without proof a few additional results on perturbed optimization. The first two of these concern the construction of the potential function $\tilde{\mu}$

of the perturbed maximizer function \tilde{M} . In fact, two constructions are available, one for each sort of perturbation.

If we define \tilde{M} in terms of an admissible deterministic perturbation v , then one can verify (using the Envelope Theorem or a direct calculation) that the *perturbed maximum function* associated with v is a potential function for \tilde{M} .

Observation 6.C.3. *The function $\tilde{\mu} : \mathbf{R}^n \rightarrow \mathbf{R}$ defined by*

$$\tilde{\mu}(\pi) = \max_{y \in \text{int}(\Delta)} y' \pi - v(y)$$

is a potential function for \tilde{M} as defined in (6.12).

Alternatively, suppose we define \tilde{M} in terms of an admissible stochastic perturbation ε . In this case, the expectation of the maximal perturbed payoff is a potential function for \tilde{M} .

Theorem 6.C.4. *The function $\tilde{\mu} : \mathbf{R}^n \rightarrow \mathbf{R}$ defined by*

$$\tilde{\mu}(\pi) = \mathbb{E} \max_{j \in S} (\pi_j + \varepsilon_j)$$

is a potential function for \tilde{M} as defined in (6.11).

The intuition behind this result is simple. If we marginally increase the value of π_i , the value of the maximum function $\max_j \pi_j + \varepsilon_j$ goes up at a unit rate at those values of ε where strategy i is optimal. The set of ε at which strategy i is optimal also changes, but the contribution of these points to the value of the maximum function is negligible. Building on these observations, one can show that

$$\frac{\partial \tilde{\mu}}{\partial \pi_i}(\pi) = \mathbb{E} 1_{\{i = \arg \max_j \pi_j + \varepsilon_j\}} = \mathbb{P}(i = \arg \max_j \pi_j + \varepsilon_j) = \tilde{M}_i(\pi).$$

Which functions are perturbed maximizer functions? The following characterization of the perturbed maximizer functions that can be derived from admissible deterministic perturbations follows easily from the proof of Theorem 6.2.2.

Corollary 6.C.5. *A bijective function $\tilde{M} : \mathbf{R}^n \rightarrow \text{int}(\Delta)$ can be derived from an admissible deterministic perturbation if and only if $D\tilde{M}(\pi)$ is symmetric, positive definite on \mathbf{R}_0^n , and satisfies $D\tilde{M}(\pi)\mathbf{1} = \mathbf{0}$.*

The counterpart of this result for stochastic perturbations is known as the *Williams-Daly-Zachary Theorem*.

Theorem 6.C.6. A bijective function $\tilde{M} : \mathbf{R}^n \rightarrow \text{int}(\Delta)$ can be derived from an admissible stochastic perturbation if and only if $D\tilde{M}(\pi)$ is symmetric, positive definite on \mathbf{R}_0^n , and satisfies $D\tilde{M}(\pi)\mathbf{1} = \mathbf{0}$, as well as the additional requirement that the partial derivatives of \tilde{M} satisfy

$$(-1)^k \frac{\partial^k \tilde{M}_{i_0}}{\partial \pi_{i_1} \cdots \partial \pi_{i_k}} > 0$$

for each $k = 1, \dots, n - 1$ and each set of $k + 1$ distinct indices $\{i_0, i_1, \dots, i_k\} \subseteq S$.

To establish the necessity of the k th order derivative conditions, one repeatedly differentiates the definition of \tilde{M} . The first order derivative condition is derived in this way in the proof of Theorem 6.2.2. These two results show that deterministic perturbations generate a strictly larger class of perturbed maximizer functions than stochastic perturbations; see Exercise 6.2.3 for an explicit example.

6.N Notes

Section 6.1: The best response dynamic was introduced by Gilboa and Matsui (1991) and further studied by Matsui (1992), Hofbauer (1995b), and Gaunersdorfer and Hofbauer (1995). Hofbauer (1995b) introduced the interpretation of the best response dynamic as a differential inclusion. In the context of normal form games, the best response dynamic can be viewed as a continuous-time version of *fictitious play*, introduced by Brown (1949, 1951) as a method of computing equilibria. See Robinson (1951), Miyazawa (1961), and Shapley (1964) for other early work on this process; more recent references include Milgrom and Roberts (1991), Monderer and Shapley (1996a), Monderer and Sela (1997), Foster and Young (1998), Krishna and Sjöström (1998), and Berger (2005, 2007a,b, 2008). A “refinement” of the best response dynamic under which nonrobust best responses are not chosen is studied by Balkenborg et al. (2008).

A complete analysis of best response dynamics in Rock-Paper-Scissors games (Example 6.1.2 and Exercise 6.1.6) can be found in Gaunersdorfer and Hofbauer (1995). Example 6.1.7 is introduced by Zeeman (1980), who shows that the interior Nash equilibrium of this game is not an ESS but is nevertheless asymptotically stable under the replicator dynamic—see Example 8.5.5. The properties of the best response dynamic described in the example are explained by Hofbauer (1995b). Example 6.1.8 is due to Golman and Page (2009).

An approximation theorem for collections of Markov processes whose mean dynamics are differential inclusions is proved by Benaïm et al. (2005), who work in a setting in which

the step size of the increments of the Markov processes shrinks over time. Results for the constant step size setting are proved by Gorodeisky (2008, 2009).

Section 6.2: This section is based on Hofbauer and Sandholm (2002, 2007).

The perturbed best response dynamic first appears in the work of Fudenberg and Kreps (1993), where it is used to study the discrete time stochastic process called *stochastic fictitious play*; see the Notes to Chapter 12 for further references to work on this process. The logit dynamic first appears in Fudenberg and Levine (1998). For the Hausdorff metric mentioned in Example 6.2.1, see Ok (2007). For further references on logit models in game theory, see the Notes to Chapter 11.

In the experimental economics literature, perturbed equilibrium goes by the name of *quantal response equilibrium*, a term introduced by McKelvey and Palfrey (1995). Some authors use this term more narrowly to refer to logit equilibrium. For more on the use of these concepts in the experimental literature, see Goeree et al. (2008) and the references therein.

The properties of the derivative matrix $D\tilde{M}(\pi)$ have long been known in the discrete choice literature—see McFadden (1981) or Anderson et al. (1992). The control cost interpretation of deterministic perturbations is suggested by van Damme (1991, Chapter 4). That independent ε_i with bounded densities generate a continuously differentiable \tilde{M} follows from standard results on convolutions; see Hewitt and Stromberg (1965, Theorem 21.33).

An intuitive discussion of the Poincaré-Hopf Theorem can be found in Hofbauer and Sigmund (1988, Section 19); see Milnor (1965) for a formal treatment. See Ritzberger (1994), Demichelis and Germano (2000, 2002), and Demichelis and Ritzberger (2003) for intriguing uses of topological ideas to study the global properties of evolutionary game dynamics.

Section 6.3: Nagurney and Zhang (1996, 1997), building on work of Dupuis and Nagurney (1993), introduce the projection dynamic in the context of congestion games. Earlier, Friedman (1991) introduced an evolutionary dynamic that is equivalent to the projection dynamic on $\text{int}(X)$, but that is different at states in $\text{bd}(X)$. The presentation in this section follows Lahkar and Sandholm (2008) and Sandholm et al. (2008). A related dynamic called the *target projection dynamic*, best viewed as a model of learning in normal form games, is introduced in Friesz et al. (1994), reintroduced and named in Sandholm (2005a), and provided with microfoundations and further analyzed in Tsakas and Voorneveld (2009).

Appendix 6.A: Smirnov (2002) provides a readable introduction to the theory of differential inclusions. A more comprehensive but less readable reference is Aubin and Cellina (1984).

The existence of solutions to differential inclusions defined by projections of multivalued maps was proved by Henry (1973); the approach described here follows Aubin and Cellina (1984, Section 5.6). Restricting attention to differential equations defined by projections of Lipschitz continuous functions allows one to establish uniqueness and continuity results, a point noted, e.g., by Dupuis and Nagurney (1993).

Appendix 6.B: Formal treatments of the Legendre transform can be found in Rockafellar (1970) and Hiriart-Urruty and Lemaréchal (2001). Example 6.B.3 is borrowed from Roberts and Varberg (1973, Section 15). For the Global Inverse Function Theorem, see Gordon (1972).

Appendix 6.C : Theorem 6.2.2 is due to Hofbauer and Sandholm (2002). For proofs of Theorem 6.C.4 and 6.C.6, see McFadden (1981) or Anderson et al. (1992). The latter source is a good general reference on discrete choice theory.

Part III

Convergence and Nonconvergence of Deterministic Dynamics

Global Convergence of Evolutionary Dynamics

7.0 Introduction

In the preceding chapters, we introduced a variety of classes of evolutionary dynamics and exhibited their basic properties. Most conspicuously, we established links between the rest points of each dynamic and the Nash equilibria of the underlying game, links that are valid regardless of the nature of the game at hand. This connection is expressed in its strongest form by dynamics satisfying Nash stationarity (NS), under which rest points and Nash equilibria coincide.

Still, once one specifies an explicitly dynamic model of behavior, the most natural approach to prediction is not to focus immediately on equilibrium points, but to determine where the dynamic leads when set in motion from various initial conditions. If equilibrium occurs as the limiting state of this adjustment process, we can feel some confidence in predicting equilibrium play. If instead our dynamics lead to limit cycles or other more complicated limit sets, then these sets rather than the unstable rest points provide superior predictions of behavior.

In this chapter, we seek conditions on games and dynamics under which behavior converges to equilibrium from all or nearly all initial population states. We therefore reconsider the three classes of population games introduced in Chapter 3—potential games, stable games, and supermodular games—and derive conditions on evolutionary dynamics that ensure convergence in each class of games. We also establish convergence results for dominance solvable games, but we shall see in Chapter 9 that these results are not robust to small changes in the dynamics for which they hold.

The most common method for proving global convergence in a dynamical system is by

constructing a *strict Lyapunov function*: a scalar-valued function that the dynamic ascends whenever it is not at rest. When the underlying game is a potential game, the game's potential function provides a natural candidate Lyapunov function for evolutionary dynamics. We verify in Section 7.1 that a potential functions serve as Lyapunov functions under any evolutionary dynamic that satisfies our basic monotonicity condition, positive correlation (PC). We then use this fact to prove global convergence in potential games under all of the evolutionary dynamics studied in Chapters 5 and 6.

Unlike potential games, stable games do not come equipped with a scalar-valued function that is an obvious candidate Lyapunov function for evolutionary dynamics. But the structure of payoffs in these games—already reflected in the fact that their sets of Nash equilibria are convex—makes it natural to expect convergence results to hold.

We develop this intuition in Section 7.2, where we develop approaches to constructing Lyapunov functions for stable games. We find that distance-like functions serve as Lyapunov functions for the replicator and projection dynamics, allowing us to establish global convergence results for these dynamics in strictly stable games. For target dynamics, including excess payoff, best response, and perturbed best response dynamics, we find that integrability of the revision protocol is the key to establishing convergence results. We argue in Section 7.2.2 that in the presence of payoff monotonicity, integrability of the protocol ensures that on average, the vector of motion deviates from the vector of payoffs in the direction of the equilibrium; given the geometry of equilibrium in stable games, this is enough to ensure convergence to equilibrium. All told, we prove global convergence results for all six of our fundamental dynamics.

In Section 7.3, we turn our attention to supermodular games. As these game's essential property is the monotonicity of their best response correspondences, it is not surprising that our convergence results address dynamics that respect this monotone structure. We begin by considering the best response dynamic, using elementary methods to prove a convergence result for supermodular games generated by two-player normal form games that satisfy a “diminishing returns” condition. To obtain convergence results that demand less structure of the game, we appeal to methods from the theory of cooperative differential equations: these are smooth differential equations under which increasing the value of one component of the state variable increases the growth rates of all other components. The smoothness requirement precludes applying these methods to the best response dynamic, but we are able to use them to study perturbed best response dynamics. We prove that after a natural change of coordinates, perturbed best response functions generated by *stochastic* perturbations of payoffs are monotone. Ultimately, this allows us to show that the corresponding perturbed best response dynamics converge to perturbed equilibrium

from almost all initial conditions.

In Section 7.4, we study evolution in games with strictly dominated strategies. We find that under the best response dynamic and under imitative dynamics, strictly dominated strategies are eliminated; so are strategies ruled out by iterative removal of strictly dominated strategies. It follows that in games that are dominance solvable—that is, in games where this iterative procedure leaves only one strategy for each population—the best response dynamic and all imitative dynamics converge to the dominance solution. We should emphasize, however, that these elimination results are *not* robust: we will see in Chapter 9 that under many small modifications of the dynamics covered by our elimination results, strictly dominated strategies can survive.

The definitions and tools from dynamical systems theory needed for our analyses are treated in the Appendix. Appendix 7.A introduces notions of stability, limit behavior, and recurrence for deterministic dynamics. Appendix 7.B presents stability and convergence results for dynamics that admit Lyapunov functions. Finally, Appendix 7.C introduces the theory of cooperative differential equations and monotone dynamical systems.

7.1 Potential Games

7.1.1 Potential Functions as Lyapunov Functions

In a potential game $F : X \rightarrow \mathbf{R}^n$, all information about incentives is captured by the potential function $f : X \rightarrow \mathbf{R}$, in that

$$(7.1) \quad \nabla f(x) = \Phi F(x) \text{ for all } x \in X.$$

In Chapter 3, we characterized Nash equilibria of F as those states that satisfy the Kuhn-Tucker first order conditions for maximizing f on X . We now take a further step, using the potential function to describe disequilibrium adjustment. In Lemma 7.1.1, we show that any evolutionary dynamic that satisfies positive correlation,

$$(PC) \quad V_F^p(x) \neq \mathbf{0} \text{ implies that } V_F^p(x)' F^p(x) > 0,$$

must ascend the potential function f .

To state this result, we introduce the notion of a Lyapunov function. The C^1 function $L : X \rightarrow \mathbf{R}$ is an *(increasing) strict Lyapunov function* for the differential equation $\dot{x} = V_F(x)$ if $\dot{L}(x) \equiv \nabla L(x)' V_F(x) \geq 0$ for all $x \in X$, with equality only at rest points of V_F .

Lemma 7.1.1. *Let F be a potential game with potential function f . Suppose the evolutionary dynamic $\dot{x} = V_F(x)$ satisfies positive correlation (PC). Then f is a strict Lyapunov function for V_F .*

Proof. Follows immediately from condition (PC) and the fact that

$$\dot{f}(x) = \nabla f(x)' \dot{x} = (\Phi F(x))' V_F(x) = \sum_{p \in \mathcal{P}} F^p(x)' V_F^p(x). \blacksquare$$

The initial equality in the expression above follows from an application of the chain rule (Section 3.A.4) to the composition $t \mapsto x_t \mapsto f(x_t)$. Versions of this argument will be used often in the proofs of the results to come.

If a dynamic admits a strict Lyapunov function, all solution trajectories of the dynamic converge to equilibrium. Combining this fact with Lemma 7.1.1 allows us to prove a global convergence result for potential games. To state this result, we briefly present some definitions concerning limit behavior of deterministic trajectories; for more on these notions and on Lyapunov functions, see Appendices 7.A and 7.B.

The ω -limit of trajectory $\{x_t\}_{t \geq 0}$ is the set of all points that the trajectory approaches arbitrarily closely infinitely often:

$$\omega(\{x_t\}) = \left\{ y \in X : \text{there exists } \{t_k\}_{k=1}^{\infty} \text{ with } \lim_{k \rightarrow \infty} t_k = \infty \text{ such that } \lim_{k \rightarrow \infty} x_{t_k} = y \right\}.$$

For dynamics $\dot{x} = V_F(x)$ that admit a unique solution trajectory from each initial condition, we write $\omega(\xi)$ for the ω -limit set of the trajectory starting from state ξ , and we let

$$\Omega(V_F) = \bigcup_{\xi \in X} \omega(\xi)$$

denote the set of all ω -limit points of all solution trajectories. The set $\Omega(V_F)$ (or its closure, when $\Omega(V_F)$ is not closed) provides a basic notion of recurrence for deterministic dynamics.

Theorem 7.1.2. *Let F be a potential game, and let $\dot{x} = V_F(x)$ be an evolutionary dynamic for F that admits a unique solution from each initial condition and that satisfies positive correlation (PC). Then $\Omega(V_F) = RP(V_F)$. In particular,*

- (i) *If V_F is an imitative dynamic, then $\Omega(V_F) = RE(F)$, the set of restricted equilibria of F .*
- (ii) *If V_F is an excess payoff dynamic, a pairwise comparison dynamic, or the projection dynamic, then $\Omega(V_F) = NE(F)$.*

Proof. Immediate from Lemma 7.1.1, Theorem 7.B.4, and the characterizations of rest points from Chapters 5 and 6. \blacksquare

Example 7.1.3. 123 Coordination. Figure 7.1.1 presents phase diagrams for the six fundamental dynamics in 123 Coordination:

$$(7.2) \quad F(x) = Ax = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix}.$$

In the first five cases, the phase diagram is plotted atop the potential function

$$(7.3) \quad f(x) = \frac{1}{2} \left((x_1)^2 + 2(x_2)^2 + 3(x_3)^2 \right).$$

Of these, the first four cases (replicator, projection, BNN, Smith) are covered by Theorem 7.1.2; evidently, every solution trajectory in diagrams (i)–(iv) ascends the potential function, ultimately converging to one of the seven Nash equilibria of F .

It is worth noting that these equilibria are not all locally stable. The interior equilibrium is a *source*, with all nearby solution trajectories moving away from the equilibrium. The three equilibria with two-strategy supports are *saddles*: for each of these, there is one solution trajectory that converges to the equilibrium, while all other nearby trajectories eventually move away from the equilibrium. Only the three remaining equilibria—the three pure equilibria—are locally stable. We defer further discussion of local stability to Chapter 8, which is devoted to this topic. §

Our convergence results for best response and perturbed best response dynamics require additional work. In the case of the best response dynamic

$$(BR) \quad \dot{x}^p \in m^p M^p(F^p(x)) - x^p, \quad \text{where } M^p(\pi^p) = \operatorname{argmax}_{i \in S^p} \pi_i^p,$$

we must account for the fact that the dynamic is multivalued.

Theorem 7.1.4. *Let F be a potential game with potential function f , and let $\dot{x} \in V_F(x)$ be the best response dynamic for F . Then*

$$\frac{\partial f}{\partial z}(x) = \sum_{p \in \mathcal{P}} m^p \max_{j \in S^p} \hat{F}_j^p(x) \quad \text{for all } z \in V_F(x) \text{ and } x \in X.$$

Therefore, every solution trajectory $\{x_t\}$ of V_F satisfies $\omega(\{x_t\}) \subseteq NE(F)$.

Proof. Recall from Theorem 6.1.9 that the best response dynamic satisfies the following

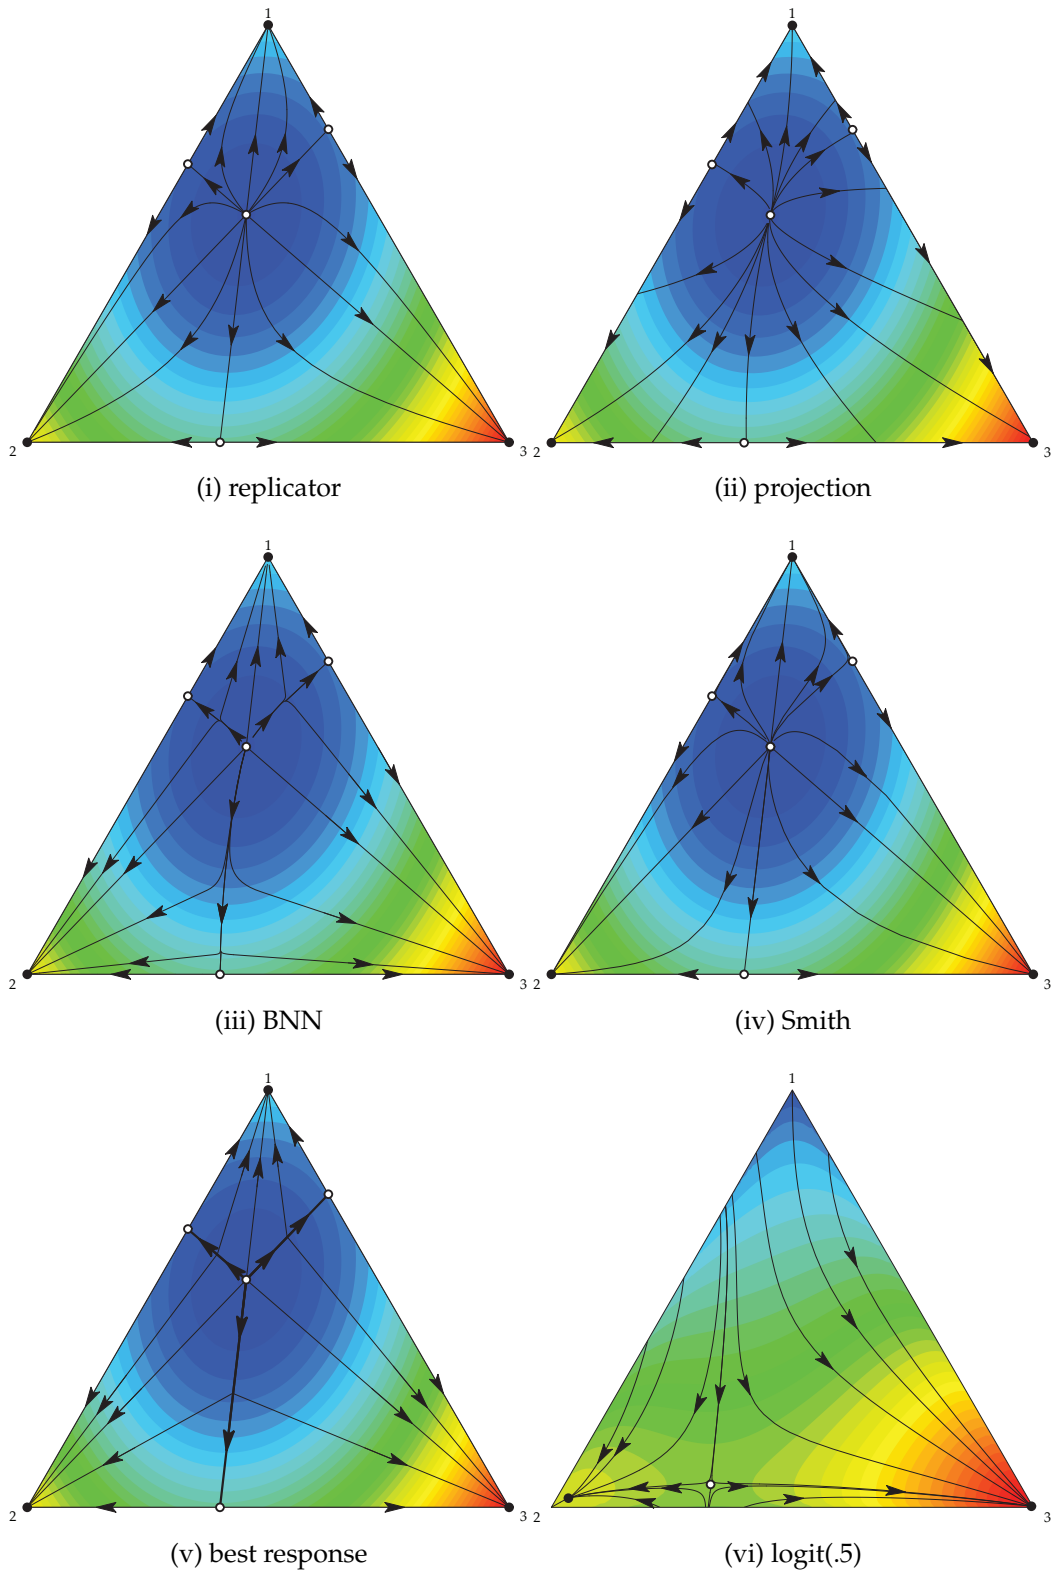


Figure 7.1.1: Six basic dynamics in 123 Coordination. The contour plots are the potential function in (i)-(v), and the logit potential function in (vi).

refinement of condition (PC):

$$(z^p)'F^p(x) = m^p \max_{j \in S^p} \hat{F}_j^p(x) \text{ for all } z^p \in V_F^p(x).$$

This condition immediately implies that

$$\frac{\partial f}{\partial z}(x) \equiv \nabla f(x)'z = (\Phi F(x))'z = \sum_{p \in \mathcal{P}} F^p(x)'z^p = \sum_{p \in \mathcal{P}} m^p \max_{j \in S^p} \hat{F}_j^p(x).$$

Thus, $\frac{\partial f}{\partial z}(x) \geq 0$ for all $x \in X$ and $z \in V_F(x)$, and Lemma 5.5.4 implies that equality holds if and only if $x \in NE(F)$. The convergence result now follows from Theorem 7.B.5. ■

Example 7.1.5. 123 Coordination revisited. Figure 7.1.1(v) presents the phase diagram of the best response dynamic in 123 Coordination (7.2), again atop the potential function (7.3). As in Example 6.1.5, there are multiple solutions starting from each initial condition on the Y-shaped set of boundaries between the best response regions. It is not hard to verify that each of these solutions converges to a Nash equilibrium. §

Finally, we turn to perturbed best response dynamics, considering the (more general) definition of these dynamics via admissible deterministic perturbations $v^p : \text{int}(\Delta^p) \rightarrow \mathbf{R}$.

$$\dot{x}^p = m^p \tilde{M}^p(F(x)) - x^p, \text{ where } \tilde{M}^p(\pi^p) = \underset{y^p \in \text{int}(\Delta^p)}{\text{argmax}} (y^p)' \pi^p - v^p(y^p),$$

While these dynamics do not satisfy positive correlation (PC), Theorem 6.2.13 showed that these dynamics do satisfy a perturbed analogue called virtual positive correlation:

$$V^p(x) \neq \mathbf{0} \text{ implies that } V^p(x)' \tilde{F}^p(x) > 0 \text{ for all } p \in \mathcal{P},$$

where the *virtual payoffs* $\tilde{F} : \text{int}(X) \rightarrow \mathbf{R}^n$ for the pair (F, v) are defined by

$$\tilde{F}^p(x) = F^p(x) - \nabla v^p\left(\frac{1}{m^p} x^p\right).$$

Accordingly, the Lyapunov function for a perturbed best response dynamic is not the potential function f , but a perturbed version thereof.

Theorem 7.1.6. *Let F be a potential game with potential function f , and let $\dot{x} = V_{F,v}(x)$ be the perturbed best response dynamic for F generated by the admissible deterministic perturbations*

$v = (v^1, \dots, v^p)$. Define the perturbed potential function $\tilde{f} : \text{int}(X) \rightarrow \mathbf{R}$ by

$$(7.4) \quad \tilde{f}(x) = f(x) - \sum_{p \in \mathcal{P}} m^p v^p \left(\frac{1}{m^p} x^p \right).$$

Then \tilde{f} is a strict Lyapunov function for $V_{F,v}$, and so $\Omega(V_{F,v}) = PE(F, v)$.

Proof. That \tilde{f} is a strict Lyapunov function for $V_{F,v}$ follows immediately from virtual positive correlation and the fact that

$$\dot{\tilde{f}}(x) \equiv \nabla \tilde{f}(x)' \dot{x} = \sum_{p \in \mathcal{P}} \left(F^p(x) - \nabla v^p \left(\frac{1}{m^p} x^p \right) \right)' V_{F,v}^p(x) = \sum_{p \in \mathcal{P}} \tilde{F}^p(x)' V_{F,v}^p(x).$$

Since $PE(F, v) \equiv RP(V_{F,v})$, that $\Omega(V_{F,v}) = PE(F, v)$ follows from Theorem 7.B.4. ■

In the case of the $\text{logit}(\eta)$ dynamic, the Lyapunov function from equation (7.4) takes the form

$$(7.5) \quad f^\eta(x) = f(x) - \eta \sum_{p \in \mathcal{P}} \sum_{i \in S} x_i^p \log \left(\frac{1}{m^p} x_i^p \right),$$

called the *logit potential function*. Theorem 7.1.6, combined with our results in Chapter 10, tells us that the logit potential function captures the *finite horizon behavior* of agents who play a potential game using the logit choice protocol. In Chapter 12, the logit potential function will be used to obtain a precise characterization of *infinite horizon behavior* in this setting.

Example 7.1.7. 123 Coordination rerevisited. Figure 7.1.1(vi) presents the phase diagram for the $\text{logit}(.5)$ dynamic in 123 Coordination (7.2). Here the contour plot is the logit potential function

$$f^\eta(x) = \frac{1}{2} \left((x_1)^2 + 2(x_2)^2 + (3(x_3)^2) \right) - .5 \sum_{i=1}^3 x_i \log x_i.$$

Because the noise level is rather high, this phase diagram looks very different than the others—in particular, it includes only three rest points (two stable and one unstable) rather than seven. Nevertheless, every solution trajectory ascends the relevant Lyapunov function f^η , ultimately converging to a perturbed equilibrium. §

7.1.2 Gradient Systems for Potential Games

Lemma 7.1.1 tells us that in potential games, any dynamic that satisfies condition (PC) must ascend the potential function f . We now turn to a more refined question: is there an evolutionary dynamic that ascends f in the fastest possible way?

A first answer to this question is suggested by Figure 7.1.1(ii): in 123 Coordination, solution trajectories of the projection dynamic,

$$(P) \quad \dot{x} = \Pi_{TX(x)}(F(x)),$$

cross the level sets of the potential function orthogonally. In fact, we have

Observation 7.1.8. *Let $F : X \rightarrow \mathbf{R}^n$ be a potential game with potential function $f : X \rightarrow \mathbf{R}$. On $\text{int}(X)$, the projection dynamic (P) is the gradient system for f :*

$$(7.6) \quad \dot{x} = \nabla f(x) \text{ on } \text{int}(X).$$

Surprisingly, there is an alternative answer to our question: it turns out that the replicator dynamic,

$$(R) \quad \dot{x}_i^p = x_i^p \hat{F}_i^p(x),$$

also defines a gradient system for the potential function f ; however, this is only true after we apply a clever change of variable. In addition to its inherent interest, this fact demonstrates a close connection between the replicator and projection dynamics; another such connection will be made in Section 7.2.1 below.

We restrict our analysis to the single population case. Define the set $\mathcal{X} = \{\chi \in \mathbf{R}_+^n : \sum_{i \in S} \chi_i^2 = 4\}$ to be the portion of the radius 2 sphere lying in the positive orthant. Our change of variable is given by the *Akin transformation* $H : \text{int}(\mathbf{R}_+^n) \rightarrow \text{int}(\mathbf{R}_+^n)$, where $H_i(x) = 2\sqrt{x_i}$. Evidently, H is a diffeomorphism that maps the interior of the simplex X onto the interior of the set \mathcal{X} . The transformation makes changes in component x_i look large when x_i itself is small.

Theorem 7.1.9 tells us that the replicator dynamic is a gradient dynamic on $\text{int}(X)$ after a change of variable that makes changes in the use of rare strategies look important relative to changes in the use of common ones. Intuitively, this reweighting accounts for the fact that under imitative dynamics, changes in the use of rare strategies are necessarily slow.

Theorem 7.1.9. *Let $F : X \rightarrow \mathbf{R}^n$ be a potential game with potential function $f : X \rightarrow \mathbf{R}$. Suppose*

we transport the replicator dynamic for F from $\text{int}(X)$ to $\text{int}(\mathcal{X})$ using the Akin transformation H . Then the resulting dynamic is the gradient dynamic for the transported potential function $\phi = f \circ H^{-1}$.

Proof. We prove Theorem 7.1.9 in two steps: first, we derive the transported version of the replicator dynamic; then we derive the gradient system for the transported version of the potential function, and show that it is the same dynamic on \mathcal{X} . The following notation will simplify our calculations: when $y \in \mathbf{R}_+^n$ and $a \in \mathbf{R}$, we let $[y^a] \in \mathbf{R}^n$ be the vector whose i th component is $(y_i)^a$.

We can express the replicator dynamic on X as

$$\dot{x} = R(x) = \text{diag}(x) (F(x) - \mathbf{1}x'F(x)) = (\text{diag}(x) - xx')F(x).$$

The transported version of this dynamic can be computed as

$$\dot{\chi} = \mathcal{R}(\chi) = DH(H^{-1}(\chi))R(H^{-1}(\chi)).$$

In words: given a state $\chi \in \mathcal{X}$, we first find the corresponding state $x = H^{-1}(\chi) \in X$ and direction of motion $R(x)$. Since $R(x)$ represents a displacement from state x , we transport it to \mathcal{X} by premultiplying it by $DH(x)$, the derivative of H evaluated at x .

Since $\chi = H(x) = 2[x^{1/2}]$, the derivative of H at x is given by $DH(x) = \text{diag}([x^{-1/2}])$. Using this fact, we derive a primitive expression for $\mathcal{R}(\chi)$ in terms of $x = H^{-1}(\chi) = \frac{1}{4}[\chi^2]$:

$$\begin{aligned} (7.7) \quad \dot{\chi} &= \mathcal{R}(\chi) \\ &= DH(x)R(x) \\ &= \text{diag}([x^{-1/2}]) (\text{diag}(x) - xx')F(x) \\ &= (\text{diag}([x^{1/2}]) - [x^{1/2}]x')F(x). \end{aligned}$$

Now, we derive the gradient system on \mathcal{X} generated by $\phi = f \circ H^{-1}$. To compute $\nabla\phi(\chi)$, we need to define an extension of ϕ to all of \mathbf{R}_+^n , compute its gradient, and then project the result onto the tangent space of \mathcal{X} at χ . The easiest way to proceed is to let $\tilde{f} : \text{int}(\mathbf{R}_+^n) \rightarrow \mathbf{R}$ be an arbitrary C^1 extension of f , and to define the extension $\tilde{\phi} : \text{int}(\mathbf{R}_+^n) \rightarrow \mathbf{R}$ by $\tilde{\phi} = \tilde{f} \circ H^{-1}$.

Since \mathcal{X} is a portion of a sphere centered at the origin, the tangent space of \mathcal{X} at χ is the subspace $T\mathcal{X}(\chi) = \{z \in \mathbf{R}^n : \chi'z = 0\}$. The orthogonal projection onto this set is represented by the $n \times n$ matrix

$$P_{T\mathcal{X}(\chi)} = I - \frac{1}{\chi'\chi} \chi\chi' = I - \frac{1}{4} \chi\chi' = I - [x^{1/2}][x^{1/2}]'.$$

Also, since $\Phi \nabla \tilde{f}(x) = \nabla f(x) = \Phi F(x)$ by construction, it follows that $\nabla \tilde{f}(x) = F(x) + c(x)\mathbf{1}$ for some scalar-valued function $c : X \rightarrow \mathbf{R}$.

Using the chain rule (Section 3.A.4), we compute that

$$\nabla \tilde{\phi}(\chi) = D(\tilde{f} \circ H^{-1})(\chi)' = (Df(H^{-1}(\chi))DH^{-1}(\chi))' = DH^{-1}(\chi)' \nabla \tilde{f}(x),$$

while applying the chain rule to the identity $H^{-1}(H(x)) \equiv x$ and then rearranging yields

$$DH^{-1}(\chi) = DH(x)^{-1}.$$


Marshaling these observations, we find that the gradient system on X generated by ϕ is

$$\begin{aligned} \dot{\chi} &= \nabla \phi(\chi) \\ &= P_{TX(\chi)} \nabla \tilde{\phi}(\chi) \\ &= P_{TX(\chi)} DH^{-1}(\chi)' \nabla \tilde{f}(x) \\ &= P_{TX(\chi)} (DH(x)^{-1})' (F(x) + c(x)\mathbf{1}) \\ &= (I - [x^{1/2}][x^{1/2}]') \text{diag}([x^{1/2}]) (F(x) + c(x)\mathbf{1}) \\ &= (\text{diag}([x^{1/2}]) - [x^{1/2}]x') (F(x) + c(x)\mathbf{1}) \\ &= (\text{diag}([x^{1/2}]) - [x^{1/2}]x') F(x). \end{aligned}$$


This agrees with equation (7.7), completing the proof of the theorem. ■

Example 7.1.10. 123 Coordination one last time. Figure 7.1.2 illustrates Theorem 7.1.9 by presenting phase diagrams of the transported replicator dynamic $\dot{\chi} = \mathcal{R}(\chi)$ for 123 Coordination (cf Example 7.1.3). These phase diagrams on X are drawn atop contour plots of the transported potential function $\phi(\chi) = (f \circ H^{-1})(\chi) = \frac{1}{32}((\chi_1)^4 + 2(\chi_2)^4 + 3(\chi_3)^4)$. According to Theorem 7.1.9, the solution trajectories of \mathcal{R} should cross the level sets of ϕ orthogonally.

Looking at Figure 7.1.2, we find that the crossings look orthogonal at the center of the figure, but not by the boundaries. This is an artifact of our drawing a portion of the sphere in \mathbf{R}^3 by projecting it orthogonally onto a sheet of paper. (For exactly the same reason, latitude and longitude lines in an orthographic projection of the Earth only appear to cross at right angles in the center of the projection, not on the left and right sides.) To check whether the crossings near a given state $\chi \in X$ are truly orthogonal, we must minimize the distortion of angles near χ by making χ the origin of the projection—that is, the point where the sphere touches the sheet of paper. In the phase diagrams in Figure 7.1.2, we



(i) origin = $H(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$



(ii) origin = $H(\frac{1}{7}, \frac{1}{7}, \frac{5}{7})$

Figure 7.1.2: The phase diagram of the transported replicator dynamic $\dot{\chi} = \mathcal{R}(\chi)$ for a coordination game. The pink dots represent the positions of the projection origins.

mark the projection origins with pink dots; evidently, the crossings are orthogonal near these points. §

7.2 Stable Games

Recall that the population game F is stable if it satisfies

$$(7.8) \quad (y - x)'(F(y) - F(x)) \leq 0 \text{ for all } x, y \in X.$$

When F is C^1 , this condition is equivalent to *self-defeating externalities*:

$$(7.9) \quad z'DF(x)z \leq 0 \text{ for all } z \in TX \text{ and } x \in X.$$

The set of Nash equilibria of a stable game is convex, and most often a singleton.

In general, uniqueness of equilibrium is not enough to ensure convergence of evolutionary dynamics. As we shall see in Chapter 9, there are many simple examples of games with a unique Nash equilibrium in which dynamics fail to converge. Nevertheless, we show in this section that under many evolutionary dynamics, the structure provided by self-defeating externalities is enough to ensure convergence. While fewer dynamics converge here than in potential games, convergence does obtain under all six fundamental dynamics.

Our convergence proofs for stable games again rely on the construction of Lyapunov functions, but here we will need to construct a distinct Lyapunov function for each dynamic we consider. It will be natural to write these Lyapunov functions so that their values fall over time: thus, we say that a C^1 function L is a (*decreasing*) *strict Lyapunov function* for the dynamic $\dot{x} = V_F(x)$ if $\dot{L}(x) \leq 0$ for all $x \in X$, with equality only at rest points of V_F . Apart from those for perturbed best response dynamics, the Lyapunov functions introduced below are also nonnegative, with zeros coinciding with the Nash equilibria of the underlying game F .

7.2.1 The Projection and Replicator Dynamics in Strictly Stable Games

To obtain convergence results for the projection and replicator dynamics, we must restrict attention to *strictly stable games*: that is, games in which condition (7.8) holds strictly for all $x, y \in X$. The Lyapunov functions for these dynamics are based on explicit notions of “distance” from the the game’s unique Nash equilibrium x^* .

Theorem 7.2.1 shows that under the projection dynamic, x^* is *globally asymptotically*

stable: all solution trajectories converge to x^* , and solutions that start near x^* never move too far away from x^* (see Appendix 7.A.2).

Theorem 7.2.1. *Let F be a strictly stable game with unique Nash equilibrium x^* , and let $\dot{x} = V_F(x)$ be the projection dynamic for F . Let the function $E_{x^*} : X \rightarrow \mathbf{R}_+$, defined by*

$$E_{x^*}(x) = |x - x^*|^2,$$

represent squared Euclidean distance from x^ . Then E_{x^*} is a strict Lyapunov function for V_F , and so x^* is globally asymptotically stable under V_F .*

Proof. Since F is a strictly stable game, its unique Nash equilibrium x^* is also its unique GESS:

$$(x - x^*)'F(x) < 0 \text{ for all } x \in X - \{x^*\}.$$

This fact and the Moreau Decomposition Theorem imply that

$$\begin{aligned} \dot{E}_{x^*}(x) &= \nabla E_{x^*}(x)' \dot{x} \\ &= 2(x - x^*)' \Pi_{TX(x)}(F(x)) \\ &= 2(x - x^*)'F(x) + 2(x^* - x)' \Pi_{NX(x)}(F(x)) \\ &\leq 2(x^* - x)' \Pi_{NX(x)}(F(x)) \\ &\leq 0, \end{aligned}$$

where the penultimate inequality is strict whenever $x \neq x^*$. Global asymptotic stability of $NE(F)$ then follows from Corollary 7.B.7. ■

Exercise 7.2.2. Let F be a stable game, and let x^* be a Nash equilibrium of F .

- (i) Show that x^* is Lyapunov stable under (P).
- (ii) Suppose that F is a null stable game (i.e., that $(y - x)'(F(y) - F(x)) = 0$ for all $x, y \in X$). Show that if $x^* \in \text{int}(X)$, then E_{x^*} defines a *constant of motion* for (P) on $\text{int}(X)$: the value of E_{x^*} is constant along interior portions of solution trajectories of (P).

Exercise 7.2.3. Show that if F is a C^1 stable game, then the squared speed of motion $L(x) = |\Phi F(x)|^2$ is a Lyapunov function for (P) on $\text{int}(X)$. Show that if F is null stable, then L defines a constant of motion for (P) on $\text{int}(X)$. (Notice that unlike that of E_{x^*} , the definition of L does not directly incorporate the Nash equilibrium x^* .)

Under the replicator dynamic (R), as under any imitative dynamic, strategies that are initially unused remain unused for all time. Therefore, if state x places no mass on a strategy in the support of the Nash equilibrium x^* , the solution to (R) starting from x cannot converge to x^* . Thus, in stating our convergence result for the replicator dynamic, we need to be careful to specify the set of states from which convergence to equilibrium occurs.

With this motivation, let $S^p(x^p) = \{i \in S^p : x_i^p > 0\}$ denote the support of x^p . Then $X_{y^p}^p = \{x^p \in X^p : S^p(y^p) \subseteq S^p(x^p)\}$ is the set of states in X^p whose supports contain the support of y^p , and $X_y = \prod_{p \in \mathcal{P}} X_{y^p}^p$ is the set of states in X whose supports contain the support of y . To construct our Lyapunov function, we introduce the function $h_{y^p}^p : X_{y^p}^p \rightarrow \mathbf{R}$, defined by

$$h_{y^p}^p(x^p) = \sum_{i \in S^p(y^p)} y_i^p \log \frac{y_i^p}{x_i^p}.$$

If population p is of unit mass, so that y^p and x^p are probability distributions, $h_{y^p}^p(x^p)$ is known as the *relative entropy* of y^p given x^p .

Theorem 7.2.4. *Let F be a strictly stable game with unique Nash equilibrium x^* , and let $\dot{x} = V_F(x)$ be the replicator dynamic for F . Define the function $H_{x^*} : X_{x^*} \rightarrow \mathbf{R}_+$ by*

$$H_{x^*}(x) = \sum_{p \in \mathcal{P}} h_{(x^*)^p}^p(x^p).$$

Then $H_{x^}^{-1}(0) = \{x^*\}$, and $H_{x^*}(x)$ approaches infinity whenever x approaches $X - X_{x^*}$. Moreover, $\dot{H}_{x^*}(x) \leq 0$, with equality only when $x = x^*$. Therefore, x^* is globally asymptotically stable with respect to X_{x^*} .*

Proof. ($p = 1$) To see that H_{x^*} is nonnegative with $H_{x^*}^{-1}(0) = \{x^*\}$, observe that by Jensen's inequality,

$$-H_{x^*}(x) = \sum_{i \in S(x^*)} x_i^* \log \frac{x_i}{x_i^*} \leq \log \left(\sum_{i \in S(x^*)} x_i^* \cdot \frac{x_i}{x_i^*} \right) = \log \left(\sum_{i \in S(x^*)} x_i \right) \leq \log 1 = 0,$$

with equality if and only if $x = x^*$. The second claim is immediate. For the third claim, note that since F is strictly stable, x^* is a GESS, so for all $x \in X_{x^*}$ we have that

$$\begin{aligned} \dot{H}_{x^*}(x) &= \nabla H_{x^*}(x)' \dot{x} \\ &= - \sum_{i \in S(x^*)} \frac{x_i^*}{x_i} \cdot x_i \hat{F}_i(x) \end{aligned}$$

$$\begin{aligned}
&= - \sum_{i \in S} x_i^* \hat{F}_i(x) \\
&= -(x^*)' (F(x) - \mathbf{1} x' F(x)) \\
&= -(x^* - x)' F(x) \\
&\leq 0,
\end{aligned}$$

where the inequality binds precisely when $x = x^*$. The conclusions about stability then follow from Theorems 7.B.2 and 7.B.4. ■

Exercise 7.2.5. Let F be a stable game, and let x^* be a Nash equilibrium of F .

- (i) Show that x^* is Lyapunov stable under (R).
- (ii) Show that if F is a null stable game and $x^* \in \text{int}(X)$, then H_{x^*} defines a constant of motion for (R) on $\text{int}(X)$.

7.2.2 Integrable Target Dynamics

Of our six fundamental dynamics, three of them—the BNN, best response, and logit dynamics, can be expressed as target dynamics of the form

$$\tau^p(\pi^p, x^p) = \tau^p(\hat{\pi}^p),$$

under which conditional switch rates only depend on on the vector of excess payoffs $\hat{\pi}^p = \pi^p - \frac{1}{m^p} \mathbf{1}(x^p)' \pi^p$. This is obviously true of the BNN dynamic. For the other two cases, note that shifting all components of the payoff vector by the same constant has no effect on either exact or perturbed best responses: in particular, the definitions (6.2) and (6.12) of the maximizer correspondence $M^p : \mathbf{R}^{n^p} \Rightarrow \Delta^p$ and the perturbed maximizer function $\tilde{M}^p : \mathbf{R}^{n^p} \rightarrow \Delta^p$ satisfy $M^p(\hat{\pi}^p) = M^p(\pi^p)$ and $\tilde{M}^p(\hat{\pi}^p) = \tilde{M}^p(\pi^p)$.

In this section, we show that these three dynamics converge to equilibrium in all stable games, as do all close enough relatives of these dynamics. Unlike in the context of potential games, monotonicity properties alone are not enough to ensure that a dynamic converges: in addition, integrability of the revision protocol plays a key role in establishing convergence results.

To begin, we provide an example to illustrate that monotonicity properties alone do not ensure convergence of target dynamics in stable games.

Example 7.2.6. Cycling in good RPS. Fix $\varepsilon > 0$, and let $g^\varepsilon : \mathbf{R} \rightarrow \mathbf{R}$ be a continuous decreasing function that equals 1 on $(-\infty, 0]$, equals ε^2 on $[\varepsilon, \infty)$, and is linear on $[0, \varepsilon]$. Then define

the revision protocol τ for Rock-Paper-Scissors games by

$$(7.10) \quad \begin{pmatrix} \tau_R(\hat{\pi}) \\ \tau_P(\hat{\pi}) \\ \tau_S(\hat{\pi}) \end{pmatrix} = \begin{pmatrix} [\hat{\pi}_R]_+ g^\varepsilon(\hat{\pi}_S) \\ [\hat{\pi}_P]_+ g^\varepsilon(\hat{\pi}_R) \\ [\hat{\pi}_S]_+ g^\varepsilon(\hat{\pi}_P) \end{pmatrix}.$$

Under this protocol, the weight placed on a strategy is proportional to positive part of the strategy's excess payoff, as in the protocol for the BNN dynamic; however, this weight is only of order ε^2 if the strategy it beats in RPS has an excess payoff greater than ε .

It is easy to verify that protocol (7.10) satisfies acuteness (5.22):

$$\tau(\hat{\pi})' \hat{\pi} = [\hat{\pi}_R]_+^2 g^\varepsilon(\hat{\pi}_S) + [\hat{\pi}_P]_+^2 g^\varepsilon(\hat{\pi}_R) + [\hat{\pi}_S]_+^2 g^\varepsilon(\hat{\pi}_P),$$

which is positive when $\hat{\pi} \in \text{int}(\mathbf{R}_*^n)$. Thus, the target dynamic induced by τ is an excess payoff dynamic. In Figure 7.2.1 we presents a phase diagram for this dynamic in the good RPS game

$$F(x) = Ax = \begin{pmatrix} 0 & -2 & 3 \\ 3 & 0 & -2 \\ -2 & 3 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \end{pmatrix}.$$

Evidently, solutions from many initial conditions lead to a limit cycle. §

To explain why cycling occurs in the example above, we review some ideas about the geometry of stable games and target dynamics. By Theorem 3.3.16, every Nash equilibrium x^* of a stable game is a GNSS. Geometrically, this means that at every nonequilibrium state x , the projected payoff vector $\Phi F(x)$ forms an acute or right angle with the line segment from x back to x^* (Figures 3.3.3, 3.3.5, and 3.3.6). Meanwhile, our monotonicity condition for dynamics, positive correlation (PC), requires that away from equilibrium, each vector of motion $V_F(x)$ forms an acute angle with the projected payoff vector $\Phi F(x)$ (Figures 5.2.1 and 5.2.2). Combining these observations, we conclude that if the law of motion $\dot{x} = V_F(x)$ tends to deviate from the projected payoffs ΦF in “outward” directions—that is, in directions heading away from equilibrium—then cycling will occur (compare Figure 3.3.6 with Figure 7.2.1). On the other hand, if the deviations of V_F from ΦF tend to be “inward”, then solutions should converge to equilibrium.

By this logic, we should be able to guarantee convergence of target dynamics in stable games by ensuring that the deviations of V_F from ΦF are toward the equilibrium, at least in some average sense. To accomplish this, we introduce an additional condition for revision




Figure 7.2.1: An excess payoff dynamic in good RPS ($w = 3, l = 2$).

protocols: *integrability*.

$$(7.11) \quad \text{There exists a } C^1 \text{ function } \gamma^p : \mathbf{R}^{n^p} \rightarrow \mathbf{R} \text{ such that } \tau^p \equiv \nabla \gamma^p.$$

We call the functions γ^p introduced in this condition *revision potentials*.

To give this condition a behavioral interpretation, it is useful to compare it to *separability*:

$$(7.12) \quad \tau_i^p(\hat{\pi}^p) \text{ is independent of } \hat{\pi}_{-i}^p.$$

The latter condition is stronger than the former: if τ^p satisfies (7.12), then it satisfies (7.11) with

$$(7.13) \quad \gamma^p(\hat{\pi}^p) = \sum_{i \in S^p} \int_0^{\hat{\pi}_i^p} \tau_i^p(s) ds.$$

In Example 7.2.6, the protocol (7.10) that generated cycling has the following noteworthy feature: the weights agents place on each strategy depend systematically on the payoffs of the next strategy in the best response cycle. Building on this motivation, one can obtain a game-theoretic interpretation of integrability. Roughly speaking, integrability (7.11) is equivalent to a requirement that in expectation, learning the weight placed on strategy j does not convey information about other strategies' excess payoffs. It thus

generalizes separability (7.12), which requires that learning the weight placed on strategy j conveys no information at all about other strategies' excess payoffs (see the Notes).

Before turning to our convergence theorems, we address a missing step in the motivating argument above: how does integrability ensure that the law of motion V_F tends to deviate from the projected payoffs ΦF in the direction of equilibrium? To make this link, let us recall a characterization of integrability from Section 3.A.9: the map $\tau : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is integrable if and only if its line integral over any piecewise smooth closed curve $C \subset \mathbf{R}^n$ evaluates to zero:

$$(7.14) \quad \oint_C \tau(\hat{\pi}) \cdot d\hat{\pi} = 0.$$

Example 7.2.7. Let the population game F be generated by matching in standard RPS:

$$F(x) = Ax = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \end{pmatrix}.$$

The unique Nash equilibrium of F is the GNSS $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Game F has the convenient property that at each state $x \in X$, the payoff vector $F(x)$, the projected payoff vector $\Phi F(x)$, and the excess payoff vector $\hat{F}(x)$ are all the same, a fact that will simplify the notation in the argument to follow.

Since F is null stable, we know that at each state $x \neq x^*$, the payoff vector $F(x)$ is orthogonal to the vector $x^* - x$. In Figure 3.3.6, these payoff vectors point counterclockwise relative to x^* . Since positive correlation (PC) requires that the direction of motion $V_F(x)$ form an acute angle with $F(x)$, dynamics satisfying (PC) also travel counterclockwise around the equilibrium.

To address whether the deviations of V_F from F tend to be inward or outward, let $C \subset X$ be a circle of radius $c \in (0, \frac{1}{\sqrt{6}}]$ centered at the equilibrium x^* . This circle is parameterized by the function $\xi : [0, 2\pi] \rightarrow X$, where

$$(7.15) \quad \xi_\alpha = \frac{c}{\sqrt{6}} \begin{pmatrix} -2 \sin \alpha \\ \sqrt{3} \cos \alpha + \sin \alpha \\ -\sqrt{3} \cos \alpha + \sin \alpha \end{pmatrix} + x^*.$$

Here α is the counterclockwise angle between the vector $\xi_\alpha - x^*$ and a rightward horizontal vector (see Figure 7.2.2).

Since state ξ_α lies on the circle C , the vector $x^* - \xi_\alpha$ can be drawn as a radius of C ; thus,




Figure 7.2.2: Integrability and inward motion of target dynamics in standard RPS.

the payoff vector $\pi_\alpha \equiv F(\xi_\alpha)$, which is orthogonal to $x^* - \xi_\alpha$, must be tangent to C at ξ_α , as shown in Figure 7.2.2. This observation is easy to verify analytically:

$$(7.16) \quad \pi_\alpha = F(\xi_\alpha) = \frac{c}{\sqrt{6}} \begin{pmatrix} -2\sqrt{3}\cos\alpha \\ -3\sin\alpha + \sqrt{3}\cos\alpha \\ 3\sin\alpha + \sqrt{3}\cos\alpha \end{pmatrix} = \sqrt{3} \frac{d}{d\alpha} \xi_\alpha.$$

If we differentiate both sides of identity (7.16) with respect to the angle α , and note that $\frac{d^2}{(d\alpha)^2} \xi_\alpha = -(\xi_\alpha - x^*)$, we can link the rate of change of the payoff vector $\pi_\alpha = F(\xi_\alpha)$ to the displacement of state ξ_α from x^* :

$$(7.17) \quad \frac{d}{d\alpha} \pi_\alpha = \sqrt{3} \frac{d^2}{(d\alpha)^2} \xi_\alpha = -\sqrt{3}(\xi_\alpha - x^*).$$

Now introduce an acute, integrable revision protocol τ . By combining integrability condition (7.14) with equation (7.17), we obtain

$$(7.18) \quad 0 = \oint_C \tau(\pi) \cdot d\pi \equiv \int_0^{2\pi} \tau(\pi_\alpha)' \left(\frac{d}{d\alpha} \pi_\alpha \right) d\alpha = -\sqrt{3} \int_0^{2\pi} \tau(\pi_\alpha)' (\xi_\alpha - x^*) d\alpha.$$

Let us write $\lambda(\pi) = \sum_{i \in S} \tau_i(\pi)$ and $\sigma_i(\pi) = \frac{\tau_i(\pi)}{\lambda(\pi)}$ to express the dynamic in target form. Then because $\xi_\alpha - x^* \in TX$ is orthogonal to $x^* = \frac{1}{3}\mathbf{1}$, we can conclude from equation (7.18) that

$$(7.19) \quad \int_0^{2\pi} \lambda(F(\xi_\alpha)) (\sigma(F(\xi_\alpha)) - x^*)' (\xi_\alpha - x^*) d\alpha = 0.$$

Equation (7.19) is a form of the requirement described at the start of this section: it asks that at states on the circle C , the vector of motion under the target dynamic

$$(7.20) \quad \dot{x} = V_F(x) = \lambda(F(x)) (\sigma(F(x)) - x).$$

typically deviates from the payoff vector $F(x)$ in an inward direction—that is, in the direction of the equilibrium x^* .

To reach this interpretation of equation (7.19), note first that if the target state $\sigma(F(\xi_\alpha))$ lies on or even near line $L^\perp(\xi_\alpha)$, then motion from ξ_α toward $\sigma(F(\xi_\alpha))$ is initially inward, as shown in Figure 7.2.2. (Of course, target state $\sigma(F(\xi_\alpha))$ lies above $L(\xi_\alpha)$ by virtue of positive correlation (PC).) Now, the integrand in (7.19) contains the inner product of the vectors $\sigma(F(\xi_\alpha)) - x^*$ and $\xi_\alpha - x^*$. This inner product is zero precisely when the two vectors are orthogonal, or, equivalently, when target state $\sigma(F(\xi_\alpha))$ lies on $L^\perp(\xi_\alpha)$. While equation (7.19) does not require the two vectors to be orthogonal, it asks that this be true on average, where the average is taken over states $\xi_\alpha \in C$, and weighted by the rates $\lambda(F(\xi_\alpha))$ at which ξ_α approaches $\sigma(F(\xi_\alpha))$. Thus, in the presence of acuteness, integrability implies that on average, the dynamic (7.20) tends to point inward, toward the equilibrium x^* . §

The foregoing arguments suggest that together, monotonicity and integrability are enough to ensure global convergence of target dynamics in stable games. We now develop this intuition into formal results by constructing suitable Lyapunov functions.

As a point of comparison, recall from Section 7.1.1 that in the case of dynamics for potential games, monotonicity conditions alone are sufficient to prove global convergence results: as the game's potential function serves as a Lyapunov function for any dynamic satisfying positive correlation (PC). Unlike potential games, stable games do not come equipped with candidate Lyapunov functions. But if the revision protocol agents follow is integrable, then the revision potential of this protocol provides a building block for constructing a suitable Lyapunov function. Evidently, this Lyapunov function will vary with the dynamic under study, even when the game under consideration is fixed.

Our first result concerns *integrable excess payoff dynamics*: that is, target dynamics whose protocols τ^p are Lipschitz continuous, acute (5.22), and integrable (7.11). The prototype

for this class is the BNN dynamic: its protocol $\tau_i^p(\hat{\pi}^p) = [\hat{\pi}_i^p]_+$ is not only acute and integrable, but also separable (7.12), and so admits potential function $\gamma^p(\hat{\pi}^p) = \frac{1}{2} \sum_{i \in S^p} [\hat{\pi}_i^p]_+^2$ (cf equation (7.13)).

Theorem 7.2.8. *Let F be a C^1 stable game, and let $\dot{x} = V_F(x)$ be the integrable excess payoff dynamic for F based on revision protocols τ^p with revision potentials γ^p . Define the C^1 function $\Gamma : X \rightarrow \mathbf{R}$ by*

$$\Gamma(x) = \sum_{p \in \mathcal{P}} m^p \gamma^p(\hat{F}^p(x)).$$

Then Γ is a strict Lyapunov function for V_F , and $NE(F)$ is globally attracting. In addition, if F admits a unique Nash equilibrium, or if the protocols τ^p also satisfy separability (7.12), then we can choose the functions γ^p in such a way that Γ is nonnegative with $\Gamma^{-1}(0) = NE(F)$, implying that $NE(F)$ is globally asymptotically stable.

For future reference, observe that the value of the Lyapunov function Γ at state x is the (m^p -weighted) sum of the values of the revision potentials γ^p evaluated at the excess payoff vectors $\hat{F}^p(x)$.

The conditions introduced in the last sentence of the theorem are needed to ensure that the Lyapunov function Γ is constant on the set $NE(F)$ of Nash equilibria. Were this not the case, the set $NE(F)$ could be globally attracting without being Lyapunov stable—see Example 7.B.3.

The proof of Theorem 7.2.8 and those to come make heavy use of multivariate product and chain rules, which we review in Section 3.A.4.

Proof of Theorem 7.2.8. ($p = 1$) Recall that the excess payoff vector $\hat{F}(x)$ is equal to $F(x) - \mathbf{1}\bar{F}(x)$, where $\bar{F}(x) = x'F(x)$ is the population's average payoff. By the product rule, the derivative of \bar{F} is

$$D\bar{F}(x) = x'DF(x) + F(x)'.$$

Therefore, the derivative matrix for the excess payoff function $\hat{F}(x) = F(x) - \mathbf{1}\bar{F}(x)$ is

$$\begin{aligned} D\hat{F}(x) &= D(F(x) - \mathbf{1}\bar{F}(x)) \\ &= DF(x) - \mathbf{1}D\bar{F}(x) \\ (7.21) \quad &= DF(x) - \mathbf{1}(x'DF(x) + F(x)'). \end{aligned}$$

Using (7.21) and the chain rule, we can compute the time derivative of Γ :

$$\begin{aligned}
\dot{\Gamma}(x) &= \nabla\Gamma(x)' \dot{x} \\
&= \nabla\gamma(\hat{F}(x))' D\hat{F}(x) \dot{x} \\
&= \tau(\hat{F}(x))' \left(DF(x) - \mathbf{1}(x' DF(x) + F(x)') \right) \dot{x} \\
&= \left(\tau(\hat{F}(x)) - \tau(\hat{F}(x))' \mathbf{1} x \right)' DF(x) \dot{x} - \tau(\hat{F}(x))' \mathbf{1} F(x)' \dot{x} \\
&= \dot{x}' DF(x) \dot{x} - (\tau(\hat{F}(x))' \mathbf{1})(F(x)' \dot{x}) \\
&\leq 0,
\end{aligned}$$

where the inequality follows from the facts that F is stable and V_F satisfies positive correlation (PC).

We now show that this inequality binds precisely on the set $NE(F)$. To begin, note that if $x \in RP(V_F)$ (i.e., if $\dot{x} = \mathbf{0}$), then $\dot{\Gamma}(x) = 0$. On the other hand, if $x \notin RP(V_F)$, then $\tau(\hat{F}(x))' \mathbf{1} > 0$ and $F(x)' \dot{x} > 0$ (by condition (PC)), implying that $\dot{\Gamma}(x) < 0$. Since $NE(F) = RP(V_F)$, the claim is proved. That $NE(F)$ is globally attracting then follows from Theorem 7.B.4.

If F admits a unique Nash equilibrium x^* , then the foregoing argument implies that x^* is the unique minimizer of Γ : since the value of Γ is nonincreasing over time, a solution starting from a state x with $\Gamma(x) < \Gamma(x^*)$ could not converge to x^* , contradicting that x^* is globally attracting. Thus, after normalizing by an additive constant, we find that Γ is nonnegative with $\Gamma^{-1}(0) = \{x^*\}$, so the global asymptotic stability of x^* follows from Corollary 7.B.7.

If instead τ satisfies separability (7.12), we can define the revision potential γ as in equation (7.13). It then follows from Exercise 5.5.7 that Γ is nonnegative, with $\Gamma(x) = 0$ if and only if $\hat{F}(x) \in \text{bd}(\mathbf{R}_*^n)$. Thus, Lemma 5.5.4 implies that $\Gamma(x) = 0$ if and only if $x \in NE(F)$, and so the global asymptotic stability of $NE(F)$ again follows from Corollary 7.B.7. ■

Next we consider the best response dynamic, which we here express by applying the maximizer correspondence

$$M^p(\hat{\pi}^p) = \operatorname{argmax}_{y^p \in \Delta^p} (y^p)' \hat{\pi}^p$$

to the vector of excess payoffs, yielding the exact target dynamic

$$(BR) \quad \dot{x}^p \in m^p M^p(\hat{F}^p(x)) - x^p.$$

Following the previous logic, we can assess the possibilities for convergence in stable

games by checking monotonicity and integrability. Monotonicity was established in Theorem 6.1.9, which showed that (BR) satisfies an analogue of positive correlation (PC) appropriate for differential inclusions. For integrability, one can argue that the protocol M^p , despite being multivalued, is integrable in a suitably defined sense, with its “potential function” being given by the *maximum function*

$$\mu^p(\pi^p) = \max_{y^p \in \Delta^p} (y^p)' \pi^p = \max_{i \in S^p} \pi_i^p.$$

Note that if the payoff vector π^p , and hence the excess payoff vector $\hat{\pi}^p$, have a unique maximizing component $i \in S^p$, then the gradient of μ^p at $\hat{\pi}^p$ is the standard basis vector e_i^p . But this vector corresponds to the unique mixed best response to $\hat{\pi}^p$, and so

$$\nabla \mu^p(\hat{\pi}^p) = e_i^p = M^p(\hat{\pi}^p).$$

One can account for multiple optimal components using a broader notion of differentiation: for all $\hat{\pi}^p \in \mathbf{R}^n$, $M^p(\hat{\pi}^p)$ is the subdifferential of the convex function μ^p at $\hat{\pi}^p$ (see the Notes).

Having verified monotonicity and integrability, we again construct our candidate Lyapunov function by plugging the excess payoff vectors into the revision potentials μ^p . The resulting function G is very simple: it measures the difference between the payoffs agents could obtain by choosing optimal strategies and their actual aggregate payoffs.

Theorem 7.2.9. *Let F be a C^1 stable game, and let $\dot{x} \in V_F(x)$ be the best response dynamic for F . Define the Lipschitz continuous function $G : X \rightarrow \mathbf{R}_+$ by*

$$G(x) = \max_{y \in X} (y - x)' F(x) = \max_{i \in S} \hat{F}_i(x).$$

Then $G^{-1}(0) = NE(F)$. Moreover, if $\{x_t\}_{t \geq 0}$ is a solution to V_F , then for almost all $t \geq 0$ we have that $\dot{G}(x_t) \leq -G(x_t)$, and so $NE(F)$ is globally asymptotically stable under V_F .

Proof. ($p = 1$) That $G^{-1}(0) = NE(F)$ follows from Lemma 5.5.4. To prove the second claim, let $\{x_t\}_{t \geq 0}$ be a solution to V_F , and let $S^*(t) \subseteq S$ be the set of pure best responses to state x_t . Since $\{x_t\}_{t \geq 0}$ is Lipschitz continuous, the map $t \mapsto \hat{F}_i(x_t)$ is also Lipschitz continuous for each strategy $i \in S$. Thus, since $G(x) = \max_{y \in X} (y - x)' F(x) = \max_{i \in S} \hat{F}_i(x)$, it follows from Danskin’s Envelope Theorem (see the Notes) that the map $t \mapsto G(x_t)$ is Lipschitz continuous, and that at almost all $t \in [0, \infty)$,

$$(7.22) \quad \dot{G}(x_t) \equiv \frac{d}{dt} \max_{i \in S} \hat{F}_i(x_t) = \frac{d}{dt} \hat{F}_{i^*}(x_t) \quad \text{for all } i^* \in S^*(t).$$

Applying equation (7.21), we find that for t satisfying equation (7.22) and at which \dot{x}_t exists, we have that

$$\begin{aligned}
\dot{G}(x_t) &= \frac{d}{dt} \hat{F}_{i^*}(x_t) && \text{for all } i^* \in S^*(t) \\
&= \left(e'_{i^*} DF(x_t) - x'_t DF(x_t) - F(x_t)' \right) \dot{x}_t && \text{for all } i^* \in S^*(t) \\
&= (y^* - x_t)' DF(x_t) \dot{x}_t - F(x_t)' \dot{x}_t && \text{for all } y^* \in \operatorname{argmax}_{y \in \Delta} y' \hat{F}(x_t) \\
&= \dot{x}'_t DF(x_t) \dot{x}_t - F(x_t)' \dot{x}_t \\
&\leq -F(x_t)' \dot{x}_t \\
&= -\max_{y \in X} F(x_t)'(y - x_t) \\
&= -G(x_t),
\end{aligned}$$

where the inequality follows from the fact that F is a stable game. (Note that the equality of the third to last and last expressions is also implied by Theorem 6.1.9.) The global asymptotic stability of $NE(F)$ then follows from Theorems 7.B.2 and 7.B.6. ■

Finally, we consider convergence under perturbed best response dynamics. These are exact target dynamics of the form

$$\dot{x}^p = m^p \tilde{M}^p(\hat{F}^p(x)) - x^p;$$

here, the target protocol is the perturbed maximizer function

$$\tilde{M}^p(\hat{\pi}^p) = \operatorname{argmax}_{y^p \in \operatorname{int}(\Delta^p)} (y^p)' \hat{\pi}^p - v^p(y^p),$$

where $v^p : \operatorname{int}(\Delta^p) \rightarrow \mathbf{R}$ is an admissible deterministic perturbation (see Section 6.2.2).

Once again, we verify the two conditions that underlie convergence. Theorem 6.2.13 showed that all perturbed best response dynamics satisfy *virtual positive correlation* (6.17), establishing the required monotonicity. As for integrability, Observation 6.C.3 showed that the protocol \tilde{M}^p is integrable; its revision potential,

$$(7.23) \quad \tilde{\mu}^p(\pi^p) = \max_{y^p \in \operatorname{int}(\Delta^p)} (y^p)' \pi^p - v^p(y^p),$$

is the *perturbed maximum function* induced by v^p . Now, mimicking Theorem 7.2.8, we construct our Lyapunov function by composing the revision potentials $\tilde{\mu}^p$ with the excess payoff functions \hat{F}^p .

Theorem 7.2.10. Let F be a C^1 stable game, and let $\dot{x} = V_{F,v}(x)$ be the perturbed best response dynamic for F generated by the admissible deterministic perturbations v . Define the function $\tilde{G} : \text{int}(X) \rightarrow \mathbf{R}_+$ by

$$\tilde{G}(x) = \sum_{p \in \mathcal{P}} m^p \left(\tilde{\mu}^p(\hat{F}^p(x)) + v^p\left(\frac{1}{m^p}x^p\right) \right),$$

Then $G^{-1}(0) = PE(F, v)$, and this set is a singleton. Moreover, \tilde{G} is a strict Lyapunov function for $V_{F,v}$, and so $PE(F, v)$ is globally asymptotically stable under $V_{F,v}$.

Proof. ($p = 1$) As in Section 6.2, let $\tilde{F}(x) = F(x) - \nabla v(x)$ be the virtual payoff function generated by (F, v) . Then

$$x \in PE(F, v) \Leftrightarrow \Phi \tilde{F}(x) = \mathbf{0} \Leftrightarrow x = \operatorname{argmax}_{y \in \text{int}(\Delta)} (y'F(x) - v(y)) \Leftrightarrow \tilde{G}(x) = 0.$$

To prove that \tilde{G} is a strict Lyapunov function, recall from Observation 6.C.3 that the perturbed maximum function $\tilde{\mu}$ defined in equation (7.23) is a potential function for the perturbed maximizer function \tilde{M} : that is, $\nabla \tilde{\mu} \equiv \tilde{M}$. Therefore, since F is stable, virtual positive correlation (6.17) implies that

$$\begin{aligned} \dot{\tilde{G}}(x) &= \frac{d}{dt} \left(\tilde{\mu}(\hat{F}(x)) + v(x) \right) \\ &= \frac{d}{dt} \left(\tilde{\mu}(F(x)) - (x'F(x) - v(x)) \right) \\ &= \tilde{M}(F(x))' DF(x) \dot{x} - (x' DF(x) \dot{x} + \dot{x}' F(x) - \dot{x}' \nabla v(x)) \\ &= (\tilde{M}(F(x)) - x)' DF(x) \dot{x} - \dot{x}' (F(x) - \nabla v(x)) \\ &= \dot{x}' DF(x) \dot{x} - \dot{x}' \tilde{F}(x) \\ &\leq 0, \end{aligned}$$

with equality if and only if x is a rest point. But $RP(V_{F,v}) = PE(F, v)$ by definition, so Corollary 7.B.7 implies that $PE(F, v)$ is globally asymptotically stable.

Finally, we prove that $PE(F, v)$ is a singleton. Let

$$\phi_{x,h}(t) = h' \tilde{F}(x + th)$$

for all $x \in X$, $h \in TX - \{0\}$, and $t \in \mathbf{R}$ such that $x + th \in \text{int}(X)$. Since F is stable and $D^2v(x + th)$ is positive definite with respect to $TX \times TX$, we have that

$$(7.24) \quad \phi'_{x,h}(t) = h' D\tilde{F}(x + th) h = h' DF(x + th) h - h' D^2v((x + th)) h < 0,$$

and so $\phi_{x,h}(t)$ is decreasing in t . Moreover,

$$(7.25) \quad x \in PE(F, v) \Leftrightarrow \tilde{F}(x) \text{ is a constant vector} \Leftrightarrow \phi_{x,h}(0) = 0 \text{ for all } h \in TX - \{0\}.$$

Now let $x \in PE(F, v)$ and $y \in X - \{x\}$. Then $y = x + t_y h_y$ for some $t_y > 0$ and $h_y \in TX - \{0\}$. Statements (7.24) and (7.25) imply that

$$\phi_{y,h_y}(0) = h'_y \tilde{F}(y) = h'_y \tilde{F}(x + t_y h_y) = \phi_{x,h_y}(t_y) < \phi_{x,h_y}(0) = 0.$$

Therefore, statement (7.25) implies that $y \notin PE(F, v)$, and hence that $PE(F, v) = \{x\}$. ■

7.2.3 Impartial Pairwise Comparison Dynamics

In Section 5.6, we defined pairwise comparison dynamics by considering Lipschitz continuous revision protocols ρ^p that only condition on payoffs and that are sign preserving:

$$\text{sgn}(\rho_{ij}^p(\pi^p)) = \text{sgn}([\pi_j^p - \pi_i^p]_+) \quad \text{for all } i, j \in S^p \text{ and } p \in \mathcal{P}.$$

To obtain a general convergence result for stable games, we require an additional condition called *impartiality*:

$$(7.26) \quad \rho_{ij}^p(\pi^p) = \phi_j^p(\pi_j^p - \pi_i^p) \quad \text{for some functions } \phi_j^p : \mathbf{R} \rightarrow \mathbf{R}_+.$$

Combining this restriction with mean dynamic equation (M), we see that impartial pairwise comparison dynamics take the form

$$\dot{x}_i^p = \sum_{j \in S^p} x_j^p \phi_i^p(F_j^p(x) - F_i^p(x)) - x_i^p \sum_{j \in S^p} \phi_j^p(F_j^p(x) - F_i^p(x)).$$

Under impartiality (7.26), the function of the payoff difference $\pi_j^p - \pi_i^p$ that describes the conditional switch rate from i to j does not depend on an agent's current strategy i . This restriction introduces at least a superficial connection with the target dynamics studied in Section 7.2.2, as both restrict the dependence of agents' decisions on their current choices of strategy.

Theorem 7.2.11 shows that together, sign preservation and impartiality ensure global convergence to Nash equilibrium in stable games.

Theorem 7.2.11. Let F be a C^1 stable game, and let $\dot{x} = V_F(x)$ be an impartial pairwise comparison dynamic for F . Define the Lipschitz continuous function $\Psi : X \rightarrow \mathbf{R}_+$ by

$$\Psi(x) = \sum_{p \in \mathcal{P}} \sum_{i \in S^p} \sum_{j \in S^p} x_i^p \psi_j^p(F_j^p(x) - F_i^p(x)), \quad \text{where } \psi_k^p(d) = \int_0^d \phi_k^p(s) ds$$

is the definite integral of ϕ_k^p . Then $\Psi^{-1}(0) = NE(F)$. Moreover, $\dot{\Psi}(x) \leq 0$ for all $x \in X$, with equality if and only if $x \in NE(F)$, and so $NE(F)$ is globally asymptotically stable.

To understand the role played by impartiality (7.26), recall the general formula for the mean dynamic:

$$(M) \quad \dot{x}_i^p = \sum_{j \in S^p} x_j^p \rho_{ji}^p(F^p(x), x^p) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(F^p(x), x^p).$$

According to the second term of this expression, the rate of outflow from strategy i is $x_i^p \sum_{k \in S^p} \rho_{ik}^p$; thus, the *percentage* rate of outflow from i , $\sum_{k \in S^p} \rho_{ik}^p$, varies with i . It follows that strategies with high payoffs can nevertheless have high percentage outflow rates: even if $\pi_i^p > \pi_j^p$, one can still have $\rho_{ik}^p > \rho_{jk}^p$ for $k \neq i, j$. Having good strategies lose players more quickly than bad strategies is an obvious impediment to convergence to Nash equilibrium.

Impartiality (7.26) places controls on these percentage outflow rates. If the conditional switch rates ϕ_j^p are monotone in payoffs, then condition (7.26) ensures that better strategies have lower percentage outflow rates. If the conditional switch rates are not monotone, but merely sign-preserving, condition (7.26) still implies that the *integrated* conditional switch rates ψ_k^p are ordered by payoffs. According to the analysis below, this control is enough to ensure convergence of pairwise comparison dynamics to Nash equilibrium in stable games.

Proof. ($p = 1$) The first claim is proved as follows:

$$\begin{aligned} \Psi(x) = 0 &\Leftrightarrow [x_i = 0 \text{ or } \psi_j(F_j(x) - F_i(x)) = 0] \text{ for all } i, j \in S \\ &\Leftrightarrow [x_i = 0 \text{ or } F_i(x) \geq F_j(x)] \text{ for all } i, j \in S \\ &\Leftrightarrow [x_i = 0 \text{ or } F_i(x) \geq \max_{j \in S} F_j(x)] \text{ for all } i, j \in S \\ &\Leftrightarrow x \in NE(F). \end{aligned}$$

To begin the proof of the second claim, we compute the partial derivatives of Ψ :

$$\frac{\partial \Psi}{\partial x_l}(x) = \sum_{i \in S} \sum_{j \in S} x_i \rho_{ij} \left(\frac{\partial F_j}{\partial x_l}(x) - \frac{\partial F_i}{\partial x_l}(x) \right) + \sum_{k \in S} \psi_k (F_k(x) - F_l(x))$$

$$\begin{aligned}
&= \sum_{i \in S} \sum_{j \in S} (x_i \rho_{ij} - x_j \rho_{ji}) \frac{\partial F_j}{\partial x_i}(x) + \sum_{k \in S} \psi_k (F_k(x) - F_i(x)) \\
&= \sum_{j \in S} \dot{x}_j \frac{\partial F_j}{\partial x_i}(x) + \sum_{k \in S} \psi_k (F_k(x) - F_i(x)).
\end{aligned}$$

Using this expression, we find the rate of change of Ψ over time along solutions to (M):

$$\begin{aligned}
\dot{\Psi}(x) &= \nabla \Psi(x)' \dot{x} \\
&= \dot{x}' DF(x) \dot{x} + \sum_{i \in S} \dot{x}_i \sum_{k \in S} \psi_k (F_k - F_i) \\
&= \dot{x}' DF(x) \dot{x} + \sum_{i \in S} \sum_{j \in S} (x_j \rho_{ji} - x_i \rho_{ij}) \sum_{k \in S} \psi_k (F_k - F_i) \\
&= \dot{x}' DF(x) \dot{x} + \sum_{i \in S} \sum_{j \in S} \left(x_j \rho_{ji} \sum_{k \in S} (\psi_k (F_k - F_i) - \psi_k (F_k - F_j)) \right).
\end{aligned}$$

To evaluate the summation, first observe that if $F_i(x) > F_j(x)$, then $\rho_{ji}(F(x)) \equiv \phi_i(F_i(x) - F_j(x)) > 0$ and $F_k(x) - F_i(x) < F_k(x) - F_j(x)$; since each ψ_k is nondecreasing, it follows that $\psi_k(F_k - F_i) - \psi_k(F_k - F_j) \leq 0$. In fact, when $k = i$, the comparison between payoff differences becomes $0 < F_i(x) - F_j(x)$; since each ψ_i is increasing on $[0, \infty)$, it follows that $\psi_i(0) - \psi_i(F_i - F_j) < 0$. We therefore conclude that if $F_i(x) > F_j(x)$, then $\rho_{ji}(F(x)) > 0$ and $\sum_{k \in S} (\psi_k (F_k - F_i) - \psi_k (F_k - F_j)) < 0$. On the other hand, if $F_j(x) \geq F_i(x)$, we have immediately that $\rho_{ji}(F(x)) = 0$. And of course, $\dot{x}' DF(x) \dot{x} \leq 0$ since F is stable.

Marshaling these facts, we find that $\dot{\Psi}(x) \leq 0$, and that

$$(7.27) \quad \dot{\Psi}(x) = 0 \text{ if and only if } x_j \rho_{ji}(F(x)) = 0 \text{ for all } i, j \in S.$$

Lemma 5.6.5 shows that the second condition in (7.27) defines the set $RP(V_F)$, which is equal to $NE(F)$ by Theorem 5.6.3; this proves the second claim. Finally, the global asymptotic stability of $NE(F)$ follows from Corollary 7.B.7. ■

Exercise 7.2.12. Construct a pairwise comparison dynamic that generates cycling in the good RPS game from Example 7.2.6.

7.2.4 Summary

In Table 7.1, we summarize the results in this section by presenting the Lyapunov functions for single-population stable games for the six fundamental evolutionary dynamics.

Dynamic	Formula	Lyapunov function
projection	$\dot{x} = \Pi_{\Gamma_X(x)}(F(x))$	$E_{x^*}(x) = x - x^* ^2$
replicator	$\dot{x}_i = x_i \hat{F}_i(x)$	$H_{x^*}(x) = \sum_{i \in S(x^*)} x_i^* \log \frac{x_i^*}{x_i}$
best response	$\dot{x} \in M(\hat{F}(x)) - x$	$G(x) = \mu(\hat{F}(x))$
logit	$\dot{x} = \tilde{M}(\hat{F}(x)) - x$	$\tilde{G}(x) = \tilde{\mu}(\hat{F}(x)) + v(x)$
BNN	$\dot{x}_i = [\hat{F}_i(x)]_+ - x_i \sum_{j \in S} [\hat{F}_j(x)]_+$	$\Gamma(x) = \frac{1}{2} \sum_{i \in S} [\hat{F}_i(x)]_+^2$
Smith	$\dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+$	$\Psi(x) = \frac{1}{2} \sum_{i \in S} \sum_{j \in S} x_i [F_j(x) - F_i(x)]_+^2$

Table 7.1: Lyapunov functions for the six fundamental dynamics in stable games.

The Lyapunov functions divide into three classes: those based on an explicit notion of “distance” from equilibrium, those based on revision potentials for target protocols, and the Lyapunov function for the Smith dynamic, which stands alone.

Example 7.2.13. Matching Pennies. In Figure 7.2.3, we present phase diagrams of the six fundamental dynamics in two-population Matching Pennies:

$$\begin{pmatrix} F_H^1(x) \\ F_T^1(x) \\ F_h^2(x) \\ F_t^2(x) \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 & -1 \\ 0 & 0 & -1 & 1 \\ -1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_H^1 \\ x_T^1 \\ x_h^2 \\ x_t^2 \end{pmatrix} = \begin{pmatrix} x_h^2 - x_t^2 \\ x_t^2 - x_h^2 \\ x_T^1 - x_H^1 \\ x_H^1 - x_T^1 \end{pmatrix}.$$

Each phase diagram is drawn atop a contour plot of the relevant Lyapunov function. Since Matching Pennies is a zero-sum game, F is null stable; thus, the Lyapunov functions for the replicator and projection dynamics define constants of motion for these dynamics, with solution trajectories cycling along level curves. In the remaining cases, all solutions converge to the unique Nash equilibrium, $x^* = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$. §

7.3 Supermodular Games

In a supermodular game, higher choices by one’s opponents make one’s own higher strategies look relatively more desirable. In Section 3.4, we used this property to show that the best response correspondences of supermodular games are monotone in the stochastic

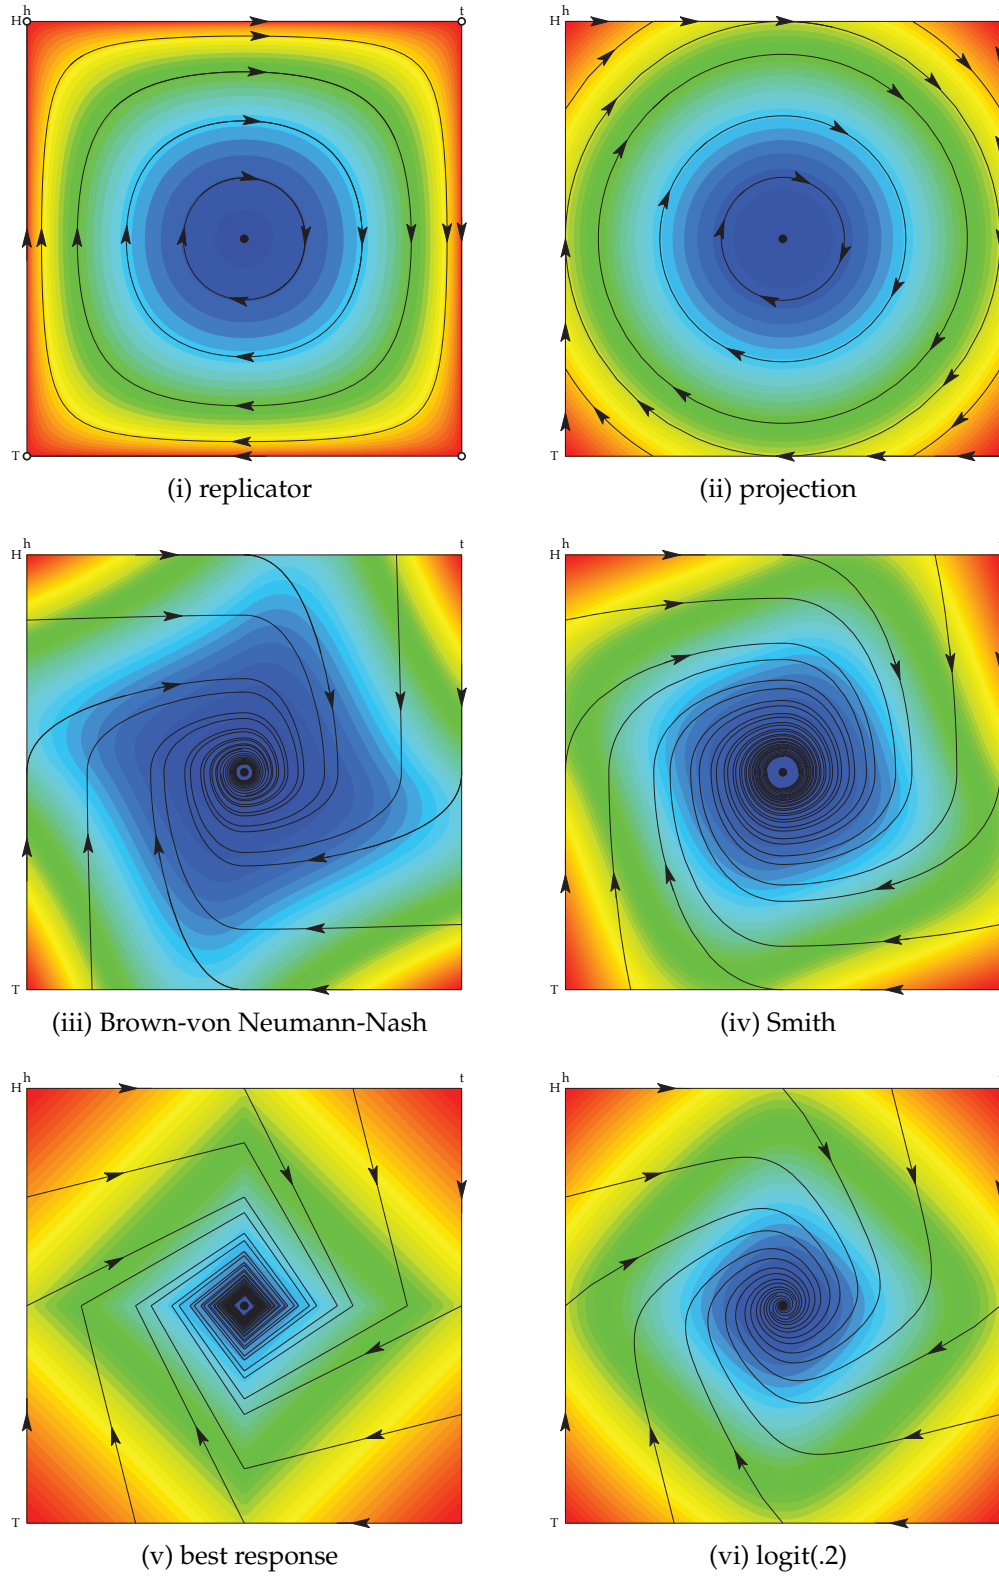


Figure 7.2.3: Six basic dynamics in Matching Pennies. The contour plots are the corresponding Lyapunov functions.

dominance order; this implies in turn that these games admit minimal and maximal Nash equilibria.

Given this monotone structure on best response correspondences, it is natural to look for convergence results for supermodular games under the best response dynamic (BR). In Section 7.3.1, we use elementary methods to establish a global convergence result for (BR) under some strong additional assumptions on the underlying game: in particular, it must be derived from a two-player normal form game that satisfies both supermodularity and “diminishing returns” conditions.

To prove more general convergence results, we appeal to the theory of cooperative differential equations. These are smooth differential equations under which increasing the value any component of the state variable increases the growth rates of all other components. Under some mild regularity conditions, almost all solutions of these equations converge to rest points.

Because of the smoothness requirement, these techniques cannot be applied to the best response dynamic itself. Happily, the needed monotonicity carries over from exact best responses to perturbed best responses, although only those that can be generated from *stochastic* perturbations of payoffs. In Section 7.3.2, we use this idea to prove almost global convergence of stochastically perturbed best response dynamics in supermodular games.

7.3.1 The Best Response Dynamic in Two-Player Normal Form Games

Let $U = (U^1, U^2)$ be a two-player normal form game, and let F be the population game obtained when members of two populations are matched to play U (cf Example 2.2.2). Then the best response dynamic (BR) for F takes the form

$$(BR) \quad \begin{aligned} \dot{x}^1 &\in B^1(x) - x^1 = M^1(F^1(x)) - x^1 = M^1(U^1 x^2) - x^1, \\ \dot{x}^2 &\in B^2(x) - x^2 = M^2(F^2(x)) - x^2 = M^2((U^2)'x^1) - x^2. \end{aligned}$$

Our convergence result for supermodular games concerns simple solutions of this dynamic. A solution $\{x_t\}_{t \geq 0}$ of (BR) is *simple* if the set of times at which it is not differentiable has no accumulation point, and if at other times, all elements of the sets of target states $B^p(x_t)$ are pure (i.e., vertices of X^p).

- Exercise 7.3.1.* (i) Provide an example of a 2×2 game with a Nash equilibrium x^* such that no solution to (BR) starting from x^* is simple.
- (ii) Show that there exists a simple solution to (BR) from every initial condition in game $U = (U^1, U^2)$ if for all nonempty sets $\hat{S}^1 \subseteq S^1$ and $\hat{S}^2 \subseteq S^2$, the game in which players

are restricted to strategies in \hat{S}^1 and \hat{S}^2 admits a pure Nash equilibrium. (Theorem 3.4.12 implies that U has this property if it is supermodular.)

If $\{x_t\}_{t \geq 0}$ is a simple solution trajectory of (BR), we can list the sequence of times $\{t_k\}$ at which the solution is not differentiable (i.e., at which the target state for at least one population changes). During each open interval of times $I_k = (t_{k-1}, t_k)$, the pure strategies $i_k \in S^1$ and $j_k \in S^2$ selected by revising agents are fixed; by continuity, these strategies are optimal not only at times in (t_{k-1}, t_k) , but also at the initial and terminal times t_{k-1} and t_k . We call i_k and j_k the *interval k selections* for populations 1 and 2.

The following lemma links shows that i_{k+1} must perform at least as well as i_k against both j_k and j_{k+1} , and that the analogous comparisons between the payoffs of j_k and j_{k+1} also hold.

Lemma 7.3.2. *Suppose that revising agents select strategies $i = i_k$ and $j = j_k$ during interval I_k , and strategies $i' = i_{k+1}$ and $j' = j_{k+1}$ during interval I_{k+1} . Then*

- (i) $U_{i'j}^1 \geq U_{ij}^1$ and $U_{i'j'}^2 \geq U_{ij'}^2$, and
- (ii) $U_{i'j'}^1 \geq U_{ij'}^1$ and $U_{i'j}^2 \geq U_{ij}^2$.

Exercise 7.3.3. Prove Lemma 7.3.2. (Hint: Notice that $x_{t_k}^2$ is a convex combination of $x_{t_{k-1}}^2$ and the vertex v_j^2 , and that $x_{t_{k+1}}^2$ is a convex combination of $x_{t_k}^2$ and $v_{j'}^2$. For part (ii), use the facts that strategies i and i' are both optimal at time t_k , and that strategy i' is optimal at time t_{k+1} .)

Now, recall from Exercise 3.4.4 that $U = (U^1, U^2)$ is supermodular if

$$(7.28) \quad U_{i+1,j+1}^1 - U_{i,j+1}^1 \geq U_{i+1,j}^1 - U_{i,j}^1 \text{ and } U_{i+1,j+1}^2 - U_{i+1,j}^2 \geq U_{i,j+1}^2 - U_{i,j}^2 \text{ for all } i < n^1, j < n^2.$$

(When (7.28) holds, the population game F induced by U is supermodular as well.) If the inequalities in (7.28) always hold strictly, we say that U is *strictly supermodular*.

Our convergence result requires two additional conditions on U . We say that U exhibits *strictly diminishing returns* if for each fixed strategy of the opponent, the benefit a player obtains by increasing his strategy is decreasing—in other words, if payoffs are “concave in own strategy”:

$$U_{i+2,j}^1 - U_{i+1,j}^1 < U_{i+1,j}^1 - U_{i,j}^1 \text{ for all } i \leq n^1 - 2 \text{ and } j \in S^2, \text{ and}$$

$$U_{i,j+2}^2 - U_{i,j+1}^2 < U_{i,j+1}^2 - U_{i,j}^2 \text{ for all } i \in S^1 \text{ and } j \leq n^2 - 2.$$

Finally, we say that U is *nondegenerate* if for each fixed pure strategy of the opponent, a player is not indifferent among any of his pure strategies.

Theorem 7.3.4. *Suppose that F is generated by matching in a two-player normal form game U that is strictly supermodular, exhibits strictly diminishing returns, and is nondegenerate. Then every simple solution trajectory of the best response dynamic (BR) converges to a pure Nash equilibrium.*

Proof. To begin, suppose that the sequence of times $\{t_k\}$ is finite, with final element t_K . Let i^* and j^* be the selections made by revising agents after time t_K . Then the pure state $x^* = (v_{i^*}^1, v_{j^*}^2)$ is in $B(x_t)$ for all $t \geq t_K$, and $\{x_t\}$ converges to x^* . Since payoffs are continuous, it follows that $x^* \in B(x^*)$, and so that x^* is a Nash equilibrium. To complete the proof of the theorem, we establish by contradiction that the sequence of times $\{t_k\}$ cannot be infinite.

To begin, note that at time t_k , agents in each population p are indifferent between their interval k and interval $k + 1$ selections. Moreover, since U exhibits strictly decreasing returns, it is easy to verify that whenever such an indifference occurs, it must be between two consecutive strategies in S^p . Putting these observations together, we find that each transition in the sequence $\{(i_k, j_k)\}$ is of length 1, in the sense that

$$|i_{k+1} - i_k| \leq 1 \text{ and } |j_{k+1} - j_k| \leq 1 \text{ for all } k.$$

Next, we say that there is an *improvement step* from $(i, j) \in S$ to $(i', j') \in S$, denoted $(i, j) \nearrow (i', j')$, if either (i) $U_{i'j}^1 > U_{ij}^1$ and $j' = j$, or (ii) $i' = i$ and $U_{ij'}^2 > U_{ij}^2$. Lemma 7.3.2(i) and the fact that U is nondegenerate imply that $(i_k, j_k) \nearrow (i_{k+1}, j_{k+1})$ if either $i_k = i_{k+1}$ or $j_k = j_{k+1}$. Moreover, applying both parts of the lemma, we find that if $i_k \neq i_{k+1}$ and $j_k \neq j_{k+1}$, we have that $(i_k, j_k) \nearrow (i_{k+1}, j_k) \nearrow (i_{k+1}, j_{k+1})$, and also that $(i_k, j_k) \nearrow (i_k, j_{k+1}) \nearrow (i_{k+1}, j_{k+1})$.

Now suppose that the sequence $\{t_k\}$ is infinite. Then since S is finite, there must be a strategy profile that is the interval k selection for more than one k . In this case, the arguments in the previous two paragraphs imply that there is a length 1 improvement cycle: that is, a sequence of length 1 improvement steps beginning and ending with the same strategy profile.

Evidently, this cycle must contain an improvement step of the form $(\tilde{i}, \tilde{j}) \nearrow (\tilde{i}, \tilde{j} + 1)$ for some $(\tilde{i}, \tilde{j}) \in S$. Strict supermodularity of U then implies that

$$(7.29) \quad (i, \tilde{j}) \nearrow (i, \tilde{j} + 1) \text{ for all } i \geq \tilde{i}.$$

It follows that for the sequence of length 1 improvement steps to return to (\tilde{i}, \tilde{j}) , there must be an improvement step of the form $(\tilde{i}, \hat{j}) \nearrow (\tilde{i} - 1, \hat{j})$ for some $\hat{j} > \tilde{j}$ (see Figure 7.3.1). This time, strict supermodularity of U implies that

$$(7.30) \quad (\tilde{i}, j) \nearrow (\tilde{i} - 1, j) \text{ for all } j \leq \hat{j}.$$




Figure 7.3.1: The proof of Theorem 7.3.4.

From (7.29) and (7.30), it follows that no cycle of length 1 improvement steps containing $(\tilde{i}, \tilde{j}) \nearrow (\tilde{i}, \tilde{j} + 1)$ can reach any strategy profile (i, j) with $i \geq \tilde{i}$ and $j \leq \tilde{j}$. In particular, the cycle cannot return to (\tilde{i}, \tilde{j}) , which is a contradiction. This completes the proof of the theorem. ■

7.3.2 Stochastically Perturbed Best Response Dynamics

While Theorem 7.3.4 was proved using elementary techniques, it was not as general as one might hope: it restricted attention to two-player normal form games, and required not only the assumption of supermodularity, but also that of decreasing returns. In order to obtain a more general convergence result, we turn from exact best response dynamics to perturbed best response dynamics. Doing so allows us to avail ourselves of a powerful set of techniques for smooth dynamics with a monotone structure: the theory of cooperative differential equations.

To begin, let us recall the transformations used to discuss the stochastic dominance order. In Section 3.4, we defined the matrices $\Sigma \in \mathbf{R}^{(n^p-1) \times n^p}$, $\tilde{\Sigma} \in \mathbf{R}^{n^p \times (n^p-1)}$, and $\Omega \in \mathbf{R}^{n^p \times n^p}$ by

$$\Sigma = \begin{pmatrix} 0 & 1 & \cdots & 1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 \end{pmatrix}, \quad \tilde{\Sigma} = \begin{pmatrix} -1 & 0 & \cdots & 0 \\ 1 & -1 & \ddots & \vdots \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & -1 \\ 0 & \ddots & 0 & 1 \end{pmatrix}, \quad \text{and } \Omega = \begin{pmatrix} 1 & 1 & \cdots & \cdots & 1 \\ 0 & 0 & \cdots & \cdots & 0 \\ 0 & 0 & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \cdots & 0 \end{pmatrix}.$$

We saw that $y^p \in X^p$ stochastically dominates $x^p \in X^p$ if and only if $\Sigma y^p \geq \Sigma x^p$. We also

verified that

$$(7.31) \quad \tilde{\Sigma}\Sigma = I - \Omega.$$

Since Ω is the null operator on TX^p , equation (7.31) describes a sense in which the stochastic dominance operator Σ is “inverted” by the difference operator $\tilde{\Sigma}$.

Applying the change of coordinates Σ to the set X^p yields the set of transformed population states

$$\mathcal{X}^p \equiv \Sigma X^p = \left\{ \chi^p \in \mathbf{R}^{m^p-1} : m^p \geq \chi_1^p \geq \dots \geq \chi_{m^p-1}^p \geq 0 \right\}.$$

By postmultiplying both sides of (7.31) by x^p and letting $\underline{x}^p = (m^p, 0, \dots, 0)$ denote the minimal state in X^p , we find that the inverse of the map $\Sigma : X^p \rightarrow \mathcal{X}^p$ is described as follows:

$$(7.32) \quad \chi^p = \Sigma x^p \Leftrightarrow x^p = \tilde{\Sigma} \chi^p + \underline{x}^p.$$

To work with full social states $x \in X$, we introduce the block diagonal matrices $\Sigma = \text{diag}(\Sigma, \dots, \Sigma)$ and $\tilde{\Sigma} = \text{diag}(\tilde{\Sigma}, \dots, \tilde{\Sigma})$, and let $\mathcal{X} \equiv \Sigma X = \prod_{p \in \mathcal{P}} \mathcal{X}^p$. If we let $\underline{x} = (\underline{x}^1, \dots, \underline{x}^P)$ be the minimal state in X , then the inverse of the map $\Sigma : X \rightarrow \mathcal{X}$ is described by

$$(7.33) \quad \chi = \Sigma x \Leftrightarrow x = \tilde{\Sigma} \chi + \underline{x}.$$

To simplify the discussion to follow, let us assume for convenience that each population is of mass 1. Then our stochastically perturbed best response dynamics take the form

$$(7.34) \quad \dot{x}^p = \tilde{M}^p(F^p(x)) - x^p,$$

where

$$\tilde{M}_i^p(\pi^p) = P\left(i = \operatorname{argmax}_{j \in S^p} \pi_j^p + \varepsilon_j^p\right)$$

for some admissible stochastic perturbations $\varepsilon = (\varepsilon^1, \dots, \varepsilon^P)$. Rather than study this dynamic directly, we apply the change of variable (7.33) to obtain a new dynamic on the set \mathcal{X} :

$$(7.35) \quad \dot{\chi}^p = \Sigma \tilde{M}^p(F^p(\tilde{\Sigma} \chi + \underline{x})) - \chi^p.$$

Given the current state $\chi \in \mathcal{X}$, we use the inverse transformation $\chi \mapsto x \equiv \tilde{\Sigma} \chi + \underline{x}$

to obtain the input for the payoff function F^p , and we use the original transformation $\tilde{M}^p(F^p(x)) \mapsto \Sigma \tilde{M}^p(F^p(x))$ to convert the perturbed best response into an element of \mathcal{X}^p . The next observation verifies the relationship between solutions to the transformed dynamic (7.35) and solutions to the original dynamic (7.34).

Observation 7.3.5. (7.34) and (7.35) are affinely conjugate: $\{x_t\} = \{\tilde{\Sigma}\chi_t + \underline{x}\}$ solves (7.34) if and only if $\{\chi_t\} = \{\Sigma x_t\}$ solves (7.35).

Our next task is to show that if F is a supermodular game, then (7.35) is a *cooperative differential equation*: writing this dynamic as $\dot{\chi} = \mathcal{V}(\chi)$, we want to show that

$$\frac{\partial \mathcal{V}_i^p}{\partial \chi_j^q}(\chi) \geq 0 \text{ for all } \chi \in \mathcal{X} \text{ whenever } (i, p) \neq (j, q).$$

If this inequality is always satisfied strictly, we say that (7.35) is *strongly cooperative*. As we explain in Section 7.C, strongly cooperative differential equations converge to rest points from almost all initial conditions. Thus, if we can prove that equation (7.35) is strongly cooperative, we can conclude that almost all solutions of our original dynamic (7.34) converge to perturbed equilibria.

To prove that (7.35) is strongly cooperative, we marshal our facts about supermodular games and stochastically perturbed best responses. Recall from Chapter 3 that if the population game F is C^1 , then F is a *supermodular* if and only if

$$\tilde{\Sigma}' DF(x) \tilde{\Sigma} \geq \mathbf{0} \text{ for all } x \in X.$$

Our result requires an additional nondegeneracy condition: we say that F is *irreducible* if each column of $\tilde{\Sigma}' DF(x) \tilde{\Sigma}$ contains a strictly positive element.

Next, we recall from Lemma 6.C.1 the basic properties of $D\tilde{M}(\pi)$, the derivative matrix of the stochastically perturbed best response function \tilde{M} .

Lemma 7.3.6. Fix $\pi \in \mathbf{R}^n$, and suppose that the perturbed best response function \tilde{M} is derived from admissible stochastic payoff perturbations. Then the derivative matrix $D\tilde{M}(\pi)$ is symmetric, has negative off-diagonal elements, and satisfies $D\tilde{M}(\pi)\mathbf{1} = \mathbf{0}$.

Combining these facts yields the desired result:

Theorem 7.3.7. Let F be a C^1 irreducible supermodular game, and let (7.34) be a stochastically perturbed best response dynamic for F . Then the transformed dynamic (7.35) is strongly cooperative.

Proof. ($p = 1$) Write the dynamic (7.35) as $\dot{\chi} = \mathcal{V}'(\chi)$. Then

$$(7.36) \quad D\mathcal{V}'(\chi) = D(\Sigma\tilde{M}(F(\tilde{\Sigma}\chi + \underline{x}))) - I.$$

Since all off-diagonal elements of I equal zero, it is enough to show that the first term on the right hand side of (7.36) has all positive components.

Let $x = \tilde{\Sigma}\chi + \underline{x}$ and $\pi = F(x)$. Using the facts that $\tilde{\Sigma}\Sigma = I - \Omega$ and $D\tilde{M}(\pi)\mathbf{1} = \mathbf{0}$, we express the first term of the right hand side of (7.36) as follows:

$$\begin{aligned} D(\Sigma\tilde{M}(F(\tilde{\Sigma}\chi + \underline{x}))) &= \Sigma D\tilde{M}(\pi) DF(x) \tilde{\Sigma} \\ &= \Sigma D\tilde{M}(\pi) (\Sigma' \tilde{\Sigma}' + \Omega') DF(x) \tilde{\Sigma} \\ &= (\Sigma D\tilde{M}(\pi) \Sigma') (\tilde{\Sigma}' DF(x) \tilde{\Sigma}). \end{aligned}$$

Lemma 7.3.6 and the fact that

$$(\Sigma D\tilde{M}(\pi) \Sigma')_{ij} = \sum_{k>i} \sum_{l>j} D\tilde{M}(\pi)_{kl}$$

imply that every component of $\Sigma D\tilde{M}(\pi) \Sigma'$ is positive (see Exercise 7.3.8). Since F is supermodular and irreducible, $\tilde{\Sigma}' DF(x) \tilde{\Sigma}$ is nonnegative, with each column containing a positive element. Thus, the product of these two matrices has all positive elements. This completes the proof of the theorem. ■

Exercise 7.3.8. (i) Prove that every component of $\Sigma D\tilde{M}(\pi) \Sigma'$ is positive.

(ii) Explain why Theorem 7.3.7 need not hold when \tilde{M} is generated by deterministic perturbations.

Observation 7.3.5, Theorem 7.3.7, and Theorems 7.C.1, 7.C.2, and 7.C.3 immediately imply the following “almost global” convergence result. In part (i) of the theorem, $\underline{x} = (\underline{x}^1, \dots, \underline{x}^p)$ is the minimal state in X introduced above; similarly, $\bar{x}^p = (0, \dots, m^p)$ is the maximal state in X^p , and $\bar{x} = (\bar{x}^1, \dots, \bar{x}^p)$ is the maximal state in X .

Theorem 7.3.9. *Let F be a C^1 irreducible supermodular game, and let $\dot{x} = V_{F,\varepsilon}(x)$ be a stochastically perturbed best response dynamic for F . Then*

- (i) *States $\underline{x}^* \equiv \omega(\underline{x})$ and $\bar{x}^* \equiv \omega(\bar{x})$ exist and are the minimal and maximal elements of $PE(F, \varepsilon)$. Moreover, $[\underline{x}^*, \bar{x}^*]$ contains all ω -limit points of $V_{F,\varepsilon}$ and is globally asymptotically stable.*
- (ii) *Solutions to $\dot{x} = V_{F,\varepsilon}(x)$ from an open, dense, full measure set of initial conditions in X converge to states in $PE(F, \varepsilon)$.*

Our final example shows that the conclusion of Theorem 7.3.9 cannot be extended from convergence from almost all initial conditions to convergence from all initial conditions.

Example 7.3.10. Let U be a normal form game with $p \geq 5$ players and two strategies per player. Each player p in U obtains a payoff of 1 if he chooses the same strategy as player $p + 1$ (with the convention that $p + 1 = 1$) and obtains a payoff of 0 otherwise. U has three Nash equilibria: two strict equilibria in which all players coordinate on the same strategy, and the mixed equilibrium $x^* = ((\frac{1}{2}, \frac{1}{2}), \dots, (\frac{1}{2}, \frac{1}{2}))$. If F is the p population game generated by matching in U , it can be shown that F is supermodular and irreducible (see Exercise 7.3.11(i)).

We now introduce random perturbations $\varepsilon^p = (\varepsilon_1^p, \varepsilon_2^p)$ to each player's payoffs. These perturbations are such that the differences $\varepsilon_2^p - \varepsilon_1^p$ admit a common density g that is symmetric about 0, is decreasing on \mathbf{R}_+ , and satisfies $g(0) > \frac{1}{2}$. It can be shown that the resulting perturbed best response dynamic (7.34) possesses exactly three rest points: the mixed equilibrium x^* , and two stable symmetric rest points that approximate the two pure Nash equilibria (see Exercise 7.3.11(ii)).

One can show that the rest point x^* is unstable under (7.34). It then follows from Theorem 7.3.9 that the two stable rest points of (7.34) attract almost all initial conditions in X , and that the basins of attraction for these rest points are separated by a $p - 1$ dimensional invariant manifold \mathcal{M} that contains x^* . Furthermore, one can show that when $p \geq 5$, the rest point x^* is unstable with respect to the manifold \mathcal{M} . Thus, solutions from all states in $\mathcal{M} - \{x^*\}$ fail to converge to a rest point. §

The details of these last arguments require techniques for determining the local stability of rest points. This is the topic of the next chapter.

Exercise 7.3.11. (i) Prove that the game F introduced in Example 7.3.10 is supermodular and irreducible.

(ii) Prove that under the assumption on payoff perturbations stated in the example, there are exactly three perturbed equilibria, all of which are symmetric.

7.4 Dominance Solvable Games

The elimination of strictly dominated strategies is the mildest requirement employed in standard game-theoretic analyses, and so it seems natural to expect evolutionary dynamics obey this dictum. In this section, we provide some positive results on the elimination of dominated strategies: under the best response dynamic, any strictly dominated strategy must vanish in the limit; the same is true under any imitative dynamic so long as we focus

on interior initial conditions. Arguing inductively, we show next that any strategy that does not survive iterated elimination of strictly dominated strategies vanishes as well. In particular, if a game is dominance solvable—that is, if removing iteratively dominated strategies leaves only one strategy for each population, then best response and imitative dynamics select this strategy.

These results may seem unsurprising. However, we argue in Chapter 9 that they are actually borderline cases: under “typical” evolutionary dynamics, strictly dominated strategies can survive in perpetuity.

7.4.1 Dominated and Iteratively Dominated Strategies

Let F be a population game. We say that strategy $i \in S^p$ is *strictly dominated* if there exists a strategy $j \in S^p$ such that $F_j(x) > F_i(x)$ for all $x \in X$: that is, if there is a strategy j that outperforms strategy i regardless of the population state. Similarly, if \hat{S}^p is a nonempty subset of S^p and $\hat{S} = \prod_{p \in \mathcal{P}} \hat{S}^p$, we say that $i \in S^p$ is *strictly dominated relative to \hat{S}* , denoted $i \in \mathcal{D}^p(\hat{S})$, if there exists a strategy $j \in \hat{S}^p$ such that $F_j(x) > F_i(x)$ for all $x \in X$ that satisfy $\text{support}(x^p) \subseteq \hat{S}^p$ for all $p \in \mathcal{P}$.

We can use these definitions to introduce the notion of iterative dominance. Set $S_0 = S$. Then $\mathcal{D}^p(S_0)$ is the set of strictly dominated strategies for population p , and so $S_1^p = S_0^p - \mathcal{D}^p(S_0)$ is the set of strategies that are not strictly dominated. Proceeding inductively, we define $\mathcal{D}^p(S_k)$ to be the set of strategies that are eliminated during the $(k+1)$ st round of removal of iteratively dominated strategies, and we let $S_{k+1}^p = S_k^p - \mathcal{D}^p(S_k)$ be the set of strategies that survive $k+1$ rounds of removal of such strategies.

Since the number of strategies is finite, this iterative procedure must converge, leaving us with nonempty sets S_*^1, \dots, S_*^p . Strategies in these sets are said to *survive iterative removal of strictly dominated strategies*. If each of these sets is a singleton, then the game F is said to be *dominance solvable*. In this case, the pure social state at which each agent plays his population’s sole surviving strategy is the game’s unique Nash equilibrium; we call this state the *dominance solution* of F .

7.4.2 The Best Response Dynamic

Under the best response dynamic, revising agents always switch to optimal strategies. Since strictly dominated strategies are never optimal, such strategies cannot persist:

Observation 7.4.1. *Let $\{x_t\}$ be a solution trajectory of (BR) for population game F , in which strategy $i \in S^p$ is strictly dominated. Then $\lim_{t \rightarrow \infty} (x_t)_i^p = 0$.*

Indeed, since i is never a best response, we have that $(\dot{x}_t)_i^p \equiv -(x_t)_i^p$, and hence that $(x_t)_i^p = (x_0)_i^p e^{-t}$: the mass playing the dominated strategy converges to zero exponentially quickly.

An inductive argument takes us from the observation above to the following result.

Theorem 7.4.2. *Let $\{x_t\}$ be a solution trajectory of (BR) for population game F , in which strategy $i \in S^p$ does not survive iterative elimination of strictly dominated strategies. Then $\lim_{t \rightarrow \infty} (x_t)_i^p = 0$. In particular, if F is dominance solvable, then all solutions of (BR) converge to the dominance solution.*

Proof. Observation 7.4.1 provides the basis for this induction: if $i \notin S_1^p$, then $\lim_{t \rightarrow \infty} (x_t)_i^p = 0$. As the inductive hypothesis, suppose that this same equality holds for all $i \notin S_k^p$. Now let $j \in S_k^p - S_{k+1}^p$. Then by definition, there exists a $j' \in S_{k+1}^p$ such that $F_{j'}^p(x) > F_j^p(x)$ whenever $x \in X_k$, where $X_k = \{x \in X : x_i^p > 0 \Rightarrow i \in S_k^p\}$ is the set of social states in which all agents in each population p choose a strategies in S_k^p . Since X_k is compact and F is continuous, it follows that for some $c > 0$, we have that $F_{j'}^p(x) > F_j^p(x) + c$ whenever $x \in X_k$, and so that for some $\varepsilon > 0$, we have that $F_{j'}^p(x) > F_j^p(x)$ whenever $x \in X_{k,\varepsilon} = \{x \in X : x_i^p > \varepsilon \Rightarrow i \in S_k^p\}$. By the inductive hypothesis, there exists a $T > 0$ such that $x_t \in X_{k,\varepsilon}$ for all $t \geq T$. Thus, for such t , j is not a best response to x_t . This implies that $(\dot{x}_t)_j^p = -(x_t)_j^p$ for $t \geq T$, and hence that $(x_t)_j^p = (x_T)_j^p e^{T-t}$, which converges to 0 as t approaches infinity. ■

Exercise 7.4.3. Show that under (BR), the time until convergence to the set $X_{*,\varepsilon} = \{x \in X : x_i^p > \varepsilon \Rightarrow i \in S_*^p\}$ is uniform over initial conditions in X . In other words, show that there is a $T < \infty$ such that all solutions to (BR) from all initial conditions in X enter $X_{*,\varepsilon}$ by time T .

7.4.3 Imitative Dynamics

We now establish analogous results for imitative dynamics. Since these dynamics leave the boundary of the state space invariant, the elimination results can only hold for solutions starting from interior initial conditions.

Theorem 7.4.4. *Let $\{x_t\}$ be an interior solution trajectory of an imitative dynamic for population game F , in which strategy $i \in S^p$ is strictly dominated. Then $\lim_{t \rightarrow \infty} (x_t)_i^p = 0$.*

Proof. ($p = 1$) Observation 5.4.16 tells us that all imitative dynamics $\dot{x} = V_F(x)$ exhibit monotone percentage growth rates (5.18): we can write the dynamic as

$$(7.37) \quad \dot{x}_i = x_i G_i(x),$$

where the continuous function $G : X \rightarrow \mathbf{R}^n$ satisfies

$$(7.38) \quad G_k(x) \leq G_l(x) \text{ if and only if } F_k(x) \leq F_l(x) \text{ for all } x \in \text{int}(X).$$

Now suppose strategy i is strictly dominated by strategy $j \in S$. Since X is compact and F is continuous, we can find a $c > 0$ such that $F_j(x) - F_i(x) > c$ for all $x \in X$. Since G is continuous as well, equation (7.38) implies that for some $C > 0$, we have that $G_j(x) - G_i(x) > C$ for all $x \in X$.

Now write $r = x_i/x_j$. Equation (7.37) and the quotient rule imply that

$$(7.39) \quad \frac{d}{dt} r = \frac{d}{dt} \frac{x_i}{x_j} = \frac{\dot{x}_i x_j - \dot{x}_j x_i}{(x_j)^2} = \frac{x_i G_i(x) x_j - x_j G_j(x) x_i}{(x_j)^2} = r (G_i(x) - G_j(x)).$$

Thus, along every interior solution trajectory $\{x_t\}$ of $\dot{x} = V_F(x)$ we have that

$$r_t = r_0 + \int_0^t r_s (G_i(x) - G_j(x)) ds \leq r_0 - C \int_0^t r_s ds.$$

Grönwall's Inequality (Lemma 4.A.7) then tells us that $r_t \leq r_0 \exp(-Ct)$, and hence that r_t vanishes as t approaches infinity. Since $(x_t)_j$ is bounded above by 1, $(x_t)_i$ must approach 0 as t approaches infinity. ■

An argument similar to the one used to prove Theorem 7.4.2 can be used to prove that iteratively dominated strategies are eliminated by imitative dynamics.

Theorem 7.4.5. *Let $\{x_t\}$ be an interior solution trajectory of an imitative dynamic for population game F , in which strategy $i \in S^p$ does not survive iterative elimination of strictly dominated strategies. Then $\lim_{t \rightarrow \infty} (x_t)_i^p = 0$. In particular, if F is dominance solvable, then all interior solutions of any imitative dynamic converge to the dominance solution.*

Exercise 7.4.6. (i) Prove Theorem 7.4.5.

(ii) Is the time until convergence to $X_{*,\varepsilon} = \{x \in X : x_i^p > \varepsilon \Rightarrow i \in S_*^p\}$ uniform over initial conditions in $\text{int}(X)$? Explain.

Appendix

7.A Limit and Stability Notions for Deterministic Dynamics

We consider differential equations and differential inclusions that are forward invariant on the compact set $X \subset \mathbf{R}^n$.

(D) $\dot{x} = V(x)$, a unique forward solution exists from each $\xi \in X$.

(DI) $\dot{x} \in V(x)$, V is nonempty, convex-valued, bounded, and upper hemicontinuous.

When V is discontinuous, we allow solutions to be of the Carathéodory type—that is, to satisfy $\dot{x}_t = V(x_t)$ (or $\dot{x}_t \in V(x_t)$) at almost all $t \in [0, \infty)$.

7.A.1 ω -Limits and Notions of Recurrence

Let $\{x_t\} = \{x_t\}_{t \geq 0}$ be a solution trajectory to (D) or (DI). The ω -limit of $\{x_t\}$ is the set of all points that the trajectory approaches arbitrarily closely infinitely often:

$$\omega(\{x_t\}) = \left\{ y \in X : \text{there exists } \{t_k\}_{k=1}^{\infty} \text{ with } \lim_{k \rightarrow \infty} t_k = \infty \text{ such that } \lim_{k \rightarrow \infty} x_{t_k} = y \right\}.$$

If $\{x_t\}$ is the unique solution to (D) with initial condition $x_0 = \xi$, we write $\omega(\xi)$ in place of $\omega(\{x_t\})$. The following proposition lists some basic properties of ω -limit sets for this case.

Proposition 7.A.1. *Let $\{x_t\}$ be a solution to (D) from initial condition ξ . Then*

- (i) $\omega(\xi)$ is non-empty and connected.
- (ii) $\omega(\xi)$ is closed. In fact, $\omega(\{x_t\}) = \bigcap_{t \geq 0} \text{cl}(\{x_s : s \geq t\})$.
- (iii) $\omega(\xi)$ is invariant under (D).

The properties of ω -limit sets for differential inclusions are somewhat more subtle—see the notes for references.

In the case of the differential equation (D), the set

$$\Omega = \bigcup_{\xi \in X} \omega(\xi),$$

contains all points that are approached arbitrarily closely infinitely often by some solution of (D). Among other things, Ω contains all rest points, periodic orbits, and chaotic attractors of (D). Since Ω need not be closed, its closure $\bar{\Omega} = \text{cl}(\Omega)$ is used to define a standard notion of recurrence for differential equations.

Example 7.A.2. To see that Ω need not be closed, consider the replicator dynamic in standard Rock-Paper-Scissors (Figure 5.3.1(i)). The unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is a rest point, and solution trajectories from all other interior initial conditions form closed orbits around x^* . The vertices $e_R, e_P,$ and e_S are also rest points, and each trajectory starting from a boundary point that is not a vertex converges to a vertex. Thus, $\Omega = \text{int}(X) \cup \{e_R, e_P, e_S\}$, but $\bar{\Omega} = X$. §

While we will not make much use of them here, many other notions of recurrence besides $\bar{\Omega}$ are available. To obtain a more demanding notion of recurrence for (D), call the state ξ *recurrent*, denoted $\xi \in \mathcal{R}$, if the solution from (D) returns arbitrarily close to ξ infinitely often—in other words, if $\xi \in \omega(\xi)$. The *Birkhoff center* of (D) is the closure $\text{cl}(\mathcal{R})$ of the set of all recurrent points of (D).

More inclusive notions of recurrence can be obtained by allowing occasional short jumps between nearby solution trajectories. Given a differential equation (D) with flow ϕ , an ε -*chain* of length T from x to y is a sequence of states $x = x_0, x_1, \dots, x_k = y$ such that for some sequence of times $t_1, \dots, t_n \geq 1$ satisfying $\sum_{j=1}^k t_j = T$, we have $|\phi_{t_i}(x_{i-1}) - x_j| < \varepsilon$ for all $i \in \{1, \dots, k\}$. State x is said to be *chain recurrent*, denoted $x \in \mathcal{CR}$, if for all $\varepsilon > 0$ there is an ε -chain from x to itself.

The primacy of the notion of chain recurrence is captured by Theorem 9.B.4, known as the *Fundamental Theorem of Dynamical Systems*: if ϕ is a smooth flow on a compact set X , then X can be decomposed into two sets: a set on which the flow admits a Lyapunov function, and the set \mathcal{CR} of chain recurrent points. Chain recurrence also plays a basic role in characterizing long run behavior in models of learning in games (see the Notes).

7.A.2 Stability of Sets of States

Let $A \subseteq X$ be a closed set, and call $O \subseteq X$ a *neighborhood* of A if it is open relative to X and contains A . We say that A is *Lyapunov stable* under (D) (or (DI)) if for every neighborhood O of A there exists a neighborhood O' of A such that every solution $\{x_t\}$ that starts in O' is contained in O : that is, $x_0 \in O'$ implies that $x_t \in O$ for all $t \geq 0$. A is *attracting* if there is a neighborhood Y of A such that every solution that starts in Y converges to A : that is, $x_0 \in Y$ implies that $\omega(\{x_t\}) \subseteq A$. A is *globally attracting* if it is attracting with $Y = X$. Finally, the set A is *asymptotically stable* if it is Lyapunov stable and attracting, and it is *globally asymptotically stable* if it is Lyapunov stable and globally attracting.

Example 7.A.3. Attracting sets need not be asymptotically stable. A counterexample is provided by a flow on the unit circle that moves clockwise except at a single point. The

fact that the domain is the unit circle is unimportant, since one can embed this flow as a limit cycle in a flow on the plane. §

Example 7.A.4. Invariance is not included in the definition of asymptotic stability. Thus, under the dynamic $\dot{x} = -x$ on \mathbf{R} , any closed interval containing the origin is asymptotically stable. §

7.B Stability Analysis via Lyapunov Functions

Let $Y \subseteq X$. The function $L : Y \rightarrow \mathbf{R}$ is a *Lyapunov function* for (D) or (DI) if its value changes monotonically along every solution trajectory. We state the results to follow for the case in which the value of L decreases along solution trajectories; of course, the obvious analogues of these results hold for the opposite case.

The following lemma will prove useful in a number of the analyses to come.

Lemma 7.B.1. *Suppose that the function $L : Y \rightarrow \mathbf{R}$ and the trajectory $\{x_t\}_{t \geq 0}$ are Lipschitz continuous.*

- (i) *If $\dot{L}(x_t) \leq 0$ for almost all $t \geq 0$, then the map $t \mapsto L(x_t)$ is nonincreasing.*
- (ii) *If in addition $\dot{L}(x_s) < 0$, then $L(x_t) < L(x_s)$ for all $t > s$.*

Proof. The composition $t \mapsto L(x_t)$ is Lipschitz continuous, and therefore absolutely continuous. It thus follows from the Fundamental Theorem of Calculus that when $t > s$, we have

$$L(x_t) - L(x_s) = \int_s^t \dot{L}(x_u) \, du \leq 0,$$

where the inequality is strict if $\dot{L}(x_s) < 0$. ■

7.B.1 Lyapunov Stable Sets

The basic theorem on Lyapunov stability applies both to differential equations (D) and differential inclusions (DI).

Theorem 7.B.2 (Lyapunov stability). *Let $A \subseteq X$ be closed, and let $Y \subseteq X$ be a neighborhood of A . Let $L : Y \rightarrow \mathbf{R}_+$ be Lipschitz continuous with $L^{-1}(0) = A$. If each solution $\{x_t\}$ of (D) (or (DI)) satisfies $\dot{L}(x_t) \leq 0$ for almost all $t \geq 0$, then A is Lyapunov stable under (D) (or (DI)).*

Proof. Let O be a neighborhood of A such that $\text{cl}(O) \subset Y$. Let $c = \min_{x \in \text{bd}(O)} L(x)$, so that $c > 0$. Finally, let $O' = \{x \in O : L(x) < c\}$. Lemma 7.B.1 implies that solution trajectories that start in O' do not leave O , and hence that A is Lyapunov stable. ■

Example 7.B.3. The requirement that the function L be constant on A cannot be dispensed with. Consider a flow on the unit circle $C = \{x \in \mathbf{R}^2 : (x_1)^2 + (x_2)^2 = 1\}$ that moves clockwise at states x with $x_1 > 0$ and is at rest at states in the semicircle $A = \{x \in C : x_1 \leq 0\}$. If we let $L(x) = x_2$, then $\dot{L}(x) \leq 0$ for all $x \in C$, and A is attracting (see Theorem 7.B.4 below), but A is not Lyapunov stable.

We can extend this example so that the flow is defined on the unit disk $D = \{x \in \mathbf{R}^2 : (x_1)^2 + (x_2)^2 \leq 1\}$. Suppose that when $x_1 > 0$, the flow travels clockwise along the circles centered at the origin, and that the half disk $A' = \{x \in D : x_1 \leq 0\}$ consists entirely of rest points. Then $L(x) = x_2$ satisfies $\dot{L}(x) \leq 0$ for all $x \in D$, and A' is attracting, but A' is not Lyapunov stable. §

7.B.2 ω -Limits and Attracting Sets

We now provide some results that use Lyapunov functions to characterize ω -limits of solution trajectories that begin in the Lyapunov function's domain. These results immediately yield sufficient conditions for a set to be attracting. To state our results, we call the (relatively) open set $Y \subset X$ *inescapable* if for each solution trajectory $\{x_t\}_{t \geq 0}$ with $x_0 \in Y$, we have that $\text{cl}(\{x_t\}) \cap \text{bd}(Y) = \emptyset$.

Our first result focuses on the differential equation (D).

Theorem 7.B.4. *Let $Y \subset X$ be relatively open and inescapable under (D). Let $L : Y \rightarrow \mathbf{R}$ be C^1 , and suppose that $\dot{L}(x) \equiv \nabla L(x)'V(x) \leq 0$ for all $x \in Y$. Then $\omega(x_0) \subseteq \{x \in Y : \dot{L}(x) = 0\}$ for all $x_0 \in Y$. Thus, if $\dot{L}(x) = 0$ implies that $V(x) = \mathbf{0}$, then $\omega(x_0) \subseteq \text{RP}(V) \cap Y$.*

Proof. Let $\{x_t\}$ be the solution to (D) with initial condition $x_0 = \xi \in Y$, let $\chi \in \omega(\xi)$, and let $\{y_t\}$ be the solution to (D) with $y_0 = \chi$. Since Y is inescapable, the closures of trajectories $\{x_t\}$ and $\{y_t\}$ are contained in Y .

Suppose by way of contradiction that $\dot{L}(\chi) \neq 0$. Since $\chi \in \omega(\xi)$, we can find a divergent sequence of times $\{t_k\}_{k=1}^{\infty}$ such that $\lim_{k \rightarrow \infty} x_{t_k} = \chi = y_0$. Since solutions to (D) are unique, and hence continuous in their initial conditions, we have that

$$(7.40) \quad \lim_{k \rightarrow \infty} x_{t_k+1} = y_1, \text{ and hence that } \lim_{k \rightarrow \infty} L(x_{t_k+1}) = L(y_1).$$

But since $y_0 = \chi \in \omega(\xi)$ and $\dot{L}(\chi) \neq 0$, applying Lemma 7.B.1 to both $\{x_t\}$ and $\{y_t\}$ yields

$$L(x_t) \geq L(\chi) > L(y_1)$$

for all $t \geq 0$, contradicting the second limit in (7.40). This proves the first claim of the theorem, and the second claim follows immediately from the first. ■

Theorem 7.B.5 is an analogue of Theorem 7.B.4 for upper hemicontinuous differential inclusions. Where the proof of Theorem 7.B.4 relied on the continuity of solutions to (D) in their initial conditions, the proof of Theorem 7.B.5 takes advantage of the upper hemicontinuity of the map from initial conditions ξ to solutions of (DI) starting from ξ .

Theorem 7.B.5. *Let $Y \subset X$ be relatively open and inescapable under (DI). Let $L : Y \rightarrow \mathbf{R}$ be C^1 and satisfy (i) $\frac{\partial L}{\partial v}(x) \equiv \nabla L(x)'v \leq 0$ for all $v \in V(x)$ and $x \in Y$, and (ii) $[\mathbf{0} \notin V(x)$ implies that $\frac{\partial L}{\partial v}(x) < 0]$ for all $v \in V(x)$ and $x \in Y$. Then for all solutions $\{x_t\}$ of (DI) with $x_0 \in Y$, we have that $\omega(\{x_t\}) \subseteq \{x \in Y : \mathbf{0} \in V(x)\}$.*

Proof. Suppose that $\chi \in \omega(\{x_t\})$, but that $\mathbf{0} \notin V(\chi)$. Then $\frac{\partial L}{\partial v}(\chi) < 0$ for all $v \in V(\chi)$. Thus, since $V(\chi)$ is compact by assumption, there exists a $b > 0$ such that $\frac{\partial L}{\partial v}(\chi) < -b$ for all $v \in V(\chi)$. Because V is upper hemicontinuous and L is C^1 , it follows that $\frac{\partial L}{\partial v}(\hat{\chi}) < -\frac{b}{2}$ for all $\hat{v} \in V(\hat{\chi})$ and all $\hat{\chi}$ sufficiently close to χ . So since V is bounded, there is a time $u \in (0, 1]$ such that all solutions $\{y_t\}$ of (DI) with $y_0 = \chi$ satisfy

$$(7.41) \quad L(y_t) \leq L(y_s) \leq L(\chi) - \frac{bs}{2} \quad \text{for all } s \in [0, u] \text{ and } t > s.$$

Now let $\{t_k\}_{k=1}^\infty$ be a divergent sequence of times such that $\lim_{k \rightarrow \infty} x_{t_k} = \chi$, and for each k , define the trajectory $\{x_t^k\}_{t \geq 0}$ by $x_t^k = x_{t+t_k}$. Since the set of continuous trajectories $C_{[0,T]}(X)$ is compact in the sup norm topology, the sequence of trajectories $\{\{x_t^k\}\}_{k=1}^\infty$ has a convergent subsequence, which we take without loss of generality to be $\{\{x_t^k\}\}_{k=1}^\infty$ itself. We call the limit of this subsequence $\{\hat{y}_t\}$. Evidently, $\hat{y}_0 = \chi$.

Given our conditions on the correspondence V , the set-valued map

$$\hat{\chi} \mapsto \{\{x_t\} : \{x_t\} \text{ is a solution to (DI) with } x_0 = \hat{\chi}\}$$

is upper hemicontinuous with respect to the sup norm topology on $C_{[0,T]}(X)$ (see Appendix 6.A). It follows that $\{\hat{y}_t\}$ is a solution to (DI). Consequently,

$$(7.42) \quad \lim_{k \rightarrow \infty} x_{t_k+1} = \hat{y}_1, \quad \text{and so} \quad \lim_{k \rightarrow \infty} L(x_{t_k+1}) = L(\hat{y}_1).$$

But Lemma 7.B.1 and inequality (7.41) imply that

$$L(x_t) \geq L(x) > L(\hat{y}_1)$$

for all $t \geq 0$, contradicting the second limit in (7.42). ■

Theorem 7.B.6 is a simple convergence result for differential inclusions. Here the Lyapunov function need only be Lipschitz continuous (rather than C^1), but the condition on the rate of decrease of this function is stronger than in the previous results.

Theorem 7.B.6. *Let $Y \subset X$ be relatively open and inescapable under (DI), and let $L : Y \rightarrow \mathbf{R}_+$ be Lipschitz continuous. Suppose that along each solution $\{x_t\}$ of (DI) with $x_0 \in Y$, we have that $\dot{L}(x_t) \leq -L(x_t)$ for almost all $t \geq 0$. Then $\omega(\{x_t\}) \subset \{x \in Y : L(x) = 0\}$.*

Proof. Observe that

$$L(x_t) = L(x_0) + \int_0^t \dot{L}(x_u) \, du \leq L(x_0) + \int_0^t (-L(x_u)) \, du = L(x_0) e^{-t},$$

where the final equality follows from the fact that $\alpha_0 + \int_0^t (-\alpha_u) \, du$ is the value at time t of the solution to the linear ODE $\dot{\alpha}_t = -\alpha_t$ with initial condition $\alpha_0 \in \mathbf{R}$. It follows immediately that $\lim_{t \rightarrow \infty} L(x_t) = 0$. ■

7.B.3 Asymptotically Stable and Globally Asymptotically Stable Sets

Combining Theorem 7.B.2 with Theorem 7.B.4, 7.B.5, or 7.B.6 yields asymptotic stability and global asymptotic stability results for deterministic dynamics. Corollary 7.B.7 offers such a result for the differential equation (D).

Corollary 7.B.7. *Let $A \subseteq X$ be closed, and let $Y \subseteq X$ be a neighborhood of A . Let $L : Y \rightarrow \mathbf{R}_+$ be C^1 with $L^{-1}(0) = A$. If $\dot{L}(x) \equiv \nabla L(x)' V(x) < 0$ for all $x \in Y - A$, then A is asymptotically stable under (D). If in addition $Y = X$, then A is globally asymptotically stable under (D).*

7.C Cooperative Differential Equations

Cooperative differential equations are defined by the property that increases in the value of one component of the state variable increase the growth rates of all other components. Their solutions have appealing monotonicity and convergence properties.

Let \leq denote the standard partial order on \mathbf{R}^n : that is, $x \leq y$ if and only if $x_i \leq y_i$ for all $i \in \{1, \dots, n\}$. We write $x < y$ when $x \leq y$ and $x \neq y$, so that $x_j < y_j$ for some j . Finally, we write $x \ll y$ when $x_i < y_i$ for all $i \in \{1, \dots, n\}$. We call a vector or a matrix *strongly positive* if all of its components are positive; thus, $x \in \mathbf{R}^n$ is strongly positive if $x \gg \mathbf{0}$.

Let $X \subset \mathbf{R}^n$ be a compact convex set that possesses a minimal and a maximal element with respect to the partial order \leq . Let $V : X \rightarrow \mathbf{R}^n$ be a C^1 vector field with $V(x) \in TX(x)$ for all $x \in X$, so that the differential equation

$$(7.43) \quad \dot{x} = V(x)$$

is forward invariant on X . We call the differential equation (7.43) *cooperative* if

$$(7.44) \quad \frac{\partial V_i}{\partial x_j}(x) \geq 0 \quad \text{for all } i \neq j \text{ and } x \in X.$$

Equation (7.43) is *irreducible* if for every $x \in X$ and every nonempty proper subset I of the index set $\{1, \dots, n\}$, there exist indices $i \in I$ and $j \in \{1, \dots, n\} - I$ such that $\frac{\partial V_i}{\partial x_j}(x) \neq 0$. An obvious sufficient condition for (7.43) to be irreducible is that it be *strongly cooperative*, meaning that the inequality in condition (7.44) is strict for all $i \neq j$ and $x \in X$.

In Appendix 4.A.3, we saw how to represent all solutions to the dynamic (7.43) simultaneously via the *semiflow* $\phi : \mathbf{R}_+ \times X \rightarrow X$, defined by $\phi_t(\xi) = x_t$, where $\{x_t\}_{t \geq 0}$ is the solution to (7.43) with initial condition $x_0 = \xi$. We say that the semiflow ϕ is *monotone* if $x \leq y$ implies that $\phi_t(x) \leq \phi_t(y)$ for all $t \geq 0$: that is, weakly ordered initial conditions induce weakly ordered solution trajectories. If in addition $x < y$ implies that $\phi_t(x) \ll \phi_t(y)$ for all $t > 0$, we say that ϕ is *strongly monotone*.

Theorem 7.C.1 tells us that cooperative irreducible differential equations generate strongly monotone semiflows.

Theorem 7.C.1. *Suppose that $\dot{x} = V(x)$ is cooperative and irreducible. Then*

- (i) *For all $t > 0$, the derivative matrix of its semiflow ϕ is strongly positive: $D\phi_t(x) \gg \mathbf{0}$.*
- (ii) *The semiflow ϕ is strongly monotone.*

For the intuition behind Theorem 7.C.1, let $\{x_t\}$ and $\{y_t\}$ be solutions to (7.43) with $x_0 < y_0$. Suppose that at some time $t > 0$, we have that $x_t \leq y_t$ and $(x_t)_i = (y_t)_i$. If we could show that $V_i(x_t) \leq V_i(y_t)$, then it seems reasonable to expect that $(x_{t+\varepsilon})_i$ will not be able to surpass $(y_{t+\varepsilon})_i$. But since x_t and y_t only differ in components other than i , the vector

$z = y_t - x_t \geq \mathbf{0}$ has $z_i = 0$, and so

$$V_i(y_t) - V_i(x_t) = \int_0^1 \nabla V_i(x_t + \alpha z)' z \, d\alpha = \int_0^1 \sum_{j \neq i} \frac{\partial V_i}{\partial x_j}(x_t + \alpha z) z_j \, d\alpha.$$

The final expression is nonnegative as long as $\frac{\partial V_i}{\partial x_j} \geq 0$ whenever $j \neq i$.

The next theorem sets out the basic properties of strongly monotone semiflows on X . To state this result, we let $C(\phi) = \{x \in X : \omega(x) = \{x^*\} \text{ for some } x^* \in RP(\phi)\}$ denote the set of initial conditions from which the semiflow ϕ converges to a rest point. Also, let $\Omega(\phi) = \bigcup_{x \in X} \omega(x)$ be the set of ω -limit points under ϕ .

Theorem 7.C.2. *Suppose that the semiflow ϕ on X is strongly monotone. Then*

- (i) *(Convergence criteria) If $\phi_T(x) \geq x$ for some $T > 0$, then $\omega(x)$ is periodic with period T . If $\phi_t(x) \geq x$ over some nonempty open interval of times, then $x \in C(\phi)$.*
- (ii) *(Unordered ω -limit sets) If $x, y \in \omega(z)$, then $x \not\prec y$ and $y \not\prec x$.*
- (iii) *(Minimal and maximal rest points) Let $\underline{x} = \min X$ and $\bar{x} = \max X$. Then $\underline{x}^* = \min RP(\phi)$ and $\bar{x}^* = \max RP(\phi)$ exist; in fact, $\omega(\underline{x}) = \underline{x}^*$ and $\omega(\bar{x}) = \bar{x}^*$. Moreover, $[\underline{x}^*, \bar{x}^*]$ contains $\Omega(\phi)$ and is globally asymptotically stable.*

Proof. (i) If $\phi_T(x) \geq x$, then $\phi_{(n+1)T}(x) \geq \phi_{nT}(x)$ for all positive integers n , so monotonicity and the compactness of X imply that $\lim_{n \rightarrow \infty} \phi_{nT}(x) = y$ for some $y \in X$. By the continuity and group properties of the flow,

$$\phi_{t+T}(y) = \phi_{t+T} \left(\lim_{n \rightarrow \infty} \phi_{nT}(x) \right) = \lim_{n \rightarrow \infty} \phi_{t+(n+1)T}(x) = \lim_{n \rightarrow \infty} \phi_t(\phi_{(n+1)T}(x)) = \phi_t(y),$$

so the flow from y is T -periodic. A continuity argument shows that the orbit from y is none other than $\omega(x)$. The proof of the second claim is omitted.

(ii) Suppose that $x, y \in \omega(z)$ and that $x < y$. Since ϕ is strongly monotone, and by the continuity of $\phi_t(\xi)$ in ξ , there are neighborhoods $N_x, N_y \subset X$ of x and y and a time $T > 0$ such that $\phi_T(N_x) \ll \phi_T(N_y)$. Choose $\tau_y > \tau_x > 0$ such that $\phi_{\tau_x}(z) \in N_x$ and $\phi_{\tau_y}(z) \in N_y$. Then for all t close enough to τ_y ,

$$\phi_{\tau_x+T}(z) \ll \phi_{t+T}(z) = \phi_{t-\tau_x}(\phi_{\tau_x+T}(z)).$$

Therefore, part (i) implies that $\omega(z)$ is a singleton.

(iii) Since \underline{x} and \bar{x} are the minimal and maximal points in X , part (i) implies that $\omega(\underline{x}) = \underline{x}^*$ and $\omega(\bar{x}) = \bar{x}^*$ for some $\underline{x}^*, \bar{x}^* \in RP(\phi)$. Hence, if $x \in X \subseteq [\underline{x}, \bar{x}]$, then $\phi_t(\underline{x}) \leq \phi_t(x) \leq \phi_t(\bar{x})$ for all $t \geq 0$, so taking limits yields $\underline{x}^* \leq \omega(x) \leq \bar{x}^*$; thus, $\Omega(\phi) \subseteq [\underline{x}^*, \bar{x}^*]$. Finally, if

$[\underline{x}^*, \bar{x}^*] \subseteq [y, z] \subseteq X$, then $x \in [y, z]$ implies that $\phi_t(x) \in [\phi_t(y), \phi_t(z)] \subseteq [y, z]$, so $[\underline{x}^*, \bar{x}^*]$ is Lyapunov stable, and hence globally asymptotically stable by the previous argument. ■

If the derivative matrices of the semiflow are strongly positive, one can obtain even stronger results, including the convergence of solution trajectories from generic initial conditions to rest points.

Theorem 7.C.3. *Suppose that the semiflow ϕ on X is strongly monotone, and that its derivative matrices $D\phi_t(x)$ are strongly positive for all $t > 0$. Then*

- (i) *(Limit set dichotomy) If $x < y$, then either $\omega(x) < \omega(y)$, or $\omega(x) = \omega(y) = \{x^*\}$ for some $x^* \in RP(\phi)$.*
- (ii) *(Generic convergence to equilibrium) $C(\phi)$ is an open, dense, full measure subset of X .*

7.N Notes

Section 7.1. The results in Section 7.1.1 are proved for symmetric normal form games in Hofbauer (2000), the seminal reference on Lyapunov functions for evolutionary dynamics. Global convergence in all potential games of dynamics satisfying positive correlation is proved in Sandholm (2001a), building on earlier work of Hofbauer and Sigmund (1988) and Monderer and Shapley (1996b). Convergence of perturbed best response dynamics in potential games is proved by Hofbauer and Sandholm (2007).

Shahshahani (1979), building on the early work of Kimura (1958), showed that the replicator dynamic for a potential game is a gradient dynamic after a “change in geometry”—that is, after the introduction of an appropriate Riemannian metric on $\text{int}(X)$. Subsequently, Akin (1979, 1990) proved that Shahshahani’s (1979) result can also be represented using the change of variable presented in Theorem 7.1.9. The direct proof offered in the text is from Sandholm et al. (2008).

Section 7.2. Theorem 7.2.1 is due to Nagurney and Zhang (1997); the proof in the text is from Sandholm et al. (2008). Theorem 7.2.4 was first proved for normal form games with an interior ESS by Hofbauer et al. (1979) and Zeeman (1980). Akin (1990, Theorem 6.4) and Aubin (1991, Section 1.4) extend this result to nonlinear single population games, while Cressman et al. (2001) extend it to linear multipopulation games.

Section 7.2.2 follows Hofbauer and Sandholm (2007, 2009). These papers take inspiration from Hart and Mas-Colell (2001), which points out the role of integrability in models of regret-based learning in repeated normal form games. Hofbauer (2000) proves the convergence of the BNN, best response, and perturbed best response dynamics in normal form games with an interior ESS. A proof of the existence of a cycle in Example 7.2.6 can

be found in Hofbauer and Sandholm (2009); this reference also contains a statement and proof of the version of Danskin's Envelope Theorem cited in the text. The probabilistic characterization of integrability alluded to the text is presented in Sandholm (2006). For subdifferentials of convex functions, see Hiriart-Urruty and Lemaréchal (2001); their Example D.3.4 is especially relevant to our discussion in the text.

Smith (1984) proves Theorem 7.2.11 for his dynamic; the general result presented here is due to Hofbauer and Sandholm (2009).

Kojima and Takahashi (2007) consider a class of normal form games called *anti-coordination games*, in which at each state x , the worst response to x is always in the support of x . They prove (see also Hofbauer (1995b)) that such games must have a unique equilibrium, that this equilibrium is interior, and that it is globally asymptotically stable under the best response dynamic. However, also they present an example (due to Hofbauer) showing that neither the replicator dynamic nor the logit dynamic need converge in these games, the latter even at arbitrarily low noise levels.

Section 7.3. Section 7.3.1 follows Berger (2007b). Exercise 7.3.1(ii) is due to Hofbauer (1995b), and Lemma 7.3.2(i) is due to Monderer and Sela (1997). It is worth noting that Theorem 7.3.4 extends immediately to ordinal supermodular games (also known as quasi-supermodular games; see Milgrom and Shannon (1994)). Moreover, since ordinal potential games (Monderer and Shapley (1996b)) are defined by the absence of cycles of improvement steps, a portion of the proof of Theorem 7.3.4 establishes the convergence of simple solutions of (BR) in nondegenerate ordinal potential games.

Section 7.3.2 follows Hofbauer and Sandholm (2002, 2007).

Section 7.4. Akin (1980) shows that starting from any interior population state, the replicator dynamic eliminates strategies that are strictly dominated by a pure strategy. Versions of Theorems 7.4.4 and 7.4.5 can be found in Nachbar (1990) and Samuelson and Zhang (1992); see also Hofbauer and Weibull (1996).

Section 7.A. For properties of ω -limit sets of differential equations, see Robinson (1995); for ω -limit sets of differential inclusions, see Benaïm et al. (2005). For applications of chain recurrence in the theory of learning in games, see Benaïm and Hirsch (1999a), Hofbauer and Sandholm (2002), and Benaïm et al. (2005, 2006). The Fundamental Theorem of Dynamical Systems is due to Conley (1978); see Robinson (1995) for a textbook treatment. Other good general references on notions of recurrence for differential equations include Nemytskii and Stepanov (1960), Akin (1993), and Benaïm (1998, 1999).

Section 7.B. The standard reference on Lyapunov functions for flows is Bhatia and Szegő (1970).

Section 7.C. The standard reference on cooperative differential equations and monotone

dynamical systems is Smith (1995). Theorems 7.C.1, 7.C.2(i), 7.C.2(ii), and 7.C.3(i) in the text are Smith's (1995) Theorems 4.1.1, 1.2.1, 1.2.3, and 2.4.5, respectively. Theorem 7.C.3(ii) combines Theorem 2.4.7 of Smith (1995) with Theorem 1.1 of Hirsch (1988), the latter after a reversal of time.

Local Stability under Evolutionary Dynamics

8.0 Introduction

In Chapter 7, we analyzed classes of games in which many evolutionary dynamics converge to equilibrium from all or most initial conditions. While we argued in Chapter 3 that games from many applications lie in these classes, it is certain that at least as many interesting games do not.

In cases where global convergence results are not available, one can turn instead to analyses of local stability. If a society somehow finds itself playing a particular equilibrium, how can we tell whether this equilibrium will persist in the face of occasional, small disturbances in behavior? This chapter introduces a refinement of Nash equilibrium—that of an *evolutionarily stable state* (or *ESS*)—that captures the robustness of the equilibrium to invasion by small groups exhibiting different aggregate behaviors. The main results in this chapter show that states satisfying a version of this concept called *regular Taylor ESS* are locally stable under many evolutionary dynamics.

We will see that games with an ESS share some structural properties with stable games, at least in the neighborhood of the ESS. Taking advantage of this connection, we show in Section 8.4 how to establish local stability of ESS under some dynamics through the use of local Lyapunov functions. Our results here build on our analyses in Section 7.2, where we constructed Lyapunov functions for many dynamics for use in stable games.

The other leading approach to local stability analysis is linearization. Given a rest point of a nonlinear (but smooth) dynamic, one can approximate the behavior of the dynamic in a neighborhood of the rest point by studying an appropriate linear dynamic: namely, the one defined by the derivative matrix of the nonlinear dynamic, evaluated at the rest

point in question. In Sections 8.5 and 8.6, we use linearization to study the two families of smooth dynamics introduced in Chapters 5 and 6: the imitative dynamics, and the perturbed best response dynamics. Surprisingly, this analysis will lead us to a deep and powerful connection between the replicator and logit dynamics, one that seems difficult to reach by other means.

It is worth noting now that linearization is also very useful for establishing *instability* results. For this reason, the techniques we develop in this chapter will be a very important ingredient of our analyses in Chapter 9, where we study nonconvergent dynamics.

The first two sections of the chapter formally establish some results that were hinted at in earlier chapters. In Section 8.1, we indicate two senses in which a non-Nash rest point of an imitative dynamic cannot be stable. In Section 8.2, we show that under most dynamics, a Nash equilibrium of a potential game is locally stable if and only if it is a local maximizer of potential.

The linearization techniques used in Sections 8.5 and 8.6 and in Chapter 9 require a working knowledge of matrix analysis and linear differential equations; we present these topics in detail in Appendices 8.A and 8.B. The main theorems of linearization theory are themselves presented in Appendix 8.C.

8.1 Non-Nash Rest Points of Imitative Dynamics

We saw in Chapters 5 and 6 that under five of our six classes of evolutionary dynamics, rest points are identical to Nash equilibria (or to perturbed versions thereof). The lone exception is the imitative dynamics: Theorem 5.4.21 shows that the rest points of these dynamics are the restricted equilibria, a set that includes not only the Nash equilibria, but also any state that would be Nash equilibria were the strategies unused at that state removed from the game. Theorem 5.7.1 established one sense in which these extra rest points are fragile: by combining a small amount of a “better behaved” dynamic with an imitative dynamic, one obtains a new dynamic that satisfies Nash stationarity. But we mentioned in Section 5.4.6 that this fragility can be expressed more directly: there we claimed that non-Nash rest points of imitative dynamics cannot be locally stable, and so are not plausible predictions of play.

We are now in a position to formally establish this last claim. Recall from Observation 5.4.16 that imitative dynamics exhibit *monotone percentage growth rates*: they can be expressed in the form

$$(8.1) \quad \dot{x}_i^p = x_i^p G_i^p(x),$$

with the percentage growth rates $G_i^p(x)$ ordered by payoffs $F_i^p(x)$ as in equation (5.18). This fact drives our instability result.

Theorem 8.1.1. *Let V_F be an imitative dynamic for population game F , and let \hat{x} be a non-Nash rest point of V_F . Then \hat{x} is not Lyapunov stable under V_F , and no interior solution trajectory of V_F converges to \hat{x} .*

Proof. ($p = 1$) Since \hat{x} is a restricted equilibrium that is not a Nash equilibrium, each strategy j in the support of \hat{x} satisfies $F_j(\hat{x}) = \bar{F}(\hat{x})$, and any best response i to \hat{x} is an unused strategy that satisfies $F_i(\hat{x}) > \bar{F}(\hat{x})$. Also, since \hat{x} is a rest point of V_F , equation (8.1) implies that each j in the support of \hat{x} has $G_j(\hat{x}) = 0$. Thus, monotonicity of percentage growth rates implies that $G_i(\hat{x}) > G_j(\hat{x}) = 0$, and so the continuity of G_i implies $G_i(x) \geq k > 0$ on some small neighborhood O of \hat{x} .

Now let $\{x_t\}$ be an interior solution trajectory of V_F (see Theorem 5.4.14). Then if $x_s \in O$ for all $s \in (t, u)$, it follows that

$$\log((x_u)_i) - \log((x_t)_i) = \int_t^u \left(\frac{d}{ds} \log((x_s)_i) \right) ds = \int_t^u \frac{(\dot{x}_s)_i}{(x_s)_i} ds = \int_t^u G_i(x_s) ds \geq k(u-t).$$

Rearranging and exponentiating yields

$$(x_u)_i \geq (x_t)_i \exp(k(u-t)).$$

Thus, during intervals that x_s is in O , $(x_s)_i$ is increasing. This immediately implies that there is no neighborhood O' of \hat{x} such that solutions starting in O' stay in O , and so \hat{x} is not Lyapunov stable. Also, since $(x_t)_i$ cannot decrease inside $O \cap \text{int}(X)$, no interior solution trajectory can converge to \hat{x} . ■

8.2 Local Stability in Potential Games

We saw in Section 7.1 that in potential games, the potential function serves as a strict Lyapunov function for any evolutionary dynamic satisfying positive correlation (PC); solution trajectories of such dynamics ascend the potential function and converge to connected sets of rest points. For dynamics that also satisfy Nash stationarity (NS), these sets consist entirely of Nash equilibria.

That the potential function is a strict Lyapunov function has important implications for local stability of sets of rest points. Call $A \subseteq X$ a *local maximizer set* of the potential function $f : X \rightarrow \mathbf{R}$ if it is connected, if f is constant on A , and if there exists a neighborhood O of

A such that $f(x) > f(y)$ for all $x \in A$ and all $y \in O - A$. Theorem 3.1.7 implies that such a set consists entirely of Nash equilibria. We call the set $A \subseteq NE(F)$ *isolated* if there is a neighborhood of A that does not contain any Nash equilibria other than those in A .

If the value of f is nondecreasing along solutions of a dynamic, then they cannot escape from a neighborhood of a local maximizer set. If the value of f is actually increasing in this neighborhood, then solutions in the neighborhood should converge to the set. This is the content of the following theorem.

Theorem 8.2.1. *Let F be a potential game with potential function f , let V_F be an evolutionary dynamic for F , and suppose that $A \subseteq NE(F)$ is a local maximizer set of f .*

- (i) *If V_F satisfies positive correlation (PC), then A is Lyapunov stable under V_F .*
- (ii) *If in addition V_F satisfies Nash stationarity (NS) and A is isolated, then A is an asymptotically stable set under V_F .*

Proof. Part (i) of the theorem follows immediately from Lemma 7.1.1 and Theorem 7.B.2. To prove part (ii), note that (NS), (PC), and the fact that A is isolated imply that there is a neighborhood O of A such that $\dot{f}(x) = \nabla f(x)'V_F(x) > 0$ for all $x \in O - A$. Corollary 7.B.7 then implies that A is asymptotically stable. ■

For dynamics satisfying (PC) and (NS), being an isolated local maximizer set is not only a sufficient condition for being asymptotically stable; it is also necessary.

Theorem 8.2.2. *Let F be a potential game with potential function f , let V_F be an evolutionary dynamic for F that satisfies (PC) and (NS). Suppose that $A \subseteq NE(F)$ is a smoothly connected asymptotically stable set under V_F . Then A is an isolated local maximizer set of f .*

Proof. Since A is a smoothly connected set of Nash equilibria, Exercise 3.1.15 implies that f takes some fixed value c throughout A . Now let ξ be an initial condition in $O - A$, where O is the basin of attraction of A . Then $\omega(\xi) \subseteq A$. But since f is a strict Lyapunov function for V_F , it follows that $f(\xi) < c$. Since $\xi \in O - A$ was arbitrary, we conclude that A is an isolated local maximizer set. ■

Theorems 8.2.1 and 8.2.2 allow us to characterize locally stable rest points for dynamics satisfying positive correlation (PC). Since the best response and perturbed best response dynamics do not satisfy this condition, the former because of lack of smoothness and the latter because of the perturbations, Theorems 8.2.1 and 8.2.2 do not apply.

In the case of the best response dynamic, Theorem 6.1.9 establishes analogues of (NS) and (PC), which in turn imply that solution trajectories ascend the potential function and converge to Nash equilibrium (Theorem 7.1.4). By using these results along with the arguments above, we obtain

Theorem 8.2.3. *Let F be a potential game with potential function f , let V_F be the best response dynamic for F , and let $A \subseteq NE(F)$ be smoothly connected. Then A is an isolated local maximizer set of f if and only if A is asymptotically stable under V_F .*

In the case of perturbed best response dynamics, the roles of conditions (PC) and (NS) are played by virtual positive correlation and perturbed stationarity (Theorem 6.2.13 and Observation 6.2.10). These in turn ensure that the dynamics ascend the perturbed potential function

$$\tilde{f}(x) = f(x) - \sum_{p \in \mathcal{P}} m^p v^p \left(\frac{1}{m^p} x^p \right)$$

(Theorem 7.1.6). Substituting these results into the arguments above yields

Theorem 8.2.4. *Let F be a potential game with potential function f , let $V_{F,v}$ be the perturbed best response dynamic for F generated by the admissible deterministic perturbations $v = (v^1, \dots, v^p)$, and let $A \subseteq PE(F, v)$ be smoothly connected. Then A is an isolated local maximizer set of \tilde{f} if and only if A is asymptotically stable under $V_{F,v}$.*

8.3 Evolutionarily Stable States

In the remainder of this chapter, we introduce the notion of an *evolutionarily stable state* and show that it provides a sufficient condition for local stability under a wide range of evolutionary dynamics, including all of the dynamics that were shown to be globally convergent in stable games in Chapter 7.

The birth of evolutionary game theory can be dated to the definition of an *evolutionarily stable strategy* in single-population matching environments by Maynard Smith and Price. In introducing this concept, these authors envisioned model of evolution quite different from the dynamic models considered in this book. Maynard Smith and Price focused on monomorphic populations—populations whose members all choose the same strategy—but allowed this common strategy to be a mixed strategy. Their notion of an evolutionarily stable strategy is meant to capture the capacity of a monomorphic population to resist invasion by a monomorphic mutant group whose members play some alternative mixed strategy.

Clearly, Maynard Smith and Price’s model is different from the polymorphic-population, pure-strategist model studied in this book. Even so, Maynard Smith and Price’s formal definition of an evolutionarily stable strategy remains coherent in our setting: indeed, we

adopt it as our definition of an evolutionarily stable state, allowing us use the acronym “ESS” without regard for context.

The interpretation of ESS in the polymorphic population setting is somewhat strained, since it relies on comparisons between the aggregate payoffs of the incumbent and invading populations, rather than on direct comparisons of the payoffs to different strategies. Still, the fact that the ESS condition is sufficient for stability under many population dynamics amply justifies its use.

Even restricting attention to single-population settings, the literature offers many alternate definitions of ESS and of related concepts. As we shall see, the differences between these definitions are rather subtle, and in normal form environments, all of the concepts we consider below are equivalent. But in multipopulation settings, there are two quite distinct definitions of ESS. One is most natural for studying monomorphic models à la Maynard Smith and Price, while the other is most natural for the polymorphic models considered here. It is a definition of the latter type, which we dub *Taylor ESS*, that provides our sufficient condition for stability for multipopulation dynamics.

8.3.1 Single-Population Games

Let us focus first on the case of single-population games. In this context, we call $x \in X$ an *evolutionarily stable state (ESS)* of F if

$$(8.2) \quad \text{There is a neighborhood } O \text{ of } x \text{ such that } (y - x)'F(y) < 0 \text{ for all } y \in O - \{x\}.$$

Two equivalent characterizations of ESS are provided in Theorems 8.3.1 and 8.3.5 below. While the conditions from Theorem 8.3.5 are the closest to the original definition of ESS (see the Notes), we use the definition above because it is the most concise.

We can interpret definition (8.2) using the notion of invasion from Section 3.3.3. Fix $x \in X$, and let $y \in O$ be a population state near x . Condition (8.2) requires that if there is an incumbent population whose aggregate behavior is described by y , and an infinitesimal group of invaders whose aggregate behavior is described by x , then the average payoff of the invaders must exceed the average payoff in the incumbent population.

While it is appealingly simple, it is not immediately clear why definition (8.2) should be viewed as a stability condition for x : it considers invasions of other states y by x , while one would expect a stability condition for x to address invasions of x by other states. We argue next that (8.2) is equivalent to a condition that is somewhat more cumbersome, but that is stated in terms of invasions of x by other states.

This new condition is stated in terms of inequalities of the form

$$(8.3) \quad (y - x)'F(\varepsilon y + (1 - \varepsilon)x) < 0.$$

Suppose that an incumbent population with aggregate behavior x is invaded by group with aggregate behavior y , with the invaders making up an ε share of the post-entry population. Inequality (8.3) requires that the average payoffs of the incumbents exceed the average payoffs of the invaders.

Now, consider the following requirement on the population state x :

$$(8.4) \quad \text{There is an } \bar{\varepsilon} > 0 \text{ such that (8.3) holds for all } y \in X - \{x\} \text{ and } \varepsilon \in (0, \bar{\varepsilon}).$$

Condition (8.4) says that x admits a *uniform invasion barrier*: so long as the proportion of invaders in the post-entry population is less than $\bar{\varepsilon}$, the incumbents will receive higher payoffs than the invaders in aggregate. If we define the *invasion barrier* of x against y by

$$b_x(y) = \inf(\{\varepsilon \in (0, 1) : (y - x)'F(\varepsilon y + (1 - \varepsilon)x) \geq 0\} \cup \{1\}),$$

then condition (8.4) says that $b_x(y) \geq \bar{\varepsilon} > 0$ for all $y \in X - \{x\}$.

Theorem 8.3.1 shows that ESSs are characterized by the existence of a uniform invasion barrier.

Theorem 8.3.1. *State $x \in X$ is an ESS if and only if condition (8.4) holds.*

Exercise 8.3.2. Prove Theorem 8.3.1 via the following three steps.

- (i) Show that if condition (8.4) holds, then there is a $\delta > 0$ such that the set $U = \{(1 - \lambda)x + \lambda y : y \in X, \lambda \in [0, b_x(y)]\}$ contains a δ -neighborhood of x . (Hint: Start by representing each point in $X - \{x\}$ as a convex combination of x and a point in the set $C = \{y \in X : y_i = 0 \text{ for some } i \in \text{support}(x)\}$.)

- (ii) Use part (i) to show that condition (8.4) is equivalent to the following condition:

$$(8.5) \quad \text{There is a neighborhood } O \text{ of } x \text{ such that (8.3) holds for all } y \in O - \{x\} \text{ and } \varepsilon \in (0, 1).$$

- (iii) Finally, show that x is an ESS if and only if condition (8.5) holds.

Exercise 8.3.3. Consider the following variant of condition (8.4):

$$(8.6) \quad \text{For each } y \in X - \{x\}, \text{ there is an } \bar{\varepsilon} > 0 \text{ such that (8.3) holds for all } \varepsilon \in (0, \bar{\varepsilon}).$$

This condition requires a positive invasion barrier $b_x(y)$ for each $y \in X - \{x\}$, but does not require uniformity: invasion barriers need not be bounded away from 0.

- (i) Show that if $F(x) = Ax$ is linear, then condition (8.6) is equivalent to condition (8.4), and so characterizes ESS in these games. (Hint: Use the fact that

$$b_x(y) = \begin{cases} \min \left\{ \frac{(x-y)'Ax}{(x-y)'A(x-y)}, 1 \right\} & \text{if } (y-x)'Ax < 0 \text{ and } (y-x)'Ay > 0, \\ 0 & \text{otherwise,} \end{cases}$$

along with the hint from Exercise 8.3.2(i) and a compactness argument.)

- (ii) Construct a three-strategy game with a state x^* that satisfies condition (8.6) but that is not an ESS. (Hint: Let $x^* = (0, \frac{1}{2}, \frac{1}{2})$, and let D_1 and D_2 be closed disks in $X \subset \mathbf{R}^3$ that are tangent to $\text{bd}(X)$ at x^* and whose radii are r_1 and $r_2 > r_1$. Introduce a payoff function of the form $F(x) = -c(x)(x - x^*)$, where $c(x)$ is positive on $\text{int}(D_1) \cup (X - D_2)$ and negative on $\text{int}(D_2) - D_1$. Then use Proposition 8.3.4 below.)

We now turn to the relationship between ESS and Nash equilibrium. As a first step, we show that the former is a refinement of the latter:

Proposition 8.3.4. *Every ESS is an isolated Nash equilibrium.*

Proof. Let x be an ESS of F , let O be the neighborhood posited in condition (8.2), and let $y \in X - \{x\}$. Then for all small enough $\varepsilon > 0$, $x_\varepsilon = \varepsilon y + (1 - \varepsilon)x$ is in O , and so satisfies $(x - x_\varepsilon)'F(x_\varepsilon) > 0$. Simplifying and dividing by ε yields $(x - y)'F(x_\varepsilon) > 0$, so taking ε to zero yields $(y - x)'F(x) \leq 0$. That is, x is a Nash equilibrium. To see that x is isolated, note that if $w \in O - \{x\}$ were a Nash equilibrium, we would have $(w - x)'F(w) \geq 0$, contradicting that x satisfies (8.2). ■

The converse of Proposition 8.3.4 is false: the mixed equilibrium of a two-strategy coordination game provides a simple counterexample.

Theorem 8.3.5 shows precisely what restrictions ESS adds to those already imposed by Nash equilibrium.

Theorem 8.3.5. *Suppose that F is Lipschitz continuous. Then state x is an ESS if and only if*

$$(8.7) \quad x \text{ is a Nash equilibrium: } (y - x)'F(x) \leq 0 \text{ for all } y \in X.$$

$$(8.8) \quad \text{There is a neighborhood } O \text{ of } x \text{ such that for all } y \in O - \{x\}, \\ (y - x)'F(x) = 0 \text{ implies that } (y - x)'F(y) < 0.$$

According to the theorem, an ESS x is a Nash equilibrium that satisfies the following additional property: if a state y near x is an alternative best response to x , then an

infinitesimal group of invaders whose aggregate behavior is described by x can invade an incumbent population playing y .

Exercise 8.3.6. Show that if $F(x) = Ax$ is linear, then condition (8.8) is equivalent to the global condition

$$(8.9) \quad \text{For all } y \in X - \{x\}, (y - x)'F(x) = 0 \text{ implies that } (y - x)'F(y) < 0.$$

Proof of Theorem 8.3.5. That condition (8.2) implies conditions (8.7) and (8.8) follows easily from Proposition 8.3.4. While the converse is immediate if $x \in \text{int}(X)$, in general the proof of the converse requires a delicate argument.

To begin, suppose that $x \in X$ satisfies conditions (8.7) and (8.8). Let $S^* = \text{argmax}_{i \in S} F_i(x)$ be the set of pure best responses to x , let $S^0 = S - S^*$, let $X^* = \{v^* \in X : \text{support}(v^*) \subseteq S^*\}$ be the set of mixed best responses to x , and let $X^0 = \{v^0 \in X : \text{support}(v^0) \subseteq S^0\}$ be the set of states at which all best responses to x are unused. Since x is a Nash equilibrium, we have that $\text{support}(x) \subseteq S^*$ and $x \in X^*$. Moreover, each $y \in X$ can be represented as a convex combination of the form

$$y = (1 - \lambda)y^* + \lambda y^0,$$

where $y^* \in X^*$, $y^0 \in X^0$, and $\lambda = \sum_{j \in S^0} y_j \in [0, 1]$. Evidently, we have $y^* = y$ when $\lambda = 0$, and $y^0 = y$ when $\lambda = 1$; in all other cases, y uniquely determines both y^* and y^0 .

We want to show that for some neighborhood U of x , we have $(y - x)'F(y) < 0$ for all $y \in U - \{x\}$. To do so, we introduce the decomposition

$$(8.10) \quad \begin{aligned} (y - x)'F(y) &= ((1 - \lambda)y^* + \lambda y^0 - x)'F(y) \\ &= \lambda(y^0 - x)'F(y) + (1 - \lambda)(y^* - x)'(F(y) - F(y^*)) + (1 - \lambda)(y^* - x)'F(y^*) \\ &\equiv T_1(y) + T_2(y) + T_3(y). \end{aligned}$$

To prove the theorem, it is enough to establish these two claims:

$$(8.11) \quad T_1(y) + T_2(y) \leq 0 \text{ when } y \in U, \text{ and } T_1(y) + T_2(y) < 0 \text{ when } y \in U - X^*;$$

$$(8.12) \quad T_3(y) \leq 0 \text{ when } y \in U, \text{ and } T_3(y) < 0 \text{ when } y \in U \cap X^* - \{x\}.$$

To accomplish this, let $\|v\| = \sum_{i=1}^n |v_i|$ denote the L^1 norm of $v \in \mathbf{R}^n$, and let K be the Lipschitz constant for F with respect to this norm. Also, let O be the neighborhood from condition (8.8), and choose $r > 0$ so that O contains a ball of radius r centered at x .

Since $x \in X^*$ and $y^0 \in X^0$, we have that $(y^0 - x)'F(x) < 0$. Therefore, since F is continuous,

we can pick a neighborhood U' of x and a positive constant ε such that

$$(8.13) \quad \lambda(y^0 - x)'F(y) < -\lambda\varepsilon \text{ for all } y \in U'.$$

Moreover, we can choose a neighborhood $U \subseteq U'$ of x such that

$$(8.14) \quad \|y - x\| < \min\{\frac{\varepsilon}{20K}, \frac{r}{5}, \frac{1}{2}\} \text{ whenever } y \in U.$$

Suppose that $\|y - x\| \leq \frac{1}{2}$. Since $\text{support}(x) \cap S^0 = \emptyset$, we have that

$$\lambda = \sum_{j \in S^0} |y_j| = \sum_{j \in S^0} |y_j - x_j| \leq \sum_{i \in S^*} |y_i - x_i| + \sum_{j \in S^0} |y_j - x_j| = \|y - x\| \leq \frac{1}{2},$$

and so that $\frac{1}{1-\lambda} \leq 2$. Since $\|v - y\| \leq 2$ for any $v \in X$, it follows that

$$\begin{aligned} \|y^* - x\| &= \left\| \frac{1}{1-\lambda}(y - \lambda y^0) - x \right\| \\ &\leq \|y - x\| + \frac{\lambda}{1-\lambda} \|y^0 - y\| \\ &\leq \|y - x\| + 4\lambda \\ &\leq 5 \|y - x\|. \end{aligned}$$

Equation (8.14) therefore yields

$$(8.15) \quad \|y^* - x\| < \min\{\frac{\varepsilon}{4K}, r\} \text{ whenever } y \in U.$$

Now suppose that $y \in U$. Applying the first bound in (8.15), we find that

$$\begin{aligned} (1 - \lambda)(y^* - x)'(F(y) - F(y^*)) &\leq \|y^* - x\| K \|y - y^*\| \\ &= \|y^* - x\| K \|\lambda(y^0 - y^*)\| \\ &= 2\lambda K \|y^* - x\| \\ &\leq \frac{1}{2}\lambda\varepsilon. \end{aligned}$$

Since $\lambda > 0$ when $y \notin X^*$, combining this bound with (8.13) yields claim (8.11).

The second bound in (8.15) tells us that $y^* \in O$. Since $y^* \in X^*$ as well, we have that $(y^* - x)'F(x) = 0$, so condition (8.8) tells us that

$$(y^* - x)'F(y^*) \leq 0, \text{ with equality if and only if } y^* = x.$$

Since y^* is distinct from x whenever $y \in X^* - \{x\}$, we have claim (8.12). This completes the

proof of the theorem. ■

8.3.2 Multipopulation Games

As we noted at the start of this section, there are two rather different ways of extending the definition of ESS from a single-population setting to a multipopulation setting. If F is a game played by $p \geq 1$ populations, we call $x \in X$ a *Taylor ESS* of F if

$$(8.16) \quad \text{There is a neighborhood } O \text{ of } x \text{ such that } (y - x)'F(y) < 0 \text{ for all } y \in O - \{x\}.$$

(The text of this condition is identical that in (8.2); the only difference is that F may now be a multipopulation game.) We call x a *Cressman ESS* of F if

$$(8.17) \quad \begin{aligned} &\text{There is a neighborhood } O \text{ of } x \text{ such that for all } y \in O - \{x\}, \\ &\text{there is a } p \in \mathcal{P} \text{ such that } (y^p - x^p)'F^p(y) < 0. \end{aligned}$$

Condition (8.17) is identical to condition (8.16) in single-population settings. But in multipopulation settings, it is weaker. Both conditions consider invasions of a collection of incumbent populations $y = (y^1, \dots, y^p) \in O - \{x\}$ by a collection of invading populations $x = (x^1, \dots, x^p)$. For x to be a Taylor ESS, the aggregate payoff of the invading populations must exceed the aggregate payoff of the incumbent populations. But for x to be a Cressman ESS, it is enough that one of the invading populations x^p earn a higher average payoff than the corresponding incumbent population y^p .

If we focus on the evolutionary setting studied by Maynard Smith and Price—that of monomorphic populations of mixed strategists—then the appropriate extension of the ESS concept to multiple populations is Cressman ESS. (See the Notes for an extended discussion.) But to understand the dynamics of behavior in polymorphic populations of pure strategists, Taylor ESS turns out to be the more useful condition. Because of this, we will write “ESS” in place of “Taylor ESS” in some of the discussions below.

Exercise 8.3.7. Let $F = (F^1, F^2)$ be a population game, and let $\hat{F} = (F^1, 2F^2)$ be the game obtained by doubling the payoffs of population 2's strategies. Show that F and \hat{F} have the same Nash equilibria and Cressman ESSs, but that their Taylor ESSs may differ.

Exercise 8.3.8. Suppose that F has no own-population interactions: $F^p(x)$ is independent of x^p for all $p \in \mathcal{P}$. Show that if x^* is a Cressman ESS of F , then it is a pure social state: $(x^*)^p = m^p e_i^p$ for some $i \in S^p$ and $p \in \mathcal{P}$. Of course, this implies that any Taylor ESS is a pure social state as well. See Proposition 3.3.10 for a closely related result in the context

of stable games. (Hint: If x^p is not pure, consider an invasion by $y = (y^p, x^{-p})$, where y^p is an alternate best response to x .)

8.3.3 Regular Taylor ESS

While in Section 8.3.2 we focused on the distinction between Taylor ESS and Cressman ESS, it is also the case that each of these concepts admits a number of characterizations corresponding to those of single population ESS. For example, Theorem 8.3.5 has the following analogue for Taylor ESS

Corollary 8.3.9. *Suppose that F is Lipschitz continuous. Then x is a Taylor ESS if and only if*

$$(8.18) \quad x \text{ is a Nash equilibrium: } (y - x)'F(x) \leq 0 \text{ for all } y \in X.$$

$$(8.19) \quad \begin{aligned} &\text{There is a neighborhood } O \text{ of } x \text{ such that for all } y \in O - \{x\}, \\ &(y - x)'F(x) = 0 \text{ implies that } (y - x)'F(y) < 0. \end{aligned}$$

The text of these conditions is identical to that in conditions (8.7) and (8.8), but F is now allowed to be a multipopulation game.

Some of our local stability results require a slight strengthening of these conditions. To strengthen the Nash equilibrium condition (8.18), we suppose that x is a *quasistrict equilibrium*: within each population, all strategies in use earn the same payoff, a payoff that is strictly greater than that of each unused strategy. (This is a generalization of *strict equilibrium*, which in addition requires x to be a pure state.) To strengthen (8.19), we replace the inequality in this condition with a differential version. All told, we call $x \in X$ a *regular Taylor ESS* if

$$(8.20) \quad x \text{ is a quasistrict equilibrium: } F_i^p(x) = \bar{F}^p(x) > F_j^p(x) \text{ whenever } x_i^p > 0 \text{ and } x_j^p = 0.$$

$$(8.21) \quad \text{For all } y \in X - \{x\}, (y - x)'F(x) = 0 \text{ implies that } (y - x)'DF(x)(y - x) < 0.$$

When $p = 1$, we can call a state x satisfying (8.20) and (8.21) a *regular ESS* for short.

Exercise 8.3.10. (i) Confirm that every regular Taylor ESS is a Taylor ESS.

(ii) Show that condition (8.21) does not change if the implication is only checked for $y \neq x$ in a neighborhood of x , as in condition (8.19).

(iii) Show that if F is linear, then condition (8.21) is equivalent to condition (8.19).

It is useful to have a more concise characterization of regular Taylor ESS. For any set of strategies $I \subset \bigcup_{p \in \mathcal{P}} S^p$, let $\mathbf{R}_I^n = \{y \in \mathbf{R}^n : y_j^p = 0 \text{ whenever } j \notin I\}$ denote the set of vectors in \mathbf{R}^n whose components corresponding to strategies outside I equal zero. Also, let $S(x) \subseteq \bigcup_{p \in \mathcal{P}} S^p$ denote the support of state x . Then it is easy to verify

Observation 8.3.11. *State x is a regular Taylor ESS if and only if it is a quasistrict equilibrium that satisfies*

$$(8.22) \quad z'DF(x)z < 0 \text{ for all nonzero } z \in TX \cap \mathbf{R}_{S(x)}^n.$$

Condition (8.22) resembles the derivative condition we associate with strictly stable games. However, the condition need only hold at the equilibrium, and negative definiteness is only required to hold in directions that move along the face of X on which the equilibrium lies.

8.4 Local Stability via Lyapunov Functions

In this remainder of this chapter, we show that any regular Taylor ESS x^* is locally stable under many evolutionary dynamics. In this section, our approach is to construct a *strict local Lyapunov function* for each dynamic in question: that is, a nonnegative function defined in a neighborhood of x^* that vanishes precisely at x^* and whose value decreases along every solution of the dynamic other than the stationary one at x^* . The results presented in Appendix 7.B show that the existence of such a function ensures the asymptotic stability of x^* .

The similarity between the definitions of regular ESS and of stable games—in particular, the negative semidefiniteness conditions that play a central role in both contexts—suggests the Lyapunov functions for stable games from Section 7.2 as the natural starting points for our stability analyses of ESSs. In some cases—under the projection and replicator dynamics, and whenever the ESS is interior—we will be able to use the Lyapunov functions from Section 7.2 without amendment. But more generally, we will need to modify these functions in order to make them local Lyapunov functions for ESSs.

8.4.1 The Replicator and Projection Dynamics

The analysis is simplest in the cases of the replicator and projection dynamics. In Section 7.2, we proved global convergence of these dynamics in every strictly stable game by showing that measures of “distance” from the game’s unique Nash equilibrium served as global Lyapunov functions. The proofs of these convergence results relied on nothing about the payoff structure of the game apart from the fact that the game’s unique Nash equilibrium is also a GESS.

This observation suggests that if state x^* is a Taylor ESS of an arbitrary population game, the same “distance” functions will serve as *local* Lyapunov functions for x^* under

the two dynamics. We confirm this logic in the following theorem.

Theorem 8.4.1. *Let x^* be a Taylor ESS of F . Then x^* is asymptotically stable under*

- (i) *the replicator dynamic for F ;*
- (ii) *the projection dynamic for F .*

Exercise 8.4.2. Prove Theorem 8.4.1 by showing that the functions H_{x^*} and E_{x^*} from Theorems 7.2.4 and 7.2.1 define strict local Lyapunov functions for the two dynamics.

8.4.2 Target and Pairwise Comparison Dynamics: Interior ESS

In proving convergence results for other classes of dynamics in Section 7.2, we relied directly on the negative semidefiniteness condition (3.16) that characterizes stable games. If a game admits an interior ESS that satisfies the strict inequalities in (8.22), then condition (3.16) holds in a neighborhood of the ESS. This allows us again to use the Lyapunov functions from Section 7.2 without amendment to prove local stability results.

Theorem 8.4.3. *Let $x^* \in \text{int}(X)$ be a regular Taylor ESS of F . Then x^* is asymptotically stable under*

- (i) *any separable excess payoff dynamic for F ;*
- (ii) *the best response dynamic for F ;*
- (iii) *any impartial pairwise comparison dynamic for F .*

Exercise 8.4.4. Prove Theorem 8.4.3 by showing that the functions Γ , G , and Ψ from Theorems 7.2.8, 7.2.9, and 7.2.11 define strict local Lyapunov functions for an ESS x^* under the three dynamics.

Rest points of perturbed best response dynamics generally do not coincide with Nash equilibria, and hence with ESSs. Nevertheless, the next exercise indicates that an appropriate negative definiteness condition is still enough to ensure local stability.

Exercise 8.4.5. Let \tilde{x} be a perturbed equilibrium of (F, v) for some admissible deterministic perturbations $v = (v^1, \dots, v^p)$, and suppose that $z'DF(\tilde{x})z < 0$ for all nonzero $z \in TX$. Show that \tilde{x} is isolated in the set of perturbed equilibria, and that the function \tilde{G} from Theorem 7.2.10 defines a strict local Lyapunov function for \tilde{x} . (Hint: To show that \tilde{x} is isolated, use the argument at the end of the proof of Theorem 7.2.10.)

For consistency with our previous results, it is natural to try to prove stability results for perturbed best response dynamics for games with an interior ESS. To do so, we need to assume that the size of the perturbations is “small”, in the hopes that there will be a perturbed equilibrium that is “close” to the ESS. Since the logit dynamic is parameterized by a noise level η , it provides a natural setting for the result we seek.

Theorem 8.4.6. *Let $x^* \in \text{int}(X)$ be a regular Taylor ESS of F . Then for some neighborhood O of x^* and each $\eta > 0$ less than some $\hat{\eta} > 0$, there is a unique $\text{logit}(\eta)$ equilibrium \tilde{x}^η in O , and this equilibrium is asymptotically stable under the $\text{logit}(\eta)$ dynamic. Finally, \tilde{x}^η varies continuously in η , and $\lim_{\eta \rightarrow 0} \tilde{x}^\eta = x^*$.*

Proof. ($p = 1$) Theorem 7.2.10 and Exercises 6.2.6 and 6.2.7 show that for $\eta > 0$, the function

$$\tilde{G}^\eta(x) = \eta \log \left(\sum_{j \in S} \exp(\eta^{-1} \hat{F}_j(x)) \right) + \eta \sum_{j \in S} x_j \log x_j.$$

(with $0 \log 0 \equiv 0$) is a Lyapunov function for the $\text{logit}(\eta)$ dynamic when F is a stable game. If we define

$$\tilde{G}^0(x) \equiv G(x) = \max_{j \in S} \hat{F}_j(x).$$

to be the Lyapunov function for the best response dynamic in stable games, then $\tilde{G}^\eta(x)$ is continuous in (x, η) on $X \times [0, \infty)$,

By Exercise 8.4.4, G defines a strict local Lyapunov function for the best response dynamic at the interior ESS x^* . In particular, x^* is a local minimizer of G : there is an open, convex neighborhood $O \subset X$ of x^* such that $G(x) > G(x^*)$ for all $x \in O - \{x^*\}$. Moreover, since F is C^1 and satisfies $z' DF(x^*) z < 0$ for all nonzero $z \in TX$, we can choose O in such a way that $z' DF(x) z < 0$ for all nonzero $z \in TX$ and $x \in O$.

Because $\tilde{G}^\eta(x)$ is continuous in (x, η) , the Theorem of the Maximum (see the Notes) implies that the map

$$\eta \mapsto \tilde{\beta}(\eta) \equiv \underset{x \in \text{cl}(O)}{\text{argmin}} \tilde{G}^\eta(x).$$

is upper hemicontinuous on $[0, \infty)$. Thus, since $\tilde{\beta}(0) = \{x^*\} \subset O$ (in particular, since $x^* \notin \text{bd}(O)$), there is an $\hat{\eta} > 0$ such that $\tilde{\beta}(\eta) \subset O$ for all $\eta < \hat{\eta}$. This implies that each $\tilde{x}^\eta \in \tilde{\beta}(\eta)$ is a local minimizer of \tilde{G}^η with respect to $\text{cl}(O)$, and so with respect to the full state space X .

Exercise 8.4.5 implies that the value of \tilde{G}^η is decreasing along solutions to the $\text{logit}(\eta)$ dynamic in the set O , implying that each local minimizer \tilde{x}^η is a rest point of this dynamic—indeed, \tilde{x}^η must be an asymptotically stable rest point. Finally, since O is convex, the last paragraph of the proof of Theorem 7.2.10 shows that when $\eta < \hat{\eta}$, $\tilde{\beta}(\eta) \subset O$ is a singleton. This completes the proof of the theorem. ■

8.4.3 Target and Pairwise Comparison Dynamics: Boundary ESS

It remains for us to prove local stability results for boundary ESSs for the dynamics considered in Theorem 8.4.3.

Theorem 8.4.7. *Let x^* be a regular Taylor ESS of F . Then x^* is asymptotically stable under*

- (i) *any separable excess payoff dynamic for F ;*
- (ii) *the best response dynamic for F ;*
- (iii) *any impartial pairwise comparison dynamic for F .*

To prove Theorem 8.4.7, we show that suitably modified versions of the Lyapunov functions for stable games serve as local Lyapunov functions here. Letting $S^p(x^*) = \text{support}((x^*)^p)$ and $C > 0$, we augment the functions Γ , G , and Ψ from Section 7.2 by the function

$$(8.23) \quad \Upsilon_{x^*}(x) = C \sum_{p \in \mathcal{P}} \sum_{j \notin S^p(x^*)} x_j^p,$$

which is proportional to the number of agents using strategies outside the support of x^* .

For intuition, let F be a single population game, and suppose that $x^* \notin \text{int}(X)$ is a regular ESS of F in which at least one pure strategy is unused. Because x^* is a quasistrict equilibrium, all unused strategies earn strictly lower payoffs than strategies in the support of x^* ; because payoffs are continuous, this is true not only at x^* itself, but also at states near x^* . Under any dynamic that respects these payoff differences, solutions from initial conditions near x^* should converge to $X_{x^*} = \{x \in X : S(x) = S(x^*)\}$, the face of X that contains x^* . If x^* is a pure state (i.e., a vertex of X), and thus a strict equilibrium, then X_{x^*} is simply the singleton $\{x^*\}$, and the foregoing argument suffices to prove local stability: see Exercise 8.4.9 below.

Now assume that x^* is neither in the interior nor at a vertex of X , and suppose that we confine our attention to the behavior of the dynamic on face X_{x^*} in the vicinity of x^* . One way to do so is to consider a restricted game in which only the strategies in the support of x^* are available. Observation 8.3.11 tells us that this restricted game resembles a stable game, at least in a neighborhood of x^* . It then follows from Theorem 8.4.3 that under the dynamic for our original game, solutions from initial conditions that are on face X_{x^*} and near x^* must converge to x^* .

To prove Theorem 8.4.7, we need to construct local Lyapunov functions whose values decrease not only along solutions on face X_{x^*} , but also along solutions in the interior of X starting near x^* . To do so, we augment the Lyapunov functions for stable games, Γ , G , and Ψ by the function Υ_{x^*} defined in (8.23), which captures the mass placed on strategies

outside the support of x^* . For instance, in the case of impartial pairwise comparison dynamics, we replace the original Lyapunov function Ψ with the function

$$(8.24) \quad \Psi_{x^*}(x) = \Psi(x) + \Upsilon_{x^*}(x) = \Psi(x) + C \sum_{j \notin S(x^*)} x_j.$$

When the current state x is in the interior of X , the value of Ψ need not decrease. But if x is close to x^* , then the value of Υ_{x^*} does decrease, as agents playing strategies outside of $S(x^*)$ switch to strategies in $S(x^*)$, though as the state approaches the boundary of X , the rate of decrease of Υ_{x^*} approaches zero. To prove the theorem, we must show that if the constant $C > 0$ is large enough, then the value of $C \Upsilon_{x^*}$ always falls fast enough to compensate for any growth in the value of Ψ , so that all told, the value of Ψ_{x^*} falls: see Lemma 8.4.8 and Exercise 8.4.11.

Below provide a detailed proof of Theorem 8.4.7 for the case of impartial pairwise comparison dynamics. We leave the proofs of the other two cases as exercises.

Proof of Theorem 8.4.7(iii). ($p = 1$) Let $\dot{x} = V_F(x)$ be an impartial pairwise comparison dynamic for F . Let $\Psi_{x^*} : X \rightarrow \mathbf{R}$ be the C^1 function defined in equation (8.24), where Ψ is defined in Theorem 7.2.11, and where the constant $C > 0$ will be determined below.

Since V_F is an impartial pairwise comparison dynamic, Theorem 7.2.11 shows that the function Ψ is nonnegative, with $\Psi(x) = 0$ if and only if $x \in NE(F)$. It follows that Ψ_{x^*} too is nonnegative, with $\Psi_{x^*}(x) = 0$ if and only if x is a Nash equilibrium of F with $\text{support}(x) \subseteq \text{support}(x^*)$. Thus, since x^* is an ESS, it is isolated in the set of Nash equilibria (Proposition 8.3.4), so there is a neighborhood O of x^* on which x^* is the unique zero of Ψ_{x^*} . If we can show that there is also a neighborhood O' of x^* such that $\dot{\Psi}_{x^*}(x) < 0$ for all $x \in O' - \{x^*\}$, then Ψ_{x^*} is a strict local Lyapunov function for x^* , so the conclusion of the theorem will follow from Corollary 7.B.7.

To reduce the amount of notation in the analysis to come, let $\mathbf{1}^0 \in \mathbf{R}^n$ be the vector whose j th component equals 0 if $j \in \text{support}(x^*)$ and equals 1 otherwise, so that $(\mathbf{1}^0)'x$ is the mass of agents who use strategies outside the support of x^* at state x . Then we can write $\Psi_{x^*}(x) = \Psi(x) + C (\mathbf{1}^0)'x$, and so can express the time derivative of Ψ_{x^*} as

$$\dot{\Psi}_{x^*}(x) = \dot{\Psi}(x) + C (\mathbf{1}^0)'\dot{x}.$$

Now the proof of Theorem 7.2.11 shows that the time derivative of Ψ satisfies

$$\dot{\Psi}(x) \leq \dot{x}'DF(x)\dot{x},$$

with equality holding precisely at the Nash equilibria of F . To finish the proof, it is enough to show that

$$\dot{x}'DF(x)\dot{x} + C(\mathbf{1}^0)'\dot{x} \leq 0$$

for all $x \in O' - \{x^*\}$. This follows directly from the following lemma, choosing $C \geq M/N$.

Lemma 8.4.8. *Let $\dot{x} = V_F(x)$ be a pairwise comparison dynamic for F , and let x^* be a regular ESS of F . Then there is a neighborhood O' of x^* and constants $M, N > 0$ such that for all $x \in O'$,*

- (i) $\dot{x}'DF(x)\dot{x} \leq M(\mathbf{1}^0)'\dot{x}$;
- (ii) $(\mathbf{1}^0)'\dot{x} \leq -N(\mathbf{1}^0)'\dot{x}$.

Proof. Suppose without loss of generality that $S(x^*) = \text{support}(x^*)$ is given by $\{1, \dots, n^*\}$. Then to complement $\mathbf{1}^0 \in \mathbf{R}^n$, let $\mathbf{1}^* \in \mathbf{R}^n$ be the vector whose first n^* components equal 1 and whose remaining components equal 0, so that $\mathbf{1} = \mathbf{1}^* + \mathbf{1}^0$. Next, decompose the identity matrix I as $I^* + I^0$, where $I^* = \text{diag}(\mathbf{1}^*)$ and $I^0 = \text{diag}(\mathbf{1}^0)$, and finally, decompose I^* as $\Phi^* + \Xi^*$, where $\Xi^* = \frac{1}{n^*}\mathbf{1}^*(\mathbf{1}^*)'$ and $\Phi^* = I^* - \Xi^*$. Notice that Φ^* is the orthogonal projection of \mathbf{R}^n onto $\mathbf{R}_0^n \cap \mathbf{R}_{S(x^*)}^n$, where $\mathbf{R}_0^n = \{z \in \mathbf{R}^n : \mathbf{1}'z = 0\}$ and $\mathbf{R}_{S(x^*)}^n = \{z \in \mathbf{R}^n : z_j = 0 \text{ whenever } j \notin S(x^*)\}$, and that $I = \Phi^* + \Xi^* + I^0$.

Using this decomposition of the identity matrix, we can write

$$\begin{aligned} \dot{x}'DF(x)\dot{x} &= ((\Phi^* + \Xi^* + I^0)\dot{x})'DF(x)((\Phi^* + \Xi^* + I^0)\dot{x}) \\ (8.25) \quad &= (\Phi^*\dot{x})'DF(x)(\Phi^*\dot{x}) + ((\Xi^* + I^0)\dot{x})'DF(x)\dot{x} + (\Phi^*\dot{x})'DF(x)((\Xi^* + I^0)\dot{x}). \end{aligned}$$

Since x^* is a regular ESS, we know that $z'DF(x^*)z < 0$ for all nonzero $z \in TX \cup \mathbf{R}_{S(x^*)}^n$. Thus, since $DF(x)$ is continuous in x , there is a neighborhood \hat{O} of x^* on which the first term of (8.25) is nonpositive.

Turning to the second term, note that since $\mathbf{1}'\dot{x} = 0$ and $(\mathbf{1}^0)' = \mathbf{1}'I^0$, we have that

$$(\Xi^* + I^0)\dot{x} = (\frac{1}{n^*}\mathbf{1}^*(\mathbf{1}^*)' + I^0)\dot{x} = (-\frac{1}{n^*}\mathbf{1}^*(\mathbf{1}^0)' + I^0)\dot{x} = ((I - \frac{1}{n^*}\mathbf{1}^*\mathbf{1}')I^0)\dot{x}.$$

Let $\|A\|$ denote the spectral norm of the matrix A (see Appendix 8.A.6). Then applying spectral norm inequalities and the Cauchy-Schwarz inequality, we find that

$$(8.26) \quad ((\Xi^* + I^0)\dot{x})'DF(x)\dot{x} = ((I - \frac{1}{n^*}\mathbf{1}^*\mathbf{1}')I^0\dot{x})'DF(x)\dot{x} \leq |I^0\dot{x}| \left\| I - \frac{1}{n^*}\mathbf{1}^*\mathbf{1}' \right\| \|DF(x)\| |\dot{x}|.$$

Since $DF(x)$, $V_F(x)$, and $\rho_{ij}(F(x), x)$ are continuous in x on the compact set X , we can

find constants K and R such that

$$(8.27) \quad \left\| I - \frac{1}{n^*} \mathbf{1}(\mathbf{1}^*)' \right\| \|DF(x)\| |\dot{x}| \leq K \quad \text{and} \quad \max_{i,j \in S} \rho_{ij}(F(x), x) \leq R \quad \text{for all } x \in X.$$

Now since x^* is a quasistrict equilibrium, we have that $F_i(x^*) = \bar{F}(x^*) > F_j(x^*)$ for all $i \in \text{support}(x^*) = \{1, \dots, n^*\}$ and all $j \notin \text{support}(x^*)$. Thus, since the pairwise comparison dynamic satisfies sign preservation (5.25), we have $\rho_{ij}(F(x^*), x^*) = 0$ for such i and j , and because F is continuous, there is a neighborhood $O' \subseteq \hat{O}$ of x^* on which for such i and j we have $F_i(x) > F_j(x)$, and hence $\rho_{ij}(F(x), x) = 0$. From this argument and the bound on ρ_{ij} in (8.27), it follows that for $x \in O'$, we have that

$$\begin{aligned} |I^0 \dot{x}| &= \sqrt{\sum_{j > n^*} |\dot{x}_j|^2} \\ &\leq \sum_{j > n^*} |\dot{x}_j| \\ &= \sum_{j > n^*} \left| \sum_{k \in S} x_k \rho_{kj}(F(x), x) - x_j \sum_{k \in S} \rho_{jk}(F(x), x) \right| \\ &\leq \sum_{j > n^*} \left(\sum_{k \in S} x_k \rho_{kj}(F(x), x) + x_j \sum_{k \in S} \rho_{jk}(F(x), x) \right) \\ &= \sum_{j > n^*} \sum_{k > n^*} x_k \rho_{kj}(F(x), x) + \sum_{j > n^*} x_j \sum_{k \in S} \rho_{jk}(F(x), x) \\ &\leq 2Rn \sum_{j > n^*} x_j \\ &= 2Rn (\mathbf{1}^0)' x. \end{aligned}$$

We therefore conclude that at all $x \in O'$,

$$((\Xi^* + I^0)\dot{x})' DF(x)\dot{x} \leq 2KRn (\mathbf{1}^0)' x.$$

Essentially the same argument provides a similar bound on the third term of (8.25), completing the proof of part (i) of the lemma.

We proceed with the proof of part (ii). Following the line of argument after equation (8.27) above, we note that since x^* is quasistrict and since the pairwise comparison dynamic satisfies sign preservation, we have $\rho_{ji}(F(x^*), x^*) > 0$ and $\rho_{ij}(F(x^*), x^*) = 0$ whenever $i \in \text{support}(x^*) = \{1, \dots, n^*\}$ and $j \notin \text{support}(x^*)$. So, since F and ρ are continuous, sign preservation implies that there is a neighborhood O' of x^* and an $r > 0$ such that

$\rho_{ji}(F(x), x) > r$ and $\rho_{ij}(F(x), x) = 0$ for all $i \leq n^*$, $j > n^*$, and $x \in O'$. Applying this observation and then canceling like terms when both j and k are greater than n^* in the sums below, we find that for all $x \in O'$,

$$\begin{aligned}
(\mathbf{1}^0)' \dot{x} &= \sum_{j>n^*} \dot{x}_j \\
&= \sum_{j>n^*} \left(\sum_{k \in S} x_k \rho_{kj}(F(x), x) - x_j \sum_{k \in S} \rho_{jk}(F(x), x) \right) \\
&= \sum_{j>n^*} \left(\sum_{k>n^*} x_k \rho_{kj}(F(x), x) - x_j \sum_{k \in S} \rho_{jk}(F(x), x) \right) \\
&= - \sum_{j>n^*} x_j \sum_{i \leq n^*} \rho_{ji}(F(x), x) \\
&\leq -r n^* (\mathbf{1}^0)' x.
\end{aligned}$$

This completes the proof of the lemma, and thus the proof of Theorem 8.4.7(i). ■

Exercise 8.4.9. Let e_i be a strict equilibrium of the single population game F : $F_i(e_i) > F_j(e_i)$ for all $j \neq i$. Show that x^* is asymptotically stable under

- (i) any excess payoff dynamic for F ;
- (ii) the best response dynamic for F ;
- (iii) any pairwise comparison dynamic for F .

Note that statements (i) and (iii) are not implied by Theorem 8.4.7, since the excess payoff dynamic in (i) is not assumed to be separable, and the pairwise comparison dynamic in (ii) is not assumed to be impartial. (Hint: In all cases, show that the function Υ_{e_i} defined in (8.23) is a strict local Lyapunov function.)

Exercise 8.4.10. Prove Theorem 8.4.7(ii) (for $p = 1$) by showing that under the best response dynamic, the function

$$G_{x^*}(x) = G(x) + \Upsilon_{x^*}(x) = \max_{y \in X} (y - x)' F(x) + C \sum_{j \notin S(x^*)} x_j$$

is a strict local Lyapunov function for any regular ESS x^* . (Hint: The proof is nearly the same as the one above, but building on the proof of Theorem 7.2.9 instead of the proof of Theorem 7.2.11, and using Theorems 7.B.2 and 7.B.6 in place of Corollary 7.B.7.)

Exercise 8.4.11. Prove Theorem 8.4.7(i) (for $p = 1$) by showing that under the separable excess payoff dynamic with revision protocol τ , the function

$$\Gamma_{x^*}(x) = \Gamma(x) + \Upsilon_{x^*}(x) = \sum_{i \in S} \int_0^{\hat{F}_i(x)} \tau_i(s) ds + C \sum_{j \notin S(x^*)} x_j$$

is a strict local Lyapunov function for any regular ESS x^* . (Hint: Establish this variant of Lemma 8.4.8: under the excess payoff dynamic generated by τ , there is a neighborhood O' of x^* such that

- (i) $\dot{x}'DF(x)\dot{x} \leq K T(x) (\mathbf{1}^0)'x$ and
- (ii) $(\mathbf{1}^0)' \dot{x} = -T(x) (\mathbf{1}^0)'x$,

for all $x \in O'$, where $T(x) = \sum_{i \in S} \tau_i(\hat{F}_i(x))$.)

8.5 Linearization of Imitative Dynamics

In this section and the next, we study the stability of rest points of evolutionary dynamics using linearization. This technique requires the dynamic in question to be smooth, at least near the rest point in question, and it can be inconclusive in borderline cases. But, more optimistically, it does not require the guesswork needed to find Lyapunov functions. Furthermore, instead of establishing just asymptotic stability, a rest point found stable via linearization (that is, one that is *linearly stable*) must attract solutions from all nearby initial conditions at an exponential rate (see Appendix 8.C). Linearization is also very useful for proving that a rest point is unstable, a fact we will avail ourselves of repeatedly when studying nonconvergence in Chapter 9. Finally, linearization techniques allow us to prove local stability results for imitative dynamics other than the replicator dynamic, for which no Lyapunov functions have been proposed.

In the Appendix, we explain the techniques from matrix analysis (Appendix 8.A), linear differential equation theory (Appendix 8.B), and linearization theory (Appendix 8.C) used in this chapter and the next. We assume in the remainder of this chapter and in the next chapter that payoffs are defined on the positive orthant (see Appendix 3.A.7), as doing so will allow us to avoid using affine calculus. Reviewing multivariate product and chain rules from Appendix 3.A.4 may be helpful for following the arguments to come.

We begin the analysis with some general background on linearization of evolutionary dynamics. Recall that a single population dynamic

$$(D) \quad \dot{x} = V(x)$$

describes the evolution of the population state through the simplex X . In evaluating the stability of the rest point x^* using linearization, we are relying on the fact that near x^* , the dynamic (D) can typically be well approximated by the linear dynamic

$$(L) \quad \dot{y} = DV(x^*)y.$$

Because we are only interested in how (D) behaves on the simplex, we only care about how (L) behaves on the tangent space TX . Indeed, it is only because (D) defines a dynamic on X that it makes sense to think of (L) as a dynamic on TX . At each state $x \in X$, $V(x) \in TX$ describes the current direction of motion through the simplex. It follows that the derivative $DV(x)$ must map any tangent vector z into TX , as one can verify by writing

$$V(x+z) = V(x) + DV(x)z + o(|z|)$$

and noting that $V(x)$ and $V(x+z)$ are both in TX . Thus, in (L), \dot{y} lies in TX whenever y lies in TX , implying that TX is invariant under (L).

Keeping this argument in mind is important when using linearization to study stability under the dynamic (D): rather than looking at all the eigenvalues of $DV(x^*)$, we should only consider those associated with the restricted linear map $DV(x^*) : TX \rightarrow TX$, which sends each tangent vector $z \in TX$ to a new tangent vector $DV(x^*)z \in TX$. The scalar $\lambda = a + ib$ is an eigenvalue of this restricted map if $DV(x^*)z = \lambda z$ for some vector z whose real and imaginary parts are both in TX . If all eigenvalues of this restricted map have negative real part, then the rest point x^* is linearly stable under (D) (cf Corollary 8.C.2).

Hines's Lemma, stated next and proved in Appendix 8.A.7, is often the key to making these determinations. In stating this result, we let $\mathbf{R}_0^n = \{z \in \mathbf{R}^n : \mathbf{1}'z = 0\}$ denote the tangent space of the simplex. In the single population case, TX and \mathbf{R}_0^n are the same, but it is useful to separate these two notations in multipopulation cases, where $TX = \prod_{p \in \mathcal{P}} \mathbf{R}_0^{n^p}$

Lemma 8.5.1. *Suppose that $Q \in \mathbf{R}^{n \times n}$ is symmetric, satisfies $Q\mathbf{1} = \mathbf{0}$, and is positive definite with respect to \mathbf{R}_0^n , and that $A \in \mathbf{R}^{n \times n}$ is negative definite with respect to \mathbf{R}_0^n . Then each eigenvalue of the linear map $QA : \mathbf{R}_0^n \rightarrow \mathbf{R}_0^n$ has negative real part.*

8.5.1 The Replicator Dynamic

In this section, we show that any regular Taylor ESS x^* is linearly stable under the replicator dynamic. To begin, we focus on the case in which x^* is interior.

Theorem 8.5.2. *Let $x^* \in \text{int}(X)$ be a regular Taylor ESS of F . Then x^* is linearly stable under the replicator dynamic.*

Proof. ($p = 1$) The single population replicator dynamic is given by

$$(R) \quad \dot{x}_i = V_i(x) = x_i \hat{F}_i(x).$$

To compute $DV(x)$, recall from equation (7.21) that the derivative of the excess payoff function $\hat{F}(x) = F(x) - \mathbf{1}\bar{F}(x)$ is given by

$$D\hat{F}(x) = DF(x) - \mathbf{1}(x'DF(x) + F(x)') = (I - \mathbf{1}x')DF(x) - \mathbf{1}F(x)'$$

Then applying the product rule for componentwise products (see Appendix 3.A.4), we find that

$$\begin{aligned} (8.28) \quad DV(x) &= D(\text{diag}(x)\hat{F}(x)) \\ &= \text{diag}(x)D\hat{F}(x) + \text{diag}(\hat{F}(x)) \\ &= \text{diag}(x)((I - \mathbf{1}x')DF(x) - \mathbf{1}F(x)') + \text{diag}(\hat{F}(x)) \\ &= Q(x)DF(x) - xF(x)' + \text{diag}(\hat{F}(x)), \end{aligned}$$

where we write $Q(x) = \text{diag}(x) - xx'$.

Since x^* is an interior Nash equilibrium, $F(x^*)$ is a constant vector, implying that $F(x^*)'\Phi = \mathbf{0}'$ and that $\hat{F}(x^*) = \mathbf{0}$. Thus, equation (8.28) becomes

$$(8.29) \quad DV(x^*)\Phi = Q(x^*)DF(x^*)\Phi.$$

Since the matrices $Q(x^*)$ and $DF(x^*)\Phi$ satisfy the conditions of Hines's Lemma, the eigenvalues of $DV(x^*)\Phi$ (and hence of $DV(x^*)$) corresponding to directions in \mathbf{R}_0^n have negative real part. This completes the proof of the theorem. ■

Exercise 8.5.3. Let x^* be an interior Nash equilibrium of F that satisfies $z'DF(x^*)z > 0$ for all nonzero $z \in TX$. Show that x^* is a source under the replicator dynamic: all relevant eigenvalues of $DV(x^*)$ have positive real part, implying that all solutions of the replicator dynamic that start near x^* are repelled. (Hint: See the discussion in Appendix 8.A.7.) Also, construct a game with an equilibrium that satisfies the conditions of this result.

Exercise 8.5.4. Show that if $x^* \in \text{int}(X)$ is a regular Taylor ESS, then x is linearly stable under the projection dynamic.

The next example highlights the fact that being a regular ESS is only a sufficient condition for an interior equilibrium to be locally stable under the replicator dynamic, not a necessary condition.

Example 8.5.5. Zeeman's game revisited. In Example 6.1.7, we introduced the single population game $F(x) = Ax$ generated by matching in the symmetric normal form game

$$A = \begin{pmatrix} 0 & 6 & -4 \\ -3 & 0 & 5 \\ -1 & 3 & 0 \end{pmatrix}.$$

This game admits Nash equilibria at states $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $(\frac{4}{5}, 0, \frac{1}{5})$ and e_1 ; the replicator dynamic has rest points at these states, as well as at the restricted equilibria $(0, \frac{5}{8}, \frac{3}{8})$, e_2 , and e_3 . Examining the phase diagram in Figure 8.5.1, we see that the behavior of the dynamic near the non-Nash rest points is consistent with Theorem 8.1.1.

Since F is not a stable game (why not?), Theorem 8.5.2 does not tell us whether x^* is stable. But we can check this directly: following the proof of Theorem 8.5.2, we compute

$$DV(x^*)\Phi = Q(x^*)DF(x^*)\Phi = Q(x^*)A\Phi = \frac{1}{9} \begin{pmatrix} 4 & 9 & -13 \\ -5 & -9 & 14 \\ 1 & 0 & -1 \end{pmatrix}.$$

In addition to the irrelevant eigenvalue of 0 corresponding to eigenvector $\mathbf{1}$, this matrix has pair of complex eigenvalues, $-\frac{1}{3} \pm i\frac{\sqrt{2}}{3}$, corresponding to eigenvectors $(-2 \pm i(3\sqrt{2}), 1 \mp i(3\sqrt{2}), 1)'$ whose real and complex parts lie in \mathbf{R}_0^n . Since the real parts of the relevant eigenvalues are both $-\frac{1}{3}$, the Nash equilibrium x^* is linearly stable under the replicator dynamic. §

We now establish the stability of all regular Taylor ESSs.

Theorem 8.5.6. *Let x^* be a regular Taylor ESS of F . Then x^* is linearly stable under the replicator dynamic.*

Proof. ($p = 1$) Suppose without loss of generality that the support of x^* is $\{1, \dots, n^*\}$, so that the number of unused strategies at x^* is $n^0 = n - n^*$. For any matrix $M \in \mathbf{R}^{n \times n}$, we let $M^{++} \in \mathbf{R}^{n^* \times n^*}$ denote the upper left $n^* \times n^*$ block of M , and we define the blocks $M^{+0} \in \mathbf{R}^{n^* \times n^0}$, $M^{0+} \in \mathbf{R}^{n^0 \times n^*}$, and $M^{00} \in \mathbf{R}^{n^0 \times n^0}$ similarly. Also, for each vector $v \in \mathbf{R}^n$, we let $v^+ \in \mathbf{R}^{n^*}$ and $v^0 \in \mathbf{R}^{n^0}$ denote the upper and lower "blocks" of v .

Recall our expression (8.28) for the derivative matrix of the replicator dynamic:

$$DV(x) = Q(x)DF(x) - xF(x)' + \text{diag}(\hat{F}(x)),$$

where $Q(x) = \text{diag}(x) - xx'$. Now observe that $x_j^* = 0$ for all $j > n^*$, that $\hat{F}_i(x^*) = 0$ for all $i \leq n^*$, and, since x^* is quasistrict, that $\hat{F}_j(x^*) < 0$ for all $j > n^*$ (see the proof of Lemma




Figure 8.5.1: The replicator dynamic in Zeeman's game.

5.5.4). Therefore, by writing $Q = Q(x^*)$, $D = DF(x^*)$, $\pi = F(x^*)$, and $\hat{\pi} = \hat{F}(x^*)$, we can express $DV(x^*)$ in the block diagonal form

$$(8.30) \quad DV(x^*) = \begin{pmatrix} Q^{++}D^{++} - (x^*)^+(\pi^+) & Q^{++}D^{+0} - (x^*)^+(\pi^0)' \\ \mathbf{0} & \text{diag}(\hat{\pi}^0) \end{pmatrix}.$$

To complete the proof of the theorem, we need to show that if $v + iw$ with $v, w \in \mathbf{R}_0^n$ is an eigenvector of $DV(x^*)$ with eigenvalue $a + ib$, then $a < 0$.

We split the analysis into two cases. Suppose first that $(v + iw)^0 = \mathbf{0}$ (i.e., that $v_j = w_j = 0$ whenever $j > n^*$). Then it is easy to see that $(v + iw)^+$ must be an eigenvector of $DV(x^*)^{++} = Q^{++}D^{++} - (x^*)^+(\pi^+)'$. Now because x^* is a Nash equilibrium with support $\{1, \dots, n^*\}$, π^+ is a constant vector, and since $v, w \in \mathbf{R}_0^n$ and $(v + iw)^0 = \mathbf{0}$, the components of $(v + iw)^+$ sum to zero. Together, these observations imply that $(x^*)^+(\pi^+)'(v + iw)^+ = \mathbf{0}$. Finally, $Q^{++} \in \mathbf{R}^{n^* \times n^*}$ and $D^{++} \in \mathbf{R}^{n^* \times n^*}$ satisfy the conditions of Hines's Lemma, the latter by requirement (8.22) for regular ESSs, and so this lemma enables us to conclude that $a < 0$.

Now suppose that $(v + iw)^0 \neq \mathbf{0}$, so that $v_j + iw_j \neq 0$ for some $j > n^*$. Then since the lower right block of $DV(x^*)$ is the diagonal matrix $\text{diag}(\hat{\pi}^0)$, the j th component of the eigenvector equation for $DV(x^*)$ is $\hat{\pi}_j(v_j + iw_j) = (a + ib)(v_j + iw_j)$, implying that $a = \hat{\pi}_j$ (and

also that $b = w_j = 0$). But as we noted above, the fact that x^* is a quasistrict equilibrium implies that $\hat{\pi}_j < 0$, and so that $a < 0$. This completes the proof of the theorem. ■

Exercise 8.5.7. Suppose that $x^* = e_i$ is a strict equilibrium of F . Show that for each $j \neq i$, the vector $e_j - e_i$ is an eigenvector of $DV(x^*)$ with eigenvalue $F_j(x^*) - F_i(x^*)$.

Exercise 8.5.8. Suppose that x^* is a quasistrict Nash equilibrium of F . We saw in Theorem 8.5.6 that for each unused strategy j , the excess payoff $\hat{F}_j(x^*)$ is an eigenvalue of $DV(x^*)$ corresponding to an eigenvector in TX . Assume that $\hat{F}_j(x^*)$ is not an eigenvalue of $DV(x^*)$ corresponding to an eigenvector in $TX \cap \mathbf{R}_{S(x^*)}^n$. Show that

$$\begin{pmatrix} \zeta + \frac{1}{n^*} \mathbf{1} \\ -\iota_j \end{pmatrix} \in TX$$

is an eigenvector of $DV(x^*)$ corresponding to eigenvalue $\hat{F}_j(x^*)$, where ι_j is the appropriate standard basis vector in \mathbf{R}^{n^0} , and where ζ is the unique vector in \mathbf{R}^{n^*} satisfying $\mathbf{1}'\zeta = 0$ and

$$(Q^{++}D^{++} - \hat{\pi}_j I)\zeta = \hat{\pi}_j(\frac{1}{n^*}\mathbf{1} - (x^*)^+) + Q^{++}(D^{+0}\iota_j - \frac{1}{n^*}D^{++}\mathbf{1}).$$

Why is there exactly one vector that satisfies these conditions? What goes wrong if the restriction on $\hat{F}_j(x^*)$ does not hold?

8.5.2 General Imitative Dynamics

Theorem 8.5.6 established the local stability of all regular Taylor ESSs under the replicator dynamic. Theorem 8.5.9 parlays the previous analysis into a local stability result for all imitative dynamics.

Theorem 8.5.9. *Assume that x^* is a hyperbolic rest point of both the replicator dynamic (R) and a given imitative dynamic (5.6). Then x^* is linearly stable under (R) if and only if it is linearly stable under (5.6). Thus, if x^* is a regular Taylor ESS that satisfies the hyperbolicity assumptions, it is linearly stable under (5.6).*

While the proof of this result requires a somewhat involved argument, the key step in the analysis is presented in Lemma 8.5.10, which shows that at interior equilibria, the linearization of any monotone percentage growth rate function G is a multiple of the linearization of the payoff function F .

Proof of Theorem 8.5.9. ($p = 1$) We only consider the case in which x^* is interior; for boundary cases, see Exercise 8.5.12.

Recall from Observation 5.4.16 that any imitative dynamic (5.6) has monotone percentage growth rates: we can express the dynamic as

$$(8.31) \quad \dot{x}_i = x_i G_i(x), \text{ where}$$

$$(8.32) \quad G_i(x) \geq G_j(x) \text{ if and only if } F_i(x) \geq F_j(x).$$

Lemma 8.5.10 shows that property (8.32) imposes a remarkable amount of structure on the derivative matrix of the percentage growth rate function G at the equilibrium x^* .

Lemma 8.5.10. *Let x^* be an interior Nash equilibrium, and suppose that $\Phi DF(x^*)$ and $\Phi DG(x^*)$ define invertible maps from TX to itself. Then $\Phi DG(x^*)\Phi = c \Phi DF(x^*)\Phi$ for some $c > 0$.*

Proof. Since x^* is a Nash equilibrium, and hence a rest point of (8.31), we have that $\Phi F(x^*) = \Phi G(x^*) = \mathbf{0}$. It follows that

$$(8.33) \quad \Phi F(x^* + \varepsilon z) = \varepsilon \Phi DF(x^*)z + o(\varepsilon) \quad \text{and} \quad \Phi G(x^* + \varepsilon z) = \varepsilon \Phi DG(x^*)z + o(\varepsilon).$$

for all $z \in TX$. Since we can rewrite condition (8.32) as

$$(e_i - e_j)'G(x) \geq 0 \text{ if and only if } (e_i - e_j)'F(x) \geq 0,$$

and since $e_i - e_j \in TX$, equation (8.33) implies that for all $i, j \in S$ and $z \in TX$,

$$(8.34) \quad (e_i - e_j)' \Phi DG(x^*)z \geq 0 \text{ if and only if } (e_i - e_j)' \Phi DF(x^*)z \geq 0.$$

(This observation is trivial when $z = \mathbf{0}$, and when $z \neq \mathbf{0}$ it follows from the fact that the linear terms dominate in (8.33) when ε is small.) By Proposition 3.B.6, condition (8.34) is equivalent to the requirement that for all $i, j \in S$, there is a $c_{ij} > 0$ such that

$$(8.35) \quad (e_i - e_j)' \Phi DG(x^*)\Phi = c_{ij}(e_i - e_j)' \Phi DF(x^*)\Phi.$$

Now write $g_{ij} = (e_i - e_j)' \Phi DG(x^*)\Phi$ and $f_{ij} = (e_i - e_j)' \Phi DF(x^*)\Phi$. Since by assumption $\Phi DF(x^*)\Phi$ is an invertible map from TX to itself, so is its transpose (see Exercise 8.5.11 below). Therefore, when i, j , and k are distinct, the unique decomposition of f_{ik} as a linear combination of f_{ij} and f_{jk} is as $f_{ij} + f_{jk}$. But equation (8.35) reveals that

$$c_{ij}f_{ij} + c_{jk}f_{jk} = g_{ij} + g_{jk} = g_{ik} = c_{ik}f_{ik},$$

and so $c_{ij} = c_{jk} = c_{ik}$. This and the fact that $c_{ij} = c_{ji}$ imply that c_{ij} is independent of i and j . So, since vectors of the form $e_i - e_j$ span TX , we conclude from equation (8.35) that

$\Phi DG(x^*)\Phi = c \Phi DF(x^*)\Phi$, where c is the common value of the constants c_{ij} . This completes the proof of the lemma. ■

We proceed with the proof of Theorem 8.5.9. Let $V(x) = \text{diag}(x)\hat{F}(x)$ and $W(x) = \text{diag}(x)G(x)$ denote the replicator dynamic (R) and the dynamic (8.31), respectively. Since $W(x) \in TX$, we have that $\mathbf{1}'W(x) = x'G(x) = 0$, and hence that $\hat{G}(x) \equiv G(x) - \mathbf{1}x'G(x) = G(x)$. Thus (8.31) can be rewritten as $W(x) = \text{diag}(x)\hat{G}(x)$.

Now, repeating calculation (8.28) reveals that

$$DW(x) = Q(x)DG(x) - x G(x)' + \text{diag}(\hat{G}(x)).$$

Since x^* is an interior rest point of W , $G(x^*)$ is a constant vector, and so

$$DW(x^*)\Phi = Q(x^*)DG(x^*)\Phi = Q(x^*)\Phi DG(x^*)\Phi,$$

where the second equality follows from the fact that $Q(x^*)\mathbf{1} = \mathbf{0}$. Similar reasoning for the replicator dynamic V shows that

$$DV(x^*)\Phi = Q(x^*)\Phi DF(x^*)\Phi$$

Lemma 8.5.10 tells us that $\Phi DG(x^*)\Phi = c\Phi DF(x^*)\Phi$ for some $c > 0$. We therefore conclude from the previous two equations that if x^* is a hyperbolic rest point under V and W , its stability properties under the two dynamics are the same. ■

Exercise 8.5.11. Suppose that $A \in \mathbf{R}^{n \times n}$ defines an invertible map from \mathbf{R}_0^n to itself and maps the vector $\mathbf{1}$ to the origin. Show that A' must also have these properties. (Hint: Use the Fundamental Theorem of Linear Algebra (8.42).)

Exercise 8.5.12. Extend the proof of Theorem 8.5.9 above to the case of boundary equilibria. (Hint: Combine Lemma 8.5.10 with the proof of Theorem 8.5.6.)

8.6 Linearization of Perturbed Best Response Dynamics

Linearization is also a useful tool for studying perturbed best response dynamics, our other main class of differentiable evolutionary dynamics.

8.6.1 Deterministically Perturbed Best Response Dynamics

In Chapter 6, we saw that perturbed best response dynamics can be defined in terms of either stochastic or deterministic payoff perturbations. But Theorem 6.2.2 showed that there is no loss of generality in focusing on the latter case, and so we will do so here.

Our first result shows that a negative definiteness condition on the payoff derivative is a sufficient condition for stability. The conclusion here is similar to that from Exercise 8.4.5, but the analysis is much simpler, and establishes not only asymptotic stability, but also linear stability.

Theorem 8.6.1. *Consider the perturbed best response dynamic for the pair (F, v) , and let \tilde{x} be a perturbed equilibrium of this pair. If $DF(\tilde{x})$ is negative definite with respect to TX , then \tilde{x} is linearly stable.*

Proof. ($p = 1$) In the single population case, the stochastically perturbed best response dynamic takes the form

$$(8.36) \quad \dot{x} = \tilde{M}(F(x)) - x,$$

where the perturbed maximizer function \tilde{M} is defined in equation (6.12). By the chain rule, the derivative of law of motion (8.36) is

$$(8.37) \quad DV(x) = D\tilde{M}(F(x))DF(x) - I.$$

To determine the eigenvalues of the product $D\tilde{M}(F(x))DF(x)$, let us recall the properties of the derivative matrix $D\tilde{M}(\pi)$ from Corollary 6.C.5: it is symmetric, positive definite on \mathbf{R}_0^n , and satisfies $D\tilde{M}(\pi)\mathbf{1} = \mathbf{0}$. Since we have assumed that $DF(\tilde{x})$ is negative definite with respect to \mathbf{R}_0^n , Hines's Lemma implies that the eigenvalues of $D\tilde{M}(F(\tilde{x}))DF(\tilde{x})$ (as a map from \mathbf{R}_0^n to itself) have negative real part. Subtracting the identity matrix I from the matrix product reduces each of these eigenvalues by 1, so the theorem is proved. ■

Exercise 8.6.2. Show that the conclusion of the theorem continues to hold if $DF(x)$ is only negative semidefinite with respect to TX . (Hint: See the discussion in Appendix 8.A.7.)

Exercise 8.6.3. Let \tilde{x} be a perturbed equilibrium for (F, v) . Let $\bar{\lambda}$ be the largest eigenvalue of $D\tilde{M}(F(\tilde{x}))$, and let \bar{s} be the largest singular value of $\Phi DF(\tilde{x})\Phi$ (see Section 8.A.6). Show that if $\bar{\lambda}\bar{s} < 1$, then \tilde{x} is linearly stable: that is, \tilde{x} is stable whenever choice probabilities are not too sensitive to changes in payoffs, or payoffs are not too sensitive to changes in the state.

8.6.2 The Logit Dynamic

Imposing the additional structure provided by logit choice allows us to carry our local stability analysis further. First, building on Theorem 8.4.6, we argue that any regular interior ESS must have a linearly stable $\text{logit}(\eta)$ equilibrium nearby whenever the noise level η is sufficiently small.

Corollary 8.6.4. *Let $x^* \in \text{int}(X)$ be a regular Taylor ESS of F . Then for some neighborhood O of x^* and all $\eta > 0$ less than some $\hat{\eta} > 0$, there is a unique and linearly stable $\text{logit}(\eta)$ equilibrium \tilde{x}^η in O .*

Proof. ($p = 1$) Theorem 8.4.6 tells us that for η small enough, the equilibrium \tilde{x}^η exists and is unique, and that $\lim_{\eta \rightarrow 0} \tilde{x}^\eta = x^*$. Since x^* is a regular interior ESS, $DF(x^*)$ is negative definite with respect to TX , so by continuity, $DF(\tilde{x}^\eta)$ is negative definite with respect to TX for all η close enough to 0. The result therefore follows from Theorem 8.6.1. ■

The derivative matrix for the logit dynamic takes an especially appealing form. Recall from Exercise 6.2.7 that the derivative matrix of the $\text{logit}(\eta)$ choice function is

$$(8.38) \quad D\tilde{M}^\eta(\pi) = \eta^{-1} \left(\text{diag}(\tilde{M}^\eta(\pi)) - \tilde{M}^\eta(\pi)\tilde{M}^\eta(\pi)' \right) = \eta^{-1}Q(\tilde{M}^\eta(\pi)).$$

Now by definition, the logit equilibrium \tilde{x}^η satisfies $\tilde{M}^\eta(F(\tilde{x}^\eta)) = \tilde{x}^\eta$. Substituting this fact into equations (8.37) and (8.38) yields

$$(8.39) \quad DV^\eta(\tilde{x}^\eta) = \eta^{-1}Q(\tilde{x}^\eta)DF(\tilde{x}^\eta) - I.$$

To see the importance of this equation, recall from equation (8.29) that at interior rest points, the derivative matrix for the replicator dynamic satisfies

$$(8.40) \quad DV(x^*)\Phi = Q(x^*)DF(x^*)\Phi.$$

Together, equations (8.39) and (8.40) show that when evaluated at their respective rest points and in the relevant tangent directions, the linearizations of the replicator and logit dynamics at their interior rest points differ only by a positive affine transformation!

Example 8.6.5. To obtain the cleanest connections between the two dynamics, consider a game that admits a Nash equilibrium $x^* = \frac{1}{n}\mathbf{1}$ at the barycenter of the simplex. Then by symmetry, $\tilde{x}^\eta = x^*$ is also a $\text{logit}(\eta)$ equilibrium for every $\eta > 0$. By the logic above, λ is a relevant eigenvalue of (8.40) if and only if $\eta^{-1}\lambda - 1$ is a relevant eigenvalue of (8.39). It follows that if x^* is linearly stable under the replicator dynamic, then it is also linearly stable under the $\text{logit}(\eta)$ dynamic for any $\eta > 0$. §

The foregoing discussion shows how analyses of local stability under the replicator and logit dynamics can be closely linked. Pushing these arguments further, one can use equations (8.39) and (8.40) to connect the long run behaviors of the replicator and best response dynamics starting from arbitrary initial conditions—see the Notes for further discussion.

Appendix

8.A Matrix Analysis

In this section we review some basic ideas from matrix analysis. In doing so, we lay the groundwork for our introduction to linear differential equations in Appendix 8.B; this in turn underlies our introduction to local linearization of nonlinear differential equations in Appendix 8.C. The techniques presented here are also used to perform the explicit calculations that arise when using linearization to analyze evolutionary dynamics.

8.A.1 Rank and Invertibility

While in most of this section we focus on square matrices, we start by considering matrices $A \in \mathbf{R}^{m \times n}$ of arbitrary dimensions. The *rank* of A is the number of linearly independent columns of A , or, equivalently, the dimension of its *range*, $\text{range}(A) = \{y \in \mathbf{R}^m : Ax = y \text{ for some } x \in \mathbf{R}^n\}$. The *nullspace* (or *kernel*) of A , $\text{nullspace}(A) = \{x \in \mathbf{R}^n : Ax = \mathbf{0}\}$ is the set of vectors in that the matrix maps to the origin, and the dimension of this set is called the *nullity* of A . The rank and nullity of a matrix must sum to its number of columns:

$$\begin{aligned} \dim(\text{nullspace}(A)) + \dim(\text{range}(A)) &= n; \\ (8.41) \quad \dim(\text{nullspace}(A')) + \dim(\text{range}(A')) &= m. \end{aligned}$$

The foregoing definitions identify the matrices A and A' with the linear maps $x \in \mathbf{R}^n \mapsto Ax \in \mathbf{R}^m$ and $y \in \mathbf{R}^m \mapsto A'y \in \mathbf{R}^n$. This identification is used implicitly in many of the discussions to follow.

In Appendix 3.B.2, we introduced the *Fundamental Theorem of Linear Algebra*:

$$(8.42) \quad \text{range}(A) = (\text{nullspace}(A'))^\perp.$$

To derive a key implication of (8.42) for the ranks of matrices, first recall that any subspace

$V \subseteq \mathbf{R}^m$ satisfies $\dim(V) + \dim(V^\perp) = m$. Letting $V = \text{nullspace}(A')$ and then combining the result with equation (8.41), we obtain

$$\dim(\text{range}(A')) = \dim((\text{nullspace}(A'))^\perp).$$

Therefore, (8.42) yields

$$\dim(\text{range}(A')) = \dim(\text{range}(A)).$$

In words: every matrix has the same rank as its transpose.

From this point forward, we suppose that $A \in \mathbf{R}^{n \times n}$ is a square matrix. We say that A is *invertible* if it admits an *inverse matrix* A^{-1} : that is, a matrix satisfying $A^{-1}A = I$. Such a matrix also satisfies $AA^{-1} = I$, and when an inverse matrix exists, it is unique. Invertible matrices can be characterized in a variety of ways: for instance, a matrix is invertible if and only if it has *full rank* (i.e., if $A \in \mathbf{R}^{n \times n}$ has rank n); alternatively, a matrix is invertible if and only if its determinant is nonzero.

8.A.2 Eigenvectors and Eigenvalues

Let $A \in \mathbf{R}^{n \times n}$, and suppose that

$$(8.43) \quad Ax = \lambda x$$

for some complex scalar $\lambda \in \mathbf{C}$ and some nonzero complex vector $x \in \mathbf{C}^n$. Then we call λ an *eigenvalue* of A , and x an *eigenvector* of A associated with λ ; sometimes, the pair (λ, x) is referred to as an *eigenpair*.

The eigenvector equation (8.43) can be rewritten as $(\lambda I - A)x = \mathbf{0}$. This equation can only be satisfied by a nonzero vector if $(\lambda I - A)$ is not invertible, or, equivalently, if $\det(\lambda I - A) = 0$. It follows that λ is an eigenvalue of A if and only if λ is a root of the *characteristic polynomial* $\det(tI - A)$.

Since $\det(tI - A)$ is a polynomial of degree n in t , the *Fundamental Theorem of Algebra* ensures that it has n complex roots:

$$(8.44) \quad \det(tI - A) = (t - \lambda_1)(t - \lambda_2) \dots (t - \lambda_n).$$

To be sure to obtain n roots, we must “count multiplicities”: if the values of λ_i in the above expression are not all distinct, the repeated values must be tallied each time they appear. Evidently, each λ_i in (8.44) is an eigenvalue of A ; if the value λ is repeated k times in (8.44),

we say that λ is an eigenvalue of A of (algebraic) multiplicity k .

We note in passing that the sum and the product of the eigenvalues of A can be described very simply:

$$\sum_{i=1}^n \lambda_i = \operatorname{tr}(A); \quad \prod_{i=1}^n \lambda_i = \det(A).$$

(Here, the *trace* $\operatorname{tr}(A)$ of the matrix A is the sum of its diagonal elements.) To remember these formulas, notice that they are trivially true if A is a diagonal matrix, since in this case the eigenvalues of A are its diagonal entries.

Each eigenvalue of A corresponds to at least one eigenvector of A , and if an eigenvalue λ is of algebraic multiplicity k , then there can be as many as k linearly independent eigenvectors of A corresponding to this eigenvalue. This number of linearly independent eigenvectors is called the *geometric multiplicity* of λ . The collection of all eigenvectors corresponding to λ , the *eigenspace* of λ , is a subspace of \mathbf{C}^n of dimension equal to the geometric multiplicity of λ .

Example 8.A.1. Let $a, b \in \mathbf{R}$ be nonzero, and consider these three 2×2 matrices:

$$A = \begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}; \quad B = \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}; \quad C = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}.$$

The matrix A has just one eigenvalue, a , which therefore has algebraic multiplicity 2. It also has geometric multiplicity 2, as its eigenspace is all of $\mathbf{C}^2 = \operatorname{span}(\{e_1, e_2\})$. (This description of \mathbf{C}^2 relies on our allowing complex scalars when taking linear combinations of e_1 and e_2 .)

The matrix B also has a lone eigenvalue of a of algebraic multiplicity 2. But here the geometric multiplicity of a is just 1, since its eigenspace is $\operatorname{span}(\{e_1\})$.

The matrix C has no real eigenvalues or eigenvectors; however, it has complex eigenvalues $a \pm ib \in \mathbf{C}$ corresponding to the complex eigenvectors $e_1 \pm ie_2 \in \mathbf{C}^2$.

Let us explain for future reference the geometry of the linear map $x \mapsto Cx$. By writing $r = \sqrt{a^2 + b^2}$ and $\theta = \cos^{-1}(\frac{a}{r})$, we can express the matrix C as

$$C = r \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}.$$

Computing Cx for various values of x (try $x = e_1$ and $x = e_2$), reveals that the map $x \mapsto Cx$ first rotates the vector x around the origin clockwise by an angle of θ , and then rescales the result by a factor of r . §

8.A.3 Similarity, (Block) Diagonalization, and the Spectral Theorem

The matrix $A \in \mathbf{R}^{n \times n}$ is *similar* to matrix $B \in \mathbf{R}^{n \times n}$ if there exists an invertible matrix $S \in \mathbf{C}^{n \times n}$, called a *similarity matrix*, such that

$$B = S^{-1}AS.$$

When A is similar to B , the linear transformations $x \mapsto Ax$ and $y \mapsto By$ are equivalent up to a linear change of variable. Similarity defines an equivalence relation on the set of $n \times n$ matrices, and matrices that are similar have the same characteristic polynomial and the same eigenvalues, counting either algebraic or geometric multiplicities.

If A is similar to a diagonal matrix D —that is, if A is *diagonalizable*—then the eigenvalues of A are simply the diagonal elements of D . In this definition the similarity matrix is allowed to be complex; if the similarity can be achieved via a real similarity matrix $S \in \mathbf{R}^{n \times n}$, then the diagonal matrix D is also real, and we call A *real diagonalizable*.

It follows easily from our definitions that a matrix A is diagonalizable if and only if the sum of the geometric multiplicities of the eigenvalues of A is n . Equivalently, A is diagonalizable if and only if each of its eigenvalues has equal algebraic and geometric multiplicities. It is simple to verify that in this case, a similarity matrix S can be constructed by choosing n linearly independent eigenvectors of A to be its columns.

It is especially convenient when similarity can be achieved using similarity matrix that is itself of a simple form. The most important instance occurs when this matrix is an *orthogonal matrix*, meaning that its columns form an *orthonormal basis* for \mathbf{R}^n : each column is of length 1, and distinct columns are orthogonal. (It would make more sense to call such a matrix an “orthonormal matrix”, but the term “orthogonal matrix” is traditional.)

Orthogonal matrices can be characterized in a variety of ways:

Theorem 8.A.2. *The following are equivalent:*

- (i) R is an orthogonal matrix.
- (ii) $RR' = I$.
- (iii) $R' = R^{-1}$.
- (iv) The map $x \mapsto Rx$ preserves lengths: $|Rx| = |x|$ for all $x \in \mathbf{R}^n$.
- (v) The map $x \mapsto Rx$ preserves inner products: $(Rx)'(Ry) = x'y$ for all $x, y \in \mathbf{R}^n$.
- (vi) The map $x \mapsto Rx$ is a composition of rotations and reflections.

The last three items are summarized by saying that the linear transformation $x \mapsto Rx$ defined by an orthogonal matrix R is a *Euclidean isometry*.

While showing that a matrix is similar to a diagonal matrix is quite useful, showing similarity to a block diagonal matrix often serves just as well. We focus on block diagonal

matrices with diagonal blocks of these two types:

$$J_1 = (\lambda); \quad J_2 = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}.$$

For reasons that will become clear in Section 8.A.5, we call block diagonal matrices of this form *simple Jordan matrices*. Calculations with simple Jordan matrices are often little more difficult than those with diagonal matrices: for instance, multiplying such a matrix by itself retains its block diagonal structure.

To muster these ideas, let us call the matrix $A \in \mathbf{R}^{n \times n}$ *normal* if it commutes with itself: that is, if $A'A = AA'$.

Theorem 8.A.3 (The Spectral Theorem for Real Normal Matrices). *The matrix $A \in \mathbf{R}^{n \times n}$ is normal if and only if it is similar via an orthogonal matrix R to a simple Jordan matrix $B = R^{-1}AR$. The matrix B is unique up to the ordering of the diagonal blocks.*

The spectral decomposition of A provides a full account of the eigenvalues and eigenvectors of A . Each J_1 block (λ) contains a real eigenvalue of A , and the pair of complex numbers $a \pm ib$ derived from each J_2 block are complex eigenvalues of A . Moreover, columns of the orthogonal similarity matrix R either are real eigenvectors of A , or are real and imaginary parts of complex eigenvectors of A .

The spectral theorem tells us that if A is normal, the behavior of the linear map $x \mapsto Ax = RBR^{-1}x$ can be decomposed into three simple steps. First, one applies the orthogonal transformation $R^{-1} = R'$ to x , obtaining $y = R'x$. Second, one applies the block diagonal matrix B to y : each J_1 block rescales a component of y , while each J_2 block rotates and rescales a pair of components of y (cf Example 8.A.1). Third, one applies R to $BR'x$ to undo the initial orthogonal transformation.

Additional restrictions on the J_1 and J_2 blocks yield characterizations of important subclasses of the normal matrices.

Corollary 8.A.4. (i) *The matrix $A \in \mathbf{R}^{n \times n}$ is symmetric ($A' = A$) if and only if it is similar via an orthogonal matrix R to a simple Jordan matrix containing only J_1 blocks. Thus, the symmetric matrices are the normal matrices with real eigenvalues.*

(ii) *The matrix $A \in \mathbf{R}^{n \times n}$ is skew-symmetric ($A' = -A$) if and only if it is similar via an orthogonal matrix R to a simple Jordan matrix whose J_1 blocks all have $\lambda = 0$ and whose J_2 blocks all have $a = 0$. Thus, the skew-symmetric matrices are the normal matrices with purely imaginary eigenvalues.*

(iii) *The matrix $A \in \mathbf{R}^{n \times n}$ is orthogonal ($A' = A^{-1}$) if and only if it is similar via an orthogonal matrix R to a simple Jordan matrix whose J_1 blocks all have $\lambda^2 = 1$ and whose J_2 blocks all*

have $a^2 + b^2 = 1$. Thus, the orthogonal matrices are the normal matrices whose eigenvalues have modulus 1.

8.A.4 Symmetric Matrices

Which matrices are real diagonalizable by an orthogonal matrix? The *spectral theorem for symmetric matrices* tells us that A is real diagonalizable by an orthogonal matrix if and only if it is symmetric. (This is just a restatement of Corollary 8.A.4.) Among other things, the spectral theorem implies that the eigenvalues of a symmetric matrix are real.

While we often associate a matrix A with the linear transformation $x \mapsto Ax$, a symmetric matrix is naturally associated with a quadratic form, $x \mapsto x'Ax$. In fact, the eigenvalues of a symmetric matrix can be characterized in terms of its quadratic form. The *Rayleigh-Ritz Theorem* provides simple descriptions of the $\bar{\lambda}$ and $\underline{\lambda}$, the maximal and minimal eigenvalues of A :

$$\bar{\lambda} = \max_{x \in \mathbf{R}^n: |x|=1} x'Ax; \quad \underline{\lambda} = \min_{x \in \mathbf{R}^n: |x|=1} x'Ax.$$

The *Courant-Fischer Theorem* shows how the remaining eigenvalues of A can be expressed in terms of a related sequence of minmax problems.

We say that the matrices $A, B \in \mathbf{R}^{n \times n}$ are *congruent* if there is an invertible matrix $Q \in \mathbf{R}^{n \times n}$ such that

$$B = QAQ'.$$

Congruence plays the same role for quadratic forms as similarity does for linear transformations: if two symmetric matrices are congruent, they define the same quadratic form up to a linear change of variable. Like similarity, congruence defines an equivalence relation on the set of $n \times n$ matrices. Lastly, note that two symmetric matrices that are similar by an orthogonal matrix Q are also congruent, since in this case $Q' = Q^{-1}$.

The eigenvalues of congruent symmetric matrices are closely linked. Define the *inertia* of a symmetric matrix to be the ordered triple consisting of the numbers of positive, negative, and zero eigenvalues of the matrix. *Sylvester's Law of Inertia* tells us that congruent symmetric matrices have the same inertia. *Ostrowski's Theorem* provides a quantitative extension of this result: if we list the eigenvalues of A and the eigenvalues of B in increasing order, then the ratios between pairs of corresponding eigenvalues are bounded by the minimal and maximal eigenvalues of $Q'Q$.

8.A.5 The Real Jordan Canonical Form

How can we tell if two matrices are similar? If the matrices are diagonalizable, then one can check for similarity by diagonalizing the two matrices and seeing whether the same diagonal matrix is obtained in each case. To apply this logic beyond the diagonalizable case, we would need to find a simple class of matrices with the property that every matrix is similar to a unique representative from this class. Such a class of matrices would also provide a powerful computational aid, since calculations involving arbitrary matrices could be reduced by similarity to calculations with these simple matrices.

With this motivation, we define a *real Jordan matrix* to be a block diagonal matrix whose diagonal blocks, known as *Jordan blocks*, are of these four types:

$$J_1 = (\lambda); \quad J_2 = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}; \quad J_3 = \begin{pmatrix} \lambda & 1 & 0 & 0 & 0 \\ 0 & \lambda & 1 & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & 0 & \lambda \end{pmatrix}; \quad J_4 = \begin{pmatrix} J_2 & I & 0 & 0 & 0 \\ 0 & J_2 & I & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & J_2 & I \\ 0 & 0 & 0 & 0 & J_2 \end{pmatrix}.$$

Theorem 8.A.5. *Every matrix $A \in \mathbf{R}^{n \times n}$ is similar via a real similarity matrix S to a real Jordan matrix $J = S^{-1}AS$. The latter matrix is unique up to the ordering of the Jordan blocks.*

The real Jordan matrix in the statement of the theorem is called the *real Jordan canonical form* of A .

The blocks in the real Jordan form of A provide detailed information about the eigenvalues of A : each J_1 block corresponds to a real eigenvalue λ ; each J_2 block corresponds to a pair of complex eigenvalues $a \pm ib$; each J_3 block corresponds to a real eigenvalue with less than full geometric multiplicity; and each J_4 block corresponds to a pair of complex eigenvalues with less than full geometric multiplicities. (We can say more if each Jordan block represents a distinct eigenvalue: then each eigenvalue has geometric multiplicity 1; the J_1 and J_2 blocks correspond to eigenvalues whose algebraic multiplicities are also 1; and the J_3 and J_4 blocks correspond to eigenvalues with higher algebraic multiplicities, with these multiplicities being given by the number of appearances of λ (in a J_3 block) or of J_2 blocks (in a J_4 block).)

Example 8.A.6. Suppose that $A \in \mathbf{R}^{2 \times 2}$ has complex eigenvalues $a \pm ib$ with complex eigenvectors $v \pm iw$. Then $A(v + iw) = (a + ib)(v + iw)$. Equating the real and imaginary

parts of this equation yields

$$A \begin{pmatrix} v & w \end{pmatrix} = \begin{pmatrix} v & w \end{pmatrix} \begin{pmatrix} a & b \\ -b & a \end{pmatrix}.$$

Premultiplying by $(v \ w)^{-1}$ reveals that the real Jordan form of A is a single J_2 block. §

Example 8.A.7. Suppose that $A \in \mathbf{R}^{2 \times 2}$ has a lone eigenvalue, $\lambda \in \mathbf{R}$, which is of algebraic multiplicity 2 but geometric multiplicity 1. Let $x \in \mathbf{R}^2$ be an eigenvector of A , so that $(A - \lambda I)x = \mathbf{0}$. It can be shown that there exists a vector y that is linearly independent of x and that satisfies $(A - \lambda I)y = x$. (Such a vector (and, more generally, vectors that satisfy higher iterates of this equation) is called a *generalized eigenvector* of A .) Rewriting the two equations above, we obtain

$$A \begin{pmatrix} x & y \end{pmatrix} = \begin{pmatrix} \lambda x & x + \lambda y \end{pmatrix} = \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}.$$

Premultiplying the first and last expressions by $(x \ y)^{-1}$ shows that A has a real Jordan form consisting of a single J_3 block. §

8.A.6 The Spectral Norm and Singular Values

It is often useful to be able to place bounds on the amount of “expansion” generated by a linear map $x \mapsto Ax$, or by a composite linear map $x \mapsto Bx \mapsto ABx$. One can obtain such bounds by introducing the *spectral norm* of a matrix $A \in \mathbf{R}^{n \times n}$, defined by

$$\|A\| = \max_{x: |x|=1} |Ax|.$$

(As always in this book, $|x|$ denotes the Euclidean norm of the vector x .) It is not difficult to check that the spectral norm is submultiplicative, in the following two senses:

$$|Ax| \leq \|A\| |x|; \text{ and}$$

$$\|AB\| \leq \|A\| \|B\|.$$

These inequalities often work hand in hand with the *Cauchy-Schwarz inequality*, which expresses the submultiplicativity of inner products of vectors:

$$|x'y| \leq |x| |y|.$$

To compute the spectral norm of a matrix, it is best to describe it in a different way. We start with the observation that the product $A'A$ generated by any matrix A is symmetric. This product therefore has n real eigenvalues (see Section 8.A.4), and it can be shown that these eigenvalues are nonnegative. The square roots of the eigenvalues of $A'A$ are called the *singular values* of A .

One can show that the spectral norm of A equals the largest singular value of A :

$$\|A\| = \max \{ \sqrt{\lambda} : \lambda \text{ is an eigenvalue of } A'A \}.$$

It makes no difference here if we replace $A'A$ with AA' , since for any $A, B \in \mathbf{R}^{n \times n}$, AB and BA have the same eigenvalues.

The notion of a singular value also underpins the *singular value decomposition*

Theorem 8.A.8. *Every matrix $A \in \mathbf{R}^{n \times n}$ can be expressed as $A = V\Sigma W'$, where V and W are orthogonal matrices, and where Σ is a diagonal matrix whose diagonal entries are the singular values of A .*

In this decomposition, the columns of V are eigenvectors of AA' , and the columns of W are eigenvectors of $A'A$.

8.A.7 Hines's Lemma

In Section 8.5, we introduced *Hines's Lemma*:

Lemma 8.5.1. *Suppose that $Q \in \mathbf{R}^{n \times n}$ is symmetric, satisfies $Q\mathbf{1} = \mathbf{0}$, and is positive definite with respect to \mathbf{R}_0^n , and that $A \in \mathbf{R}^{n \times n}$ is negative definite with respect to \mathbf{R}_0^n . Then each eigenvalue of the linear map $QA : \mathbf{R}_0^n \rightarrow \mathbf{R}_0^n$ has negative real part.*

If we ignored the complications caused by the fact that our dynamics are restricted to the simplex, Lemma 8.5.1 would reduce to

Lemma 8.A.9. *If Q is symmetric positive definite and A is negative definite, then the eigenvalues of QA have negative real parts.*

The proof of Lemma 8.A.9 is a simpler version of the proof below.

The argument below can also be used when other definiteness conditions are imposed on A . In particular, if A is only negative semidefinite with respect to \mathbf{R}_0^n , then the relevant eigenvalues of QA have nonpositive real parts, and if A is positive definite with respect to \mathbf{R}_0^n , the relevant eigenvalues of QA have positive real parts.

Proof of Lemma 8.5.1. Since Q is positive definite with respect to \mathbf{R}_0^n , since $Q\mathbf{1} = \mathbf{0}$, and since $\mathbf{R}^n = \mathbf{R}_0^n \oplus \text{span}(\{\mathbf{1}\})$, we have that $\text{nullspace}(Q) = \text{span}(\{\mathbf{1}\})$. Thus, because Q is symmetric, the Fundamental Theorem of Linear Algebra (8.42) tells us that

$$\text{range}(Q) = (\text{nullspace}(Q'))^\perp = (\text{nullspace}(Q))^\perp = (\text{span}(\{\mathbf{1}\}))^\perp = \mathbf{R}_0^n.$$

In other words, Q maps \mathbf{R}_0^n onto itself, and so is invertible on this space.

Now suppose that

$$(8.45) \quad QA(v + iw) = (a + ib)(v + iw)$$

for some $v, w \in \mathbf{R}_0^n$ with $v + iw \neq \mathbf{0}$ and some a, b in \mathbf{R} . Since Q is invertible on \mathbf{R}_0^n , there exist $y, z \in \mathbf{R}_0^n$, at least one of which is not $\mathbf{0}$, such that $Qy = v$ and $Qz = w$. We can thus rewrite equation (8.45) as

$$QA(v + iw) = (a + ib)Q(y + iz).$$

Since Q is invertible on \mathbf{R}_0^n , this implies that

$$A(v + iw) = (a + ib)(y + iz).$$

Premultiplying by $(v - iw)' = (Q(y - iz))'$ yields

$$(v - iw)'A(v + iw) = (a + ib)(y - iz)'Q(y + iz).$$

Equating the real parts of each side yields

$$v'Av + w'Aw = a(y'Qy + z'Qz).$$

Since Q is positive definite with respect to \mathbf{R}_0^n and A is negative definite with respect to \mathbf{R}_0^n , we conclude that $a < 0$. ■

8.B Linear Differential Equations

The simplest ordinary differential equations on \mathbf{R}^n are *linear differential equations*:

$$(L) \quad \dot{x} = Ax,$$

where $A \in \mathbf{R}^{n \times n}$. Although our main interest in this book is in nonlinear differential equations, linear differential equations are still very important to us: as we explain in Section 8.C, the behavior of a nonlinear equation in the neighborhood of a rest point is often well approximated by the behavior of linear equation in a neighborhood of the origin.

8.B.1 Examples

Example 8.B.1. Linear dynamics on the line. In the one-dimensional case, equation (L) becomes $\dot{x} = ax$. We described the solution to this equation from initial condition $x_0 = \xi$ in Example 4.A.1: they are of the form $x_t = \xi \exp(at)$. Thus, if $a \neq 0$, the equation has its unique rest point at the origin. If $a > 0$, all solutions other than the stationary one move away from the origin, while if $a < 0$, all solutions converge to the origin. §

One can always apply a linear change of variable to (L) to reduce it to a simpler form. In particular, if $B = SAS^{-1}$ is similar to A , let $y = Sx$; then since $\dot{y} = S\dot{x}$, we can rewrite (L) as $S^{-1}\dot{y} = AS^{-1}y$, and hence as $\dot{y} = By$. It follows from this observation and from Theorem 8.A.5 that to understand linear differential equations, it is enough to understand linear differential equations defined by real Jordan matrices.

Example 8.B.2. Linear dynamics on the plane. There are three generic types of 2×2 matrices: diagonalizable matrices with two real eigenvalues, diagonalizable matrices with two complex eigenvalues, and nondiagonalizable matrices with a single real eigenvalue. The corresponding real Jordan forms are a diagonal matrix (which contains two J_1 blocks), a J_2 matrix, and a J_3 matrix, respectively. We therefore consider linear differential equations based on these three types of real Jordan matrices.

When A is diagonal, the linear equation (L) and its solution from initial condition $x_0 = \xi$ are of the following form:

$$\dot{x} = Ax = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}; \quad x_t = \begin{pmatrix} \xi_1 e^{\lambda t} \\ \xi_2 e^{\mu t} \end{pmatrix}.$$

The phase diagrams in Figure 8.B.1 show that the behavior of this dynamic depends on the values of the eigenvalues λ and μ : if both are negative, the origin is a *stable node*, if their signs differ, the origin is a *saddle*, and if both are positive, the origin is an *unstable node*.

Now suppose that A is the real Jordan form of a matrix with complex eigenvalues

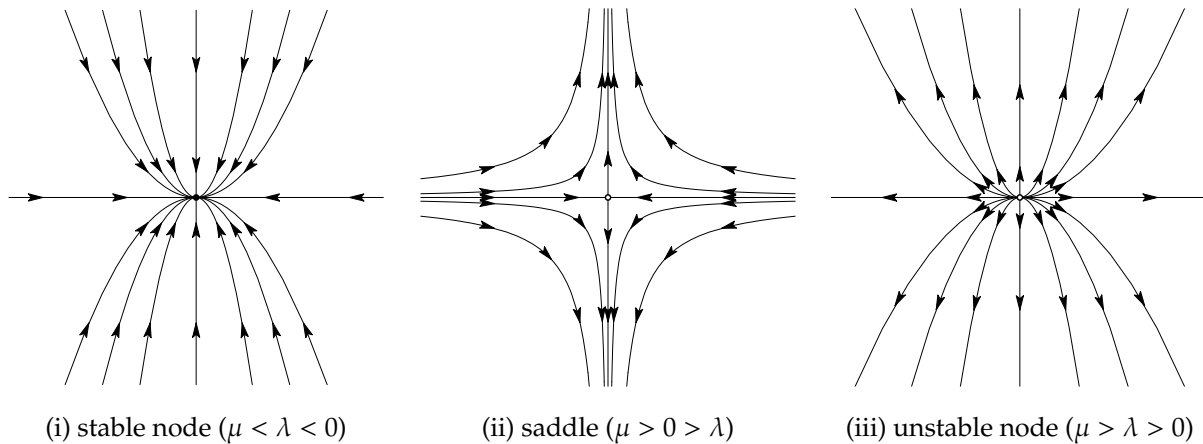


Figure 8.B.1: Linear dynamics on the plane: two real eigenvalues λ, μ .





Figure 8.B.2: Linear dynamics on the plane: complex eigenvalues $a \pm ib, b < 0$.

$a \pm ib$. Then we have

$$\dot{x} = Ax = \begin{pmatrix} a & b \\ -b & a \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}; \quad x_t = \begin{pmatrix} \xi_1 e^{at} \cos bt + \xi_2 e^{at} \sin bt \\ -\xi_1 e^{at} \sin bt + \xi_2 e^{at} \cos bt \end{pmatrix}.$$

Phase diagrams for this equation are presented in Figure 8.B.2. Evidently, the stability of the origin is determined by the real part of the eigenvalues: if $a < 0$, the origin is a *stable spiral*, while if $a > 0$, the origin is an *unstable spiral*. In the nongeneric case where $a = 0$, the origin is a *center*, with each solution following a closed orbit around the origin. The value of b determines the orientation of the cycles. The diagrams in Figure 8.B.2 use $b < 0$, which causes solutions to cycle counterclockwise; had we chosen $b > 0$, these orientations would have been reversed.



(i) stable improper node ($\lambda < 0$) (ii) unstable improper node ($\lambda > 0$)

Figure 8.B.3: Linear dynamics on the plane: A not diagonalizable, one real eigenvalue λ .

Finally, suppose that A is the real Jordan form of a nondiagonalizable matrix with lone eigenvalue λ . Then we obtain

$$\dot{x} = Ax = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}; \quad x_t = \begin{pmatrix} \xi_1 e^{\lambda t} + \xi_2 t e^{\lambda t} \\ \xi_1 e^{\lambda t} \end{pmatrix}.$$

The phase diagrams in Figure 8.B.3 reveal the origin to be an *improper* (or *degenerate*) *node*. It is stable if the eigenvalue λ is negative and unstable if λ is positive. §

8.B.2 Solutions

The Picard-Lindelöf Theorem (Theorem 4.A.2) implies that for any matrix $A \in \mathbf{R}^{n \times n}$ there is a unique solution to the linear equation (L) starting from each initial condition $\xi \in \mathbf{R}^n$. While solutions of nonlinear differential equations generally cannot be expressed in closed form, the solutions to linear equations can always be described explicitly. In the planar case, Example 8.B.2 provided explicit formulas when A is a Jordan matrix, and the solutions for other matrices can be obtained through a change of variable. Similar logic can be employed in the general case, yielding the following result:

Theorem 8.B.3. *Let $\{x_t\}_{t \in (-\infty, \infty)}$ be the solution to (L) from initial condition x_0 . Then each coordinate of x_t is a linear combination of terms of the form $t^k e^{at} \cos(bt)$ and $t^k e^{at} \sin(bt)$, where $a + ib \in \mathbf{C}$ is an eigenvalue of A and $k \in \mathbf{Z}_+$ is less than the algebraic multiplicity of this eigenvalue.*

For analytic purposes, it is often convenient to express solutions of the linear equation (L) in terms of *matrix exponentials*. Given a matrix $A \in \mathbf{R}^{n \times n}$, we define $e^A \in \mathbf{R}^{n \times n}$ by

applying the series definition of the exponential function to the matrix A : that is,

$$(8.46) \quad e^A = \sum_{k=0}^{\infty} \frac{A^k}{k!},$$

where A^k denotes the k th power of A and $A^0 \equiv I$ is the identity matrix.

Recall that the flow $\phi : (-\infty, \infty) \times \mathbf{R}^n \rightarrow \mathbf{R}^n$ generated by (L) is defined by $\phi_t(\xi) = x_t$, where $\{x_t\}_{t \in (-\infty, \infty)}$ is the solution to (L) with initial condition $x_0 = \xi$. Theorem 8.B.4 provides a concise expression for solutions to (L) in terms of matrix exponentials.

Theorem 8.B.4. *The flow of (L) is $\phi_t(\xi) = e^{At}\xi$.*

A benefit of representing solutions to (L) in this way is that properties established for matrix exponentials can be given immediate interpretations in terms of solutions to (L). For examples, consider these properties:

Proposition 8.B.5.

- (i) *If A and B commute, then $e^{A+B} = e^B e^A$.*
- (ii) *If $B = S^{-1}AS$, then $e^B = S^{-1}e^A S$.*
- (iii) *$e^{(A')} = (e^A)'$.*

Applying part (i) of the proposition to matrices As and At yields the group property of the flow of (L): $\phi_{s+t}(\xi) = \phi_t(\phi_s(\xi))$. Part (ii) shows that linear flows generated by similar matrices are linearly conjugate (i.e., that they are equivalent up to a linear change of variables), as we discussed before Example 8.B.2. Applying parts (iii) and (i) to At when A is skew-symmetric shows that in this case, e^{At} is an orthogonal matrix: thus, for each fixed time t , the map $\xi \mapsto \phi_t(\xi)$ is a Euclidean isometry (cf Figure 8.B.2(ii)).

8.B.3 Stability and Hyperbolicity

Theorem 8.B.3 shows in generic cases, the stability of the origin under the linear equation (L) is determined by the eigenvalues $\{a_1 + i b_1, \dots, a_n + i b_n\}$ of A : more precisely, by the real parts a_i of these eigenvalues. If each a_i is negative, then all solutions to (L) converge to the origin; in this case, the origin is called a *sink*, and the flow $\phi_t(x) = e^{At}x$ is called a *contraction*. If instead each a_i is positive, then all solutions besides the stationary solution at the origin move away from the origin; in this case, the origin is called a *source*, and the flow of (L) is called an *expansion*.

When the origin is a sink, solutions to (L) converge to the origin at an exponential rate. Define a norm on \mathbf{R}^n by $\|x\|_S = |S^{-1}x|$, where S is the similarity matrix from the Jordan

decomposition $J = S^{-1}AS$ of A . Then for any $a > 0$ satisfying $a < |a_i|$ for all $i \in \{1, \dots, n\}$, the flow ϕ of (L) satisfies

$$\mathbf{0} \text{ is a sink} \Leftrightarrow \|\phi_t(\xi)\|_S \leq e^{-at} \|\xi\|_S \quad \text{for all } t \geq 0 \text{ and all } \xi \in \mathbf{R}^n.$$

A similar statement in terms of the Euclidean norm holds if one introduces an appropriate multiplicative constant $C = C(a) \geq 1$:

$$(8.47) \quad \mathbf{0} \text{ is a sink} \Leftrightarrow |\phi_t(\xi)| \leq Ce^{-at} |\xi| \quad \text{for all } t \geq 0 \text{ and all } \xi \in \mathbf{R}^n.$$

If the origin is the source, analogous statements hold if time is run backward: for instance,

$$(8.48) \quad \mathbf{0} \text{ is a source} \Leftrightarrow |\phi_t(\xi)| \leq Ce^{-a|t|} |\xi| \quad \text{for all } t \leq 0 \text{ and all } \xi \in \mathbf{R}^n.$$

More generally, the flow of (L) may be contracting in some directions and expanding in others. In the generic case in which each real part a_i of an eigenvalue of A is nonzero, the differential equation $\dot{x} = Ax$, its rest point at the origin, and its flow $\phi_t(x) = e^{At}x$ are all said to be *hyperbolic*. Hyperbolic linear flows come in three varieties: contractions (if all a_i are negative, as in (8.47)), expansions (if all a_i are positive, as in (8.48)), and *saddles* (if there is at least one a_i of each sign). If a linear flow is hyperbolic, then the origin is globally asymptotically stable if it is a sink, and it is unstable otherwise.

If the flow of (L) is hyperbolic, then A has k eigenvalues with negative real part (counting algebraic multiplicities) and $n - k$ eigenvalues with positive real part. In this case, we can view $\mathbf{R}^n = E^s \oplus E^u$ as the direct sum of subspaces of dimensions $\dim(E^s) = k$ and $\dim(E^u) = n - k$, where the *stable subspace* E^s contains all solutions of (L) that converge to the origin at an exponential rate (as in (8.47)), while the *unstable subspace* E^u contains all solutions of (L) that converge to the origin at an exponential rate if time is run backward (as in (8.48)).

If A is real diagonalizable, then it follows easily from Theorem 8.B.3 that E^s and E^u are the spans of the eigenvectors of A corresponding to the eigenvalues of A with negative and positive real part, respectively. More generally, E^s and E^u can be computed by way of the real Jordan form $J = S^{-1}AS$ of A . Arrange S and J so that the Jordan blocks of J corresponding to eigenvalues of A with negative real parts appear in the first k rows and columns, while the blocks corresponding to eigenvalues with positive real parts appear in the remaining $n - k$ rows and columns. Then E^s is the span of the first k columns of the similarity matrix S , and E^u is the span of the remaining $n - k$ columns of S . (The columns of S are the real and imaginary parts of the so-called *generalized eigenvectors* of A —see

Example 8.A.7.)

8.C Linearization of Nonlinear Differential Equations

Virtually all of the differential equations we study in this book are nonlinear. Nevertheless, when studying the behavior of nonlinear equations in the neighborhood of a rest point, the theory of linear equations takes on a central role.

Consider the C^1 differential equation

$$(D) \quad \dot{x} = V(x)$$

with rest point x^* . By the definition of the derivative, we can approximate the value of V in the neighborhood of x^* via

$$V(y) = \mathbf{0} + DV(x^*)(y - x^*) + o(|y - x^*|).$$

This suggests that the behavior of the dynamic (D) near x^* can be approximated by the behavior near the origin of the linear equation

$$(L) \quad \dot{y} = DV(x^*)y.$$

Making this idea precise requires some new definitions. Let X and Y be subsets of \mathbf{R}^n . The function $h : X \rightarrow Y$ is called a *homeomorphism* if it is bijective (i.e., one-to-one and onto) and continuous with a continuous inverse. Now let $I \subseteq \mathbf{R}$ be an interval containing 0, and let $\phi : I \times X \rightarrow X$ and $\psi : I \times Y \rightarrow Y$ be two flows. We say that ϕ and ψ are *topologically conjugate* on X and Y if there is a homeomorphism $h : X \rightarrow Y$ such that $\phi_t(x_0) = h^{-1} \circ \psi_t \circ h(x_0)$ for all times $t \in I$. In other words, ϕ and ψ are topologically conjugate if there is a continuous map with continuous inverse that sends trajectories of ϕ to trajectories of ψ (and vice versa), preserving the rate of passage of time. Therefore, to find $\phi_t(x_0)$, the position at time t under flow ϕ when the initial state is $x_0 \in X$, one can apply $h : X \rightarrow Y$ to x_0 to obtain the transformed initial condition $y_0 = h(x_0) \in Y$, then run the flow ψ from y_0 for t time units, and finally apply h^{-1} to the result. We summarize this construction in the diagram below:

$$\begin{array}{ccc} x_0 & \xrightarrow{h} & h(x_0) \\ \phi_t \downarrow & & \downarrow \psi_t \\ \phi_t(x_0) & \xleftarrow{h^{-1}} & \psi_t(h(x_0)) \end{array}$$

The use of linearization to study the behavior of nonlinear differential equations around fixed points is justified by the *Hartman-Grobman Theorem*.

Theorem 8.C.1 (The Hartman-Grobman Theorem). *Let ϕ and ψ be the flows of the C^1 equation (D) and the linear equation (L), where x^* is a hyperbolic rest point of (D). Then there exist neighborhoods O_{x^*} of x^* and O_0 of the origin $\mathbf{0}$ on which ϕ and ψ are topologically conjugate.*

Combining the Hartman-Grobman Theorem with our analysis in Section 8.B.3 provides a simple characterization of the stability of hyperbolic rest points of (D).

Corollary 8.C.2. *Let x^* be a hyperbolic rest point of (D). Then x^* is asymptotically stable if all eigenvalues of $DV(x^*)$ have negative real parts, and x^* is unstable otherwise.*

By virtue of these results, we say that x^* is *linearly stable* if the eigenvalues of $DV(x^*)$ all have negative real part. While the Hartman-Grobman Theorem implies that a linearly stable rest point is asymptotically stable, it can be shown further that solutions starting near a linearly stable rest point converge to it at an exponential rate, as in equation (8.47).

We say that x^* is *linearly unstable* if $DV(x^*)$ has at least one eigenvalue with positive real part. (We do not require x^* to be hyperbolic.) It can be shown that as long as one eigenvalue of $DV(x^*)$ has positive real part, most solutions of (D) will move away from x^* at an exponential rate.

While the topological conjugacy established in Theorem 8.C.1 is sufficient for local stability analysis, one should understand that topological conjugacy need not preserve the *geometry* of a flow. The following result for linear equations makes this point clear.

Theorem 8.C.3. *Let $\dot{x} = Ax$ and $\dot{y} = By$ be hyperbolic linear differential equations on \mathbf{R}^n with flows ϕ and ψ . If A and B have the same numbers of eigenvalues with negative real part (counting algebraic multiplicities), then ϕ and ψ are topologically conjugate throughout \mathbf{R}^n .*

Looking back at Example 8.B.2, we see that the phase diagrams of stable nodes (Figure 8.B.1(i)), stable spirals (Figure 8.B.2(i)), and stable improper nodes (Figure 8.B.3(i)) have very different appearances. Nevertheless, Theorem 8.C.3 reveals that the flows described in these figures are topologically conjugate—that is, they can be continuously transformed into one another! To ensure that the geometry of phase diagrams is preserved, one needs not only topological conjugacy, but rather *differentiable conjugacy*: that is, conjugacy under a *diffeomorphism* (a differentiable transformation with differentiable inverse). As it turns out, it is possible to establish a local differentiable conjugacy between (D) near x^* and (L) near $\mathbf{0}$ if V is sufficiently smooth, and if the eigenvalues of $DV(x^*)$ are distinct and satisfy a mild nonresonance condition (see the Notes).

Much additional information about the flow of (D) can be surmised from the derivative matrix $DV(x^*)$ at a hyperbolic rest point x^* . Suppose that $DV(x^*)$ has k eigenvalues with negative real part and $n - k$ eigenvalues with positive real part, counting algebraic multiplicities. The *Stable Manifold Theorem* tells us that within some neighborhood of x^* , there is k dimensional *local stable manifold* M_{loc}^s on which solutions converge to x^* at an exponential rate (as in (8.47)), and an $n - k$ dimensional *local unstable manifold* M_{loc}^u on which solutions converge to x^* at an exponential rate if time is run backward (as in (8.48)).

Moreover, both of these manifolds can be extended globally: the k dimensional (*global*) *stable manifold* M^s includes all solutions of (D) that converge to x^* , while the $n - k$ dimensional (*global*) *unstable manifold* M^u includes all solutions that converge to x^* as time runs backward. Among other implications of the existence of these manifolds, it follows that if x^* is hyperbolic and unstable, then the set M^s of states from which solutions converge to x^* is of measure zero, while the complement of this set is open, dense, and of full measure.

8.N Notes

Section 8.1: Theorem 8.1.1 is established by Bomze (1986) for the replicator dynamic and by Nachbar (1990) for general imitative dynamics; see also Weibull (1995).

Section 8.2: This section follows Sandholm (2001a). Bomze (2002) provides an exhaustive treatment of local stability under the replicator dynamic for single-population linear potential games (which are generated by matching in common interest games), and the connections between this stability analysis and quadratic programming.

Section 8.3: The notion of an evolutionarily stable strategy was introduced by Maynard Smith and Price (1973). The distinction between evolutionarily stable strategies and evolutionarily stable states is emphasized by Thomas (1984). General references on ESS theory include the survey of Hines (1987) and the monographs of Bomze and Pötscher (1989) and Cressman (1992).

Most early work on evolutionarily stable strategies considers single populations matched to play symmetric normal form games. The original ESS definition of Maynard Smith and Price (1973) (see also Maynard Smith (1974)) is via conditions (8.7) and (8.9). The characterizations of ESS in terms of invasion barriers (8.6) and invasion of nearby states (8.2) in this linear setting are due to Taylor and Jonker (1978) and Hofbauer et al. (1979), respectively. That invasion barriers are uniform in this setting (Exercise 8.3.3(ii)) is pointed out explicitly by Vickers and Cannings (1987); see also Bomze (1986), Zeeman (1980), and Hofbauer and Sigmund (1988).

Turning now to single-population games with nonlinear payoffs, Theorem 8.3.5 is

announced in Pohley and Thomas (1983), where a state satisfying condition (8.2) is called a “local ESS”. The theorem is proved in Thomas (1985), and our proof in the text is a streamlined version of Thomas’s (1985) proof. Theorem 8.3.1 is due to Bomze (1991), who calls a state satisfying our ESS condition (8.2) “strongly uninvadable”, and a state satisfying the uniform invasion barrier condition (8.4) “uninvadable”; see also Bomze and Weibull (1995). Exercise 8.3.3(ii) is Example 18 of Bomze and Pötscher (1989).

Definition (8.16) of Taylor ESS is introduced by Taylor (1979); also see Schuster et al. (1981a). Exercise 8.3.8 is essentially due to Selten (1980); see also van Damme (1991) and Swinkels (1992).

Definition (8.17) of Cressman ESS is due to Cressman (1992, 1995, 1996b, 2006) and Cressman et al. (2001), who refer to it as *monomorphic ESS* and as *p-species ESS*. To show that this definition is the heir of the single-population ESS of Maynard Smith and Price (1973), these papers study a collection of p -dimensional replicator systems, with one system for each strategy profile $y = (y^1, \dots, y^p)$ other than the candidate for stability, $x = (x^1, \dots, x^p)$. The p th component of the state variable in the p -dimensional system describes the fraction of the p th species using mixed strategy y^p ; the remainder of the species uses the incumbent mixed strategy x^p . Results in Cressman et al. (2001) and Cressman (2006) imply that the origin (i.e., the state at which all members of each species p choose mixed strategy x^p) is asymptotically stable in each such system if and only if x satisfies condition (8.17). (To see this, use equations (1) and (6) in Cressman et al. (2001) to show that the B -matrix conditions appearing in Theorems 3 and 5 of that paper are equivalent to condition (8.17) here.) Interestingly, Cressman (1992) shows that in two-population linear games, any Cressman ESS is asymptotically stable under the replicator dynamic, but that this is not true in games played by more than two populations.

The notion of regular ESS is introduced in a single-population setting by Taylor and Jonker (1978), who prove that a regular ESS is asymptotically stable under the replicator dynamic. Taylor (1979) extends this notion to the multi-population case, and observes that the stability result for the replicator dynamic extends to this setting.

Thomas (1985), Swinkels (1992), Balkenborg and Schlag (2001, 2007), and Cressman (2003) consider set-valued generalizations of the ESS concept, which are particularly useful in the context of matching in extensive form games.

Section 8.4: Theorem 8.4.1(i) on the local stability of ESS under the replicator dynamic is one of the earliest results on evolutionary game dynamics; see Taylor and Jonker (1978), Taylor (1979), Hofbauer et al. (1979), Zeeman (1980), and Schuster et al. (1981a). Theorem 8.4.1(ii) follows easily from results of Nagurney and Zhang (1997); see also Sandholm et al. (2008). The results in Section 8.4.2 are extensions of ones from Hofbauer and

Sandholm (2009). For the Theorem of the Maximum, see Ok (2007). Theorem 8.4.7 is due to Sandholm (2009d). Hofbauer (1995b) establishes the asymptotic stability of ESS under the best response dynamic in a single population normal form context using a construction different from the one presented here.

Section 8.5: Lemma 8.5.1 is due to Hines (1980); see also Hofbauer and Sigmund (1988), Hopkins (1999), and Sandholm (2007a). Versions of Theorems 8.5.2 and 8.5.6 can be found in Taylor and Jonker (1978), Taylor (1979), Hines (1980), and Cressman (1992, 1997). Example 8.5.5 is taken from Zeeman (1980). Theorem 8.5.9 is due to Cressman (1997).

Section 8.6: Linearization of perturbed best response dynamics is studied by Hopkins (1999, 2002), Hofbauer (2000), Hofbauer and Sandholm (2002, 2007), Hofbauer and Hopkins (2005), and Sandholm (2007a). Exercise 8.6.3 is used in Sandholm (2007a) to show that Nash equilibria of normal form games can always be purified (in the sense of Harsanyi (1973)) in an evolutionarily stable fashion through an appropriate choice of payoff noise. See Ellison and Fudenberg (2000) and Ely and Sandholm (2005) for related results. Example 8.6.5 is due to Hopkins (1999). Hopkins (2002) uses this result to show that the replicator dynamic closely approximates the evolution of choice probabilities under stochastic fictitious play. Hofbauer et al. (2007) use similar ideas to establish an exact relationship between the long run time averaged behavior of the replicator dynamic and the long run behavior of the best response dynamic.

Appendix 8.A: Horn and Johnson (1985) is an outstanding general reference on matrix analysis. Many of the results we described are also presented in Hirsch and Smale (1974).

Appendix 8.B: Both Hirsch and Smale (1974) and Robinson (1995) provide thorough treatments of linear differential equations at the undergraduate and graduate levels, respectively.

Appendix 8.C: Robinson (1995) is an excellent reference on dynamical systems in general and on linearization in particular. For more on differentiable conjugacy around rest points, see Hartman (1964).

Nonconvergence of Evolutionary Dynamics

9.0 Introduction

We began our study of the global behavior of evolutionary dynamics in Chapter 7, focusing on combinations of games and dynamics generating global or almost global convergence to equilibrium. The analysis there demonstrated that global payoff structure—in particular, the structure captured in the definitions of potential, stable, and supermodular games—makes compelling evolutionary justifications of the Nash prediction possible.

On the other hand, once we move beyond these classes of well-behaved games, it is not clear how often convergence will occur. The present chapter counterbalances Chapter 7 by investigating nonconvergence of evolutionary dynamics for games, describing a variety of environments in which cycling or chaos offer the best predictions of long run behavior.

Section 9.1 leads with a study of conservative properties of evolutionary dynamics, focusing on the existence of constants of motion and on the preservation of volume under the replicator and projection dynamics. Section 9.2 continues with a panoply of examples of nonconvergence. Among other things, this section offers games in which no reasonable evolutionary dynamic converges to equilibrium, demonstrating that no evolutionary dynamic can provide a blanket justification for the prediction of Nash equilibrium play. Section 9.3 proceeds by offering examples of chaotic evolutionary dynamics—that is, dynamics exhibiting complicated attracting sets and sensitive dependence on initial conditions.

The possibility of nonconvergence has surprising implications for evolutionary support of traditional solution concepts. Under dynamics that satisfy Nash stationarity (NS), solution trajectories that converge necessarily converge to Nash equilibria. But since no

reasonable evolutionary dynamic converges in all games, general support for standard solution concepts is not assured.

Since the Nash prediction is not always supported by an evolutionary analysis, it is natural to turn to a less demanding notion—namely, the elimination of strategies that are strictly dominated by a pure strategy. As this requirement is the mildest employed in standard game theoretic analyses, it is natural to expect to find support for this requirement via an evolutionary approach.

In Section 9.4, we present the striking finding that evolutionary dynamics satisfying four mild conditions—continuity, Nash stationarity, positive correlation, and innovation—do not eliminate strictly dominated strategies in all games. Moreover, while we saw in Chapter 7 that imitative dynamics and the best response dynamic eliminate strictly dominated strategies, we show here that small perturbations of these dynamics do not. This analysis demonstrates that evolutionary dynamics provide surprisingly little support for a basic rationality criterion.

As always, the appendices provide the mathematical background necessary for our analysis. Appendix 9.A describes some classical theorems on nonconvergence used throughout the chapter. Appendix 9.B introduces the notion of an attractor of a dynamic, and establishes the continuity properties of attractors that underlie our analysis of dominated strategies.

9.1 Conservative Properties of Evolutionary Dynamics

It is often impossible to provide precise descriptions of long run behavior under nonconvergent dynamics. An important exception occurs in cases where the dynamics lead certain quantities to be preserved. We explore this idea in the current section, where we argue that in certain strategic environments, the replicator and projection dynamics exhibit noteworthy conservative properties.

9.1.1 Constants of Motion in Null Stable Games

In Section 7.2.1, we introduced *null stable* population games. These games are defined by the requirement that

$$(y - x)'(F(y) - F(x)) = 0 \text{ for all } x, y \in X,$$

and include zero-sum games (Example 3.3.7) and multi-zero-sum games (Exercise 3.3.9) as special cases.

In Exercise 7.2.2, we saw if x^* is an interior Nash equilibrium of a null stable game $F : X \rightarrow \mathbf{R}^n$, then the value of the function

$$E_{x^*}(x) = |x - x^*|^2$$

is preserved along interior segments of solution trajectories of the projection dynamic: thus, as these segments are traversed, Euclidean distance from the equilibrium x^* is fixed. Similar conclusions hold for interior solutions of the replicator dynamic: Exercise 7.2.5 shows that such solutions preserve the value of the function

$$H_{x^*}(x) = \sum_{p \in \mathcal{P}} h_{(x^*)^p}^p(x^p), \quad \text{where } h_{y^p}^p(x^p) = \sum_{i \in S^p(y^p)} y_i^p \log \frac{y_i^p}{x_i^p}$$

is a relative entropy function.

When x^* is interior, the level sets of E_{x^*} and H_{x^*} foliate from x^* like the layers of an onion. Each solution trajectory is limited to one of these layers, a manifold whose dimension is one less than that of X .

Example 9.1.1. In Figure 5.3.1, we presented phase diagrams of the six basic evolutionary dynamics for standard Rock-Paper-Scissors,

$$F(x) = \begin{pmatrix} F_R(x) \\ F_P(x) \\ F_S(x) \end{pmatrix} = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \end{pmatrix} = \begin{pmatrix} x_S - x_P \\ x_R - x_S \\ x_P - x_R \end{pmatrix},$$

a zero-sum game with unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Figures 5.3.1(i) and 5.3.1(ii) show that interior solutions of the replicator and projection dynamics form closed orbits around x^* . These orbits describe the level sets of the functions E_{x^*} and H_{x^*} . Note that an affine transformation of H_{x^*} yields a simpler constant of motion for the replicator dynamic, $\mathcal{H}(x) = -\sum_{i \in S} \log x_i$. §

When $\dim(X) > 2$, the level sets of E_{x^*} and H_{x^*} need not pin down the locations of interior solutions of (P) and (R). But if the null stable game F has multiple Nash equilibria, then there are multiple collections of level sets, and intersections of these sets do determine the positions of interior solutions.

Example 9.1.2. Consider the population game F generated by matching in the symmetric




Figure 9.1.1: Solutions of the projection dynamic on level set $E_{x^*}(x) = \frac{\sqrt{3}}{12}$, $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$.

zero-sum game A :

$$(9.1) \quad F(x) = Ax = \begin{pmatrix} 0 & -1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ -1 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} x_4 - x_2 \\ x_1 - x_3 \\ x_2 - x_4 \\ x_3 - x_1 \end{pmatrix}.$$

The Nash equilibria of F are the points on line segment NE connecting states $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $(0, \frac{1}{2}, 0, \frac{1}{2})$.

The arguments above show that interior solutions to the projection dynamic maintain a constant distance from every Nash equilibrium of F . This is illustrated in Figure 9.1.1, which presents solutions on the sphere inscribed in the pyramid X ; this is the level set on which E_{x^*} takes the value $\frac{\sqrt{3}}{12}$, where $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Each solution drawn in the figure is a circular closed orbit orthogonal to line segment NE .

Figure 9.1.2 presents solution trajectories of the replicator dynamic for game F . Diagrams (i) and (ii) show solutions on level sets of H_{x^*} where $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$; the first (smaller) level set is nearly spherical, while the second approximates the shape of the pyramid X . Diagrams (iii) and (iv) present solutions on level sets of H_{x^*} with $x^* = (\frac{3}{8}, \frac{1}{8}, \frac{3}{8}, \frac{1}{8})$ and $x^* = (\frac{1}{8}, \frac{3}{8}, \frac{1}{8}, \frac{3}{8})$. By our previous discussion, the intersection of the two level sets is a closed curve describing a single orbit of the dynamic. §

Example 9.3.2 will show that even in zero-sum games, very complicated dynamics can arise within the level sets of H_{x^*} .

Exercise 9.1.3. (i) Suppose that $A \in \mathbf{R}^{n \times n}$ is skew-symmetric. Show that the eigenvalues of A all have zero real part, and so that the number of nonzero eigenvalues is even.




Figure 9.1.2: Solutions of the replicator dynamic on level sets of H_{x^*} .

- (ii) Suppose that $A \in \mathbf{R}^{n \times n}$ is a symmetric zero-sum game that admits an interior Nash equilibrium x^* . Show that if n is even, then x^* is contained in a line segment consisting entirely of Nash equilibria. (Hint: Consider the matrix $\Phi A \Phi$.)

The previous analysis shows that in zero-sum games, typical solutions of the replicator dynamic do not converge. The next exercise shows that the *time averages* of these solutions do converge, and that the limits of the time averages are Nash equilibria.

Exercise 9.1.4. Convergence of time averages under the replicator dynamic. Let $F(x) = Ax$ be the population game generated by matching in the symmetric normal form game $A \in \mathbf{R}^{n \times n}$, and let $\dot{x} = V_F(x)$ be the replicator dynamic for this game. Suppose that $\{x_t\}_{t \geq 0}$ is a solution to V_F that is bounded away from $\text{bd}(X)$ (i.e., that there is an $\varepsilon > 0$ such that $(x_t)_i \geq \varepsilon$ for all $t \geq 0$ and $i \in S$). Let

$$\bar{x}_t = \frac{1}{t} \int_0^t x_s \, ds$$

be the average value of the state over the time interval $[0, t]$. Following the steps below, prove that $\{\bar{x}_t\}_{t \geq 0}$ converges to the set of (interior) Nash equilibria of F as t approaches infinity:

$$(9.2) \quad \lim_{t \rightarrow \infty} \min_{x^* \in NE(F)} |\bar{x}_t - x^*| = 0.$$

In particular, if F has a unique interior Nash equilibrium x^* , then $\{\bar{x}_t\}$ converges to x^* .

- (i) Define $y_t \in \mathbf{R}^n$ by $(y_t)_i = \log(x_t)_i$. Compute $\frac{d}{dt} y_t$.
(ii) Show that

$$\frac{1}{t} (y_t - y_0) = \frac{1}{t} \int_0^t (Ax_s - \mathbf{1}x'_s Ax_s) \, ds.$$

- (iii) Let \bar{x}^* be an ω -limit point of the trajectory $\{\bar{x}_t\}$. Show that $A\bar{x}^*$ is a constant vector, and hence that \bar{x}^* is a Nash equilibrium. (Hint: Use the fact that the trajectory $\{y_t\}$ is constrained to a compact set.)
(iv) Conclude that (9.2) holds. (Hint: Use the fact that the trajectory $\{\bar{x}_t\}$ is constrained to a compact set.)

Exercise 9.1.5. Prove that the conclusion of Exercise 9.1.4 continues to hold in a two-population matching setting.

Exercise 9.1.6. Explain why the argument in Exercise 9.1.4 does not allow its conclusion to be extended to matching in $p \geq 3$ populations.

9.1.2 Preservation of Volume

Let $\dot{x} = V(x)$ be differential equation on X with flow $\phi : \mathbf{R} \times X \rightarrow X$, and let μ denote Lebesgue measure on X . The differential equation is said to *volume preserving* (or *incompressible*) on $Y \subseteq X$ if for any measurable set $A \subseteq Y$, we have $\mu(\phi_t(A)) = \mu(A)$ for all $t \in \mathbf{R}$. Preservation of volume has strong implications for local stability of rest points: since an asymptotically stable rest point must draw in all nearby initial condition, no such rest points can exist in regions where volume is preserved (see Theorem 9.A.4).

We now show that in single population zero-sum games, the replicator dynamic is volume preserving after a well-chosen change in speed. Compared to the standard replicator dynamic, the *speed-adjusted replicator dynamic* on $\text{int}(X)$,

$$(9.3) \quad \dot{x}_i^p = q(x) x_i^p \hat{F}_i^p(x), \quad \text{where } q(x) = \prod_{r \in \mathcal{P}} \prod_{j \in S} \frac{1}{x_j^r},$$

moves relatively faster at states closer to the boundary of the simplex, with speeds approaching infinity as the boundary is approached. The solution trajectories of (9.3) have the same locations as those of the standard replicator dynamic (see Exercise 5.4.10), so the implications of volume preservation for stability of rest points extend immediately to the latter dynamic.

Theorem 9.1.7. *Let $F(x) = Ax$ be generated by matching in the symmetric zero-sum game $A = -A' \in \mathbf{R}^{n \times n}$. Then the dynamic (9.3) for F is volume preserving on $\text{int}(X)$. Therefore, no interior Nash equilibrium of F is asymptotically stable under the replicator dynamic.*

The proof of Theorem 9.1.7 is based on *Liouville's Theorem*, which tells us that the rate at which the dynamic $\dot{x} = V(x)$ expands or contracts volume near state x is given by the *divergence* $\text{div}V(x) \equiv \text{tr}(DV(x))$. More precisely, Liouville's Theorem tells us that

$$\frac{d}{dt} \mu(\phi_t(A)) = \int_{\phi_t(A)} \text{div}V(x) d\mu(x).$$

for each Lebesgue measurable set A . Thus, if $\text{div}V \equiv 0$, so that V is *divergence free*, then the flow ϕ is volume preserving. See Section 9.A.1 for a proof and further discussion of this result.

Proof. The replicator dynamic is described by the vector field $R : X \rightarrow TX$, where

$$R(x) = \text{diag}(x)(F(x) - \mathbf{1}x'F(x)).$$

Since $F(x) = Ax$, and since $x'Ax \equiv 0$ (because A is symmetric zero-sum), we can simplify the previous expression to

$$(9.4) \quad R(x) = \text{diag}(x)Ax.$$

The dynamic (9.3) can be written as

$$V(x) = q(x)R(x),$$

where q is the function from $\text{int}(X) \rightarrow \mathbf{R}_+$ defined in equation (9.3). If we can show that V is divergence free on $\text{int}(X)$, then our result will follow from Liouville's Theorem.

To compute $DV(x)$, let $\hat{q} : \text{int}(\mathbf{R}_+^n) \rightarrow \mathbf{R}_+$ and $\hat{R} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ be the natural extensions of q and R , so that $\nabla q(x) = \Phi \nabla \hat{q}(x)$ and $DR(x) = D\hat{R}(x)\Phi$. Then the chain rule implies that

$$(9.5) \quad DV(x) = q(x)DR(x) + R(x)\nabla q(x)' = \left(q(x)D\hat{R}(x) + R(x)\nabla \hat{q}(x)' \right) \Phi.$$

To evaluate this expression, write $[x^{-1}] = (\frac{1}{x_1}, \dots, \frac{1}{x_n})'$, and compute from equations (9.3) and (9.4) that

$$\nabla \hat{q}(x) = -\hat{q}(x)[x^{-1}] \quad \text{and} \quad D\hat{R}(x) = \text{diag}(x)A + \text{diag}(Ax).$$

Substituting into equation (9.5) yields

$$\begin{aligned} DV(x) &= q(x) \left((\text{diag}(x)A + \text{diag}(Ax) - \text{diag}(x)Ax[x^{-1}]') \right) \Phi \\ &= q(x) \left[(\text{diag}(x)A + \text{diag}(Ax) - \text{diag}(x)Ax[x^{-1}]') \right. \\ &\quad \left. - \frac{1}{n} (\text{diag}(x)A + \text{diag}(Ax) - \text{diag}(x)Ax[x^{-1}]') \mathbf{1}\mathbf{1}' \right]. \end{aligned}$$

Therefore,

$$\begin{aligned} \text{div} V(x) &= q(x) \left[\sum_{i \in S} x_i A_{ii} + \sum_{i \in S} (Ax)_i - \sum_{i \in S} x_i (Ax)_i \frac{1}{x_i} \right. \\ &\quad \left. - \frac{1}{n} \sum_{i \in S} x_i \sum_{j \in S} A_{ij} - \frac{1}{n} \sum_{i \in S} \sum_{j \in S} A_{ij} x_j + \frac{1}{n} \sum_{i \in S} \sum_{j \in S} x_i A_{ij} x_j \sum_{k \in S} \frac{1}{x_k} \right]. \end{aligned}$$

The first term in the brackets equals 0 since $A_{ii} = 0$; the second and third terms cancel; the fourth and fifth terms cancel since $A_{ij} = -A_{ji}$; and the sixth term is 0 since $x'Ax = 0$. We therefore conclude that $\text{div} V(x) = 0$ on $\text{int}(X)$, and hence that the flow of (9.3) is volume

preserving. The conclusion about asymptotic stability follows from Theorem 9.A.4. ■

Under single population matching, volume preservation under the replicator dynamic is only assured in zero-sum games. Remarkably, moving to multipopulation matching ensures volume preservation regardless of the payoffs in the underlying normal form game.

Suppose the population game F is generated by matching of members of $p \geq 2$ populations to play a p player normal form game. Since each agent's opponents in a match will be members of the other populations, the agent's payoffs do not depend on his own population's state: $F^p(x) \equiv F^p(x^{-p})$. Theorem 9.1.8 shows that this last condition is sufficient to prove that the flow of the replicator dynamic for F is volume preserving.

Theorem 9.1.8. *Let F be a game played by $p \geq 2$ populations that satisfies $F^p(x) \equiv F^p(x^{-p})$. Then the dynamic (9.3) for F is volume preserving on $\text{int}(X)$. Therefore, no interior Nash equilibrium of F is asymptotically stable under the replicator dynamic.*

Exercise 9.1.9. Prove Theorem 9.1.8. To simplify the notation, assume that each population is of unit mass. (Hint: To prove that the vector field V from equation (9.3) is divergence free, start by showing that the derivative matrix of V^p at x with respect to directions in TX^p is the $n^p \times n^p$ matrix

$$D_{TX^p} V^p(x) = q(x) \left(\text{diag}(\pi^p) - \bar{\pi}^p I - x^p (\pi^p)' - \text{diag}(x^p) \pi^p [(x^p)^{-1}]' + \bar{\pi}^p x^p [(x^p)^{-1}]' \right) \Phi,$$

where $\pi^p = F^p(x^{-p})$ and $\bar{\pi}^p = \bar{F}^p(x^{-p}) = (x^p)' \pi^p$.)

Analogues of Theorems 9.1.7 and 9.1.8 can be established for the projection dynamics via much simpler calculations, and without introducing a change in speed.

Exercise 9.1.10. Let $F(x) = Ax$ be generated by matching in the symmetric zero sum game $A = -A' \in \mathbf{R}^{n \times n}$. Show that the projection dynamic for F is volume preserving on $\text{int}(X)$.

Exercise 9.1.11. Let F be a game played by $p \geq 2$ unit-mass populations that satisfies $F^p(x) \equiv F^p(x^{-p})$. Show that the projection dynamic for F is volume preserving on $\text{int}(X)$.

9.2 Games with Nonconvergent Evolutionary Dynamics

In this section, we introduce examples of games for which many evolutionary dynamics fail to converge to equilibrium.

9.2.1 Circulant Games

The matrix $A \in \mathbf{R}^{n \times n}$ is called a *circulant matrix* if it is of the form

$$A = \begin{pmatrix} a_0 & a_1 & \cdots & a_{n-2} & a_{n-1} \\ a_{n-1} & a_0 & a_1 & \cdots & a_{n-2} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ a_2 & \cdots & a_{n-1} & a_0 & a_1 \\ a_1 & a_2 & \cdots & a_{n-1} & a_0 \end{pmatrix}.$$

When we view A as the payoff matrix for a symmetric normal form game, we refer to A as a *circulant game*. Such games always include the central state $x^* = \frac{1}{n}\mathbf{1}$ among their Nash equilibria. Note that Rock-Paper-Scissors games are circulant games with $n = 3$, $a_0 = 0$, $a_1 = -l$, and $a_2 = w$. Most of the specific games considered below will also have diagonal payoffs equal to 0.

Their symmetric structure make circulant games simple to analyze. In doing so, we will find it convenient to refer to strategies modulo n .

Exercise 9.2.1. Verify that the eigenvalue/eigenvector pairs of the circulant matrix A are

$$(9.6) \quad (\lambda_k, v_k) = \left(\sum_{j=0}^{n-1} a_j \iota_n^{jk}, (1, \iota_n^k, \dots, \iota_n^{(n-1)k})' \right), \quad k = 0, \dots, n-1,$$

where $\iota_n = \exp(\frac{2\pi i}{n}) = \cos(\frac{2\pi}{n}) + i \sin(\frac{2\pi}{n})$ is the n th root of unity.

Exercise 9.2.2. Let $F(x) = Ax$ be generated by matching in the circulant game A , and let $\dot{x} = R(x) = \text{diag}(x)(Ax - \mathbf{1}x'Ax)$ be the replicator dynamic for F . Show that the derivative matrix of R at the Nash equilibrium $x^* = \frac{1}{n}\mathbf{1}$ is the circulant matrix

$$DR(x^*) = \frac{1}{n}(A - 2\mathbf{1}\mathbf{1}'\bar{a}),$$

where $\bar{a} = \frac{1}{n}\mathbf{1}'a$ is the average of the components of the vector $a = (a_0, a_1, \dots, a_{n-1})'$. It then follows from the previous exercise that the eigenvalue/eigenvector pairs (λ_k, v_k) of $DR(x^*)$ are given by

$$(9.7) \quad (\lambda_k, v_k) = \left(\frac{1}{n} \sum_{j=0}^{n-1} (a_j - 2\bar{a}) \iota_n^{jk}, (1, \iota_n^k, \dots, \iota_n^{(n-1)k})' \right), \quad k = 0, \dots, n-1.$$

Example 9.2.3. The hypercycle system. Suppose that $a_0 = \dots = a_{n-2} = 0$ and that $a_{n-1} = 1$, so that each strategy yields a positive payoff only against the strategy that precedes it




Figure 9.2.1: Eigenvalues of the hypercycle system.

(modulo n). In this case, $x^* = \frac{1}{n}\mathbf{1}$ is the unique Nash equilibrium of F , and the replicator dynamic for A is known as the *hypercycle system*.

We determine the local stability of the rest point x^* by considering the eigenvalues of $DR(x^*)$. Substituting into equations (9.6) and (9.7) shows that the eigenvector/eigenvalue pairs are of the form

$$(\lambda_k, v_k) = \left(\frac{1}{n} l_n^{(n-1)k} - \frac{2}{n^2} \sum_{j=0}^{n-1} l_n^{jk}, (1, l_n^k, \dots, l_n^{(n-1)k})' \right), \quad k = 0, \dots, n-1.$$

Eigenvalue $\lambda_0 = \frac{1}{n} - \frac{2}{n} = -\frac{1}{n}$ corresponds to eigenvector $v_0 = \mathbf{1}$ and so has no bearing on the stability analysis. For $k \geq 1$, the sum in the formula for λ_k vanishes (why?), leaving us with $\lambda_k = \frac{1}{n} l_n^{(n-1)k} = \frac{1}{n} l_n^{-k}$. The stability of x^* therefore depends on whether any λ_k with $k > 0$ has positive real part. As Figure 9.2.1 illustrates, this largest real part is negative when $n \leq 3$, zero when $n = 4$, and positive when $n \geq 5$. It follows that x^* is asymptotically stable when $n \leq 3$, but unstable when $n \geq 5$. Exercise 9.2.4 shows that the local stability results can be extended to global stability results, and that global stability can also be proved when $n = 4$. When $n \geq 5$, it is possible to show that the boundary of X is repelling, as it is in the lower dimensional cases, and that the dynamic admits a stable periodic orbit (see the Notes). §

Exercise 9.2.4. Consider the function $\mathcal{H} : \text{int}(X) \rightarrow \mathbf{R}$ defined by $\mathcal{H}(x) = -\sum_{i \in S} \log x_i$ (cf Example 9.1.1.)

- (i) Show that under the hypercycle equation with $n = 2$ or 3 , \mathcal{H} is a strict Lyapunov function on $\text{int}(X)$, and hence that x^* is globally asymptotically stable with respect to $\text{int}(X)$.

- (ii) Show that under the hypercycle equation with $n = 4$ we have $\dot{H}(x) \leq 0$ on $\text{int}(X)$, with equality if and only if x lies in $Y = \{y \in \text{int}(X) : y_1 + y_3 = y_2 + y_4\}$. Show that the sole invariant subset of Y is $\{x^*\}$. Then use Theorems 7.B.2 and 7.B.4 and Proposition 7.A.1(iii) to conclude that x^* is globally asymptotically stable with respect to $\text{int}(X)$.

Example 9.2.5. Monocyclic games. A circulant game A is *monocyclic* if $a_0 = 0, a_1, \dots, a_{n-2} \leq 0$, and $a_{n-1} > 0$. Let $\bar{a} = \frac{1}{n} \sum_i a_i$. If we assume that $\bar{a} < 0$, then the Nash equilibrium $x^* = \frac{1}{n} \mathbf{1}$, which yields a payoff of \bar{a} for each strategy, is the unique interior Nash equilibrium of $F(x) = Ax$. More importantly, there is an open, dense, full measure set of initial conditions from which the best response dynamic for $F(x) = Ax$ converges to a limit cycle; this limit cycle is contained in the set where $M(x) = \max_{i \in S} F_i(x)$ equals 0.

Here is a sketch of the proof. Consider a solution trajectory $\{x_t\}$ of the best response dynamic that lies in set $B_1 = \{x \in X : \text{argmax}_{i \in S} F_i(x) = \{1\}\}$ during time interval $[0, T)$. For any $t \in [0, T)$, we have that

$$x_t = e^{-t} x_0 + (1 - e^{-t}) e_1.$$

Since the diagonal elements of A all equal zero, it follows that

$$(9.8) \quad M(x_t) = F_1(x_t) = e^{-t} F_1(x_0) = e^{-t} M(x_0).$$

For $j \notin \{1, 2\}$ we have that

$$(9.9) \quad F_j(x_t) = e^{-t} F_j(x_0) + (1 - e^{-t}) A_{j1} \leq e^{-t} F_j(x_0) < e^{-t} F_1(x_0) = F_1(x_t).$$

Equations (9.8) and (9.9) and the fact that

$$F_1(e_1) = 0 < a_{n-1} = F_2(e_1)$$

imply that a solution starting in region B_1 must hit the set $B_{12} = \{x \in X : \text{argmax}_{i \in S} F_i(x) = \{1, 2\}\}$, and then immediately enter region $B_2 = \{x \in X : \text{argmax}_{i \in S} F_i(x) = \{2\}\}$.

Repeating the foregoing argument shows that the trajectory next enters best response regions B_3, B_4, \dots, B_0 in succession before returning to region B_1 . Therefore, if we denote by B the set of states at which there are at most two best responses, then B is forward invariant under the best response dynamic. Moreover, equation (9.8) implies that the maximal payoff $M(x_t)$ approaches 0 along all solution trajectories in B .

In light of this discussion, we can define the *return map* $r : B_{12} \rightarrow B_{12}$, where $r(x)$ is the position at which a solution starting at $x \in B_{12}$ first returns to B_{12} . All fixed points of r lie in

$M^{-1}(0)$. In fact, it can be shown that r is a contraction on $M^{-1}(0)$ for an appropriate choice of metric, and so that r has a unique fixed point (see the Notes). We therefore conclude that any solution trajectory starting in the open, dense, full measure set B converges to the closed orbit that passes through the unique fixed point of the return map r . §

9.2.2 Continuation of Attractors for Parameterized Games

The games we construct in the examples to come will generate nonconvergent behavior for large classes of evolutionary dynamics. Recall our general formulation of evolutionary dynamics from Chapter 4: each revision protocol ρ defined a map from population games F to differential equations $\dot{x} = V_F(x)$ via

$$(9.10) \quad \dot{x}_i^p = (V_F)_i^p(x) = \sum_{j \in S^p} x_j^p \rho_{ji}^p(F^p(x), x^p) - x_i^p \sum_{j \in S^p} \rho_{ij}^p(F^p(x), x^p).$$

In Chapter 5, we introduced the following three desiderata for ρ and V .

- (C) *Continuity:* ρ^p is Lipschitz continuous.
- (NS) *Nash stationarity:* $V_F(x) = \mathbf{0}$ if and only if $x \in NE(F)$.
- (PC) *Positive correlation:* $V_F^p(x) \neq \mathbf{0}$ implies that $V_F^p(x)' F^p(x) > 0$.

We have seen that under continuity condition (C), any Lipschitz continuous population game F will generate a Lipschitz continuous differential equation (9.10), an equation that admits unique solutions from every initial condition in X . But a distinct consequence of condition (C)—one involving comparisons of dynamics across games—is equally important for the analyses to come.

Suppose we have a collection of population games $\{F^\varepsilon\}_{\varepsilon \in (-\bar{\varepsilon}, \bar{\varepsilon})}$ that have identical strategy sets and whose payoffs vary continuously in ε . Then under condition (C), the law of motion $\dot{x} = V_{F^\varepsilon}(x)$ varies continuously in ε . Moreover, if we let $\phi^\varepsilon : \mathbf{R}_+ \times X \rightarrow X$ denote the semiflow under V_{F^ε} , then the map $(\varepsilon, t, x) \mapsto \phi_t^\varepsilon(x)$ is continuous as well. This fact is important for understanding how evolution under $V_{(\cdot)}$ changes as we vary the underlying game. To capture the effects on long run behavior under $V_{(\cdot)}$, we must introduce the notion of an attractor. We keep the introduction here brief; additional details can be found in Appendix 9.B.

A set $\mathcal{A} \subseteq X$ is an *attractor* of the flow ϕ if it is nonempty, compact, and invariant under ϕ , and if there is a neighborhood U of \mathcal{A} such that

$$(9.11) \quad \limsup_{t \rightarrow \infty} \sup_{x \in U} \text{dist}(\phi_t(x), \mathcal{A}) = 0.$$

The set $B(\mathcal{A}) = \{x \in X : \omega(x) \subseteq \mathcal{A}\}$ is called the *basin* of \mathcal{A} . Put differently, attractors are asymptotically stable sets that are also invariant under the flow.

A key property of attractors for the current context is known as *continuation*. Fix an attractor $\mathcal{A} = \mathcal{A}^0$ of the flow ϕ^0 . Then as ε varies continuously from 0, there exist attractors \mathcal{A}^ε of the flows ϕ^ε that vary upper hemicontinuously from \mathcal{A} ; their basins $B(\mathcal{A}^\varepsilon)$ vary lower hemicontinuously from $B(\mathcal{A})$. Thus, if we slightly change the parameter ε , the attractors that exist under ϕ^0 continue to exist, and they do not explode.

Exercise 9.2.6. In defining an attractor via equation (9.11), we require that it attract solutions from all nearby states uniformly in time. To understand the role of uniformity in this definition, let ϕ be a flow on the unit circle that moves clockwise except at the topmost point x^* (cf Example 7.A.3). Explain why $\{x^*\}$ is not an attractor under this flow.


As a first application of these ideas, consider the 4×4 circulant game

$$(9.12) \quad F^\varepsilon(x) = A^\varepsilon x = \begin{pmatrix} 0 & 0 & -1 & \varepsilon \\ \varepsilon & 0 & 0 & -1 \\ -1 & \varepsilon & 0 & 0 \\ 0 & -1 & \varepsilon & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$


When $\varepsilon = 0$, the payoff matrix $A^\varepsilon = A^0$ is symmetric, so F^0 is a potential game with potential function $f(x) = \frac{1}{2}x'A^0x = -x_1x_3 - x_2x_4$. The function f attains its minimum of $-\frac{1}{4}$ at states $v = (\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $w = (0, \frac{1}{2}, 0, \frac{1}{2})$, has a saddle point with value $-\frac{1}{8}$ at the Nash equilibrium $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, and attains its maximum of 0 along the closed path of Nash equilibria γ consisting of edges $\overline{e_1e_2}$, $\overline{e_2e_3}$, $\overline{e_3e_4}$, and $\overline{e_4e_1}$. It follows from results in Section 7.1 that if $\dot{x} = V_{F^0}(x)$ satisfies (NS) and (PC), then all solutions whose initial conditions ξ satisfy $f(\xi) > -\frac{1}{8}$ converge to γ . (In fact, if x^* is a hyperbolic rest point of V_{F^ε} , then the Stable Manifold Theorem (see Appendix 8.C) tells us that the set of initial conditions from which solutions converge to x^* is a manifold of dimension no greater than 2, and hence has measure zero.) The phase diagram for the Smith dynamic game F^0 is presented in Figure 9.2.2(i).

Now suppose that $\varepsilon > 0$. If our revision protocol satisfies continuity (C), then the attractor γ of V_{F^0} continues to an attractor γ^ε of V_{F^ε} ; γ^ε is contained in a neighborhood of γ , and its basin approximates that of γ (see Figure 9.2.2(ii)). At the same time, the unique Nash equilibrium of F^ε is the central state x^* . We have therefore proved

Proposition 9.2.7. *Let $V_{(\cdot)}$ be an evolutionary dynamic that satisfies (C), (PC), and (NS), let F^ε be given by (9.12), and let $\delta > 0$. Then for $\varepsilon > 0$ sufficiently small, solutions to $\dot{x} = V_{F^\varepsilon}(x)$ from*



(i) $\varepsilon = 0$



(ii) $\varepsilon = \frac{1}{10}$

Figure 9.2.2: The Smith dynamic in game F^ε .




Figure 9.2.3: The replicator dynamic in Mismatching Pennies.

all initial conditions x with $f(x) > -\frac{1}{8} + \delta$ converge to an attractor γ^ε on which f exceeds $-\delta$; in particular, γ^ε contains neither Nash equilibria nor rest points.

9.2.3 Mismatching Pennies

Mismatching Pennies is a three-player normal form game in which each player has two strategies, Heads and Tails. Player p receives a payoff of 1 for choosing a different strategy than player $p + 1$ and a payoff of 0 otherwise, where players are indexed modulo 3.

If we let F be the population game generated by matching in Mismatching Pennies, then for each population $p \in \mathcal{P} = \{1, 2, 3\}$ we have that

$$F^p(x) = \begin{pmatrix} F_H^p(x) \\ F_T^p(x) \end{pmatrix} = \begin{pmatrix} x_T^{p+1} \\ x_H^{p+1} \end{pmatrix}.$$

The unique Nash equilibrium of F is the central state $x^* = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$. Since there are two strategies per player, it will simplify our analysis to let $y^p = x_H^p$ be the proportion of population p players choosing Heads, and to focus on the new state variable $y = (y^1, y^2, y^3) \in Y = [0, 1]^3$ (see Exercise 9.2.12 for details).

Example 9.2.8. The replicator dynamic for Mismatching Pennies. After our change of variable,

the replicator dynamic $\dot{y} = \hat{V}_F(y)$ for Mismatching Pennies takes the form

$$\dot{y} = \begin{pmatrix} \dot{y}^1 \\ \dot{y}^2 \\ \dot{y}^3 \end{pmatrix} = \begin{pmatrix} y^1(1-y^1)(1-2y^2) \\ y^2(1-y^2)(1-2y^3) \\ y^3(1-y^3)(1-2y^1) \end{pmatrix}.$$

The derivative matrix for an arbitrary state y and the equilibrium state $y^* = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ are

$$D\hat{V}(y) = \begin{pmatrix} (1-2y^1)(1-2y^2) & -2y^1(1-y^1) & 0 \\ 0 & (1-2y^2)(1-2y^3) & -2y^2(1-y^2) \\ -2y^3(1-y^3) & 0 & (1-2y^3)(1-2y^1) \end{pmatrix} \quad \text{and} \quad D\hat{V}(y^*) = \begin{pmatrix} 0 & -\frac{1}{2} & 0 \\ 0 & 0 & -\frac{1}{2} \\ -\frac{1}{2} & 0 & 0 \end{pmatrix}.$$

$D\hat{V}(y^*)$ is a circulant matrix with an eigenvalue of $-\frac{1}{2}$ corresponding to eigenvector $\mathbf{1}$, and eigenvalues of $\frac{1}{4} \pm \frac{\sqrt{3}}{4}i$ corresponding to eigenvectors $(-1, -1, 2)' \pm i(-\sqrt{3}, \sqrt{3}, 0)'$; note that $\mathbf{1}$, $(-1, -1, 2)'$, and $(-\sqrt{3}, \sqrt{3}, 0)'$ are mutually orthogonal. The phase diagram for the replicator dynamic is a *spiral saddle*: interior solutions on the diagonal where $y^1 = y^2 = y^3$ head directly toward y^* , while all other orbits are attracted to a two-dimensional manifold containing an unstable spiral. This is depicted in Figure 9.2.3, where behavior in populations 1, 2, and 3 is measured on the left-right, front-back, and top-bottom axes, respectively. Solutions on the manifold containing the unstable spiral converge to a six-segment heteroclinic cycle; this cycle agrees with the best response cycle of the underlying normal form game. §

Example 9.2.9. The best response dynamic in Mismatching Pennies. The analysis of the best response dynamic in Mismatching Pennies is very similar to the corresponding analysis in monocyclic games (Example 9.2.5). Divide the state space $Y = [0, 1]^3$ into eight octants in the natural way. Then the two octants corresponding to vertices HHH and TTT are backward invariant, while solutions starting in any of the remaining six octants proceed through those octants according to the best response cycle of the underlying game (see Exercise 9.2.10). As Figure 9.2.4 illustrates, almost all solutions to the best response dynamic converge to a six-sided closed orbit in the interior of Y . §

- Exercise 9.2.10.* (i) Give an explicit formula for the best response dynamic for Mismatching Pennies in terms of the state variable $y \in Y = [0, 1]^3$.
(ii) Prove that octants HHH and TTT described in the previous example are backward invariant.
(iii) Prove that solutions starting in any of the remaining octants proceed through those octants according to the best response cycle of the underlying game.

The following proposition shows that the previous two examples are not exceptional.

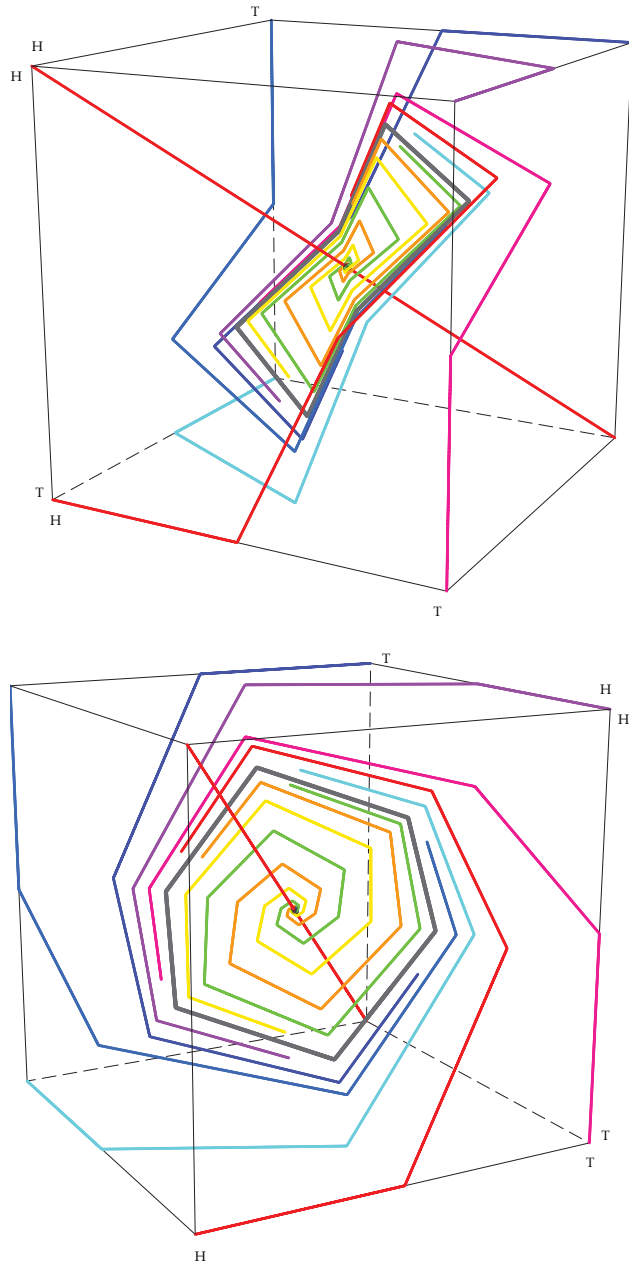


Figure 9.2.4: The best response dynamic in Mismatching Pennies (two viewpoints).

Proposition 9.2.11. *Let $V_{(\cdot)}$ be an evolutionary dynamic that is generated by a C^1 revision protocol ρ and that satisfies Nash stationarity (NS). Let F be Mismatching Pennies, and suppose that the unique Nash equilibrium x^* of F is a hyperbolic rest point of $\dot{x} = V_F(x)$. Then x^* is unstable under V_F , and there is an open, dense, full measure set of initial conditions from which solutions to V_F do not converge.*

Proposition 9.2.11 is remarkable in that it does not require the dynamic to satisfy a payoff monotonicity condition. Instead, it takes advantage of the fact that by definition, the revision protocol for population p does not condition on the payoffs of other populations. In fact, the specific payoffs of Mismatching Pennies are not important to obtain the instability result; any three-player game whose unique Nash equilibrium is interior works equally well. The proof of the theorem makes these points clear.

Proof. For ε close to 0, let F^ε be generated by a perturbed version of Mismatching Pennies in which player 3's payoff for playing H when player 1 plays T is not 1, but $\frac{1+2\varepsilon}{1-2\varepsilon}$. Then like Mismatching Pennies itself, F^ε has a unique Nash equilibrium, here given by $((\frac{1}{2} + \varepsilon, \frac{1}{2} - \varepsilon), (\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$.

For convenience, let us argue in terms of the state variable $y = (x_H^1, x_H^2, x_H^3) \in Y = [0, 1]^3$ (see Exercise 9.2.12). If $\dot{y} = \hat{V}_{F^\varepsilon}(y)$ is the dynamic $\dot{x} = V_{F^\varepsilon}(x)$ expressed in terms of y , then Nash stationarity (NS) tells us that

$$(9.13) \quad \hat{V}_{F^\varepsilon}(\frac{1}{2} + \varepsilon, \frac{1}{2}, \frac{1}{2}) = \mathbf{0}$$

whenever $|\varepsilon|$ is small. Now by definition, the law of motion for population 1 does not depend directly on payoffs in the other populations, regardless of the game at hand (cf equation (9.10)). Therefore, since changing the game from F^ε to F^0 does not alter population 1's payoff function, equation (9.13) implies that

$$\hat{V}_{F^0}^1(\frac{1}{2} + \varepsilon, \frac{1}{2}, \frac{1}{2}) = 0$$

whenever $|\varepsilon|$ is small. This observation and the fact that the dynamic is differentiable at $y^* = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ imply that

$$\frac{\partial \hat{V}_{F^0}^1}{\partial y^1}(y^*) = 0.$$

Repeating this argument for the other populations shows that the trace of $D\hat{V}_{F^0}(y^*)$, and hence the sum of the eigenvalues of $D\hat{V}_{F^0}(y^*)$, is 0. Since y^* is a hyperbolic rest point of \hat{V}_{F^0} , it follows that some eigenvalue of $D\hat{V}_{F^0}(y^*)$ has positive real part, and thus that y^*

is unstable under \hat{V}_{F0} . Thus, the Stable Manifold Theorem (see Appendix 8.C) tells us that the set of initial conditions from which solutions converge to y^* is of dimension at most 2, and that its complement is open, dense, and of full measure in Y . ■

Exercise 9.2.12. Let X be the state space for a p population game with two strategies per population, and let $Y = [0, 1]^p$, so that $TY = \mathbf{R}^p$.

- (i) Show that the change of variable $h : X \rightarrow Y$ has inverse $h^{-1} : Y \rightarrow X$, where

$$h(x) = \begin{pmatrix} x_1^1 \\ \vdots \\ x_1^p \end{pmatrix} \quad \text{and} \quad h^{-1}(y) = \begin{pmatrix} y^1 \\ 1-y^1 \\ \vdots \\ y^p \\ 1-y^p \end{pmatrix}.$$

- (ii) Show that the derivative of h at x , $Dh(x) : TX \rightarrow TY$, and the derivative of h^{-1} at y , $Dh^{-1}(y) : TY \rightarrow TX$, can be written as $Dh(x)z = Mz$ and $Dh^{-1}(y)\zeta = \tilde{M}\zeta$ for some matrices $M \in \mathbf{R}^{p \times 2p}$ and $\tilde{M} \in \mathbf{R}^{2p \times p}$. Show that if M is viewed as a linear map from TX to TY , then its inverse is \tilde{M} .
- (iii) Fix a C^1 vector field $V : X \rightarrow TX$, and define the new vector field $\hat{V} : Y \rightarrow TY$ by $\hat{V}(y) = h(V(h^{-1}(y)))$. Show that the dynamics $\dot{x} = V(x)$ and $\dot{y} = \hat{V}(y)$ are linearly conjugate under H : that is, that $\{x_t\}$ solves the former equation if and only if $\{h(x_t)\}$ solves the latter.
- (iv) Let x^* be a rest point of V , and let $y^* = h(x^*)$ be the corresponding rest point of \hat{V} . Show that the eigenvalues of $DV(x^*)$ with respect to TX are identical to the eigenvalues of $D\hat{V}(y^*)$ with respect to TY . What is the relationship between the corresponding pairs of eigenvectors?

9.2.4 The Hypnodisk Game

The virtue of Proposition 9.2.11 is that apart from hyperbolicity of equilibrium, virtually no assumptions about the evolutionary dynamic $V_{(\cdot)}$ were needed to establish nonconvergence. We now show that if one is willing to introduce a payoff monotonicity condition—namely, positive correlation (PC)—then one can obtain a nonconvergence without smoothness conditions, and using a two-dimensional state variable, rather than a three-dimensional one as in Mismatching Pennies. This low dimensionality will turn out to be crucial when we study survival of dominated strategies in Section 9.4.

Our construction will be based on potential games. In Figure 9.2.5, we present the

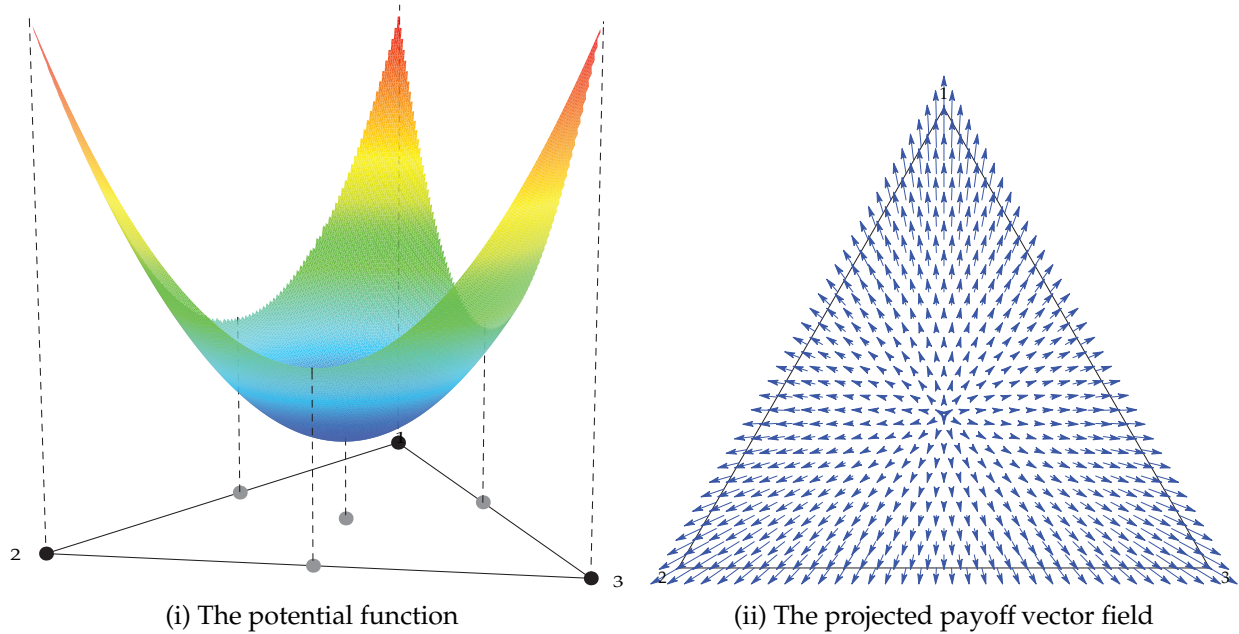


Figure 9.2.5: A coordination game.

potential function and projected payoff vector field of the coordination game

$$F^C(x) = Cx = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}.$$

By our analysis in Chapter 3, solutions to any evolutionary dynamic $\dot{x} = V_{F^C}(x)$ satisfying conditions (NS) and (PC) ascend the potential function $f^C(x) = \frac{1}{2}x'Cx = \frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3)^2)$ drawn in diagram (i), or, equivalently, travel at acute angles to the projected payoff vectors in diagram (ii). It follows that solutions to V_{F^C} from most initial conditions converge to the strict Nash equilibria at the vertices of X .

As a second example, suppose that agents are matched to play the antcoordination game $-C$. In Figure 9.2.6, we draw the resulting population game $F^{-C}(x) = -Cx = -x$ and its concave potential function $f^{-C}(x) = -\frac{1}{2}x'Cx = -\frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3)^2)$. Both pictures reveal that under any evolutionary dynamic satisfying conditions (NS) and (PC), all solution trajectories converge to the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

The construction of the hypnodisk game $H : X \rightarrow \mathbf{R}^3$ is easiest to describe in geometric terms. Begin with the coordination game $F^C(x) = Cx$ pictured in Figure 9.2.5(ii). Then draw two circles centered at state $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ with radii $0 < r < R < \frac{1}{\sqrt{6}}$, as shown in Figure 9.2.7(i); the second inequality ensures that both circles are contained in the simplex. Twist the portion of the vector field lying outside of the inner circle in a clockwise direction,

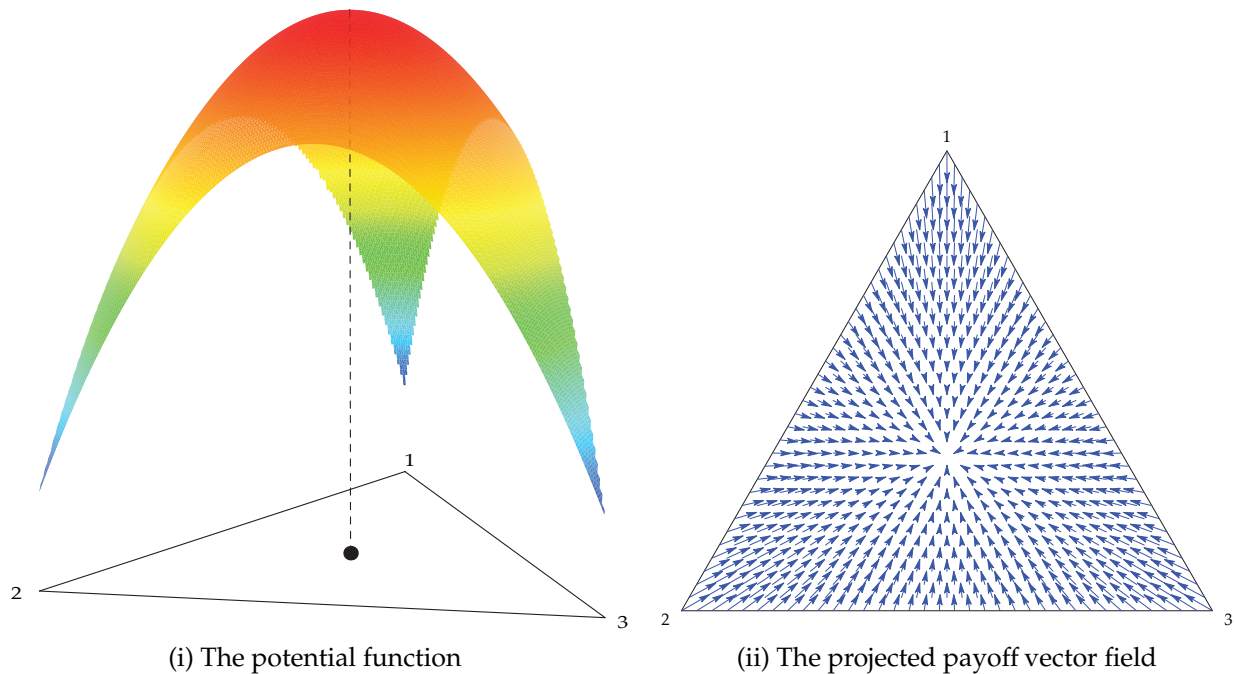


Figure 9.2.6: An anticoordination game.


excluding larger and larger circles as the twisting proceeds, so that the outer circle is reached when the total twist is 180° (Figure 9.2.7(ii)).

Exercise 9.2.13. Provide an explicit formula for the resulting population game $H(x)$.


What does this construction accomplish? Examining Figure 9.2.7(ii), we see that inside the inner circle, H is identical to the coordination game F^C . Thus, solutions to dynamics satisfying (NS) and (PC) starting at states other than x^* in the inner circle must leave the inner circle. At states outside the outer circle, H is identical to the anticoordination game F^{-C} , so solutions to dynamics satisfying (NS) and (PC) starting at states outside the outer circle must enter the outer circle. Finally, at each state x in the annulus bounded by the two circles, $H(x)$ is not a componentwise constant vector. Therefore, states in the annulus are not Nash equilibria, and so are not rest points of dynamics satisfying (NS). We assemble these observations in the following proposition.

Proposition 9.2.14. *Let $V_{(\cdot)}$ be an evolutionary dynamic that satisfies (C), (NS), and (PC), and let H be the hypnodisk game. Then every solution to $\dot{x} = V_H(x)$ other than the stationary solution at x^* enters the annulus with radii r and R and never leaves, ultimately converging to a cycle therein.*

The claim of convergence to limit cycles in the final sentence of the proposition follows from the Poincaré-Bendixson Theorem (Theorem 9.A.5).



(i) Projected payoff vector field for the coordination game



(ii) Projected payoff vector field for the hypnodisk game

Figure 9.2.7: Construction of the hypnodisk game.

9.3 Chaotic Evolutionary Dynamics

In all of the phase diagrams we have seen so far, ω -limit sets have taken a fairly simple form: solution trajectories have converged to rest points, closed orbits, or chains of rest points and connecting orbits. When we consider games with just two or three strategies, this is unavoidable: clearly, all solution trajectories of continuous time dynamics in one dimension converge to equilibrium, while in two-dimensional systems, the Poincaré-Bendixson Theorem (Theorem 9.A.5) tells us that the three types of ω -limit sets described above exhaust all possibilities.

Once we move to flows in three or more dimensions, ω -limit sets can be much more complicated sets often referred to as *chaotic* (or *strange*) *attractors*. Central to most definitions of chaos is *sensitive dependence on initial conditions*: solution trajectories starting from close together points on the attractor move apart at an exponential rate. Chaotic attractors can also be recognized in phase diagrams by their rather intricate appearance. Rather than delving deeply into these ideas, we content ourselves by presenting a few examples.

Example 9.3.1. Consider the single population game F generated by matching in the normal form game A below:

$$F(x) = Ax = \begin{pmatrix} 0 & -12 & 0 & 22 \\ 20 & 0 & 0 & -10 \\ -21 & -4 & 0 & 35 \\ 10 & -2 & 2 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

The lone interior Nash equilibrium of this game is the central state $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$.

Let $\dot{x} = V_F(x)$ be the replicator dynamic for game F . One can calculate that the eigenvalues of $DV_F(x^*)$ are approximately -3.18 and $.34 \pm 1.98i$, so like the Nash equilibrium of Mismatching Pennies (Example 9.2.8), the interior equilibrium x^* here is a spiral saddle with an unstable spiral.

Figure 9.3.1 presents the initial portion of the solution of $\dot{x} = V_F(x)$ from initial condition $x_0 = (.24, .26, .25, .25)$. This solution spirals clockwise about x^* . Near the rightmost point of each circuit, where the value of x_3 gets close to zero, solutions sometimes proceed along an “outside” path on which the value of x_3 surpasses .6. But they sometimes follow an “inside” path on which x_3 remains below .4, and at other times do something in between. Which of these alternatives occurs is difficult to predict from approximate information about the previous behavior of the system.

Sensitive dependence on initial conditions is illustrated directly in Figure 9.3.2, which tracks the solutions from two nearby initial conditions, $(.47, .31, .11, .11)$ and $(.46999, .31,$




Figure 9.3.1: A chaotic attractor under the replicator dynamic.

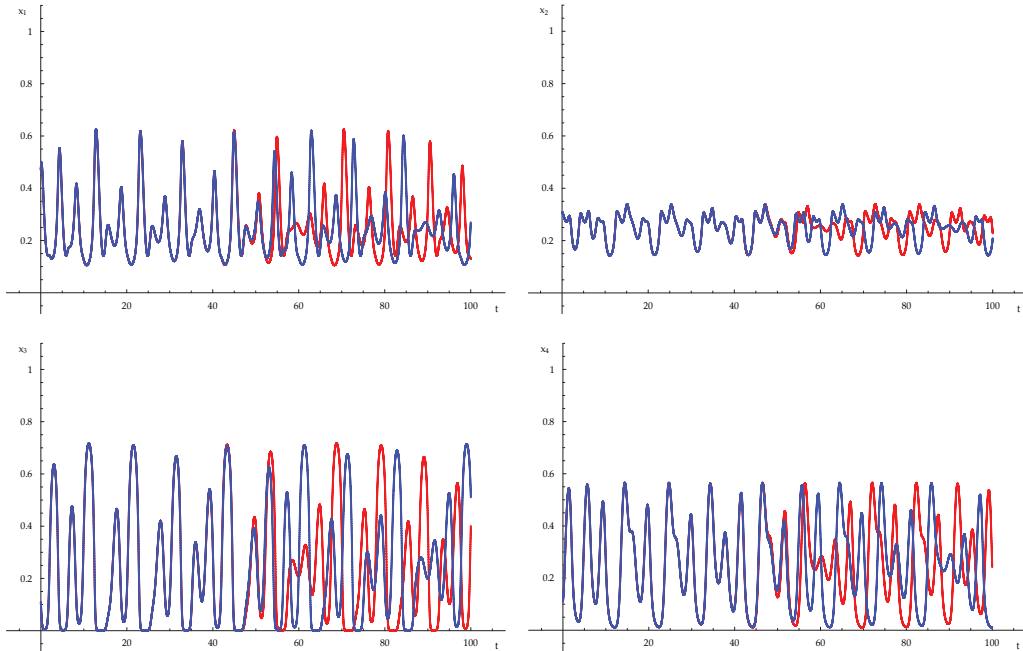


Figure 9.3.2: Sensitive dependence on initial conditions under the replicator dynamic.

.11, .11001). Apparently, the two solutions stay close together through time $t = 50$ but diverge thereafter; after time $t = 60$, the current position of one of the solutions provides little hint about the current position of the other. §

The scattered payoff entries in the previous example may seem to suggest that chaos only occurs in “artificial” examples. To dispute this view, we now show that chaotic behavior can occur in very simple games.

Example 9.3.2. Asymmetric Rock-Paper-Scissors. Suppose that two populations of agents are matched to play the two-player zero-sum game $U = (U^1, U^2)$:

		II		
		r	p	s
I	R	$\frac{1}{2}, -\frac{1}{2}$	$-1, 1$	$1, -1$
	P	$1, -1$	$\frac{1}{2}, -\frac{1}{2}$	$-1, 1$
	S	$-1, 1$	$1, -1$	$\frac{1}{2}, -\frac{1}{2}$

U is an asymmetric version of Rock-Paper-Scissors in which a “draw” results in a half-credit win for player 1.

Figures 9.3.3 and 9.3.4 each present a single solution trajectory of the replicator dynamic for F_U . Since the social state $x = (x^1, x^2)$ is four-dimensional, we draw it in two pieces,

with x^1 represented on the left hand side of each figure and x^2 represented on the right. Because U is a zero-sum game with Nash equilibrium $((\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}))$, each solution of the replicator dynamic lives within a level set of $\mathcal{H}(x) = -\sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{S}^p} \log x_i^p$. In Figure 9.3.3, whose initial condition is $((.5, .25, .25), (.5, .25, .25))$, the solution trajectory appears to follow a periodic orbit, much like those in our examples from Section 9.1.1. But in Figure 9.3.4, whose initial condition $((.5, .01, .49), (.5, .25, .25))$ is closer to the boundary of X , the solution trajectory travels around the level set of \mathcal{H} in a seemingly haphazard way. Thus, despite the regularity provided by the constant of motion, the evolution of behavior in this simple game is complicated indeed. §

9.4 Survival of Dominated Strategies

By now we have considered in detail the extent to which predictions of Nash equilibrium play are supported by the behavior of evolutionary dynamics. On the positive side, Chapters 5 and 6 show that there are many dynamics whose rest points are always identical to the Nash equilibria of the underlying game, and Chapter 7 shows that convergence to Nash equilibrium can be assured under many of these dynamics in particular classes of games. But the examples in Section 9.2 reveal that no general guarantee of convergence to Nash equilibrium is possible, as there are games in which no reasonable dynamic leads to Nash equilibrium play.

This negative result leads us to consider a more modest question. Rather than seek evolutionary support for equilibrium play, we instead turn our attention to a more basic rationality requirement: namely, the avoidance of strategies that are strictly dominated. Section 7.4 provided some positive results on this question, showing that the best response dynamic and all imitative dynamics eliminate dominated strategies, at least along solutions starting from most initial conditions.

Since the elimination of strategies strictly dominated by a pure strategy is the mildest requirement employed in standard game-theoretic analyses, the results from Section 7.4 may seem unremarkable. But some reflection reveals that the case for elimination of dominated strategies by evolutionary dynamics is rather tenuous. Evolutionary dynamics describe the aggregate behavior of agents who employ simple revision protocols, switching to strategies whose current payoffs are good, though not necessarily optimal. In some cases—in particular, when solutions converge—these simple rules are enough to ensure individually optimal long run behavior. When a solution trajectory of an evolutionary dynamic converges, the payoffs to each strategy converge as well; because payoffs become fixed, even simple rules are enough to ensure that only optimal strategies are chosen.

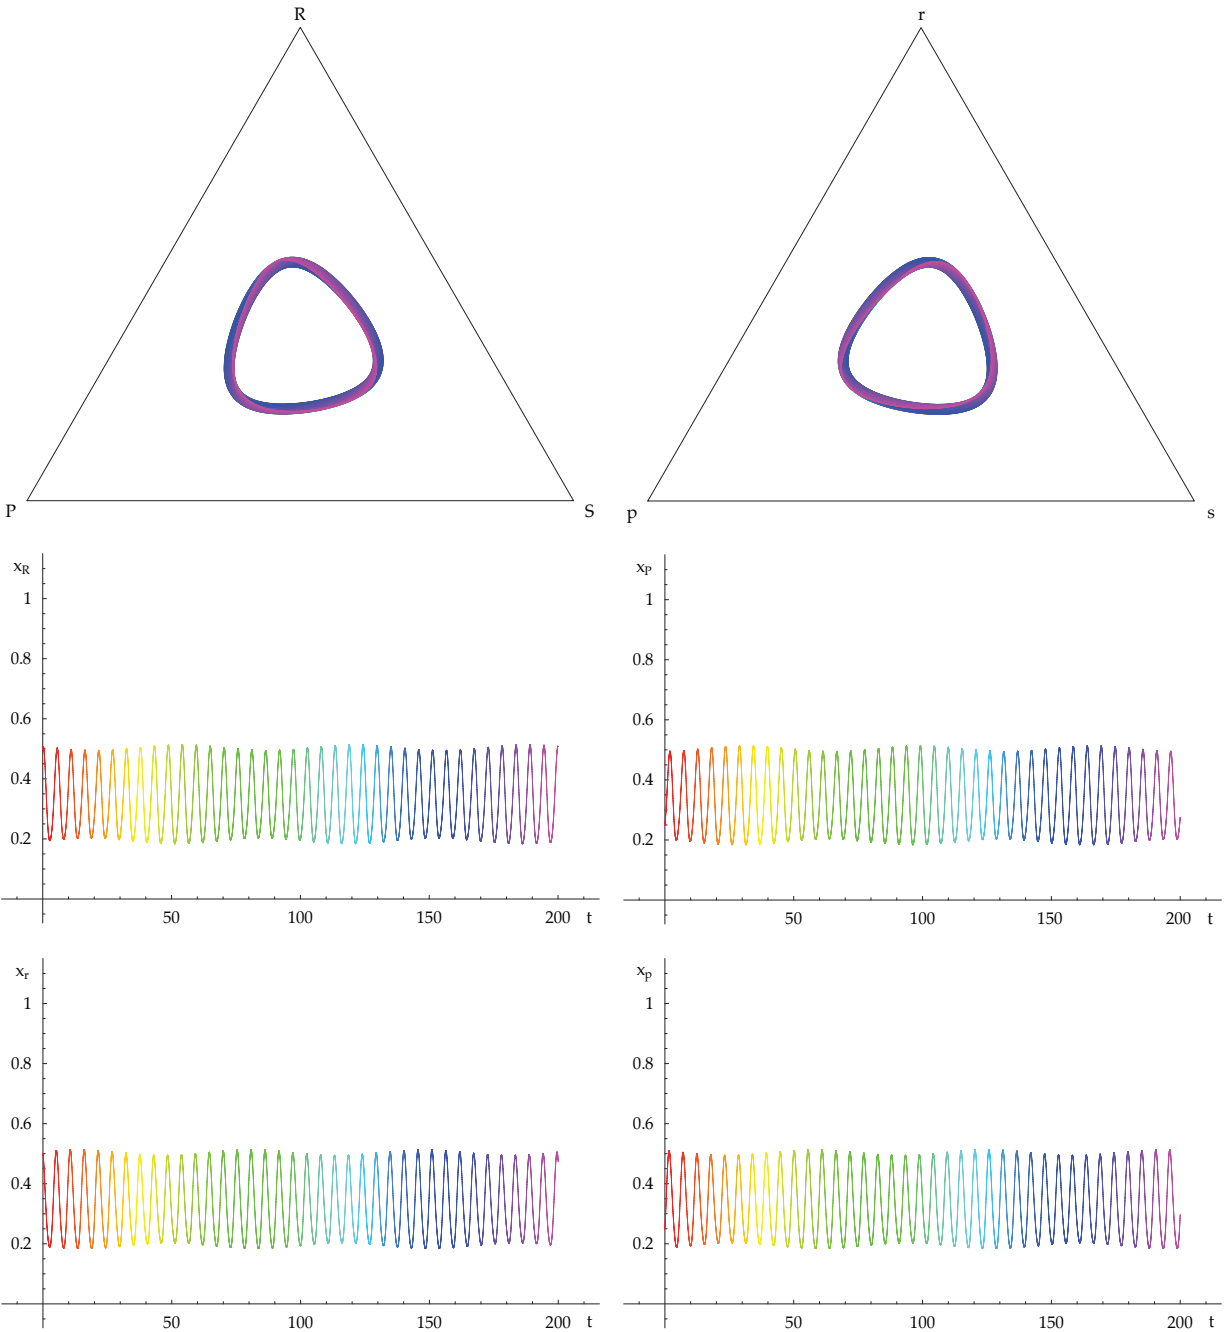


Figure 9.3.3: Cycling in asymmetric Rock-Paper-Scissors.




Figure 9.3.4: Chaos in asymmetric Rock-Paper-Scissors.

In formal terms, the limits of convergent solution trajectories must be Nash equilibria, implying that strictly dominated strategies are not chosen.

We have argued in this chapter that solutions of evolutionary dynamics need not converge, but instead may enter limit cycles or more complicated limit sets. When solutions do not converge, payoffs remain in flux. In this situation, it is not obvious whether choice rules favoring strategies whose current payoffs are relatively high will necessarily eliminate strategies that perform well at many states, but that are never optimal. To the contrary, the analysis to follow demonstrates that if play remains in disequilibrium, even strategies that are strictly dominated by other pure strategies can persist indefinitely.

9.4.1 A General Survival Theorem

To turn this intuition into a formal result, we introduce one further condition on evolutionary dynamics.

(IN) *Innovation* If $x \notin NE(F)$, $x_i = 0$, and $i \in \operatorname{argmax}_{j \in S} F_j(x)$, then $(V_F)_i(x) > 0$.

Innovation (IN) requires that when a non-Nash population state includes an unused optimal strategy, this strategy's growth rate must be positive. In other words, if an unplayed strategy is sufficiently rewarding, some members of the population will discover it and select it. This condition excludes dynamics based purely on imitation; however, it includes dynamics that combine imitation with a small amount of direct selection of candidate strategies—see Section 5.7, as well as Example 9.4.8 below.

We are now prepared to state our survival theorem.

Theorem 9.4.1. *Suppose the evolutionary dynamic $V_{(\cdot)}$ satisfies (C), (NS), (PC), and (IN). Then there is a game F such that under $\dot{x} = V_F(x)$, along solutions from most initial conditions, there is a strictly dominated strategy played by a fraction of the population bounded away from 0.*

Theorem 9.4.1 states that evolutionary dynamics satisfying four natural requirements fail to eliminate dominated strategies in some games. In essence, the theorem shows that the dynamics studied in Section 7.4 are essentially the only ones that should be expected to eliminate dominated strategies in all games, and that even the elimination results for these dynamics are knife-edge cases.

The proof of Theorem 9.4.1 makes use of the hypnodisk game from Section 9.2.4. While introducing this unusual game simplifies the analysis, the survival of dominated strategies is a broader consequence of nonconvergence, and does not depend on the introduction of complicated game in any essential way. See Section 9.4.2 for further discussion.

Proof of Theorem 9.4.1. Let H be the hypnodisk game introduced in Section 9.2.4. Let F be the four-strategy game obtained from H by adding a twin to strategy 3:

$$F_i(x_1, x_2, x_3, x_4) = H_i(x_1, x_2, x_3 + x_4) \quad \text{for } i \in \{1, 2, 3\};$$

$$F_4(x) = F_3(x).$$

Strategies 3 and 4 are identical, in that they always yield the same payoff and always have the same payoff consequences for other strategies. The set of Nash equilibria of F is the line segment

$$NE = \left\{ x^* \in X : x_1^* = x_2^* = x_3^* + x_4^* = \frac{1}{3} \right\}.$$

Let

$$I = \left\{ x \in X : (x_1 - \frac{1}{3})^2 + (x_2 - \frac{1}{3})^2 + (x_3 + x_4 - \frac{1}{3})^2 \leq r^2 \right\} \text{ and}$$

$$O = \left\{ x \in X : (x_1 - \frac{1}{3})^2 + (x_2 - \frac{1}{3})^2 + (x_3 + x_4 - \frac{1}{3})^2 \leq R^2 \right\}$$

be concentric cylindrical regions in X surrounding NE , as pictured in Figure 9.4.1. By construction, we have that

$$F(x) = \tilde{C}x = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

at all $x \in I$, so under any dynamic satisfying (PC) and (NS), solutions starting in $I - NE$ ascend the potential function $f^{\tilde{C}}(x) = \frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3 + x_4)^2)$ until leaving the set I . At states outside the set O , we have that $F(x) = -\tilde{C}x$, so solutions starting in $X - O$ ascend $f^{-\tilde{C}}(x) = -f^{\tilde{C}}(x)$ until entering O . In summary:

Lemma 9.4.2. *Suppose that $V_{(\cdot)}$ is an evolutionary dynamic that satisfies conditions (C), (NC) and (PC), and let F be the “hypnodisk with a twin” game. Then every solution to $\dot{x} = V_F(x)$ other than the stationary solutions in NE enter region $D = O - I$ and never leave.*

Define the flow from the set $U \subseteq X$ under the dynamic V_F by

$$\phi_t(U) = \{ \xi \in X : \text{there is a solution } \{x_s\} \text{ to } \dot{x} = V_F(x) \text{ with } x_0 \in U \text{ and } x_t = \xi. \}$$




Figure 9.4.1: Regions O , I , and $D = O - I$.

In words, $\phi_t(U)$ contains the time t positions of solutions to V_F whose initial conditions are in U .

Since solutions to V_F starting in $I - NE$ ascend the function $f^{\tilde{C}}$ until leaving the set I , the reverse time flow is well-defined from all such states, and NE is a *repellor* under V_F . This means that all backward-time solutions to V_F that begin in some neighborhood U of NE converge to NE uniformly over time, or, equivalently, that NE is an attractor of the time-reversed equation $\dot{y} = -V_F(y)$ (see Appendix 9.B). The *dual attractor* \mathcal{A} of the repellor NE is the forward-time limit of the flow of V_F starting from the complement of $\text{cl}(U)$:

$$\mathcal{A} = \bigcap_{t \geq 0} \phi_t(X - \text{cl}(U)).$$

\mathcal{A} is nonempty, compact, and (both forward and backward) invariant under V_F , and Lemma 9.4.2 tells us that $\mathcal{A} \subset D$.

We now show that the twin strategy is used by a positive mass of agents throughout the attractor \mathcal{A} . Let $Z = \{x \in X : x_4 = 0\}$ be the face of X on which the twin strategy is unused; we prove

Lemma 9.4.3. *The attractor \mathcal{A} and the face Z are disjoint.*

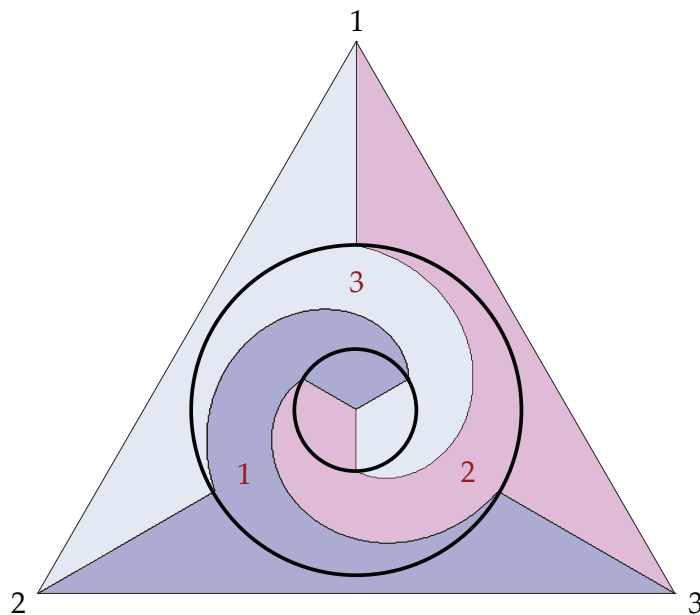


Figure 9.4.2: The best response correspondence of the hypnodisk game.

Proof. Since V_F is Lipschitz continuous and satisfies $(V_F)_i(x) \geq 0$ whenever $x_i = 0$, solutions to V_F that start in $X - Z$ cannot approach Z more than exponentially quickly, and in particular cannot reach Z in finite time (see Exercise 9.4.4). Equivalently, backward solutions to V_F starting from states in Z cannot enter $\text{int}(X)$.

Now suppose by way of contradiction that there exists a state ξ in $\mathcal{A} \cap Z$. Then by our previous arguments, the entire backward orbit from ξ is also contained in $\mathcal{A} \cap Z$, and hence in $D \cap Z$. Since the latter set contains no rest points by condition (PC), the Poincaré-Bendixson Theorem (Theorem 9.A.5) implies that the backward orbit from ξ converges to a closed orbit γ in $D \cap Z$ that circumnavigates $I \cap Z$.

By construction, the annulus $D \cap Z$ can be split into three regions: one in which strategy 1 is the best response, one in which strategy 2 is the best response, and one in which strategy 3 (and hence strategy 4) is a best response (Figure 9.4.2). Each of these regions is bounded by a simple closed curve that intersects the inner and outer boundaries of the annulus. Therefore, the closed orbit γ , on which strategy 4 is unused, passes through the region in which strategy 4 is optimal. This contradicts innovation (IN). ■

Exercise 9.4.4. Use Grönwall's Inequality (Lemma 4.A.7) to check the initial claim in the proof of the lemma.




Figure 9.4.3: The Smith dynamic in bad RPS.

To complete the proof, we now make the twin strategy “feeble”: we uniformly reduce its payoff by $d > 0$, creating the new game

$$F^d(x) = F(x) - de_4.$$

Observe that strategy 4 is strictly dominated by strategy 3 in game F^d .

As increasing d from 0 continuously changes the game from F to F^d , doing so also continuously changes the dynamic from V_F to V_{F^d} . Thus, by Theorem 9.B.5 on continuation of attractors, we have that for small d , the attractor \mathcal{A} of V_F continues to an attractor \mathcal{A}^d of V_{F^d} on which $x_4 > 0$: thus, the dominated strategy survives throughout \mathcal{A}^d . The basin of the attractor \mathcal{A}^d contains all points outside of a thin tube around the set NE of Nash equilibria of F . This completes the proof of Theorem 9.4.1. ■

9.4.2 Examples and Discussion

We conclude this chapter with some examples and exercises that demonstrate the robustness of the analysis above.

Example 9.4.5. We use the hypnodisk game as the basis for the proof of Theorem 9.4.1 because it generates cycling under any dynamic that satisfies (NS) and (PC). But the use

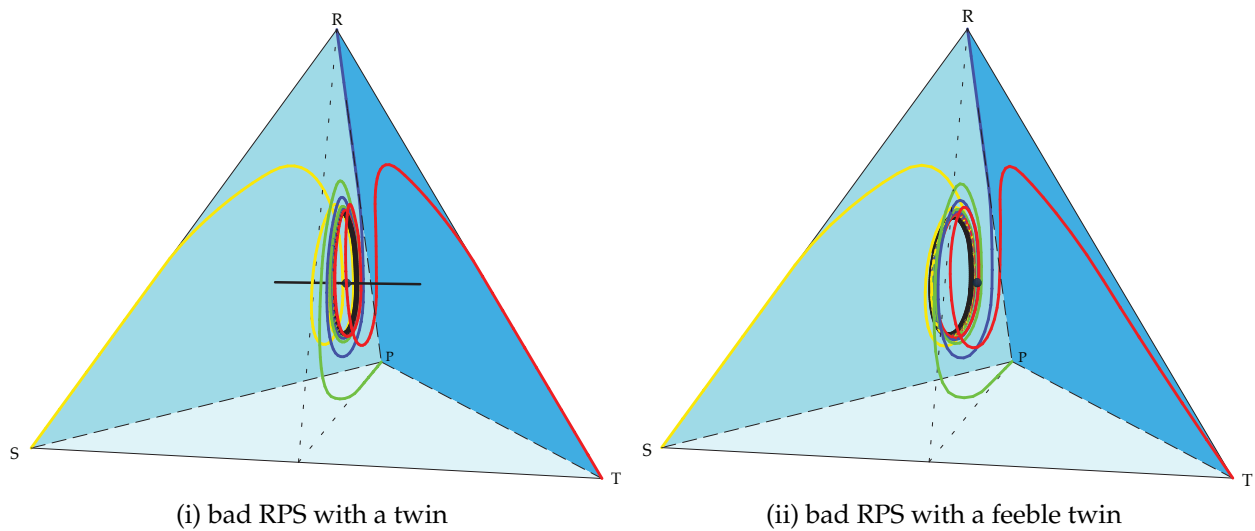


Figure 9.4.4: The Smith dynamic in two games.

of this game is not essential: once we fix the dynamic under consideration, we can find a simpler game that leads to cycling; then the argument based on the introduction of twin strategies can proceed as above.

We illustrate this point by constructing an example of survival under the Smith dynamic. Figure 9.4.3 contains the phase diagram for the Smith dynamic in the bad Rock-Paper-Scissors game

$$G(x) = Ax = \begin{pmatrix} 0 & -2 & 1 \\ 1 & 0 & -2 \\ -2 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}.$$

Evidently, the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is unstable, and most solution trajectories converge to a cycle located in $\text{int}(X)$.

Figure 9.4.4(i) presents the Smith dynamic in “bad RPS with a twin”,

$$(9.14) \quad F(x) = \tilde{A}x = \begin{pmatrix} 0 & -2 & 1 & 1 \\ 1 & 0 & -2 & -2 \\ -2 & 1 & 0 & 0 \\ -2 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

The Nash equilibria of F are the states on line segment $NE = \{x^* \in X : x^* = (\frac{1}{3}, \frac{1}{3}, c, \frac{1}{3} - c)\}$, which is a repeller under the Smith dynamic. Furthermore, since Scissors and Twin always earn the same payoffs ($F_3(x) \equiv F_4(x)$), we can derive a simple expression for the rate of

change of the difference between their utilization levels:

$$\begin{aligned}
 (9.15) \quad \dot{x}_3 - \dot{x}_4 &= \left(\sum_{j \in S} x_j [F_3(x) - F_j(x)]_+ - x_3 \sum_{j \in S} [F_j(x) - F_3(x)]_+ \right) \\
 &\quad - \left(\sum_{j \in S} x_j [F_4(x) - F_j(x)]_+ - x_4 \sum_{j \in S} [F_j(x) - F_4(x)]_+ \right) \\
 &= -(x_3 - x_4) \sum_{j \in S} [F_j(x) - F_4(x)]_+.
 \end{aligned}$$

Intuitively, strategies lose agents at rates proportional to their current levels of use, but gain strategies at rates that depend on their payoffs; thus, when the dynamics are not at rest, the weights x_3 and x_4 move closer together. It follows that except at Nash equilibrium states, the dynamic moves toward the plane $P = \{x \in X : x_3 = x_4\}$ on which the identical twins receive equal weight (see Exercise 9.4.6).

Figure 9.4.4(ii) presents the Smith dynamic in “bad RPS with a feeble twin”,

$$(9.16) \quad F^d(x) = \tilde{A}^d x = \begin{pmatrix} 0 & -2 & 1 & 1 \\ 1 & 0 & -2 & -2 \\ -2 & 1 & 0 & 0 \\ -2-d & 1-d & -d & -d \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix},$$

where $d = \frac{1}{10}$. Evidently, the attractor from the previous figure moves slightly to the left, and the strictly dominated strategy Twin survives. Indeed, since the Nash equilibrium of “RPS with a twin” on plane P puts mass $\frac{1}{6}$ on Twin, when d is small solutions to the Smith dynamic in “RPS with a feeble twin” place mass greater than $\frac{1}{6}$ on the strictly dominated strategy Twin infinitely often. This lower bound is driven by the fact that in the game with an exact twin, solutions converge to plane P ; thus, the bound will obtain under any dynamic that treats different strategies symmetrically. §

Exercise 9.4.6. Show that under the Smith dynamic in “RPS with a twin”, solutions from states not on the line of Nash equilibria NE converge to the plane P where the weights on Scissors and Twin are equalized. (Hint: Use equation (9.15) and the Poincaré-Bendixson Theorem. You may take as given that the set NE is a repeller.)

Example 9.4.7. The analyses of survival of dominated strategies presented above do not quantify the level of domination that is consistent with a dominated strategy maintaining a significant presence in the population. We can provide a sense of this magnitude through a numerical analysis.

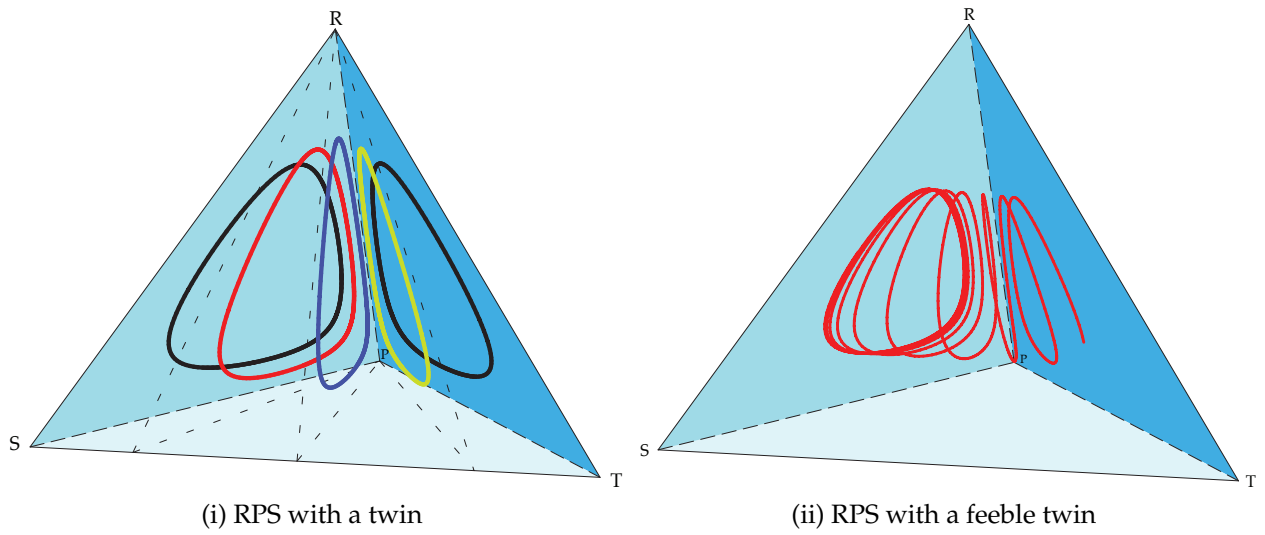


Figure 9.4.5: The replicator dynamic in two games.

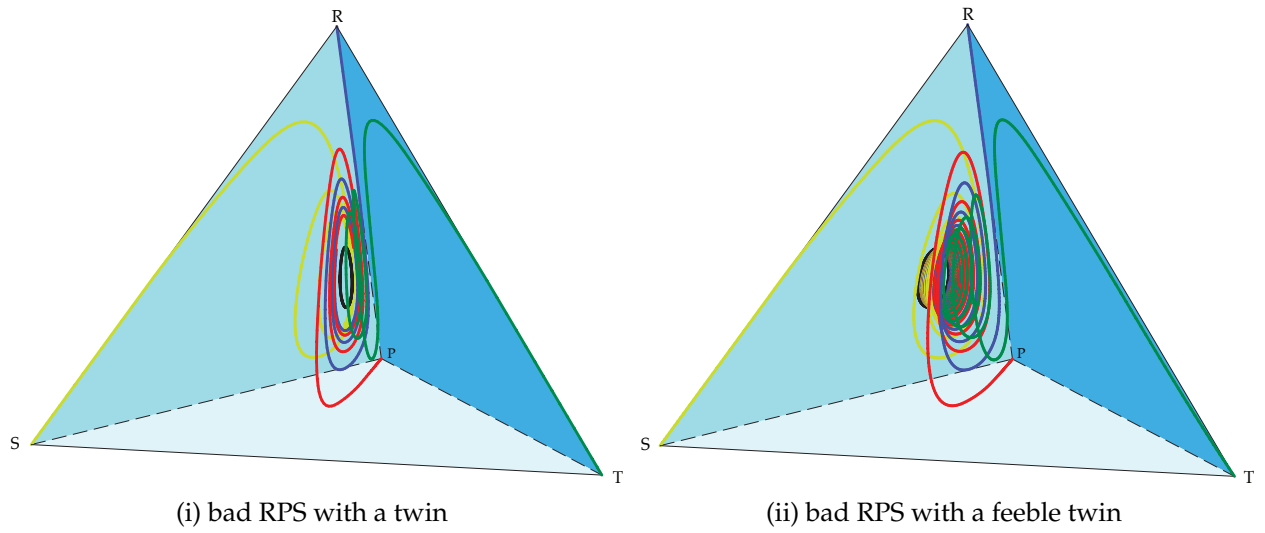


Figure 9.4.6: The $\frac{9}{10}$ replicator + $\frac{1}{10}$ Smith dynamic in two games.




Figure 9.4.7: The maximum, time-average, and minimum weight on the dominated strategy under the Smith dynamic as functions of the domination level d .

We consider the behavior of the Smith dynamic in “bad RPS with a feeble twin” (equation (9.16)): In Figure 9.4.7, we present the maximum, time-average, and minimum weight on the dominated strategy in the limit cycle of the Smith dynamic, with these weights being presented as functions of the domination level d . The figure shows that Twin is recurrently played by at least 10% of the population when $d \leq .31$, by at least 5% of the population when $d \leq .47$, and by at least 1% of the population when $d \leq .66$. These values of d are surprisingly large relative to the base payoff values of 0, -2 , and 1: even strategies that are dominated by a significant margin can be played in perpetuity under common evolutionary dynamics. §

Example 9.4.8. Theorem 7.4.4 showed that dominated strategies are eliminated along interior solutions of imitative dynamics. But Theorem 9.4.1 shows that this result is not robust to small changes in these dynamics.

To understand why, consider evolution under the replicator dynamic in “(standard) RPS with a twin”. In standard Rock-Paper-Scissors, interior solutions of the replicator dynamic are closed orbits (see, e.g., Section 9.1.1). When we introduce an exact twin, equation (7.39) tells us that the ratio $\frac{x_S}{x_T}$ is constant along every solution trajectory. This is evident in Figure 9.4.5(i), which shows that the planes on which the ratio $\frac{x_S}{x_T}$ is constant

are all invariant sets. If we make the twin feeble by lowering its payoff uniformly by ε , we obtain the dynamics pictured in Figure 9.4.5(ii): now the ratio $\frac{x_S}{x_T}$ increases monotonically, and the dominated strategy is eliminated.

The existence of a continuum of invariant hyperplanes under imitative dynamics in games with identical twins is crucial to this argument. At the same time, dynamics with a continuum of invariant hyperplanes are structurally unstable. If we fix the game but slightly alter the agents' revision protocol, these invariant sets can collapse, overturning the elimination result.

To make this argument concrete, suppose that instead of always following an imitative protocol, agents occasionally use a protocol based on the direct selection of candidate strategies. Such a situation is illustrated in Figure 9.4.6(i), which contains the phase diagram for "bad RPS with a twin" (with benefit of winning $w = 1$ and cost of losing $l = \frac{11}{10}$) under a $(\frac{9}{10}, \frac{1}{10})$ convex combination of the replicator and Smith dynamics. While Figure 9.4.5(i) displayed a continuum of invariant hyperplanes, Figure 9.4.6(i) shows almost all solution trajectories converging to a limit cycle on the plane where $x_S = x_T$. If we make the twin feeble, the limit cycle moves slightly to the left, as in Figure 9.4.6(ii), and the dominated strategy survives. §

Exercise 9.4.9. Show that an analogue of equation (7.39) holds for the projection dynamic on $\text{int}(X)$. Explain why this does not imply that dominated strategies are eliminated along all solutions to the projection dynamic starting from interior initial conditions.

Exercise 9.4.10. We saw in Section 7.4 that dominated strategies are eliminated under the best response dynamic. Explain intuitively why this result does not extend to perturbed best response dynamics, including the $\text{logit}(\eta)$ dynamic with η small but positive. Using a computer, perform a numerical analysis of a well-chosen example to verify your intuition.

Appendix

9.A Three Classical Theorems on Nonconvergent Dynamics

9.A.1 Liouville's Theorem

Let $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ be a C^1 vector field, and consider the differential equation $\dot{x} = V(x)$ with flow $\phi : \mathbf{R} \times \mathbf{R}^n \rightarrow \mathbf{R}^n$. Let the set $A \subset \mathbf{R}^n$ be measurable with respect to Lebesgue measure μ on \mathbf{R}^n . Liouville's Theorem concerns the time evolution of $\mu(\phi_t(A))$, the measure (or *volume*) of the time t image of A under ϕ .

Theorem 9.A.1 (Liouville's Theorem). $\frac{d}{dt}\mu(\phi_t(A)) = \int_{\phi_t(A)} \text{tr}(DV(x)) d\mu(x)$.

The quantity $\text{tr}(DV(x)) = \sum_i \frac{\partial V_i}{\partial x_i}(x) \equiv \text{div}V(x)$ is known as the *divergence* of V at x . According to Liouville's Theorem, $\text{div}V$ governs the local rates of change in volume under the flow ϕ of $\dot{x} = V(x)$. In particular, if $\text{div}V = 0$ on an open set $O \subseteq \mathbf{R}^n$ —that is, if V is *divergence-free* on this set—then the flow ϕ conserves volume on O .

Before proceeding with the proof of Liouville's Theorem, let us note that it extends immediately to cases in which the law of motion $V : X \rightarrow TX$ defined on an affine set $X \subset \mathbf{R}^n$ with tangent space TX . In this case, μ represents Lebesgue measure on (the affine hull of) X . The only cautionary note is that the derivative of V at state $x \in X$ must be represented using the derivative matrix $DV(x) \in \mathbf{R}^{n \times n}$, which by definition has rows in TX . We showed how to compute this matrix in Appendix 3.B.3: if $\hat{V} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is a C^1 extension of V , then $DV(x) = D\hat{V}(x)P_{TX}$, where $P_{TX} \in \mathbf{R}^{n \times n}$ is the orthogonal projection of \mathbf{R}^n onto the subspace TX .

Proof. Using the standard multivariate change of variable, we express the measure of the set $\phi_t(A)$ as

$$(9.17) \quad \mu(\phi_t(A)) = \int_{\phi_t(A)} 1 d\mu(x_t) = \int_A |\det(D\phi_t(x_0))| d\mu(x_0).$$

The derivative matrix $D\phi_t(x_0)$ in equation (9.17) captures changes in $\phi_t(x_0)$, the time t position of the solution to $\dot{x} = V(x)$ from initial condition x_0 , as this initial condition is varied. It follows from arguments below that $\det(D\phi_t(x_0)) > 0$, so that the absolute value taken in equation (9.17) is unnecessary. Taking the time derivative of this equation and then differentiating under the integral sign thus yields

$$(9.18) \quad \frac{d}{dt}\mu(\phi_t(A)) = \int_A \frac{d}{dt} \det(D\phi_t(x_0)) d\mu(x_0).$$

Evaluating the right hand side of equation (9.18) requires two lemmas. The first of these is stated in terms of the time inhomogeneous linear equation

$$(9.19) \quad \dot{y}_t = DV(x_t)y_t,$$

where $\{x_t\}$ is the solution to $\dot{x} = V(x)$ from initial condition x_0 . Equation (9.19) is known as the (*first*) *variation equation* associated with $\dot{x} = V(x)$.

Lemma 9.A.2. *The matrix trajectory $\{D\phi_t(x_0)\}_{t \geq 0}$ is the matrix solution to the first variation*

equation from initial condition $D\phi_0(x_0) = I \in \mathbf{R}^{n \times n}$. More explicitly,

$$(9.20) \quad \frac{d}{dt} D\phi_t(x_0) = DV(\phi_t(x_0)) D\phi_t(x_0).$$

In words, Lemma 9.A.2 tells us that the column trajectories of $\{D\phi_t(x_0)\}_{t \geq 0}$ are the solutions to the first variation equation whose initial conditions are the standard basis vectors $e_1, \dots, e_n \in \mathbf{R}^n$.

Proof. By definition, the time derivative of the flow from x_0 satisfies $\frac{d}{dt} \phi_t(x_0) = V(\phi_t(x_0))$. Differentiating with respect to x_0 and then reversing the order of differentiation yields (9.20). ■

Lemma 9.A.3 provides two basic matrix identities, the first of which is sometimes called *Liouville's formula*.

Lemma 9.A.3. *Let $M \in \mathbf{R}^{n \times n}$. Then*

- (i) $\det(\exp(M)) = \exp(\text{tr}(M))$.
- (ii) $\left. \frac{d}{dt} \det(\exp(Mt)) \right|_{t=0} = \text{tr}(M)$.

Proving part (i) of the lemma is not difficult, but the intuition is clearest when M is a diagonal matrix:

$$\det \left(\exp \begin{pmatrix} \lambda_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \lambda_n \end{pmatrix} \right) = \det \begin{pmatrix} e^{\lambda_1} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & e^{\lambda_n} \end{pmatrix} = \prod_i e^{\lambda_i} = \exp \left(\sum_i \lambda_i \right) = \exp \left(\text{tr} \begin{pmatrix} \lambda_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \lambda_n \end{pmatrix} \right).$$

Part (ii) follows from part (i) by replacing M with Mt and differentiating.

Lemmas 9.A.2 and 9.A.3(ii) enable us to evaluate equation (9.18). First, note that Lemma 9.A.2 and the linearity of the first variation equation imply that

$$D\phi_t(x_0) \approx \exp(DV(x_0)t)$$

when t is close to 0. Combining this observation with Lemma 9.A.3(ii) shows that

$$\left. \frac{d}{dt} \det(D\phi_t(x_0)) \right|_{t=0} \approx \left. \frac{d}{dt} \det(\exp(DV(x_0)t)) \right|_{t=0} = \text{tr}(DV(x_0)).$$

By substituting this equality into equation (9.18) and noting that our focus on time $t = 0$ has been arbitrary, we obtain Liouville's Theorem. ■

Liouville's Theorem can be used to prove the nonexistence of asymptotically stable sets. Since solutions in a neighborhood of such a set all approach the set, volume must

be contracted in this neighborhood. It follows that a region in which divergence is nonnegative cannot contain an asymptotically stable set.

Theorem 9.A.4. *Suppose $\operatorname{div}V \geq 0$ on the open set $O \subseteq \mathbf{R}^n$, and let $A \subset O$ be compact. Then A is not asymptotically stable under $\dot{x} = V(x)$.*

Theorem 9.A.4 does not rule out the existence of Lyapunov stable sets. In fact, the example of the replicator dynamic in standard Rock-Paper-Scissors shows that such sets are not unusual when V is divergence-free.

9.A.2 The Poincaré-Bendixson and Bendixson-Dulac Theorems

We now present two classical results concerning differential equations on the plane.

The celebrated *Poincaré-Bendixson Theorem* characterizes the possible long run behaviors of such dynamics, and provides a simple way of establishing the existence of periodic orbits. Recall that a *periodic* (or *closed*) *orbit* of a differential equation is a nonconstant solution $\{x_t\}_{t \geq 0}$ such that $x_T = x_0$ for some $T > 0$.

Theorem 9.A.5 (The Poincaré-Bendixson Theorem). *Let $V : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ be Lipschitz continuous, and consider the differential equation $\dot{x} = V(x)$.*

- (i) *Let $x \in \mathbf{R}^2$. If $\omega(x)$ is compact, nonempty, and contains no rest points, then it is a periodic orbit.*
- (ii) *Let $Y \subset \mathbf{R}^2$. If Y is nonempty, compact, forward invariant, and contains no rest points, then it contains a periodic orbit.*

Theorem 9.A.5 tells us that in planar systems, the only possible ω -limit sets are rest points, sequences of trajectories leading from one rest point to another (called *heteroclinic orbits* where there are multiple rest points in the sequence and *homoclinic orbits* when there is just one), and *periodic orbits*. In part (i) of the theorem, the requirement that $\omega(x)$ be compact and nonempty are automatically satisfied when the dynamic is forward invariant on a compact set—see Proposition 7.A.1.

The next result, the *Bendixson-Dulac Theorem*, provides a method of ruling out the existence of closed orbits in planar systems. To state this theorem, we recall that a set $Y \subset \mathbf{R}^2$ is *simply connected* if it contains no holes: more precisely, if every closed curve in Y can be continuously contracted within Y to a single point.

Theorem 9.A.6 (The Bendixson-Dulac Theorem). *Let $V : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ be C^1 , and consider the differential equation $\dot{x} = V(x)$. If $\operatorname{div}V \neq 0$ throughout the simply connected set Y , then Y does not contain a closed orbit.*

Proof. If γ is a closed orbit in Y , then the region R bounded by γ is invariant under ϕ . Thus

$$\frac{d}{dt}\mu(\phi_t(R)) = \int_{\phi_t(R)} \operatorname{div}V(x) d\mu(x)$$

by Liouville's Theorem. Since $\operatorname{div}V$ is continuous and nonzero throughout Y , its sign must be constant throughout Y . If this sign is negative, then the volume of R contracts under ϕ ; if it is positive, then the volume of R expands under ϕ . Either conclusion contradicts the invariance of R under ϕ . ■

Both of the results above extend to dynamics defined on two-dimensional affine spaces in the obvious way.

9.B Attractors and Continuation

9.B.1 Attractors and Repellers

Let ϕ be a *semiflow* on the compact set $X \subset \mathbf{R}^n$: that is, $\phi : [0, \infty) \times X \rightarrow X$ is a continuous map with $\phi_0(x) = x$ that satisfies the group property $\phi_t(\phi_s(x)) = \phi_{t+s}(x)$ for all $s, t \geq 0$ and $x \in X$. We call the set $\mathcal{A} \subseteq X$ *forward invariant* under ϕ if $\phi_t(\mathcal{A}) = \mathcal{A}$ for all $t \geq 0$. Note that in this case, the sets $\{\phi_t(\mathcal{A})\}_{t \geq 0}$ are nested. We call \mathcal{A} *invariant* under ϕ if $\phi_t(\mathcal{A}) = \mathcal{A}$ for all $t \in \mathbf{R}$. It is implicit in this definition that on the set \mathcal{A} we have not only a semiflow, but also a *flow*: on \mathcal{A} , we can extend the map ϕ to be well-defined and satisfy the group property not just for times in $[0, \infty)$, but also for times in $(-\infty, \infty)$.

A set $\mathcal{A} \subseteq X$ is an *attractor* of ϕ if it is nonempty, compact, and invariant under ϕ , and if there is a neighborhood U of \mathcal{A} such that

$$(9.21) \quad \limsup_{t \rightarrow \infty} \sup_{x \in U} \operatorname{dist}(\phi_t(x), \mathcal{A}) = 0.$$

The set $B(\mathcal{A}) = \{x \in X : \omega(x) \subseteq \mathcal{A}\}$ is called the *basin* of \mathcal{A} .

The supremum operation in condition (9.21) ensures that states in the neighborhood U are attracted to \mathcal{A} uniformly in time. This uniformity is important: without it, the rest point in the flow on the circle from Example 7.A.3 would be an attractor. Because of the uniformity requirement, attractors differ from asymptotically stable sets (defined in Appendix 7.A.2) only in that the latter need not be invariant.

Attractors can be defined in a number of equivalent ways. In the following proposition,

the ω -limit of the set $U \subseteq X$ is defined as

$$\omega(U) = \bigcap_{t \geq 0} \text{cl} \left(\bigcup_{s \geq t} \phi_s(U) \right).$$

Proposition 9.B.1. *The following statements are equivalent:*

- (i) \mathcal{A} is an attractor of ϕ .
- (ii) $\mathcal{A} = \omega(U)$ for some neighborhood U of \mathcal{A} .
- (iii) $\mathcal{A} = \bigcap_{t \geq 0} \phi_t(O)$ for some open set O that satisfies $\phi_T(\text{cl}(O)) \subset O$ for some $T > 0$.
- (iv) $\mathcal{A} = \bigcap_{t \geq 0} \phi_t(O)$ for some open, forward invariant set O that satisfies $\phi_T(\text{cl}(O)) \subset O$ for some $T > 0$.
- (v) $\mathcal{A} = \bigcap_{t \geq 0} \phi_t(O)$ for some open, forward invariant set O that satisfies $\phi_t(\text{cl}(O)) \subset O$ for all $t > 0$.

In parts (iii), (iv), and (v), the set O is known as a *weak trapping region*, a *trapping region*, and a *strongly forward invariant trapping region*, respectively.

Now suppose that $\phi : (-\infty, \infty) \times X \rightarrow X$ is a flow on X with attractor \mathcal{A} , and let U be a trapping region for \mathcal{A} . The set $\mathcal{A}^* = \bigcap_{t \leq 0} \phi_t(X - U)$ is known as the *dual repeller* of \mathcal{A} . \mathcal{A}^* is the α -limit of a neighborhood of itself (i.e., it is the ω -limit of a neighborhood of itself under the time-reversed version of ϕ); it is also nonempty, compact, and invariant under ϕ .

The set $C(\mathcal{A}, \mathcal{A}^*) = X - (\mathcal{A} \cup \mathcal{A}^*)$ is called the *set of connecting orbits* of the attractor-repeller pair $(\mathcal{A}, \mathcal{A}^*)$. Theorem 9.B.2 shows that the behavior of the flow on this set is very simple: it admits a strict Lyapunov function.

Theorem 9.B.2. *Let $(\mathcal{A}, \mathcal{A}^*)$ be an attractor-repeller pair of the flow ϕ on the compact set X . Then there exists a continuous function $L : X \rightarrow [0, 1]$ with $L^{-1}(1) = \mathcal{A}^*$ and $L^{-1}(0) = \mathcal{A}$ such that L is decreasing on $C(\mathcal{A}, \mathcal{A}^*)$ under ϕ .*

If ϕ is only a semiflow, one can still find a continuous Lyapunov function $L : X \rightarrow [0, 1]$ with $L^{-1}(0) = \mathcal{A}$ that is decreasing on $C(\mathcal{A}, \mathcal{A}^*)$.

Interestingly, the notion of chain recurrence introduced in Appendix 7.A.1 can be characterized in terms of attractor-repeller pairs. In particular, the set \mathcal{CR} of chain recurrent states of the flow ϕ consists of those states found in every attractor-repeller pair of ϕ :

Theorem 9.B.3. $\mathcal{CR} = \bigcap_{(\mathcal{A}, \mathcal{A}^*)} (\mathcal{A} \cup \mathcal{A}^*)$.

By combining Theorems 9.B.2 and 9.B.3 with the fact that the number of attractor-repeller pairs of ϕ is countable, one can establish the *Fundamental Theorem of Dynamical Systems*:

Theorem 9.B.4. *Let ϕ be a flow on a compact metric space X . Then ϕ admits a Lyapunov function $L : X \rightarrow \mathbf{R}$ that is decreasing off $C\mathcal{R}$ and such that $L(C\mathcal{R})$ is a nowhere dense subset of \mathbf{R} .*

9.B.2 Continuation of Attractors

Consider now a one-parameter family of differential equations $\dot{x} = V^\varepsilon(x)$ in \mathbf{R}^n with unique solutions $x_t = \phi_t^\varepsilon(x_0)$ such that $(\varepsilon, x) \mapsto V^\varepsilon(x)$ is continuous. Then $(\varepsilon, t, x) \mapsto \phi_t^\varepsilon(x)$ is continuous as well. Suppose that $X \subset \mathbf{R}^n$ is compact and forward invariant under the semi-flows ϕ^ε . For $\varepsilon = 0$ we omit the superscript in ϕ .

Roughly speaking, Theorem 9.B.5 shows that each attractor of the differential equation $\dot{x} = V^0(x)$ will “survive” any small enough perturbation of the differential equation.

Theorem 9.B.5. *Let \mathcal{A} be an attractor for ϕ with basin $B(\mathcal{A})$. Then for each small enough $\varepsilon > 0$ there exists an attractor \mathcal{A}^ε of ϕ^ε with basin $B(\mathcal{A}^\varepsilon)$, such that the map $\varepsilon \mapsto \mathcal{A}^\varepsilon$ is upper hemicontinuous and the map $\varepsilon \mapsto B(\mathcal{A}^\varepsilon)$ is lower hemicontinuous.*

Upper hemicontinuity cannot be replaced by continuity in this result. Consider the family of differential equations $\dot{x} = (\varepsilon + x^2)(1 - x)$ on the real line. The semi-flow ϕ corresponding to $\varepsilon = 0$ admits $\mathcal{A} = [0, 1]$ as an attractor, but when $\varepsilon > 0$ the unique attractor of ϕ^ε is $\mathcal{A}^\varepsilon = \{1\}$. This example shows that perturbations can cause attractors to implode; the theorem shows that perturbations cannot cause attractors to explode.

Theorem 9.B.5 is a direct consequence of the following lemma.

Lemma 9.B.6. *Let \mathcal{A} be an attractor for ϕ with basin $B(\mathcal{A})$, and let U_1 and U_2 be open sets satisfying $\mathcal{A} \subset U_1 \subseteq U_2 \subseteq \text{cl}(U_2) \subseteq B(\mathcal{A})$. Then for each small enough $\varepsilon > 0$ there exists an attractor \mathcal{A}^ε of ϕ_ε with basin $B(\mathcal{A}^\varepsilon)$, such that $\mathcal{A}^\varepsilon \subset U_1$ and $U_2 \subset B(\mathcal{A}^\varepsilon)$.*

In this lemma, one can always set $U_1 = \{x : \text{dist}(x, \mathcal{A}) < \delta\}$ and $U_2 = \{x \in B(\mathcal{A}) : \text{dist}(x, X - B(\mathcal{A})) > \delta\}$ for some small enough $\delta > 0$.

Proof of Lemma 9.B.6. Since \mathcal{A} is an attractor and $\omega(\text{cl}(U_2)) = \mathcal{A}$, there is a $T > 0$ such that $\phi_t(\text{cl}(U_2)) \subset U_1$ for $t \geq T$. By the continuous dependence of the flow on the parameter ε and the compactness of $\phi_T(\text{cl}(U_2))$, we have that $\phi_T^\varepsilon(\text{cl}(U_2)) \subset U_1 \subseteq U_2$ for all small enough ε . Thus, U_2 is a weak trapping region for the semi-flow ϕ^ε , and so $\mathcal{A}^\varepsilon \equiv \omega(U_2)$ is an attractor for ϕ^ε . In addition, $\mathcal{A}^\varepsilon \subset U_1$ (since $\mathcal{A}^\varepsilon = \phi_T^\varepsilon(\mathcal{A}^\varepsilon) \subseteq \phi_T^\varepsilon(\text{cl}(U_2)) \subset U_1$) and $U_2 \subset B(\mathcal{A}^\varepsilon)$. ■

9.N Notes

Section 9.1. The conservative properties of dynamics studied in this chapter—the existence of a constant of motion and the preservation of volume—are basic properties of Hamiltonian systems. For more on this connection, see Akin and Losert (1984) and Hofbauer (1995a, 1996); for a general introduction to Hamiltonian systems, see Marsden and Ratiu (2002). Exercises 9.1.4 and 9.1.5 are due to Schuster et al. (1981b,c). Theorem 9.1.7 is due to Akin and Losert (1984) and Hofbauer (1995a), while Theorem 9.1.8 is due to Hofbauer and Sigmund (1988) and Ritzberger and Weibull (1995) (also see Weibull (1995)).

Section 9.2. Circulant games were introduced by Hofbauer and Sigmund (1988), who call them “cyclically symmetric games”; also see Schuster et al. (1981c). The hypercycle system was proposed by Eigen and Schuster (1979) to model of cyclical catalysis in a collection of polynucleotides during prebiotic evolution. Schuster et al. (1979) introduce the term *permanent* to describe dynamics on the simplex under which the boundary is a repeller. The permanence of the hypercycle system when $n \geq 5$ was established by Hofbauer et al. (1981); the existence of stable limit cycles in this context was proved by Hofbauer et al. (1991). For further results on permanence under the replicator and related dynamics, see Hofbauer (1981a), Hutson (1984), Hofbauer and Sigmund (1988), Hutson and Schmitt (1992), Schreiber (2000), and Garay and Hofbauer (2003).

Monocyclic games are studied in the context of the replicator dynamic by Hofbauer and Sigmund (1988), who call them “essentially hypercyclic” games. The uniqueness of the interior Nash equilibrium in Example 9.2.5 follows from the fact that the replicator dynamic for this game is permanent; see Hofbauer and Sigmund (1988). The analysis of the best response dynamic in this example is due to Hofbauer (1995b), Gaunersdorfer and Hofbauer (1995), and Benaïm et al. (2009). Proposition 9.2.7 is due to Hofbauer and Swinkels (1996); also see Hofbauer and Sigmund (1998, Section 8.6).

The Mismatching Pennies game was introduced by Jordan (1993), and was inspired by a 3×3 example due to Shapley (1964); see Sparrow et al. (2008) and van Strien and Sparrow (2009) for recent analyses that identify chaotic behavior in Shapley’s (1964) example. The analyses of the replicator and best response dynamics in Mismatching Pennies are due to Gaunersdorfer and Hofbauer (1995). Proposition 9.2.11 is due to Hart and Mas-Colell (2003). The hypnodisk game is introduced in Hofbauer and Sandholm (2006). Viossat (2007, 2008) constructs examples in which all pure strategies used in correlated equilibria vanish along some (nonconvergent) solution trajectories of various dynamics.

Section 9.3. For introductions to chaotic differential equations, see Hirsch et al. (2004) at the undergraduate level or Guckenheimer and Holmes (1983) at the graduate level. Example 9.3.1 is due to Arneodo et al. (1980), who introduce it in the context of the Lotka-

Volterra equations; also see Skyrms (1992). The attractor in this example is known as a Shilnikov attractor; see Hirsch et al. (2004, Chapter 16). Example 9.3.2 is due to Sato et al. (2002). An early analysis of chaotic behavior under the best response dynamic is Cowan (1992); see Sparrow et al. (2008) and van Strien and Sparrow (2009) for more recent analyses.

Section 9.4. This section follows Hofbauer and Sandholm (2006). That paper builds on the work of Berger and Hofbauer (2006), who show that strictly dominated strategies can survive under the BNN dynamic. For a survival result for the projection dynamic, see Sandholm et al. (2008).

Section 9.A. For further details on Liouville's Theorem, see Sections 4.1 and 5.3 of Hartman (1964). Theorem 9.A.4 in the text is Proposition 6.6 of Weibull (1995). For treatments of the Poincaré-Bendixson Theorem, see Hirsch and Smale (1974) and Robinson (1995).

Section 9.B. The definition of attractor we use is from Benaïm (1999). Definition (ii) in Proposition 9.B.1 is from Conley (1978), and definitions (iii), (iv), and (v) are from Robinson (1995). Theorems 9.B.2, 9.B.3, and 9.B.4 are due to Conley (1978).

Theorem 9.B.5 is part of the folklore of dynamical systems theory; compare Proposition 8.1 of Smale (1967). The analysis presented here is from Hofbauer and Sandholm (2006).

Part IV

Stochastic Evolutionary Models

Stochastic Evolution and Deterministic Approximation

10.0 Introduction

In Parts II and III of this book, we investigated the evolution of aggregate behavior under deterministic dynamics. We provided foundations for these dynamics in Chapter 4: there we showed that given any revision protocol ρ and population game F , we can derive a mean dynamic $\dot{x} = V_F(x)$. This differential equation describes expected motion under the stochastic process that ρ and F implicitly define. We justified our focus on this deterministic equation through an informal appeal to a law of large numbers: since all of the randomness in our evolutionary model is idiosyncratic, it should be averaged away in the aggregate so long as the population size is sufficiently large.

Our goal in this chapter is to make this argument rigorous. To do so, we explicitly derive a stochastic evolutionary process—a Markov process—from a given population game F , revision protocol ρ , and finite population size N . Our main result in this chapter, Theorem 10.2.3, is a finite horizon deterministic approximation theorem. Building on our earlier intuition, the theorem shows that over any finite time span, the behavior of the stochastic evolutionary process is indeed nearly deterministic: if the population size is large enough, the stochastic process closely follows a solution trajectory of the mean dynamic with probability close to one.

The Markov process we introduce in this chapter provides a precise description of the stochastic evolution of aggregate behavior. Theorem 10.2.3 tells us that over time horizons of moderate length, we can do without studying this Markov process directly, as the deterministic approximation is adequate to address most questions of interest. But if we want to understand behavior in a society over very long time spans, then the

deterministic approximation theorem no longer applies, and we must study the Markov process directly. This infinite horizon analysis is the subject of our final chapter.

The remaining three chapters employ a variety of techniques from the theory of probability and stochastic processes. These techniques are reviewed in Appendices 10.A, 10.B, and 10.C and in the appendices to Chapters 11 and 12.

In all of the remaining chapters, we will conserve on notation by only considering single-population models. All of the results we present can be extended to multipopulation models with little difficulty.

10.1 The Stochastic Evolutionary Process

We begin by developing the model of stochastic evolution introduced in Section 4.1. We consider a population of agents who recurrently play a population game $F : X \rightarrow \mathbf{R}^n$ with pure strategy set $S = \{1, \dots, n\}$. The agents' choice procedure is described by a revision protocol $\rho : \mathbf{R}^n \times X \rightarrow \mathbf{R}_+^{n \times n}$ that takes current payoffs and population states as inputs and returns collections of conditional switch rate $\rho_{ij}(F(x), x)$ as outputs.

To set the stage for our limiting analysis, we suppose now that the population size is large but finite, with N members. The feasible social states therefore lie in the discrete grid $X^N = X \cap \frac{1}{N}\mathbf{Z}^n = \{x \in X : Nx \in \mathbf{Z}^n\}$.

Exercise 10.1.1. Show that X^N has cardinality $\binom{N+n-1}{n-1}$. (Hint: Use induction on the number of strategies n .) ‡

The stochastic process $\{X_t^N\}$ generated by F, ρ, N is described as follows. Each agent in the society is equipped with a rate R Poisson alarm clock, where $R < \infty$ is an upper bound on the row sums of ρ :

$$(10.1) \quad R \geq \max_{x,i} \sum_{j \in S} \rho_{ij}(F(x), x).$$

We will see that subject to satisfying this constraint, the choice of R is irrelevant to our approximation results.

The ringing of a clock signals the arrival of a revision opportunity for the clock's owner: if the owner is currently playing strategy $i \in S$, he switches to strategy $j \neq i$ with probability ρ_{ij}/R . (Remember that the diagonal elements of ρ , though sometimes useful as placeholders (see equation (4.1)), play no formal role in the model.) Finally, the model respects independence assumptions that ensure that "the future is independent of the past except through the present": different agents' clocks ring independently of one

another, strategy choices are made independently of the timing of the clocks' rings, and as evolution proceeds, the clocks and the agents are only influenced by the history of the process by way of the current value of the social state.

In Chapter 4, we argued informally that the stochastic process described above is well approximated by solutions to the mean dynamic

$$(M) \quad \dot{x}_i = \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x).$$

The rest of this chapter provides a formal defense of this approximation result.

We begin by giving a more formal account of the stochastic evolutionary process $\{X_t^N\}$. The independence assumptions above ensure that $\{X_t^N\}$ is a continuous-time Markov process on the finite state space \mathcal{X}^N . To describe this process explicitly, it is enough to specify its jump rates $\{\lambda_x^N\}_{x \in \mathcal{X}^N}$ and transition probabilities $\{P_{xy}^N\}_{x, y \in \mathcal{X}^N}$ (see Appendix 10.C).

If the current social state is $x \in \mathcal{X}^N$, then Nx_i of the N agents are playing strategy $i \in S$. Since agents receive revision opportunities independently at exponential rate R , the basic properties of the exponential distribution (see Proposition 10.A.1) imply that revision opportunities arrive in the society as a whole at exponential rate NR .

When an agent playing strategy $i \in S$ receives a revision opportunity, he switches to strategy $j \neq i$ with probability ρ_{ij}/R . Since this choice is independent of the arrivals of revision opportunities, the probability that the next revision opportunity goes to an agent playing strategy i who then switches to strategy j is

$$\frac{Nx_i}{N} \times \frac{\rho_{ij}}{R} = \frac{x_i \rho_{ij}}{R}.$$

This switch decreases the number of agents playing strategy i by one and increases the number playing j by one, shifting the state by $\frac{1}{N}(e_j - e_i)$.

Summarizing this analysis yields the following observation, which specifies the parameters of the Markov process $\{X_t^N\}$.

Observation 10.1.2. *A population game F , a revision protocol ρ , a constant R satisfying (10.1), and a population size N define a Markov process $\{X_t^N\}$ on the state space \mathcal{X}^N . This process is*

described by some initial state $X_0^N = x_0^N$, the jump rates $\lambda_x^N = NR$, and the transition probabilities

$$P_{x,x+z}^N = \begin{cases} \frac{x_i \rho_{ij}(F(x), x)}{R} & \text{if } z = \frac{1}{N}(e_j - e_i), i, j \in S, i \neq j, \\ 1 - \sum_{i \in S} \sum_{j \neq i} \frac{x_i \rho_{ij}(F(x), x)}{R} & \text{if } z = \mathbf{0}, \\ 0 & \text{otherwise.} \end{cases}$$

10.2 Finite Horizon Deterministic Approximation

In the previous section, we formally defined the Markov process $\{X_t^N\}$ generated by a population game F , a revision protocol ρ , and a population size N . Earlier, we introduced the mean dynamic (M), an ordinary differential equation that captures the expected motion of this process; solutions $\{x_t\}$ of (M) are continuous paths through the set of social states. Can we say more precisely how the stochastic and deterministic processes are linked?

The main result in this chapter, Theorem 10.2.3, shows that when the population size N is sufficiently large, the Markov process $\{X_t^N\}$ is well approximated over finite time spans by the deterministic trajectory $\{x_t\}$.

10.2.1 Kurtz's Theorem

To begin, we state a general result on the convergence of a sequence $\{\{X_t^N\}_{N=N_0}^\infty\}$ of Markov processes with decreasing step sizes. We suppose that the process indexed by N takes values in the state space $\mathcal{X}^N = \{x \in X : Nx \in \mathbf{Z}^n\}$, and we let $\lambda^N \in \mathbf{R}_+^{\mathcal{X}^N}$ and $P^N \in \mathbf{R}_+^{\mathcal{X}^N \times \mathcal{X}^N}$ denote the jump rate vector and transition matrix of this process.

To simplify the definitions to follow, we let ζ_x^N be a random variable (defined on an arbitrary probability space) whose distribution describes the stochastic increment of $\{X_t^N\}$ from state x :

$$(10.2) \quad \mathbb{P}(\zeta_x^N = z) = P_{x,x+z}^N,$$

where τ_k is the time of the process's k th jump. We then define the functions $V^N : \mathcal{X}^N \rightarrow TX$, $A^N : \mathcal{X}^N \rightarrow \mathbf{R}$, and $A_\delta^N : \mathcal{X}^N \rightarrow \mathbf{R}$ by

$$\begin{aligned} V^N(x) &= \lambda_x^N \mathbb{E} \zeta_x^N, \\ A^N(x) &= \lambda_x^N \mathbb{E} |\zeta_x^N|, \\ A_\delta^N(x) &= \lambda_x^N \mathbb{E} \left| \zeta_x^N \mathbf{1}_{\{|\zeta_x^N| > \delta\}} \right|. \end{aligned}$$

$V^N(x)$, the product of the jump rate at state x and the expected increment per jump at x , represents the expected increment per time unit from x under $\{X_t^N\}$. V^N is thus an alternate definition of the mean dynamic of $\{X_t^N\}$. In a similar vein, $A^N(x)$ is the expected absolute displacement per time unit, while $A_\delta^N(x)$ is the expected absolute displacement per time unit due to jumps traveling further than δ .

With these definitions in hand, we can state the basic approximation result.

Theorem 10.2.1 (Kurtz's Theorem). *Let $V : X \rightarrow TX$ be a Lipschitz continuous vector field. Suppose that for some sequence $\{\delta^N\}_{N=N_0}^\infty$ converging to 0, we have*

$$(10.3) \quad \lim_{N \rightarrow \infty} \sup_{x \in \mathcal{X}^N} |V^N(x) - V(x)| = 0,$$

$$(10.4) \quad \sup_N \sup_{x \in \mathcal{X}^N} A^N(x) < \infty, \text{ and}$$

$$(10.5) \quad \lim_{N \rightarrow \infty} \sup_{x \in \mathcal{X}^N} A_{\delta^N}^N(x) = 0,$$

and that the initial conditions $X_0^N = x_0^N$ converge to $x_0 \in X$. Let $\{x_t\}_{t \geq 0}$ be the solution to the mean dynamic

$$(M) \quad \dot{x} = V(x)$$

starting from x_0 . Then for each $T < \infty$ and $\varepsilon > 0$, we have that

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} |X_t^N - x_t| < \varepsilon \right) = 1.$$

Fix a finite time horizon $T < \infty$ and an error bound $\varepsilon > 0$. Kurtz's Theorem tells us that when the index N is large, nearly all sample paths of the Markov process $\{X_t^N\}$ stay within ε of a solution of the mean dynamic (M) through time T . By making N large enough, we can ensure that with probability close to one, X_t^N and x_t differ by no more than ε for all t between 0 and T (Figure 10.2.1).

What conditions do we need to reach this conclusion? Condition (10.3) demands that as N grows large, the expected displacements per time unit V^N converge uniformly to a Lipschitz continuous vector field V . Lipschitz continuity of V ensures the existence and uniqueness of solutions of the mean dynamic $\dot{x} = V(x)$. Condition (10.4) requires that the expected absolute displacement per time unit is bounded. Finally, condition (10.5) demands that jumps larger than δ^N make vanishing contributions to the motion of the processes, where $\{\delta^N\}_{N=N_0}^\infty$ is a sequence of constants that approaches zero.

The intuition behind Kurtz's Theorem can be explained as follows. At each revision




Figure 10.2.1: Kurtz's Theorem.

opportunity, the increment in the process $\{X_t^N\}$ is stochastic. However, the expected number of revision opportunities that arrive during the brief time interval $I = [t, t + dt]$ is of order $\lambda_x^N dt$. Whenever it does not vanish, this quantity grows without bound as the population size N becomes large. Conditions (10.4) and (10.5) ensure that when N is large, each increment in the state is likely to be small. This ensures that the total change in the state during time interval I is small, so that jump rates and transition probabilities vary little during this interval. Since during I there are a very large number of revision opportunities, each generating nearly the same expected increment, intuition from the law of large numbers suggests that the change in $\{X_t^N\}$ during the interval should be almost completely determined by the expected motion of $\{X_t^N\}$. This expected motion is captured by the limiting mean dynamic V , whose solutions approximate the stochastic process $\{X_t^N\}$ over finite time spans with probability close to one.

Exercise 10.2.2. Suppose that $\{X_t^N\}$ is a Markov process on $\mathcal{X}^N = \{0, \frac{1}{N}, \frac{2}{N}, \dots, 1\}$ with $\lambda_x^N \equiv N$. To ensure that $\frac{1}{2}$ is always a state, restrict attention to even N . Give examples of sequences of transition probabilities from state $\frac{1}{2}$ that

- (i) satisfy condition (10.3) of Kurtz's Theorem, but not conditions (10.4) or (10.5);
- (ii) satisfy conditions (10.3) and (10.5), but not condition (10.4);
- (iii) satisfy conditions (10.3) and (10.4), but not condition (10.5).

(Hint: it is enough to consider transition probabilities under which $V^N(\frac{1}{2}) = 0$.)

Guided by your answers to parts (ii) and (iii), explain intuitively what conditions (10.4)

and (10.5) require.

10.2.2 Deterministic Approximation of the Stochastic Evolutionary Process

Returning to our model of evolution, we now use Kurtz's Theorem to show that the Markov processes $\{\{X_t^N\}_{N=N_0}^\infty\}$ defined in Section 10.1 can be approximated by solutions to the mean dynamic (M) derived in Section 4.2.1.

To begin, we compute the expected increment per time unit $V^N(x)$ of the process $\{X_t^N\}$. Defining the random variable ζ_x^N as in equation (10.2) above, we find that

$$\begin{aligned}
V^N(x) &= \lambda_x^N \mathbb{E} \zeta_x^N \\
&= NR \sum_{i \in S} \sum_{j \neq i} \frac{1}{N} (e_j - e_i) \mathbb{P} \left(\zeta_x^N = \frac{1}{N} (e_j - e_i) \right) \\
&= NR \sum_{i \in S} \sum_{j \neq i} \frac{1}{N} (e_j - e_i) \frac{x_i \rho_{ij}}{R} \\
&= \sum_{i \in S} \sum_{j \in S} (e_j - e_i) x_i \rho_{ij} \\
&= \sum_{j \in S} e_j \sum_{i \in S} x_i \rho_{ij} - \sum_{i \in S} e_i x_i \sum_{j \in S} \rho_{ij} \\
&= \sum_{i \in S} e_i \left(\sum_{j \in S} x_j \rho_{ji} - x_i \sum_{j \in S} \rho_{ij} \right).
\end{aligned}$$

Thus, the vector field $V^N = V$ is independent of N , and is expressed more concisely as

$$(M) \quad V_i(x) = \sum_{j \in S} x_j \rho_{ji} - x_i \sum_{j \in S} \rho_{ij},$$

as we established using a different calculation in Section 4.2.1.

Conditions (10.4) and (10.5) of Kurtz' Theorem require that the motions of the processes $\{X_t^N\}$ not be too abrupt. To verify these conditions, observe that since $|e_j - e_i| = \sqrt{2}$ for any distinct $i, j \in S^p$, the increments of $\{X_t^N\}$ are always either of length $\frac{\sqrt{2}}{N}$ or of length zero. If we choose $\delta^N = \frac{\sqrt{2}}{N}$, this observation immediately implies condition (10.5):

$$A_{\frac{\sqrt{2}}{N}}^N(x) = \lambda_x^N \mathbb{E} \left| \zeta_x^N 1_{\left\{ |\zeta_x^N| > \frac{\sqrt{2}}{N} \right\}} \right| = 0.$$

The observation also helps us verify condition (10.4):

$$A^N(x) = \lambda_x^N \mathbb{E} |\zeta_x^N| \leq RN \times \frac{\sqrt{2}}{N} = \sqrt{2}R.$$

With these calculations in hand, we can present the deterministic approximation theorem.

Theorem 10.2.3 (Deterministic Approximation of $\{X_t^N\}$). *Let $\{\{X_t^N\}_{N=N_0}^\infty\}$ be the sequence of stochastic evolutionary processes defined in Observation 10.1.2. Suppose that $V = V^N$ is Lipschitz continuous. Let the initial conditions $X_0^N = x_0^N$ converge to state $x_0 \in X$, and let $\{x_t\}_{t \geq 0}$ be the solution to the mean dynamic (M) starting from x_0 . Then for all $T < \infty$ and $\varepsilon > 0$,*

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} |X_t^N - x_t| < \varepsilon \right) = 1.$$

Choose a finite time span T and two small constants δ and ε . Then for all large enough population sizes N , the probability that the process $\{X_t^N\}$ stays within ε of the deterministic trajectory $\{x_t\}$ through time T is at least $1 - \delta$.

A key requirement of Theorem 10.2.3 is that V must be Lipschitz continuous, ensuring that the mean dynamic (M) admits a unique solution from every initial condition in X . This requirement is satisfied by members of the families of dynamics (imitative, excess payoff, pairwise comparison) studied in Chapter 5, as well as the perturbed best response dynamics from Chapter 6. The best response and projection dynamics, being discontinuous, are not covered by Theorem 10.2.3, but it seems likely that deterministic approximation results that apply to these dynamics can be proved (see the Notes).

It is well worth emphasizing that Theorem 10.2.3 is a *finite horizon* approximation result, and that it cannot be extended to an infinite horizon result. To see why not, consider the logit choice protocol (Example 4.3.5). Under this protocol, switches between all pairs of strategies occur with positive probability regardless of the current state. It follows that the induced Markov process $\{X_t^N\}$ is *irreducible*: there is a positive probability path between each ordered pair of states in \mathcal{X}^N . As we shall see in Chapter 11, irreducibility implies that every state in \mathcal{X}^N is visited infinitely often with probability one. This fact clearly precludes an infinite horizon analogue of Theorem 10.2.3. Indeed, infinite horizon analysis of $\{X_t^N\}$ requires a different set of tools, which we present in the next two chapters.

Example 10.2.4. Toss and switch. Suppose that agents play a game with strategy set $S = \{L, R\}$ using the constant revision protocol ρ , where $\rho_{LL} = \rho_{LR} = \rho_{RL} = \rho_{RR} = \frac{1}{2}$. Under the simplest interpretation of this protocol, each agent receives revision opportunities at rate 1; upon receiving an opportunity, an agent flips a fair coin, switching strategies if the coin

comes up Heads.

For each population size N , the protocol generates a Markov process $\{X_t^N\}$ with common jump rate $\lambda_x^N \equiv N$ and transition probabilities

$$P_{x,x+z}^N = \begin{cases} \frac{1}{2}x_R & \text{if } z = \frac{1}{N}(e_L - e_R), \\ \frac{1}{2}x_L & \text{if } z = \frac{1}{N}(e_R - e_L), \\ \frac{1}{2} & \text{if } z = \mathbf{0}. \end{cases}$$

We can simplify the notation by replacing the vector state variable $x = (x_L, x_R) \in X$ with the scalar state variable $\chi = x_R \in [0, 1]$. The resulting Markov process has common jump rate $\lambda_\chi^N \equiv N$ and transition probabilities

$$P_{\chi,\chi+z}^N = \begin{cases} \frac{1}{2}\chi & \text{if } z = -\frac{1}{N}, \\ \frac{1}{2}(1 - \chi) & \text{if } z = \frac{1}{N}, \\ \frac{1}{2} & \text{if } z = 0. \end{cases}$$

Its mean dynamic is thus

$$\begin{aligned} V^N(\chi) &= \lambda_\chi^N \mathbb{E}\zeta_\chi^N \\ &= N \left(\left(-\frac{1}{N} \cdot \frac{1}{2}\chi\right) + \left(\frac{1}{N} \cdot \frac{1}{2}(1 - \chi)\right) \right) \\ &= \frac{1}{2} - \chi, \end{aligned}$$

regardless of the population size N . To solve this differential equation, we move the rest point $\chi = \frac{1}{2}$ to the origin using change of variable $v = \chi - \frac{1}{2}$. The equation

$$\dot{v} = \dot{\chi} = \frac{1}{2} - \chi = \frac{1}{2} - (v + \frac{1}{2}) = -v$$

has the general solution $v_t = v_0 e^{-t}$, implying that

$$\chi_t = \frac{1}{2} + (\chi_0 - \frac{1}{2}) e^{-t}.$$

Fix a time horizon $T < \infty$. Theorem 10.2.3 tells us that when N is sufficiently large, the evolutionary process is very likely to stay very close to an almost deterministic trajectory; this trajectory converges to state $\chi = \frac{1}{2}$, with convergence occurring at exponential rate 1.

If we instead fix the population size N and look at behavior over the infinite time horizon ($T = \infty$), the process eventually splits off from the deterministic trajectory, visiting all states in $\{0, \frac{1}{N}, \dots, 1\}$ infinitely often. We will consider the infinite horizon behavior of

this process in more detail in Example 11.2.1. §

Exercise 10.2.5. Consider a population playing game F using revision protocol ρ .

- (i) Show that the resulting mean dynamic can be expressed as

$$\dot{x} = R(x)'x,$$

where $R(x) \in \mathbf{R}^{n \times n}$ is given by

$$R_{ij}(x) = \begin{cases} \rho_{ij}(x, F(x)) & \text{if } i \neq j, \\ -\sum_{k \neq i} \rho_{ik}(x, F(x)) & \text{if } i = j. \end{cases}$$

Note that when ρ is independent of $F(x)$ and x as in the previous example, the matrix R is independent of x as well. In this case we obtain the linear dynamic $\dot{x} = R'x$, whose solutions can be expressed in closed form (see Appendix 8.B).

- (ii) Suppose that $\rho_{ij} = 1$ for all i and j . Describe the parameters of the resulting Markov process $\{X_t^N\}$, and write down the corresponding mean dynamic. Show that solutions to the mean dynamic take the form $x_t = x^* + (x_0 - x^*)e^{-nt}$, where $x^* = \frac{1}{n}\mathbf{1}$.

10.3 Extensions

10.3.1 Discrete-Time Models

It is also possible to prove deterministic approximation results for discrete time models of stochastic evolution. To do so, we assume that the number of discrete periods that pass per unit of clock time grows with the population size N . In this situation, one can employ a discrete time version of Kurtz's Theorem, the requirements of which are direct analogues of those from Theorem 10.2.1 above.

To obtain a deterministic approximation theorem for discrete time Markov chains, we must assume that the length of a period with respect to clock time becomes vanishingly small as the population size N increases. Let d^N be the duration of a period under the Markov chain $\{X_t^N\}$, so that this chain is initialized at time 0 and has transitions at times $d^N, 2d^N, \dots$. We can define $\{X_t^N\}$ at all times in $[0, \infty)$ by letting $X_t^N = X_{kd^N}^N$ when $t \in [kd^N, (k+1)d^N)$, making each sample path $\{X_t^N(\omega)\} = \{X_t^N(\omega)\}_{t \geq 0}$ a step function whose jumps occur at multiples of d^N .

Theorem 10.3.1 (Kurtz's Theorem in Discrete Time). *Suppose that $\lim_{N \rightarrow \infty} d^N = 0$. Define the distributions of the random variables ζ_x^N by*

$$\mathbb{P}(\zeta_x^N = z) = \mathbb{P}\left(X_{(k+1)d^N}^N = x + z \mid X_{kd^N}^N = x\right),$$

and define the functions V^N, A^N , and A_δ^N by

$$V^N(x) = \frac{1}{d^N} \mathbb{E} \zeta_x^N, \quad A^N(x) = \frac{1}{d^N} \mathbb{E} |\zeta_x^N|, \quad \text{and} \quad A_\delta^N(x) = \frac{1}{d^N} \mathbb{E} \left[\zeta_x^N \mathbf{1}_{\{|\zeta_x^N| > \delta\}} \right].$$

Then, mutatis mutandis, Theorem 10.2.1 holds for the sequence of Markov chains $\{\{X_t^N\}_{N=N_0}^\infty\}$.

To apply the theorem, let us suppose that when the population size is N , each discrete time period is of duration $d^N = \frac{1}{NR}$, so that periods begin at times in the set $T^N = \{0, d^N, 2d^N, \dots\}$. The exercises to follow propose two specifications of the discrete time evolutionary process $\{X_t^N\}_{t \in T^N}$.

Exercise 10.3.2. Discrete time model I: One revision opportunity per period. Suppose that during each period, exactly one agent is selected at random and granted a revision opportunity, with each agent being equally likely to be chosen. The chosen agent's choices are then governed by the conditional switch probabilities ρ_{ij}/R . Using Theorem 10.3.1, show that Theorem 10.2.3 extends to this discrete time model.

Discrete time models can allow a possibility that our continuous time model cannot: they permit many agents to switch strategies simultaneously. The next exercise shows that deterministic approximation is still possible even when simultaneous revisions by many agents are possible, so long as they are sufficiently unlikely.

Exercise 10.3.3. Discrete time model II: Random numbers of revision opportunities in each period. Suppose that during each period, each agent tosses a coin that comes up heads with probability $\frac{1}{N}$. Every agent who tosses a head receives a revision opportunity; choices for such agents are again governed by the conditional switch probabilities ρ_{ij}/R . Use the Poisson Limit Theorem (Propositions 10.A.4(ii) and 10.A.5) and Theorem 10.3.1 to show that Theorem 10.2.3 extends to this model. (Hint: In any given period, the number of agents whose tosses come up heads is binomially distributed with parameters N and $\frac{1}{N}$.)

10.3.2 Finite-Population Adjustments

In our model of individual choice, the revision protocol ρ was defined independently of the population size. In some cases, it is more appropriate to allow the revision protocol to

depend on N in some vanishing way—for example, to account for the effects of sampling from a finite population, or for the fact that an agent whose choices are based on imitation will not imitate himself. If we include these effects, then ρ varies with N , so the normalized expected increments V^N vary with N as well. Fortunately, Kurtz’s Theorem allows for these sorts of effects so long as they are vanishing in size: examining condition (10.3), we see that as long as the functions V^N converge uniformly to a limiting mean dynamic V , the finite horizon approximation continues to hold.

Finite population adjustments will play a more important role when we consider infinite horizon behavior, where these adjustments can greatly simplify our calculations; see especially Sections 11.4 and 11.5.

Appendix

10.A The Exponential and Poisson Distributions

10.A.1 Basic Properties

The random variable T with support $[0, \infty)$ has an *exponential distribution with rate λ* , denoted $T \sim \text{exponential}(\lambda)$, if its decumulative distribution is $\mathbb{P}(T \geq t) = e^{-\lambda t}$, so that its density function is $f(t) = \lambda e^{-\lambda t}$. A Taylor approximation shows that for small $dt > 0$,

$$(10.6) \quad \mathbb{P}(T \leq dt) = 1 - e^{-\lambda dt} = 0 + \lambda e^{-\lambda \cdot 0} dt + O((dt)^2) \approx \lambda dt.$$

Exponential random variables are often used to model the random amount of time that passes before a certain occurrence: the arrival of a customer at a queue, the decay of a particle, and so on. We often describe the behavior of exponential random variables using the rhetorical device of a “stochastic alarm clock” that rings after an exponentially distributed amount of time has passed.

Some basic properties of the exponential distribution are listed next.

Proposition 10.A.1. *Let T_1, \dots, T_n be independent with $T_i \sim \text{exponential}(\lambda_i)$. Then*

- (i) $\mathbb{E}T_i = \lambda_i^{-1}$;
- (ii) $\mathbb{P}(T_i \geq u + t | T_i \geq u) = \mathbb{P}(T_i \geq t) = e^{-\lambda_i t}$;
- (iii) *If $M_n = \min\{T_1, \dots, T_n\}$ and $I_n = \operatorname{argmin}_j T_j$, then $M_n \sim \text{exponential}(\sum_{i=1}^n \lambda_i)$, $\mathbb{P}(I_n = i) = \lambda_i / \sum_{j=1}^n \lambda_j$, and M_n and I_n are independent.*

Property (ii), *memorylessness*, says that if the time before one’s alarm clock rings is exponentially distributed, then one’s beliefs about how long from now the clock will ring

do not depend on how long one has already been waiting. Together, this property and equation (10.6) above tell us that until the time when the clock rings, the conditional probability that it rings during the next dt times units is proportional to dt :

$$\mathbb{P}(T_i \leq t + dt | T_i \geq t) = \mathbb{P}(T_i \leq dt) \approx \lambda_i dt$$

The exponential distributions are the only continuous distributions with these properties.

Property (iii) says that given a collection of independent exponential alarm clocks, then the time until the first clock rings is itself exponentially distributed, the probability that a particular clock rings first is proportional to its rate, and the time until the first ring and the ringing clock's identity are independent random variables. These facts are essential to the workings of our stochastic evolutionary model.

Proof. Parts (i) and (ii) are easily verified. To establish part (iii), set $\lambda = \sum_{i=1}^n \lambda_i$, and compute the distribution of M_n as follows:

$$\mathbb{P}(M_n \geq t) = \mathbb{P}\left(\bigcap_{i=1}^n \{T_i \geq t\}\right) = \prod_{i=1}^n \mathbb{P}(T_i \geq t) = \prod_{i=1}^n e^{-\lambda_i t} = e^{-\lambda t}.$$

To prove the remaining claims from part (iii), observe that

$$\begin{aligned} (10.7) \quad \mathbb{P}\left(\bigcap_{j \neq i} \{T_j \leq T_i\} \cap \{T_i \geq t\}\right) &= \int_t^\infty \left(\prod_{j \neq i} \int_{t_i}^\infty \lambda_j e^{-\lambda_j t_j} dt_j\right) \lambda_i e^{-\lambda_i t_i} dt_i \\ &= \int_t^\infty \left(\prod_{j \neq i} e^{-\lambda_j t_i}\right) \lambda_i e^{-\lambda_i t_i} dt_i \\ &= \int_t^\infty \lambda_i e^{-\lambda t_i} dt_i \\ &= (\lambda_i / \lambda) e^{-\lambda t}. \end{aligned}$$

Setting $t = 0$ in equation (10.7) yields $\mathbb{P}(I_n = i) = \lambda_i / \lambda$, and an arbitrary choice of t shows that $\mathbb{P}(M_n \geq t, I_n = i) = \mathbb{P}(M_n \geq t) \mathbb{P}(I_n = i)$. ■

A random variable R has a *Poisson distribution with rate λ* , denoted $R \sim \text{Poisson}(\lambda)$, if $\mathbb{P}(R = r) = e^{-\lambda} \lambda^r / r!$ for all $r \in \{0, 1, 2, \dots\}$. Poisson random variables are used to model the number of occurrences of rare events (see Propositions 10.A.3 and 10.A.4). Two of their basic properties are listed below.

Proposition 10.A.2. *If R_1, \dots, R_n are independent with $R_i \sim \text{Poisson}(\lambda_i)$, then*

- (i) $\mathbb{E}(R_i) = \lambda_i$;
- (ii) $\sum_{j=1}^n R_j \sim \text{Poisson}(\sum_{j=1}^n \lambda_j)$.




Figure 10.A.1: Ring times S_n and numbers of rings R_t of an exponential alarm clock.

Proof. (i) $\mathbb{E}(R_i) = \sum_{r=1}^{\infty} r e^{-\lambda_i} \frac{(\lambda_i)^r}{r!} = \sum_{r=1}^{\infty} \lambda_i e^{-\lambda_i} \frac{(\lambda_i)^{r-1}}{(r-1)!} = \lambda_i \sum_{s=0}^{\infty} e^{-\lambda_i} \frac{(\lambda_i)^s}{s!} = \lambda_i.$

(ii) When $n = 2$, we can compute that

$$\begin{aligned} \mathbb{P}(R_1 + R_2 = r) &= \sum_{r_1=0}^r \mathbb{P}(R_1 = r_1) \mathbb{P}(R_2 = r - r_1) = \sum_{r_1=0}^r e^{-\lambda_1} \frac{(\lambda_1)^{r_1}}{r_1!} e^{-\lambda_2} \frac{(\lambda_2)^{r-r_1}}{(r-r_1)!} \\ &= e^{-(\lambda_1+\lambda_2)} \sum_{r_1=0}^r \frac{(\lambda_1)^{r_1} (\lambda_2)^{r-r_1}}{r_1! (r-r_1)!} = e^{-(\lambda_1+\lambda_2)} \frac{(\lambda_1 + \lambda_2)^r}{r!}, \end{aligned}$$

where the final equality follows from the binomial expansion

$$(\lambda_1 + \lambda_2)^r = \sum_{r_1=0}^r \frac{r!}{r_1! (r-r_1)!} (\lambda_1)^{r_1} (\lambda_2)^{r-r_1}.$$

Iterating yields the general result. ■

The exponential and Poisson distributions are fundamentally linked. Let $\{T_i\}_{i=1}^{\infty}$ be a sequence of i.i.d. *exponential*(λ) random variables. We can interpret T_1 as the first time that an exponential alarm clock rings, T_2 as the interval between the first and second rings, and T_k as the interval between the $(k-1)$ st and k th rings. In this interpretation, the sum $S_n = \sum_{k=1}^n T_k$ represents the time of the n th ring, while $R_t = \max\{n : S_n \leq t\}$ represents the number of rings through time t . Figure 10.A.1 presents a single realization of the ring time sequence $\{S_n\}_{n=1}^{\infty}$ and the number-of-rings process $\{R_t\}_{t \geq 0}$.

Proposition 10.A.3 derives the distribution of R_t , establishing a key connection between the exponential and Poisson distributions.

Proposition 10.A.3. $R_t \sim \text{Poisson}(\lambda t)$.

Proof. To begin, we prove that S_n has density

$$(10.8) \quad f_n(t) = \lambda e^{-\lambda t} \frac{(\lambda t)^{n-1}}{(n-1)!}.$$

This formula is obviously correct when $n = 1$. Suppose it is true for some arbitrary n . Then using the convolution formula $f_{X+Y}(z) = \int_{-\infty}^{\infty} f_Y(z-x)f_X(x) dx$, we find that

$$f_{n+1}(t) = \int_0^t f_n(t-s)f_1(s)ds = \int_0^t \lambda \frac{(\lambda(t-s))^{n-1}}{(n-1)!} e^{-\lambda(t-s)} \times \lambda e^{-\lambda s} ds = \lambda e^{-\lambda t} \frac{(\lambda t)^n}{n!}.$$

Next, we show that this equation implies that S_n has cumulative distribution

$$\mathbb{P}(S_n \leq t) = \sum_{m=n}^{\infty} e^{-\lambda t} \frac{(\lambda t)^m}{m!}.$$

Since $\sum_{m=0}^{\infty} \frac{(\lambda t)^m}{m!} = e^{\lambda t}$, this statement is equivalent to

$$\mathbb{P}(S_n \leq t) = 1 - \sum_{m=0}^{n-1} e^{-\lambda t} \frac{(\lambda t)^m}{m!}.$$

Differentiating shows that this expression is in turn equivalent to the density of S_n taking form (10.8), as established above.

To complete the proof, we express the event that at least n rings have occurred by time t in two equivalent ways: $\{R_t \geq n\} = \{S_n \leq t\}$. This observation and the expression for $\mathbb{P}(S_n \leq t)$ above imply that

$$\mathbb{P}(R_t = n) = \mathbb{P}(R_t \geq n) - \mathbb{P}(R_t \geq n+1) = \mathbb{P}(S_n \leq t) - \mathbb{P}(S_{n+1} \leq t) = e^{-\lambda t} \frac{(\lambda t)^n}{n!}. \blacksquare$$

10.A.2 The Poisson Limit Theorem

Proposition 10.A.3 shows that the Poisson distribution describes the number of rings of an exponential alarm clock during a fixed time span. We now establish a discrete analogue of this result.

The random variable X^p has a *Bernoulli distribution* with parameter $p \in [0, 1]$, denoted $X^p \sim \text{Bernoulli}(p)$, if $\mathbb{P}(X^p = 1) = p$ and $\mathbb{P}(X^p = 0) = 1 - p$. Let $\{X_i^p\}_{i=1}^n$ be a sequence of i.i.d. *Bernoulli*(p) random variables (e.g. coin tosses), and let $S_n^p = \sum_{i=1}^n X_i^p$ denote their sum

(the number of heads in n tosses). Then S_n^p has a *binomial distribution* with parameters n and p ($S_n^p \sim \text{binomial}(n, p)$):

$$\mathbb{P}(S_n^p = s) = \binom{n}{s} p^s (1-p)^{n-s} \text{ for all } s \in \{0, 1, \dots, n\}.$$

Finally, the random variable Z has a *standard normal distribution* ($Z \sim N(0, 1)$) if its density function is $f(z) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{z^2}{2})$.

Proposition 10.A.4 considers the behavior of the binomial random variables S_n^p when the number of tosses n becomes large. Recall that the sequence of random variables $\{Y_n\}_{n=1}^\infty$ with distribution functions $\{F_n\}_{n=1}^\infty$ converges in distribution (or converges weakly) to the random variable Y with distribution function F (denoted $Y_n \Rightarrow Y$, or $F_n \Rightarrow F$) if $\lim_{n \rightarrow \infty} F_n(x) = F(x)$ at all points $x \in \mathbf{R}$ at which F is continuous.

Proposition 10.A.4. Let $S_n^p \sim \text{binomial}(n, p)$. Then as $n \rightarrow \infty$,

- (i) $\frac{S_n^p - np}{\sqrt{np(1-p)}} \Rightarrow Z$, where $Z \sim N(0, 1)$.
- (ii) $S_n^{\lambda/n} \Rightarrow R^\lambda$ where $R^\lambda \sim \text{Poisson}(\lambda)$.

If we increase the number of tosses n of a coin whose bias p is fixed, the *Central Limit Theorem* tells us that the distribution of the number of heads S_n^p approaches a normal distribution. (In statement (i), we subtract the mean $\mathbb{E}S_n^p = np$ off of S_n^p and then divide by the standard deviation $SD(S_n^p) = \sqrt{np(1-p)}$ to obtain convergence to a fixed distribution.)

Suppose instead that as we increase the number of tosses n , we decrease the probability of heads p in such a way that the expected number of heads $np = \lambda$ remains fixed. Then statement (ii), the *Poisson Limit Theorem*, tells us that the distribution of S_n^p approaches a Poisson distribution. The basic calculation needed to prove this is as follows:

$$\begin{aligned} \mathbb{P}(S_n^{\lambda/n} = s) &= \frac{n!}{s!(n-s)!} \left(\frac{\lambda}{n}\right)^s \left(1 - \frac{\lambda}{n}\right)^{n-s} = \left(1 - \frac{\lambda}{n}\right)^n \frac{\lambda^s}{s!} \times \left(1 - \frac{\lambda}{n}\right)^{-s} \frac{n!}{(n-s)! n^s} \\ &= P(R^\lambda = s) \times \frac{\left(1 - \frac{\lambda}{n}\right)^n}{e^{-\lambda}} \times \prod_{r=0}^{s-1} \frac{n-r}{n-\lambda} \rightarrow \mathbb{P}(R^\lambda = s) \text{ as } n \rightarrow \infty. \end{aligned}$$

The second term of the penultimate expression above is independent of s and is less than 1 (because $(1 - \frac{\lambda}{n})^n$ increases to $e^{-\lambda}$), while the final term attains its maximum over s when $s = \lfloor \lambda + 1 \rfloor$ and decreases to 1 as n grows large. Together, these observations yield the following upper bound, which is needed in Exercise 10.3.3.

Proposition 10.A.5. $\mathbb{P}(S_n^{\lambda/n} = s) \leq C^\lambda \mathbb{P}(R^\lambda = s)$ for some $C^\lambda \in \mathbf{R}$ independent of n and s .

10.B Probability Models and their Interpretation

Appendix 10.C describes the construction and some basic properties of continuous time Markov processes on countable state spaces, our main modeling tool in the remainder of the book. To prepare for this, we now give a brief review of some basic concepts from probability theory, including some comments on the interpretation of a few of the fundamental theorems about i.i.d. random variables.

10.B.1 Countable Probability Models

We begin our review of probability theory by discussing probability models with a countable sample space. A *countable probability model* is a pair (Ω, \mathbb{P}) , where the *sample space* Ω is a finite or countable set, 2^Ω is the set of all *events* (i.e., subsets of Ω), and $\mathbb{P} : 2^\Omega \rightarrow [0, 1]$ is a *probability measure*: that is, a function satisfying $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(\Omega) = 1$, and *countable additivity*: if $\{A_k\}$ is a finite or countable collection of disjoint *events* (i.e., subsets of Ω), then $\mathbb{P}(\bigcup_k A_k) = \sum_k \mathbb{P}(A_k)$.

A *random variable* X is a function whose domain is Ω . The *distribution* of X is defined by $\mathbb{P}(X \in B) = \mathbb{P}(\omega \in \Omega : X(\omega) \in B)$ for all subsets B of the range of X . To define a finite collection of discrete random variables $\{X_k\}_{k=1}^n$, we specify a probability model (Ω, \mathbb{P}) and then define the random variables as functions on Ω . To interpret this construction, imagine picking an ω at random from the sample space Ω according to the probability distribution \mathbb{P} . The value of ω so selected determines the realizations $X_1(\omega), X_2(\omega), \dots, X_n(\omega)$ of the entire sequence of random variables X_1, X_2, \dots, X_n .

Frequently one is interested in studying a sequence of random variables $\{X_k\}_{k=1}^n$ such that each X_k takes values in the same countable set R . To construct such a sequence, one can let the sample space be the set of n -vectors $\Omega = R^n$, with typical element $\omega = (\omega_1, \dots, \omega_n)$. The random variables X_k can then be defined as *coordinate functions*: $X_k(\omega) = \omega_k$ for all $\omega \in \Omega$ and $k \in \{1, \dots, n\}$. By choosing the probability measure \mathbb{P} on Ω , one determines the joint distribution of the sequence $\{X_k\}_{k=1}^n$.

Example 10.B.1. Repeated rolls of a fair die. Suppose we would like to construct a sequence of random variables $\{X_k\}_{k=1}^n$, where X_k is to represent the k th roll of a fair die. To accomplish this, we let $R = \{1, 2, 3, 4, 5, 6\}$ be the set of possible results of an individual roll, and let $\Omega = R^n$. To define the probability measure \mathbb{P} , it is enough to let

$$(10.9) \quad \mathbb{P}(\{\omega\}) = \left(\frac{1}{6}\right)^n \text{ for all } \omega \in \Omega;$$

additivity then determines the probabilities of all other events in 2^Ω . Since

$$\mathbb{P}(X_k = x_k) = \mathbb{P}(\omega \in \Omega : X_k(\omega) = x_k) = \mathbb{P}(\omega \in \Omega : \omega_k = x_k) = \frac{1}{6}$$

for all $x_k \in R$, the random variables X_k have the correct marginal distributions. Moreover, if $A_k \subseteq R$ for $k \in \{1, \dots, n\}$, it is easy to confirm that

$$\mathbb{P}\left(\bigcap_{k=1}^n \{X_k \in A_k\}\right) = \prod_{k=1}^n \mathbb{P}(X_k \in A_k),$$

so the X_k are *independent*, as desired. §

Example 10.B.2. Repeated rolls of a die with uncertain bias. To construct a sequence of dependent random variables, one needs to choose a probability measure \mathbb{P} that is not of the product form (10.9). Suppose that it is equally likely that a die is fair, or that 3s and 4s occur four times as frequently as the other outcomes. To model this situation, one uses the sample space $\Omega = R^n$ above, but replaces the probability measure defined by (10.9) with

$$\mathbb{P}(\{\omega\}) = \frac{1}{2} \cdot \left(\frac{1}{6}\right)^n + \frac{1}{2} \cdot \left(\left(\frac{1}{3}\right)^{N(\omega)} \left(\frac{1}{12}\right)^{n-N(\omega)}\right),$$

where $N(\omega) = \sum_k 1_{\omega_k \in \{3,4\}}$ represents the number of 3s and 4s in the sequence ω . §

The *expected value* of a random variable is its integral with respect to the probability measure \mathbb{P} . In the case of the k th roll of a fair die,

$$\mathbb{E}X_k = \int_{\Omega} X_k(\omega) d\mathbb{P}(\omega) = \sum_{\omega \in \Omega} \omega_k \mathbb{P}(\{\omega\}) = \sum_{\omega_k \in R} \omega_k \left(\sum_{\omega_{-k}} \mathbb{P}(\{(\omega_k, \omega_{-k})\}) \right) = \sum_{i=1}^6 i \times \frac{1}{6} = 3\frac{1}{2}.$$

We can create new random variables out of old ones using functional operations. For instance, the total of the results of the n die rolls is a new random variable S_n defined by $S_n = \sum_{k=1}^n X_k$, or, more explicitly, by $S_n(\omega) = \sum_{k=1}^n X_k(\omega)$ for all $\omega \in \Omega$.

10.B.2 Uncountable Probability Models and Measure Theory

While the constructions above are sufficient for finite collections of discrete random variables, they do not suffice when individual random variables take an uncountable number of values, or when we are interested in infinite numbers of random variables. To handle these situations, we need the sample space Ω to be *uncountable*: that is, not expressible as a sequence of elements.

Unfortunately, uncountable sample spaces introduce a serious new technical difficulty. As an illustration, suppose we want to construct a random variable representing a uniform draw from the unit interval. It is natural to choose $\Omega = [0, 1]$ as our sample space and to define our random variable as the identity function on Ω : that is, $X(\omega) = \omega$. But then we encounter a major difficulty: it is impossible to define a countably additive probability measure \mathbb{P} that specifies the probability of *every* subset of Ω .

To resolve this problem, one chooses a set of subsets $\mathcal{F} \subseteq 2^\Omega$ whose probabilities will be specified, and then introduces corresponding restrictions on the definition of a random variable. A random variable satisfying these restrictions is said to be *measurable*, and this general approach to studying functions defined on uncountable domains is known as *measure theory*.

In summary, an *uncountable probability model* consists of a triple $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a sample space, $\mathcal{F} \subseteq 2^\Omega$ is a collection (more specifically, a σ -algebra) of subsets of Ω , and $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ is a countably additive probability measure.

Example 10.B.3. Defining a uniform random variable. To define a random variable representing a uniform draw from the unit interval, we can let $\Omega = (0, 1)$, let \mathcal{F} be the set of Borel sets in $(0, 1)$ (which we will not define here), and let $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ be *Lebesgue measure* on Ω ; \mathbb{P} is the unique measure on (Ω, \mathcal{F}) that agrees with the usual notion of length, in the sense that satisfying $\mathbb{P}((a, b]) = b - a$ for every interval $(a, b] \subset [0, 1]$. If we define $U : \Omega \rightarrow (0, 1)$ by $U(\omega) = \omega$, then U is a uniform random variable on the unit interval.

The standard proof of the existence and uniqueness of Lebesgue measure is by means of a result from measure theory called the *Carathéodory Extension Theorem*. In the present context, this theorem shows that there is a unique way of extending \mathbb{P} from the collection of half-open intervals $(a, b] \subset (0, 1)$ to the collection of all Borel sets in $(0, 1)$. §

Example 10.B.4. Defining a random variable with a given cumulative distribution function. The function F is a *cumulative distribution function* if it is nondecreasing, right continuous, and satisfies $\lim_{x \rightarrow -\infty} F(x) = 0$ and $\lim_{x \rightarrow \infty} F(x) = 1$. How can we define a random variable X which has distribution F , meaning that $\mathbb{P}(X \leq x) = F(x)$ for all $x \in \mathbf{R}$?

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the probability model from the previous example, and define $X(\omega) = \sup\{y : F(y) < \omega\}$; roughly speaking, $X : (0, 1) \rightarrow \mathbf{R}$ is the inverse of $F : \mathbf{R} \rightarrow [0, 1]$. Using the right continuity of F , one can verify that $\{\omega \in \Omega : X(\omega) \leq x\} = \{\omega \in \Omega : \omega \leq F(x)\}$. Since the latter set is a Borel set, these observations imply that $\mathbb{P}(X \leq x) = \mathbb{P}(\omega \leq F(x)) = F(x)$. §

More generally, suppose we would like to define a (possibly infinite) collection of random variables described by some prespecified joint distributions. How do we know whether it is possible to construct these random variables on some well-chosen probability

space? As long as the marginal and joint distributions satisfy certain obviously necessary consistency conditions, the *Kolmogorov Extension Theorem* ensures the existence of a probability model on which one can define random variables with the given distributions.

10.B.3 Distributional Properties and Sample Path Properties

The reader may wonder why we bother with the explicit construction of random variables. After all, once we specify the joint distributions of the basic random variables of interest, we also determine the joint distributions of any random variables that can be derived from our original collection. Why not work entirely in terms of these distributions and avoid the explicit construction of the random variables altogether?

If we are only interested in *distributional properties* of our random variables, explicit construction of the random variables is not essential. However, many key results in probability theory concern not the distributional properties of random variables, but rather their *sample path properties*. These are properties of realization sequences: i.e., the sequences of values $X_1(\omega), X_2(\omega), X_3(\omega), \dots$ that arise for each choice of $\omega \in \Omega$. We can begin to illustrate the differences between the two sorts of properties through a simple example.

Example 10.B.5. Consider the probability model (Ω, \mathbb{P}) with sample space $\Omega = \{-1, 1\}$ and probability measure $\mathbb{P}(\{-1\}) = \mathbb{P}(\{1\}) = \frac{1}{2}$. Define the sequences of random variables $\{X_t\}_{t=1}^\infty$ and $\{\hat{X}_t\}_{t=1}^\infty$ as follows:

$$X_t(\omega) = \omega;$$

$$\hat{X}_t(\omega) = \begin{cases} -\omega & \text{if } t \text{ is odd,} \\ \omega & \text{if } t \text{ is even.} \end{cases}$$

If we look only at time t marginal distributions, $\{X_t\}_{t=1}^\infty$ and $\{\hat{X}_t\}_{t=1}^\infty$ look identical, as both sequences consist of random variables equally likely to have realizations -1 and 1 . But from the sample path point of view, the two sequences are different: for either choice of ω , the sequence $\{X_t(\omega)\}_{t=1}^\infty$ is constant, while the sequence $\{\hat{X}_t(\omega)\}_{t=1}^\infty$ alternates between 1 and -1 forever.

We illustrate these ideas in Figure 10.B.1, which provides graphical representations of our two sequences of random variables. In these pictures, the vertical axis represents the sample space Ω , the horizontal axis represents indices (or “times”) of the trials, and the interiors of the figures contain the realizations $X_t(\omega)$ and $\hat{X}_t(\omega)$.

To focus on properties of the time t marginal distributions of a sequence of random

variables, we look at the *collection* of outcomes in each *vertical* section (Figure 10.B.1(i)). In this respect, each X_t is identical to its partner \hat{X}_t , and in fact all of the random variables in both sequences share the same distribution.

It is possible to distinguish between the two sequences from a distributional point of view, however, by considering their joint distributions at pairs of consecutive times t and $t + 1$. From this vantage, we see that X_t and X_{t+1} are perfectly correlated, while \hat{X}_t and \hat{X}_{t+1} are perfectly negatively correlated. This is a statement about the collection of outcome pairs in vertical sections of width 2 (Figure 10.B.1(ii)).

Finally, to focus on sample path properties, we look at the *sequences* of outcomes in each *horizontal* slice of each picture (Figure 10.B.1(iii)). Doing so, we see that for each ω , the sample path $\{X_t(\omega)\}_{t=1}^{\infty}$ is quite different from the sample path $\{\hat{X}_t(\omega)\}_{t=1}^{\infty}$: the former is constant, while the later alternates forever. §

Intuitively, distributional properties are ones that are only meaningful from an *ex ante* point of view, before the realization of ω (and hence of the random variables) is known. Theorems on distributional properties constrain the probabilities of certain events, often in the limit as some parameter (e.g., the number of trials) grows large. These probability assignments are only relevant before the values of the random variables have been observed.

For their part, theorems on sample path properties typically state that with probability one, the infinite sequence of realizations of a process must satisfy certain properties. These theorems can be interpreted as *ex post* statements about the random variables, since they provide information about the infinite sequence of realizations $\{X_t(\omega)\}_{t=1}^{\infty}$ that we actually observe. (To be precise, this is “only” true for a set of ω s that has probability one: we cannot completely avoid referring to the *ex ante* point of view.)

Example 10.B.6. Properties of i.i.d. random variables. The distinction between distributional properties and sample path properties is important to understand the fundamental theorems about i.i.d. random variables. Let $\{X_t\}_{t=1}^{\infty}$ be a sequence of i.i.d. random variables, each of which is a function on the (uncountable) probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For simplicity, assume that each X_t has mean zero and variance one. Then the sum $S_n = \sum_{t=1}^n X_t$ has mean zero and variance n , while the sample average $\bar{X}_n = S_n/n$ has mean zero and variance $\frac{1}{n}$.

The laws of large numbers concern the convergence of the sample averages \bar{X}_n as the number of trials n grows large.

The Weak Law of Large Numbers: For all $\varepsilon > 0$, $\lim_{n \rightarrow \infty} \mathbb{P}(\omega \in \Omega : \bar{X}_n(\omega) \in [-\varepsilon, \varepsilon]) = 1$.

The Strong Law of Large Numbers: $\mathbb{P}\left(\omega \in \Omega : \lim_{n \rightarrow \infty} \bar{X}_n(\omega) = 0\right) = 1.$

The WLLN is a distributional result. It tells us that for any $\varepsilon > 0$ and $\delta > 0$, we can find a number N with this property: if the sample size n is at least N , the probability we assign to the event $\bar{X}_n \in [-\varepsilon, \varepsilon]$ is at least $1 - \delta$. The SLLN is a sample path result: with probability one, the sequence of realizations of the sample average that we actually observe, $\{\bar{X}_n(\omega)\}_{n=1}^{\infty}$, is a sequence of numbers that converges to zero. (To build intuition, it may be helpful to describe both of these results using pictures like those in Figure 10.B.1.)

While the WLLN can be stated directly in terms of distributions, the SLLN only makes sense if our random variables are defined as functions on a probability space. Also, the distributional and sample path results can be distinguished by the order in which the probability \mathbb{P} and the limit $\lim_{n \rightarrow \infty}$ appear: the WLLN concerns a “limit of probabilities”, while the SLLN concerns a “probability of a limit”.

The SLLN is genuinely stronger than the WLLN, in that the conclusion of the former implies the conclusion of the latter. In general, though, distributional and sample path results address distinct aspects of a process’s behavior. Consider the following pair of results, which focus on variation:

The Central Limit Theorem: $\lim_{n \rightarrow \infty} \mathbb{P}\left(\omega \in \Omega : \frac{S_n(\omega)}{\sqrt{n}} \in [a, b]\right) = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-x^2/2} dx.$

The Law of the Iterated Logarithm: $\mathbb{P}\left(\omega \in \Omega : \limsup_{n \rightarrow \infty} \frac{S_n(\omega)}{\sqrt{2n \log \log n}} = 1\right) = 1.$

The CLT is a distributional result: as n goes to infinity, the distributions of the normalized sums S_n/\sqrt{n} converge to the standard normal distribution. This theorem allows us to approximate *ex ante* the probability with which the random variable S_n/\sqrt{n} will take a value in any given interval $[a, b]$. The LIL looks instead looks at variation within individual sample paths. It tells us that with probability one, the sequence of realizations $\{S_n(\omega)\}_{n=1}^{\infty}$ that we actually observe exceeds $(1 - \varepsilon) \sqrt{2n \log \log n}$ infinitely often, but exceeds $(1 + \varepsilon) \sqrt{2n \log \log n}$ only finitely often. §

Theorems 11.A.19 and 11.A.22 in Chapter 11 are generalizations of the WLLN and the SLLN for Markov processes. These results are the key to understanding the infinite horizon behavior of our stochastic evolutionary process.

While so far we have focused on sequences of random variables, one should also be cognizant of the distinction between distributional and sample path properties when we studying limits of sequences of stochastic processes, as we do in Section 10.2.

Example 10.B.7. Kurtz’s Theorem (Theorem 10.2.1) provides conditions ensuring that over an initial finite time span $[0, T]$, the sequence of Markov processes $\{\{X_t^N\}_{N=N_0}^\infty\}$ converges to a deterministic trajectory $\{x_t\}$, in the sense that

$$(10.10) \quad \lim_{N \rightarrow \infty} \mathbb{P} \left(\omega \in \Omega : \sup_{t \in [0, T]} |X_t^N(\omega) - x_t| < \varepsilon \right) = 1.$$

While (10.10) might first appear to be sample path property, it is better regarded as a distributional property. It is a statement about a limit of probabilities, where the “random variables” $\{X_t^N\}_{t \in [0, T]}$ whose distributions are at issue are infinite-dimensional, taking values in the set of step functions from the time interval $[0, T]$ to the simplex X . §

10.C Countable State Markov Chains and Processes

10.C.1 Countable State Markov Chains

Markov chains and Markov processes are collections of random variables $\{X_t\}_{t \in T}$ with the property that “the future only depends on the past through the present”. We focus on settings where these random variables take values in some finite or countable state space \mathcal{X} . (Even if the *state space* \mathcal{X} is countable, the random variables $X_t : \Omega \rightarrow \mathcal{X}$ must be defined on a probability model with an uncountable *sample space* Ω if the set of times T is infinite: whenever \mathcal{X} has at least two elements, the set of infinite sequences (x_1, x_2, \dots) with $x_t \in \mathcal{X}$ is an uncountable set.) We use the terms “Markov chain” and “Markov process” to distinguish between the discrete time ($T = \{0, 1, \dots\}$) and continuous time ($T = [0, \infty)$) frameworks. (Some authors use these terms to distinguish between discrete and continuous state spaces.)

The sequence of random variables $\{X_t\} = \{X_t\}_{t=0}^\infty$ is a *Markov chain* if it satisfies the *Markov property*:

$$\mathbb{P}(X_{t+1} = x_{t+1} | X_0 = x_0, \dots, X_t = x_t) = \mathbb{P}(X_{t+1} = x_{t+1} | X_t = x_t)$$

for all times $t \in \{0, 1, \dots\}$ and all collections of states $x_0, \dots, x_{t+1} \in \mathcal{X}$ for which the conditional expectations are well defined. We only consider *temporally homogeneous* Markov chains, which are Markov chains whose one-step transition probabilities are independent of time:

$$\mathbb{P}(X_{t+1} = y | X_t = x) = P_{xy}.$$

We call the matrix $P \in \mathbf{R}_+^{X \times X}$ the *transition matrix* for the Markov chain $\{X_t\}$. The vector $\pi \in \mathbf{R}_+^X$ defined by $\mathbb{P}(X_0 = x) = \pi_x$ is the *initial distribution* of $\{X_t\}$; when π puts all of its mass on a single state x_0 , we call x_0 the *initial condition* or the *initial state*. The vector π and the matrix P fully determine the joint distributions of $\{X_t\}$ via

$$\mathbb{P}(X_0 = x_0, \dots, X_t = x_t) = \pi_{x_0} \prod_{s=1}^t P_{x_{s-1}x_s}.$$

Since certain properties of Markov chains do not depend on the initial distribution π , it is sometimes left unspecified.

By definition, the one-step transition probabilities of the Markov chain $\{X_t\}$ are the elements of the matrix P :

$$\mathbb{P}(X_1 = y | X_0 = x) = P_{xy}.$$

The two-step transition probabilities of $\{X_t\}$ are obtained by multiplying P by itself

$$\begin{aligned} \mathbb{P}(X_2 = y | X_0 = x) &= \sum_{z \in X} \mathbb{P}(X_2 = y, X_1 = z | X_0 = x) \\ &= \sum_{z \in X} \mathbb{P}(X_1 = z | X_0 = x) \mathbb{P}(X_2 = y | X_1 = z, X_0 = x) \\ &= \sum_{z \in X} P_{xz} P_{zy} \\ &= (P^2)_{xy}. \end{aligned}$$

Continuing inductively, we conclude that the the t step transition probabilities of $\{X_t\}$ are given by the entries of the t th power of the transition matrix:

$$\mathbb{P}(X_t = y | X_0 = x) = (P^t)_{xy}.$$

10.C.2 Countable State Markov Processes: Definition and Construction

A (*temporally homogeneous*) Markov process on the countable state space X is a collection of random variables $\{X_t\} = \{X_t\}_{t \geq 0}$ with continuous time index t . This collection must satisfy the following three properties:

(MP) *The (continuous time) Markov property:*

$$\mathbb{P}(X_{t_{k+1}} = x_{t_{k+1}} | X_{t_0} = x_{t_0}, \dots, X_{t_k} = x_{t_k}) = \mathbb{P}(X_{t_{k+1}} = x_{t_{k+1}} | X_{t_k} = x_{t_k}) \text{ for all } 0 \leq t_0 < \dots < t_{k+1} \text{ and } x_{t_0}, \dots, x_{t_{k+1}} \in X \text{ with } \mathbb{P}(X_{t_0} = x_{t_0}, \dots, X_{t_k} = x_{t_k}) > 0.$$

(TH) *Temporal homogeneity:*

$$\mathbb{P}(X_{t+u} = y | X_t = x) = P_{xy}(u) \text{ for all } t, u \geq 0.$$

(RCLL) *Right continuity and left limits:*

For every $\omega \in \Omega$, the sample path $\{X_t(\omega)\}_{t \geq 0}$ is continuous from the right and has left limits. That is, $\lim_{s \downarrow t} X_s(\omega) = X_t(\omega)$ for all $t \in [0, \infty)$, and $\lim_{s \uparrow t} X_s(\omega)$ exists for all $t \in (0, \infty)$.

While conditions (MP) and (TH) are restrictions on the (joint) distributions of $\{X_t\}$, condition (RCLL) is a restriction on the sample paths of $\{X_t\}$.

There are two reasons for introducing restrictions on the behavior of sample paths. First, when we use Markov processes as models, restrictions like (RCLL) are very natural requirements for the sample paths of the process. Second, (RCLL) is often needed to ensure that certain events whose probabilities we might want to assess are in fact measurable. The trouble arises here because the set of times $[0, \infty)$ is uncountable, while the probability measure \mathbb{P} is only satisfies countable additivity. Thus, if (RCLL) does not hold, sets like $Z = \{\omega \in \Omega : X_t(\omega) \neq 0 \text{ for some } t \geq 0\}$ need not be measurable. But if (RCLL) does hold, then Z is essentially determined by the behavior of the process $\{X_t\}$ at rational times $t \in \mathbf{Q}$; because \mathbf{Q} is countable, Z can be shown to be measurable.

Processes satisfying the distributional requirements (MP) and (TH) must take this form: there must be an initial distribution $\pi \in \mathbf{R}_+^X$, a *jump rate vector* $\lambda \in \mathbf{R}_+^X$, and a transition matrix $P \in \mathbf{R}_+^{X \times X}$ such that

- (i) The initial distribution of the process is given by $\mathbb{P}(X_0 = x) = \pi_x$.
- (ii) When the process is in state x , the random time before the next jump is exponentially distributed with rate λ_x .
- (iii) The state at which a jump from x lands follows the distribution $\{P_{xy}\}_{y \in X}$. (Note that the landing state can be x itself if $P_{xx} > 0$.)
- (iv) Times between jumps are independent of each other, and are also independent of the past conditional on the current state.

The objects π , λ , and P implicitly define the joint distributions of the random variables $\{X_t\}$, so the Kolmogorov Extension Theorem (Section 10.B.2) tells us that a collection of random variables with these joint distributions exists (i.e., can be defined as functions on some well chosen probability space). However, Kolmogorov's Theorem does not ensure that the random variables so constructed satisfy the sample path continuity property (RCLL).

Fortunately, it is not too difficult to construct the process $\{X_t\}$ explicitly. Let $\{Y_k\}_{k=0}^\infty$ be a discrete time Markov chain with initial distribution π and transition matrix P , and let $\{T_k\}_{k=1}^\infty$ be a sequence of i.i.d. *exponential*(1) random variables that are independent of the

Markov chain $\{Y_k\}$. (Since both of these collections are countable, questions of sample path continuity do not arise; the existence of these random variables as functions defined on a common probability space is ensured by Kolmogorov's Theorem.)

Define the random *jump times* $\{\tau_n\}_{n=0}^\infty$ by $\tau_0 = 0$ and

$$\tau_n = \sum_{k=1}^n \frac{T_k}{\lambda_{Y_{k-1}}}, \text{ so that } \tau_n - \tau_{n-1} = \frac{T_n}{\lambda_{Y_{n-1}}}.$$

Finally, define the process $\{X_t\}_{t \geq 0}$ by

$$X_t = Y_n \text{ when } t \in [\tau_n, \tau_{n+1}).$$

The process $\{X_t\}$ begins at some initial state $X_0 = Y_0 = y_0$. It remains there for the random duration $\tau_1 \sim \text{exponential}(\lambda_{y_0})$, at which point a transition to some new state $X_{\tau_1} = Y_1 = y_1$ occurs; the process then remains at y_1 for the random duration $\tau_2 - \tau_1 \sim \text{exponential}(\lambda_{y_1})$, at which point a transition to $X_{\tau_2} = Y_2 = y_2$ occurs; and so on. By construction, the sample paths of $\{X_t\}$ are right continuous with left limits, and it is easy to check that the joint distributions of $\{X_t\}$ are the ones we desire.

Example 10.C.1. The Poisson Process. Consider a Markov process $\{X_t\}$ with state space $\mathcal{X} = \mathbf{Z}_+$, initial condition $X_0 = 0$, jump rates $\lambda_x = \lambda > 0$ for all $x \in \mathcal{X}$, and transition matrix $P_{xy} = 1_{\{y=x+1\}}$ for all $x, y \in \mathcal{X}$. Under this process, jumps arrive randomly at the fixed rate λ , and every jump increases the state by exactly one unit. A Markov process fitting this description is called a *Poisson process*.

By the definition of this process,

- (P1) The waiting times $\tau_n - \tau_{n-1}$ are i.i.d. with $\tau_n - \tau_{n-1} \sim \text{exponential}(\lambda)$
($n \in \{1, 2, \dots\}$).

In fact, it can be shown that under the sample path continuity condition (RCLL) and other mild technical requirements, condition (P1) is equivalent to

- (P2) The increments $X_{t_k} - X_{t_{k-1}}$ are independent random variables,
and $(X_{t_k} - X_{t_{k-1}}) \sim \text{Poisson}(\lambda(t_k - t_{k-1}))$ ($0 < t_1 < \dots < t_k$).

Proposition 10.A.3 established part of this result: it showed that if condition (P1) holds, then $X_t \sim \text{Poisson}(\lambda t)$ for all $t > 0$. But the present result says much more: a “pure birth process” whose waiting times are i.i.d. exponentials is not only Poisson distributed at each time t ; in fact, all *increments* of the process are Poisson distributed, and nonoverlapping increments are stochastically independent. Conversely, if one begins with the assumption

that the increments of the process are independent and Poisson, then the waits between jumps must be i.i.d. and exponential. §

10.C.3 Countable State Markov Processes: Transition Probabilities

The time t transition probabilities of a countable state Markov process can be expressed in an appealingly simple form. To avoid the use of infinite-dimensional matrices, we focus here on the case of a finite state space \mathcal{X} ; however, versions of the results below also hold when \mathcal{X} is countably infinite.

Let $\{X_t\}$ be a Markov process with jump rates $\lambda \in \mathbf{R}_+^{\mathcal{X}}$ and transition matrix $P \in \mathbf{R}_+^{\mathcal{X} \times \mathcal{X}}$, and define

$$P_{xy}(t) = \mathbb{P}(X_t = y | X_0 = x),$$

to be the time t transition probability from state x to state y . (To prevent the notations for the transition matrix P and the time t transition probabilities $P(t)$ from overlapping, we always refer the collection of all of the latter as $\{P(t)\}_{t \geq 0}$.) It is clear that $P(0) = I$, the identity matrix. Note as well that since $\{X_t\}$ is a temporally homogeneous Markov process, we have that

$$\begin{aligned} P_{xy}(s+t) &= \mathbb{P}(X_{s+t} = y | X_0 = x) \\ &= \sum_{z \in \mathcal{X}} \mathbb{P}(X_s = z | X_0 = x) \mathbb{P}(X_{s+t} = y | X_s = z, X_0 = x) \\ &= \sum_{z \in \mathcal{X}} \mathbb{P}(X_s = z | X_0 = x) \mathbb{P}(X_{s+t} = y | X_s = z) \\ &= \sum_{z \in \mathcal{X}} \mathbb{P}(X_s = z | X_0 = x) \mathbb{P}(X_t = y | X_0 = z). \end{aligned}$$

Expressing this result in matrix form, we conclude that the matrix trajectory $\{P(t)\}_{t \geq 0}$ is a *semigroup*:

$$(10.11) \quad P(s+t) = P(s)P(t) \text{ for all } s, t \geq 0.$$

The key tool for studying the transition probabilities of a Markov process is its *generator*,

$$Q = \text{diag}(\lambda)(P - I) \in \mathbf{R}^{\mathcal{X} \times \mathcal{X}}.$$

Using the fact that the row sums of P are 1, we can express the entries of Q as

$$Q_{xy} = \begin{cases} \lambda_x P_{xy} & \text{if } x \neq y, \\ -\lambda_x \sum_{y \neq x} P_{xy} & \text{if } x = y. \end{cases}$$

Thus, for $x \neq y$, Q_{xy} represents the rate of transitions from x to y , while $-Q_{xx}$ represents the total rate of transitions away from x .

To put the generator to work, notice that when $\tau > 0$ is small, so that it is very unlikely that more than one transition has occurred by time τ , we have

$$P_{xy}(\tau) = \mathbb{P}(X_\tau = y | X_0 = x) \approx \begin{cases} \tau \lambda_x P_{xy} & \text{if } x \neq y, \\ 1 - \tau \lambda_x \sum_{z \neq x} P_{xz} & \text{if } x = y. \end{cases}$$

Expressing this in matrix notation, we obtain

$$(10.12) \quad P(\tau) \approx I + \tau Q.$$

This equation explains why Q is sometimes referred to as the *infinitesimal generator* of $\{X_t\}$.

Using the semigroup property (10.11) and expression (10.12), we can informally derive an exact expression for the transition probabilities of $\{X_t\}$. If we divide the time interval $[0, t]$ into n subintervals of length $\frac{t}{n}$, then applying (10.11) and (10.12) in turn yields

$$(10.13) \quad P(t) = P\left(\frac{1}{n}t\right)^n \approx \left(I + \frac{1}{n}Q\right)^n.$$

Taking n to infinity, the analogy with the scalar formula $\lim_{n \rightarrow \infty} \left(1 + \frac{q}{n}\right)^n = e^q$ suggests that the right hand side of equation (10.13) should converge to the matrix exponential e^{Qt} (see Appendix 8.B.2). And indeed, one can establish rigorously that the transition probabilities can be expressed as

$$(10.14) \quad P(t) = e^{Qt} \equiv \sum_{k=0}^{\infty} \frac{(Qt)^k}{k!}.$$

There are two ways of expressing (10.14) in differential form. Theorem 8.B.4 tells us that (10.14) is a solution of the *backward equation*, which is the matrix differential equation $\dot{P}(t) = QP(t)$. At the same time, since the matrices Q and $P(t) = e^{Qt}$ commute, we can rewrite the previous equation as $\dot{P}(t) = P(t)Q$. This is known as the *forward equation*. The trajectory $P(t) = e^{Qt}$ is the unique solution to each of these linear differential equations

from initial condition $P(0) = I$. In Section 11.A.7, we will use the backward equation to characterize stationary distributions of Markov processes.

10.N Notes

Sections 10.2 and 10.3. Theorems 10.2.1 and 10.3.1 first appeared in Kurtz (1970). See Ethier and Kurtz (1986, Chapter 11) for an advanced textbook treatment.

The first formal results in the game theory literature akin to Theorem 10.2.3 focus on specific revision protocols. Boylan (1995) shows how evolutionary processes based on random matching schemes converge to deterministic trajectories when the population size grows large. Binmore et al. (1995b), Börgers and Sarin (1997), and Schlag (1998) consider particular models of evolution that converge to the replicator dynamic. Binmore and Samuelson (1999) prove a general deterministic approximation result for discrete time models of evolution under a somewhat restrictive timing assumption. Sandholm (2003) uses Kurtz's Theorem to prove a general finite horizon convergence result. This paper also shows that after spatial normalization, the behavior of $\{X_t^N\}$ near rest points of the mean dynamic can be approximated by a diffusion. The strongest deterministic approximation results can be found in Benaïm and Weibull (2003, 2009). These authors establish an exponential bound on the probability of deviations of $\{X_t^N\}$ from solutions of the mean dynamic. They also establish results relating the infinite horizon behavior of $\{X_t^N\}$ to the mean dynamic; we introduce these results in Chapter 12.

While the results described above rely on the assumption that the mean dynamic is Lipschitz continuous, we conjecture that that analogous results can be established in more general settings—in particular, when the mean dynamic is not a differential equation at all, but rather a differential inclusion. For related results in a somewhat different context, see Benaïm et al. (2005).

While we have focused here on the evolution of the distribution of behavior, Tanabe (2006) proves results about the evolution of the strategy profile: i.e., about the joint distribution of individual agents' choice trajectories. Suppose that at time 0, the N agents' choices of strategies from S are i.i.d. Then as N grows large, each agent's random choice trajectory converges in distribution to ν , the distribution of a certain time-inhomogeneous Markov process—a so-called *McKean process*—taking values in S . Furthermore, the joint distribution of any k individuals' choice trajectories converges to the k -fold product of the measure ν . This means that the independence of the k individuals' choices at time 0 persists over any finite time span, a phenomenon sometimes called *propagation of chaos*; see Sznitman (1991). One can further show that the empirical distribution of the N agents'

choice trajectories also converges to the measure ν . Since the time t marginal of this empirical distribution is none other than our state variable X_t^N , Theorem 10.2.3 tells us that the collection of time t marginals of ν is none other than the solution to our mean dynamic (M).

Appendices 10.A, 10.B, and 10.C. Billingsley (1995) and Durrett (2005) are excellent graduate level probability texts. The former book provides more thorough coverage of the topics considered in this chapter, and contains an especially clear treatment of the Poisson process. Norris (1997), Brémaud (1999), and Stroock (2005) are all excellent books on Markov chains and Markov processes. The first of these is at an undergraduate level, the last at a graduate level, and the middle one somewhere in between.

Stationary Distributions and Infinite Horizon Behavior

11.0 Introduction

The central result of Chapter 10 established that over finite time spans, when the population size is sufficiently large, the stochastic evolutionary process $\{X_t^N\}$ follows a nearly deterministic path, closely shadowing a solution trajectory of the corresponding mean dynamic (M). But if we look at longer time spans—that is, if we fix the population size N of interest and consider the position of the process at large values of t —the random nature of the process must assert itself. In particular, we will soon see that if the process is generated by a *full support revision protocol*, one that always assigns positive probabilities to transitions to all neighboring states in \mathcal{X}^N , then $\{X_t^N\}$ must visit all states in \mathcal{X}^N infinitely often. Evidently, an *infinite horizon approximation theorem* along the lines of Theorem 10.2.3 cannot hold. To make predictions about play over very long time spans, we need new techniques for characterizing the infinite horizon behavior of the stochastic evolutionary process.

In finite horizon analyses the basic object of study is the mean dynamic (M), an ordinary differential equation derived from the Markov process $\{X_t^N\}$. In infinite horizon analyses, the corresponding object is the *stationary distribution* μ^N of the process $\{X_t^N\}$. A stationary distribution is defined by the property that a process whose initial condition is described by this distribution will continue to be described by this distribution at all future times. If $\{X_t^N\}$ is generated by a full support revision protocol, then its stationary distribution μ^N is not only unique, but also describes the infinite horizon behavior of $\{X_t^N\}$ regardless of this process's initial distribution. In principle, this fact allows us to use the stationary distribution to form predictions about a population's very long run behavior that do not

depend on its initial behavior. This contrasts sharply with predictions based on the mean dynamic (M), which generally require knowledge of the initial state.

We begin the formal development of these ideas in Section 11.1 by introducing full support revision protocols, as well as the related notion of *irreducibility* for the stochastic evolutionary process $\{X_t^N\}$. A typical feature of full support revision protocols is their inclusion of at least a small level of noise, which ensures that revising agents have at least some small probability of choosing each available strategy. Next, we define the stationary distribution μ^N of the process $\{X_t^N\}$, and we review results from probability theory linking this distribution to the infinite-horizon behavior of the process. Finally, we introduce a condition on Markov processes called *reversibility*, which requires that the process “look the same” whether it is run forward or backward in time.

When a Markov process is reversible, its stationary distribution often takes an especially simple form. Taking advantage of this fact, the remainder of the chapter focuses on the two settings in which the evolutionary process $\{X_t^N\}$ is known to be reversible: two-strategy games under arbitrary revision protocols, studied in Section 11.2, and potential games under exponential revision protocols, studied in Section 11.5, after some preparations in Section 11.4.

In principle, the stationary distribution μ^N could spread its mass over a wide range of states in the state space \mathcal{X}^N , in which case the prediction of play that it offers, while independent of initial behavior, would be rather diffuse. But the analyses in this chapter show that if the population size is not too small, and the amount of noise in agents’ decisions not too large, then μ^N will typically concentrate its mass on a single region in \mathcal{X}^N —for instance, the states corresponding to a neighborhood of a stable rest point of the relevant mean dynamic. It is in such cases that the infinite horizon analysis truly provides a unique prediction of play.

Summing up, we now have two approaches to forecasting behavior in population games: finite horizon predictions, based on the mean dynamic (M), depend on the initial population state; infinite horizon predictions, using the stationary distribution μ^N , are independent of the initial population state. In applications, the choice between these two approaches should depend on the time span of interest, with long enough time horizons pointing toward the infinite horizon analysis. But how long is “long enough”? We address this question in Section 11.3, where we investigate how much time is required for infinite horizon analyses to become useful for predictions. We find that if the population is even of moderate size, or if the level of noise in agents choices is even somewhat small, then the amounts of time needed before infinite horizon analysis will yield meaningful predictions can be of astronomical magnitudes. For this reason, we feel that in typical economic

applications, the history dependent predictions provided by the mean dynamic are most appropriate. Further discussion of this point, including descriptions of environments in which infinite horizon analysis may be apt, is offered in the text.

The analyses in this chapter are developed using the theory of finite-state Markov processes, particularly those tools related to stationary distributions and infinite horizon behavior of these processes. In Appendix 11.A, we offer a detailed presentation of the relevant mathematical techniques. It can be read as an independent unit, or used as a reference while working through the main text.

11.1 Irreducible Evolutionary Processes

11.1.1 Full Support Revision Protocols

Let us briefly review the construction of the stochastic evolutionary process $\{X_t^N\}$ presented in Section 10.1, again focusing on the single-population setting. A population of N agents recurrently plays the population game $F : X \rightarrow \mathbf{R}^n$. The agents are equipped with independent rate R Poisson alarm clocks, and employ the revision protocol $\rho : \mathbf{R}^n \times X \rightarrow \mathbf{R}_+^{n \times n}$. When an i player's clock rings, he switches to strategy $j \neq i$ with probability $\rho_{ij}(F(x), x)/R$.

This model defines a Markov process $\{X_t^N\}$ on the discrete state space $\mathcal{X}^N = X \cap \frac{1}{N}\mathbf{Z}^n = \{x \in X : Nx \in \mathbf{Z}^n\}$. The process is characterized by the common jump rate $\lambda_x^N \equiv NR$ and the transition probabilities

$$(11.1) \quad P_{xy}^N = \begin{cases} \frac{x_i \rho_{ij}(F(x), x)}{R} & \text{if } y = x + \frac{1}{N}(e_j - e_i), j \neq i, \\ 1 - \sum_{i \in S} \sum_{j \neq i} \frac{x_i \rho_{ij}(F(x), x)}{R} & \text{if } y = x, \\ 0 & \text{otherwise.} \end{cases}$$

To ensure that transition probabilities are well-defined, we assume that the clock rate R is an upper bound on the row sums of the revision protocol ρ (see equation (10.1)). To introduce the possibility of unique infinite-horizon predictions, we now assume in addition that the conditional switch rates are bounded away from zero: there is a positive constant \underline{R} such that

$$(11.2) \quad \rho_{ij}(F(x), x) \geq \underline{R} \text{ for all } i, j \in S \text{ and } x \in X.$$

We refer to a revision protocol that satisfies condition (11.2) as having *full support*.

Full support protocols usually include some form of perturbation to ensure that all strategies are always chosen with positive probability. The next two examples consider two full-support extensions of best response protocols. In each case, one can interpret the choice of a suboptimal strategy as the result of a mistake, of a noisy observation of payoffs, or of conscious experimentation.

Example 11.1.1. Best response with mutations. Under *best response with mutations* at mutation rate $\varepsilon > 0$, called *BRM*(ε) for short, a revising agent switches to his current best response with probability $1 - \varepsilon$, but chooses a strategy uniformly at random (or *mutates*) with probability $\varepsilon > 0$. Thus, if the game has two strategies, each yielding different payoffs, a revising agent will choose the optimal strategy with probability $1 - \frac{\varepsilon}{2}$ and will choose the suboptimal strategy with probability $\frac{\varepsilon}{2}$.

To complete the specification of the protocol, one must specify what a non-mutating agent does when there are multiple optimal strategies. A common setup has an agent who does not mutate stick with his current strategy if it is optimal, and choose at random among the optimal strategies otherwise. §

Exercise 11.1.2. Compute the mean dynamic for the BRM protocol, focusing on states at which the best response is unique. ‡

Example 11.1.3. Logit choice. Of the revision protocols underlying the six basic dynamics studied in Chapters 5 and 6, the only one satisfying the full support condition (11.2) is the logit choice protocol with noise level $\eta > 0$, introduced in Example 6.2.1:

$$(11.3) \quad \rho_{ij}(\pi) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}.$$

For intuition, it is useful to rewrite the logit choice protocol as

$$(11.4) \quad \rho_{ij}(\pi) = \frac{\exp(\eta^{-1}(\pi_j - \pi_{k^*}))}{\sum_{k \in S} \exp(\eta^{-1}(\pi_k - \pi_{k^*}))}.$$

where k^* is an optimal strategy under π . Then as η approaches zero, the denominator of (11.4) converges to a constant (namely, the number of optimal strategies under π), so as η^{-1} approaches infinity, $\rho_{ij}(\pi, x)$ vanishes at exponential rate $\pi_{k^*} - \pi_j$. §

The mutation rate $\varepsilon \in (0, 1)$ from the BRM protocol directly specifies the probability of a suboptimal choice, while the noise level $\eta \in (0, \infty)$ of the logit protocol appears as an exponent in the formula for the choice probability. One can move between these

parameterizations using the transformation $\varepsilon \mapsto \eta = -(\log \varepsilon)^{-1}$, an increasing function from $(0, 1)$ onto $(0, \infty)$, and its inverse $\eta \mapsto \varepsilon = \exp(-\eta^{-1})$. As the next exercise shows, these transformations shuttle the noise parameter between the base and the exponent of the expression for $\rho_{ij}(\pi)$. We will use these transformations frequently when studying stochastic stability in Chapter 12; see Sections 12.3.1 and 12.A.5 for additional discussions.

- Exercise 11.1.4.* (i) Express the BRM rule in terms of the noise parameter η .
(ii) Express the logit choice rule in terms of the noise parameter ε . ‡

As their noise parameters approach zero, both the BRM and logit protocols come to resemble the exact best response protocol, the protocol underlying the deterministic best response dynamic (Section 6.1). But this similarity masks a fundamental qualitative difference between the two protocols. Under best response with mutations, the probability of choosing a particular suboptimal strategy is independent of the payoff consequences of doing so: mutations do not favor alternative strategies with higher payoffs over those with lower payoffs. In contrast, since the logit protocol is defined using payoff perturbations that are symmetric across strategies, more costly “mistakes” are less likely to be made.

The probabilities of suboptimal choices under both the BRM and logit protocols are small when the noise level is small. One might expect the precise specification of these probabilities to be of little consequence. If our interest is in finite-horizon behavior, so that the mean dynamic (M) forms the basis for our predictions, this impression is largely correct (see, for instance, Theorem 9.B.5). But as we shall see in Section 11.2.3 and throughout Chapter 12, predictions of infinite horizon behavior hinge on the relative probabilities of rare events. As a consequence, seemingly minor differences in choice probabilities can lead to entirely different predictions of behavior.

While mutations are commonly employed in combination with a best response rule, they are an equally natural complement to imitative rules.

Example 11.1.5. Imitative logit choice with (and without) mutations. The protocol

$$(11.5) \quad \rho_{ij}(\pi, x) = \frac{x_j \exp(\eta^{-1}\pi_j)}{\sum_{k \in S} x_k \exp(\eta^{-1}\pi_k)} + \varepsilon$$

augments the i-logit protocol (Example 5.4.9) by adding rare mutations. Since $\exp(\eta^{-1}\pi_j)$ is always positive, the conditional switch rates to all strategies currently in use are positive even when $\varepsilon = 0$. We take advantage of this property in Section 11.4.3, where we introduce a modification of our basic model under which the process $\{X_t^N\}$ generated by protocol (11.5) is irreducible even without mutations. §

11.1.2 Stationary Distributions and Infinite Horizon Behavior

The full support assumption (11.2) ensures that at each revision opportunity, every strategy in S has a positive probability of being chosen by the revising agent. Therefore, there is a positive probability that the process $\{X_t^N\}$ will transit from any given current state x to any other given state y within a finite number of periods. A Markov process with this property is said to be *irreducible*.

Below we summarize some basic results on the infinite-horizon behavior of irreducible Markov processes. These results provide the foundation for all of our subsequent analyses. A detailed presentation of the relevant theory is offered in Appendix 11.A.

Suppose that $\{X_t\}_{t \geq 0}$ is an irreducible Markov process on the finite state space \mathcal{X} , where the process has equal jump rates $\lambda_x \equiv l$ and transition matrix P . Theorem 11.A.18 shows that there is a unique probability vector $\mu \in \mathbf{R}_+^{\mathcal{X}}$ satisfying

$$(11.6) \quad \sum_{x \in \mathcal{X}} \mu_x P_{xy} = \mu_y \quad \text{for all } y \in \mathcal{X}.$$

The vector μ is called the *stationary distribution* of the process $\{X_t\}$. Equation (11.6) tells us that if we run the process $\{X_t\}$ from initial distribution μ , then at the random time of the first jump, the distribution of the process is also μ . Moreover, if we use the notation $\mathbb{P}_\pi(\cdot)$ to represent $\{X_t\}$ being run from initial distribution π , then equation (11.52) shows that

$$(11.7) \quad \mathbb{P}_\mu(X_t = x) = \mu_x \quad \text{for all } x \in \mathcal{X} \text{ and } t \geq 0.$$

In other words, if the process starts off in its stationary distribution, it remains in this distribution at all subsequent times t .

While equation (11.7) tells us what happens if $\{X_t\}$ starts off in its stationary distribution, our main interest is in what happens to this process in the very long run if it starts in an arbitrary initial distribution π . Theorem 11.A.19 shows that as t grows large, the time t distribution of $\{X_t\}$ converges to μ :

$$(11.8) \quad \lim_{t \rightarrow \infty} \mathbb{P}_\pi(X_t = x) = \mu_x \quad \text{for all } x \in \mathcal{X}.$$

Thus, looking at the process $\{X_t\}$ from the ex ante point of view, the probable locations of the process at sufficiently distant future times are essentially determined by μ .

To describe long run behavior from an ex post point of view, we need to consider the behavior of the process's sample paths. Here again, the stationary distribution plays the central role. Theorem 11.A.22 states that along almost every sample path, the proportion

of time spent at each state in the long run is described by μ :

$$(11.9) \quad \mathbb{P}_\pi \left(\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T 1_{\{X_t=x\}} dt = \mu_x \right) = 1 \quad \text{for all } x \in \mathcal{X}.$$

We can also summarize equation (11.9) by saying that the limiting empirical distribution of $\{X_t\}$ is almost surely equal to μ .

11.1.3 Reversibility

In general, computing the stationary distribution of a Markov process means finding an eigenvector of a matrix, a task that is computationally daunting unless the state space, and hence the dimension of the matrix, is small. But there is a special class of Markov processes whose stationary distributions are easy to compute. A constant jump rate Markov process $\{X_t\}$ is said to be *reversible* if it admits a *reversible distribution*: a probability distribution μ on \mathcal{X} that satisfies the *detailed balance conditions*:

$$(11.10) \quad \mu_x P_{xy} = \mu_y P_{yx} \quad \text{for all } x, y \in \mathcal{X}.$$

A process satisfying this condition is called reversible because, probabilistically speaking, it “looks the same” whether time is run forward or backward; see Appendix 11.A.6 for a discussion. Since summing the equality in (11.10) over x yields condition (11.6), a reversible distribution is also a stationary distribution.

While in general reversible Markov processes are rather special, in the present context there are two environments in which the evolutionary process $\{X_t^N\}$ is reversible: in two-strategy games with an arbitrary revision protocol, and in potential games under the logit choice protocol. These two situations are studied in the next section and in Section 11.5.

11.2 Stationary Distributions for Two-Strategy Games

When the population plays a game with just two strategies, the state space \mathcal{X}^N is a grid in the simplex in \mathbf{R}^2 , and so is linearly ordered. In this case, regardless of the (full support) revision protocol the agents employ, we can compute the stationary distribution of the process $\{X_t^N\}$ explicitly. After deriving this distribution in Theorem 11.2.4, we apply it to a few examples, providing a preview of the equilibrium selection results that are the focus of Chapter 12.

Let $F : X \rightarrow \mathbf{R}^2$ be a two strategy game with strategy set $S = \{0, 1\}$, let $\rho : \mathbf{R}^2 \times X \rightarrow \mathbf{R}^{2 \times 2}$

be a full support revision protocol, and let N be a finite population size. As we saw in Section 11.1, these objects define a Markov process $\{X_t^N\}$ on the state space \mathcal{X}^N .

While population states in this game are elements of $X = \{x \in \mathbf{R}_+^2 : x_0 + x_1 = 1\}$, the simplex in \mathbf{R}^2 , it is convenient to identify state x with the weight $\chi \equiv x_1$ that it places on strategy 1. Under this notational device, the state space of the Markov process $\{X_t^N\}$ becomes $\mathcal{X}^N = \{0, \frac{1}{N}, \dots, 1\}$, a uniformly-spaced grid in the unit interval. We will also write $F(\chi)$ for $F(x)$ and $\rho(\pi, \chi)$ for $\rho(\pi, x)$ whenever it is convenient to do so.

11.2.1 Birth and Death Processes

Because agents in our model switch strategies sequentially, transitions of the process $\{X_t^N\}$ are always between adjacent states. Since in addition states are linearly ordered, $\{X_t^N\}$ falls into a class of Markov processes called birth and death processes. These processes are quite amenable to explicit calculations, as we now illustrate by deriving a simple expression for the stationary distribution μ^N .

A constant jump rate Markov process $\{X_t^N\}$ on the state space $\mathcal{X}^N = \{0, \frac{1}{N}, \dots, 1\}$ is a *birth and death process* if the only positive probability transitions move one step to the right, move one step to the left, or remain still. This implies that there are vectors $p^N, q^N \in \mathbf{R}^{\mathcal{X}^N}$ with $p_1^N = q_0^N = 0$ such that the transition matrix of $\{X_t^N\}$ takes the form

$$P_{\chi y}^N \equiv \begin{cases} p_\chi^N & \text{if } y = \chi + \frac{1}{N}, \\ q_\chi^N & \text{if } y = \chi - \frac{1}{N}, \\ 1 - p_\chi^N - q_\chi^N & \text{if } y = \chi, \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, the process $\{X_t^N\}$ is irreducible if $p_\chi^N > 0$ for $\chi < 1$ and $q_\chi^N > 0$ for $\chi > 0$, as we henceforth assume.

Because of their simple transition structure, birth and death chains are reversible. For the transition matrix above, the reversibility conditions (11.10) reduce to

$$\mu_\chi^N q_\chi^N = \mu_{\chi-1/N}^N p_{\chi-1/N}^N \quad \text{for } \chi \in \{\frac{1}{N}, \dots, 1\}.$$

Applying this formula inductively, we find that the stationary distribution of $\{X_t^N\}$ satisfies

$$(11.11) \quad \frac{\mu_\chi^N}{\mu_0^N} = \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N} \quad \text{for } \chi \in \{\frac{1}{N}, \dots, 1\}.$$

Since the weights in μ^N must sum to one, we can compute μ_0 as

$$(11.12) \quad \mu_0^N = \left(\sum_{i=0}^N \prod_{j=1}^i \frac{p_{(j-1)/N}^N}{q_{j/N}^N} \right)^{-1},$$

where the empty product equals 1.

Example 11.2.1. Toss and switch revisited. In Example 10.2.4, we considered a population of size N whose members are equipped with rate 1 Poisson alarm clocks. Each agent responds to the ringing of his clock by flipping a fair coin, switching strategies if the coin comes up Heads. The resulting Markov process $\{X_t^N\}$ is irreducible, with constant jump rate $\lambda_\chi^N \equiv N$ and positive transition probabilities $p_\chi^N = \frac{1}{2}(1 - \chi)$ and $q_\chi^N = \frac{1}{2}\chi$.

In Example 10.2.4, we showed that the mean dynamic of this process is $\dot{\chi} = \frac{1}{2} - \chi$. Solutions of this dynamic are of the form $\chi_t = \frac{1}{2} + (\chi_0 - \frac{1}{2})e^{-t}$, so the dynamic is an exponential contraction toward the central state $\chi^* = \frac{1}{2}$. Now suppose we fix a time horizon $T < \infty$ and an error bound $\varepsilon > 0$. Then Theorem 10.2.3 tells us that if the population size N is large enough, the value of the random variable X_t^N will stay within ε of $\chi_t = \frac{1}{2} + (\chi_0 - \frac{1}{2})e^{-t}$ for all times t in the interval $[0, T]$ with probability at least $1 - \varepsilon$ (see Figure 10.2.1).

Now suppose instead that we fix the population size N and consider the behavior of the process $\{X_t^N\}$ over a very long time horizon. As discussed in Section 11.1.2, the limiting distribution and the limiting empirical distribution of $\{X_t^N\}$ are given by its stationary distribution μ^N . Using formulas (11.11) and (11.12), it is easy to show that this distribution is given by

$$\mu_\chi^N = \frac{1}{2^N} \binom{N}{N\chi} \text{ for all } \chi \in \mathcal{X}^N = \{0, \frac{1}{N}, \dots, 1\}.$$

In words, μ^N describes a binomial distribution with parameters N and $\frac{1}{2}$, but with outcomes (in \mathcal{X}^N) representing the proportion rather than the number of successful trials.

If N is not small, the Central Limit Theorem tells us that μ^N is approximately normal with mean $\frac{1}{2}$ and variance $\frac{1}{4N}$. Figure 11.2.1 illustrates μ^N for population sizes $N = 100$ and $N = 10,000$. If the population is not too small and if enough time has passed, all states in \mathcal{X}^N will have been visited many times, but the vast majority of periods will have been spent at states where the two strategies are used in nearly proportions. §

Exercise 11.2.2. (i) Suppose that in the previous example, revising agents flip a coin that comes up Heads with probability $h \in (0, 1)$, switching strategies when Heads





Figure 11.2.1: Mean dynamics and stationary distributions for Toss and Switch.

occurs. What are the mean dynamic of $\{X_t^N\}$ and the stationary distribution of $\{X_t^N\}$ in this case?

- (ii) Now suppose that revising agents flip a coin that comes up Heads with probability $h \in (0, 1)$, choosing strategy 1 when Heads occurs and strategy 0 otherwise. What are the mean dynamic of $\{X_t^N\}$ and the stationary distribution of $\{X_t^N\}$ now? ‡

Exercise 11.2.3. Toss and Switch with n strategies. Consider the n -strategy version of the Toss and Switch process: an agent whose clock rings randomizes uniformly over the n available strategies. Verify that this process is reversible with stationary distribution

$$\mu_x^N = \frac{1}{n^N} \frac{N!}{\prod_{k=1}^n (Nx_k)!} \text{ for all } x \in \mathcal{X}^N. \quad \ddagger$$

11.2.2 The Stationary Distribution of the Evolutionary Process

Let us now use formula (11.11) to compute the stationary distribution of our stochastic evolutionary process, maintaining the assumption that the process is generated by a full support revision protocol. Referring back to Section 11.1, we find that the process $\{X_t^N\}$ has constant jump rates $\lambda_\chi^N = NR$, and that its upward and downward transition probabilities are given by

$$(11.13) \quad p_\chi^N = (1 - \chi) \cdot \frac{1}{R} \rho_{01}(F(\chi), \chi) \text{ and}$$

$$(11.14) \quad q_\chi^N = \chi \cdot \frac{1}{R} \rho_{10}(F(\chi), \chi).$$

Substituting formulas (11.13) and (11.14) into equation (11.11), we see that for $\chi \in \{\frac{1}{N}, \frac{2}{N}, \dots, 1\}$, we have

$$\frac{\mu_\chi^N}{\mu_0^N} = \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N} = \prod_{j=1}^{N\chi} \frac{(1 - \frac{j-1}{N})}{\frac{j}{N}} \cdot \frac{\frac{1}{R} \rho_{01}(F(\frac{j-1}{N}), \frac{j-1}{N})}{\frac{1}{R} \rho_{10}(F(\frac{j}{N}), \frac{j}{N})}.$$

Simplifying this expression yields the following result.

Theorem 11.2.4. *Suppose that a population of N agents plays the two-strategy game F using the full support revision protocol ρ . Then the stationary distribution for the evolutionary process $\{X_t^N\}$ on \mathcal{X}^N is given by*

$$\frac{\mu_\chi^N}{\mu_0^N} = \prod_{j=1}^{N\chi} \frac{(N - j + 1)}{j} \cdot \frac{\rho_{01}(F(\frac{j-1}{N}), \frac{j-1}{N})}{\rho_{10}(F(\frac{j}{N}), \frac{j}{N})} \text{ for } \chi \in \{\frac{1}{N}, \frac{2}{N}, \dots, 1\},$$

with μ_0^N determined by the requirement that $\sum_{\chi \in \mathcal{X}^N} \mu_\chi^N = 1$.

11.2.3 Examples

The power of infinite horizon analysis lies in its ability to generate unique predictions of play even in games with multiple strict equilibria. We now illustrate this idea by computing some stationary distributions for two-strategy coordination games under the BRM and logit rules. In all cases, we find that these distributions place most of their weight near a single equilibrium. But we also find that the two rules need not select the same equilibrium.

To obtain unique predictions of infinite horizon behavior, it is generally enough either that the population size not be too small, or that the noise level in agents' choices not be too large. But one can obtain cleaner and more general results by studying the limiting behavior of the stationary distribution as the population size approaches infinity, the noise level approaches zero, or both. This approach to studying infinite horizon behavior, known as *stochastic stability theory*, is the subject of Chapter 12.

Example 11.2.5. Stag Hunt. The symmetric normal form coordination game

$$A = \begin{pmatrix} h & h \\ 0 & s \end{pmatrix}$$

with $s > h > 0$ is known as *Stag Hunt*. By way of interpretation, we imagine that each agent in a match must decide whether to hunt for hare or for stag. Hunting for hare ensures a payoff of h regardless of the match partner's choice. Hunting for stag can generate a payoff of $s > h$ if the opponent does the same, but results in a zero payoff otherwise. Each of the two strategies has distinct merits. Coordinating on Stag yields higher payoffs than coordinating on Hare. But the payoff to Hare is certain, while the payoff to Stag depends on the choice of one's partner.

Suppose that a population of agents is repeatedly matched to play Stag Hunt. If we let χ denote the proportion of agents playing Stag, then with our usual abuse of notation, the payoffs in the resulting population game are $F_H(\chi) = h$ and $F_S(\chi) = s\chi$. This population game has three Nash equilibria: the two pure equilibria, and the mixed equilibrium $\chi^* = \frac{h}{s}$. We henceforth suppose that $h = 2$ and $s = 3$, so that the mixed equilibrium places mass $\chi^* = \frac{2}{3}$ on Stag.

Suppose that agents follow the best response with mutations protocol, with mutation rate $\varepsilon = .10$. The resulting mean dynamic,

$$\dot{\chi} = \begin{cases} \frac{\varepsilon}{2} - \chi & \text{if } \chi < \frac{2}{3}, \\ (1 - \frac{\varepsilon}{2}) - \chi & \text{if } \chi > \frac{2}{3}, \end{cases}$$


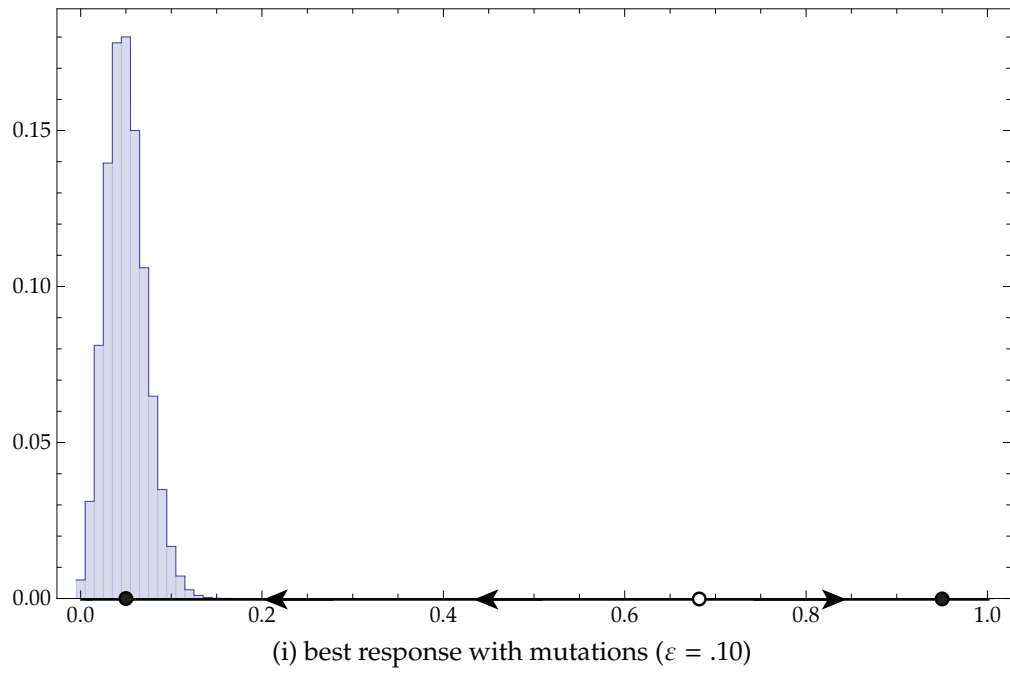


Figure 11.2.2: Stationary distribution weights μ_χ for Stag Hunt ($h = 2, s = 3, N = 100$).

has stable rest points at $\chi = .05$ and $\chi = .95$. The basins of attraction of these rest points meet at the mixed equilibrium $\chi^* = \frac{2}{3}$. Note that the rest point that approximates the all-Hare equilibrium has the larger basin of attraction.

In Figure 11.2.2(i), we present this mean dynamic underneath the stationary distribution μ^N for $N = 100$, which we computed using the formula derived in Theorem 11.2.4. While the mean dynamic has two stable equilibria, nearly all of the mass in the stationary distribution is concentrated at states where between 88 and 100 agents choose Hare. Thus, while coordinating on Stag is efficient, the “safe” strategy Hare is selected by the stochastic evolutionary process.

Suppose instead that agents use the logit rule with noise level $\eta = .25$. The mean dynamic is then the logit dynamic,

$$\dot{\chi} = \frac{\exp(3\chi\eta^{-1})}{\exp(2\eta^{-1}) + \exp(3\chi\eta^{-1})} - \chi,$$

which has stable rest points at $\chi = .0003$ and $\chi = .9762$, and an unstable rest point at $\chi = .7650$, so that the basin of attraction of the “almost all-Hare” rest point $\chi = .0003$ is even larger than under BRM. Examining the resulting stationary distribution (Figure 11.2.2(ii)), we see that virtually all of its mass is placed on states where either 99 or 100 agents choose Hare, in rough agreement with the result for the BRM(.10) rule. §

Why does most of the mass in the stationary distribution becomes concentrated around a single equilibrium? The stochastic evolutionary process $\{X_t^N\}$ typically moves in the direction indicated by the mean dynamic. If the process begins in the basin of attraction of a rest point or other attractor of this dynamic, then the initial period of evolution generally results in convergence to and lingering near this locally stable set.

However, since BRM and logit choice lead to irreducible evolutionary processes, this cannot be the end of the story. Indeed, we know that the process $\{X_t^N\}$ eventually reaches all states in \mathcal{X}^N ; in fact, it visits all states in \mathcal{X}^N infinitely often. This means that the process at some point must leave the basin of the stable set visited first; it then enters the basin of a new stable set, at which point it is extremely likely to head directly the set itself. The evolution of the process continues in this fashion, with long periods of visits to each attractor punctuated by sudden jumps between the stable set.

Which states are visited most often over the infinite horizon is determined by the *relative* unlikelihoods of these rare but inevitable transitions between stable sets. In the examples above, the transitions from the Stag rest point to the Hare rest point and from the Hare rest point to the Stag rest point are both very unlikely events. But for purposes of determining the stationary distribution, what matters is that in relative terms, the former

transitions are much more likely than the latter. This enables us to conclude that over very long time spans, the evolutionary process will spend most periods at states where most agents play Hare.

While Theorem 11.2.4 makes it very easy to compute the stationary distributions in the example above, it obscured the forces that underlie equilibrium selection. When we study stochastic stability in Chapter 12, we will be able to determine the limit behavior of the stationary distribution as the noise level becomes small even in cases in which there is no simple expression for the distribution itself. We accomplish this in Section 12.5 by studying the asymptotic probabilities of transitions between stable sets, as suggested by the discussions above.

At the same time, the fact that transitions out of the basin of any stable set are very low probability events suggests that the predictions provided by the stationary distribution may only be relevant in extremely long-running interactions. We verify and discuss this point in Section 11.3 below.

Example 11.2.6. A nonlinear Stag Hunt. We now consider a version of the Stag Hunt game in which payoffs depend nonlinearly on the population state. With our usual abuse of notation, we define payoffs in this game by $F_H(\chi) = h$ and $F_S(\chi) = s\chi^2$, with χ representing the proportion of agents playing Stag. The population game F has three Nash equilibria: the pure equilibria $\chi = 0$ and $\chi = 1$, and the mixed equilibrium $\chi^* = \sqrt{h/s}$. We focus on the case in which $h = 2$ and $s = 7$, so that $\chi^* = \sqrt{2/7} \approx .5345$.

Suppose first that a population of 100 agents play this game using the BRM(.10) rule. In Figure 11.2.3(i) we present the resulting mean dynamic beneath a graph of the stationary distribution μ^{100} . The mean dynamic has rest points at $\chi = .05$, $\chi = .95$, and $\chi^* \approx .5345$, so the the “almost all Hare” rest point again has the larger basin of attraction. As was true in the linear Stag Hunt from Example 11.2.5, the stationary distribution generated by the BRM(.10) rule in this nonlinear Stag Hunt places nearly all of its mass on states where at least 88 agents choose Hare.

Figure 11.2.3(ii) presents the mean dynamic and the stationary distribution μ^{100} for the logit rule with $\eta = .25$. The rest points of the logit(.25) dynamic are $\chi = .0003$, $\chi = 1$, and $\chi = .5398$, so the “almost all Hare” rest point once again has the larger basin of attraction. Nevertheless, the stationary distribution μ^{100} places virtually all of its mass on the state in which all 100 agents choose Stag.

To summarize, our prediction for very long run behavior under the BRM(.10) rule is inefficient coordination on Hare, while our prediction under the logit(.25) rule is efficient coordination on Stag. §

For the intuition behind this discrepancy in predictions, recall the discussion from





Figure 11.2.3: Stationary distribution weights μ_χ for a nonlinear Stag Hunt ($h = 2, s = 7, N = 100$).

Section 11.1.1 about the basic distinction between the logit and BRM protocols: under logit choice, the probability of a “mistake” depends on its payoff consequences, while under BRM, it does not. As we shall see in Chapter 12, the latter observation implies that under BRM, the probabilities of escaping from the basins of attraction of stable sets, and hence the identities of the states predominate in the very long run, depend only on the size and the shapes of the basins. In the current one-dimensional example, these shapes are always line segments, so that only the size of the basins matters; since the “almost all-Hare” state has the larger basin, it is selected under the BRM rule.

On the contrary, the probability of escaping a stable equilibrium under logit choice depends not only on the shape and size of its basin, but also on the payoff differences that must be overcome during the journey. In the nonlinear Stag Hunt game, the basin of the “almost all-Stag” equilibrium is smaller than that of the all-Hare equilibrium. But because the payoff advantage of Stag over Hare in the former’s basin tends to be much larger than the payoff advantage of Hare over Stag in the latter’s, it is more difficult for the population to escape the all-Stag equilibrium than the all-Hare equilibrium; as a result, the population spends virtually all periods coordinating on Stag over the infinite horizon.

We can compare the process of escaping from the basin of a stable rest point to an attempt to swim upstream. Under BRM, the stream is steady, so the difficulty of a given excursion is proportional to distance. Under logit choice, the strength of the stream is variable, so the difficulty of an excursion depends on how this strength varies over the distance travelled. In general, the probability of escaping from a stable set is determined by both the distance that must be travelled and the strength of the oncoming flow.

11.3 Waiting Times and Infinite Horizon Prediction

In order to place credence in the history independent prediction provided by the stationary distribution, we must believe that the time horizon of interest in our application is sufficiently long to justify an appeal to this distribution. In this section, we show that the lengths of time needed for history independent predictions to be relevant are often extremely long, too long to be of use in any economic application. After presenting some examples, we offer some suggestions about when infinite horizon predictions are likely to be of practical use.

11.3.1 Examples

In cases where the mean dynamic has a globally asymptotically stable state, the stationary distribution μ^N will concentrate its mass around this state, at least when the population size is large. (See Section 12.6 for a formal analysis.) We begin our treatment of waiting times by considering such a case.

Example 11.3.1. Toss and Switch once more. In Examples 10.2.4 and 11.2.1, we considered the Toss and Switch process $\{X_t^N\}$, which is defined by the constant jump rate $\lambda_\chi^N \equiv N$ and positive transition probabilities $p_\chi^N = \frac{1}{2}(1 - \chi)$ and $q_\chi^N = \frac{1}{2}\chi$. We have seen that the mean dynamic of this process, $\dot{\chi} = \frac{1}{2} - \chi$, is a contraction toward the rest point $\chi^* = \frac{1}{2}$, and that its stationary distribution is approximately normal with mean $\frac{1}{2}$ and variance $\frac{1}{4N}$. Thus, the larger the population size N , the more concentrated around the rest point $\chi^* = \frac{1}{2}$ the stationary distribution becomes.

Now suppose the process $\{X_t^N\}$ begins in state $\chi = 0$. How long will it take before the agents are nearly equally divided between the two strategies? Using equation (11.44) from Appendix 11.A.3, we can compute the expected time before the process reaches state .45, and the expected time before it reaches state .5, considering population sizes of 100, 1,000, and 10,000. We report the results in Table 11.1 below. Evidently, an equal distribution between the strategies is achieved very quickly, even when the population size is large.

For some insight into the numbers of the table, let us compare the numbers in the table to predictions of waiting times based on the mean dynamic. (This comparison is not always justified: see Exercise 11.3.3 below.) The general solution to the mean dynamic is $\chi_t = \frac{1}{2} + (\chi_0 - \frac{1}{2})e^{-t}$. If we run the dynamic from state $x_0 = 0$, the time T that state $x_T = .45$ is reached must satisfy $x_T = \frac{1}{2} - \frac{1}{2}e^{-T}$, and hence $T = \log 10 \approx 2.3026$. This is very close to the expected waiting time of 2.2977 for the 10,000 agent process.

Because the mean dynamic is continuous, it requires an infinite amount of time to reach the rest point $\chi^* = \frac{1}{2}$ from any state other than x^* . Nevertheless, because of the random fluctuations in the process $\{X_t^N\}$, the expected time for this process to travel from state 0 to state $\frac{1}{2}$ remains quite small even when $N = 10,000$. §

In cases where the mean dynamic (M) has multiple asymptotically stable sets, we can reach similar but somewhat weaker conclusions about the time required to reach an asymptotically stable state from an initial condition in that state's basin of attraction: Suppose that x_0 is in the basin of x^* , and the initial conditions $X_0^N = x_0^N$ converge to x_0 . Then it follows immediately from Theorem 10.2.3 that for any $\varepsilon > 0$, there is a finite time

	$N = 100$	$N = 1,000$	$N = 10,000$
$x = .45$	2.0389	2.2588	2.2977
$x = .50$	2.9378	4.0891	5.2403

Table 11.1: Expected wait to reach state x from state 0 in Toss and Switch.

	$N = 100$	$N = 1,000$	$N = 10,000$
$\varepsilon = .1$	3.03×10^{17}	7.33×10^{171}	1.59×10^{1720}
$\varepsilon = .01$	1.23×10^{50}	2.60×10^{492}	1.43×10^{4920}

Table 11.2: Expected wait before escaping the basin of attraction of “almost all Stag” under BRM(ε).

T such that

$$(11.15) \quad \lim_{N \rightarrow \infty} \mathbb{P}(|X_T^N - x^*| < \varepsilon) = 1.$$

(To see why this is weaker than the conclusions of Example 11.3.1, see Exercise 11.3.3 below.)

Now imagine again that the mean dynamic (M) has multiple stable sets, and that a neighborhood of just one of these sets receives the preponderance of the mass in the stationary distribution μ^N . If we are to use this fact as the basis for an infinite horizon prediction, we should believe that if the population begins play near some other stable set of (M), it will transit to the one selected by μ^N within some reasonable amount of time. But the lengths of time that such transitions require can be extraordinary.

Example 11.3.2. Stag Hunt once more. Example 11.2.5 presented stationary distributions for evolution in the Stag Hunt game

$$(11.16) \quad A = \begin{pmatrix} 2 & 2 \\ 0 & 3 \end{pmatrix}$$

under the BRM and logit protocols. In particular, Figure 11.2.2(i) shows that with a population size of 100 under the BRM(.10) protocol, virtually all of the weight in the stationary distribution is placed in the vicinity of state $\chi = .05$ where most agents play the “safe” strategy Hare. Still, if more than $\chi^* = \frac{2}{3}$ of the population initially plays Stag, the expected motion of the process is toward state $\chi = .95$, at which nearly all agents play Stag. The results discussed in Section 11.1.2 show that after a long enough time has passed, the empirical distribution of the process $\{X_t^N\}$ will be close to its stationary

distribution, implying that the population will have spent most periods coordinating on Hare. How much time must pass before this infinite horizon analysis is relevant?

In Table 11.2, we report the expected amount of time before a process starting at state $\chi = .95$ reaches a state less than $\chi^* = \frac{2}{3}$, so that it is out of the basin of attraction of the all-Stag equilibrium. We do this for two mutation rates (.10 and .01) and three population sizes (100, 1,000, and 10,000). All of the expected waiting times are large. To provide some feeling for the size of the numbers in the table, we note that the age of the universe is currently estimated to be 1.37×10^{10} years, or about 4.33×10^{17} seconds. §

The example above shows that the expected amount of time to leave the basin of a stable rest point, including the rest point that is not selected by the stationary distribution, can be extremely large. This calls into question whether the stationary distribution is relevant in applications. We discuss this issue at length below.

Exercise 11.3.3. The expected time to reach a stable rest point from a point *in its own basin of attraction* can also be surprisingly large. As in the previous example, suppose that 100 agents use the BRM(.10) protocol during recurrent play of the Stag Hunt game (11.16). Remember that an agent who “mutates” chooses his strategy randomly, the probability that a revising agent chooses the suboptimal strategy is .05. Also, to interpret the claims below, recall that the mean dynamic has stable rest points at $\chi = .05$ and $\chi = .95$, with the division point between the basins of attraction occurring at state $\chi^* = \frac{2}{3}$.

- (i) Using one of the formulas from Example 11.A.8 (and a computer!), show that if play begins in state $\chi = .20$, the expected time required before the “almost all Hare” state $\chi = .05$ is first reached is about 8.36. (Since the formulas in Example 11.A.8 are for Markov chains, and since the continuous-time evolutionary process $\{X_t^N\}_{t \geq 0}$ has constant jump rate $N = 100$, the result of the Markov chain calculation must be divided by 100 to obtain the correct expected hitting time for the Markov process—see the discussion following Proposition 11.A.7.)
- (ii) Now show that if play begins in state $\chi = .80$, the expected time required before the “almost all Stag” state $\chi = .95$ is first reached is about 2.44×10^{32} .
- (iii) Using equation (11.46) from Example 11.A.9, show that if play begins in state .80, then the probability of reaching state .05 before state .95 is about 1.32×10^{-7} , in accordance with the deterministic approximation theorem and equation (11.15).
- (iv) At first glance, the claims made in (ii) and (iii) may seem inconsistent with one another. Explain in words how they can both be true.

While the discussion above has focused on mean hitting times, one can also gauge the amount of time needed before the stationary distribution becomes relevant by determining

how quickly the time t distributions of the evolutionary process approach the stationary distribution. See Appendix 11.A.8 for a brief presentation of the relevant techniques.

11.3.2 Discussion

Example 11.3.2 shows that the amount of time needed to escape from the basin of attraction of a stable equilibrium can be extraordinarily large, even when the population size and noise parameters are not extreme. In fact, it is possible to show that the expected time required to escape the basin of a stable equilibrium grows exponentially both in the population size N and in the inverse noise level η^{-1} (which corresponds to $-\log \varepsilon$ in the BRM(ε) model; see Section 11.1.1). The waiting time critique will only become more acute in the next chapter, where we introduce the notion of stochastic stability. In stochastic stability analysis, one studies the limiting behavior of the stationary distributions of the stochastic evolutionary process as the population size N goes to infinity and as the noise level η goes to zero. Doing so enables one to obtain cleaner and more general selection results than are possible without taking limits; at the same time, taking limits in N and η guarantees that the waiting time problem will be especially severe. While in this book we will not pursue any formal analyses of waiting times in models of stochastic evolution, the Notes provide a number of references where such analyses can be found.

What are the implications for predicting behavior in applications? Looking at the size of the numbers in Table 11.2, it seems hard to avoid the conclusion that in economic environments with at least moderate population sizes at most moderate amounts of idiosyncratic noise in agents' decisions, the predictions of infinite horizon analysis do not hold force within any relevant time span. Therefore, the history-dependent predictions provided by the mean dynamic (M) seem the most appropriate ones to utilize. In most cases this means giving up the possibility of unique predictions that do not require knowledge of past behavior, but there is little to gain in making unique predictions that are as likely as not to be incorrect.

Can we identify environments in which infinite horizon predictions are appropriate? One possibility lies in modeling biological evolution: with simple organisms, generation lengths are very short, mutation rates can be measured precisely, and the time spans of interest may last thousands of years. If we restrict attention to human decision makers, then small populations and large noise levels will make waiting times lower, and so can make infinite horizon predictions more apt. At the same time, when populations are small, one may doubt the appropriateness of our standing assumptions that agents are anonymous and myopic. We have seen as well that large noise levels impose a strong equalizing force on the mean dynamic (cf Example 6.2.4); in cases where this dynamic

already admits a globally asymptotically stable state, infinite horizon analysis provides limited additional predictive power.

In seeking further scope for infinite horizon predictions, it is worth bearing in mind that every economic model is an abstraction from the application it represents, one that trades off descriptive accuracy for ease of analysis and interpretation. We therefore should not be too hasty to discard a model with some clearly unrealistic implications if we feel that those implications might be absent from a more detailed though possibly less tractable model.

In the case of the stochastic evolutionary process studied in this book, the long waiting times for transitions between equilibria are due to these transitions requiring the contemporaneous occurrence of many of independent, low probability events. There are plausible alternative models in which the requirements for transitions are much less unlikely to be met.

One possibility lies in replacing the global interaction structure used in this book with structures in which each agent only interacts with a small, fixed subset of his opponents. For instance, agents might live in distinct locations, and only interact with neighbors. Alternatively, agents might be linked through a network (based on friendship, say, or collaboration), and only interact with those to whom they have a direct connection. In these cases, simultaneous changes in strategy by agents in a single, small neighborhood can spread contagiously throughout the entire population. Since a small number of revisions can set off a population-wide transition between equilibria, the wait until a transition occurs need not be long. Unfortunately, local interaction and network models are beyond the scope of this book; see the Notes for a selection of references to the literature.

Waiting times could also be greatly reduced by allowing correlation in agents' randomized choices. As a simple example, suppose that the common level of noise in agents' decisions varied with time. Then during periods of high noise levels, transitions between equilibria might not be especially unlikely, even as over time the overall proportion of suboptimal choices remained very small.

Of course, for the properties of these alternate models to be relevant to our current study, we would need to know that they generated qualitatively similar predictions of strategic behavior. Since little is known about how the predictions of the various models compare, we do not know whether these models can be used to address the waiting time critique in a convincing way. The point we wish to make here is more limited: namely, that the existence of very long waiting times in our basic model does not imply that the infinite horizon predictions provided by this model cannot be useful.

11.4 Model Adjustments for Finite Populations

This chapter is the first to focus directly on behavior in populations of a finite size N without taking N to infinity. This means that individual agents are no longer negligible: a change in strategy by a single agent alters the population state. In this section, we modify our earlier definitions of games and revision protocols in order to account for finite population effects. In some cases these changes are matters of modeling precision, or even of convenience. In others—most notably, that of logit evolution in potential games—these modifications are needed to ensure reversibility, which in turn allows us to obtain exact results without recourse to large population or small noise limits.

11.4.1 Finite-Population Games

When there are N agents in the population choosing strategies from the set $S = \{1, \dots, n\}$, the population state is an element of the set $\mathcal{X}^N = \{x \in X : Nx \in \mathbf{Z}^n\}$, a uniform grid in the simplex in \mathbf{R}^n . We can therefore identify an N agent *finite-population game* with its payoff function $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$, where as usual, $F_i^N(x) \in \mathbf{R}$ is the payoff to strategy i when the population state is $x \in \mathcal{X}^N$. Notice that only the values that F_i^N takes on the set $\mathcal{X}_i^N = \{x \in \mathcal{X}^N : x_i > 0\}$ are of consequence, since at the remaining states strategy i is unplayed.

State $x \in \mathcal{X}^N$ is a *Nash equilibrium* of F^N if no player can obtain a higher payoff by switching strategies:

$$[x_i > 0 \Rightarrow F_i^N(x) \geq F_j^N(x + \frac{1}{N}(e_j - e_i))] \text{ for all } i, j \in S.$$

This condition accounts for the fact that after an agent switches from strategy i to strategy j , the population has one less i player and one more j player.

Example 11.4.1. Matching without self-matching. In Example 2.2.1, we defined the population game F generated by matching in the symmetric normal form game $A \in \mathbf{R}^{n \times n}$ by $F_i(x) = \sum_{j \in S} A_{ij}x_j = (Ax)_i$, so that $F(x) = Ax$. When agents are infinitesimal, this is the only definition that makes sense. But when there are only a finite number of agents, this definition implicitly assumes that agents can be matched against themselves. To specify payoffs without self-matching, observe that when the population state is x , each strategy i player faces Nx_j opponents playing strategy $j \neq i$, but only $Nx_i - 1$ opponents playing strategy i , making the effective population state $\frac{1}{N-1}(Nx - e_i)$. The expected payoff to a

strategy i player at population state x is therefore

$$(11.17) \quad F_i^N(x) = \frac{1}{N-1}(A(Nx - e_i))_i = (Ax)_i + \frac{1}{N-1}((Ax)_i - A_{ii}). \quad \S$$

To be able to say that a continuous-population game is close to a large finite-population game, we require a notion of convergence for a sequence of finite-population games $\{F^N\}_{N=N_0}^\infty$ to a limit game $F : X \rightarrow \mathbf{R}^n$. A natural notion of convergence for such a sequence of functions is *uniform convergence*, which requires that

$$(11.18) \quad \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} |F^N(x) - F(x)| = 0.$$

Thus, equation (11.17) implies that the sequence of finite games $\{F^N\}$ generated by matching without self-matching in A converges uniformly to the limit game $F(x) = Ax$.

Exercise 11.4.2. Suppose that the sequence of finite-population games $\{F^N\}_{N=N_0}^\infty$ converges uniformly to the limit game $F : X \rightarrow \mathbf{R}^n$. Show that the Nash equilibrium correspondence of this sequence is upper hemicontinuous at infinity: if for each $N \geq N_0$, x^N is a Nash equilibrium of F^N as defined above, and if $\lim_{N \rightarrow \infty} x^N = x$, then x is a Nash equilibrium of F as defined in Chapter 2.

11.4.2 Clever Payoff Evaluation

If a population has N members, an agent who switches from strategy i to strategy j when the state is x changes the state to $x + \frac{1}{N}(e_j - e_i)$. If this agent wants to compare his current payoff $F_i^N(x)$ to the payoff he will obtain after switching, the relevant comparison is not to $F_j^N(x)$, but rather to $F_j^N(x + \frac{1}{N}(e_j - e_i))$. It is often convenient to assume that agents account for this change when deciding whether to switch strategies. Agents who do so are said to use *clever payoff evaluation*, while those who do not are said to use *simple payoff evaluation*.

To formalize this idea, we define the set of *diminished population states* by $\mathcal{X}_-^N = \{z \in \mathbf{R}_+^n : \sum_{i \in S} z_i = \frac{N-1}{N} \text{ and } Nz \in \mathbf{Z}^n\}$. Each diminished population state describes the behavior of the opponents of one member of a population of size N . Then given a game $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$, we define the *clever payoff function* $\check{F}^N : \mathcal{X}_-^N \rightarrow \mathbf{R}^n$ by

$$(11.19) \quad \check{F}_k^N(z) = F_k^N(z + \frac{1}{N}e_k).$$

Thus, the clever payoff vector $\check{F}^N(z)$ describes the current payoff opportunities of an agent whose *opponents'* behavior distribution is $z \in \mathcal{X}_-^N$. Finally, we say that an agent using

revision protocol $\rho = \rho(\pi, x)$ in game F^N is *clever* if at state x , his conditional switch rate from i to j is not $\rho_{ij}(F^N(x), x)$, but rather

$$(11.20) \quad \rho_{ij}(\check{F}^N(x - \frac{1}{N}e_i), x).$$

We will use the notation for clever payoffs introduced above in Section 11.5 when study evolution in potential games. When considering two-strategy games, where we use the notation $\chi \equiv x_1$ to refer to the proportion of agents choosing strategy 1, it is easier to describe clever choice directly, without introducing expressions (11.19) and (11.20).

Example 11.4.3. Consider the symmetric normal form game with strategy set $S = \{0, 1\}$ and payoff matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Suppose that a population of N agents are matched without self-matching to play A , as in Example 11.4.1. If we follow our convention from Section 11.2 by writing $\chi \equiv x_1 = 1 - x_0$, we can express the strategies' expected payoffs as

$$\begin{aligned} F_0^N(\chi) &= \frac{N(1-\chi)-1}{N-1}a + \frac{N\chi}{N-1}b = \frac{1}{N-1} \left(N\chi(b-a) + (N-1)a \right), \\ F_1^N(\chi) &= \frac{N(1-\chi)}{N-1}c + \frac{N\chi-1}{N-1}d = \frac{1}{N-1} \left(N\chi(d-c) + Nc - d \right). \end{aligned}$$

A simple agent prefers strategy 1 to strategy 0 whenever $F_1^N(\chi) > F_0^N(\chi)$. But a clever agent currently playing strategy 0 prefers strategy 1 whenever $F_0^N(\chi) < F_1^N(\chi + \frac{1}{N})$, while a clever agent playing strategy 1 prefers strategy 0 whenever $F_1^N(\chi) < F_0^N(\chi - \frac{1}{N})$.

The comparison made by a clever strategy 0 player at state χ is the same one made by a clever strategy 1 player at state $\chi + \frac{1}{N}$, since both players have $N(1 - \chi) - 1$ opponents playing strategy 0 and $N\chi$ playing strategy 1. §

For any fixed game A , the difference between simple and clever payoff evaluation becomes inconsequential when the population size N is large enough. But the exercises to follow show that for any given N , there are games in which simple and clever payoff evaluation always lead to opposite recommendations for play.

Exercise 11.4.4. (i) Let A be a symmetric two-strategy normal form game in which strategy 1 strictly dominates strategy 0 ($c > a$ and $d > b$). Show that under matching without self-matching, clever agents always want to switch from strategy 0 to strategy 1, and never want to switch from strategy 1 to strategy 0.

(ii) Now suppose that agents are playing the game

$$A = \begin{pmatrix} 0 & 3N - 2 \\ 1 & 3N \end{pmatrix}.$$

How will simple agents behave in this game? Provide intuition for your result. ‡

Exercise 11.4.5. Suppose that agents are matched *with* self-matching in the following Prisoner's Dilemma game

$$A = \begin{pmatrix} 3N & 0 \\ 3N + 2 & 1 \end{pmatrix}.$$

What happens in this game if agents use simple payoff evaluation? What if they use clever payoff evaluation? Provide intuition for your results. ‡

11.4.3 Committed Agents and Imitative Protocols

Our general infinite-population specification of an imitative revision protocol is

$$(11.21) \quad \rho_{ij}(\pi, x) = x_j r_{ij}(\pi, x).$$

The x_j term in (11.21) represents the random choice of whom to imitate; the x appearing as an argument of the conditional imitation rate $r_{ij}(\pi, x)$ can be used to allow for repeated sampling (see Example 5.4.7).

In a finite population setting, using the actual state x as the argument of ρ has the effect of allowing each player to imitate himself. To avoid self-imitation, we should change the second input of protocol ρ_{ij} from $x \in \mathcal{X}^N$ to $\frac{Nx - e_i}{N-1} \in \mathcal{X}^{N-1}$, which describes the distribution of strategies among the $N - 1$ opponents of the revising i player. The main effect of removing the revising agent himself from the population he samples is minor: it increases the rate at which agents switch strategies by a factor of $\frac{N}{N-1}$.

We now consider a small modification of our imitative model that has major effects on our predictions of play. On many occasions, we have noted that purely imitative protocols satisfy extinction: unused strategies are never subsequently chosen. In the present context, this means that the stochastic evolutionary process $\{X_t^N\}$ has multiple recurrent classes, and hence many stationary distributions. Suppose we alter our basic model so as to ensure that there is always at least one agent playing each strategy. Then if conditional imitation rates r_{ij} are always positive, then the process $\{X_t^N\}$ will be irreducible, and so will possess a unique stationary distribution.

The simplest and least intrusive way to accomplish this is to assume that in addition to the N standard agents, there are also n committed agents, one for each of the n strategies in S . The i th committed agent always plays strategy i ; he never receives revision opportunities. If we let the state variable $x \in \mathcal{X}^N$ represent the behavior of the standard (uncommitted) agents and exclude self-imitation, we obtain a revision protocol of the form $\rho_{ij}(\pi, \frac{Nx+1-e_i}{N+n-1})$. We will continue to express payoffs $F^N: \mathcal{X}^N \rightarrow \mathbf{R}^n$ as a function of the behavior of the standard agents. Were we to start with a function $G^N: \mathcal{X}^{N+n} \rightarrow \mathbf{R}^n$ that expresses payoffs as a function of the aggregate behavior both kinds of agents, we could obtain F^N as $F^N(x) = G^N(\frac{Nx+1}{N+n})$.

While the fraction of committed agents approaches zero as the population size becomes large, the introduction of these agents leads to completely different predictions of infinite horizon behavior.

Example 11.4.6. Suppose that a standard agent who receives a revision opportunity picks an opponent at random and imitates him: that is, let $r_{ij} \equiv 1$, so that $\rho_{ij}(\pi, x) = x_j$. Without committed agents, the resulting stochastic evolutionary process converges with probability 1 to one of the n pure states e_1, \dots, e_n . But with a single committed agent for each strategy, the process $\{X_t^N\}$ is irreducible, and so admits a unique stationary distribution μ^N . In fact, the process is reversible, and its stationary distribution is the uniform distribution on \mathcal{X}^N .

Since the only positive probability transitions are between pairs of adjacent states, establishing reversibility with a uniform stationary distribution, reduces to verifying that $P_{xy}^N = P_{yx}^N$ for all such pairs $x \in \mathcal{X}^N$ and $y = x + \frac{1}{N}(e_j - e_i)$. Letting $z = x - \frac{1}{N}e_i = y - \frac{1}{N}e_j \in \mathcal{X}_-^N$ represent the behavior of the revising player's opponents, we compute as follows:

$$\begin{aligned}
(11.22) \quad P_{xy}^N &= x_i \cdot \rho_{ij}(\pi, \frac{Nx+1-e_i}{N+n-1}) \\
&= x_i \cdot \frac{Nx_j + 1}{N + n - 1} \\
&= \frac{Nz_i + 1}{N} \cdot \frac{Nz_j + 1}{N + n - 1} \\
&= \frac{Nz_j + 1}{N} \cdot \frac{Nz_i + 1}{N + n - 1} \\
&= y_j \cdot \frac{Ny_i + 1}{N + n - 1} \\
&= y_j \cdot \rho_{ji}(\pi, \frac{Ny+1-e_j}{N+n-1}) \\
&= P_{yx}^N.
\end{aligned}$$

Moreover, it follows from Exercise 10.1.1 that the mass on each state $x \in \mathcal{X}^N$ is $\mu_x^N = \frac{1}{\#\mathcal{X}^N} =$

$(\frac{N+n-1}{n-1})^{-1}$. §

Example 11.4.7. Suppose that N standard agents and $n = 2$ committed agents play the two-strategy game F^N , employing an imitative protocol of the form

$$\begin{aligned}\rho_{01}(\pi, \chi) &= (1 - \chi) r_{01}(\pi, \chi), \\ \rho_{10}(\pi, \chi) &= \chi r_{10}(\pi, \chi).\end{aligned}$$

As usual, $\chi \in [0, 1]$ denotes the proportion of agents choosing strategy 1. If r_{01} and r_{10} are positive-valued, the resulting evolutionary process is irreducible. Adjusting equation (11.11) to account for the presence of the committed agents, we find that the stationary distribution of the evolutionary process is described by

$$\begin{aligned}(11.23) \quad \frac{\mu_\chi^N}{\mu_0^N} &= \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N} \\ &= \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N}}{\frac{j}{N}} \cdot \frac{\frac{1}{R} \rho_{01}(F^N(\frac{j-1}{N}), \frac{j}{N+1})}{\frac{1}{R} \rho_{10}(F^N(\frac{j}{N}), \frac{j}{N+1})} \\ &= \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N}}{\frac{j}{N}} \cdot \frac{\frac{j}{N+1}}{\frac{N-j+1}{N+1}} \cdot \frac{r_{01}(F^N(\frac{j-1}{N}), \frac{j}{N+1})}{r_{10}(F^N(\frac{j}{N}), \frac{j}{N+1})} \\ &= \prod_{j=1}^{N\chi} \frac{r_{01}(F^N(\frac{j-1}{N}), \frac{j}{N+1})}{r_{10}(F^N(\frac{j}{N}), \frac{j}{N+1})}.\end{aligned}$$

Thus, the stationary distribution weights can be expressed in terms of the conditional imitation rates alone. §

We will consider further consequences of adding committed agents to stochastic imitative models in the next section.

Exercise 11.4.8. Exercise 11.2.3 presented the stationary distribution for the n strategy Toss and Switch process, under which revising agents randomize uniformly over the available strategies. Write down the mean dynamics generated by this process and the one defined in Example 11.4.6, and use these dynamics to provide an intuitive explanation for the differences between the stationary distributions of the two processes. ‡

11.5 Exponential Protocols and Potential Games

The analysis of infinite-horizon behavior in two-strategy games is greatly simplified by the fact that in this context, the Markov process $\{X_t^N\}$ is a birth and death chain, and hence reversible. Beyond two-strategy settings, the only other context in which $\{X_t^N\}$ is known to be reversible is that of potential games when agents employ one of a number of revision protocols under which conditional switch rates are exponential functions of payoffs.

11.5.1 Finite-Population Potential Games

Establishing reversibility requires that we introduce a finite population definition of potential games. Our definitions and characterizations of continuous-population potential games in Chapter 3 rely on tools from calculus. In this section, we introduce definitions of full potential games and potential games for finite populations, and establish basic connections with the continuous-population definitions from Chapter 3 by means of convergence results. Much of this section is comprised of exercises establishing various links between potential games defined in the two contexts.

In Section 3.1, we introduced the notion of a full potential function, a function whose partial derivatives equal the payoffs of the underlying game. To ensure the existence of these partial derivatives, the full potential function was defined on the positive orthant, a full-dimensional set in \mathbf{R}^n .

To define full potential games for finite populations, we introduce a discrete analogue of this device. We define the set of *diminished population states* by $\mathcal{X}_-^N = \{z \in \mathbf{R}_+^n : \sum_{i \in S} z_i = \frac{N-1}{N} \text{ and } Nz \in \mathbf{Z}^n\}$. Each diminished population state describes the behavior of the opponents of one member of a population of size N . We then call a finite-population game $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$ a *full potential game* if it admits a *full potential function*: a function $f^N : \mathcal{X}^N \cup \mathcal{X}_-^N \rightarrow \mathbf{R}$ such that

$$(11.24) \quad F_i^N(x) = f^N(x) - f^N(x - \frac{1}{N}e_i) \text{ for all } x \in \mathcal{X}^N \text{ and } i \in S.$$

One can describe condition (11.24) as requiring the payoff to strategy i to be determined by the i th “discrete partial derivative” of the function $\frac{1}{N}f^N$.

Exercise 11.5.1. Formalize and verify this last statement. ‡

Exercise 11.5.2. Matching in common interest games. Let the symmetric normal form game $A \in \mathbf{R}^{n \times n}$ be a common interest game (i.e., $A = A'$). In Example 3.1.2, we saw that when

an infinite population is matched to play A , the resulting population game $F(x) = Ax$ is a full potential game with full potential function $f(x) = \frac{1}{2}x'Ax$.

- (i) Let $F^N(x) = Ax$ be the finite-population game obtained via matching with self-matching. Verify that F^N is a full potential game with full potential function

$$f^N(x) = \frac{1}{2} \left(Nx'Ax + \sum_{k \in S} A_{kk}x_k \right).$$

- (ii) Now let F^N be the finite-population game obtained via matching without self-matching, as defined in equation (11.17). Verify that this F^N is also a full potential game, but with full potential function

$$f^N(x) = \frac{1}{2} \frac{N}{N-1} \left(Nx'Ax - \sum_{k \in S} A_{kk}x_k \right).$$

Note that in each case, the functions $\frac{1}{N}f^N$ converge to f as N grows large. ‡

Exercise 11.5.3. Congestion games. In Example 2.2.5, we defined a continuous-population congestion game $F : \mathbf{R}_+^n \rightarrow \mathbf{R}$ by

$$F_i(x) = - \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x)).$$

where Φ_i is the set of facilities (or links) used by strategy (or path) i , $c_\phi : \mathbf{R}_+ \rightarrow \mathbf{R}$ is the cost function of facility ϕ , and $u_\phi = \sum_{i \in S: \phi \in \Phi_i} x_i$ is the utilization level of facility ϕ . Example 3.1.4 showed that this game is a full potential game with full potential function

$$f(x) = - \sum_{\phi \in \Phi} \int_0^{u_\phi(x)} c_\phi(z) dz.$$

Apart from the change in domain, the definition of the finite-population congestion game $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$ is identical to that of F . Show that F^N is also a full potential game, with full potential function

$$f^N(x) = - \sum_{\phi \in \Phi_i} \sum_{k=1}^{Nu_\phi(x)} c_\phi\left(\frac{k}{N}\right). \ddagger$$

The continuous-population definition of a potential game from Section 3.2 uses a potential function that is only defined on the original set of population states, and that only determines the game's relative payoffs. Extending this notion to the present setting,

we call a finite-population game $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$ a *potential game* if it admits a *potential function*: in this case, a function $f^N : \mathcal{X}^N \rightarrow \mathbf{R}$ such that

$$(11.25) \quad F_j^N(x + \frac{1}{N}(e_j - e_i)) - F_i^N(x) = f^N(x + \frac{1}{N}(e_j - e_i)) - f^N(x) \text{ for all } x \in \mathcal{X}_i^N \text{ and } i, j \in S.$$

Equation (11.25) requires that when an agent switches from strategy i to strategy j , the change in his payoff is equal to the change in potential. This definition closely resembles the definition of normal form potential games from Exercise 3.2.10. We investigate this connection in Exercise 11.5.17(ii) below.

Exercise 11.5.4. Two-strategy games. Let F^N be a two-strategy game with strategy set $S = \{0, 1\}$. Verify that F^N is a finite-population potential game with potential function

$$f^N(x) = \sum_{j=1}^{Nx_1} \left(F_1^N(x) - F_0^N(x + \frac{1}{N}(e_0 - e_1)) \right). \quad \ddagger$$

Exercise 11.5.5. Equilibrium and evolutionary dynamics for finite-population potential games.

- (i) Suppose that F^N is a potential game with potential function f^N . Show that $x \in \mathcal{X}^N$ is a Nash equilibrium of F^N (as defined in Section 11.4.1) if and only if x is a local maximizer of f^N . (Be sure to define what it means to be a “local maximizer of f^N ”.)
- (ii) Argue that if agents in a finite-population potential game F^N switch to better-performing strategies sequentially, the population state will converge to a Nash equilibrium of F^N after a finite number of switches. \ddagger

Theorem 3.2.12 showed that despite first appearances, the continuous-population definitions of full potential games and potential games are essentially equivalent. Theorem 11.5.6 shows that this equivalence persists when populations are finite. Fortunately, the proof in this discrete setting requires much less effort.

Theorem 11.5.6. *Let F^N be a potential game with potential function $f^N : \mathcal{X}^N \rightarrow \mathbf{R}$. Then there is an extension $\tilde{f}^N : \mathcal{X}^N \cup \mathcal{X}_-^N \rightarrow \mathbf{R}$ of f that is a full potential function for F^N . Therefore, F^N is a potential game if and only if it is a full potential game.*

Proof. Let $\tilde{f}(x) = f(x)$ when $x \in \mathcal{X}^N$, and for each $z \in \mathcal{X}_-^N$, let

$$(11.26) \quad \tilde{f}^N(z) = \tilde{f}^N(z + \frac{1}{N}e_1) - F_1^N(z + \frac{1}{N}e_1).$$

Rearranging this expression shows that the condition (11.24) holds when $i = 1$. To verify that it holds for an arbitrary strategy i , use equations (11.25) and (11.26) (with $z = x - \frac{1}{N}e_i$)

to compute as follows:

$$\begin{aligned}
F_i^N(x) &= F_1^N(x + \frac{1}{N}(e_1 - e_i)) - \tilde{f}^N(x + \frac{1}{N}(e_1 - e_i)) + \tilde{f}^N(x) \\
&= (\tilde{f}^N(x + \frac{1}{N}(e_1 - e_i)) - \tilde{f}^N(x - \frac{1}{N}e_i)) - \tilde{f}^N(x + \frac{1}{N}(e_1 - e_i)) + \tilde{f}^N(x) \\
&= \tilde{f}^N(x) - \tilde{f}^N(x - \frac{1}{N}e_i).
\end{aligned}$$

Since it is clear that every full potential game is a potential game, the result is proved. ■

By virtue of Theorem 11.5.6, there is no need for the term “full potential game” in the finite-population setting. (We cannot quite do this in the continuous-population setting, where potential games and full potential games have different domains.) Of course, potential functions and full potential functions must still be distinguished.

All of the developments in this section suggest that the finite-population and continuous-population definitions of potential games are different expressions of the same idea. To formalize this intuition, let $\{F^N\}_{N=N_0}^\infty$ be a sequence of finite-population potential games with full potential functions $\{f^N\}_{N=N_0}^\infty$, and let the function $f: \mathbf{R}_+^n \rightarrow \mathbf{R}$ be C^1 . We say that the sequence of rescaled full potential functions $\{\frac{1}{N}f^N\}$ is *Lipschitz convergent* with limit f if there is a vanishing sequence $\{K^N\}_{N=N_0}^\infty$ such that

$$(11.27) \quad \left| \left(\frac{1}{N}f^N(x) - f(x) \right) - \left(\frac{1}{N}f^N(y) - f(y) \right) \right| \leq K^N |x - y| \quad \text{for all } x, y \in X^N \cup X_-^N.$$

In words, condition (11.27) requires the differences $\frac{1}{N}f^N - f$ to be Lipschitz continuous, with Lipschitz constants that approach zero as N grows large. The rescaling of the potential functions reflects the fact that the values of f^N are of order N , and so must be shrunk by a factor of $\frac{1}{N}$ for convergence to be feasible.

Theorem 11.5.7 shows that a sequence of finite-population potential games $\{F^N\}$ converges uniformly to a continuous full potential game if and only if the corresponding sequence of rescaled full potential functions $\{\frac{1}{N}f^N\}$ is Lipschitz convergent. Exercise 11.5.10 provides the analogous result for potential functions defined on X^N only.

Theorem 11.5.7. *Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of finite-population potential games with full potential functions $\{f^N\}_{N=N_0}^\infty$.*

- (i) *Suppose that $\{\frac{1}{N}f^N\}$ is Lipschitz convergent with C^1 limit $f: \mathbf{R}_+^n \rightarrow \mathbf{R}$, and define $F: \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ by $F(x) = \nabla f(x)$. Then $\{F^N\}$ converges uniformly to $F|_X$.*
- (ii) *Suppose that $\{F^N\}$ converges uniformly to $F: X \rightarrow \mathbf{R}^n$, and that F admits a full potential function $f: \mathbf{R}_+^n \rightarrow \mathbf{R}$, in the sense that $F(x) = \nabla f(x)$ for all $x \in X$. Then $\{\frac{1}{N}f^N\}$ is Lipschitz convergent with limit f .*

Exercise 11.5.8. Prove Theorem 11.5.7. (Hint: If $f: \mathbf{R}_+^n \rightarrow \mathbf{R}$ is C^1 , the Mean Value Theorem implies that there is a point y^N on the line segment from x to $x - \frac{e_i}{N}$ such that $f(x) - f(x - \frac{e_i}{N}) = \frac{1}{N} \frac{\partial f}{\partial x_i}(y^N)$. Combine this fact with compactness arguments.) ‡

Exercise 11.5.9. (i) Lipschitz convergence is unaffected by the addition of constant terms to the potential functions f^N and f . Show that if we normalize these functions so that $f^N(e_1) = f(e_1) = 0$, and if the resulting sequence $\{\frac{1}{N}f^N\}$ is Lipschitz convergent with limit f , then the sequence $\{\frac{1}{N}f^N\}$ converges uniformly to f .

(ii) Show that the uniform convergence of $\{\frac{1}{N}f^N\}$ to f does not imply the Lipschitz convergence of $\{\frac{1}{N}f^N\}$ to f . (Hint: Suppose that $f^N(e_1) = 0$ and that $f^N(x) = \sqrt{N}$ for $x \in \mathcal{X}^N \cup \mathcal{X}^N - \{e_1\}$.) ‡

Exercise 11.5.10. Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of finite-population potential games with potential functions $\{f^N\}_{N=N_0}^\infty$. (Thus f^N has domain \mathcal{X}^N .)

(i) Suppose that $\{\frac{1}{N}f^N\}$ is Lipschitz convergent with C^1 limit $f: X \rightarrow \mathbf{R}$, and define $F: X \rightarrow \mathbf{R}^n$ by $F(x) = \nabla f(x)$. Show that $\{\Phi F^N\}$ converges uniformly to F , where $\Phi = I - \frac{1}{n}\mathbf{1}\mathbf{1}'$ is the orthogonal projection of \mathbf{R}^n onto TX .

(ii) Suppose that $\{\Phi F^N\}$ converges uniformly to $F: X \rightarrow \mathbf{R}^n$, and that F admits potential function $f: X \rightarrow \mathbf{R}$. Show that $\{\frac{1}{N}f^N\}$ is Lipschitz convergent with limit f . ‡

11.5.2 Exponential Revision Protocols

We now introduce exponential revision protocols. These protocols will be divided into two classes according to whether agents choose candidate strategies directly or through imitation. We consider the direct version first.

Definition. We call $\rho: \mathbf{R}^n \rightarrow \mathbf{R}_+^{n \times n}$ a direct exponential protocol with noise level η if

$$(11.28) \quad \rho_{ij}(\pi) = \frac{\exp(\eta^{-1}\psi(\pi_i, \pi_j))}{d_{ij}(\pi)},$$

where the functions $\psi: \mathbf{R}^2 \rightarrow \mathbf{R}$ and $d: \mathbf{R}^n \rightarrow (0, \infty)^{n \times n}$ satisfy

$$(11.29) \quad \psi(\pi_i, \pi_j) - \psi(\pi_j, \pi_i) = \pi_j - \pi_i, \text{ and}$$

$$(11.30) \quad d_{ij}(\pi) = d_{ji}(\pi).$$

The definition of a direct exponential protocol is broad, in that it allows agents choice procedures to take a broad range of qualitative forms, depending on the specification of the functions ψ and d . For instance, under condition (11.29), the function ψ can reflect these following sorts of dependence of switch rates on payoffs:

positive dependence on candidate payoff:	$\psi(\pi_i, \pi_j) = \pi_j$
negative dependence on current payoff:	$\psi(\pi_i, \pi_j) = -\pi_i$
positive dependence on payoff difference:	$\psi(\pi_i, \pi_j) = \frac{1}{2}(\pi_j - \pi_i)$
positive dependence on positive payoff difference:	$\psi(\pi_i, \pi_j) = [\pi_j - \pi_i]_+$
negative dependence on negative payoff difference:	$\psi(\pi_i, \pi_j) = -[\pi_j - \pi_i]_-$

Thus, exponential protocols allow switching rates to be determined by the desirability of the payoff of the candidate strategy, by dissatisfaction with the payoff of the current strategy, or with comparisons of the payoffs of both strategies. The information requirements of an exponential protocol can be quite modest: for instance, the second option above only requires the agent to know his own current payoff.

Different choices of the function d can be used to reflect different reference groups that agents employ when considering a switch. The symmetry condition (11.30) requires that when an i player considers switches to strategy j , he employs the same comparison group as a j player who considers switching to i . Focus, for instance, on the case in which $\psi(\pi_i, \pi_j) = \pi_j$. If we set

$$d_{ij}(\pi) = \sum_{k \in S} \exp(\eta^{-1} \pi_k),$$

so that agents use the full set of strategies as the comparison group, then we obtain the logit protocol

$$\rho_{ij}(\pi) = \frac{\exp(\eta^{-1} \pi_j)}{\sum_{k \in S} \exp(\eta^{-1} \pi_k)}.$$

If instead $d_{ij}(\pi) = \exp(\eta^{-1} \pi_i) + \exp(\eta^{-1} \pi_j)$, so that the comparison group only contains the current and candidate strategies, we instead obtain the *pairwise logit protocol*,

$$\rho_{ij}(\pi) = \frac{\exp(\eta^{-1} \pi_j)}{\exp(\eta^{-1} \pi_i) + \exp(\eta^{-1} \pi_j)}.$$

Of course, there are many further choices of ψ and ρ that satisfy the requirements above.

Exercise 11.5.11. Show that ρ is a direct exponential protocol if and only if

$$(11.31) \quad \eta \log \frac{\rho_{ij}(\pi)}{\rho_{ji}(\pi)} = \pi_j - \pi_i \quad \text{for all } \pi \in \mathbf{R}^n. \quad \ddagger$$

Under a direct exponential protocol, it is as if agents choose candidate strategies from a list of all the available strategies. Under an imitative exponential protocol, agents instead

choose candidate strategies by observing the choice of an opponent.

Definition. We call $\rho : \mathbf{R}^n \times X \rightarrow \mathbf{R}_+^{n \times n}$ an imitative exponential protocol with noise level η if

$$(11.32) \quad \rho_{ij}(\pi, x) = x_j \frac{\exp(\eta^{-1}\psi(\pi_i, \pi_j))}{d_{ij}(\pi, x)},$$

where the functions $\psi : \mathbf{R}^2 \rightarrow \mathbf{R}$ and $d : \mathbf{R}^n \times X \rightarrow (0, \infty)^{n \times n}$ satisfy conditions (11.29) and

$$(11.33) \quad d_{ij}(\pi, x) = d_{ji}(\pi, x).$$

Of course, we have the same flexibility in specifying imitative exponential protocols as we did in specifying the direct protocols considered earlier. For instance, the case in which $\psi(\pi_i, \pi_j) = \pi_j$. If we set

$$d_{ij}(\pi, x) = \sum_{k \in S} x_k \exp(\eta^{-1}\pi_k),$$

the i-logit protocol from Example 11.1.5:

$$\rho_{ij}(\pi, x) = \frac{x_j \exp(\eta^{-1}\pi_j)}{\sum_{k \in S} x_k \exp(\eta^{-1}\pi_k)}.$$

If instead $d_{ij}(\pi) = x_i \exp(\eta^{-1}\pi_i) + x_j \exp(\eta^{-1}\pi_j)$, we obtain the *pairwise i-logit protocol*,

$$\rho_{ij}(\pi, x) = \frac{x_j \exp(\eta^{-1}\pi_j)}{x_i \exp(\eta^{-1}\pi_i) + x_j \exp(\eta^{-1}\pi_j)}.$$

11.5.3 Reversibility and Stationary Distributions

We now prove our main results on infinite horizon behavior in potential games. The first of these, Theorem 11.5.12, shows that if clever agents employ a direct exponential protocol, the stochastic evolutionary process $\{X_t^N\}$ is reversible. The stationary distribution weight $\mu^N(x)$ is proportional to the product of two terms. The first term is a multinomial coefficient, and represents the number of ways of assigning the N agents to strategies in S so that each strategy i is played by precisely Nx_i agents. The second term is an exponential function of the value of potential at state x . Thus, the value of μ_x^N balances the value of potential at state x with the likelihood that state x would arise were agents assigned to strategies at random.

Theorem 11.5.12. Let F^N be a finite population potential game with potential function f^N , and suppose that agents are clever and follow a direct exponential protocol with noise level η . Then the stochastic evolutionary process $\{X_t^N\}$ is reversible with stationary distribution

$$(11.34) \quad \mu_x^N = \frac{1}{K^N} \frac{N!}{\prod_{k \in S} (Nx_k)!} \exp(\eta^{-1} f^N(x))$$

for $x \in \mathcal{X}^N$, where K^N is determined by the requirement that $\sum_{x \in \mathcal{X}^N} \mu_x^N = 1$.

Proof. To verify the reversibility condition (11.10), it is enough to check that the equality $\mu_x^N P_{xy}^N = \mu_y^N P_{yx}^N$ holds for pairs of states that are adjacent, in the sense that $y = x + \frac{1}{N}(e_j - e_i)$. To do so, note first that $z = x - \frac{1}{N}e_i = y - \frac{1}{N}e_j \in \mathcal{X}_-^N$ represents both the distribution of opponents for an i player at state x , and the distribution of opponents of a j player at state y . Thus, in both cases, a clever player who is revising will consider the payoff vector $\check{F}^N(z)$ defined by $\check{F}_k^N(z) = F_k^N(z + \frac{1}{N}e_k)$ for all $k \in S$, as introduced in equation (11.19).

Now by the definition of the potential function f^N ,

$$(11.35) \quad f^N(y) - f^N(x) = F_j^N(y) - F_i^N(x) = \check{F}_j^N(z) - \check{F}_i^N(z).$$

Therefore, writing $\check{\pi} = \check{F}^N(z)$, and using equations (11.35), (11.29), and (11.30) to obtain the third equality below, we obtain

$$\begin{aligned} \mu_x^N P_{xy}^N &= \mu_x^N \cdot x_i \rho_{ij}(\check{\pi}) \\ &= \frac{1}{K^N} \frac{N!}{\prod_{k \in S} (Nx_k)!} \exp(\eta^{-1} f^N(x)) \cdot x_i \frac{\exp(\eta^{-1} \psi(\check{\pi}_i, \check{\pi}_j))}{d_{ij}(\check{\pi})} \\ &= \frac{1}{K^N} \frac{(N-1)!}{\prod_{k \in S} (Nz_k)!} \exp(\eta^{-1} (f^N(y) - \check{\pi}_j + \check{\pi}_i)) \frac{\exp(\eta^{-1} (\psi(\check{\pi}_j, \check{\pi}_i) + \check{\pi}_j - \check{\pi}_i))}{d_{ji}(\check{\pi})} \\ &= \frac{1}{K^N} \frac{N!}{\prod_{k \in S} (Ny_k)!} \exp(\eta^{-1} f^N(y)) \cdot y_j \frac{\exp(\eta^{-1} \psi(\check{\pi}_j, \check{\pi}_i))}{d_{ji}(\check{\pi})} \\ &= \mu_y^N \cdot y_j \rho_{ji}(\check{\pi}) \\ &= \mu_y^N P_{yx}^N. \blacksquare \end{aligned}$$

Theorem 11.5.13 considers evolution in potential games when clever agents employ imitative exponential protocols, assuming that there is one committed agent playing each of the n strategies in S . The stochastic evolutionary process $\{X_t^N\}$ is again reversible. This time, the stationary distribution weight μ_x^N is directly proportional to $\exp(\eta^{-1} f^N(x))$; no

additional terms are required. Section 12.2 will explore the consequences of this distinction for predictions about infinite horizon behavior.

Theorem 11.5.13. *Let F^N be a finite population potential game with potential function f^N . Suppose that there are N clever agents who follow an imitative exponential protocol with noise level η , and that there is one committed agent for each strategy. Then the stochastic evolutionary process $\{X_t^N\}$ is reversible with stationary distribution*

$$(11.36) \quad \mu_x^N = \frac{1}{\kappa^N} \exp(\eta^{-1} f^N(x))$$

for $x \in \mathcal{X}^N$, where κ^N is determined by the requirement that $\sum_{x \in \mathcal{X}^N} \mu_x^N = 1$.

Proof. As before, let $y = x + \frac{1}{N}(e_j - e_i)$, let $z = x - \frac{1}{N}e_i = y - \frac{1}{N}e_j \in \mathcal{X}_-^N$, and let $\check{\pi} = \check{F}^N(z)$. Then noting as in equation (11.22) that

$$x_i \frac{Nx_j + 1}{N + n - 1} = \frac{Nz_i + 1}{N} \cdot \frac{Nz_j + 1}{N + n - 1} = \frac{Nz_j + 1}{N} \cdot \frac{Nz_i + 1}{N + n - 1} = y_j \cdot \frac{Ny_i + 1}{N + n - 1},$$

equations (11.35), (11.29), and (11.33) yield

$$\begin{aligned} \mu_x^N P_{xy}^N &= \mu_x^N \cdot x_i \rho_{ij}(\check{\pi}, \frac{Nx+1-e_i}{N+n-1}) \\ &= \frac{1}{\kappa^N} \exp(\eta^{-1} f^N(x)) \cdot x_i \frac{Nx_j + 1}{N + n - 1} \frac{\exp(\eta^{-1} \psi(\check{\pi}_i, \check{\pi}_j))}{d_{ij}(\check{\pi}, \frac{Nx+1-e_i}{N+n-1})} \\ &= \frac{1}{\kappa^N} \exp(\eta^{-1} (f^N(y) - \check{\pi}_j + \check{\pi}_i)) \cdot x_i \frac{Nx_j + 1}{N + n - 1} \frac{\exp(\eta^{-1} (\psi(\check{\pi}_j, \check{\pi}_i) + \check{\pi}_j - \check{\pi}_i))}{d_{ij}(\check{\pi}, \frac{Nz+1}{N+n-1})} \\ &= \frac{1}{\kappa^N} \exp(\eta^{-1} f^N(y)) \cdot y_j \frac{Ny_i + 1}{N + n - 1} \frac{\exp(\eta^{-1} \psi(\check{\pi}_j, \check{\pi}_i))}{d_{ji}(\check{\pi}, \frac{Ny+1-e_j}{N+n-1})} \\ &= \mu_y^N \cdot y_j \rho_{ji}(\check{\pi}, \frac{Ny+1-e_j}{N+n-1}) \\ &= \mu_y^N P_{yx}^N. \blacksquare \end{aligned}$$

Exercise 11.5.14. Describe the qualitative differences between stationary distributions (11.34) and (11.36) when the noise level η is large. Building on your answer to Exercise 11.4.8, relate the differences between the stationary distributions to differences between the process's mean dynamics. ‡

Exercise 11.5.15. Show that if we replace clever payoff evaluation with simple payoff evaluation in either of the previous two theorems, the resulting stochastic evolutionary process is not reversible. ‡

Exercise 11.5.16. Use the birth and death chain formula from Theorem 11.2.4 to give alternate proofs of Theorems 11.5.12 and 11.5.13 in the case where F^N is a two-strategy game. (Recall from Exercise 11.5.4 that every two-strategy game is a potential game.) ‡

Exercise 11.5.17. This exercise provides an alternate proof of Theorem 11.5.12 by way of normal form potential games. Recall from Exercise 3.2.10 that the N player normal form game $U = (U^1, \dots, U^N)$ with strategy sets S^1, \dots, S^N is a potential game if it admits a potential function $V : \prod_p S^p \rightarrow \mathbf{R}$, which is a function satisfying

$$U^p(\hat{s}^p, s^{-p}) - U^p(s) = V(\hat{s}^p, s^{-p}) - V(s) \text{ for all } s \in \prod_p S^p, \hat{s}^p \in S^p, \text{ and } p \in \mathcal{P}.$$

- (i) Define a stochastic evolutionary process $\{Y_t^N\}$ on the set of pure strategy profiles $\prod_p S^p$ by assuming that each of the N players receives revision opportunities at rate 1, and that each follows a direct exponential protocol with noise level η when such opportunities arise. Show that if U is a normal form potential game with potential function V , then the process $\{Y_t^N\}$ is reversible, and that its stationary distribution ν^N is defined by $\nu^N(s) \propto \exp(\eta^{-1}V(s))$.

Now suppose that players in the game U are *indistinguishable*: all N players share the same strategy set $S^p = S \equiv \{1, \dots, n\}$, and each player's payoffs are defined by the same function of the player's own strategy and the overall distribution of strategies $\xi(s) \in \mathcal{X}^N$, defined by $\xi_i(s) = \#\{p \in \{1, \dots, N\} : s^p = i\}/N$. This means that there is a finite-population game $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$ that represents U , in the sense that $U^p(s) = F_{s^p}^N(\xi(s))$ for all $s \in S$.

- (ii) Suppose again that U is a normal form potential game with potential function V . Show that if players in U are indistinguishable, then V is measurable with respect to \mathcal{X}^N , in the sense that $V(\hat{s}) = V(s)$ whenever $\xi(\hat{s}) = \xi(s)$. Furthermore, show that the finite-population game F^N introduced above is a potential game, and that its potential function $f^N : \mathcal{X}^N \rightarrow \mathbf{R}$ is given by $f^N(x) = V(s)$, where s is any strategy profile in $\xi^{-1}(x)$.
- (iii) Let $\{Y_t^N\}$ be the evolutionary process on \mathcal{X}^N defined in part (i). Show that under the assumptions of part (ii), the process $\{\xi(Y_t^N)\}$ is a Markov process on \mathcal{X}^N with the same jump rates and transition probabilities as the process $\{X_t^N\}$ from Theorem 11.5.12. Finally, using part (i), prove directly that the stationary distribution of $\{\xi(Y_t^N)\}$ is given by equation (11.34). ‡

In Chapter 12, we turn our attention to stochastic stability analysis, which concerns the limiting behavior of the stationary distributions as the noise level becomes small or the population size becomes large. Since the reversible settings we focused on in this chapter allow the stationary distribution to be described explicitly, we not only will compute the

limiting stationary distributions for these settings, but also will establish precise results on the asymptotics of these distributions. Once we move beyond the reversible cases, simple expressions for the stationary distributions are no longer available, and we will need to introduce new techniques to determine the stochastically stable states.

Appendix

11.A Long Run Behavior of Markov Chains and Processes

In Appendix 10.C, we defined (discrete-time) Markov chains and (continuous-time) Markov processes on a countable state space, and showed how to construct them on an appropriate probability space $(\Omega, \mathcal{F}, \mathbb{P})$. In the present appendix, we turn to the long run behavior of Markov chains and processes that take values in a *finite* state space. Apart from the issue of periodicity, which is particular to the discrete-time setting, the discrete-time and continuous-time theories are quite similar, although certain concepts have somewhat different formulations in the two cases.

In the remainder of this section, $\{X_t\}_{t=0}^{\infty}$ denotes a Markov chain with transition matrix $P \in \mathbf{R}^{\mathcal{X} \times \mathcal{X}}$ on the finite state space \mathcal{X} , while $\{X_t\}_{t \geq 0}$ denotes a Markov process on \mathcal{X} with transition matrix $P \in \mathbf{R}^{\mathcal{X} \times \mathcal{X}}$ and jump rate vector $\lambda \in \mathbf{R}^{\mathcal{X}}$. The shorthand $\{X_t\}$ is used to refer to both cases. When we need to refer to initial conditions, we use the notations $\mathbb{P}_x(\cdot) \equiv \mathbb{P}(\cdot | X_0 = x)$ and $\mathbb{P}_\pi(\cdot) \equiv \sum_{x \in \mathcal{X}} \pi_x \mathbb{P}(\cdot | X_0 = x)$ to describe the behavior of $\{X_t\}$ when it is run from initial state $x \in \mathcal{X}$ or from initial distribution $\pi \in \mathbf{R}_+^{\mathcal{X}}$.

11.A.1 Communication, Recurrence, and Irreducibility

To begin our study of long run behavior of the Markov chains and processes, we introduce a partial order on the finite set \mathcal{X} that describes feasible multi-step transitions under the transition matrix P . We say that state y is *accessible* from state x , and write $x \rightsquigarrow y$, if for some $n \geq 0$ there is a sequence of states $x = x_0, x_1, \dots, x_n = y$ such that $P_{x_{i-1}, x_i} > 0$ for all $i \in \{1, \dots, n\}$. We allow $n = 0$ to ensure that each state is accessible from itself. We write $x \leftrightarrow y$ to indicate that x and y are mutually accessible, in the sense that $x \rightsquigarrow y$ and $y \rightsquigarrow x$.

Accessibility defines a partial order on the set \mathcal{X} . The equivalence classes under this order, referred to as *communication classes*, are the maximal sets of (pairwise) mutually accessible states.

To identify the states which $\{X_t\}$ can visit in the long run, we first define a set of states




Figure 11.A.1: Feasible transitions between the states of a Markov chain. Members of the same communication class are the same color. Members of recurrent classes are shades of blue.

$R \subseteq X$ to be *closed* if the process cannot leave it:

$$[x \in R, x \rightsquigarrow y] \Rightarrow y \in R.$$

If we assume that R is a communication class, then R is closed if and only if it is minimal under the partial order \rightsquigarrow , where we view this order as “pointing downward” (see Figure 11.A.1). Once $\{X_t\}$ enters a closed communication class, it remains in the class forever.

Example 11.A.1. Suppose that $\{X_t\}$ has state space $X = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$ and transition matrix

$$P = \begin{pmatrix} .4 & .3 & .2 & .1 & 0 & 0 & 0 & 0 & 0 \\ .5 & .1 & 0 & 0 & .3 & .1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & .8 & .1 & .1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & .9 & 0 & .1 \end{pmatrix}.$$

The feasible transitions of $\{X_t\}$ are represented in Figure 11.A.1. $\{X_t\}$ has six communication classes: $\{1, 2\}$, $\{3\}$, $\{4, 5\}$, $\{6\}$, and $\{7, 8, 9\}$. Of these, $\{3\}$, $\{4, 5\}$, and $\{7, 8, 9\}$ are closed. §

If the process $\{X_t\}$ begins in a communication class that is not closed (for example, the class $\{1, 2\}$ in Example 11.A.1), it is possible for it to remain in this class for an arbitrarily long finite amount of time. But with probability 1, the process will leave the class after

some finite amount of time, and once it leaves it can never return. Of course, once $\{X_t\}$ enters a closed communication class, it remains in the class forever.

To express these ideas more precisely, call state x *transient* if $\mathbb{P}_x(\{t : X_t = x\} \text{ is unbounded}) = 0$, and call state x *recurrent* if $\mathbb{P}_x(\{t : X_t = x\} \text{ is unbounded}) = 1$. Then we have

Theorem 11.A.2. *Let $\{X_t\}$ be a Markov chain or Markov process on a finite set X . Then*

- (i) *Every state in X is either transient or recurrent.*
- (ii) *A state is recurrent if and only if it is a member of a closed communication class.*

In light of this result, a closed communication class R is commonly referred to as a *recurrent class*.

By virtue of Theorem 11.A.2, we can understand the infinite horizon behavior of the process $\{X_t\}$ by focusing on cases in which the entire state space X forms a single recurrent class. (In effect, we are focusing on the behavior of our original process after it enters a recurrent class.) When all of X forms a single recurrent class, $\{X_t\}$ is said to be *irreducible*.

As a simple illustration of irreducibility, we introduce an example that we will revisit throughout this appendix.

Example 11.A.3. Birth and death chains. A *birth and death chain* is a Markov chain $\{X_t\}_{t=0}^{\infty}$ on the state space $X = \{0, 1, \dots, N\}$ (or some other finite set endowed with a linear order) under which all transitions move the state one step to the right, move the state one step to the left, or leave the state unchanged. It follows that there are vectors $p, q \in \mathbf{R}^X$ with $p_N = q_0 = 0$ such that

$$(11.37) \quad P_{ij} = \begin{cases} p_i & \text{if } j = i + 1, \\ q_i & \text{if } j = i - 1, \\ 1 - p_i - q_i & \text{if } j = i, \\ 0 & \text{otherwise.} \end{cases}$$

A birth and death chain is thus a Markov chain with a tridiagonal transition matrix; it is irreducible if and only if $p_k > 0$ for $k < N$ and that $q_k > 0$ for $k > 0$. We study the behavior of irreducible birth and death chains in Examples 11.A.6, 11.A.8, 11.A.9, and 11.A.16 below. §

In Section 11.1, we introduced the notion of a full support revision protocol in order to ensure that under the stochastic evolutionary process $\{X_t^N\}$, every population state in X is accessible from every other. Put differently, the full support condition ensures that the stochastic evolutionary process is irreducible.

11.A.2 Periodicity

In the discrete-time case, the behavior of the Markov chain $\{X_t\}_{t=0}^{\infty}$ within a recurrent class depends on the period structure of that class. The *period* of the recurrent state x is defined as $\text{gcd}(\{t \geq 1 : \mathbb{P}_x(X_t = x) > 0\})$: that is, it is the greatest common divisor of the set of times at which the chain can revisit x if it is run from x . It can be shown that states from the same recurrent class have the same period, and so it makes sense to speak of the period of the class itself.

Example 11.A.4. The Markov chain defined in Example 11.A.1 has three recurrent classes: $\{3\}$, $\{4, 5\}$, and $\{7, 8, 9\}$. Clearly, the period of state 3 is 1, and the period of states 4 and 5 is 2. Since $P_{99} > 0$, a process run from initial state 9 can remain there for any finite number of periods; thus state 9 has period 1, as do the other states in its class. §

This example illustrates a simple and useful fact: if any state x in a recurrent class has $P_{xx} > 0$, then all states in this class are of period 1.

If the Markov chain $\{X_t\}$ is irreducible, then all of its states are of the same period. If this common period is greater than 1, we say that $\{X_t\}$ is *periodic*; if instead the common period is 1, $\{X_t\}$ is *aperiodic*.

Theorem 11.A.5 describes how periodicity constrains the t step transition probabilities of an irreducible Markov chain. Part (i) of the theorem observes that the state space can be partitioned into sets that the chain must pass through sequentially. Part (ii) of the theorem implies that any t step transition not forbidden by part (i) will have positive probability if t is large enough.

Theorem 11.A.5. Let $\{X_t\}_{t=0}^{\infty}$ be an irreducible Markov chain with period d . Then there is a partition C_0, C_1, \dots, C_{d-1} of X such that (i) $[x \in C_i \text{ and } P_{xy}^t > 0]$ implies that $y \in C_{(i+t) \bmod d}$; and (ii) $x, y \in C_i$ implies that $P_{xy}^{n \cdot d} > 0$ for all n large enough.

Example 11.A.6. Birth and death chains: periodicity. Let $\{X_t\}_{t=0}^{\infty}$ be a birth and death chain on the state space $X = \{0, 1, \dots, N\}$ (Example 11.A.3). If the probability of staying still is zero at every state (i.e., if $p_0 = q_N = 1$ and $p_k + q_k = 1$ for $k \in \{1, \dots, N - 1\}$), then the chain has period 2, with C_0 containing the even states and C_1 the odd states. Otherwise, the chain is aperiodic. §

Let us emphasize once more that periodicity is a discrete-time phenomenon: since the jump times of continuous-time Markov processes follow exponential distributions, periodic behavior is impossible.

11.A.3 Hitting Times and Hitting Probabilities

To this point, our discussion of the behavior of Markov chains and processes has been “qualitative”, focusing only on the possibility or impossibility of transitions between given pairs of states over various lengths of time. We now turn to the “quantitative” analysis of these processes, which addresses the likelihoods of these transitions.

We start by considering hitting times. Let $\{X_t\}$ be an irreducible Markov chain or Markov process on the finite state space X , and let Z be a subset of X . The *hitting time* of Z , which we denote by T_Z , is the random time at which the process first takes a value in Z : that is, $T_Z = \inf\{t \geq 0 : X_t \in Z\}$.

Fix the set $Z \subseteq X$, and let $w_x = \mathbb{E}_x T_Z$ denote the expected time to hit Z if $\{X_t\}$ is run from state x . By definition, $w_z = 0$ whenever $z \in Z$. The expected hitting times for other initial conditions admit a pleasingly simple characterization.

Proposition 11.A.7. *Let $\{X_t\}$ be an irreducible Markov chain or process, let $Z \subseteq X$, and let $\{w_x\}_{x \notin Z}$ be the collection of expected times to hit Z starting from states outside Z .*

(i) *If $\{X_t\}_{t=0}^\infty$ is a Markov chain with transition matrix P , then $\{w_x\}_{x \notin Z}$ is the unique solution to the linear equations*

$$(11.38) \quad w_x = 1 + \sum_{y \notin Z} P_{xy} w_y \quad \text{for all } x \notin Z.$$

(ii) *If $\{X_t\}_{t \geq 0}$ is a Markov process with jump rate vector λ and transition matrix P , then $\{w_x\}_{x \notin Z}$ is the unique solution to the linear equations*

$$(11.39) \quad w_x = \lambda_x^{-1} + \sum_{y \notin Z} P_{xy} w_y \quad \text{for all } x \notin Z.$$

We first provide intuition for the discrete-time case (Proposition 11.A.7(i)). If we begin at state $x \notin Z$, then the expected time to reach the set Z is the sum of 1, representing the period now being spent at x , and the expected time to reach Z from the state that occurs next. But a transition into Z means that the wait is over, we need only account for transitions to states outside of Z , including the “null transition” that keeps us at x . This is precisely what condition (11.38) requires.

In the continuous-time case (Proposition 11.A.7(ii)), the time until the first transition from state x is exponentially distributed with rate λ_x . This generates a mean contribution of λ_x^{-1} to the expected time to reach Z , accounting for the difference between conditions (11.38) and (11.39). If the jump rate vector λ is constant, then the mean hitting times are those from the discrete-time case multiplied by the common value of λ_x^{-1} , representing the mean time between transitions. If we introduce the generator $Q = \text{diag}(\lambda)(P - I)$ of

the Markov process (see Section 10.C.3 or 11.A.7), we can express condition (11.39) in the concise form

$$-\sum_{y \notin Z} Q_{xy} w_y = 1 \quad \text{for all } x \notin Z.$$

Example 11.A.8. Birth and death chains: expected hitting times. Let $\{X_t\}_{t=0}^{\infty}$ be an irreducible birth and death chain on $\mathcal{X} = \{0, 1, \dots, N\}$ as defined in Example 11.A.3. To compute the expected hitting times to reach a given state $z \in \mathcal{X}$, substitute the definition (11.37) of the transition probabilities into condition (11.38); some rearranging yields

$$(11.40) \quad (p_k + q_k)w_k - p_k w_{k+1} - q_k w_{k-1} = 1.$$

To find the expected hitting time of z from an initial state x below z , we must solve the collection of equalities (11.40) with k ranging from 0 to $z-1$. This problem is a second-order linear difference equation with variable coefficients and boundary conditions $w_z = 0$ and $q_0 = 0$.

To solve it, we define for $k \in \{1, \dots, z\}$ the first differences $d_k = w_k - w_{k-1}$, so that

$$(11.41) \quad w_x = w_0 + \sum_{k=1}^x d_k.$$

Equation (11.40) tells us that

$$(11.42) \quad d_k = \frac{q_{k-1}}{p_{k-1}} d_{k-1} - \frac{1}{p_{k-1}},$$

so the terminal condition $q_0 = 0$ implies that $d_1 = -1/p_0$. Repeated substitution of the expressions for $d_{k-1}, d_{k-2}, \dots, d_1$ given by (11.42) into (11.42) itself yields

$$(11.43) \quad d_k = -\sum_{j=1}^k \frac{1}{p_{j-1}} \prod_{i=j}^{k-1} \frac{q_i}{p_i},$$

where the empty product equals 1. Equation (11.41), the terminal condition $w_z = 0$, and equation (11.43) together imply that

$$w_0 = -\sum_{k=1}^z d_k = \sum_{k=1}^z \sum_{j=1}^k \frac{1}{p_{j-1}} \prod_{i=j}^{k-1} \frac{q_i}{p_i}.$$

Combining this with equation (11.41) yields the general formula for the expected hitting times,

$$(11.44) \quad w_x = \sum_{k=x+1}^z \sum_{j=1}^k \frac{1}{p_{j-1}} \prod_{i=j}^{k-1} \frac{q_i}{p_i} \text{ for } x \in \{0, \dots, z-1\}.$$

A similar argument combining (11.40) and the boundary conditions $w_z = 0$ and $p_N = 0$ yields an expression for the expected hitting time of z from a state y above z :

$$(11.45) \quad w_y = \sum_{k=z+1}^y \sum_{j=k}^N \frac{1}{q_j} \prod_{i=k}^{j-1} \frac{p_i}{q_i} \text{ for } y \in \{z+1, \dots, N\}.$$

One can also derive an expression for the expected hitting time of the set $\{a, b\}$ from an initial state in between (see the next example), but this expression is not as simple as the two given above. §

Example 11.A.9. Birth and death chains: Hitting probabilities. Birth and death chains also allow one to express the probability that the chain reaches state a before state b using a simple formula. To state this formula concisely, $\{X_t\}$ be an irreducible birth and death chain on $\mathcal{X} = \{0, 1, \dots, N\}$, and define the increasing function $\phi : \mathcal{X} \rightarrow \mathbf{R}_+$ inductively via

$$\phi_0 = 0, \quad \phi_1 = 1, \quad \text{and} \quad \phi_{k+1} = \phi_k + \frac{\prod_{i=1}^k q_i}{\prod_{i=1}^k p_i} = \sum_{j=0}^k \prod_{i=1}^j \frac{q_i}{p_i},$$

where the empty product again equals 1. If $a < x < b$, one can show that

$$(11.46) \quad \mathbb{P}_x(T_a < T_b) = \frac{\phi_b - \phi_x}{\phi_b - \phi_a}.$$

For intuition, note that ϕ is defined so that

$$\mathbb{E}(\phi_{X_{t+1}} | \phi_{X_t} = \phi_k) = E(\phi_{X_{t+1}} | X_t = k) = p_k \phi_{k+1} + (1 - p_k - q_k) \phi_k + q_k \phi_{k-1} = \phi_k.$$

(This equality says that $\{\phi_{X_t}\}$ is a *martingale*.) If we let $T = T_{\{a,b\}}$ denote the hitting time of the set $\{a, b\}$, one can prove (using the *Optional Stopping Theorem*) that

$$\mathbb{E}_x(\phi_{X_T}) = \phi_x.$$

In words: expected value of the process $\{\phi_{X_t}\}$ at the (random) time it hits the set $\{a, b\}$ is

the same as the process's initial value. Evaluating the expectation above yields

$$\phi_a \mathbb{P}_x(T_a < T_b) + \phi_b \mathbb{P}_x(T_b < T_a) = \phi_x,$$

which can be rearranged to obtain (11.46). §

11.A.4 The Perron-Frobenius Theorem

Many key properties of irreducible finite-state Markov chains can be viewed as consequences of the *Perron-Frobenius Theorem*, a basic result from matrix analysis. To present this theorem in its usual language, we refer to a transition matrix $P \in \mathbf{R}^{X \times X}$ (i.e., a non-negative row matrix with row sums equal to 1) as a *stochastic matrix*, and call such a matrix *irreducible* or *aperiodic* according to whether the induced Markov chain has these properties. Nonnegative matrices that are both irreducible and aperiodic are sometimes referred to as *primitive*.

Theorem 11.A.10 (Perron-Frobenius).

Suppose that the matrix $P \in \mathbf{R}^{X \times X}$ is stochastic and irreducible. Then:

- (i) 1 is an eigenvalue of P of algebraic multiplicity 1, and no eigenvalue of P has modulus greater than 1.
- (ii) The vector $\mathbf{1}$ is a right eigenvector of P corresponding to eigenvalue 1. That is, $P\mathbf{1} = \mathbf{1}$.
- (iii) There is a probability vector μ with positive components that is a left eigenvector of P corresponding to eigenvalue 1; thus, $\mu'P = \mu'$.

Suppose in addition that P is aperiodic, and hence primitive. Then:

- (iv) All eigenvalues of P other than 1 have modulus less than 1.
- (v) The matrix powers P^t converge to the matrix $\mathbf{1}\mu'$ as t approaches infinity.
- (vi) Indeed, let λ_2 be the eigenvalue of P with the second-largest modulus, and let $r \in (|\lambda_2|, 1)$. Then for some $c > 0$, we have that

$$\max_{ij} |(P^t - \mathbf{1}\mu')_{ij}| \leq c r^t \text{ for all } t \geq 1.$$

If P is (real or complex) diagonalizable, this statement remains true when $r = |\lambda_2|$.

In Sections 11.A.5 and 11.A.8, we will interpret all of these statements in the context of Markov chains.

11.A.5 Stationary Distributions for Markov Chains

This section and the next consider stationary distributions for discrete-time Markov chains. We will see in Section 11.A.7 that the basic formulas for the continuous-time case are identical to those for the discrete-time case when the jump rates λ_x are the same for each state, as is the case for the stochastic evolutionary process studied in the text.

Let $\{X_t\}_{t=0}^{\infty}$ be a Markov chain with finite state space \mathcal{X} and transition matrix P . We call a probability distribution $\mu \in \mathbf{R}^{\mathcal{X}}$ a *stationary distribution* of $\{X_t\}_{t=0}^{\infty}$ if

$$(11.47) \quad \mu'P = \mu'.$$

More explicitly, μ is a stationary distribution if

$$(11.48) \quad \sum_{x \in \mathcal{X}} \mu_x P_{xy} = \mu_y \quad \text{for all } y \in \mathcal{X}.$$

To interpret these conditions, decompose the probability of the chain being at state y at time 1 as follows:

$$\sum_{x \in \mathcal{X}} \mathbb{P}(X_0 = x) \mathbb{P}(X_1 = y | X_0 = x) = \mathbb{P}(X_1 = y).$$

Comparing the previous two equations, we see that if X_0 is distributed according to the stationary distribution μ , then X_1 is also distributed according to μ , and so, by the Markov property, is every subsequent X_t .

In general, a finite-state Markov chain must admit at least one stationary distribution, and it may admit many. For instance, if the transition matrix P is the identity matrix, so that each state defines its own recurrent class, then every probability distribution on \mathcal{X} is a stationary distribution. But if a Markov chain is irreducible, definition (11.47) and Theorem 11.A.10(i) and (iii) imply that its stationary distribution is unique.

Theorem 11.A.11. *If the Markov chain $\{X_t\}$ is irreducible, it admits a unique stationary distribution.*

Example 11.A.12. Focusing on the three irreducible closed sets from Example 11.A.1, we can define three irreducible Markov chains with transition matrices

$$P = (1), \quad P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \text{and} \quad P = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ .9 & 0 & .1 \end{pmatrix}.$$

In the first case the stationary distribution is trivial, in the second case it is $\mu = (\frac{1}{2}, \frac{1}{2})$, and in the final case it is $\mu = (\frac{9}{28}, \frac{9}{28}, \frac{10}{28})$. §

If the Markov chain $\{X_t\}$ is not irreducible, then we can find a unique stationary distribution μ^R corresponding to each of the chain's recurrent classes R , and the set of stationary distributions of $\{X_t\}$ is the convex hull of these μ^R . Intuitively, by specifying the weight placed on μ^R in the convex combination, we choose the probability that the process begins in recurrent class R ; to maintain stationarity, the relative weights on states in R must be in the proportions given by μ^R .

One can interpret the stationary distribution weights of irreducible Markov chains in terms of expected return times. In Section 11.A.3, we defined the *hitting time* of state x under the Markov chain $\{X_t\}$ to be the random variable $T_x = \inf\{t \geq 1 : X_t = x\}$. When $\{X_t\}$ is run from initial condition x , T_x is called the *return time* of state x . (Note that if $X_1 = x$, then the return time is 1, even though the chain has not actually left state x .) One can establish the following connection between expected return times and the stationary distribution weights.

Proposition 11.A.13. *If the Markov chain $\{X_t\}$ is irreducible with stationary distribution μ , then $\mathbb{E}_x T_x = \mu_x^{-1}$ for all $x \in \mathcal{X}$.*

Thus, the higher is the weight on x in the stationary distribution, the less time we expect will pass before a chain starting at x returns to this state.

11.A.6 Reversible Markov Chains

In general, computing the stationary distribution of an irreducible Markov chain means finding an eigenvector of its transition matrix, a task that is computationally demanding when the state space is large. Fortunately, there are many interesting examples in which the stationary distribution satisfies a property that is both stronger and easier to check. We say that the Markov chain $\{X_t\}$ is *reversible* if it admits a *reversible distribution*: that is, a probability distribution μ that satisfies the *detailed balance conditions*

$$(11.49) \quad \mu_x P_{xy} = \mu_y P_{yx} \quad \text{for all } x, y \in \mathcal{X}.$$

If we sum this equation over $x \in \mathcal{X}$, we obtain equation (11.48), so a reversible distribution is also a stationary distribution.

To understand why reversibility is so named, imagine running the Markov chain $\{X_t\}$ from initial distribution μ , and observing two consecutive frames from a film of the chain. Suppose that these frames display the realizations x and y . According to equation (11.49),

the probability that the chain is first at state x and next at state y , $\mu_x P_{xy}$, is equal to the probability that the chain is first at state y and next at state x , $\mu_y P_{yx}$. Put differently, the probability of observing (x, y) is the same whether the film is running forward or backward. It is easy to verify that this property extends to any finite sequence of realizations.

Example 11.A.14. In Example 11.A.12, the first two Markov chains are reversible, but the third is not: if the states are labeled 7, 8, and 9 and our film shows state 7 and then state 8, then we can conclude that it is running forward: compare Figure 11.A.1. §

Example 11.A.15. Random walks on graphs. Let $\mathcal{G} = (\mathcal{X}, \mathcal{E})$ be a connected undirected graph with node set \mathcal{X} and edge set $\mathcal{E} \subseteq \{\{x, y\} : x, y \in \mathcal{X}, x \neq y\}$. Let $d_x = \#\{E \in \mathcal{E} : x \in E\}$ be the degree of node x (i.e., the number of edges containing x), and let $D = \sum_{x \in \mathcal{X}} d_x$. Let $\{X_t\}$ be a Markov chain on \mathcal{X} with transition probabilities

$$P_{xy} = \begin{cases} 1/d_x & \text{if } \{x, y\} \in \mathcal{E}, \\ 0 & \text{otherwise.} \end{cases}$$

Evidently, $\{X_t\}$ is irreducible and reversible with stationary distribution $\mu_x = d_x/D$. §

Example 11.A.16. Birth and death chains: Reversibility and the stationary distribution. Birth and death chains (Example 11.A.3) are always reversible. Since only adjacent transitions are possible, the detailed balance conditions (11.49) reduce to $\mu_k q_k = \mu_{k-1} p_{k-1}$ for $k \in \{1, \dots, N\}$. Applying this formula inductively and noting that stationary distribution weights sum to 1, we obtain an explicit formula for this distribution:

$$\mu_k = \mu_0 \prod_{i=1}^k \frac{p_{i-1}}{q_i} \text{ for } k \in \{1, \dots, N\}, \text{ and } \mu_0 = \left(\sum_{k=0}^N \prod_{i=1}^k \frac{p_{i-1}}{q_i} \right)^{-1},$$

where as always the empty product equals 1. §

11.A.7 Stationary Distributions and Reversibility for Markov Processes

The notions of stationarity and reversibility are equally important in the continuous-time setting, where they are most easily studied using the generator of the Markov process. In Section 10.C.3, we described the transition probabilities of the Markov process $\{X_t\}_{t \geq 0}$ using the matrix semigroup $\{P(t)\}_{t \geq 0}$, where $P_{xy}(t) = \mathbb{P}(X_t = y | X_0 = x)$. The generator of $\{X_t\}_{t \geq 0}$ is the matrix $Q = \text{diag}(\lambda)(P - I)$, where $\lambda \in \mathbf{R}_+^X$ and $P \in \mathbf{R}_+^{X \times X}$ are the vector of jump rates and the transition matrix of the Markov process, respectively. Equation (10.14)

showed that the generator and the transition probabilities are related by $P(t) = e^{Qt}$, which is the matrix solution to the *backward equation* $\dot{P}(t) = QP(t)$ from initial condition $P(0) = I$.

Proposition 11.A.17 introduces two equivalent definitions of the stationary distribution for a Markov process.

Proposition 11.A.17. *Let $\{X_t\}_{t \geq 0}$ be a Markov process on the finite state space X with semigroup $\{P(t)\}_{t \geq 0}$ and generator Q , and let μ be a probability distribution on X . Then the following conditions are equivalent:*

$$(11.50) \quad \mu'Q = \mathbf{0}';$$

$$(11.51) \quad \mu'P(t) = \mu' \text{ for all } t \geq 0.$$

When these conditions hold, we call μ a stationary distribution of $\{X_t\}_{t \geq 0}$.

Condition (11.51) can be written explicitly as

$$(11.52) \quad \mathbb{P}_\mu(X_t = x) = \mu_x \text{ for all } x \in X \text{ and } t \geq 0.$$

In other words, if $\{X_t\}$ is run with the stationary distribution μ as its initial distribution, then it continues to follow distribution μ at all times $t \geq 0$.

Proof of Proposition 11.A.17. Premultiplying the backward equation by the probability distribution μ , we obtain

$$(11.53) \quad \frac{d}{dt}(\mu'P(t)) = \mu'\dot{P}(t) = \mu'QP(t).$$

Thus, if $\mu'Q = \mathbf{0}'$, then $\frac{d}{dt}(\mu'P(t)) = \mathbf{0}'$, which with $\mu'P(0) = \mu'I = \mu'$ implies (11.51). To establish the other direction, note that (11.51) and (11.53) imply that $\mathbf{0}' = \frac{d}{dt}(\mu'P(t)) = \mu'QP(t)$; since $P(0) = I$, we conclude that $\mu'Q = \mathbf{0}'$. ■

To obtain a condition that more closely resembles the discrete-time stationarity condition (11.47), one can substitute the definition of Q into equation (11.50) to obtain

$$(11.54) \quad (\mu \bullet \lambda)'P = (\mu \bullet \lambda),$$

where \bullet denotes componentwise products: $(\mu \bullet \lambda)_i = \mu_i \lambda_i$. Notice that if the jump rate vector λ is constant, then equation (11.54) reduces to the discrete-time condition $\mu'P = \mu'$ from equation (11.47).

Equation (11.54) suggests that the basic properties of the stationary distribution in the continuous-time setting should parallel those in the discrete-time setting so long as the

jump rates are accounted for correctly. Suppose, for instance, that the Markov process $\{X_t\}$ is irreducible. If we fix $\alpha \in (0, (\max_x \lambda_x)^{-1}]$, it is easy to see that $\alpha Q + I$ is an irreducible stochastic matrix. Theorem 11.A.10(iii) then tells us that there is a unique probability vector μ satisfying $\mu'(\alpha Q + I) = \mu'$, or, equivalently, $\mu'Q = \mathbf{0}$. We therefore have

Theorem 11.A.18. *If the Markov process $\{X_t\}$ is irreducible, it admits a unique stationary distribution.*

To link stationary distribution weights and expected return times, define the hitting time of state x as $T_x = \inf \{t \geq \tau_1 : X_t = x\}$, where τ_1 is the first jump time of the process $\{X_t\}_{t \geq 0}$. (Notice again that the state need not actually change at time τ_1). If the Markov process $\{X_t\}_{t \geq 0}$ is irreducible, then it can be shown that $\mathbb{E}_x T_x = (\mu_x \lambda_x)^{-1}$.

Finally, we say that a Markov process is *reversible* if it admits a *reversible distribution* μ , which here means that

$$(11.55) \quad \mu_x \lambda_x P_{xy} = \mu_y \lambda_y P_{yx} \quad \text{for all } x, y \in \mathcal{X}.$$

Summing this equation over x yields the stationarity condition (11.54). Again, equation (11.55) reduces to the discrete-time condition $\mu_x P_{xy} = \mu_y P_{yx}$ from equation (11.49) if jump rates are constant.

11.A.8 Convergence in Distribution

The fundamental results of the theory of finite-state Markov chains and processes establish various senses in which the infinite-horizon behavior of irreducible processes is independent of their initial conditions, and characterizable in terms of the process's stationary distribution. As with i.i.d. random variables, there are distinct results on distributional properties and on sample path properties; we consider the former here, and the latter in Appendix 11.A.9. For a general discussion of the distinction between distributional and sample path properties, see Appendix 10.B.3.

The basic distributional result, Theorem 11.A.19, provides conditions under which the time t distributions of a Markov chain or process converge to its stationary distribution. In the continuous-time setting, irreducibility is sufficient for convergence in distribution; in the discrete-time setting, aperiodicity is also required.

Theorem 11.A.19 (Convergence in distribution). *Suppose that $\{X_t\}$ is either an irreducible aperiodic Markov chain or an irreducible Markov process, and that its stationary distribution is μ .*

Then for any initial distribution π , we have that

$$\lim_{t \rightarrow \infty} \mathbb{P}_\pi(X_t = x) = \mu_x \text{ for all } x \in \mathcal{X}.$$

In the discrete-time case, Theorem 11.A.19 is an immediate consequence of Theorem 11.A.10(v):

$$\lim_{t \rightarrow \infty} \mathbb{P}_\pi(X_t = x) = \lim_{t \rightarrow \infty} (\pi' P^t)_x = (\pi' \mathbf{1} \mu')_x = \mu_x.$$

Theorem 11.A.19 shows that the after a long enough time has passed, the time t distribution of $\{X_t\}$ will be close to its stationary distribution. But it does not tell us how long we must wait before this limit result becomes relevant. The answer to this question can be found in Theorem 11.A.10(vi), which shows that the rate of convergence to the stationary distribution is determined by the second-largest eigenvalue modulus of the transition matrix P . Rather than restate this point formally, we illustrate this conclusion and build intuition by way of two examples.

Example 11.A.20. Consider the two-state Markov chain $\{X_t\}$ with transition matrix

$$P = \begin{pmatrix} 1-a & a \\ b & 1-b \end{pmatrix},$$

To ensure that this chain is irreducible and aperiodic, we suppose that $a, b > 0$ and that $a + b < 2$. The eigenvalues of P are $\lambda_1 = 1$ and $\lambda_2 = 1 - a - b$, and the associated left eigenvectors are

$$\mu \equiv \begin{pmatrix} \frac{b}{a+b} \\ \frac{a}{a+b} \end{pmatrix} \text{ and } \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

respectively. Of course, since $\mu' P = \mu'$, μ is the stationary distribution of $\{X_t\}$.

Now let

$$L = \begin{pmatrix} \frac{b}{a+b} & \frac{a}{a+b} \\ -1 & 1 \end{pmatrix}$$

be the matrix whose rows are the left eigenvectors of P , and let $\Lambda = \text{diag}(\lambda)$. Then the matrix form of the left eigenvector equation for P is $LP = \Lambda L$. Since L is invertible,

$P = L^{-1}\Lambda L$ is diagonalizable. In fact, since $PL^{-1} = L^{-1}\Lambda$, the columns of

$$R \equiv L^{-1} = \begin{pmatrix} 1 & \frac{-a}{a+b} \\ 1 & \frac{b}{a+b} \end{pmatrix}$$

are the right eigenvectors of P .

The t step transitions of the Markov chain are described by the t th power of P , which we can evaluate as follows:

$$\begin{aligned} (11.56) \quad P^t &= (R\Lambda L)^t \\ &= R\Lambda^t L \\ &= \begin{pmatrix} 1 & \frac{-a}{a+b} \\ 1 & \frac{b}{a+b} \end{pmatrix} \begin{pmatrix} (\lambda_1)^t & 0 \\ 0 & (\lambda_2)^t \end{pmatrix} \begin{pmatrix} \frac{b}{a+b} & \frac{a}{a+b} \\ -1 & 1 \end{pmatrix} \\ &= (\lambda_1)^t \mathbf{1}\mu' + (\lambda_2)^t \begin{pmatrix} \frac{-a}{a+b} \\ \frac{b}{a+b} \end{pmatrix} \begin{pmatrix} -1 & 1 \end{pmatrix} \\ &= \mathbf{1}\mu' + (\lambda_2)^t \begin{pmatrix} \frac{a}{a+b} & \frac{-a}{a+b} \\ -\frac{b}{a+b} & \frac{b}{a+b} \end{pmatrix}. \end{aligned}$$

Thus, if the chain begins at initial distribution π , its time t distribution is

$$\pi' P^t = \mu' + (\lambda_2)^t \begin{pmatrix} \frac{\pi_1 a - \pi_2 b}{a+b} & \frac{-\pi_1 a + \pi_2 b}{a+b} \end{pmatrix}.$$

We conclude that the distribution of the Markov chain converges to the stationary distribution μ at geometric rate $|\lambda_2| = |1 - a - b|$. §

Some simple observations about matrix multiplication will be useful in the next example. For any matrices $A, B \in \mathbf{R}^{n \times n}$, we can write

$$(11.57) \quad AB = AIB = \sum_{k=1}^n A(e_k e_k') B = \sum_{k=1}^n (Ae_k)(e_k' B).$$

In words, to compute the product AB , we take the “outer product” of the k th column of A and the k th row of B , and then sum over k . (For a different derivation of (11.57), recall that the ij th term of AB is $\sum_k A_{ik} B_{kj}$; (11.57) expresses this fact in matrix form.) If we interpose a diagonal matrix $D = \text{diag}(d)$ between A and B , we can obtain a similar expression, used

implicitly in (11.56):

$$(11.58) \quad ADB = \sum_{k=1}^n d_k(Ae_k)(e'_k B).$$

Example 11.A.21. Let P be the transition matrix of an irreducible aperiodic Markov chain, and suppose that P is (real or complex) diagonalizable. Then, as in the previous example, we can write $P = R\Lambda L = R \operatorname{diag}(\lambda)L$, where λ is the vector of eigenvalues of P , the rows of L are left eigenvectors, and the columns of $R = L^{-1}$ are right eigenvectors. As before, we can write the t step transition probabilities as $P^t = R\Lambda^t L$, which we can rewrite using equation (11.58) as

$$(11.59) \quad P^t = R\Lambda^t L = \sum_{k=1}^n (\lambda_k)^t (Re_k)(e'_k L).$$

The Perron-Frobenius Theorem tells us that 1 is an eigenvalue of P with algebraic multiplicity 1, and that all other eigenvalues of P have modulus less than 1. Moreover, the left and right eigenvectors corresponding to eigenvalue 1 are the stationary distribution μ and the constant vector $\mathbf{1}$. Thus, if we order the eigenvectors so that $1 = \lambda_1 > |\lambda_2| \geq |\lambda_3| \geq \dots \geq |\lambda_n| > -1$, then $e'_1 L = \mu'$ and $Re_1 = \mathbf{1}$. (To see that Re_1 is $\mathbf{1}$ rather than some multiple of $\mathbf{1}$, note that $LR = I$, implying that $\mu' Re_1 = e'_1 L Re_1 = I_{11} = 1$.)

Using these observations, we can rewrite equation (11.59) as

$$P^t = \mathbf{1}\mu' + \sum_{k=2}^n (\lambda_k)^t (Re_k)(e'_k L).$$

Evidently, the rate of convergence of $\{X_t\}$ to its stationary distribution is determined by the second largest eigenvalue modulus of P . §

11.A.9 Ergodicity

Having considered distributional properties, we turn to sample path properties, which tell us about the behavior of almost every sequence of realizations of the Markov chain or process. (What distinguishes these two kinds of properties is discussed in Section 10.B.3.) Theorem 11.A.22 shows that if $\{X_t\}$ is an irreducible Markov chain or process, then for almost all realizations of $\omega \in \Omega$, the proportion of time that the sample path $\{X_t(\omega)\}$ spends in each state x converges; the limit is the weight μ_x that state x holds in the stationary distribution. If $\{X_t\}_{t=0}^{\infty}$ is an i.i.d. sequence, with each X_t having distribution

μ , this conclusion can be obtained by applying the Strong Law of Large Numbers to the sequence of indicator random variables $\{1_{\{X_t=x\}}\}_{t=0}^{\infty}$. Theorem 11.A.22 reveals that Markov dependence is enough for this implication to hold.

Theorem 11.A.22 (Ergodicity).

- (i) Suppose that $\{X_t\}_{t=0}^{\infty}$ is an irreducible Markov chain with stationary distribution μ . Then for any initial distribution π , we have that

$$\mathbb{P}_{\pi} \left(\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} 1_{\{X_t=x\}} = \mu_x \right) = 1 \text{ for all } x \in \mathcal{X}.$$

- (ii) Suppose that $\{X_t\}_{t \geq 0}$ is an irreducible Markov process with stationary distribution μ . Then for any initial distribution π , we have that

$$\mathbb{P}_{\pi} \left(\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T 1_{\{X_t=x\}} dt = \mu_x \right) = 1 \text{ for all } x \in \mathcal{X}.$$

While irreducible, aperiodic Markov chains converge in distribution and are ergodic, we should emphasize that the latter two properties are distinct in general. For instance, a Markov chain that is irreducible but not aperiodic will still be ergodic, but will not converge in distribution in the sense described in Theorem 11.A.19. Conversely, if we consider a Markov chain whose transition matrix is the identity matrix, then this chain trivially converges in distribution to its initial distribution, but the chain is not ergodic in the sense of Theorem 11.A.22.

11.N Notes

Section 11.1. The best response with mutations model is introduced in Kandori et al. (1993). Models of imitation with mutations are considered by Binmore and Samuelson (1997) and Fudenberg and Imhof (2006, 2008). See the Notes to Chapter 12 for further references.

While the analysis in this chapter and the next focuses on irreducible processes, we should mention a distinct branch of the literature on stochastic evolution that looks at reducible processes, and focuses on whether these processes converge to sets that are closed under better or best replies: see Basu and Weibull (1991) and Voorneveld (2004, 2005) for the relevant set-valued solution concepts, and for models of dynamics see Hurkens (1995), Monderer and Shapley (1996b), Friedman and Mezzetti (2001), Kukushkin (2004),

Josephson and Matros (2004), Dindoš and Mezzetti (2006), Josephson (2008) and Marden et al. (2009b).

Section 11.2. Stationary distributions for models of stochastic evolution in two-strategy games are studied by Binmore et al. (1995b), Binmore and Samuelson (1997), Young (1998b), Blume (2003), Benaïm and Weibull (2003), and Sandholm (2007c, 2009e,f). Toss and Switch (Examples 11.2.1 and 11.3.1) is a slight modification of the Ehrenfest urn model from statistical physics; see Brémaud (1999) for a discussion of the latter. The “swimming upstream” analogy following Example 11.2.6 appears in related contexts in Fudenberg and Harris (1992), Kandori (1997), and Binmore and Samuelson (1997).

Section 11.3. Ellison (1993, 2000) and Beggs (2005) (also see Kandori et al. (1993) and Binmore et al. (1995b)) study the rate of growth of waiting times as the noise level becomes small, while Benaïm and Weibull (2003, Lemma 3) (also see Benaïm and Sandholm (2007)) considers their rate of growth as the population size becomes large.

For local interaction models, see Ellison (1993, 2000), Blume (1993, 1995), Morris (2000), Nowak and May (1992, 1993), Nowak (2006), Szabó and Fáth (2007), and the references therein. For evolution in games on networks, see Blume (1997), as well as the general references of Goyal (2007), Vega-Redondo (2007), and Jackson (2008).

Section 11.4. The distinction between simple and clever payoff evaluation is noted in Sandholm (1998). Exercise 11.4.4(ii) is due to Rhode and Stegeman (1996). Committed agents for imitative dynamics are introduced in Sandholm (2009b,f)

Section 11.5. This section follows Sandholm (2009b). That paper builds on earlier work by Blume (1993, 1997, 2003), who demonstrated the usefulness of condition (11.31), whose origins are in statistical mechanics, in evolutionary game theory. Normal form potential games are introduced in Monderer and Shapley (1996b); see also Hofbauer and Sigmund (1988). The definition of finite-population full potential games and Theorem 11.5.7(i) are from Sandholm (2001a). Exercise 11.5.17(i) is due to Blume (1997).

Appendix 11.A. Most of the material presented here (excepting equations (11.44) and (11.45)) can be found in standard references on Markov chains and processes; see Norris (1997), Brémaud (1999), and Stroock (2005). Durrett (2005, Chapter 5) also offers a clear presentation of Markov chains. For more on the Perron-Frobenius Theorem, see Seneta (1981) or Horn and Johnson (1985).

Limiting Stationary Distributions and Stochastic Stability

12.0 Introduction

Chapter 11 began our study of the infinite horizon behavior of stochastic evolutionary processes. We saw that any process generated by a full support revision protocols admits a unique stationary distribution, which describes the behavior of the process over very long time spans regardless of the initial population state. The examples showed that even when the underlying game has multiple strict equilibria, the stationary distribution is often concentrated in the vicinity of just one of them if the noise level η is small or the population size N is large. In these cases, the population state so selected provides a unique prediction of infinite horizon play.

The analyses in the previous chapter considered fixed evolutionary processes $\{X_t^{N,\eta}\}$, defined by particular choices of η and N . The stationary distribution $\mu^{N,\eta}$ of such a process necessarily places positive mass on every state in \mathcal{X}^N . Thus, if the values of η and N are held fixed, it is not possible to obtain clean selection results. Furthermore, actually computing the stationary distribution for set choices of η and N is generally a difficult task. Unless the evolutionary process is reversible, one cannot usually expect to do more than calculate this distribution numerically, making it difficult to prove analytical results about the underlying process.

In this chapter, we circumvent both of these difficulties by allowing the parameters η and N to approach their limiting values. While each fixed stationary distribution $\mu^{N,\eta}$ has full support on \mathcal{X}^N , the limit of a sequence of stationary distributions may converge to a point mass at a single state; thus, taking limits in η and N allows us to obtain exact equilibrium selection results. And while computing a particular stationary distribution

requires solving a large collection of linear equalities, the limiting stationary distribution can often be found without explicitly computing any of the stationary distributions along the sequence. Moving to a limiting analysis not only allows us to reach novel conclusions about the reversible environments from the previous chapter, but also allows us to extend our analyses beyond those convenient environments.

While taking η and N to their limiting values is useful for obtaining equilibrium selection results, doing so also exacerbates the waiting time problem studied in Section 11.3. Since in practice these parameters are not at their limits, the rationale for the limiting analyses performed in this chapter is simply analytical convenience. Before applying the selection results described here, one should assess whether the time span of interest in the application are commensurate with the waiting times needed to generate equilibrium selection results at the relevant parameter values. If not, the use of the infinite horizon analysis is not warranted.

Population states that retain mass in a limiting stationary distribution are said to be *stochastically stable*. There are a number of different definitions of stochastic stability, depending on which limits are taken—just η , just N , η followed by N , or N followed by η —and on what should count as “retaining mass”. In Section 12.1, we introduce a variety of definitions of stochastic stability, and use examples to illustrate the differences among them. We also explain how the modeler’s choice of which limits to take, and in which order, is governed by the application at hand. Taking only the small noise limit, or taking this limit first, emphasizes the rarity of suboptimal play as the key force behind equilibrium selection. Taking only the large population limit, or taking it first, emphasizes the effects of large numbers of conditionally independent decisions in driving equilibrium selection.

Since it is not always easy to know which of these forces should be viewed as the primary one, a main goal in this chapter is to identify settings in which the small noise and large population limits agree. When this is the case, the modeler is spared the difficulty of adjudicating between discrepant predictions about very long run behavior.

In Sections 12.2–12.4, we revisit the reversible environments first studied in Chapter 11. Using the explicit formulas for stationary distributions developed there, we are not only able to find the limiting stationary distribution, but also are able to characterize the asymptotics of the stationary distribution as the parameters approach their limits.

In Section 11.5, we saw that when agents play a potential game using an exponential revision protocol, the stationary distribution $\mu^{N,\eta}$ must take one of two forms, with the form depending only on whether the protocol is direct or imitative. The asymptotics of the stationary distribution in these environments is the subject of Section 12.2.

Section 12.3 considers evolution in two-strategy games under noisy best response protocols. The examples in the last chapter showed that even in this simple strategic setting, seemingly similar revision protocols can lead to different equilibrium selections. Here we derive simple formulas describing the asymptotics of the stationary distribution, and use them to identify conditions under which all noisy best response protocols select the same stochastically stable state. We also describe general conditions under which the small noise and large population asymptotics agree, so that the choice of which limit to emphasize is inconsequential.

Section 12.4 studies evolution in two-strategy games under imitative protocols, both with and without mutations. In this context, the asymptotics of the stationary distribution can depend on the order in which limits in the population size N and mutation rate ϵ are taken. However, we show that introducing just two committed agents, one for each strategy, ensures agreement between the double limits. We also illustrate how equilibrium selection can be surprisingly sensitive to the specification of the revision protocol: we offer two imitative protocols whose common mean dynamic is the replicator dynamic, but that generate different stochastically stable states in a range of simple games.

The remainder of the chapter considers limiting stationary distributions in environments that are not necessarily reversible, so that tractable formulas for the stationary distribution are not available. Section 12.5 and Appendix 12.A focus on small noise limits. The methods explored here identify the stochastically stable states by constructing certain graphs on the state space \mathcal{X}^N . Each edge is assigned a *cost*, which captures the rate of decay of the edge's traversal probability as the noise level approaches zero. Questions about stochastic stability can be answered by comparing the overall costs of particular well-chosen graphs, obviating the need to compute stationary distributions.

In Section 12.6 and Appendix 12.B we consider large population limits, providing infinite-horizon complements to the finite-horizon approximation results from Chapter 10. We show first that as the population size grows large, all mass in the stationary distribution becomes concentrated around the so-called Birkhoff center of the mean dynamic (M): only those states satisfying a natural notion of recurrence under (M) are candidates for stochastic stability in the large population limit. We then show that under some global restrictions on the mean dynamic—that (M) admits a global Lyapunov function, or that it defines a cooperative differential equation—combined with nondegeneracy conditions on the stochastic process itself, we can refine the previous result: the mass in the stationary distribution can only become concentrated around stable rest points of the mean dynamic.

Compared to those in the rest of the book, the developments presented in Sections 12.5 and 12.6 are rather incomplete. While the graph-theoretic tools for studying small noise

stochastic stability are quite general, to date they have served as the basis for relatively few general selection results, at least in the model studied in this book (see the Notes to Section 12.5 for applications in other game-theoretic frameworks). In a similar vein, our results in Section 12.6 do not tell us which state or states are selected in the large population limit, despite the fact that in typical cases, we expect the mass in the stationary distribution to become concentrated on a single asymptotically stable set. In short, there are many interesting open questions about stochastic stability analysis for games, and much work remains to be done to answer them.

12.1 Definitions of Stochastic Stability

We begin by introducing a variety of definitions of stochastic stability. As in the last two chapters, we keep the notation manageable by focusing on the single population case.

Let us first review the basic stochastic evolutionary model from Sections 10.1 and 11.1. A population of size N recurrently plays a population game F^N , with agents receiving revision opportunities via independent Poisson alarm clocks. An agent's choice upon receiving a revision opportunity is governed by a full support revision protocol ρ^η , where $\eta > 0$ denotes the noise level. The full support assumption ensures that the resulting stochastic evolutionary process $\{X_t^{N,\eta}\}$ is irreducible, and so admits a unique stationary distribution $\mu^{N,\eta}$ on the state space \mathcal{X}^N . As explained in Section 11.1.2, the stationary distribution captures the infinite horizon behavior of the process $\{X_t^{N,\eta}\}$ in two distinct senses: it is both the process's limiting distribution and its limiting empirical distribution.

When the parameters η and N are fixed, the stationary distribution $\mu^{N,\eta}$ places positive mass on every state in $\mathcal{X}^N \subset X$. But the examples in Chapter 11 show that as either or both of these parameters approach their extremes, the stationary distribution may concentrate its mass on a small set of states, even in cases where the mean dynamic of the process admits multiple stable equilibria. The population states around which the mass in the stationary distribution accumulates are said to be *stochastically stable*. We formalize this concept now.

12.1.1 Small Noise Limits

We begin with the mathematically simplest setting for stochastic stability analysis: we fix the population size N , and take the noise level η to zero. Taking the small noise limit emphasizes the rarity of suboptimal choices over the magnitude of the population as the force behind equilibrium selection.

Because N is fixed, each stationary distribution in the collection $\{\mu^{N,\eta}\}_{\eta \in [0,\bar{\eta}]}$ is a probability measure on the same finite state space, \mathcal{X}^N , making stochastic stability easy to define in the present setting. We call state $x \in \mathcal{X}^N$ *stochastically stable in the small noise limit* if

$$(12.1) \quad \lim_{\eta \rightarrow 0} \mu_x^{N,\eta} > 0.$$

When the sequence $\{\mu^{N,\eta}\}_{\eta \in (0,\bar{\eta}]}$ converges to a point mass at state x , we call state x *uniquely stochastically stable*. (Exercise 12.1.1 notes one subtlety about the latter definition.)

Section 12.5 and Appendix 12.A provide general conditions under which the sequence $\{\mu^{N,\eta}\}_{\eta \in (0,\bar{\eta}]}$ converges to a limiting distribution $\mu^{N,*}$ as η approaches zero. When this limiting distribution exists, its support is the set of stochastically stable states.

Exercise 12.1.1. Construct a collection of stationary distributions $\{\mu^{N,\eta}\}_{\eta \in [0,\bar{\eta}]}$ under which the only stochastically stable state is not uniquely stochastically stable in the sense of the definition above. (Hint: The limiting distribution $\mu^{N,*}$ need not exist.) ‡

It sometimes is more convenient to work with a more inclusive notion of stochastic stability, one that is stated in terms of the exponential rates of decay of stationary distribution weights (cf Appendix 12.A.5). To introduce this idea, let us consider the equation

$$(12.2) \quad \lim_{\eta \rightarrow 0} \eta \log \mu_x^{N,\eta} = r_x^N,$$

where $r_x^N \leq 0$. This equation is equivalent to the requirement that

$$(12.3) \quad \mu_x^{N,\eta} = \exp\left(\eta^{-1}\left(r_x^N + o(1)\right)\right),$$

where $o(1)$ represents a term that approaches 0 as η approaches zero. Both (12.2) and (12.3) say that r_x^N is the exponential rate of decay of the stationary distribution weight on state x as η^{-1} approaches infinity. With this in mind, we say that state $x \in \mathcal{X}^N$ *weakly stochastically stable in the small noise limit* if

$$(12.4) \quad \lim_{\eta \rightarrow 0} \eta \log \mu_x^{N,\eta} = 0.$$

Put differently, state x is weakly stochastically stable if as η^{-1} approaches infinity, $\mu_x^{N,\eta}$ does not vanish at an exponential rate. This weight may remain positive, but it also may converge to zero at a subexponential pace.

The following proposition links the above definitions of stochastic stability.

Proposition 12.1.2. (i) *Every stochastically stable state is weakly stochastically stable.*

- (ii) Suppose that the limit in (12.2) exists for each state in \mathcal{X}^N , and that only state x is weakly stochastically stable. Then state x is uniquely stochastically stable.

Exercise 12.1.3. (i) Verify both parts of Proposition 12.1.2.

- (ii) Show by example that without the hypothesis about the existence of the limits in (12.2), Proposition 12.1.2(ii) is no longer true. ‡

12.1.2 Large Population Limits

By defining stochastic stability in terms of the small noise limit, one takes the vanishing probability of suboptimal behavior to be the source of equilibrium selection. An alternative approach, one more in keeping with our analyses earlier in the book, fixes the noise level and examines the behavior of the stationary distributions as the population size grows large. This approach emphasizes the role of the population's magnitude in determining equilibrium selection.

Defining stochastic stability for large population limits is complicated by the fact that as the population size N grows, the state spaces \mathcal{X}^N vary, becoming increasingly fine grids in the simplex X . We call state $x \in X$ *stochastically stable in the large population limit* if for every (relatively) open set $O \subseteq X$ containing x , we have

$$(12.5) \quad \lim_{N \rightarrow \infty} \mu^{N,\eta}(O) > 0.$$

If this limit is always equal to 1, we call x *uniquely stochastically stable*. The warning about the meaning “uniqueness” from Exercise 12.1.1 applies equally well here.

Exercise 12.1.4. The natural notion of convergence for the sequence of stationary distributions $\{\mu^{N,\eta}\}_{N=N_0}^\infty$ is *weak convergence*, which requires that there be a probability measure $\mu^{*,\eta}$ such that for every open set $O \subseteq X$, one has

$$\liminf_{N \rightarrow \infty} \mu^{N,\eta}(O) \geq \mu^{*,\eta}(O).$$

Show that state x is uniquely stochastically stable if and only if $\mu^{*,\eta}$ exists and is equal to a point mass at state x . ‡

As in the context of small noise limits, it is useful here to introduce a less demanding notion of stochastic stability based on exponential rates of decay. In Sections 12.2 and 12.3, we will establish limits of the form

$$(12.6) \quad \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - r^\eta(x) \right| = 0.$$

where r^η is a continuous function from X to \mathbf{R}_- . Condition (12.6) is equivalent to the requirement that

$$(12.7) \quad \mu_x^{N,\eta} = \exp\left(\eta^{-1}N(r^\eta(x) + o(1))\right) \quad \text{uniformly in } x \in \mathcal{X}^N,$$

where $o(1)$ represents an term that converges to zero uniformly in x as N approaches infinity. Both (12.6) and (12.7) say that as N approaches infinity, the exponential rate of decay of the stationary distribution weights on states close to x is approximately $\eta^{-1}r^\eta(x)$. If $r^\eta(x) = 0$, we call state x *weakly stochastically stable in the large population limit*.

As in the case of small noise limits, we have the following links between our notions of stochastic stability for the large population limit.

Proposition 12.1.5. *Suppose that condition (12.6) holds for some continuous function r^η . Then*

- (i) *Every stochastically stable state is weakly stochastically stable;*
- (ii) *If only one state is weakly stochastically stable, then this state is uniquely stochastically stable.*

While the corresponding results for the small noise setting were immediate, this proposition requires a proof.

Proof of Proposition 12.1.5. To prove part (i), we show that $m^\eta = \max_{x \in X} r^\eta(x)$ must equal 0. To establish this, suppose first that m^η is negative. Since Exercise 10.1.1 tells us that $\#\mathcal{X}^N = \binom{N+n-1}{n-1} < (N+n)^n$, equation (12.7) implies that for large enough N , we have

$$\mu^{N,\eta}(\mathcal{X}^N) \leq \#\mathcal{X}^N \cdot \exp\left(\eta^{-1}N \cdot \frac{m^\eta}{2}\right) < (N+n)^n \exp\left(\eta^{-1}N \cdot \frac{m^\eta}{2}\right).$$

Since m^η is negative, the last expression vanishes as N grows large, contradicting that $\mu^{N,\eta}$ is a probability measure on \mathcal{X}^N .

Now suppose that m^η is positive. Then since r^η is continuous, there is a neighborhood of the maximizer of r^η on which r^η exceeds $\frac{2m^\eta}{3}$. Once N is large enough, this neighborhood must contain at least one state y^N from \mathcal{X}^N , and equation (12.7) implies that for large enough N , $\mu_{y^N}^{N,\eta} \geq \exp\left(\eta^{-1}N \cdot \frac{m^\eta}{3}\right)$. This expression grows without bound as N grows large, again contradicting that $\mu^{N,\eta}$ is a probability measure, and allowing us to conclude that $m^\eta = \max_{x \in X} r^\eta(x) = 0$.

We now prove part (i) by establishing the contrapositive. Suppose that state x is not weakly stochastically stable: $r^\eta(x) < 0$. Since r^η is continuous, there is an open set $O \subset X$ containing x and a constant $\delta > 0$ such that $r^\eta(y) < -\delta$ for all $y \in O$. Equation (12.7) then

implies that for large enough N , we have

$$\mu_y^{N,\eta} \leq \exp\left(\eta^{-1}N\left(r^\eta(y) + \frac{\delta}{2}\right)\right) < \exp\left(-\eta^{-1}N \cdot \frac{\delta}{2}\right) \text{ whenever } y \in \mathcal{X}^N \cap O.$$

Thus

$$\mu^{N,\eta}(O) < \#\mathcal{X}^N \cdot \exp\left(-\eta^{-1}N \cdot \frac{\delta}{2}\right) < (N+n)^n \exp\left(-\eta^{-1}N \cdot \frac{\delta}{2}\right).$$

This implies that $\lim_{N \rightarrow \infty} \mu^{N,\eta}(O) = 0$. Since state x is in O , it cannot be stochastically stable.

We now prove part (ii). Suppose that state x is the only weakly stochastically stable state, and let $O \subset X$ be an open set containing x . Since $X - O$ is closed, there is a $\delta > 0$ such that $r^\eta(y) < -\delta$ for all $y \in X - O$. Repeating the argument from the proof of part (i) shows that $\mu^{N,\eta}(X - O)$ approaches zero as N grows large, and hence that $\mu^{N,\eta}(O)$ approaches one as N grows large. Therefore, state x is uniquely stochastically stable. This completes the proof of the proposition. ■

12.1.3 Double Limits

The two limits introduced above emphasize two different forces that could drive infinite horizon equilibrium selection. In applications, it may not be self-evident which of these forces is the central one. As this point is moot when both limits generate the same predictions, it is natural to look for settings in which the small noise and large population limits agree.

To answer this question, we first must account for the fact that the small noise and large population limits are not directly comparable: the former limit concerns a sequence of probability distributions on a fixed finite grid in the simplex, while the latter considers a sequence of distributions on finer and finer grids. Moreover, for there to be any chance of agreement, we must take each limit while fixing the other parameter at a value “close to its limit”: there is no reason to expect predictions generated by the small noise limit at small population sizes to agree with predictions generated by the large population limit at high noise levels.

The most convenient way of handling both of these issues is to consider double limits. For instance, we call state $x \in X$ *stochastically stable in the small noise double limit* if for every open set $O \subseteq X$ containing x , we have

$$(12.8) \quad \lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \mu^{N,\eta}(O) > 0.$$

If this limit is always equal to 1, we call state x *uniquely stochastically stable*. Since in (12.8) the population size N is held fixed while the noise level η is taken to zero, this double limit describes the behavior of the small noise limit at large population sizes.

In the same way, we call state $x \in X$ *stochastically stable in the large population double limit* if for every open set $O \subseteq X$ containing x , we have

$$(12.9) \quad \lim_{\eta \rightarrow 0} \lim_{N \rightarrow \infty} \mu^{N,\eta}(O) > 0.$$

If this limit is always equal to 1, we call x *uniquely stochastically stable*. Evidently, the double limit (12.9) describes the behavior of the large population limit at small noise levels.

As with single limits, one can consider weaker notions of stochastic stability defined in terms of exponential rates of decay. For instance, suppose that for some continuous function $r: X \rightarrow \mathbf{R}_-$ we have

$$(12.10) \quad \lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - r(x) \right| = 0.$$

Then when the population size N is sufficiently large, the exponential rate of decay of $\mu_x^{N,\eta}$ as η^{-1} approaches infinity is approximately $Nr(x)$. Similarly, if there is a continuous function $\hat{r}: X \rightarrow \mathbf{R}_-$ such that

$$(12.11) \quad \lim_{\eta \rightarrow 0} \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - \hat{r}(x) \right| = 0,$$

then for small η , as N approaches infinity, the exponential rate of decay of the stationary distributions weights on states near x is approximately $\eta^{-1}\hat{r}(x)$.

Exercise 12.1.6. Write down precise statements justifying these interpretations of conditions (12.10) and (12.11). ‡

For either double limit, we say that state x is *weakly stochastically stable* if its limiting rate of decay, $r(x)$ or $\hat{r}(x)$, is equal to zero. As in Proposition 12.1.5, one can show that if the function r or \hat{r} is continuous, then every stochastically stable state is weakly stochastically stable, and that if only one state is weakly stochastically stable, it is uniquely stochastically stable.

Exercise 12.1.7. Prove these extensions of Proposition 12.1.5 to the cases of small noise and large population double limits. ‡

Exercise 12.1.8. Consider the small noise double limit $\mu^{\dagger,*} = \lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \mu^{N,\eta}$, interpreting each limit in the sense of weak convergence of probability measures (see Exercise 12.1.4).

- (i) Show that if $\mu^{+,*}$ is a point mass at state x , then x is uniquely stochastically stable.
- (ii) Construct a collection of probability measures $\{\mu^{N,\eta}\}$ for which state x is uniquely stochastically stable in the small noise double limit, but for which $\mu^{+,*}$ does not exist. ‡

In cases where the rate of decay functions r and \hat{r} are identical, the double limits agree in a very strong sense: not only do the small noise and large population limits agree about which state will predominate in the long run, but they also agree about the degrees of unlikelihood of all of the other states.

The next two sections focus on the two reversible settings studied in the previous chapter: potential games under exponential revision protocols, and two-strategy games under arbitrary protocols. Taking the explicit formulas for the fixed-parameter stationary distributions as our starting point, we are able to derive simple expressions for the rates at which the stationary distribution weights decay. We show that when agents employ noisy best response protocols, or when they employ imitative protocols in the presence of committed agents, the rates of decay functions do not depend on the order in which the small noise and large population limits are taken. Thus, in these settings, which limit is emphasized has little bearing on infinite horizon predictions.

It is an open question whether these agreements between the double limits persist beyond these reversible settings. Our analyses of small noise limits (Section 12.5) and large population limits (Section 12.6) take some necessary first steps toward answering this question.

12.1.4 Double Limits: A Counterexample

Before studying environments in which the two double limits agree, we provide an example which illustrates that agreement cannot be taken for granted.

Example 12.1.9. Consider a population of agents who are matched to play the symmetric normal form game with strategy set $S = \{0, 1\}$ and payoff matrix

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 1 \end{pmatrix}.$$

The unique Nash equilibrium of the population game $F(x) = Ax$ is the mixed equilibrium $x^* = (x_0^*, x_1^*) = (\frac{1}{3}, \frac{2}{3})$. To simplify notation in what follows we allow self-matching, but the analysis is virtually identical without it.

Suppose that agents employ the following revision protocol, which combines imitation of successful opponents and mutations:

$$\rho_{ij}^\varepsilon(\pi, x) = x_j \pi_j + \varepsilon.$$

Here we will simplify again by allowing agents to imitate themselves; this too has a negligible effect on the results. The protocol ρ^ε generates the mean dynamic

$$(12.12) \quad \dot{x}_i = V_i^\varepsilon(x) = x_i \hat{F}_i(x) + 2\varepsilon \left(\frac{1}{2} - x_i\right),$$

which is the sum of the replicator dynamic and an order ε term that points toward the center of the simplex. When $\varepsilon = 0$, this dynamic is simply the replicator dynamic: the Nash equilibrium $x^* = (\frac{1}{3}, \frac{2}{3})$ attracts solutions from all interior initial conditions, while pure states e_0 and e_1 are unstable rest points. When $\varepsilon > 0$, the two boundary rest points disappear, leaving a globally stable rest point that is near x^* , but slightly closer to the center of the simplex.


Using the formulas from Theorem 11.2.4, we can compute the stationary distribution $\mu^{N,\varepsilon}$ of the process $\{X_t^{N,\varepsilon}\}$ generated by F and ρ^ε for any fixed values of N and ε . Four instances are presented in Figure 12.1.1. In the discussion to follow, we streamline our notation in the usual way, letting $\chi = x_1$ denote the fraction of agents choosing strategy 1.

Figure 12.1.1(i) presents the stationary distribution when $\varepsilon = .1$ and $N = 100$. This distribution is drawn above the phase diagram of the mean dynamic (12.12), whose global attractor is appears at $\hat{\chi} \approx .6296$. The stationary distribution $\mu^{N,\varepsilon}$ has its mode at state $\chi = .64$, but is dispersed rather broadly about this state.


Figure 12.1.1(ii) presents the stationary distribution and mean dynamic when $\varepsilon = .1$ and $N = 10,000$. Increasing population size moves the mode of the distribution occurs to state $\chi = .6300$, and, more importantly, causes the distribution to exhibit much less dispersion around the modal state. This numerical analysis suggests that in the large population limit, the stationary distribution $\mu^{N,\varepsilon}$ will approach a point mass at $\hat{\chi} \approx .6296$, the global attractor of the relevant mean dynamic. We will confirm this conjecture in Section 12.4.1.

As the noise level ε approaches zero, the rest point of the mean dynamic approaches the Nash equilibrium $\chi^* = \frac{2}{3}$. Therefore, if after taking N to infinity we take ε to zero, we obtain the double limit


$$(12.13) \quad \lim_{\varepsilon \rightarrow 0} \lim_{N \rightarrow \infty} \mu^{N,\varepsilon} = \delta_{\chi^*},$$




(i) $N = 100, \epsilon = .1$



(ii) $N = 10,000, \epsilon = .1$



(iii) $N = 100, \epsilon = 10^{-5}$



(iv) $N = 100, \epsilon = 10^{-7}$

Figure 12.1.1: Stationary distribution weights $\mu_{\chi}^{N,\epsilon}$ in an antcoordination game under an “imitation with mutation” protocol.

where the limits refer to weak convergence of probability measures (see Exercise 12.1.4), and δ_{χ^*} denotes the point mass at state χ^* .

The remaining pictures illustrate the effects of setting very small mutation rates. When $N = 100$ and $\varepsilon = 10^{-5}$ (Figure 12.1.1(iii)) most of the mass in $\mu^{100,\varepsilon}$ falls in a bell-shaped distribution centered at state $\chi = .68$, but a mass of $\mu_1^{100,\varepsilon} = .0460$ sits in isolation at the boundary state $\chi = 1$. When ε is reduced to 10^{-7} (Figure 12.1.1(iv)), this boundary state commands a majority of the weight in the distribution ($\mu_1^{100,\varepsilon} = .8286$).

This numerical analysis suggests that when the mutation rate approaches zero, the stationary distribution will approach a point mass at state 1. We also confirm this conjecture in Section 12.4.1. Increasing the population size does not alter this result, so for the small noise double limit we obtain

$$(12.14) \quad \lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \mu^{N,\varepsilon} = \delta_1,$$

where δ_1 denotes the unit point mass at state 1.

Comparing equations (12.13) and (12.14), we conclude that the large population double limit and the small noise double limit disagree. §

In the preceding example, the (single and double) large population limits agree with the predictions of the mean dynamic, while the small noise limits do not. Still, the behavior of the latter limits is easy to explain. Starting from any interior state, and from the boundary as well when $\varepsilon > 0$, the expected motion of the process $\{X_t^{N,\varepsilon}\}$ is toward the interior rest point of the mean dynamic V^ε . But when ε is zero, the boundary states 0 and 1 become rest points of V^ε , and are absorbing states of $\{X_t^{N,\varepsilon}\}$; in fact, it is easy to see that they are the only recurrent states of the zero-noise process. Therefore, when $\varepsilon = 0$, $\{X_t^{N,\varepsilon}\}$ reaches either state 0 or state 1 in finite time, and then remains at that state forever.

If instead ε is positive, the boundary states are no longer absorbing, and they are far from any rest point of the mean dynamic. But once the process $\{X_t^{N,\varepsilon}\}$ reaches such a state, it can only depart by way of a mutation. Thus, if we fix the population size N and make ε extremely small, then a journey from an interior state to a boundary state—here a journey against the flow of the mean dynamic—is “more likely” than an escape from a boundary state by way of a single mutation. It follows that in the small noise limit, the stationary distribution must become concentrated on the boundary states regardless of the nature of the mean dynamic. (In fact, it will typically become concentrated on just one of these states—see Section 12.4.1.)

As this discussion indicates, the prediction provided by the small noise limit does not become a good approximation of behavior at fixed values of N and ε unless ε is so small

enough that lone mutations are much more rare than excursions from the interior of \mathcal{X}^N to the boundary. In Figures 12.1.1(iii) and (iv), which consider a modest population size of $N = 100$, we see that a mutation rate of $\varepsilon = 10^{-5}$ is not small enough to yield agreement with the prediction of the small noise limit, though a mutation rate of $\varepsilon = 10^{-7}$ yields a closer match. With larger population sizes, the relevant mutation rates would be even smaller.

This example suggests that in economic contexts, where the probabilities of “mutations” may not be especially small, the large population limit is more likely to be the relevant one in cases where the predictions of the two limits disagree. In biological contexts, where mutation rates may indeed be quite small, the choice between the limits seems less clear.

12.2 Exponential Protocols and Potential Games

As our first study of the asymptotics of the stationary distribution, we consider populations of agents who play potential games using exponential protocols. In Section 11.5, we saw that for fixed values of η and N , the stationary distributions $\mu^{N,\eta}$ arising in this setting take only two forms, one for direct protocols and one for imitative protocols. For this reason, our focus will be on the mechanics of the limiting analysis, and on the differences between the predictions generated by the various limits introduced above.

Recall from Section 11.5 that $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$ is a finite-population potential game with potential function $f^N : \mathcal{X}^N \rightarrow \mathbf{R}$ if

$$F_j^N(x + \frac{1}{N}(e_j - e_i)) - F_i^N(x) = f^N(x + \frac{1}{N}(e_j - e_i)) - f^N(x) \text{ for all } x \in \mathcal{X}_i^N \text{ and } i, j \in S.$$

Theorem 11.5.12 shows that if clever agents play F^N recurrently using a direct exponential protocol, then the stationary distribution of the resulting evolutionary process takes the form

$$(12.15) \quad \mu_x^{N,\eta} = \frac{1}{K^N} \frac{N!}{\prod_{k \in S} (Nx_k)!} \exp(\eta^{-1} f^N(x)) \text{ for all } x \in \mathcal{X}^N,$$

where K^N is determined by the requirement that the stationary distribution weights sum to one. Similarly, Theorem 11.5.13 shows that if clever agents play F^N recurrently using an imitative exponential protocol, and if in addition there is one committed agent for each

strategy in S , then the stationary distribution of the evolutionary process takes the form

$$(12.16) \quad \mu_x^N = \frac{1}{\kappa^N} \exp(\eta^{-1} f^N(x)) \text{ for all } x \in \mathcal{X}^N.$$

The remainder of this section studies the asymptotic behavior of these distributions as η and N approach their limiting values.

One new piece of notation will be handy for describing the rates of decay of stationary distribution weights. If $g: C \rightarrow \mathbf{R}$ is a continuous real-valued function on the compact set C (which below will be either \mathcal{X}^N or X), we define the function $\Delta g: C \rightarrow \mathbf{R}_-$ by

$$(12.17) \quad \Delta g(x) = g(x) - \max_{y \in C} g(y).$$

Thus, Δg is obtained from g by shifting its values uniformly, doing so in such a way that the maximum value of Δg is zero.

12.2.1 Direct Exponential Protocols: The Small Noise Limit

We begin by studying the asymptotics of the stationary distribution under direct exponential protocols, starting with the simplest case of small noise asymptotics. Equation (12.15) immediately implies that

$$(12.18) \quad \frac{\mu_x^{N,\eta}}{\mu_y^{N,\eta}} = \frac{\prod_{k \in S} (Ny_k)!}{\prod_{k \in S} (Nx_k)!} \exp(\eta^{-1} (f^N(x) - f^N(y))) \text{ for all } x, y \in \mathcal{X}^N.$$

Taking logarithms of both sides of (12.18) and multiplying by the noise level η , we obtain

$$(12.19) \quad \lim_{\eta \rightarrow 0} \eta \log \frac{\mu_x^{N,\eta}}{\mu_y^{N,\eta}} = \lim_{\eta \rightarrow 0} \left((f^N(x) - f^N(y)) + \eta \log \left(\frac{\prod_{k \in S} (Ny_k)!}{\prod_{k \in S} (Nx_k)!} \right) \right) \\ = f^N(x) - f^N(y).$$

Thus, when the noise level is small, the multinomial terms become inconsequential compared to the potential function values. This observation is enough to identify the maximizers of the potential function f^N as the stochastically stable states.

Theorem 12.2.1. *Suppose that clever agents using a direct exponential protocol play a finite-population potential game F^N with potential function f^N . Then in the small noise limit, the stochastically stable states are those in $\operatorname{argmax}_{x \in \mathcal{X}^N} f^N(x)$.*

Combining equation (12.19) with the fact (proved below) that the the stationary dis-

tribution weights on stochastically stable states have a zero rate of decay, we obtain the following result characterizing the rates of decay at all states.

Theorem 12.2.2. *Under the assumptions of Theorem 12.2.1,*

$$(12.20) \quad \lim_{\eta \rightarrow 0} \eta \log \mu_x^{N,\eta} = \Delta f^N(x) \text{ for all } x \in \mathcal{X}^N.$$

Theorem 12.2.2 says that as η^{-1} approaches infinity, the stationary distribution weight $\mu_x^{N,\eta}$ vanishes at exponential rate $\Delta f^N(x)$. This rate is the difference between the value of the potential function at x and its maximal value.

Proof of Theorem 12.2.2. Let x_*^N be a maximizer of f^N on \mathcal{X}^N . If we could show that

$$(12.21) \quad \lim_{\eta \rightarrow 0} \eta \log \mu_{x_*^N}^{N,\eta} = 0,$$

it would follow from equation (12.19) that

$$(12.22) \quad \begin{aligned} \lim_{\eta \rightarrow 0} \eta \log \mu_x^{N,\eta} &= \lim_{\eta \rightarrow 0} \left(\eta \log \frac{\mu_x^{N,\eta}}{\mu_{e_1}^{N,\eta}} - \eta \log \frac{\mu_{x_*^N}^{N,\eta}}{\mu_{e_1}^{N,\eta}} + \eta \log \mu_{x_*^N}^{N,\eta} \right) \\ &= (f^N(x) - f^N(e_1)) - (f^N(x_*^N) - f^N(e_1)) \\ &= \Delta f^N(x), \end{aligned}$$

proving the theorem.

The argument used to establish (12.21) follows the same lines as the first part of the proof of Proposition 12.1.5. First, suppose contrary to (12.21) that there is a sequence $\{\eta^k\}$ converging to zero along which the limit in (12.21) is $-c < 0$. In this case, the reasoning in equation (12.22) implies that

$$\lim_{\eta^k \rightarrow 0} \eta^k \log \mu_x^{N,\eta^k} = \Delta f^N(x) - c \leq -c \text{ for all } x \in \mathcal{X}^N.$$

It follows that for η^k far enough along the sequence, we have

$$\sum_{x \in \mathcal{X}^N} \mu_x^{N,\eta^k} = \sum_{x \in \mathcal{X}^N} \exp\left(\frac{1}{\eta^k} \cdot \eta^k \log \mu_x^{N,\eta^k}\right) \leq \#\mathcal{X}^N \cdot \exp\left(-\frac{c}{2\eta^k}\right),$$

The last expression vanishes as k grows large, contradicting the fact that μ^{N,η^k} is a probability measure.

Second, suppose contrary to (12.21) that there is a sequence $\{\eta^k\}$ converging to zero along which the limit in (12.21) is $c > 0$. Then by definition, there is a sequence $\{\delta^k\}$ converging to zero such that $\mu_{x_*^N}^{N,\eta^k} = \exp((\eta^k)^{-1}(c + \delta^k))$. This expression grows without bound as k grows large, contradicting the fact that each μ^{N,η^k} is a probability measure. This completes the proof of the theorem. ■

12.2.2 Direct Exponential Protocols: The Large Population Limit

In Section 12.1.2, we noted that the analysis of the large population limit is more complicated than that of the small noise limit, owing to the dependence of the state space \mathcal{X}^N on the population size N . We must contend with this and a few new technical issues to obtain our next result, Theorem 12.2.3.

First, to obtain any sort of limit result, we need to consider a sequence of potential games $\{F^N\}_{N=N_0}^\infty$ that “settle down” in some sense. Theorem 12.2.3 requires the rescaled potential functions $\{\frac{1}{N}f^N\}_{N=N_0}^\infty$ to converge uniformly to a limit function $f: X \rightarrow \mathbf{R}$. This requirement is rather weak: Theorem 11.5.7 and the subsequent exercises show that the requirement is necessary but not sufficient for the potential games themselves to converge.

Second, taking the large population limit does not cause the multinomial term from equation (12.15) to vanish, as it did in the small noise limit. Theorem 12.2.3 shows that this term can be captured using the logit potential function $f^\eta: X \rightarrow \mathbf{R}$, which we first encountered when studying convergence of the logit dynamic in potential games in Section 7.1. This function, presented in equation (7.5), is given by

$$f^\eta(x) = f(x) - \eta \sum_{i \in S} x_i \log x_i$$

where $0 \log 0 = 0$ as usual. The second term in f^η , given by η times the *entropy function* $h(x) = -\sum_{i \in S} x_i \log x_i$, increases the value of f^η at states representing “more random” probability distributions, the more so the higher the value of the noise level η .

The original use of f^η in Theorem 7.1.6 was as a strict Lyapunov function for the logit(η) dynamic. In the present context, we show now that as N grows large, the stationary distribution weights $\mu_x^{N,\eta}$ decay at rates determined by $\frac{1}{\eta} \Delta f^\eta$.

Theorem 12.2.3. *Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of finite-population potential games whose rescaled potential functions $\{\frac{1}{N}f^N\}_{N=N_0}^\infty$ converge uniformly to the C^1 function $f: X \rightarrow \mathbf{R}$. Suppose that agents are clever and employ a direct exponential protocol with noise level $\eta > 0$. Then the sequence*

of stationary distributions $\{\mu^{N,\eta}\}_{N=N_0}^\infty$ satisfies

$$(12.23) \quad \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - \Delta f^\eta(x) \right| = 0.$$

Corollary 12.2.4. *Under the conditions of Theorem 12.2.3, state x is weakly stochastically stable in the large population limit if and only if $x \in \operatorname{argmax}_{y \in X} f^\eta(y)$.*

Proof of Theorem 12.2.3. Equation (12.18) tells us that for any $x \in \mathcal{X}^N$, we have

$$(12.24) \quad \frac{\mu_x^{N,\eta}}{\mu_{e_1}^{N,\eta}} = \frac{N!}{\prod_{i \in S} (Nx_i)!} \exp\left(\eta^{-1} \left(f^N(x) - f^N(e_1)\right)\right).$$

By Stirling's formula (see the Notes), $N! = \sqrt{2\pi N} N^N \exp(-N + \frac{\theta^N}{N})$ for some $\theta^N \in [0, \frac{1}{12}]$. Therefore, letting $S(x) = \{i \in S : x_i > 0\}$ and $n(x) = \#S(x)$, we can rewrite (12.24) as

$$\begin{aligned} \frac{\mu_x^{N,\eta}}{\mu_{e_1}^{N,\eta}} &= \frac{\sqrt{2\pi N} N^N \exp(-N + \frac{\theta^N}{N})}{\prod_{i \in S(x)} \sqrt{2\pi N x_i} (Nx_i)^{Nx_i} \exp(-Nx_i + \frac{\theta^{Nx_i}}{Nx_i})} \exp\left(\eta^{-1} \left(f^N(x) - f^N(e_1)\right)\right) \\ &= \exp\left(\eta^{-1} \left(f^N(x) - f^N(e_1)\right)\right) \left(\prod_{i \in S(x)} x_i^{-(Nx_i+1/2)} \right) \cdot \frac{1}{(2\pi N)^{n(x)-1/2}} \exp\left(\frac{\theta^N}{N} - \sum_{i \in S(x)} \frac{\theta^{Nx_i}}{Nx_i}\right). \end{aligned}$$

Taking logarithms and multiplying by $\frac{\eta}{N}$ yields

$$\begin{aligned} \frac{\eta}{N} \log \frac{\mu_x^{N,\eta}}{\mu_{e_1}^{N,\eta}} &= \frac{1}{N} f^N(x) - \frac{1}{N} f^N(e_1) - \eta \sum_{i \in S(x)} x_i \log x_i \\ &\quad + \frac{\eta}{N} \left(\frac{\theta^N}{N} - \sum_{i \in S(x)} \left(\frac{\theta^{Nx_i}}{Nx_i} + \frac{1}{2} \log x_i \right) - \frac{n(x)-1}{2} \log 2\pi N \right). \end{aligned}$$

Since any $x \in \mathcal{X}^N$ with $x_i > 0$ satisfies $x_i \geq \frac{1}{N}$, it follows that

$$\left| \frac{1}{N} \sum_{i \in S(x)} \left(\frac{\theta^{Nx_i}}{Nx_i} + \frac{1}{2} \log x_i \right) \right| \leq \left| \frac{n(x)}{12N} \right| + \left| \frac{n(x)}{2N} \log \frac{1}{N} \right|.$$

Combining the last three equations, and using the facts that $\{\frac{1}{N} f^N\}$ converges uniformly to f and that $f^\eta(e_1) = f(e_1)$, we find that

$$(12.25) \quad \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \frac{\mu_x^{N,\eta}}{\mu_{e_1}^{N,\eta}} - (f^\eta(x) - f^\eta(e_1)) \right| = 0.$$

The argument that (12.25) implies (12.23) is similar to the proof of Theorem 12.2.2. ■

Theorem 12.2.3 shows how the both the rates of decay of the stationary distribution weights and the location of the stochastically stable state depend on the noise level η . The larger is η , the more weight that f^η places on the entropy term, and hence the better the prospects for visitation for states far from the boundary of the simplex. We illustrate this point in the next example. (If the dependence described here seems inevitable, it is not: see Section 12.2.4.)

Example 12.2.5. Suppose that agents are matched to play 123 Coordination (cf Examples 3.1.8, 6.2.4, and 7.1.7):

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}.$$

Whether the population is finite or continuous, and whether self-matching is allowed or prohibited, the resulting population game is a potential game. In the continuous-population case, the potential function of the potential game is the convex function

$$f(x) = \frac{1}{2} \left((x_1)^2 + 2(x_2)^2 + 3(x_3)^2 \right),$$

which is illustrated in Figure 3.1.1. The potential functions for the finite-population games are described in Exercise 11.5.2.

Suppose that agents are clever and employ a direct exponential protocol with noise level $\eta > 0$. Theorem 12.2.3 tells us that as the population size N grows large, the rates of decay of the stationary distribution weights are determined by the logit potential function

$$\begin{aligned} f^\eta(x) &= \frac{1}{2} \left((x_1)^2 + 2(x_2)^2 + 3(x_3)^2 \right) - \eta(x_1 \log x_1 + x_2 \log x_2 + x_3 \log x_3) \\ &= f(x) + \eta h(x). \end{aligned}$$

The function f^η is the weighted sum of the convex potential function f and the concave entropy function h , with the weight on the latter given by the noise level η .

At larger values of η , f^η is dominated by its entropy term. We illustrate this in Figure 12.2.1, which presents a graph and contour plot of f^η when $\eta = 1.5$. Here f^η appears concave. Its unique maximizer, and hence the uniquely stochastically stable state in the large population limit, is $x^* \approx (.2114, .2596, .5290)$, which puts significant weight on all three strategies. Thus, if the revision protocol is noisy, then in the infinite horizon, the population tends to congregate at states at which all three strategies are common.

More moderate noise levels balance the contributions of f and h to f^η . We illustrate







Figure 12.2.1: Graph (i) and contour plot (ii) of the $\logit(1.5)$ potential function for 123 Coordination.




(i)




(ii)

Figure 12.2.2: Graph (i) and contour plot (ii) of the $\text{logit}(.6)$ potential function for 123 Coordination.



(i)



(ii)

Figure 12.2.3: Graph (i) and contour plot (ii) of the $\text{logit}(2)$ potential function for 123 Coordination.

this in Figure 12.2.2, which presents a graph and contour plot of f^η when $\eta = .6$. This logit potential function does not look especially concave or convex. The uniquely stochastically stable state, $x^* \approx (.0072, .0073, .9855)$, places nearly all weight on strategy 3. There is another local maximum of f^η at $x^2 \approx (.0502, .8864, .0634)$, which by Theorem 8.2.4 is a locally stable rest point of the deterministic logit dynamic. Thus, if the process $\{X_t^{N,\eta}\}$ starts off near x^2 , it is likely to approach x^2 and then remain there for a large amount of time. Nevertheless, x^2 is not stochastically stable, and the stationary distribution weights on nearby states decay at rates of about

$$\frac{1}{\eta} \Delta f^\eta(x^2) = \frac{1}{.6} (f^\eta(x^*) - f^\eta(x^2)) = \frac{10}{6} (1.5084 - 1.0522) = .7603.$$

When η is small, f^η closely resembles the original potential function f . Figure 12.2.3 presents a graph and contour plot of f^η when $\eta = .2$; these pictures are barely distinguishable from the pictures of f in Figure 3.1.1. In addition to the stochastically stable state $x^* \approx e_3$, the function f^η has local maxima at states $x^2 \approx e_2$ and $x^1 \approx (.9841, .0078, .0081)$. The rates of decay of the stationary distribution weights near the latter two states are

$$\begin{aligned} \frac{1}{\eta} \Delta f^\eta(x^2) &= \frac{1}{.2} (f^\eta(x^*) - f^\eta(x^2)) = 5(1.5 - 1) = 2.5 \quad \text{and} \\ \frac{1}{\eta} \Delta f^\eta(x^1) &= \frac{1}{.2} (f^\eta(x^*) - f^\eta(x^1)) = 5(1.5 - .5029) = 4.9855. \quad \S \end{aligned}$$

12.2.3 Direct Exponential Protocols: Double Limits

In Section 12.1.3, we argued that the best way to compare the predictions of the small noise and large population limits is to compare the two double limits: the small noise double limit, whose interior limit has η approaching zero, and the large population double limit, whose interior limit has N approaching infinity. In the current setting, the behavior of the double limits is easily deduced from the previous theorems, along with the facts that the sequences of functions $\{\frac{1}{N} \Delta f^N\}_{N=N_0}^\infty$ and $\{\Delta f^\eta\}_{\eta \in (0, \bar{\eta}]}$ converge uniformly to f .

Corollary 12.2.6. *Under the conditions of Theorem 12.2.3, the stationary distributions $\mu^{N,\eta}$ satisfy*

- (i) $\lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - \Delta f(x) \right| = 0$ and
- (ii) $\lim_{\eta \rightarrow 0} \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - \Delta f(x) \right| = 0.$

Corollary 12.2.6 tells us that when agents play potential games using direct exponential protocols, the predictions generated by the small noise and large population limits agree in a strong sense: in both cases, when the other parameter is held fixed at a value “close to its limit”, the rates of decay of stationary distribution weights are approximately described

by differences in the potential function f . Thus, the choice between the two analyses has limited consequences for our forecasts of infinite horizon behavior.

12.2.4 Imitative Exponential Protocols with Committed Agents

To close this section, we compute limiting stationary distributions for potential games under imitative exponential protocols. To ensure irreducibility, we assume that there is one committed agent for each available strategy. Our analysis follows a similar path to the one above, but starting from the stationary distribution formula (12.16) rather than (12.15). The only difference between these formulas is that formula (12.15) for the direct protocol contains a multinomial term, while formula (12.16) for the imitative protocol does not. This not only makes the analysis of the imitative case simpler, but also alters our prediction of behavior in the large population limit.

Equation (12.16) implies that under an imitative exponential protocol, ratios between stationary distribution weights are given by

$$\frac{\mu_x^{N,\eta}}{\mu_y^{N,\eta}} = \exp\left(\eta^{-1}\left(f^N(x) - f^N(y)\right)\right) \text{ for all } x, y \in \mathcal{X}^N.$$

This implies in turn that

$$\frac{\eta}{N} \log \frac{\mu_x^{N,\eta}}{\mu_y^{N,\eta}} = \frac{1}{N} \left(f^N(x) - f^N(y)\right).$$

From here, the same argument used to prove Theorem 12.2.2 yields

Theorem 12.2.7. *Suppose that clever agents using an imitative exponential protocol play a finite-population potential game F^N with potential function f^N , and that there is one committed agent for each strategy in F^N . Then for all $x \in \mathcal{X}^N$,*

$$\lim_{\eta \rightarrow 0} \eta \log \mu_x^{N,\eta} = \Delta f^N(x).$$

Together, Theorems 12.2.2 and 12.2.7 show that in the small noise limit, infinite horizon behavior under direct and imitative protocols looks the same. The explanation is simple: the stationary distributions (12.15) and (12.16) only differ by the inclusion in the former of the multinomial term, a term that becomes inconsequential in the small noise limit.

Turning to the large population limit, Theorem 12.2.3 showed that under direct exponential protocols, the rates of decay of the stationary distribution are described by $\frac{1}{\eta} \Delta f^\eta$.

As Example 12.2.5 showed, high noise levels lead to predictions far from what one would expect from the potential function f alone.

In the imitative case, the absence of the multinomial term from (12.16) means that an entropy term is no longer needed to describe the large population asymptotics. Theorem 12.2.8 shows that as N grows large, the stationary distribution weights of the stochastic process $\{X_t^N\}$ decay at rates determined not by $\frac{1}{\eta} \Delta f^\eta$, but by $\frac{1}{\eta} \Delta f$. Thus, while higher noise levels lead to lower exponential rates of decay, they do not introduce any distortions of potential.

Theorem 12.2.8. *Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of finite-population potential games whose rescaled potential functions $\{\frac{1}{N} f^N\}_{N=N_0}^\infty$ converge uniformly to the C^1 function $f: X \rightarrow \mathbf{R}$. Suppose that agents are clever and employ an imitative exponential protocol with noise level $\eta > 0$, and that there is one committed agent for each strategy in F^N . Then the sequence of stationary distributions $\{\mu^{N,\eta}\}_{N=N_0}^\infty$ satisfies*

$$\lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_x^{N,\eta} - \Delta f(x) \right| = 0.$$

Corollary 12.2.9. *Under the conditions of Theorem 12.2.8, state x is weakly stochastically stable in the large population limit if and only if $x \in \operatorname{argmax}_{y \in X} f(y)$.*

At a first encounter, the fact that imitative protocols lead to cleaner equilibrium selection results than direct protocols may be unexpected. Recalling some of our earlier results on deterministic dynamics makes this conclusion less surprising. For instance, while we argued in Section 6.2 that the logit dynamic violates positive correlation (PC), and that its rest points are perturbed equilibria rather than Nash equilibria, we saw in Section 5.4 that the i-logit dynamic (Example 5.4.9) satisfies (PC), and that its rest points include all Nash equilibria of the underlying game. Evidently, by enabling agents to take advantage of their opponents' experiences, imitative protocols generate aggregate behavior that respects the incentives in the underlying game. This is so even when agents employing the protocol barely distinguish between strategies with high and low payoffs.

Exercise 12.2.10. To verify that the double limits agree in this context, state and prove an analogue of Corollary 12.2.6 for imitative exponential protocols with committed agents. ‡

Of course, if we did not introduce committed agents, our results for imitative protocols would be very different, even if the irreducibility of the evolutionary process were ensured in some other way: see Example 12.1.9, as well as Section 12.4.1 below.

12.3 Noisy Best Response Protocols in Two-Strategy Games

In this section we return to the other reversible setting from Chapter 11: that of two-strategy games. Since here reversibility places no restrictions on the form of the revision protocol, the equilibrium selection results are not preordained: in many games, different protocols lead to different stochastically stable states.

To begin this section, we define a general class of noisy best response protocols. In doing so, we introduce the important notion of the *cost* of a suboptimal choice, which is defined as the rate of decay of the probability of making this choice as the noise level approaches zero. Using this notion, we derive simple formulas that characterize the asymptotics of the stationary distribution under the various limits in η and N , and offer a necessary and sufficient condition for an equilibrium to be uniquely stochastically stable under every noisy best response protocol.

In this section and the next, we employ our usual notation for two-strategy games: we identify state x with the weight $\chi \equiv x_1$ that it places on strategy 1, so that the state space becomes $\mathcal{X}^N = \{0, \frac{1}{N}, \dots, 1\} \subset [0, 1]$, and we also write $F(\chi)$ for $F(x)$ and $\rho(\pi, \chi)$ for $\rho(\pi, x)$ whenever doing so is convenient. Also, many of the proofs below use the Dominated Convergence Theorem to justify switching the order of limits and integration; see the Notes for references on this result.

12.3.1 Noisy Best Response Protocols and their Cost Functions

Until Section 12.4.1, we focus on evolution under *noisy best response protocols*. These protocols can be expressed as

$$(12.26) \quad \rho_{ij}^\eta(\pi) = \sigma^\eta(\pi_j - \pi_i),$$

for some function $\sigma^\eta: \mathbf{R} \rightarrow (0, 1)$: when a current strategy i player receives a revision opportunity, he switches to strategy $j \neq i$ with a probability that only depends on the payoff advantage of strategy j over strategy i . To justify its name, the protocol σ^η should recommend optimal strategies with high probability when the noise level is small:

$$\lim_{\eta \rightarrow 0} \sigma^\eta(a) = \begin{cases} 1 & \text{if } a > 0, \\ 0 & \text{if } a < 0. \end{cases}$$

To place further structure on the probabilities of suboptimal choices, we impose restrictions on the rates at which the probabilities $\sigma^\eta(a)$ of choosing a suboptimal strategy

approach zero as η approaches zero. To do so, we define the *cost* of switching to a strategy with payoff *disadvantage* $d \in \mathbf{R}$ as

$$(12.27) \quad \kappa(d) = -\lim_{\eta \rightarrow 0} \eta \log \sigma^\eta(-d).$$

By unpacking this expression, we can write the probability of switching to a strategy with payoff disadvantage d when the noise level is η as

$$\sigma^\eta(-d) = \exp\left(-\eta^{-1}(\kappa(d) + o(1))\right),$$

where $o(1)$ represents a term that vanishes as η approaches 0. Thus, $\kappa(d)$ is the exponential rate of decay of the choice probability $\sigma^\eta(-d)$ as η^{-1} approaches infinity. Later on, in Section 12.5.1, we will introduce a more restrictive notion of cost that is also common in the literature; see that section and Appendix 12.A.5 for further discussions.

We are now ready to define the class of protocols we will consider.

Definition. *We say that the noisy best response protocol (12.26) is regular if*

- (i) *the limit in (12.27) exists for all $d \in \mathbf{R}$, with convergence uniform on compact intervals;*
- (ii) *κ is nondecreasing;*
- (iii) *$\kappa(d) = 0$ whenever $d < 0$;*
- (iv) *$\kappa(d) > 0$ whenever $d > 0$.*

Conditions (ii)-(iv) impose constraints on the rates of decay of switching probabilities. Condition (ii) requires the rate of decay to be nondecreasing in the payoff disadvantage of the alternative strategy. Condition (iii) requires the switching probability of an agent currently playing the suboptimal strategy to have rate of decay zero; the condition is satisfied when the probability is bounded away from zero, although this is not necessary for the condition to hold. Finally, condition (iv) requires the probability of switching from the optimal strategy to the suboptimal one to have a positive rate of decay. These conditions are consistent with having either $\kappa(0) > 0$ or $\kappa(0) = 0$: thus, when both strategies earn the same payoff, the probability that a revising agent opts to switch strategies can converge to zero with a positive rate of decay, as in Example 12.3.1 below, or can be bounded away from zero, as in Examples 12.3.2 and 12.3.3.

We now present the three leading examples of noisy best response protocols.

Example 12.3.1. Best response with mutations. The BRM protocol with noise level $\eta (=$

$-(\log \varepsilon)^{-1}$), introduced in Example 11.1.1, is defined by

$$\sigma^\eta(a) = \begin{cases} 1 - \exp(-\eta^{-1}) & \text{if } a > 0, \\ \exp(-\eta^{-1}) & \text{if } a \leq 0. \end{cases}$$

In this specification, an indifferent agent only switches strategies in the event of a mutation. Since for $d \geq 0$ we have $-\eta \log \sigma^\eta(-d) = 1$, protocol σ^η is regular with cost function

$$\kappa(d) = \begin{cases} 1 & \text{if } d \geq 0, \\ 0 & \text{if } d < 0. \end{cases} \quad \S$$

Example 12.3.2. Logit choice. The *logit choice protocol* with noise level $\eta > 0$, introduced in Examples 6.2.1 and 11.1.3, is defined by

$$\sigma^\eta(a) = \frac{\exp(\eta^{-1}a)}{\exp(\eta^{-1}a) + 1}.$$

For $d \geq 0$, we have that $-\eta \log \sigma^\eta(-d) = d + \eta \log(\exp(-\eta^{-1}d) + 1)$. It follows that σ^η is regular with cost function

$$\kappa(d) = \begin{cases} d & \text{if } d > 0, \\ 0 & \text{if } d \leq 0. \end{cases} \quad \S$$

Example 12.3.3. Probit choice. In Example 6.2.5, we saw that the logit choice protocol can be derived from a random utility model in which the strategies' payoffs are perturbed by i.i.d., double exponentially distributed random variables. The *probit choice protocol* assumes instead that the payoff perturbations are i.i.d. normal random variables with mean 0 and variance η . Thus

$$\sigma^\eta(a) = \mathbb{P}(\sqrt{\eta}Z + a > \sqrt{\eta}Z'),$$

where Z and Z' are independent and standard normal. It follows easily that

$$(12.28) \quad \sigma^\eta(a) = \Phi\left(\frac{a}{\sqrt{2\eta}}\right),$$

where Φ is the standard normal distribution function.

A well-known approximation of Φ (see the Notes) tells us that when $z < 0$,

$$(12.29) \quad \Phi(z) = K(z) \exp\left(\frac{-z^2}{2}\right)$$

for some $K(z) \in (\frac{-1}{\sqrt{2\pi z}}(1 - \frac{1}{z^2}), \frac{-1}{\sqrt{2\pi z}})$. It follows that $K(z) \in (\frac{-1}{2\sqrt{2\pi z}}, \frac{-1}{\sqrt{2\pi z}})$ whenever $z < -\sqrt{2}$. And one can verify directly that (12.29) holds with $K(z) \in [e \cdot \Phi(-\sqrt{2}), \frac{1}{2}]$ whenever $z \in [-\sqrt{2}, 0]$.

Now equations (12.28) and (12.29) imply that

$$(12.30) \quad -\eta \log \sigma^\eta(-d) = -\eta \log \Phi\left(\frac{-d}{\sqrt{2\eta}}\right) = \frac{1}{4}d^2 - \eta \log K\left(\frac{-d}{\sqrt{2\eta}}\right)$$

when $d \geq 0$, with our earlier estimates showing that

$$\begin{aligned} \eta \log K\left(\frac{-d}{\sqrt{2\eta}}\right) &\in \left(\frac{1}{2}\eta \log \eta - \eta \log 2\sqrt{\pi} d, \frac{1}{2}\eta \log \eta - \eta \log \sqrt{\pi} d\right) && \text{if } d > 2\sqrt{\eta}, \text{ and} \\ \eta \log K\left(\frac{-d}{\sqrt{2\eta}}\right) &\in \left[\eta\left(1 + \log \Phi(-\sqrt{2})\right), \eta \log \frac{1}{2}\right] && \text{if } d \in [0, 2\sqrt{\eta}]. \end{aligned}$$

Thus, for any $D > 0$ and any $\delta > 0$, we have $|\eta \log K(\frac{-d}{\sqrt{2\eta}})| < \delta$ for all $d \in [0, D]$ once $\eta > 0$ is sufficiently small. We conclude from equation (12.30) that σ^η is regular with cost function

$$\kappa(d) = \begin{cases} \frac{1}{4}d^2 & \text{if } d > 0, \\ 0 & \text{if } d \leq 0. \quad \S \end{cases}$$

12.3.2 The Small Noise Limit

We now investigate the asymptotics of the stationary distribution $\mu^{N,\eta}$, starting with the small noise limit. For convenience we continue to assume clever payoff evaluation; the alternative assumption has only minor effects on the small noise limit (Exercise 12.3.5), and no effect at all on the large population limit or the double limits.

After accounting for clever payoff evaluation (see Section 11.4.2, especially equations (11.19) and (11.20)), Theorem 11.2.4 tells us that the stationary distribution $\mu^{N,\eta}$ for the game F^N under protocol σ^η satisfies

$$(12.31) \quad \begin{aligned} \eta \log \frac{\mu_\chi^{N,\eta}}{\mu_0^{N,\eta}} &= \eta \log \left(\prod_{j=1}^{N\chi} \frac{(N-j+1)}{j} \cdot \frac{\rho_{01}^\eta(\check{F}^N(\frac{j-1}{N}), \frac{j-1}{N})}{\rho_{10}^\eta(\check{F}^N(\frac{j-1}{N}), \frac{j}{N})} \right) \\ &= \sum_{j=1}^{N\chi} \left(-\eta \log \sigma^\eta \left(F_0^N\left(\frac{j-1}{N}\right) - F_1^N\left(\frac{j}{N}\right) \right) + \eta \log \sigma^\eta \left(F_1^N\left(\frac{j}{N}\right) - F_0^N\left(\frac{j-1}{N}\right) \right) + \eta \log \frac{N-j+1}{j} \right). \end{aligned}$$

With this motivation, we define the *relative cost function* $\tilde{\kappa}: \mathbf{R} \rightarrow \mathbf{R}$ by

$$(12.32) \quad \begin{aligned} \tilde{\kappa}(d) &= \lim_{\eta \rightarrow 0} \left(-\eta \log \sigma^\eta(-d) + \eta \log \sigma^\eta(d) \right) \\ &= \kappa(d) - \kappa(-d) \\ &= \begin{cases} \kappa(d) & \text{if } d > 0, \\ 0 & \text{if } d = 0, \\ -\kappa(-d) & \text{if } d \leq 0. \end{cases} \end{aligned}$$

Our assumptions on κ imply that $\tilde{\kappa}$ is nondecreasing, sign preserving ($\text{sgn}(\tilde{\kappa}(d)) = \text{sgn}(d)$), and odd ($\tilde{\kappa}(d) = -\tilde{\kappa}(-d)$). Now, equations (12.31) and (12.32) imply that

$$\begin{aligned} \lim_{\eta \rightarrow 0} \eta \log \frac{\mu_\chi^{N,\eta}}{\mu_0^{N,\eta}} &= \sum_{j=1}^{N\chi} \left(\kappa \left(F_1^N \left(\frac{j}{N} \right) - F_0^N \left(\frac{j-1}{N} \right) \right) - \kappa \left(F_0^N \left(\frac{j}{N} \right) - F_1^N \left(\frac{j-1}{N} \right) \right) \right) \\ &= \sum_{j=1}^{N\chi} \tilde{\kappa} \left(F_1^N \left(\frac{j}{N} \right) - F_0^N \left(\frac{j-1}{N} \right) \right). \end{aligned}$$

Therefore, if we define the function $I^N: \mathcal{X}^N \rightarrow \mathbf{R}$ by

$$(12.33) \quad I^N(\chi) = \sum_{j=1}^{N\chi} \tilde{\kappa} \left(F_1^N \left(\frac{j}{N} \right) - F_0^N \left(\frac{j-1}{N} \right) \right),$$

then an easy modification of the proof of Theorem 12.2.2 yields the following result.

Theorem 12.3.4. *Suppose that N clever agents play a two-strategy game F^N using a regular noisy best response protocol with cost function κ . Then for all $\chi \in \mathcal{X}^N$,*

$$(12.34) \quad \lim_{\eta \rightarrow 0} \eta \log \mu_\chi^{N,\eta} = \Delta I^N(\chi).$$

To build intuition, it is worth comparing this result to the corresponding result for potential games and direct exponential protocols, Theorem 12.2.2. In both cases, taking the small noise limit causes the combinatorial term to vanish, leaving only a term that depends on payoff differences. But in equation (12.34), the payoff differences are transformed by the relative cost function $\tilde{\kappa}$, which determines how these differences influence the asymptotics of the stationary distribution. In the case of logit choice, Example 12.3.2 shows that $\tilde{\kappa}$ is the identity function, and so the analysis here agrees with Theorem 12.2.2.

Theorem 12.3.4 can be used to prove equilibrium selection results for finite-population

games. We will postpone presenting such results until Section 12.5, and instead will study equilibrium selection in two-strategy games using our coming results on double limits (Theorem 12.3.11). Doing so allows us to delay some picky accounting for finite-population effects that is necessary to obtain exact selection results.

Exercise 12.3.5. State and prove the analogue of Theorem 12.3.4 for simple agents. ‡

12.3.3 The Large Population Limit

Stating our description of the large population limit requires a few additional assumptions and definitions. First, we suppose that the sequence of two-strategy games $\{F^N\}_{N=N_0}^\infty$ converges uniformly to a continuous-population game F . We assume throughout that $F: [0, 1] \rightarrow \mathbf{R}^2$ is a continuous function. Second, we let

$$F_\Delta(\chi) \equiv F_1(\chi) - F_0(\chi)$$

denote the payoff advantage of strategy 1 at state x in the limit game. Third, we let

$$(12.35) \quad \tilde{\sigma}^\eta(a) = \frac{\sigma^\eta(a)}{\sigma^\eta(-a)}$$

to be the ratio of the probability of switching to a strategy with payoff advantage a to the probability of switching to a strategy with payoff advantage $-a$. Finally, we define the function $I^\eta: [0, 1] \rightarrow \mathbf{R}$ by

$$(12.36) \quad I^\eta(\chi) = \int_0^\chi \eta \log \tilde{\sigma}^\eta(F_\Delta(y)) \, dy - \eta (\chi \log \chi + (1 - \chi) \log(1 - \chi)),$$

where $0 \log 0 = 0$.

Theorem 12.3.6. *Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of two-strategy games that converges uniformly to the continuous-population game F . Suppose that agents are clever and employ regular noisy best response protocol σ^η . Then the sequence of stationary distributions $\{\mu^{N,\eta}\}_{N=N_0}^\infty$ satisfies*

$$\lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_\chi^{N,\eta} - \Delta I^\eta(\chi) \right| = 0.$$

Proof. Equations (12.31) and (12.35) imply that

$$(12.37) \quad \frac{\eta}{N} \log \frac{\mu_\chi^{N,\eta}}{\mu_0^{N,\eta}} = \frac{\eta}{N} \sum_{j=1}^{N\chi} \left(\log \tilde{\sigma}^\eta \left(F_1^N \left(\frac{j}{N} \right) - F_0^N \left(\frac{j-1}{N} \right) \right) + \log \frac{N-j+1}{N} - \log \frac{j}{N} \right).$$

Since σ^η is bounded away from zero, and since $0 \geq \log\left(\frac{\lceil N\chi \rceil}{N}\right) \geq \log(\chi)$ and $0 \geq \log\left(\frac{N - \lceil N\chi \rceil + 1}{N}\right) \geq \log(1 - \chi)$ for $\chi \in (0, 1)$, the Dominated Convergence Theorem implies that the Riemann sum in (12.37) converges to an integral. In particular, we have

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{\eta}{N} \log \frac{\mu_\chi^{N,\eta}}{\mu_0^{N,\eta}} &= \int_0^\chi \eta \left(\log \tilde{\sigma}^\eta(F_1(y) - F_0(y)) + \log(1 - y) - \log(y) \right) dy \\ &= \int_0^\chi \eta \log \tilde{\sigma}^\eta(F_\Delta(y)) dy - \eta(\chi \log \chi + (1 - \chi) \log(1 - \chi)) \\ &= I^\eta(\chi), \end{aligned}$$

with the limit, which is taken over those N for which $\chi \in \mathcal{X}^N$, being uniform in χ . The remainder of the proof is similar to that of Theorem 12.2.2. ■

Once again, there are similarities between this result and the corresponding result for potential games and direct exponential protocols, Theorem 12.2.3. As in the earlier result, the function I^η describing rates of decay is the sum of two terms: a term that depends on payoff differences in the limit game, and an entropy term. Also, the function I^η serves as a Lyapunov function for the relevant mean dynamic—compare Theorem 7.1.6 to Exercise 12.3.7 below. But as was the case in the previous theorem, the payoff differences appearing in I^η are transformed to account for the influence of the protocol σ^η .

- Exercise 12.3.7.* (i) Compute the mean dynamic obtained when agents use noisy best response protocol σ^η to play population game F .
(ii) Prove that I^η is a global strict Lyapunov function for this mean dynamic. ‡

12.3.4 Double Limits

Stating our results for double limits requires one final definition. Given a continuous-population game F and a noisy best response protocol with cost function κ , we define the *ordinal potential function* $I : [0, 1] \rightarrow \mathbf{R}$ by

$$(12.38) \quad I(\chi) = \int_0^\chi \tilde{\kappa}(F_\Delta(y)) dy,$$

where the relative cost function $\tilde{\kappa}$ is defined in equation (12.32). Observe that by marginally adjusting the state χ so as to increase the mass on the optimal strategy, we increase the value of I at rate $\tilde{\kappa}(a)$, where a is the optimal strategy's payoff advantage. Thus, the ordinal potential function combines information about payoff differences with the costs of the associated suboptimal choices.

Example 12.3.8. If ρ^η represents best response with mutations (Example 12.3.1), then the ordinal potential function (12.38) becomes the *signum potential function*

$$I_{\text{sgn}}(\chi) = \int_0^\chi \text{sgn}(F_\Delta(y)) \, dy.$$

This slope of this function at state χ is 1, -1 , or 0, according to whether the optimal strategy at χ is strategy 1, strategy 0, or both. §

Example 12.3.9. If ρ^η represents logit choice (Example 12.3.2), then (12.38) becomes the usual potential function

$$I_1(\chi) = \int_0^\chi F_\Delta(y) \, dy,$$

from Example 3.2.8, whose slope at state χ is just the payoff difference at χ . §

Example 12.3.10. If ρ^η represents probit choice (Example 12.3.3), then (12.38) becomes the *quadratic potential function*

$$I_2(\chi) = \int_0^\chi \frac{1}{4} \langle F_\Delta(y) \rangle^2 \, dy,$$

where $\langle a \rangle^2 = \text{sgn}(a) a^2$ is the signed square function. The values of I_2 again depend on payoff differences, but relative to the logit case, larger payoff differences play a more important role. This contrast can be traced to the fact that at small noise levels, the double exponential distribution has fatter tails than the normal distribution—compare Example 12.3.3. §

Theorem 12.3.11 shows that under either double limit, the rates of decay of the stationary distribution are captured by the ordinal potential function I . Since the double limits agree, our predictions of infinite horizon behavior under noisy best response rules do not depend on whether equilibrium selection is driven by the the small noise or large population limit.

Theorem 12.3.11. *Under the conditions of Theorem 12.3.6, the stationary distributions $\mu^{N,\eta}$ satisfy*

- (i) $\lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \max_{\chi \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_\chi^{N,\eta} - \Delta I(\chi) \right| = 0$ and
- (ii) $\lim_{\eta \rightarrow 0} \lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_\chi^{N,\eta} - \Delta I(\chi) \right| = 0.$

Proof. To prove part (i), one uses the Dominated Convergence Theorem to show that the Riemann sums $\frac{1}{N}I^N(\chi)$ converge uniformly to the integrals $I(\chi)$, in the sense that

$$\lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N}I^N(\chi) - I(\chi) \right| = 0.$$

This uniform convergence implies that

$$\lim_{N \rightarrow \infty} \max_{y \in \mathcal{X}^N} \frac{1}{N}I^N(y) = \max_{y \in [0,1]} I(y),$$

and hence that

$$\lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N}\Delta I^N(\chi) - \Delta I(\chi) \right| = \lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \left(\frac{1}{N}I^N(\chi) - \frac{1}{N}I^N(\chi_*) \right) - \left(I(\chi) - I(\chi_*) \right) \right| = 0,$$

where χ_*^N and χ_* maximize I^N and I , respectively. Since the limit in equation (12.34) is uniform (because \mathcal{X}^N is finite), we have

$$\begin{aligned} \lim_{\eta \rightarrow 0} \max_{\chi \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_{\chi}^{N,\eta} - \Delta I(\chi) \right| \\ \leq \lim_{\eta \rightarrow 0} \max_{\chi \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_{\chi}^{N,\eta} - \frac{1}{N}\Delta I^N(\chi) \right| + \lim_{\eta \rightarrow 0} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N}\Delta I^N(\chi) - \Delta I(\chi) \right| \\ = \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N}\Delta I^N(\chi) - \Delta I(\chi) \right|. \end{aligned}$$

Combining this inequality with the previous equation yields

$$\lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \max_{\chi \in \mathcal{X}^N} \left| \frac{\eta}{N} \log \mu_{\chi}^{N,\eta} - \Delta I(\chi) \right| \leq \lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N}\Delta I^N(\chi) - \Delta I(\chi) \right| = 0,$$

which proves part (i) of the theorem. The proof of part (ii) of the theorem is similar. ■

12.3.5 Stochastic Stability: Examples

By Theorem 12.3.11 (and the definitions in Section 12.1.3), the states that are weakly stochastically stable in the double limits are those that maximize the ordinal potential function I ; if this function has a unique maximizer, that state is uniquely stochastically stable. We now investigate in greater detail how a game's payoff function and the revision protocol's cost function interact to determine the stochastically stable states.

Stochastic stability analysis is most interesting when it allows us to select among multiple strict equilibria. For this reason, we focus the analysis to come on coordination

games. The two-strategy population game $F : [0, 1] \rightarrow \mathbf{R}^2$ is a *coordination game* if there is a state $\chi^* \in (0, 1)$ such that

$$\text{sgn}(\Delta F(\chi)) = \text{sgn}(\chi - \chi^*) \text{ for all } \chi \neq \chi^*.$$

Any ordinal potential function I for a coordination game is quasiconvex, with local maximizers at each boundary state. Because $I(0) \equiv 0$ by definition, Theorem 12.3.11 implies the following result.

Corollary 12.3.12. *Assume the conditions of Theorem 12.3.6, and suppose that the limit game F is a coordination game. Then state 1 is uniquely stochastically stable in both double limits if $I(1) > 0$, while state 0 is uniquely stochastically stable in both double limits if $I(1) < 0$.*

The next two examples, which revisit two games introduced in the previous chapter, show that the identity of the stochastically stable state may or may not depend on the revision protocol the agents employ.

Example 12.3.13. Stag Hunt revisited. In Example 11.2.5, we considered stochastic evolution in the Stag Hunt game

$$A = \begin{pmatrix} h & h \\ 0 & s \end{pmatrix},$$

where $s > h > 0$. When a continuous population of agents are matched to play this game, their expected payoffs are given by $F_H(\chi) = h$ and $F_S(\chi) = s\chi$, where χ denotes the proportion of agents playing Stag. This coordination game has two pure Nash equilibria, as well as a mixed Nash equilibrium that puts weight $\chi^* = \frac{h}{s}$ on Stag.

The ordinal potentials for the BRM, logit, and probit protocols in this game are

$$\begin{aligned} I_{\text{sgn}}(\chi) &= |\chi - \chi^*| - \chi^*, \\ I_1(x) &= \frac{s}{2}\chi^2 - h\chi, \text{ and} \\ I_2(x) &= \begin{cases} -\frac{s^2}{12}\chi^3 + \frac{hs}{4}\chi^2 - \frac{h^2}{4}\chi & \text{if } \chi \leq \chi^*, \\ \frac{s^2}{12}\chi^3 - \frac{hs}{4}\chi^2 + \frac{h^2}{4}\chi - \frac{h^3}{6s} & \text{if } \chi > \chi^*. \end{cases} \end{aligned}$$

Figure 12.3.1 presents the normalized functions ΔI_{sgn} , ΔI_1 , and ΔI_2 for two specifications of payoffs: $h = 2$ and $s = 3$ (in (i)), and $h = 2$ and $s = 5$ (in (ii)). For any choices of $s > h > 0$, ΔI is symmetric about its minimizer, the mixed Nash equilibrium $\chi^* = \frac{h}{s}$. As a result, the three protocols always agree about equilibrium selection: the all-Hare equilibrium




Figure 12.3.1: The ordinal potentials ΔI_{sgn} (blue), ΔI_1 (purple), and ΔI_2 (yellow) for Stag Hunt.




Figure 12.3.2: Payoffs and mixed equilibria in Nonlinear Stag Hunt when $h = 2$ and $s = 5, 5.75, 7,$ and 8.5 .

is uniquely stochastically stable when $\chi^* > \frac{1}{2}$ (or, equivalently, when $2h > s$), while the all-Stag equilibrium is uniquely stochastically stable when the reverse inequality holds. §

Example 12.3.14. Nonlinear Stag Hunt revisited. In Example 11.2.6, we introduced the nonlinear Stag Hunt game with payoff functions $F_H(\chi) = h$ and $F_S(\chi) = s\chi^2$, with χ again representing the proportion of agents playing Stag. This game has two pure Nash equilibria and a mixed equilibrium at $\chi^* = \sqrt{h/s}$. The payoffs and mixed equilibria for $h = 2$ and various choices of s are graphed in Figure 12.3.2.

The ordinal potentials for the BRM, logit, and probit models are given by

$$I_{\text{sgn}}(\chi) = |\chi - \chi^*| - \chi^*,$$




Figure 12.3.3: The ordinal potentials ΔI_{sgn} (blue), ΔI_1 (purple), and ΔI_2 (yellow) for Nonlinear Stag Hunt.

$$I_1(x) = \frac{s}{3}\chi^3 - h\chi, \text{ and}$$

$$I_2(x) = \begin{cases} -\frac{s^2}{20}x^5 + \frac{hs}{6}x^3 - \frac{h^2}{4}x & \text{if } \chi \leq \chi^*, \\ \frac{s^2}{20}x^5 - \frac{hs}{6}x^3 + \frac{h^2}{4}x - \frac{4h^2\chi^*}{15} & \text{if } \chi > \chi^*. \end{cases}$$

Figure 12.3.3 presents the functions ΔI_{sgn} , ΔI_1 , and ΔI_2 for $h = 2$ and for various choices of s .

When s is at its lowest level of 5, coordination on Stag is at its least attractive. Since $\chi^* = \sqrt{2/5} \approx .6325$, the basin of attraction of the all-Hare equilibrium is considerably larger than that of the all-Stag equilibrium. Figure 12.3.3(i) illustrates that coordination on Hare is stochastically stable under all three protocols.

If we make coordination on Stag somewhat more attractive by increasing s to 5.75, the mixed equilibrium becomes $\chi^* = \sqrt{2/5.75} \approx .5898$. The all-Hare equilibrium remains stochastically stable under the BRM and logit rules, but all-Stag becomes stochastically stable under the probit rule (Figure 12.3.3(ii)).

Increasing s further to 7 shifts the mixed equilibrium closer to the midpoint of the unit interval ($\chi^* = \sqrt{2/7} \approx .5345$). The BRM rule continues to select all-Hare, while the probit and logit rules both select all-Stag (Figure 12.3.3(iii)).

Finally, when $s = 8.5$, the all-Stag equilibrium has the larger basin of attraction ($\chi^* = \sqrt{2/8.5} \approx .4851$). At this point, coordination on Stag becomes attractive enough that all three protocols select the all-Stag equilibrium (Figure 12.3.3(iv)).

Why as we increase the value of s does the transition to selecting all-Stag occur first for the probit rule, then for the logit rule, and finally for the BRM rule? Examining Figure 12.3.2, we see that increasing s not only shifts the mixed Nash equilibrium to the left, but also markedly increases the payoff advantage of Stag at states where it is optimal. Since the cost function of the probit rule is the most sensitive to payoff differences, its equilibrium selection changes at the lowest level of s . The next selection to change is that of the (moderately sensitive) logit rule, and the last is the selection of the (insensitive) BRM rule. §

- Exercise 12.3.15.* (i) Construct a two-strategy coordination game in which the logit protocols selects a different equilibrium than the BRM and probit protocols.
- (ii) Construct a two-strategy game in which each the BRM, logit, and probit protocols each generates a distinct equilibrium selection. (Evidently the game cannot be a coordination game.) ‡

12.3.6 Risk Dominance, Stochastic Dominance, and Stochastic Stability

Building on these examples, we now seek general conditions on payoffs that ensure stochastic stability under all noisy best response protocols.

Example 12.3.13 showed that in the Stag Hunt game with linear payoffs, the noisy best response rules we considered always selected the equilibrium with the larger basin of attraction. The reason for this is easy to explain. Linearity of payoffs, along with the fact that the relative cost function $\tilde{\kappa}$ is sign-preserving and odd (see equation (12.32)), implies that the ordinal potential function I is symmetric about the mixed equilibrium χ^* , where it attains its minimum value. If, for example, χ^* is less than $\frac{1}{2}$, so that pure equilibrium 1 has the larger basin of attraction, then $I(1)$ exceeds $I(0)$, implying that state 1 is uniquely stochastically stable. Similarly, if χ^* exceeds $\frac{1}{2}$, then $I(0)$ exceeds $I(1)$, and state 0 is uniquely stochastically stable.

With this motivation, we call strategy i *strictly risk dominant* in the two-strategy coordination game F if the set of states where it is the unique best response is larger than the corresponding set for strategy $j \neq i$. Thus, if F has mixed equilibrium $\chi^* \in (0, 1)$, then strategy 0 is strictly risk dominant if $\chi^* > \frac{1}{2}$, and strategy 1 is strictly risk dominant if $\chi^* < \frac{1}{2}$. If the relevant inequality holds weakly in either case, we call the strategy in question *weakly risk dominant*.

The foregoing arguments yield the following result, in which we denote by e_i the state at which all agents play strategy i .

Corollary 12.3.16. *Assume the conditions of Theorem 12.3.6, and suppose that the limit game F is a coordination game with linear payoffs. Then*

- (i) *State e_i is weakly stochastically stable under every noisy best response protocol if and only if strategy i is weakly risk dominant in F .*
- (ii) *If strategy i is strictly risk dominant in F , then state e_i is uniquely stochastically stable under every noisy best response protocol.*

Example 12.3.14 shows that once we turn to games with nonlinear payoffs, risk dominance only characterizes stochastic stability under the BRM rule. In any coordination game with mixed equilibrium χ^* , the ordinal potential function for the BRM rule is $I_{\text{sgn}}(\chi) = |\chi - \chi^*| - \chi^*$. This function is minimized at χ^* , and increases at a unit rate as one moves away from χ^* in either direction, reflecting the fact that under the BRM rule, the probability of a suboptimal choice is independent of its payoff consequences. Clearly, whether $I_{\text{sgn}}(1)$ is greater than $I_{\text{sgn}}(0)$ depends only on whether χ^* is less than $\frac{1}{2}$. We therefore have

Corollary 12.3.17. *Assume the conditions of Theorem 12.3.6, and suppose that the limit game F is a coordination game and that σ^n is the BRM rule. Then*

- (i) *State e_i is weakly stochastically stable if and only if strategy i is weakly risk dominant in F .*
- (ii) *If strategy i is strictly risk dominant in F , then state e_i is uniquely stochastically stable.*

Once one moves beyond the BRM rule and linear payoffs, risk dominance is no longer a necessary or sufficient condition for stochastic stability. In what follows, we introduce a natural refinement of risk dominance that serves this role.

To work toward our new definition, let us first observe that any function on the unit interval $[0, 1]$ can be viewed as a random variable by regarding the interval as a sample space endowed with Lebesgue measure λ . With this interpretation in mind, we define the *advantage distribution* of strategy i to be the cumulative distribution function of the payoff advantage of strategy i over the alternative strategy $j \neq i$:

$$G_i(a) = \lambda(\{\chi \in [0, 1] : F_i(\chi) - F_j(\chi) \leq a\}).$$

We let \bar{G}_i denote the corresponding decumulative distribution function:

$$\bar{G}_i(a) = \lambda(\{\chi \in [0, 1] : F_i(\chi) - F_j(\chi) > a\}) = 1 - G_i(a).$$

In words, $\bar{G}_i(a)$ is the measure of the set of states at which the payoff to strategy i exceeds the payoff to strategy j by more than a .

It is easy to restate the definition of risk dominance in terms of the advantage distribution.

Observation 12.3.18. *Let F be a coordination game. Then strategy i is weakly risk dominant if and only if $\bar{G}_i(0) \geq \bar{G}_j(0)$, and strategy i is strictly risk dominant if and only if $\bar{G}_i(0) > \bar{G}_j(0)$.*

To obtain our refinement of risk dominance, we require not only that strategy i be optimal at a larger set of states than strategy j , but also that strategy i have a payoff advantage of at least a at a larger set of states than strategy j for every $a \geq 0$. More precisely, we say that strategy i is *weakly stochastically dominant* in the coordination game F if $\bar{G}_i(a) \geq \bar{G}_j(a)$ for all $a \geq 0$. If in addition $\bar{G}_i(0) > \bar{G}_j(0)$, we say that strategy i is *strictly stochastically dominant*. The notion of stochastic dominance for strategies proposed here is obtained by applying the usual definition of stochastic dominance from utility theory (see the Notes) to the strategies' advantage distributions.

Exercise 12.3.19. Consider the Nonlinear Stag Hunt game from Examples 12.3.14 and 11.2.6, with payoff functions $F_H(\chi) = h$ and $F_S(\chi) = s\chi^2$.

- (i) Plot \bar{G}_H and \bar{G}_S when $h = 2$ and $s = 5$. Is either strategy stochastically dominant?
- (ii) Plot \bar{G}_H and \bar{G}_S when $h = 2$ and $s = 8.5$. Is either strategy stochastically dominant? ‡

Theorem 12.3.20 shows that stochastic dominance is both sufficient and necessary to ensure stochastic stability under every noisy best response rule.

Theorem 12.3.20. *Assume the conditions of Theorem 12.3.6, and suppose that the limit game F is a coordination game. Then*

- (i) *State e_i is weakly stochastically stable under every noisy best response protocol if and only if strategy i is weakly stochastically dominant in F .*
- (ii) *If strategy i is strictly stochastically dominant in F , then state e_i is uniquely stochastically stable under every noisy best response protocol.*

The idea behind Theorem 12.3.20 is simple. The definitions of I , $\tilde{\kappa}$, κ , F_Δ , and G_i imply that

$$\begin{aligned}
 (12.39) \quad I(1) &= \int_0^1 \tilde{\kappa}(F_\Delta(y)) \, dy \\
 &= \int_0^1 \kappa(F_1(y) - F_0(y)) \, dy - \int_0^1 \kappa(F_0(y) - F_1(y)) \, dy \\
 &= \int_{-\infty}^{\infty} \kappa(a) \, dG_1(a) - \int_{-\infty}^{\infty} \kappa(a) \, dG_0(a).
 \end{aligned}$$

As we have seen, whether state e_1 or state e_0 is stochastically stable depends on whether $I(1)$ is greater than or less than $I(0) = 0$. This in turn depends on whether the value of the first integral in the final line of (12.39) exceeds the value of the second integral. Once we recall that the cost function κ is monotone, Theorem 12.3.20 reduces to a variation on the standard characterization of first-order stochastic dominance: namely, that distribution G_1 stochastically dominates distribution G_0 if and only if $\int \kappa \, dG_1 \geq \int \kappa \, dG_0$ for every nondecreasing function κ .

Proof of Theorem 12.3.20. Again view $[0, 1]$ as a sample space by endowing it with Lebesgue measure, and define $Y_i : [0, 1] \rightarrow \mathbf{R}_+$ by

$$Y_i(\omega) = \sup\{a : G_i(a) < \omega\}.$$

Then Y_i is a random variable with distribution G_i (see Example 10.B.4). It thus follows from equation (12.39) that

$$(12.40) \quad I(1) = \int_0^1 \kappa(Y_1(\omega)) \, d\omega - \int_0^1 \kappa(Y_0(\omega)) \, d\omega.$$

By construction we have that

$$\begin{aligned} Y_i(\omega) &< 0 \text{ when } \omega \in [0, G_i(0-)), \\ Y_i(\omega) &= 0 \text{ when } \omega \in [G_i(0-), G_i(0)], \text{ and} \\ Y_i(\omega) &> 0 \text{ when } \omega \in (G_i(0), 1], \end{aligned}$$

and that

$$G_1(0) - G_1(0-) = \lambda(\{x \in [0, 1] : F_1(x) = F_0(x)\}) = G_0(0) - G_0(0-).$$

Thus, since κ equals 0 on $(-\infty, 0)$, we can rewrite (12.40) as

$$\begin{aligned} (12.41) \quad I(1) &= \int_{G_1(0-)}^1 \kappa(Y_1(\omega)) \, d\omega - \int_{G_0(0-)}^1 \kappa(Y_0(\omega)) \, d\omega \\ &= \int_{G_1(0)}^1 \kappa(Y_1(\omega)) \, d\omega - \int_{G_0(0)}^1 \kappa(Y_0(\omega)) \, d\omega. \end{aligned}$$

To prove the “if” direction of part (i), suppose without loss of generality that strategy 1 is weakly stochastically dominant in F . Then $G_1(a) \leq G_0(a)$ for all $a \geq 0$, so the definition of Y_i implies that $Y_1(\omega) \geq Y_0(\omega)$ for all $\omega \in [G_1(0), 1]$. Since κ is nondecreasing and nonnegative, it follows from equation (12.41) that $I(1) \geq I(0)$, and hence that state 1 is weakly stochastically stable.

To prove part (ii), suppose without loss of generality that strategy 1 is strictly stochastically dominant in F . Then $G_1(a) \leq G_0(a)$ for all $a \geq 0$ and $G_1(0) < G_0(0)$. In this case, we not only have that $Y_1(\omega) \geq Y_0(\omega)$ for all $\omega \in [G_1(0), 1]$, but also that $Y_1(\omega) > 0$ when $\omega \in (G_1(0), G_0(0)]$. Since κ is nondecreasing, and since it is positive on $(0, \infty)$, it follows from equation (12.41) that

$$I(1) \geq \int_{G_1(0)}^{G_0(0)} \kappa(Y_1(\omega)) \, d\omega > 0 = I(0),$$

and hence that state 1 is uniquely stochastically stable.

Finally, to prove the “only if” direction of part (i), suppose without loss of generality that strategy 1 is not weakly stochastically dominant in F . Then $G_1(b) > G_0(b)$ for some

$b \geq 0$. Now consider a noisy best response protocol with cost function

$$\kappa(a) = \begin{cases} 0 & \text{if } a \leq 0, \\ 1 & \text{if } a \in (0, b], \\ k & \text{if } a > b, \end{cases}$$

where $k > G_1(b)/(G_1(b) - G_0(b))$. Then

$$\int_{-\infty}^{\infty} \kappa(a) dG_i(a) = (G_i(b) - G_i(0)) + k(1 - G_i(b)).$$

Therefore, equation (12.39) implies that

$$\begin{aligned} I(1) &= ((G_1(b) - G_1(0)) - (G_0(b) - G_0(0))) + k((1 - G_1(b)) - (1 - G_0(b))) \\ &\leq G_1(b) + k(G_0(b) - G_1(b)) \\ &< 0, \end{aligned}$$

implying that state e_1 is not weakly stochastically stable. This completes the proof of the theorem. ■

Exercise 12.3.21. Derive Corollary 12.3.16 from Theorem 12.3.20 by showing that in linear coordination games, risk dominance and stochastic dominance are equivalent. ‡

12.4 Imitative Protocols in Two-Strategy Games

We now shift our focus from noisy best response protocols to imitative protocols. To begin, we develop the basic insight from Example 12.1.9, showing that under imitative protocols with mutations, the double limits of the stationary distribution $\mu^{N,\varepsilon}$ in the population size N and mutation rate ε need not agree: while the large population double limit can select any state as stochastically stable, the small noise double limit can only select boundary states. We then alter our basic model by introducing one committed agent for each strategy. Doing so not only yields an irreducible evolutionary process even when mutations are absent, but also ensures that the two double limits agree when mutations are present. To conclude this section, we illustrate how the stochastically stable state can vary with the choice of imitative protocol, even across protocols generating the same mean dynamic, and even in games with linear payoffs.

12.4.1 Imitative Protocols with Mutations

We start by defining the class of protocols to be studied. We call the protocol

$$(12.42) \quad \rho_{ij}^\varepsilon(\pi, x) = x_j r_{ij}(\pi, x) + \varepsilon$$

a *positive imitative protocol with mutations* if the functions r_{01} and r_{10} that determine conditional imitation rates are continuous and positive-valued (and hence bounded away from zero on compact sets), and if the mutation rate ε is positive. If instead $\varepsilon = 0$, we call (12.42) a *positive imitative protocol without mutations*. As in Section 5.4, we will assume that net conditional imitation rates are monotone. This condition, which was defined for many-strategy games in equation (5.7), reduces in the two-strategy case to

$$(12.43) \quad \pi_j \geq \pi_i \Leftrightarrow r_{ij}(\pi, x) \geq r_{ji}(\pi, x) \text{ for } i, j \in \{0, 1\}.$$

To reduce the amount of notation in what follows, we will assume that the conditional switch rates $r_{ij}(\pi, x) = r_{ij}(\pi)$ do not depend directly on the population state; however, none of our results require this restriction.

We first suppose that there are no committed agents. Then in order for the process $\{\mathcal{X}_t^{N,\varepsilon}\}$ to be irreducible, the mutation rate ε must be positive. If this is so, the stationary distribution $\mu^{N,\varepsilon}$ satisfies

$$(12.44) \quad \begin{aligned} \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} &= \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N} \\ &= \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N}}{\frac{j}{N}} \cdot \frac{\frac{1}{R} \rho_{01}(F^N(\frac{j-1}{N}), \frac{j-1}{N-1})}{\frac{1}{R} \rho_{10}(F^N(\frac{j}{N}), \frac{j-1}{N-1})} \\ &= \prod_{j=1}^{N\chi} \frac{N-j+1}{j} \cdot \frac{\frac{j-1}{N-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon} \end{aligned}$$

for $\chi \in \{\frac{1}{N}, \frac{2}{N}, \dots, 1\}$. (We have assumed here that there is no self-imitation, although the alternative assumption leads to similar results—see Example 11.4.7.)

Theorem 12.4.1 describes the asymptotics of the stationary distribution in ε and N for both orders of limits. This result is stated in terms of the ordinal potential function $J: [0, 1] \rightarrow \mathbf{R}$, defined by

$$(12.45) \quad J(\chi) = \int_0^\chi \log \frac{r_{01}(F(y))}{r_{10}(F(y))} dy.$$

As with the ordinal potential I defined in equation (12.38), marginally increasing the mass on the optimal strategy here increases the value of J ; this follows directly from monotonicity condition (12.43).

Theorem 12.4.1 also employs one new definition: we say that the collection $\{\alpha^\varepsilon\}_{\varepsilon \in (0, \bar{\varepsilon}]}$ is of exact order ε as ε approaches zero, denoted $\alpha^\varepsilon \in \Theta(\varepsilon)$, if there is an interval $[a, b] \subset (0, \infty)$ such that $\alpha^\varepsilon / \varepsilon \in [a, b]$ for all ε close enough to zero.

Theorem 12.4.1. *Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of two-strategy population games that converges uniformly to the continuous-population game F . Suppose that agents employ a positive imitative protocol with mutations. Then the stationary distributions $\mu^{N, \varepsilon}$ satisfy*

- (i) $\lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{1}{N} \log \frac{\mu_1^{N, \varepsilon}}{\mu_0^{N, \varepsilon}} = J(1)$, and $\frac{\mu_\chi^{N, \varepsilon}}{\mu_0^{N, \varepsilon}}$ and $\frac{\mu_\chi^{N, \varepsilon}}{\mu_1^{N, \varepsilon}}$ are in $\Theta(\varepsilon)$ when $\chi \in \mathcal{X}^N - \{0, 1\}$, and
- (ii) $\lim_{\varepsilon \rightarrow 0} \lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_\chi^{N, \varepsilon} - \Delta J(\chi) \right| = 0$.

Theorem 12.4.1 formalizes and develops the intuitions introduced in Example 12.1.9. Part (i) shows that in the small noise double limit, all of the mass in the stationary distribution becomes concentrated on the boundary states. In fact, as long as $J(1) \neq 0$, it becomes concentrated on just one boundary state: state 1 if $J(1) > 0$, and state 0 if $J(1) < 0$. But part (ii) shows that in the large population double limit, a stochastically stable state must be one that maximizes J on the unit interval; no favoritism is shown toward boundary states in this case.

Exercise 12.4.2. To verify that Theorem 12.4.1(ii) agrees with the large population limit, derive the mean dynamic associated with protocol (12.42) when $\varepsilon = 0$, and show that the function J is a strict Lyapunov function for this dynamic on the open interval $(0, 1)$. ‡

Proof of Theorem 12.4.1. To prove the first statement in part (i), observe that

$$\begin{aligned} \frac{\mu_1^{N, \varepsilon}}{\mu_0^{N, \varepsilon}} &= \prod_{j=1}^N \frac{N-j+1}{j} \cdot \frac{\frac{j-1}{N-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon} \\ &= \frac{N\varepsilon}{r_{10}(F^N(\frac{1}{N})) + \varepsilon} \cdot \prod_{j=2}^{N-1} \frac{N-j+1}{j} \cdot \frac{\frac{j-1}{N-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon} \cdot \frac{r_{01}(F^N(\frac{N-1}{N})) + \varepsilon}{N\varepsilon} \\ &= \prod_{j=1}^{N-1} \frac{N-j}{j} \cdot \frac{\frac{j}{N-1} r_{01}(F^N(\frac{j}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon}. \end{aligned}$$

Since we have eliminated the troublesome $N\varepsilon$ terms, none of the terms in the final product

converge to zero or infinity. Indeed, we have

$$\lim_{\varepsilon \rightarrow 0} \log \frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \log \left(\prod_{j=1}^{N-1} \frac{N-j}{j} \cdot \frac{\frac{j}{N-1} r_{01}(F^N(\frac{j}{N}))}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N}))} \right) = \sum_{j=1}^{N-1} \log \frac{r_{01}(F^N(\frac{j}{N}))}{r_{10}(F^N(\frac{j}{N}))},$$

so the Dominated Convergence Theorem implies that

$$\lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{1}{N} \log \frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \lim_{N \rightarrow \infty} \frac{N-1}{N} \cdot \frac{1}{N-1} \sum_{j=1}^{N-1} \log \frac{r_{01}(F^N(\frac{j}{N}))}{r_{10}(F^N(\frac{j}{N}))} = J(1).$$

To prove the second statement in part (i), observe that for $\chi \in \mathcal{X}^N - \{0, 1\}$, we have

$$\begin{aligned} \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} &= \prod_{j=1}^{N\chi} \frac{N-j+1}{j} \cdot \frac{\frac{j-1}{N-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon} \\ &= \frac{N\varepsilon}{r_{10}(F^N(\frac{1}{N}), 0) + \varepsilon} \cdot \prod_{j=2}^{N\chi} \frac{N-j+1}{j} \cdot \frac{\frac{j-1}{N-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon}. \end{aligned}$$

Since all terms except the initial $N\varepsilon$ approach positive constants as ε approaches zero, the claim follows. The analysis of $\mu_\chi^{N,\varepsilon} / \mu_1^{N,\varepsilon}$ is similar.

We now prove part (ii). Equation (12.44) tells us that

$$\frac{1}{N} \log \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \frac{1}{N} \sum_{j=1}^{N\chi} \left(\log \frac{N-j+1}{j} + \log \frac{\frac{j-1}{N-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{\frac{N-j}{N-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon} \right).$$

Thus, the Dominated Convergence Theorem implies that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \int_0^\chi \left(\log \frac{1-y}{y} + \log \frac{y r_{01}(F(y)) + \varepsilon}{(1-y)r_{10}(F(y)) + \varepsilon} \right) dy.$$

(For the first integrand, see the proof of Theorem 12.3.6.) A second application of the Dominated Convergence Theorem then yields

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \lim_{N \rightarrow \infty} \frac{1}{N} \log \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} &= \int_0^\chi \left(\log \frac{1-y}{y} + \log \frac{y r_{01}(F(y))}{(1-y)r_{10}(F(y))} \right) dy \\ &= \int_0^\chi \log \frac{r_{01}(F(y))}{r_{10}(F(y))} dy \\ &= J(\chi). \end{aligned}$$

Since all limits are uniform in χ , an argument similar to the proof of Theorem 12.2.2 completes the proof of part (ii). ■

This analysis confirms that in the absence of committed players, the small noise and large population double limits yield different predictions. As we noted after Example 12.1.9, we are inclined to favor the prediction of the large population double limit in most economic applications, but believe that either double limit could be useful in biological applications.

Before proceeding, we should note that in games with $n > 2$ strategies the small noise limit is far easier to evaluate than the large population limit when agents use imitative protocols with mutations. If we fix the population size N and make ε small enough, the process $\{X_t^{N,\varepsilon}\}$ will spend nearly all periods at vertices of the simplex X ; moreover, sojourns between vertices will nearly always travel along the edges of X , and on each of these $\binom{n}{2}$ edges, the process $\{X_t^{N,\varepsilon}\}$ is essentially a birth and death process. By developing this observation, one can analyze the small noise limit of $\{X_t^{N,\varepsilon}\}$ by studying an auxiliary Markov chain with just n states, with transition probabilities being governed by the behavior of the $\binom{n}{2}$ birth and death processes. See the Notes for references that carry out this analysis.

12.4.2 Imitative Protocols with Committed Agents

We now suppose that there is one committed agent playing each strategy. This assumption ensures that the stochastic evolutionary process is irreducible, even in the absence of mutations. Theorem 12.4.3 shows that the large population asymptotics of the stationary distribution take a very simple form.

Theorem 12.4.3. *Let $\{F^N\}_{N=N_0}^\infty$ be a sequence of two-strategy population games that converges uniformly to the continuous-population game F . Suppose that the N standard agents employ a positive imitative protocol without mutations, and that there is one committed agent for each strategy. Then the sequence of stationary distributions $\{\mu^N\}_{N=N_0}^\infty$ satisfies*

$$(12.46) \quad \lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_\chi^N - \Delta J(\chi) \right| = 0.$$

Proof. As we saw in Example 11.4.7, the stationary distribution μ^N satisfies

$$(12.47) \quad \frac{\mu_\chi^N}{\mu_0^N} = \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N}$$

$$\begin{aligned}
&= \prod_{j=1}^{N\chi} \frac{N-j+1}{N} \cdot \frac{\frac{1}{R} \rho_{01}(F^N(\frac{j-1}{N}), \frac{j}{N+1})}{\frac{1}{R} \rho_{10}(F^N(\frac{j}{N}), \frac{j}{N+1})} \\
&= \prod_{j=1}^{N\chi} \frac{N-j+1}{N} \cdot \frac{\frac{j}{N+1}}{\frac{N-j+1}{N+1}} \cdot \frac{r_{01}(F^N(\frac{j-1}{N}))}{r_{10}(F^N(\frac{j}{N}))} \\
&= \prod_{j=1}^{N\chi} \frac{r_{01}(F^N(\frac{j-1}{N}))}{r_{10}(F^N(\frac{j}{N}))}.
\end{aligned}$$

Since r_{01} and r_{10} are both bounded and bounded away from zero, the Dominated Convergence Theorem implies that

$$(12.48) \quad \lim_{N \rightarrow \infty} \frac{1}{N} \log \frac{\mu_{\chi}^N}{\mu_0^N} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{r_{01}(F^N(\frac{j-1}{N}))}{r_{10}(F^N(\frac{j}{N}))} = J(\chi),$$

and that convergence is uniform in χ . The remainder of the proof is similar to that of Theorem 12.2.2. ■

The following theorem, whose proof is left as an exercise, shows that when committed agents are present, introducing rare mutations has no significant impact on the conclusions of Theorem 12.4.3, regardless of the order in which the limits in ε and N are taken, making the choice between the orders of limits inconsequential.

Theorem 12.4.4. *Let $\{F^N\}_{N=N_0}^{\infty}$ be a sequence of two-strategy population games that converges uniformly to F . Suppose that the N standard agents employ a positive imitative protocol with mutations, and that there is one committed agent for each strategy. Then the stationary distributions $\mu^{N,\varepsilon}$ satisfy*

- (i) $\lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_{\chi}^{N,\varepsilon} - \Delta J(\chi) \right| = 0$ and
- (ii) $\lim_{\varepsilon \rightarrow 0} \lim_{N \rightarrow \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_{\chi}^{N,\varepsilon} - \Delta J(\chi) \right| = 0.$

Exercise 12.4.5. Prove Theorem 12.4.4.

Exercise 12.4.6. Show that Theorems 12.4.3 and 12.4.4 remain true if there are $k_0 \geq 0$ committed agents playing strategy 0 and $k_1 \geq 1$ committed agents playing strategy 1. ‡

12.4.3 Imitative Protocols, Mean Dynamics, and Stochastic Stability

Once committed agents are introduced, the asymptotics of the stationary distribution are the same for either order of limits, and agree with those of the large population limit

alone when mutations are absent. Thus, for a given game and imitative protocol, the stochastically stable states are defined unambiguously. For simplicity, let us focus on the large population limit when mutations are absent. In this case, stochastic stability is defined as in equation (12.5), and Proposition 12.1.5 and Theorem 12.4.3 imply that a unique maximizer of the ordinal potential function J is uniquely stochastically stable.

In Section 12.3, we saw that by virtue of their different levels of sensitivity to payoff disadvantages, different noisy best response protocols can generate different stochastically stable states. Before noise levels reach their limit, dissimilarities among the protocols also manifest themselves in small differences among their mean dynamics. One might therefore expect that variation among protocols' stochastically stable states can always be linked back to variation among their mean dynamics.

We now argue under imitative protocols, mean dynamics alone are not enough to identify the stochastically stable state. To do so, we consider two imitative protocols that generate the replicator dynamic as their mean dynamics, and show that in a range of linear games, the two protocols generate different stochastically stable states.

Example 12.4.7. Stochastic imitation in Boar Hunt. We call the symmetric normal form game

$$A = \begin{pmatrix} h & h \\ b & B \end{pmatrix}$$

with $b < h < B$ *Boar Hunt*. The strategies in this coordination game are hunting for hare (strategy 0) and hunting for boar (strategy 1). As in the Stag Hunt game (Examples 11.2.5 and 12.3.13), hunting for hare yields a certain payoff of h , while coordinating on hunting for the larger prey yields the highest possible payoff. But unlike in Stag Hunt, either strategy in Boar Hunt can be risk dominant. If a population of agents are matched to play Boar Hunt, the resulting population game has payoffs $F_0(\chi) = h$ and $F_1(\chi) = b(1 - \chi) + B\chi$. In addition to the all-Hare ($\chi = 0$) and all-Boar ($\chi = 1$) equilibria, there is a mixed equilibrium in which fraction $\chi^* = \frac{h-b}{B-b}$ of the population hunts for Boar.

Since Boar Hunt is a coordination game, the ordinal potential J induced by any imitative protocol satisfying monotonicity condition (12.43) must be maximized at state 0, state 1, or both. Theorem 12.4.3 then implies that the all-Hare equilibrium is uniquely stochastically stable if $J(1) < J(0) \equiv 0$, while the all-Boar equilibrium is uniquely stochastically stable if $J(1) > 0$.

Now suppose that that agents employ an imitative protocol whose conditional imitation rates take one of the following two forms:

$$(12.49) \quad r_{ij}(\pi) = M - \pi_i, \text{ or}$$

$$(12.50) \quad r_{ij}(\pi) = \pi_j - m,$$

where $M > B$ and $m < b$. These protocols are linear specifications of imitation driven by dissatisfaction (Example 5.4.4) and imitation of success (Example 5.4.5), respectively. As we saw in these earlier examples, both protocols generate the replicator dynamic as their mean dynamics regardless of the game at hand.

Because both the population game F and the revision protocols (12.49) and (12.50) are linear, the corresponding ordinal potential functions J_D and J_S are easy to evaluate. Evaluating these functions at state 1, we find that

$$J_D(1) = \int_0^1 \log \frac{M - F_0(y)}{M - F_1(y)} dy = \log \frac{M - h}{M - B} + \frac{M - b}{B - b} \log \frac{M - B}{M - b} + 1 \text{ and}$$

$$J_S(1) = \int_0^1 \log \frac{F_1(y) - m}{F_0(y) - m} dy = \log \frac{B - m}{h - m} + \frac{b - m}{B - b} \log \frac{B - m}{b - m} + 1.$$

To make this example more concrete, let us fix $m = 0$, $b = 1$, $h = 2$, $B = 3$, and $M = 4$. Then the mixed Nash equilibrium of F is $\chi^* = \frac{h-b}{B-b} = \frac{1}{2}$, implying that neither strategy is strictly risk dominant; obviously, the all-Boar equilibrium $\chi = 1$ is payoff dominant. The normalized functions ΔJ_D and ΔJ_S for this case are graphed in Figure 12.4.1. Since $J_D(1) = -\frac{3}{2} \log 3 + \log 2 + 1 \approx .0452$, imitation driven by dissatisfaction (12.49) leads to the selection of the payoff dominant all-Boar equilibrium, $\chi = 1$. But since and $J_S(1) = -J_D(1) \approx -.0452$, imitation of success (12.50) leads to the selection of the payoff dominated all-Hare equilibrium, $\chi = 0$. Clearly, the same qualitative results obtain in any game with similar payoff values, implying that either protocol can select or fail to select a strictly risk dominant equilibrium. To sum up, protocols (12.49) and (12.50) both generate the replicator dynamic as their mean dynamics, but they still can support different stochastically stable states in simple games.

This discrepancy can be traced to the different manners in which the conditional imitation rates r_{ij} influence the evolutionary process's mean dynamic and stationary distribution. Since the mean dynamic describes the expected increments of the process, its values depend on differences between the conditional imitation rates r_{01} and r_{10} :

$$\dot{\chi} = (1 - \chi)\chi r_{01}(F(\chi)) - \chi(1 - \chi) r_{10}(F(\chi)) = \chi(1 - \chi)(r_{01}(F(\chi)) - r_{10}(F(\chi))).$$

For its part, the stationary distribution is the limiting distribution of the stochastic process, and so is determined by the products of the process's one-step transition probabilities. Because of this, the stationary distribution weights depend on the ratios of r_{01} and r_{10} ,




Figure 12.4.1: The ordinal potentials ΔJ_D (green) and ΔJ_S (red) when $m = 0$, $b = 1$, $h = 2$, $B = 3$, and $M = 4$.

or, equivalently, on the differences between the logarithms of r_{01} and r_{10} : see equations (12.47) and (12.48) above.

We can use this last observation to explain the discrepant equilibrium selection results. To begin, write

$$\begin{aligned}
 J(1) &= \int_0^1 \log \frac{r_{01}(F(y))}{r_{10}(F(y))} dy \\
 (12.51) \quad &= \int_0^1 \log r_{01}(F(y)) dy - \int_0^1 \log r_{10}(F(y)) dy.
 \end{aligned}$$

The all-Boar equilibrium is stochastically stable if $J(1) > J(0) = 0$, which is true when the first integral in (12.51) is larger than the second. Similarly, the all-Hare equilibrium is stochastically stable if the second integral in (12.51) is larger than the first.

Observe that as χ varies from 0 to 1, the payoff to Hare is constant at $F_0(\chi) = 2$, while the payoff to Boar, $F_1(x) = 2\chi + 1$, is uniformly distributed between 1 and 3. As we explain next, the identity of the stochastically stable state is driven by which direction of imitation, Hare to Boar or Boar to Hare, is determined by the constant payoff.

We consider each protocol in turn. Under imitation driven by dissatisfaction (12.49), the conditional imitation rates are given by $r_{01}(\chi) = M - F_0(\chi) = 2$ and $r_{10}(\chi) = 4 - F_1(\chi)$. If we view $F_1(\chi)$ as a random variable with a *uniform*[1, 3] distribution, then $r_{10}(\chi)$ has a *uniform*[1, 3] distribution as well. Thus, since the logarithm function is concave, Jensen's inequality implies that the first integral in (12.51) is larger than the second, and hence that the all-Boar equilibrium is stochastically stable. Intuitively, the conditional imitation probabilities in each direction have the same mean, but only those from Boar to Hare are

variable. The Boar to Hare transition is therefore less likely than its opposite, and so the all-Boar equilibrium is selected.

Under imitation of success (12.50), the conditional imitation rates are given by $r_{01}(\chi) = F_1(\chi) - m = F_1(\chi)$ and $r_{10}(\chi) = F_0(\chi) - m = 2$. This time, the conditional imitation probabilities from Boar to Hare are fixed, while those from Hare to Boar are variable; the latter transition is therefore less likely, and the all-Hare equilibrium is stochastically stable. §

In order to obtain robust equilibrium selection results under imitative protocols, one must place additional structure on the manner in which conditional imitation rates depend on the payoffs of an agent's current and prospective strategies. See the Notes for references to work in this direction.

12.5 Small Noise Limits

To this point, our analyses of infinite horizon behavior have focused on settings that generate reversible processes, so that stationary distributions can be expressed as simple formulas. Beyond these settings, working directly from exact expressions for the stationary distribution becomes a much less promising way of evaluating stochastic stability. Fortunately, there are techniques that allow us to determine the limiting stationary distributions directly, without the intermediate step of computing stationary distributions for fixed parameter values.

In this section, we introduce general techniques for evaluating stochastic stability in the small noise limit, based on the construction of certain graphs on the state space \mathcal{X}^N . We then use these techniques to obtain equilibrium selection results. A full development of the methods applied here can be found in Appendix 12.A. The methods used to study large population limits are drawn from stochastic approximation theory; these are treated in Section 12.6 and Appendix 12.B.

12.5.1 Noisy Best Response Protocols and Cost Functions

We consider a fixed population of N agents who play an n strategy population game, employing a noisy best response protocol $\sigma^\varepsilon : \mathbf{R}^n \rightarrow \mathbf{R}_+^{n \times n}$. We assume throughout that $\sum_{j \in S} \sigma_{ij}^\varepsilon(\pi) = 1$ for all $i \in S$ and $\pi \in \mathbf{R}^n$, and that $\sigma_{ij}^\varepsilon(\pi) > 0$ whenever $\varepsilon > 0$.

As in Section 12.3.1, we impose restrictions on the behavior of the choice probabilities $\sigma_{ij}^\varepsilon(\pi)$ as ε approaches zero, expressing these restrictions in terms of costs. Here we follow most of the literature in using a definition of costs, (12.52), that is somewhat more

restrictive than the one introduced in Section 12.3.1. The restrictions we place on the values of costs, (12.53) and (12.54), are rather weak.

In detail, we suppose that for every ordered pair (i, j) of distinct strategies and every payoff vector $\pi \in \mathbf{R}^n$, there are constants $a_{ij}(\pi) > 0$ and $\kappa_{ij}(\pi) \geq 0$ such that

$$(12.52) \quad \sigma_{ij}^\varepsilon(\pi) = (a_{ij}(\pi) + o(1)) \varepsilon^{\kappa_{ij}(\pi)}, \text{ where}$$

$$(12.53) \quad \kappa_{ij}(\pi) = 0 \text{ when } \{j\} = \operatorname{argmax}_{k \in S} \pi_k, \text{ and}$$

$$(12.54) \quad \kappa_{ij}(\pi) > 0 \text{ when } j \notin \operatorname{argmax}_{k \in S} \pi_k.$$

We call $\kappa : \mathbf{R}^n \rightarrow \mathbf{R}_+^{n \times n}$ the *cost function* associated with protocol σ^ε .

The limit used in (12.52) to define costs $\kappa_{ij}(\pi)$ generates a more demanding requirement than the the limit (12.27) used in Section 12.3.1. To see this, apply the change of parameter $\varepsilon = \exp(-\eta^{-1})$ to rewrite (12.52) as

$$(12.55) \quad \sigma_{ij}^\eta(\pi) = (a_{ij}(\pi) + o(1)) \exp(-\eta^{-1} \kappa_{ij}(\pi)).$$

It follows from (12.55) that

$$(12.56) \quad -\lim_{\eta \rightarrow 0} \eta \log \sigma_{ij}^\eta(\pi) = \lim_{\eta \rightarrow 0} (\kappa_{ij}(\pi) - \eta \log(a_{ij}(\pi) + o(1))) = \kappa_{ij}(\pi),$$

in agreement with condition (12.27). We discuss the differences between these definitions after the examples below and in Appendix 12.A.5.

Conditions (12.53) and (12.54) are minimal requirements for (12.52) to be considered a noisy best response protocol: (12.53) says that switching to the unique optimal strategy has zero cost, while (12.54) says that switching to any suboptimal strategy has a positive cost. (As in Section 12.3.1, we allow some flexibility in specifying the cost of switching to an optimal strategy when more than one exists: compare Examples 12.5.1 and 12.5.2 below.) There are certainly other requirements that are natural to impose on κ (for instance, monotonicity and symmetry), but we will not concern ourselves with these conditions here. As with those of revision protocols, the diagonal elements of the cost function κ are just place holders, so changing their values has no effect on the analyses to come.

We now derive the cost functions for the BRM and logit protocols. Compared to the analyses in Section 12.3.1, the ones here use the more restrictive definition of costs from equation (12.52), and allow for games with more than two strategies. We write $S^*(\pi) = \operatorname{argmax}_{k \in S} \pi_k$ for the set of optimal strategies at π , $n^*(\pi) = \#S^*(\pi)$ for the number of such strategies, and $k^*(\pi) \in S^*(\pi)$ for an arbitrary optimal strategy, and we omit the argument π from these notations when doing so will not cause confusion.

Example 12.5.1. Best response with mutations. Under the BRM(ε) protocol (Example 11.1.1), a revising agent switches to a current best response with probability $1 - \varepsilon$, and chooses a strategy at random with probability ε . There are various ways to specify what happens in the event of a tie; here we assume that if a mutation does not occur, an agent playing an optimal strategy does not switch, while an agent not playing an optimal strategy chooses at random among such strategies. In sum, we have

$$\sigma_{ij}^\varepsilon(\pi) = \begin{cases} (1 - \varepsilon) + \frac{1}{n}\varepsilon & \text{if } \pi_i = \pi_{k^*} \text{ and } j = i, \\ \frac{1}{n^*}(1 - \varepsilon) + \frac{1}{n}\varepsilon & \text{if } \pi_i < \pi_j = \pi_{k^*}, \\ \frac{1}{n}\varepsilon & \text{if } \max\{\pi_i, \pi_j\} < \pi_{k^*}, \text{ or if } \pi_i = \pi_{k^*} \text{ and } j \neq i. \end{cases}$$

If we consider only its off-diagonal elements, the cost function for the BRM protocol takes a simple form:

$$\text{when } i \neq j, \kappa_{ij}(\pi) = \begin{cases} 0 & \text{if } \pi_i < \pi_j = \max_{k \in S} \pi_k, \\ 1 & \text{otherwise.} \end{cases} \quad \S$$

Example 12.5.2. Logit choice. The logit(η) choice protocol is

$$\sigma_{ij}^\eta(\pi) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}.$$

Performing the change of variable $\eta = -(\log \varepsilon)^{-1}$ and rearranging, we obtain

$$\begin{aligned} \sigma_{ij}^\varepsilon(\pi) &= \frac{\varepsilon^{-\pi_j}}{\sum_{k \in S} \varepsilon^{-\pi_k}} \\ &= \frac{\varepsilon^{\pi_{k^*} - \pi_j}}{n^* + \sum_{k \notin S^*} \varepsilon^{\pi_{k^*} - \pi_k}} \\ &= \frac{1}{n^*} \left(1 - \frac{\sum_{k \notin S^*} \varepsilon^{\pi_{k^*} - \pi_k}}{n^* + \sum_{k \notin S^*} \varepsilon^{\pi_{k^*} - \pi_k}} \right) \varepsilon^{\pi_{k^*} - \pi_j}. \end{aligned}$$

Therefore, the cost function for the logit rule is

$$\kappa_{ij}(\pi) = \max_{k \in S} \pi_k - \pi_j. \quad \S$$

We close with more comments about the two different limits used to define costs: condition (12.27) from Section 12.3.1, and condition (12.52) above. Since (12.52) is more demanding than (12.27), it allows one to prove slightly stronger results in cases where

both apply: (12.52) leads to results about stochastic stability, while (12.27) leads to results about weak stochastic stability. Of course, because definition (12.52) is more demanding, it rules out certain protocols that definition (12.27) allows—most notably, the probit choice protocol, which presents some novel properties in the n -strategy case (see Example 12.A.8 and the Notes). But one can perform versions of the analyses in this section while defining costs using (12.27); see Appendix 12.A.5 for details.

12.5.2 Limiting Stationary Distributions via Trees

Let us review our present model of stochastic evolution. A population of N agents repeatedly plays a population game $F^N : \mathcal{X}^N \rightarrow \mathbf{R}^n$. Each agent receives revision opportunities according to a rate 1 Poisson process. When such an opportunity arises, an agent decides what to do next by employing a noisy best response protocol $\sigma^\varepsilon : \mathbf{R}^n \rightarrow \mathbf{R}_+^{n \times n}$. We assume that agents are clever, in that they account for the effect of their own choices on the population state when assessing the payoffs of other strategies. Thus, if the population state is x , and an i player considers switching to strategy j , he evaluates the payoffs of the latter strategy as $\check{F}_j^N(x - \frac{1}{N}e_i) = F_j^N(x + \frac{1}{N}(e_j - e_i))$ (see Section 11.4.2). Assuming that agents are clever leads to cleaner selection results, but is not essential to the analysis.

Under these assumptions, the stochastic evolutionary process $\{X_t^\varepsilon\}$ has transition probabilities of the form

$$P_{xy}^\varepsilon = \begin{cases} x_i \sigma_{ij}^\varepsilon(\check{F}(x - \frac{1}{N}e_i)) & \text{if } y = x + \frac{1}{N}(e_j - e_i), j \neq i, \\ \sum_{i \in S} x_i \sigma_{ii}^\varepsilon(\check{F}(x - \frac{1}{N}e_i)) & \text{if } y = x, \\ 0 & \text{otherwise.} \end{cases}$$

We denote the unique stationary distribution of this process by μ^ε .

We now introduce a notion of cost for one-step transitions of the process $\{X_t^\varepsilon\}$. For distinct states $x, y \in \mathcal{X}^N$, we define the *cost* c_{xy} of the transition from x to y to be the rate of decay of the the transition probability P_{xy}^ε as ε approaches zero, with the cost defined to be infinite if the transition is impossible. Evidently, transition costs for the present model take the form

$$c_{xy} = \begin{cases} \kappa_{ij}(\check{F}(z)) & \text{if } x = z + \frac{1}{N}e_i \text{ and } y = z + \frac{1}{N}e_j \neq x, \\ \infty & \text{otherwise.} \end{cases}$$

These transition costs provide a simple way of computing the limiting stationary distribution $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$. Here we present the key result as concisely as possible;

complete details can be found in Appendices 12.A.1 and 12.A.2.

To begin, we introduce two types of directed graphs on \mathcal{X}^N . First, a *path* from $x \in \mathcal{X}^N$ to $y \neq x$ is a directed graph $\{(x, x_1), (x_1, x_2), \dots, (x_{l-1}, y)\}$ on \mathcal{X}^N whose edges lead from x to y without hitting any node more than once. Second, a *tree* with *root* $x \in \mathcal{X}^N$, also called an *x-tree*, is a directed graph on \mathcal{X}^N with no outgoing edges from x , exactly one outgoing edge from each $y \neq x$, and a path, necessarily unique, from each $y \neq x$ to x . We let T_x denote the set of x -trees on \mathcal{X}^N .

The characterization of the limiting stationary distribution, Theorem 12.5.3, requires us to define costs of trees. Given a directed graph γ on \mathcal{X}^N , we define its *cost* $C(\gamma)$ to be the sum of the costs of the edges it contains:

$$(12.57) \quad C(\gamma) = \sum_{(y,z) \in \gamma} c_{yz}.$$

We define C_x^* to be the lowest cost of any x -tree, and C^* to be the lowest cost of any tree on \mathcal{X}^N .

$$C_x^* = \min_{\tau_x \in T_x} C(\tau_x), \quad \text{and} \quad C^* = \min_{x \in \mathcal{X}} C_x^*.$$

We then have the following result, which follows from Theorem 12.A.2 and Corollary 12.A.3 in Appendix 12.A.2.

Theorem 12.5.3. *Consider the collection of processes $\{X_t^\varepsilon\}_{t \geq 0}$ for $\varepsilon \in [0, \varepsilon]$ on the state space \mathcal{X}^N , and let $\{\mu^\varepsilon\}_{\varepsilon \in [0, \varepsilon]}$ be their stationary distributions. Then*

- (i) $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$ is a stationary distribution of $\{X_t^0\}$ with support $\{x \in \mathcal{X}^N : C_x^* = C^*\}$, and
- (ii) for each $x \in \mathcal{X}^N$, there is a constant $b_x > 0$ such that $\mu_x^\varepsilon = b_x \varepsilon^{C_x^* - C^*} + o(\varepsilon^{C_x^* - C^*})$.

Part (i) of the theorem shows that the stochastically stable states are precisely those that admit minimal cost trees. Part (ii) of the theorem reveals that the stationary distribution weight on each other state decays at a rate determined by its “cost disadvantage” relative to the stochastically stable states.

According to Theorem 12.5.3, one can show that state x is uniquely stochastically stable by proving that for any other state $y \neq x$ and any y -tree τ_y , there is an x -tree whose cost is less than $C(\tau_y)$. In fact, since μ^* is a stationary distribution of the zero-noise process, only the recurrent states of this process are candidates for stochastic stability, making it sufficient to consider trees rooted at recurrent states.

In some settings this direct approach to assessing stochastic stability works perfectly well; we demonstrate this next. In other settings, it is preferable to employ simpler criteria



Figure 12.5.1: The lone χ -tree whose cost is finite.

for stochastic stability that can be deduced from Theorem 12.5.3, but that do not require the explicit construction of x -trees. We develop this approach in Sections 12.5.4 and 12.5.5.

12.5.3 Two-Strategy Games and Risk Dominance

We now illustrate the use of Theorem 12.5.3 in the simplest possible setting: that of two-strategy games.

Let $F : \mathcal{X}^N \rightarrow \mathbf{R}^2$ be an N -agent, two-strategy population game, and as usual, identify state x with the weight $\chi \equiv x_1$ that it places on strategy 1. Because the state space $\mathcal{X}^N = \{0, \frac{1}{N}, \dots, 1\}$ is linearly ordered, and since only transitions to adjacent states have positive probability, each state $\chi \in \mathcal{X}^N$ admits a unique tree with finite cost, namely

$$\tau_\chi = \{(0, \frac{1}{N}), \dots, (\chi - \frac{1}{N}, \chi), (\chi + \frac{1}{N}, \chi), \dots, (1, \frac{N-1}{N})\}.$$


(See Figure 12.5.1.) Therefore, state χ is stochastically stable if and only if $C_\chi^* = C(\tau_\chi)$ is minimal, where

$$\begin{aligned} C(\tau_\chi) &= \sum_{k=0}^{N\chi-1} c_{k/N, (k+1)/N} + \sum_{k=N\chi+1}^N c_{k/N, (k-1)/N} \\ &= \sum_{k=0}^{N\chi-1} \kappa_{01}(F_0(\frac{k}{N}), F_1(\frac{k+1}{N})) + \sum_{k=N\chi+1}^N \kappa_{10}(F_0(\frac{k-1}{N}), F_1(\frac{k}{N})). \end{aligned}$$


Example 12.5.4. More Stag Hunt. Consider the Stag Hunt game from Examples 11.2.5 and 12.3.13:

$$A = \begin{pmatrix} h & h \\ 0 & s \end{pmatrix}.$$

This game has two pure Nash equilibria and one mixed Nash equilibrium. We suppose that $h = 4$ and $s = 7$, so that the mixed equilibrium puts weight $\chi^* = \frac{4}{7}$ on Stag. If N agents are matched without self-matching to play this game, the expected payoffs to Hare and Stag are $F_H^N(\chi) = 4$ and $F_S^N(\chi) = \frac{N\chi-1}{N-1} \cdot 7$.



(i) τ_0 and its edge costs.



(ii) τ_1 and its edge costs.

Figure 12.5.2: Trees and BRM edge costs in Stag Hunt ($h = 4, s = 7, N = 15$).


Suppose that $N = 15$. Then since $F_H^N(\frac{8}{15}) = 4 = F_S^N(\frac{9}{15})$, a clever agent with exactly 8 opponents playing Stag is indifferent; other numbers of opponents playing Stag lead to a strict preference for one strategy or the other. Notice that the mixed Nash equilibrium of A lies between states $\frac{8}{15}$ and $\frac{9}{15}$.

We prove below that under any noisy best response rule, states 0 and 1 are the only recurrent states, and hence the only candidates for stochastic stability. Thus, by Theorem 12.5.3, our stochastic stability analysis reduces to a comparison of $C(\tau_0)$ and $C(\tau_1)$.


Figure 12.5.2 presents trees τ_0 and τ_1 , and labels the edges in these trees with their costs under the BRM rule. Each edge in τ_0 points toward the state below it, and so represents a switch by one agent from strategy 1 to strategy 0. Edges starting at states $\chi \geq \frac{9}{15}$ require switching away from an optimal strategy, and so have cost 1; the remaining states do not, and have cost 0. (It is because of our assumption that indifferent agents do not switch without mutating that edge $(\frac{9}{15}, \frac{8}{15})$ has cost 1.) It follows that $C(\tau_0) = \#\{\frac{9}{15}, \dots, 1\} = 7$. By similar logic, the cost of τ_1 under the BRM rule is $C(\tau_1) = \#\{0, \dots, \frac{8}{15}\} = 9$. Since $C(\tau_0) = 7 < 9 = C(\tau_1)$, state 0 is stochastically stable under the BRM rule.

Figure 12.5.3 labels the edges in τ_0 and τ_1 with their costs under the logit rule. The cost of an edge representing a switch to an optimal strategy still has cost zero, but the cost of an edge representing a switch to a suboptimal strategy now equals that strategy's payoff disadvantage. Summing the edge costs in each tree, we find that $C(\tau_0) = 10\frac{1}{2} < 18 = C(\tau_1)$. Thus, state 0 is also stochastically stable under the logit rule. §

In Section 12.3.6, we introduced the notions of weak and strict risk dominance for infinite-population games. We then showed when stochastic stability is defined in terms



(i) τ_0 and its edge costs.



(ii) τ_1 and its edge costs.

Figure 12.5.3: Trees and logit edge costs in Stag Hunt ($h = 4, s = 7, N = 15$).

of double limits, strict risk dominance implies unique stochastic stability if payoffs are linear (Corollary 12.3.16), or if agents use the BRM protocol (Corollary 12.3.17). We now establish versions of these results for the finite-population setting, where only the small noise limit is taken. Because we are working with a discrete state space, we will need to use slightly stronger definitions and evaluate a number of slightly different cases in order to obtain exact results.

We begin by introducing three definitions of risk dominance for the two-strategy game F^N . Let $y \in \mathcal{X}^{N-1}$ represent the proportion of an agent's opponents who play strategy 1. We call strategy 1 *weakly risk dominant* if

$$(12.58) \quad F_1^N \left(\frac{N-1}{N} y + \frac{1}{N} \right) \geq F_0^N \left(\frac{N-1}{N} y \right) \text{ for all } y \in \mathcal{X}^{N-1} \text{ with } y \geq \frac{1}{2},$$

We call strategy 1 *strictly risk dominant* if

$$(12.59) \quad F_1^N \left(\frac{N-1}{N} y + \frac{1}{N} \right) > F_0^N \left(\frac{N-1}{N} y \right) \text{ for all } y \in \mathcal{X}^{N-1} \text{ with } y \geq \frac{1}{2}.$$

And we call strategy 1 *strongly risk dominant* if

$$(12.60) \quad F_1^N \left(\frac{N-1}{N} y + \frac{1}{N} \right) \geq F_0^N \left(\frac{N-1}{N} y \right) \text{ for all } y \in \mathcal{X}^{N-1} \text{ with } y \geq \frac{1}{2} \frac{N-2}{N-1}.$$

Exercise 12.5.5. Define weak, strict, and strong risk dominance for strategy 0 in a manner that agrees with the definitions above. ‡

The first two definitions correspond to the definitions of weak and strict risk dominance from Section 12.3.6. The main novelties here are that the appropriate state variable is the distribution of opponent's behavior, and that the effect of the agent's own choice on the population state is taken into account in the payoff comparison employed. As we will see below, the more demanding notion of strong risk dominance is precisely what is needed to obtain an exact selection result under the BRM protocol.

Because conditions (12.58)–(12.60) only consider points y in \mathcal{X}^{N-1} (rather than in all of $[0,1]$), there is some overlap among the definitions above. The form of the overlap depends on whether N is odd or even.

- Exercise 12.5.6.* (i) Show that when N is odd, strong and strict risk dominance are equivalent, and are more demanding than weak risk dominance.
(ii) Show that when N is even, strict and weak risk dominance are equivalent, and are less demanding than strong risk dominance. ‡

Exercise 12.5.7. Matching and risk dominance in normal form games. Consider the symmetric normal form game with strategy set $S = \{0, 1\}$ and payoffs

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Suppose that $a > c$ and $d > b$, so that A is a coordination game with mixed equilibrium $\chi^* = \frac{a-c}{a-b-c+d}$. We say that strategy 1 is strictly risk dominant in the normal form game A if $\chi^* < \frac{1}{2}$, and that strategy 0 is strictly risk dominant in A if $\chi^* > \frac{1}{2}$.

Let F^N be the game obtained when N agents are matched without self-matching to play A , as described in Example 11.4.3:

$$F_0^N(\chi) = \frac{N(1-\chi)-1}{N-1}a + \frac{N\chi}{N-1}b, \text{ and}$$

$$F_1^N(\chi) = \frac{N(1-\chi)}{N-1}c + \frac{N\chi-1}{N-1}d.$$

- (i) Let the preceding equations define F_0^N and F_1^N for all $\chi \in [0, 1]$. Show that strategy 1 is strictly risk dominant in A if and only if

$$F_1^N\left(\frac{N-1}{N}y + \frac{1}{N}\right) > F_0^N\left(\frac{N-1}{N}y\right) \text{ whenever } y \in \left[\frac{1}{2}, 1\right],$$

which is condition (12.59) strengthened to allow y to be outside of \mathcal{X}^{N-1} .

- (ii) Suppose that N is odd. Using part (i), show that strategy 1 is strictly risk dominant in A if and only if it is strictly risk dominant in F^N , as defined in condition (12.59).

- (iii) Now suppose that N is even. Show that strategy 1 being strongly risk dominant in F^N implies that it is strictly risk dominant in A , and that this in turn implies that strategy 1 is strictly risk dominant in F^N . Then show that the two converse implications are false. ‡

We call the two-strategy finite-population game F^N a *coordination game* if there is a *threshold* $y^* \in (0, 1)$ such that

$$(12.61) \quad \text{sgn}\left(F_1^N\left(\frac{N-1}{N}y + \frac{1}{N}\right) - F_0^N\left(\frac{N-1}{N}y\right)\right) = \text{sgn}(y - y^*) \text{ for all } y \in \mathcal{X}^{N-1}.$$

Thus, a clever agent prefers strategy 1 if more than fraction y^* of his opponents play strategy 1, and prefers strategy 0 if less than fraction y^* play strategy 1. (In the matching context of Example 12.5.7, we can let y^* equal χ^* , the weight on strategy 1 in the mixed equilibrium of A .) If y^* is not in \mathcal{X}^{N-1} , it is not uniquely determined, and can be replaced by any other point in the interval $(y_-^*, y_+^*) = (\frac{1}{N-1}\lfloor(N-1)y^*\rfloor, \frac{1}{N-1}\lceil(N-1)y^*\rceil)$ without affecting condition (12.61). The endpoints of this interval are adjacent states in \mathcal{X}^{N-1} , and they are uniquely determined.

Theorem 12.5.8 establishes the connection between risk dominance and stochastic stability under the BRM rule. The proof of the theorem can be viewed as a discrete version of our corresponding analysis from Section 12.3: there the length of each basin of attraction was measured by differences in the signum potential I_{sgn} , while here these lengths are measured by counting states. The discreteness of the state space leads to slightly stronger conditions for stochastic stability, and also forces us to consider a number of similar cases in order to prove the result.

Theorem 12.5.8. *Let F^N be a two-strategy coordination game, and suppose that agents are clever and employ the BRM protocol.*

- (i) *If strategy 1 is weakly risk dominant, then state e_1 is stochastically stable.*
- (ii) *If strategy 1 is strongly risk dominant, then state e_1 is uniquely stochastically stable.*

Proof. We leave part (i) as an exercise and prove part (ii). Let y^* be a threshold for the game F^N , as defined in equation (12.61), and suppose first that $y^* \notin \mathcal{X}^{N-1}$. Then letting $y_+^* = \frac{1}{N-1}\lceil(N-1)y^*\rceil \in \mathcal{X}^{N-1}$, and letting $1[s]$ equal 1 if statement s is true and 0 otherwise, we have

$$\begin{aligned} c_{\chi, \chi + \frac{1}{N}} &= \kappa_{01}(F_0^N(\chi), F_1^N(\chi + \frac{1}{N})) \\ &= 1[F_0^N(\chi) \geq F_1^N(\chi + \frac{1}{N})] \\ &= 1[\chi \leq \frac{N-1}{N}y_+^* - \frac{1}{N}], \text{ and} \end{aligned}$$

$$\begin{aligned}
c_{\chi, \chi - \frac{1}{N}} &= \kappa_{10}(F_0^N(\chi - \frac{1}{N}), F_1^N(\chi)) \\
&= 1[F_0^N(\chi - \frac{1}{N}) \leq F_1^N(\chi)] \\
&= 1[\chi \geq \frac{N-1}{N} y_+^* + \frac{1}{N}],
\end{aligned}$$

where the final equalities follow from condition (12.61). This shows that every state except states 0 and 1 has an outgoing zero-cost edge (and that state $\frac{N-1}{N} y_+^*$ has two of them). Thus states 0 and 1 are the only recurrent states, and hence the only candidates for stochastic stability. We therefore compute

$$\begin{aligned}
C(\tau_1) &= \sum_{k=0}^{N-1} c_{k/N, (k+1)/N} \\
&= \#\{\chi \in \mathcal{X}^N : \chi \leq \frac{N-1}{N} y_+^* - \frac{1}{N}\} \\
&= \#\{z \in \mathbf{Z} : 0 \leq z \leq (N-1)y_+^* - 1\} \\
&= (N-1)y_+^*, \text{ and} \\
C(\tau_0) &= \sum_{k=1}^N c_{k/N, (k-1)/N} \\
&= \#\{\chi \in \mathcal{X}^N : \chi \geq \frac{N-1}{N} y_+^* + \frac{1}{N}\} \\
&= \#\{z \in \mathbf{Z} : (N-1)y_+^* + 1 \leq z \leq N\} \\
&= (N-1)(1 - y_+^*) + 1.
\end{aligned}$$

If N is odd, then strategy 1 is strongly risk dominant if and only if $y_+^* \leq \frac{1}{2}$; if N is even, the relevant inequality is $y_+^* \leq \frac{N-2}{2(N-1)}$. Both of these conditions imply that $C(\tau_1) < C(\tau_0)$. Therefore, by Theorem 12.5.3, state e_1 is uniquely stochastically stable.

Now suppose that $y^* \in \mathcal{X}^{N-1}$. Then

$$\begin{aligned}
c_{\chi, \chi + \frac{1}{N}} &= \kappa_{01}(F_0^N(\chi), F_1^N(\chi + \frac{1}{N})) = 1[\chi \leq \frac{N-1}{N} y^*], \\
c_{\chi, \chi - \frac{1}{N}} &= \kappa_{10}(F_0^N(\chi - \frac{1}{N}), F_1^N(\chi)) = 1[\chi \geq \frac{N-1}{N} y^* + \frac{1}{N}].
\end{aligned}$$

Thus, states 0 and 1 are again the only ones with no outgoing zero-cost edge (though in this case no state has two such edges). Now

$$\begin{aligned}
C(\tau_1) &= \sum_{k=0}^{N-1} c_{k/N, (k+1)/N} = (N-1)y^* + 1, \text{ and} \\
C(\tau_0) &= \sum_{k=1}^N c_{k/N, (k-1)/N} = (N-1)(1 - y^*) + 1.
\end{aligned}$$

If N is odd, then strategy 1 is strongly risk dominant if and only if $y^* \leq \frac{N-3}{2(N-1)}$; if N is even, the relevant inequality is again $y^* \leq \frac{N-2}{2(N-1)}$. Both of these conditions imply that

$C(\tau_1) < C(\tau_0)$, and hence that state e_1 is uniquely stochastically stable. This completes the proof of the theorem. ■

Exercise 12.5.9. (i) Prove part (i) of Theorem 12.5.8.

(ii) Prove that the converse of part (ii) of Theorem 12.5.8 is also true. ‡

Exercise 12.5.10. Show that Theorem 12.5.8 remains true without the assumption that F^N is a coordination game. ‡

Exercise 12.5.11. To focus on cost functions that agree with those from Section 12.3, let us suppose that

$$\kappa_{01}(c, c - d) = \kappa_{10}(c - d, c) = \hat{\kappa}(d) \text{ for all } c, d \in \mathbf{R},$$

where $\hat{\kappa} : \mathbf{R} \rightarrow \mathbf{R}_+$ is nondecreasing, zero on $(-\infty, 0)$, and positive on $(0, \infty)$. Verify that under these assumptions, Theorems 12.5.3 and 12.3.4 agree, in the sense that they single out the same states as stochastically stable. (Hint: It is enough to show that $C(\tau_\chi) - C(\tau_y) = I^N(y) - I^N(\chi)$ for all $\chi, y \in \mathcal{X}^N$.) ‡

Exercise 12.5.12. Let F^N be the population game generated by matching in the two-strategy coordination game A , and suppose that agents are clever and employ a noisy best response rule whose cost function satisfies the conditions from Exercise 12.5.11.

(i) Show that state e_i is stochastically stable if and only if strategy i is risk dominant.

(ii) Show that if $\hat{\kappa}$ is increasing on \mathbf{R}_+ , then state e_i is uniquely stochastically stable if and only if strategy i is strictly risk dominant. Explain in words why strict risk dominance is sufficient for unique stochastic stability here, while strong risk dominance was needed under the BRM rule. ‡

12.5.4 The Radius-Coradius Theorem

While we have thus far confined our analysis to two-strategy games, the real advantage of the graph-theoretic approach to stochastic stability analysis is that it can be applied in nonreversible settings. To explore this possibility, we introduce a new sufficient condition for stochastic stability that is derived from Theorem 12.5.3, but is often easier to apply.

For the intuition behind the new characterization of stochastic stability, look again at Figures 12.5.2 and 12.5.3, which present the trees τ_0 and τ_1 , along with the edge costs generated by a Stag Hunt game under the BRM and logit protocols. In order to prove that state 0 is stochastically stable under these protocols, we showed that in each case, the cost of τ_0 is lower than the cost of τ_1 .

We now suggest an alternative interpretation for this analysis. Each figure shows that the cost of traveling from state 0 to a point from which state 1 can be reached for free exceeds the cost of traveling from state 1 to a point from which state 0 can be reached for free. Put differently, the tree analysis can be recast as one that compares the cost of escaping the basin of attraction of state 0 to the cost of returning to this equilibrium from state 1. If the former cost is greater than the latter, state 0 is stochastically stable.

This intuition is the basis for the *Radius-Coradius Theorem*, which we present next. The proof and extensions of this result are provided in Appendices 12.A.3 and 12.A.4.

As in Section 12.5.2, we suppose that N clever agents repeatedly play a population game $F^N: \mathcal{X}^N \rightarrow \mathbf{R}^n$, responding to revision opportunities by applying a noisy best response protocol σ^ε with cost function κ . Let $\{X_t^\varepsilon\}$ denote the resulting Markov process on \mathcal{X}^N .

Now let $\hat{\mathcal{X}} \subseteq \mathcal{X}$ denote the set of recurrent states of the zero-noise process $\{X_t^0\}$. For distinct x and y in $\hat{\mathcal{X}}$, we define Π_{xy} to be the set of paths $\pi = \{(x, x_1), (x_1, x_2), \dots, (x_{l-1}, y)\}$ through \mathcal{X} such that $x_i \notin \hat{\mathcal{X}}$ for all $i \in \{1, \dots, l-1\}$.

Next, let \mathfrak{R} denote the set of recurrent classes of $\{X_t^0\}$. For distinct recurrent classes $R, S \in \mathfrak{R}$, we define the *cost* of a transition from R to S by

$$c_{RS} = \min_{x \in R} \min_{y \in S} \min_{\pi \in \Pi_{xy}} C(\pi),$$

where $C(\pi)$ is the sum of the edge costs along path π , as defined in equation (12.57). Evidently, c_{RS} is the lowest cost of any path from a state in R to a state in S .

We now introduce notions of difficulty of escape from and ease of return to the basin of attraction of a recurrent class. The *radius* of recurrent class R ,

$$\text{rad}(R) = \min_{S \neq R} c_{RS},$$

is the minimal cost of a transition from R to another recurrent class. It therefore measures the difficulty of exiting the basin of attraction of R .

To define our notion of ease of return, let Π_{SR} denote the set of paths $\pi = \{(S, Q_1), (Q_1, Q_2), \dots, (Q_{l-1}, R)\}$ in \mathfrak{R} from S to R . Then the *coradius* of recurrent class R is given by

$$\text{corad}(R) = \max_{S \neq R} \min_{\pi \in \Pi_{SR}} \sum_{(Q, Q') \in \pi} c_{QQ'}.$$

In words, the coradius of R describes the total cost of reaching recurrent class R from the most disadvantageous initial class S , allowing that route to involve stops in other recurrent classes, and assuming that the path taken from each class to the next has as low a cost as possible.

Theorem 12.5.13 says that if the difficulty of escaping from R , as measured by $\text{rad}(R)$, exceeds the difficulty of returning to R , as measured by $\text{corad}(R)$, then R is the set of stochastically stable states.

Theorem 12.5.13. *Let $\{\{X_t^\varepsilon\}_{t=0}^\infty\}_{\varepsilon \in [0, \bar{\varepsilon}]}$ be a regular collection of Markov chains. If R is a recurrent class of $\{X_t^0\}$ with $\text{rad}(R) > \text{corad}(R)$, then R is the set of stochastically stable states.*

This result follows directly from the Theorem 12.A.6. Appendix 12.A.4 contains the proof of this and more general results.

12.5.5 Half Dominance

Theorem 12.5.8 showed that in two-strategy games, strong risk dominance is sufficient for equilibrium selection under the BRM rule. We now use the Radius-Coradius Theorem to obtain an extension of this result for games with arbitrary numbers of strategies.

In a two-strategy game, a risk dominant strategy is one that is optimal whenever it is played by at least half of one's opponents. Our generalization of risk dominance for n -strategy games, which we call *half-dominance*, retains this crucial property.

We introduce three definitions of half dominance which parallel the definitions of risk dominance from Section 12.5.3. Let F^N be an n -strategy, finite-population game. Strategy i *weakly half dominant* in F^N if

$$(12.62) \quad F_i^N\left(\frac{N-1}{N}y + \frac{1}{N}e_i\right) \geq F_j^N\left(\frac{N-1}{N}y + \frac{1}{N}e_j\right) \text{ for all } y \in X^{N-1} \text{ with } y_i \geq \frac{1}{2} \text{ and all } j \neq i.$$

Strategy i *strictly half dominant* if

$$(12.63) \quad F_i^N\left(\frac{N-1}{N}y + \frac{1}{N}e_i\right) > F_j^N\left(\frac{N-1}{N}y + \frac{1}{N}e_j\right) \text{ for all } y \in X^{N-1} \text{ with } y_i \geq \frac{1}{2} \text{ and all } j \neq i.$$

Finally, strategy i *strongly half dominant* if

$$(12.64) \quad F_i^N\left(\frac{N-1}{N}y + \frac{1}{N}e_i\right) \geq F_j^N\left(\frac{N-1}{N}y + \frac{1}{N}e_j\right) \text{ for all } y \in X^{N-1} \text{ with } y_i \geq \frac{1}{2} \frac{N-2}{N-1} \text{ and all } j \neq i.$$

The statements in the next exercise are the analogues of those made about risk dominance in Exercise 12.5.6.

Exercise 12.5.14. (i) Show that when N is odd, strong and strict half dominance are equivalent, and are more demanding than weak half dominance.

(ii) Show that when N is even, strict and weak half dominance are equivalent, and are less demanding than strong half dominance. ‡

Exercise 12.5.15. Matching, pairwise risk dominance, and half dominance. Let $A \in \mathbf{R}^{n \times n}$ be a symmetric normal form game. We call strategy i *strictly half-dominant* in A if $(Ax)_i > (Ax)_j$ whenever $x_i \geq \frac{1}{2}$.

- (i) Fix a population size N , and let F^N be generated by matching without self-matching in A , as described in Example 11.4.1:

$$F_i^N(x) = \frac{1}{N-1}(A(Nx - e_i))_i = (Ax)_i + \frac{1}{N-1}((Ax)_i - A_{ii}).$$

Suppose we use this equation to define F^N throughout X . Show that strategy i is strictly half dominant in A if and only if

$$F_i^N\left(\frac{N-1}{N}y + \frac{1}{N}e_i\right) > F_j^N\left(\frac{N-1}{N}y + \frac{1}{N}e_j\right) \text{ for all } y \in X \text{ with } y_i \geq \frac{1}{2} \text{ and all } j \neq i.$$

which is condition (12.63) strengthened to allow y to lie outside of \mathcal{X}^{N-1} .

We say that strategy i *strictly pairwise risk dominates* strategy j in A if $A_{ii} + A_{ij} > A_{ji} + A_{jj}$.

- (ii) Construct a symmetric normal form game in which strategy i strictly pairwise risk dominates every other strategy $j \neq i$, but in which strategy i is not strictly half dominant in A .

We say that A satisfies the *marginal bandwagon property* if $A_{ii} - A_{ji} > A_{ik} - A_{jk}$ for all $i, j, k \in S$, so that the loss associated with playing strategy j rather than strategy i is worst when one's opponent in a match plays strategy i .

- (iii) Suppose that every pure strategy of A is a symmetric Nash equilibrium. Show that if strategy i strictly pairwise risk dominates every other strategy, and A satisfies the marginal bandwagon property, then strategy i is strictly half dominant in A .
- (iv) Show that the conclusion of part (iii) remains true even without the restriction that all pure strategies generate symmetric Nash equilibria. (Hint: Show that under the marginal bandwagon property, any pure strategy that does not generate a symmetric Nash equilibrium is strictly dominated by another pure strategy.) ‡

Theorem 12.5.16 shows that strong half dominance implies stochastic stability under the BRM rule.

Theorem 12.5.16. *Suppose that clever agents play a finite-population game F^N using the BRM rule. If strategy i is strongly half-dominant in F^N , then state e_i is uniquely stochastically stable.*

The intuition behind this result is simple. At a first approximation, that strategy i is strongly half dominant means that it is the unique best response whenever it is played by at least fraction α of the population, where $\alpha < \frac{1}{2}$. Since there are N agents in the

population, it follows that $N(1 - \alpha)$ is an approximate lower bound on the number of mutations need to escape the basin of attraction of equilibrium e_i . This lower bound is the radius of $\{e_i\}$. By the same token, $N\alpha$ is an approximate upper bound on the number of mutations needed to reach the basin of e_i from any outside state. This upper bound is the coradius of $\{e_i\}$. Since $\alpha < \frac{1}{2}$, the radius of $\{e_i\}$ is larger than its coradius, implying that e_i is the unique stochastically stable state.

The heuristic analysis above contains a number of loose arguments. The exact proof below makes these arguments precise.

Proof of Theorem 12.5.16. First assume that N is odd, so that there are states in \mathcal{X}^N with $x_i = \frac{N-1}{2N}$ and $x_i = \frac{N+1}{2N}$, and suppose that $j \neq i$. If x is a state with $x_i \geq \frac{N-1}{2N}$, then at least fraction $\frac{N}{N-1} \cdot \frac{N-1}{2N} = \frac{1}{2}$ of a j player's opponents play strategy i . Similarly, if x is a state with $x_i \geq \frac{N+1}{2N}$, then at least fraction $\frac{N}{N-1}(\frac{N+1}{2N} - \frac{1}{N}) = \frac{1}{2}$ of an i player's opponents play strategy i . Since i is strongly half dominant, it follows that

$$(12.65) \quad c_{xy} = 0 \text{ whenever } x_i \geq \frac{N-1}{2N} \text{ and } y = x + \frac{1}{N}(e_i - e_j), \text{ and that}$$

$$(12.66) \quad c_{xy} = 1 \text{ whenever } x_i \geq \frac{N+1}{2N} \text{ and } y = x + \frac{1}{N}(e_j - e_i).$$

Statement (12.65) implies that the only state x that satisfies $x_i \geq \frac{N-1}{2N}$ and is a recurrent state of the zero-noise process is state e_i . Therefore, if y is a recurrent state other than e_i , and $\pi \in \Pi_{e_i y}$, then π must contain edges of the form $(z, z + \frac{1}{N}(e_k - e_i))$ with z_i taking every value in the set $\{1, \frac{N-1}{N}, \dots, \frac{N+1}{2N}\}$. By (12.66), each such edge has cost 1. It follows that $\text{rad}(\{e_i\}) \geq \#\{1, \frac{N-1}{N}, \dots, \frac{N+1}{2N}\} = \frac{N+1}{2N}$.

Now let y be a recurrent state other than e_i . Let $\hat{\pi} \in \Pi_{y e_i}$ be a path in which every edge is of the form $(z, z + \frac{1}{N}(e_i - e_k))$. When $z_i \geq \frac{N-1}{2N}$, (12.65) implies that each such edge has cost zero. It follows that $C(\hat{\pi}) \leq \#\{0, \frac{1}{N}, \dots, \frac{N-3}{2N}\} = \frac{N-1}{2N}$. Since this is true of every recurrent state other than e_i , it must be that $\text{corad}(\{e_i\}) \leq \frac{N-1}{2N}$. Thus $\text{rad}(\{e_i\}) \geq \frac{N+1}{2N} > \frac{N-1}{2N} \geq \text{corad}(\{e_i\})$, and so Theorem 12.5.13 implies that state e_i is uniquely stochastically stable.

Next, assume that N is even, so that there are states in \mathcal{X}^N with $x_i = \frac{1}{2}$, and suppose that $j \neq i$. If x is a state with $x_i \geq \frac{1}{2}$, then at least fraction $\frac{N}{N-1}(\frac{1}{2} - \frac{1}{N}) = \frac{N-2}{2(N-1)}$ of an i player's opponents and at least fraction $\frac{N}{2(N-1)}$ of a j player's opponents are playing i . Since strategy i is strongly half dominant, it follows that

$$c_{xy} = 0 \text{ whenever } x_i \geq \frac{1}{2} \text{ and } y = x + \frac{1}{N}(e_i - e_j), \text{ and that}$$

$$c_{xy} = 1 \text{ whenever } x_i \geq \frac{1}{2} \text{ and } y = x + \frac{1}{N}(e_j - e_i).$$

Arguments analogous to those above then show that $\text{rad}(\{e_i\}) \geq \frac{N}{2} + 1$ and that $\text{corad}(\{e_i\}) \leq$

$\frac{N}{2}$, so Theorem 12.5.13 again implies that state e_i is uniquely stochastically stable. ■

The graph-theoretic machinery developed in this section has been fruitful for stochastic stability analysis in more settings beyond the considered here; in particular, models of extensive form games and local interaction have been studied using these techniques (see the Notes). In our present setting of global interaction, simultaneous move population games, few results exist beyond those we have presented. Still, the currently available tools seem sufficient to develop further equilibrium selection results, making this a natural direction for future research.

12.6 Large Population Limits

In this final section, we consider infinite horizon behavior in the large population limit. We proved in Theorem 10.2.3 that over any finite time interval $[0, T]$, the sample paths of the stochastic evolutionary process $\{X_t^N\}$ are very likely to stay very close to a solution trajectory of the mean dynamic (M) once the population size N is large enough. This deterministic approximation is not valid over the infinite horizon. For instance, if $\{X_t^N\}$ is irreducible, it must visit every state in \mathcal{X}^N infinitely often; this behavior cannot be approximated by a solution to an ordinary differential equation.

Although the infinite horizon analogue of Theorem 10.2.3 is false, it is still natural to expect the mean dynamic to impose restrictions on infinite horizon behavior: while all states in \mathcal{X}^N are visited infinitely often, states that are visited frequently should correspond to possible long run behaviors of the mean dynamic when N is large enough. After all, if state x is “transient” under (M), then a stochastic process visiting x is very likely to follow the solution trajectory of (M) leading away from x , and it cannot return to x without an excursion against the flow of (M).

We formalize this intuition by studying the behavior of the stationary distribution μ^N as N approaches infinity. In Section 12.6.1, we show that under quite general conditions, the mass in μ^N can only accumulate on states that are recurrent under (M). It follows that in settings where the mean dynamic has a globally attracting state—for instance, in strictly stable games under a variety of revision protocols—this state is uniquely stochastically stable. Section 12.6.2 offers more refined results, showing that in potential games and supermodular games, under perturbed best response protocols, the mass in μ^N can only accumulate on states that are locally stable rest points of (M). These results in this section bring together ideas developed throughout the book—the well-behaved classes of games from Part I, the analyses of deterministic evolutionary dynamics from Parts II and III, and

the study of infinite horizon behavior under stochastic evolutionary dynamics from Part IV.

12.6.1 Convergence to Recurrent States of the Mean Dynamic

Fix a population size N , let $\{X_t^N\}$ be the stochastic evolutionary process generated by a population game F^N and a Lipschitz continuous revision protocol ρ . As in Section 10.2.1, we let ζ_x^N be a random variable whose distribution describes the stochastic increment of $\{X_t^N\}$ from state x :

$$\mathbb{P}(\zeta_x^N = z) = \mathbb{P}\left(X_{\tau_{k+1}}^N = x + z \mid X_{\tau_k}^N = x\right),$$

where τ_k is the k th jump time of the process $\{X_t^N\}$. Until Section 12.6.2, we need not assume that the process $\{X_t^N\}$ is irreducible, and we can allow the transition probabilities of the process $\{X_t^N\}$ to be influenced by the finite-population effects (see Exercise 12.6.1 below).

We suppose that as N grows large, the sequence of games $\{F^N\}_{N=N_0}^\infty$ converges uniformly to a Lipschitz continuous, continuous-population game $F: X \rightarrow \mathbf{R}^n$. It follows that the transition laws of the processes $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ converge: there is a collection of random variables $\{\zeta_x\}_{x \in X}$ taking values in $\{z \in \mathbf{R}^n: z = e_j - e_i \text{ for some } i, j \in S\}$ whose transition probabilities are Lipschitz continuous in x , and that satisfy

$$(12.67) \quad \lim_{N \rightarrow \infty} \max_{x \in X^N} \max_{i, j \in S} \left| \mathbb{P}\left(\zeta_x^N = \frac{1}{N}(e_j - e_i)\right) - \mathbb{P}\left(\zeta_x = e_j - e_i\right) \right| = 0.$$

In other words, $N\zeta_x^N$ converges in distribution to ζ_x , with convergence uniform over X .

Exercise 12.6.1. Verify that condition (12.67) holds, even allowing for finite population effects, including clever payoff evaluation, no self-imitation, and fixed numbers of committed agents for each strategy (see Section 11.4). ‡

As in Section 10.2.1, let $V^N(x) = N \mathbb{E} \zeta_x^N$ be the expected increment per time unit of the process $\{X_t^N\}$. Then condition (12.67) implies that V^N converges uniformly to the Lipschitz continuous function $V: X \rightarrow TX$ defined by $V(x) = \mathbb{E} \zeta_x$. Moreover, Theorem 10.2.3 tells us that over any finite time interval, if N is large enough, the process $\{X_t^N\}$ is well-approximated by an appropriate solution of the mean dynamic,

$$(M) \quad \dot{x} = V(x).$$

To describe infinite horizon behavior in the large population limit, we recall from Appendix 7.A.1 that state $x \in X$ is a *recurrent point* of (M) ($x \in \mathcal{R}$) if $x \in \omega(x)$, so that x is

contained in its own ω -limit set. The closure $\text{cl}(\mathcal{R})$ of the set of recurrent points is called the *Birkhoff center* of (M). Our first result tells us that when the population size is large, any stationary distribution of $\{X_t^N\}$ must concentrate its mass near this set.

Theorem 12.6.2. *Under the assumptions above, let $\{\mu^N\}_{N=N_0}^\infty$ be any sequence of stationary distributions for the Markov processes $\{X_t^N\}_{N=N_0}^\infty$, and let $O \subseteq X$ be an open set containing $\text{cl}(\mathcal{R})$. Then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.*

Theorem 12.6.2 is a consequence of Theorem 12.B.1 in Appendix 12.B.

Theorem 12.6.2 provides an infinite horizon complement to our finite-horizon approximation result, Theorem 10.2.3. While Theorem 12.6.2 does not tell us that the process $\{X_t^N\}$ will shadow any particular solution trajectory of (M) indefinitely, it does show that when the population size is large, the vast majority of time in the very long run will be spent near recurrent points of (M). Together, Theorems 10.2.3 and 12.6.2 help to justify our single-minded focus on mean dynamic in the middle half of the book.

Since the Birkhoff center contains not only rest points, but also closed orbits and more complicated limit sets of the mean dynamic, Theorem 12.6.2 is usually not enough to identify the stochastically stable states. However, if the mean dynamic has a globally attracting state x^* , then $\mathcal{R} = \{x^*\}$, so the theorem implies that x^* is uniquely stochastically stable. We present this general result in Appendix 12.B.1 (see Corollary 12.B.2). Here in the text, we note some implications for stochastic evolution in stable games. To read part (ii) of the result below, recall that the set of perturbed equilibria $PE(F, v)$ consists of the rest points of the perturbed best response dynamic generated by F and v (see Section 6.2).

Corollary 12.6.3. *Suppose the conditions of Theorem 12.6.2 hold.*

- (i) *Let F be a strictly stable game with unique Nash equilibrium x^* , and let ρ be a revision protocol that generates an integrable excess payoff dynamic or an impartial pairwise comparison dynamic. Then μ^N converges weakly to a point mass at x^* , and so x^* is uniquely stochastically stable.*
- (ii) *Let F be a stable game, and let ρ be a perturbed best response protocol generated by an admissible deterministic perturbation v . Then μ^N converges weakly to a point mass at the unique perturbed equilibrium $\tilde{x} \in PE(F, v)$, and so \tilde{x} is uniquely stochastically stable.*

Proof. Theorems 7.2.8 and 7.2.11 show that in a strictly stable game, the unique Nash equilibrium x^* is globally asymptotically stable under any integrable excess payoff dynamic and any impartial pairwise comparison dynamic. Similarly, Theorem 7.2.10 shows that in a stable game, any perturbed best response dynamic generated by an admissible deterministic perturbation has a unique and globally asymptotically stable rest point. The corollary follows immediately from these results and Theorem 12.6.2. ■

12.6.2 Convergence to Stable Rest Points of the Mean Dynamic

In cases where the mean dynamic has multiple components of recurrent states, Theorem 12.6.2 does not tell us which of these components receive mass in the limiting stationary distribution. Still, components that are unstable seem like poor candidates for this role: if the process $\{X_t^N\}$ moves off the component, then obeying the mean dynamic (M) will draw the process away from the component, and also make it difficult for the process to return to the component.

For this argument to be sound, it must be the case that when the process visits, say, an unstable rest point of (M), it does not stop moving entirely; rather, there must be enough random motion that the process can leave the immediate neighborhood of the rest point, after which the pull of the mean dynamic will take the process still further away. In fact, we already know that without this random motion, the process can converge to an unstable rest point of the mean dynamic and stay there forever: for instance, this is precisely what happens if agents use a positive imitative protocol (without mutations or committed agents) in a game with no pure Nash equilibria (see Example 12.B.3). Of the families of processes from Chapters 5 and 6 that are defined by continuous revision protocols—imitative, excess payoff, pairwise comparison, and perturbed best response—only the last feature exhibit the nondegeneracy needed to rule out unstable rest points of the mean dynamic as limit behaviors of the stochastic process.

While one might expect such nondegeneracy to be sufficient for ruling out unstable rest points, it is not. For instance, one can construct deterministic dynamics whose recurrent states are all locally unstable (see Example 12.B.4); in this case, Theorem 12.6.2 tells us that unstable states are the only candidates for stochastic stability. To ensure that the mass in μ^N does not accumulate at unstable rest points, we need some assurance that the mean dynamic is well-behaved in some global sense—for instance, that it admits a global Lyapunov function (Appendix 7.B), or that it is a cooperative differential equation (7.C)—so that examples of the sort just mentioned are ruled out. In Appendix 12.B.2, we state general results of this sort; here in the text, we present implications for perturbed best response processes.

Existing results that rule out unstable rest points allow less flexibility defining the process $\{X_t^N\}$ than was allowed in Section 12.6.1. In particular, transition probabilities cannot depend on N . In the notation of the previous section, we require that

$$\mathbb{P}\left(\zeta_x^N = \frac{1}{N}(e_j - e_i)\right) = \mathbb{P}\left(\zeta_x = e_j - e_i\right) \text{ for all } i, j \in S, x \in \mathcal{X}^N, \text{ and } N \geq N_0.$$

For this restriction to hold, payoffs must be independent of the population size ($F^N = F$),

and there cannot be finite population effects of the sorts introduced in Section 11.4; in particular, agents must use simple payoff evaluation.

With these additional restrictions in place, we can state our result. Below, $PE(F, v)$ again denotes the set of perturbed equilibria for the pair (F, v) , and $LS(F, v)$ denotes the set of Lyapunov stable states under the perturbed best response dynamic generated by F and v .

Theorem 12.6.4. *Under the conditions of Theorem 12.6.2 and the additional restrictions noted above:*

- (i) *Let F be a potential game, and let ρ be a perturbed best response protocol generated by an admissible deterministic perturbation v . Suppose $PE(F, v)$ is finite. If O is an open set containing $LS(F, v)$, then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.*
- (ii) *Let F be an irreducible supermodular game, and that ρ is a perturbed best response protocol generated by an admissible stochastic perturbation ε . Let v be the corresponding deterministic perturbation, whose existence is ensured by Theorem 6.2.2. If O is an open set containing $LS(F, v)$, then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.*

Proof. Theorem 7.1.6 shows that under the conditions of part (i), the function $\tilde{f}(x) = f(x) - v(x)$ is a strict global Lyapunov function for the perturbed best response dynamic. Moreover, that ρ is a perturbed best response protocol implies that the process $\{X_t^N\}$ is nondegenerate at every rest point x of the perturbed best response dynamic, in the sense that $\text{span}(\text{support}(\zeta_x)) = TX$. The conclusion of part (i) therefore follows from Theorem 12.B.6.

Observation 7.3.5 and Theorem 7.3.7 show that under the conditions of part (ii), the perturbed best response dynamic is strongly cooperative, and hence irreducible (see Appendix 7.C). The conclusion of part (ii) then follows from Example 12.B.7 and Theorem 12.B.8. ■

While Theorem 12.6.4 represents an improvement on our earlier results, it still is not enough to identify the stochastically stable states in cases with multiple stable rest points. Accomplishing this would require us to evaluate the asymptotic probabilities of motions of $\{X_t^N\}$ that go against the flow of the mean dynamic, causing transitions between its stable rest points. (The situation is analogous to Theorem 12.A.5 in Appendix 12.A.3, but it is more complex, as it concerns the large population limit rather the small noise limit.) It is known that the asymptotic transition probabilities can be obtained as solutions to certain optimal control problems, which describe paths through X that minimize the “costs” of transitions between stable rest points. However, solving this optimal control

problem seems difficult, at least outside of the reversible cases studied in Sections 12.2 and 12.3; see the Notes for further discussion and references.

A larger unanswered question is that of the agreement of the small noise and large population double limits. We saw in Sections 12.2 and 12.3 that these double limits agree in many (though not all) reversible cases, allowing us to predict infinite horizon behavior without having to assess whether the rarity of mistakes or the averaging generated by large population sizes drives equilibrium selection. While it seems natural to expect this agreement to extend to nonreversible but otherwise well-behaved settings, whether and when it does so is not yet known.

Appendix

12.A Trees, Escape Paths, and Stochastic Stability

In Appendix 11.A, we saw that an irreducible Markov process $\{X_t\}$ on a finite state space \mathcal{X} admits a unique stationary distribution μ , and that μ is both the limiting distribution and the limiting empirical distribution of the process. In this Appendix, we look at limits of the stationary distribution as the parameters that describe the process $\{X_t\}$ approach their extremities.

In this section we consider small noise limits. Here the state space of the Markov chain stays fixed, but the transition probabilities approach some limiting values as the noise level vanishes. The analysis is most interesting in cases where each Markov chain corresponding to a positive noise level is irreducible, and so has a unique stationary distribution, while the chain corresponding to the zero noise limit is not, and so has multiple recurrent classes and stationary distributions. We will see that as the noise level approaches zero, the corresponding sequence of unique stationary distributions converges, and that its limit is a stationary distribution of the limiting Markov chain. In this way, the limiting analysis typically singles out a single recurrent class of the limiting process; the states in this class are said to be *stochastically stable*.

We begin our analysis by showing how the stationary distribution of an irreducible Markov chain on the finite state space \mathcal{X} can be represented in terms of trees whose nodes are in \mathcal{X} . We then introduce Markov chains with a noise parameter, and show how the tree characterization of the stationary distribution can be used to determine the set of stochastically stable states. While we present all of our results in a discrete-time setting, these results apply immediately to continuous-time Markov processes with constant jump rates—see Section 11.A.7.

12.A.1 The Markov Chain Tree Theorem

Let $\{X_t\}$ be an irreducible Markov chain on the finite state space \mathcal{X} . Let $P \in \mathbf{R}_+^{X \times X}$ be its transition matrix, and let $\mu \in \mathbf{R}_+^X$ be its unique stationary distribution: $\mu'P = \mu'$.

The Markov Chain Tree Theorem characterizes μ in terms of certain graphs defined on the state space \mathcal{X} . We view \mathcal{X} as a set of *nodes* that can be connected by *directed edges*, which are ordered pairs $(x, y) \in \mathcal{X} \times \mathcal{X}$ with $x \neq y$. A *directed graph* g on \mathcal{X} can be identified with a set of directed edges.

Five special types of directed graphs will be important in the coming analysis: we introduce four of them here, and the fifth in the proof of Theorem 12.A.1 below. A *walk* from x to y is a directed graph $\{(x, x_1), (x_1, x_2), \dots, (x_{l-1}, y)\}$ whose directed edges traverse a route connecting x to y . A *path* from x to $y \neq x$ is a walk from x to y with no repeated nodes. A *cycle* is a walk from x to itself that contains no other repeated nodes.

A *tree with root* x , also called an *x -tree*, is a directed graph with no outgoing edges from x , exactly one outgoing edge from each $y \neq x$, and a unique path from each $y \neq x$ to x . We let T_x denote the set of x -trees on \mathcal{X} .

Now, define the vector $v \in \mathbf{R}_+^X$ by

$$(12.68) \quad v_x = \sum_{\tau_x \in T_x} \prod_{(y,z) \in \tau_x} P_{yz}.$$

The scalar v_x is obtained by taking the product of the transition probabilities associated with each x -tree, and then summing the results over all x -trees. The irreducibility of $\{X_t^N\}$ implies that at least one of the summands is positive, and hence that v_x itself is positive.

In fact, much more is true. We now establish that the positive vector v is proportional to the stationary distribution of $\{X_t\}$.

Theorem 12.A.1 (The Markov Chain Tree Theorem).

Let $\{X_t\}$ be an irreducible Markov chain with transition matrix $P \in \mathbf{R}_+^{X \times X}$ and stationary distribution μ . If we define the vector v via equation (12.68), then v is a positive multiple of μ .

Proof. Let G_x be the set of directed graphs γ_x on \mathcal{X} such that (i) each $y \in \mathcal{X}$ has exactly one outgoing edge in γ_x ; (ii) γ_x contains a unique cycle; and (iii) the unique cycle contains x . The proof of theorem is based on the observation that the set G_x can be represented in two distinct ways as disjoint unions of trees augmented by a single edge:

$$(12.69) \quad G_x = \bigcup_{y \neq x} \bigcup_{\tau_y \in T_y} \tau_y \cup \{(y, x)\};$$

$$(12.70) \quad G_x = \bigcup_{\tau_x \in T_x} \bigcup_{y \neq x} \tau_x \cup \{(x, y)\}.$$

In (12.69), we take each τ_y tree with $y \neq x$ and construct a graph $\gamma_x \in G_x$ by adding an edge from y to x . In (12.70), we take each τ_x tree and each $y \neq x$, and construct a graph $\gamma_x \in G_x$ by adding an edge from x to y .

Now let ψ be the analogue of ν whose components are obtained by taking the sum over graphs in G_x rather than over trees in T_x :

$$\psi_x = \sum_{\gamma_x \in G_x} \prod_{(y,z) \in \gamma_x} P_{yz}.$$

Using (12.69) and (12.70), we obtain two new expressions for ψ in terms of ν and P :

$$\sum_{y \neq x} \nu_y P_{yx} = \psi_x = \nu_x \sum_{y \neq x} P_{xy}.$$

By adding $\nu_x P_{xx}$ to the first and last expressions, we find that

$$\sum_{y \in \mathcal{X}} \nu_y P_{yx} = \nu_x,$$

or, equivalently, that $\nu'P = \nu'$. Since the stationary distribution of $\{X_t^N\}$ is unique and since ν is positive, ν must be a positive multiple of μ . ■

12.A.2 Limiting Stationary Distributions via Trees

We now consider the long run behavior of Markov chains whose transition probabilities include small noise terms. These noise terms ensure that these Markov chains are irreducible, even in cases where the unperturbed chain is not. It is natural to expect the stationary distributions of the perturbed chains to place most of their weight on the recurrent classes of the unperturbed chain. But, more interestingly, one typically finds nearly all of this weight being placed on just one recurrent class of the unperturbed chain. The members of this recurrent class are said to be *stochastically stable*. In summary, by introducing a small amount of noise to the transition probabilities, we can obtain unique, history-independent predictions about infinite horizon behavior in settings where multiple, history-dependent predictions would otherwise prevail.

Before proceeding, we observe that it is possible to carry out the analysis below under weaker restrictions on the form of transition probabilities. This approach is pursued in Section 12.A.5 below.

Let $\{\{X_t^\varepsilon\}_{t=0}^\infty\}_{\varepsilon \in [0, \bar{\varepsilon}]}$ be a parameterized collection of Markov chains on the finite state space

\mathcal{X} . The transition matrices $P^\varepsilon \in \mathbf{R}_+^{\mathcal{X} \times \mathcal{X}}$ of these chains are assumed to vary continuously in ε . In addition, for each pair of distinct states x and $y \neq x$, we assume that if $P_{xy}^0 = 0$, and if $P_{xy}^{\hat{\varepsilon}} > 0$ for some $\hat{\varepsilon} \in (0, \bar{\varepsilon}]$, then the collection $\{P_{xy}^\varepsilon\}_{\varepsilon \in (0, \bar{\varepsilon})}$ satisfies

$$(12.71) \quad P_{xy}^\varepsilon = a_{xy}\varepsilon^k + o(\varepsilon^k) \text{ for some } a_{xy} > 0 \text{ and some } k > 0.$$

(Recall that $o(\varepsilon^k)$ represents a remainder function $r: (0, \bar{\varepsilon}) \rightarrow \mathbf{R}$ satisfying $\lim_{\varepsilon \rightarrow 0} r(\varepsilon)/\varepsilon^k = 0$, so that $r(\varepsilon)$ approaches zero faster than ε^k approaches zero.) In words, we have assumed that each positive but vanishing off-diagonal component of the transition matrix has a leading term that is monomial in ε . Finally, we assume that when ε is positive, the chain $\{X_t^\varepsilon\}$ is irreducible, and so admits a unique stationary distribution $\mu^\varepsilon \in \mathbf{R}_+^{\mathcal{X}}$. When all of these assumptions are met, we say that the collection $\{\{X_t^\varepsilon\}_{t=0}^\infty\}_{\varepsilon \in [0, \bar{\varepsilon}]}$ is *regular*.

Next, we define the *cost* c_{xy} of the transition from state x to state $y \neq x$ for the regular collection $\{\{X_t^\varepsilon\}_{t=0}^\infty\}_{\varepsilon \in [0, \bar{\varepsilon}]}$:

$$(12.72) \quad c_{xy} = \begin{cases} 0 & \text{if } P_{xy}^0 > 0, \\ k & \text{if } P_{xy}^\varepsilon = a_{xy}\varepsilon^k + o(\varepsilon^k), \\ \infty & \text{if } P_{xy}^\varepsilon = 0 \text{ for all } \varepsilon \in [0, \bar{\varepsilon}]. \end{cases}$$

By regularity, this definition covers all possible cases.

We now extend the notion of cost to graphs on \mathcal{X} . Given an x -tree τ_x , we define the *cost* $C(\tau_x)$ of this tree to be the sum of the costs of the edges it contains:

$$C(\tau_x) = \sum_{(y,z) \in \tau_x} c_{yz}.$$

Finally, we let C_x^* be the lowest cost of any x -tree, and let C^* be the minimum value of C_x^* over all states in \mathcal{X} :

$$C_x^* = \min_{\tau_x \in T_x} C(\tau_x), \text{ and } C^* = \min_{x \in \mathcal{X}} C_x^*.$$

With these preliminaries in hand, we can state the following result on the limiting behavior of the stationary distributions μ^ε .

Theorem 12.A.2. *Let $\{\{X_t^\varepsilon\}_{t=0}^\infty\}_{\varepsilon \in [0, \bar{\varepsilon}]}$ be a regular collection of Markov chains with transition costs c_{xy} , and let μ^ε be the stationary distribution of $\{X_t^\varepsilon\}$. Then*

- (i) $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$ exists.
- (ii) μ^* is a stationary distribution of $\{X_t^0\}$.

(iii) $\mu_x^* > 0$ if and only if $C_x^* = C^*$.

State $x^* \in \mathcal{X}$ is said to be *stochastically stable* if it receives positive mass in the limiting stationary distribution μ^* . According to Theorem 12.A.2, state x^* is stochastically stable if and only if there is an x^* -tree whose cost is minimal. Part (ii) of the theorem tells us that the limiting stationary distribution μ^* is a stationary distribution of the unperturbed process $\{X_t^0\}$. This tells us that the perturbations provide a selection among the recurrent classes of this unperturbed process: μ^* either puts positive mass on all states in a given recurrent class of $\{X_t^0\}$, or it puts zero mass on all of the states in this class—compare Theorem 11.A.11 and the subsequent discussion.

The argument to follow not only establishes the theorem, but also provides an explicit expression for the limiting stationary distribution.

Proof of Theorem 12.A.2. To begin, marshal the previous definitions to show that

$$\prod_{(y,z) \in \tau_x} P_{yz}^\varepsilon = \prod_{(y,z) \in \tau_x} (a_{yz} \varepsilon^{c_{yz}} + o(\varepsilon^{c_{yz}})) = A_{\tau_x} \varepsilon^{C(\tau_x)} + o(\varepsilon^{C(\tau_x)}), \quad \text{where } A_{\tau_x} = \prod_{(y,z) \in \tau_x} a_{yz}.$$

Thus, defining $v^\varepsilon \in \mathbf{R}_+^{\mathcal{X}}$ as in (12.68), we see that

$$v_x^\varepsilon = \sum_{\tau_x \in \mathcal{T}_x} \prod_{(y,z) \in \tau_x} P_{yz}^\varepsilon = \alpha_x \varepsilon^{C_x^*} + o(\varepsilon^{C_x^*}), \quad \text{where } \alpha_x = \sum_{\tau_x: C(\tau_x) = C_x^*} A_{\tau_x}.$$

Now Theorem 12.A.1 tells us that

$$(12.73) \quad \mu_x^\varepsilon = \frac{v_x^\varepsilon}{\sum_{y \in \mathcal{X}} v_y^\varepsilon} = \frac{\alpha_x \varepsilon^{C_x^*} + o(\varepsilon^{C_x^*})}{\sum_{y \in \mathcal{X}} (\alpha_y \varepsilon^{C_y^*} + o(\varepsilon^{C_y^*}))}.$$

Taking limits yields

$$\mu_x^* \equiv \lim_{\varepsilon \rightarrow 0} \mu_x^\varepsilon = \begin{cases} \frac{\alpha_x}{\sum_{y: C_y^* = C^*} \alpha_y} & \text{if } C_x^* = C^*, \\ 0 & \text{if } C_x^* > C^*, \end{cases}$$

proving parts (i) and (iii) of the theorem.

Finally, since $\mu_x^* = \lim_{\varepsilon \rightarrow 0} \mu_x^\varepsilon$ exists, since $\lim_{\varepsilon \rightarrow 0} P^\varepsilon = P^0$, and since $(\mu^\varepsilon)' P^\varepsilon = (\mu^\varepsilon)'$ for all $\varepsilon \in (0, \bar{\varepsilon}]$, a straightforward continuity argument shows that $(\mu^*)' P^0 = (\mu^*)'$, proving part (ii) of the theorem. ■

In addition to identifying the stochastically stable states, the proof of Theorem 12.A.2 establishes rates of decay for the stationary distribution weights on other states μ_x^ε .

Corollary 12.A.3. *Under the conditions of Theorem 12.A.2, there are constants $b_x > 0$ such that $\mu_x^\varepsilon = b_x \varepsilon^{C_x^* - C^*} + o(\varepsilon^{C_x^* - C^*})$.*

12.A.3 Limiting Stationary Distributions via Trees on Recurrent Classes

Since in the end, μ^* provides a selection from the recurrent classes of the limit chain $\{X_t^0\}$, it is natural to seek a version of Theorem 12.A.2 stated in terms of trees whose nodes are not the states in \mathcal{X} , but rather the recurrent classes of $\{X_t^0\}$. Since there are generally far fewer recurrent classes than states, such a result could be significantly easier to apply than Theorem 12.A.2. After obtaining an auxiliary result, we establish the desired characterization in Theorem 12.A.5 below.

To begin, let $\hat{\mathcal{X}} \subseteq \mathcal{X}$ denote the set of recurrent states of $\{X_t^0\}$. For each $\varepsilon \in (0, \bar{\varepsilon}]$, we define the *censored chain* $\{\hat{X}_k^\varepsilon\}_{k=0}^\infty$ be a process that records the values of the the process $\{X_t^\varepsilon\}$ when it visits states in $\hat{\mathcal{X}}$. More precisely, we let T_k^ε denote the (random) time that $\{X_t^\varepsilon\}$ makes its k th visit to $\hat{\mathcal{X}}$, and then set $\hat{X}_k^\varepsilon = X_{T_k^\varepsilon}^\varepsilon$.

One can verify that $\{\hat{X}_k^\varepsilon\}$ is indeed a Markov chain that inherits irreducibility from the original chain. Using Theorem 11.A.22, one can obtain an expression for the stationary distribution $\hat{\mu}^\varepsilon$ of the censored chain in terms of that of the original chain:

$$\hat{\mu}_x^\varepsilon = \frac{\mu_x^\varepsilon}{\sum_{y \in \hat{\mathcal{X}}} \mu_y^\varepsilon} \text{ for all } x \in \hat{\mathcal{X}}.$$

Since Theorem 12.A.2 implies that $\lim_{\varepsilon \rightarrow 0} \mu_z^\varepsilon = 0$ whenever $z \notin \hat{\mathcal{X}}$, we can conclude that the masses in the distributions $\hat{\mu}^\varepsilon$ converge to those in our original limiting stationary distribution μ^* :

$$(12.74) \quad \lim_{\varepsilon \rightarrow 0} \hat{\mu}_x^\varepsilon = \mu_x^* \text{ for all } x \in \hat{\mathcal{X}}.$$

In light of equation (12.74), we can determine the limiting stationary distribution μ^* by applying the analysis from Section 12.A.2 to the censored chain $\{\hat{X}_k^\varepsilon\}$. For $x, y \in \hat{\mathcal{X}}$, let W_{xy} be the set of walks $\{(x, x_1), (x_1, x_2), \dots, (x_{l-1}, y)\}$ from x to y such that $x_i \notin \hat{\mathcal{X}}$ for all $i \in \{1, \dots, l-1\}$. Then the transition probabilities of the censored chain $\{\hat{X}_k^\varepsilon\}$ are given by

$$\hat{P}_{xy}^\varepsilon = \sum_{w \in W_{xy}} \prod_{(z, z') \in w} P_{zz'}^\varepsilon.$$

By regularity, it follows that for some $\hat{a}_{xy} > 0$, we have

$$(12.75) \quad \hat{P}_{xy}^\varepsilon = \hat{a}_{xy} \varepsilon^{\hat{c}_{xy}} + o(\varepsilon^{\hat{c}_{xy}}), \quad \text{where } \hat{c}_{xy} = \min_{w \in W_{xy}} \sum_{(z,z') \in w} c_{zz'}.$$

In words, the cost of the transition from x to y in the censored chain is minimum total cost of a sequence from x to y in the original chain.

To obtain a slightly simpler characterization of the transition cost \hat{c}_{xy} in (12.75), notice that one of the walks in W_{xy} that achieves the minimum in (12.75) will be a path—that is, a walk that visits no state more than once. Thus, if we let Π_{xy} be the set of paths $\pi = \{(x, x_1), (x_1, x_2), \dots, (x_{l-1}, y)\}$ through \mathcal{X} such that $x_i \notin \hat{\mathcal{X}}$ for all $i \in \{1, \dots, l-1\}$, we can write

$$(12.76) \quad \hat{c}_{xy} = \min_{\pi \in \Pi_{xy}} \sum_{(z,z') \in \pi} c_{zz'}.$$

Now let $\hat{\tau}_x$ denote an x -tree with nodes in $\hat{\mathcal{X}}$, and let \hat{T}_x denote the set of such trees. If we define

$$\hat{C}(\hat{\tau}_x) = \sum_{(y,z) \in \hat{\tau}_x} \hat{c}_{yz}, \quad \hat{C}_x^* = \min_{\hat{\tau}_x \in \hat{T}_x} \hat{C}(\hat{\tau}_x) \quad \text{and} \quad \hat{C}^* = \min_{y \in \hat{\mathcal{X}}} \hat{C}_y^*,$$

then Theorem 12.A.2 and equation (12.74) yield

Corollary 12.A.4. *Under the assumptions of Theorem 12.A.2, $\mu_x^* > 0$ if and only if $\hat{C}_x^* = \hat{C}^*$.*

Had we attempted to define a censored process directly on the set \mathfrak{R} of recurrent classes of $\{X_t^0\}$, we would have encountered a difficulty: unless each recurrent class is a singleton, this process need not have the Markov property. Nevertheless, a few simple observations will allow us to move from Corollary 12.A.4, which characterizes stochastic stability using graphs on the set of recurrent states $\hat{\mathcal{X}}$, to a characterization via graphs on the set of recurrent classes \mathfrak{R} .

Let us note a few simple consequences of the fact that transitions between states in the same recurrent class have zero cost. First, any minimal cost x -tree $\hat{\tau}_x$ must have these two properties: (i) for each recurrent class R that does not contain x , $\hat{\tau}_x$ has exactly one edge from a state in R to a state outside R ; (ii) $\hat{\tau}_x$ has no edges from any state in R_x , the recurrent class that contains x , to a state outside R_x . For instance, if $\hat{\tau}_x$ had two states $y, z \in R$ with edges to states outside R , then replacing the outgoing edge from z with one from z to y would result in a new x -tree with lower cost than $\hat{\tau}_x$. Second, the minimal cost \hat{C}_x^* is the same for all states in the same recurrent class. For if $x, y \in R_x$ and $\hat{\tau}_x$ is a minimal cost

x -tree, we can construct a y -tree $\hat{\tau}_y$ with $C(\hat{\tau}_y) \leq C(\hat{\tau}_x)$ from $\hat{\tau}_x$ by deleting the outgoing edge from y and adding an edge from x to y . (In fact, point (ii) above tells us that both the new edge from x and the deleted edge from y have zero cost, and so that $C(\hat{\tau}_y) = C(\hat{\tau}_x)$.)

These observations imply that in applying Corollary 12.A.4, one can treat all of the nodes in a given recurrent class as a unit by replacing trees on recurrent states with trees on recurrent classes. To state this idea precisely, we define the cost c_{RS} of a transition from recurrent class R to recurrent class $S \neq R$ by

$$(12.77) \quad c_{RS} \equiv \min_{x \in R} \min_{y \in S} \hat{c}_{xy} = \min_{x \in R} \min_{y \in S} \min_{\pi \in \Pi_{xy}} \sum_{(z,z') \in \pi} c_{zz'}.$$

Let τ_R be an R -tree with nodes in \mathfrak{R} , the set of recurrent classes of $\{X_t^0\}$, and let T_R be the set of such trees. Using formulation (12.77) of transition costs, define

$$C(\tau_R) = \sum_{(Q,Q') \in \tau_R} c_{QQ'}, \quad C_R^* = \min_{\tau_R \in T_R} C(\tau_R), \quad \text{and} \quad C^* = \min_{R \in \mathfrak{R}} C_R^*.$$

Then Corollary 12.A.4, the observations above, and some bookkeeping are enough to establish the following result.

Theorem 12.A.5. *Let $\{\{X_t^\epsilon\}_{t=0}^\infty\}_{\epsilon \in [0, \bar{\epsilon}]}$ be a regular collection of Markov chains with transition costs c_{xy} , let μ^ϵ be the stationary distribution of $\{X_t^\epsilon\}$, and let $\mu^* = \lim_{\epsilon \rightarrow 0} \mu^\epsilon$. Then $\text{support}(\mu^*) = \bigcup \{R \in \mathfrak{R} : C_R^* = C^*\}$.*

The reuse of the notation C^* from Theorem 12.A.2 is intentional: one can show that the minimal cost generated by a tree on \mathcal{X} equals the minimal cost generated by a tree on \mathfrak{R} .

In applying Theorem 12.A.5, one should keep in mind that in computing the edge costs c_{RS} , one need only consider paths in the set Π_{xy} : that is, paths that only have transient states (i.e., states in $\mathcal{X} - \hat{\mathcal{X}}$) as intermediate states. While one can formulate the theorem without this constraint (see the Notes), the present formulation is typically easier to apply. To understand why, suppose that the lowest cost path from recurrent class S to recurrent class R is by way of a third recurrent class Q . Then when we apply Theorem 12.A.5, the minimal cost R -tree will not include an edge from S to R , since a lower cost tree can always be obtained by instead including the edge from S to Q . In other words, it is enough to consider “direct” paths when determining edge costs, since the advantages of any fruitful “indirect” paths are reaped when the minimal cost trees are found.

12.A.4 Radius-Coradius Theorems

Theorem 12.A.5 provides a characterization of stochastic stability in terms of minimum cost trees whose nodes are recurrent classes. In examples with few recurrent classes, this result is easy to apply, as one can find the minimum cost trees by computing the cost of every possible tree and then determining which cost is smallest. But if there are more than a few recurrent classes, exhaustive search quickly becomes cumbersome. As an alternative, one can look for sufficient conditions for stochastic stability that are easy to check, because they obviate the need to explicitly determine the minimum cost tree in the cases where they apply.

The sufficient conditions we consider rely on the notions of the radius and coradius of a recurrent class. For any pair of recurrent classes $(R, S) \in \mathfrak{R} \times \mathfrak{R}$, we define the cost c_{RS} of the transition from R to S as in equation (12.77). The *radius* of recurrent class R ,

$$\text{rad}(R) = \min_{S \neq R} c_{RS},$$

is the minimal cost of a transition from R to another recurrent class. Thus, the radius provides a simple measure of the difficulty of escaping from R .

Our sufficient conditions for stability also require a measure of the difficulty of returning to R from other recurrent classes. To define this measure, we first let Π_{SR} denote the set of (nonrepeating) paths $\pi = \{(S, Q_1), (Q_1, Q_2), \dots, (Q_{l-1}, R)\}$ through \mathfrak{R} that lead from S to R . Then the *coradius* of recurrent class R is given by

$$\text{corad}(R) = \max_{S \neq R} \min_{\pi \in \Pi_{SR}} C(\pi), \quad \text{where } C(\pi) = \sum_{(Q, Q') \in \pi} c_{QQ'}.$$

In words, the coradius of R describes the total cost of reaching R from the most disadvantageous initial class S , assuming that the route taken from S to R has as low a cost as possible.

Theorem 12.A.6 provides our first sufficient condition for stochastic stability.

Theorem 12.A.6. *Let $\{\{X_t^\epsilon\}_{t=0}^\infty\}_{\epsilon \in [0, \bar{\epsilon}]}$ be a regular collection of Markov chains with transition costs c_{xy} , let μ^ϵ be the stationary distribution of $\{X_t^\epsilon\}$, and let $\mu^* = \lim_{\epsilon \rightarrow 0} \mu^\epsilon$. If $R \in \mathfrak{R}$ is a recurrent class of $\{X_t^0\}$ with $\text{rad}(R) > \text{corad}(R)$, then $\text{support}(\mu^*) = R$.*

Theorem 12.A.6 says that if the difficulty of escaping from R , as measured by $\text{rad}(R)$, exceeds the difficulty of returning to R , as measured by $\text{corad}(R)$, then R is the set of stochastically stable states. This result follows immediately from the more general Theorem 12.A.7, whose proof will be provided below.

We obtain a more powerful sufficient condition for stochastic stability by introducing the notions of modified cost and modified coradius. The *modified cost* of edge (Q, Q') is

$$c_{QQ'}^- = c_{QQ'} - \min_{\hat{Q} \neq Q} c_{Q\hat{Q}} = c_{QQ'} - \text{rad}(Q).$$

The modified cost of path $\pi = \{(Q_0, Q_1), (Q_1, Q_2), \dots, (Q_{l-1}, Q_l)\}$ is

$$C^-(\pi) = c_{Q_0 Q_1} + \sum_{i=1}^{l-1} c_{Q_i Q_{i+1}}^-.$$

Note well the use of the original cost rather than the modified cost for the first edge in the path. Finally, the *modified coradius* of recurrent class R is

$$\text{corad}^-(R) = \max_{S \neq R} \min_{\pi \in \Pi_{SR}} C^-(\pi).$$

An intuition for these definitions is as follows: According to Theorem 12.A.5, the states in recurrent class R are stochastically stable if and only if there is a minimum cost R -tree. By definition, every R -tree τ_R must contain an edge departing from every recurrent class $Q \neq R$, and every S -tree τ_S must contain an edge departing from every recurrent class $Q \neq S$. Thus, for each $Q \notin \{R, S\}$, both τ_R and τ_S will contain an edge whose cost is at least $\text{rad}(Q)$. Since the costs represented by these radii must appear in both trees, they can be factored out when comparing the costs of τ_R and τ_S themselves. This factoring out is precisely what the notions of modified cost and modified coradius accomplish.

Theorem 12.A.7. *Let $\{\{X_t^\varepsilon\}_{t=0}^\infty\}_{\varepsilon \in [0, \bar{\varepsilon}]}$ be a regular collection of Markov chains with transition costs c_{xy} , let μ^ε be the stationary distribution of $\{X_t^\varepsilon\}$, and let $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$. If $R \in \mathfrak{R}$ is a recurrent class of $\{X_t^0\}$ with $\text{rad}(R) > \text{corad}^-(R)$, then $\text{support}(\mu^*) = R$.*

Proof. Let $S \in \mathfrak{R} - \{R\}$, and let τ_S be an S -tree. If we can show that there is an R -tree τ_R with $C(\tau_R) < C(\tau_S)$, then the result follows from Theorem 12.A.5.

Let $\pi^* = \{(S, Q_1), (Q_1, Q_2), \dots, (Q_{l-1}, R)\}$ be a path from S to R satisfying

$$\pi^* \in \underset{\pi \in \Pi_{SR}}{\text{argmin}} C^-(\pi),$$

so that $C^-(\pi^*) \leq \text{corad}^-(R)$. To construct τ_R , we start with the tree τ_S , delete each edge leading out from nodes Q_1, \dots, Q_{l-1} , and then add all edges in path π^* .

We now verify that τ_R is an R -tree. Clearly, τ_R has exactly one outgoing edge from each $Q \neq R$ and no outgoing edges from R . To verify that there is path from each $Q \neq R$ to

R , it is enough to observe that the sequence of edges leading from Q through subsequent nodes must hit path π . If it did not, it would enter a cycle consisting entirely of nodes outside of path π . The edges in this cycle would then all be in τ_S , contradicting that τ_S is a tree.

Now, since τ_R and τ_S only differ in their outgoing edges from S, R , and Q_1, \dots, Q_{l-1} , and since $C^-(\pi^*) \leq \text{corad}^-(R)$, we find that

$$\begin{aligned} c(\tau_R) - c(\tau_S) &\leq C(\pi^*) - \left(\sum_{i=1}^{l-1} \text{rad}(Q_i) + \text{rad}(R) \right) \\ &= \left(C^-(\pi^*) + \sum_{i=1}^{l-1} \text{rad}(Q_i) \right) - \left(\sum_{i=1}^{l-1} \text{rad}(Q_i) + \text{rad}(R) \right) \\ &\leq \text{corad}^-(R) - \text{rad}(R) \\ &< 0. \end{aligned}$$

Since we assumed at the outset that $\text{rad}(R) > \text{corad}^-(R)$, the proof is complete. ■

It can be shown that when $\text{rad}(R) > \text{corad}^-(R)$, the expected time to reach R from any other recurrent class is of asymptotic order no greater than $\varepsilon^{-\text{corad}^-(R)}$; see the Notes for references.

12.A.5 Lenient Transition Costs and Weak Stochastic Stability

In Section 12.A.2, we defined regular Markov chains to be those whose vanishing transition probabilities satisfy condition (12.71), a condition that can be rewritten as

$$(12.78) \quad P_{xy}^\varepsilon = (a_{xy} + o(1)) \varepsilon^k.$$

Here $k \geq 0$ is the cost of the transition, a_{xy} is a positive constant, and $o(1)$ is a term that vanishes as ε approaches zero.

In some settings condition (12.78) is too strong to apply, making it useful to replace it with the following more lenient requirement:

$$(12.79) \quad P_{xy}^\varepsilon = \varepsilon^{k+o(1)}.$$

It is easy to verify that condition (12.78) implies condition (12.79), but that the converse implication does not hold—see Example 12.A.8 below.

Condition (12.79) is more commonly expressed after the change of parameter $\varepsilon \mapsto \eta \equiv$

$-(\log \varepsilon)^{-1}$, so that the noise level η appears in an exponent. In this formulation, (12.78) becomes

$$(12.80) \quad P_{xy}^\eta = (a_{xy} + o(1)) \exp(-\eta^{-1}k),$$

where $o(1)$ now represents a term that vanishes as η approaches zero. For its part, equation (12.79) becomes

$$(12.81) \quad P_{xy}^\eta = \exp(-\eta^{-1}(k + o(1))),$$

which we can express equivalently as

$$(12.82) \quad -\lim_{\eta \rightarrow 0} \eta \log P_{xy}^\eta = k.$$

The analysis in Section 12.3 shows that the lenient definition (12.82) of transition costs is the natural one for studying the small-noise asymptotics of the stationary distribution. The lenient definition of transition costs also broadens the class of noisy best response rules that are susceptible to analysis.

Example 12.A.8. Probit choice. In Example 12.3.3, agents' decisions were based on *probit choice*: revising agents choose an optimal response after the payoffs to each strategy are subject to i.i.d. normal perturbations with mean zero and variance η . Suppose that at state x in a two-strategy game, the payoff to strategy i exceeds that of strategy j by $d > 0$. Then the probability of a transition from state x to state $y = x + \frac{1}{N}(e_j - e_i)$ under probit choice is

$$P_{xy}^\eta = x_i \cdot k^\eta(d) \exp\left(-\frac{d^2}{4\eta}\right) \text{ for some } k^\eta(d) \in \left(\frac{\sqrt{\eta}}{\sqrt{\pi}d} \left(1 - \frac{2\eta}{d^2}\right), \frac{\sqrt{\eta}}{\sqrt{\pi}d}\right).$$

This transition probability does not satisfy (12.80), but it satisfies (12.81) with transition cost $k = \frac{1}{4}d^2$. §

We now give a careful account of the implications of allowing lenient transition costs. Let $\{\{X_t^\eta\}_{t=0}^\infty\}_{\eta \in [0, \bar{\eta}]}$ be a parameterized collection of Markov chains on the finite state space \mathcal{X} . Suppose that the chain $\{X_t^\eta\}$ is irreducible when η is positive, and that the transition matrices $P^\eta \in \mathbf{R}_+^{X \times X}$ vary continuously in η . Finally, suppose that for each pair of distinct states x and $y \neq x$, if $P_{xy}^0 = 0$ and $P_{xy}^{\hat{\eta}} > 0$ for some $\hat{\eta} > 0$, then the sequence $\{P_{xy}^\eta\}_{\eta \in (0, \bar{\eta}]}$ satisfies (12.82) for some $k \geq 0$. If these conditions are met, we call the collection of Markov chains *weakly regular*, and we define the (*lenient*) *cost* of the transition from state x to

state $y \neq x$ by

$$(12.83) \quad c_{xy} = -\lim_{\eta \rightarrow 0} \eta \log P_{xy}^\eta,$$

adopting the convention that $-\log 0 = \infty$.

One can prove analogues of the results from Sections 12.A.2–12.A.4 using lenient transition costs, but at the cost of losing some borderline selection results. If we recreate the proof of Theorem 12.A.2 with transition costs defined as in equation (12.83), then in place of equation (12.73) we obtain

$$(12.84) \quad \mu_x^\eta = \frac{\exp(-\eta^{-1}(C_x^* + o(1)))}{\sum_{y \in X} \exp(-\eta^{-1}(C_y^* + o(1)))}.$$

It does not follow from this that a state x with $C_x^* = C^*$ is stochastically stable, in the sense that $\lim_{\eta \rightarrow 0} \mu_x^\eta > 0$. But equation (12.84) does imply that for all x and y in X , we have

$$-\lim_{\eta \rightarrow 0} \eta \log \frac{\mu_x^\eta}{\mu_y^\eta} = C_x^* - C_y^*;$$

then using the argument used to prove Theorem 12.2.2, we conclude that

$$(12.85) \quad -\lim_{\eta \rightarrow 0} \eta \log \mu_x^\eta = C_x^* - C^*.$$

We call state $x \in X$ *weakly stochastically stable* if the limit in (12.85) is equal to zero, or, equivalently, if as η^{-1} approaches infinity, μ_x^η does not vanish at an exponential rate. Equation (12.85) shows that x is weakly stochastically stable if and only if $C_x^* = C^*$.

Sections 12.A.2–12.A.4 presented a variety of sufficient conditions for stochastic stability in regular collections of Markov chains. The discussion above implies that in each case, there is a corresponding sufficient condition for weak stochastic stability in weakly regular collections of Markov chains, so long as we define transition costs using (12.83) rather than (12.72). In general, the limiting stationary distribution of such a collection need not exist. But Proposition 12.1.2 shows that if there is just one weakly stochastically stable state, then the limiting stationary distribution exists and is a point mass at this state; in the parlance of Section 12.1.1, such a state is *uniquely stochastically stable*.

This argument can be extended to obtain conditions for the selection of a unique recurrent class in a weakly regular collection of Markov chains. In particular, if there is a unique recurrent class R of the zero-noise chain that satisfies a sufficient condition for

stochastic stability from Theorem 12.A.5, 12.A.6, or 12.A.7, then R is the set of stochastically stable states. For instance, if $\text{rad}(R) > \text{corad}^-(R)$, with costs defined as in (12.83), then R is the set of stochastically stable states.

12.B Stochastic Approximation Theory

Stochastic approximation theory studies stochastic processes whose increments are small, and whose expected increments can be described by a mean dynamic. The basic aim of the theory is to link the infinite horizon behavior of the stochastic processes to various sorts of recurrent sets of the deterministic mean dynamic.

The theory distinguishes between two main settings. In the case of *decreasing step sizes*, one considers a single stochastic process whose increments become smaller as time passes. This is the case when the state variable represents a time-averaged quantity, or, alternatively, the proportions of balls of different colors in an urn to which balls are added in each period. A fundamental result from this theory shows that the sample paths of the stochastic process converge to almost surely to components of the *chain recurrent set* of the mean dynamic (see Appendix 7.A.1). See the Notes for references and further discussion.

In the case of *constant step sizes*, one considers a sequence of stochastic processes indexed by a parameter that determines the size of the increments of the process. Here the goal of the theory is to link the stationary distributions of the processes with small step sizes to the recurrent states of the mean dynamic (M).

The remainder of this appendix presents results from the constant step size theory. In Appendix 12.B.1, we show that under quite general conditions, when the step size is sufficiently small, the stationary distributions of the stochastic processes concentrate their mass near recurrent states of (M). If the processes exhibit enough random variation, then some recurrent states—for instance, unstable rest points—would seem to be unlikely candidates for retaining stationary distribution mass. In Appendix 12.B.2, we describe conditions on the mean dynamic and the stochastic processes themselves ensuring that mass concentrates only on stable rest points of (M).

12.B.1 Convergence to the Birkhoff Center

We follow the development of the deterministic approximation results in Section 10.2. For each $N \geq N_0$, corresponding to a step size of $\frac{1}{N}$, we suppose that the process $\{X_t^N\}_{t \geq 0}$ takes values in the state space $\mathcal{X}^N = \{x \in X : Nx \in \mathbf{Z}^n\}$. We let $\lambda^N \in \mathbf{R}_+^{\mathcal{X}^N}$ and $p^N \in \mathbf{R}_+^{\mathcal{X}^N \times \mathcal{X}^N}$ denote the jump rate vector and transition matrix of this process. We let ζ_x^N be a random

variable whose distribution describes the stochastic increment of $\{X_t^N\}$ from state x :

$$\mathbb{P}(\zeta_x^N = z) = \mathbb{P}\left(X_{\tau_{k+1}}^N = x + z \mid X_{\tau_k}^N = x\right),$$

where τ_k is the time of the process's k th jump. The expected increment per time unit of the process $\{X_t^N\}$ is described by $V^N: X \rightarrow TX$, defined by $V^N(x) = \lambda_x^N \mathbb{E}\zeta_x^N$.

As in Theorem 10.2.1, we assume that there is a Lipschitz continuous vector field $V: X \rightarrow TX$ such that the functions V^N converge uniformly to V :

$$\lim_{N \rightarrow \infty} \sup_{x \in \mathcal{X}^N} |V^N(x) - V(x)| = 0.$$

In addition, we assume that there is a constant $c < \infty$ such that $|\zeta_x^N| \leq \frac{c}{N}$ with probability one, which implies that the remaining conditions of Theorem 10.2.1 hold. Therefore, if we consider initial conditions $X_0^N = x_0^N$ that converge to $x_0 \in X$, and let $\{x_t\}_{t \geq 0}$ denote the solution to the mean dynamic

$$(M) \quad \dot{x} = V(x),$$

then for each $T < \infty$ and $\varepsilon > 0$, we have that

$$(12.86) \quad \lim_{N \rightarrow \infty} \mathbb{P}\left(\sup_{t \in [0, T]} |X_t^N - x_t| < \varepsilon\right) = 1.$$

Theorem 10.2.1 is a statement about the *finite horizon* behavior: the time span $T < \infty$ is fixed in advance, and (12.86) tells us that when N is large enough, the stochastic process $\{X_t^N\}$ is very likely to mirror the deterministic trajectory through time T .

In order to state an infinite-horizon approximation result, we must review a few definitions from Appendix 7.A.1. Let $\{x_t\} = \{x_t\}_{t \geq 0}$ be a solution trajectory of the mean dynamic (M). The ω -limit of $\{x_t\}$ is the set of all points that the trajectory approaches arbitrarily closely infinitely often:

$$(12.87) \quad \omega(x_0) = \left\{ y \in X : \text{there exists } \{t_k\}_{k=1}^{\infty} \text{ with } \lim_{k \rightarrow \infty} t_k = \infty \text{ such that } \lim_{k \rightarrow \infty} x_{t_k} = y \right\}.$$

State $x \in X$ is a *recurrent point* of (M) ($x \in \mathcal{R}$) if $x \in \omega(x)$. The *Birkhoff center* of (M) is the closure $\text{cl}(\mathcal{R})$ of the set of recurrent points.

Since we have not assumed that the processes $\{X_t^N\}$ are irreducible, each of these process may admit multiple stationary distributions. Theorem 12.B.1 says that if we fix any sequence $\{\mu^N\}_{N=N_0}^{\infty}$ of stationary distributions, then all stationary distributions far

enough along the sequence place most of their mass near the Birkhoff center of (M).

Theorem 12.B.1. *Let $\{\mu^N\}_{N=N_0}^\infty$ be any sequence of stationary distributions for the Markov processes $\{X_t^N\}_{N=N_0}^\infty$, and let $O \subseteq X$ be an open set containing $\text{cl}(\mathcal{R})$. Then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.*

Theorem 12.B.1 has strong implications in settings in which the mean dynamic has a unique recurrent point. Recall that the *point mass* at state $x \in X$, denoted δ_x , is the probability measure on X that satisfies $\delta_x(A) = 1$ if $x \in A$ and $\delta_x(A) = 0$ otherwise. Also, we say that the sequence of measures $\{\mu^N\}_{N=N_0}^\infty$ *converges weakly* to the point mass δ_x if $\lim_{N \rightarrow \infty} \mu^N(O) = 1$ for every (relatively) open set $O \subseteq X$ containing x (cf Exercise 12.1.4).

Corollary 12.B.2. *Let $\{\mu^N\}_{N=N_0}^\infty$ be any sequence of stationary distributions for the Markov processes $\{X_t^N\}_{N=N_0}^\infty$, and suppose that the mean dynamic (M) admits a unique recurrent point, x^* . Then the sequence $\{\mu^N\}$ converges weakly to the point mass δ_{x^*} .*

Under additional regularity assumptions, it can be shown that if the state space is dilated by a factor of \sqrt{N} , then the stationary distributions μ^N converge to a multivariate normal distribution centered at x^* (see the Notes).

12.B.2 Sufficient Conditions for Convergence to Stable Rest Points

Theorem 12.B.1 does not rule out the possibility that these stationary distributions μ^N concentrate their mass near unstable rest points or other unstable invariant sets of the mean dynamic. In this section, we present conditions under which the stationary distributions must become concentrate their mass near stable rest points of (M).

The results to follow require additional structure on the Markov processes $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$. We say that this collection has *homogeneous transitions* if the distribution of the rescaled stochastic increment $\lambda_x^N \zeta_x^N$ is independent of N . In this case, we can let ζ_x denote a random variable with the same distribution as $\lambda_x^N \zeta_x^N$. It follows that the expected increment per time unit $V^N(x) = \lambda_x^N \mathbb{E} \zeta_x^N = V(x)$ is independent of N . We maintain our assumption from Section 12.B.1 that V is Lipschitz continuous.

To rule out convergence to unstable rest points, we must assume in addition that the process $\{X_t^N\}$ exhibits enough random motion near these points. The next example illustrates this idea.

Example 12.B.3. Consider a birth and death processes $\{X_t^N\}$ on $\mathcal{X}^N = \{0, \frac{1}{N}, \dots, 1\}$ with common jump rate $\lambda_x^N \equiv N$ and transition probabilities

$$p_x^N = p_x = (1 - \chi)(1 - (1 - \chi)^2) \text{ and}$$

$$q_\chi^N = q_\chi = \chi(1 - \chi^2).$$

An interpretation of this process is as follows: when a jump time occurs, one agent from the population is selected at random. This agent then samples two members of the population. If at least one of them is playing the other strategy, the agent switches to that strategy.

Since the product of the jump rate ($\lambda_\chi^N = N$) and the distance between states ($\frac{1}{N}$) is 1, the mean dynamic for this process for any N is

$$(12.88) \quad \dot{\chi} = p_\chi - q_\chi = \chi(1 - \chi)(1 - 2\chi).$$

Under this dynamic, the rest point $\chi^* = \frac{1}{2}$ attracts all initial conditions in $(0, 1)$, and the boundary rest points 0 and 1 are unstable. Nevertheless, the only recurrent states of the process are states 0 and 1. Thus, for any value of N , every stationary distribution μ^N of the process $\{X_t^N\}$ places all of its mass on unstable rest points of the mean dynamic (12.88). §

The example shows that if x is a rest point of the mean dynamic (M) at which the process $\{X_t^N\}$ stops moving (that is, if $P(\zeta_x = \mathbf{0}) = 1$), then the point mass δ_x is a stationary distribution of $\{X_t^N\}$ even if x is unstable under (M). But while x being a rest point of (M) means that the *expected* increment $\mathbb{E}\zeta_x^N$ is null, the actual increment ζ_x^N need not be. As long as the process $\{X_t^N\}$ churns enough to wander slightly away from the rest point x , the force of the mean dynamic will push it still further away, and will make it very difficult for the process to return.

To formalize this idea, we say that the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ is *nondegenerate* at rest point x if $\text{span}(\text{support}(\zeta_x)) = TX$, so that the directions of random motions around the rest point x span the tangent space of X . Notice that nondegeneracy is only possible for interior rest points: if from a boundary state there is a positive probability of motion into the interior, this state cannot be a rest point.

Even if an unstable rest point is nondegenerate, it is still possible for the stationary distribution to become concentrated around this point. While the following example does not satisfy our standing assumptions about the form of the state space or the homogeneity of the processes, it illustrates our point in a simple way.

Example 12.B.4. Consider a dynamic $\dot{x} = V(x)$ on the unit circle that moves clockwise except at a single rest point x^* (cf Example 7.A.3), and let $\{X_t^N\}_{N=N_0}^\infty$ be a sequence of Markov processes that are defined on “grids” in the unit circle, and whose mean dynamics V^N converge to V . If the process $\{X_t^N\}$ is nondegenerate at x^* , then churning ensures that whenever the process reaches x^* , it eventually wanders a nonnegligible distance past

x^* . Once this happens, the process resumes its clockwise motion around the circle until returning to a neighborhood of x^* , where the churning then resumes. Since x^* is the sole recurrent point of the mean dynamic, Theorem 12.B.1 tells us that the stationary distributions μ^N converge weakly to the point mass δ_{x^*} as N approaches infinity. §

Ruling out examples of this sort requires assumptions about the global behavior of the mean dynamic, assumptions ensuring that there is no easy way for the Markov process to return to an unstable rest point. One possibility is to require the mean dynamic to generate a so-called simple flow. As a partner to the ω -limit set $\omega(x_0)$ for initial condition $x_0 \in X$ under the mean dynamic $\dot{x} = V(x)$, define the α -limit set $\alpha(x_0)$ to be the ω -limit set of the time-reversed equation $\dot{x} = -V(x)$. The set $\alpha(x_0)$ is nonempty only if the solution of the mean dynamic from x_0 exists in X for all backward time. We then say that (M) generates a *simple flow* if the set $\bigcup_{x \in X} (\alpha(x) \cup \omega(x))$ of all α - and ω -limit points of (M) is finite. In a simple flow, each solution trajectory converges to one of a finite number of rest points, as does each backward time solution that does not leave X . Simple flows do not contain closed orbits, limit cycles, chaotic attractors, or other complicated limit behaviors.

To formalize the requirement that there be no easy way for the Markov process to return to the rest point x^* after leaving its vicinity, we define an *orbit chain* to be sequence $\{x^k\}_{k=1}^K$ of rest points of (M) with connecting forward orbits $\{\gamma_k\}_{k=1}^{K-1}$, meaning that $x^k = \alpha(\gamma^k)$ and $x^{k+1} = \omega(\gamma^k)$ for all $k \in \{1, \dots, K-1\}$. We say that the rest point x^* is *orbit chain unstable* if there is an orbit chain leading from x^* to another rest point y^* and no orbit chain leading back. Otherwise, we call x^* *orbit chain stable*.

With these definitions in hand, we can state the following theorem:

Theorem 12.B.5. *Suppose that the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ has homogenous transitions and is nondegenerate at every rest point of (M) that is not Lyapunov stable. Suppose further that (M) generates a simple flow on X . If $\{\mu^N\}$ is a sequence of stationary distributions for the Markov chains $\{X_t^N\}$, and if $O \subseteq X$ is an open set containing the orbit chain stable equilibria of (M), then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.*

For an application of this result, recall from Section 7.1 that the C^1 function $L : X \rightarrow \mathbf{R}$ is a *strict Lyapunov function* for V if $\dot{L}(x) \equiv \nabla L(x) \cdot V(x) \geq 0$ for all $x \in X$, with equality only at rest points of V . It is easy to see that any dynamic with a finite number of rest points that admits a strict Lyapunov function generates a simple flow. This observation and Theorem 12.B.5 yield the following result.

Theorem 12.B.6. *Suppose that the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ has homogenous transitions and that it is nondegenerate at every rest point of the mean dynamic (M) that is not Lyapunov stable. In addition, suppose that $RP(V)$ is finite and that (M) admits a strict Lyapunov function $L : X \rightarrow \mathbf{R}$.*

Let $\{\mu^N\}_{N=N_0}^\infty$ be a sequence of stationary distributions for the Markov chains, and let $O \subseteq X$ be an open set containing all Lyapunov stable rest points of (M). Then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.

Since in this theorem the dynamic (M) has a finite number of rest points and admits a strict Lyapunov function, its Lyapunov stable rest points and its asymptotically stable rest points are identical.

One can also prove that the stationary distributions μ^N concentrate their mass near stable rest points of (M) when (M) is cooperative and irreducible. Since such dynamics can admit complicated unstable invariant sets, a global form of nondegeneracy is needed to reach the desired conclusion. For each state $x \in X$, define $\Sigma_x = \mathbb{E}\zeta_x\zeta_x' - \mathbb{E}\zeta_x\mathbb{E}\zeta_x' \in \mathbf{R}^{n \times n}$ to be the covariance matrix generated by the random vector ζ_x . We call the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ globally nondegenerate if there is a $c > 0$ such that $\text{Var}(z'\zeta_x) \equiv z'\Sigma_x z \geq c$ for all unit-length $z \in TX$ and all $x \in X$.

Example 12.B.7 verifies that stochastic evolutionary processes generated by perturbed best response protocols satisfy global nondegeneracy.

Example 12.B.7. Let $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ be the collection of stochastic evolutionary processes generated by the game F and the protocol $\rho(\pi, x) = \tilde{M}(x)$, where $\tilde{M}: \mathbf{R}^n \rightarrow \text{int}(X)$ is a continuous perturbed maximizer function. If we write $\tilde{B} \equiv \tilde{M} \circ F$, then the distribution of the normalized increment ζ_x can be expressed as

$$\mathbb{P}(\zeta_x = z) = \begin{cases} x_i \tilde{B}_j(x) & \text{if } z = e_j - e_i, \\ \sum_{i \in S} x_i \tilde{B}_i(x) & \text{if } z = \mathbf{0}, \\ 0 & \text{otherwise.} \end{cases}$$

It follows that

$$\mathbb{E}\zeta_x = \tilde{B}(x) - x \quad \text{and} \quad \mathbb{E}\zeta_x\zeta_x' = \text{diag}(\tilde{B}(x)) + \text{diag}(x) - \tilde{B}(x)x' - x\tilde{B}(x)',$$

and hence that

$$\Sigma_x = \mathbb{E}\zeta_x\zeta_x' - \mathbb{E}\zeta_x\mathbb{E}\zeta_x' = \text{diag}(\tilde{B}(x)) - \tilde{B}(x)\tilde{B}(x)' + \text{diag}(x) - xx'.$$

Now let

$$b = \min_{x \in X} \min_{i \in S} \tilde{B}_i(x).$$

Since X is compact and F and \tilde{M} are continuous, b is positive. Moreover, if z is a unit-length

vector in TX , we have that

$$\begin{aligned}
z' \Sigma_x z &= \sum_{i \in S} (z_i)^2 \tilde{B}_i(x) - \left(\sum_{i \in S} z_i \tilde{B}_i(x) \right)^2 + \sum_{i \in S} (z_i)^2 x_i - \left(\sum_{i \in S} z_i x_i \right)^2 \\
&= \sum_{i \in S} \tilde{B}_i(x) \left(z_i - \sum_{j \in S} z_j \tilde{B}_j(x) \right)^2 + \sum_{i \in S} x_i \left(z_i - \sum_{j \in S} z_j x_j \right)^2 \\
&\geq b \sum_{i \in S} \left(z_i - \sum_{j \in S} z_j \tilde{B}_j(x) \right)^2 \\
&= b \left(\sum_{i \in S} (z_i)^2 + n \left(\sum_{j \in S} z_j \tilde{B}_j(x) \right)^2 \right) \\
&> b.
\end{aligned}$$

Thus the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ is globally nondegenerate. §

With this background established, we present our final stochastic approximation result.

Theorem 12.B.8. *Suppose that the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$ has homogenous transitions and is globally nondegenerate, and suppose that the mean dynamic (M) is affinely conjugate to a cooperative, irreducible differential equation. Then if $\{\mu^N\}_{N=N_0}^\infty$ is a sequence of stationary distributions of the Markov chains $\{X_t^N\}$, and if $O \subseteq X$ is an open set containing all Lyapunov stable rest points of (M), then $\lim_{N \rightarrow \infty} \mu^N(O) = 1$.*

12.N Notes

Section 12.1. Stochastic stability analysis was introduced to game theory in papers by Foster and Young (1990), Kandori et al. (1993), and Young (1993a), which define stochastic stability using the small noise limit. Stochastic stability in the large population limit and in the double limits is studied by Binmore et al. (1995b), Binmore and Samuelson (1997), Benaïm and Weibull (2000), and Sandholm (2009e,f,b). Example 12.1.9 is essentially due to Binmore and Samuelson (1997).

Section 12.2. This section follows Sandholm (2009b), which builds on earlier work of Blume (1993, 1997). For the exact form of Stirling's formula used in the proof of Theorem 12.2.3, see Lang (1997). For recent work on evolution under variations on the logit rule in normal form games, see Alós-Ferrer and Netzer (2008), Marden and Shamma (2008), and Okada and Tercieux (2009). For further results on logit evolution in large population potential games, see Benaïm and Sandholm (2007).

Section 12.3. Sections 12.3.1–12.3.6 follow Sandholm (2009e), which builds on earlier papers of Blume (2003) and Sandholm (2007c); also see Maruta (2002). The definition

of noisy best response protocols is essentially due to Blume (2003). Evolutionary game models of probit choice include Ui (1998), Myatt and Wallace (2003), and Dokumacı and Sandholm (2007b). For the approximation of normal distribution tail probabilities used in Example 12.3.3, see Durrett (2005, Theorem 1.1.3). See Billingsley (1995) or Folland (1999) for the Dominated Convergence Theorem, and see the Appendix of Sandholm (2009e) for details on how to apply it in the proofs of Theorems 12.3.6 and 12.3.11.

Ordinal potential functions are introduced in the context of normal form games by Monderer and Shapley (1996b). The need to account for both “widths” and “depths” in stochastic stability analyses is emphasized by Fudenberg and Harris (1992), Binmore and Samuelson (1997), and Kandori (1997). Corollary 12.3.16 is due to Blume (2003), and Corollary 12.3.17 is essentially due to Kandori et al. (1993). For more on stochastic dominance in utility theory, see Border (2001).

Section 12.4. This section follows Sandholm (2009f). That paper builds on the work of Binmore et al. (1995b) and Binmore and Samuelson (1997), who introduce multiple limits to the study of stochastic stability, and who establish a version of Theorem 12.4.1.

For results on equilibrium selection in n -strategy games under imitative protocols with mutations, and on equilibrium selection under the frequency-dependent Moran process, see Fudenberg and Imhof (2006, 2008). For applications of the latter results, see Nowak et al. (2004), Taylor et al. (2004), and Fudenberg et al. (2006), as well as surveys by Nowak (2006) and Traulsen and Hauert (2009).

Binmore and Samuelson (1997) obtain equilibrium selection results along the lines described at the end of this section for imitative protocols driven by dissatisfaction, providing conditions under which an equilibrium that is both risk dominant and payoff dominant is stochastically stable. Maruta (2002) characterizes the effects of different forms of payoff dependence in direct revision protocols on stochastic stability results.

Section 12.5. The use of minimal cost trees to determine small noise stochastically stable states was introduced to game theory by Kandori et al. (1993) and Young (1993a), building on ideas developed in Freidlin and Wentzell (1998). These papers and most of their successors focus on the BRM model; exceptions include Ui (1998), Myatt and Wallace (2003), Alós-Ferrer and Netzer (2008), Marden and Shamma (2008), and Dokumacı and Sandholm (2007b).

Selection of risk dominant equilibrium in the BRM model is established in Kandori et al. (1993); the risk dominance concept itself is due to Harsanyi and Selten (1988). The importance of clever agents for obtaining exact finite population selection results is noted in Sandholm (1998), see also Rhode and Stegeman (1996). Robson and Vega-Redondo (1996) and Bergin and Lipman (1996) study alternate noisy best response models in which

Kandori et al.'s (1993) selection result fails to hold. Robles (1998) and Sandholm and Pauzner (1998) show that decreasing mutation rates and growing populations can abrogate history-independent predictions in stochastic evolutionary models.

The Radius-Coradius Theorem is due to Ellison (2000). The notion of half-dominance first appears in Morris et al. (1995). Exercise 12.5.15 is due to Kandori and Rob (1998). Theorem 12.5.16 is due to Maruta (1997), Kandori and Rob (1998), and Ellison (2000). Algorithms for determining the stochastically stable state in supermodular games are presented in Kandori and Rob (1995) and Beggs (2005).

There is a considerable body of work on stochastic stability in extensive form games that we do not consider here. See the Notes to the Preface for references.

Section 12.6. Theorem 12.6.2 is due to Benaïm (1998) and Benaïm and Weibull (2009). Theorem 12.6.4 is taken from Hofbauer and Sandholm (2007).

The evaluation of the probabilities of transitions between stable rest points of the mean dynamic in the large population limit falls is part of the study of sample path large deviations, and is considered by Freidlin (1978), Azencott and Ruget (1977), and Dupuis (1988); books with related results include Freidlin and Wentzell (1998) and Dembo and Zeitouni (1998). These transition probabilities are evaluated in the two-strategy case by Benaïm and Weibull (2000), and in the case of potential games under the logit dynamic by Benaïm and Sandholm (2007).

Appendix 12.A. Theorem 12.A.1 is an old and often-resdiscovered result from Markov chain theory; see the notes to Chapter 9 of Aldous and Fill (2001) for some historical remarks. This theorem was first used to compute limiting stationary distributions by Freidlin and Wentzell (1998) in the context of stochastically perturbed differential equations. The results presented in Section 12.A.2 for Markov chains on finite state spaces approach are due to Kandori et al. (1993) and Young (1993a).

Young (1993a, 1998b) introduces the notion of a regular collection of Markov chains that we employ here, and also provides characterizations of limiting stationary distributions by way of trees on recurrent classes, as considered in Section 12.A.3. The analogue of Theorem 12.A.5 stated in Young (1993a, 1998b) requires one to consider “indirect” paths between recurrent classes (see the discussion after Theorem 12.A.5), and is proved using a “tree-surgery” argument that is rather different than the argument employed here. The statement of Theorem 12.A.5 we provide, which uses only “direct” paths, is proposed in Kandori and Rob (1995).

The radius-coradius theorems in Section 12.A.4 are modifications of results of Ellison (2000), whose analysis also considers waiting times for transitions between recurrent classes. More refined result on waiting times can be found in Beggs (2005).

Appendix 12.B. Basic references on stochastic approximation theory with decreasing step sizes and fixed step sizes are Benaïm (1999) and Benaïm (1998), respectively. The decreasing step size results are the basic tool for studying the model of learning in games known as *stochastic fictitious play*, in which the state variable describes the time-averaged play of n players who recurrently play an n -player normal form game. References include Fudenberg and Kreps (1993), Kaniovski and Young (1995), Benaïm and Hirsch (1999a), Hofbauer and Sandholm (2002), Benaïm et al. (2005, 2006), and Benaïm et al. (2009). Pemantle (2007) surveys results on this and related processes.

Theorem 12.B.1 follows from Corollary 3.2 of Benaïm (1998). For asymptotic normality, see Kurtz (1976, Theorem 2.7) and Ethier and Kurtz (1986, Section 11.2). Theorem 12.B.5 follows from Theorem 4.3 of Benaïm (1998), and Theorem 12.B.8 from Theorem 1.5 of Benaïm and Hirsch (1999b).

BIBLIOGRAPHY

- Abraham, R. and Robbin, J. (1967). *Transversal Mappings and Flows*. W. A. Benjamin, New York.
- Abreu, D. and Sethi, R. (2003). Evolutionary stability in a reputational model of bargaining. *Games and Economic Behavior*, 44:195–216.
- Agastya, M. (2004). Stochastic stability in a double auction. *Games and Economic Behavior*, 48:203–222.
- Akin, E. (1979). *The Geometry of Population Genetics*. Springer, Berlin.
- Akin, E. (1980). Domination or equilibrium. *Mathematical Biosciences*, 50:239–250.
- Akin, E. (1990). The differential geometry of population genetics and evolutionary games. In Lessard, S., editor, *Mathematical and Statistical Developments of Evolutionary Theory*, pages 1–93. Kluwer, Dordrecht.
- Akin, E. (1993). *The General Topology of Dynamical Systems*. American Mathematical Society, Providence, RI.
- Akin, E. and Losert, V. (1984). Evolutionary dynamics of zero-sum games. *Journal of Mathematical Biology*, 20:231–258.
- Aldous, D. and Fill, J. A. (2001). Reversible Markov chains and random walks on graphs. Unpublished manuscript, UC Berkeley and Johns Hopkins University.
- Alós-Ferrer, C. and Ania, A. B. (2005). The evolutionary stability of perfectly competitive behavior. *Economic Theory*, 26:497–516.
- Alós-Ferrer, C., Ania, A. B., and Schenk-Hoppé, K. R. (2000). An evolutionary model of Bertrand oligopoly. *Games and Economic Behavior*, 33:1–19.
- Alós-Ferrer, C., Kirchsteiger, G., and Walzl, M. (2006). On the evolution of market institutions: The platform design paradox. *Economic Journal*, forthcoming.
- Alós-Ferrer, C. and Netzer, N. (2008). The logit response dynamics. *Games and Economic Behavior*, forthcoming.

- Alós-Ferrer, C. and Weidenholzer, S. (2006). Imitation, local interactions, and efficiency. *Economics Letters*, 93:163–168.
- Alós-Ferrer, C. and Weidenholzer, S. (2007). Partial bandwagon effects and local interactions. *Games and Economic Behavior*, 61:179–197.
- Alós-Ferrer, C. and Weidenholzer, S. (2008). Contagion and efficiency. *Journal of Economic Theory*, 143:251–274.
- Anderlini, L. and Ianni, A. (1996). Path dependence and learning from neighbors. *Games and Economic Behavior*, 13:141–177.
- Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge.
- Ania, A. B., Tröger, T., and Wambach, A. (2002). An evolutionary analysis of insurance markets with adverse selection. *Games and Economic Behavior*, 40:153–184.
- Arneodo, A., Couillet, P., and Tresser, C. (1980). Occurrence of strange attractors in three-dimensional Volterra equations. *Physics Letters*, 79A:259–263.
- Arslan, G., Marden, J. R., and Shamma, J. S. (2007). Autonomous vehicle target assignment: A game-theoretical formulation. *ASME Journal of Dynamic Systems, Measurement and Control*, 129:584–596.
- Aubin, J.-P. (1991). *Viability Theory*. Birkhäuser, Boston.
- Aubin, J.-P. and Cellina, A. (1984). *Differential Inclusions*. Springer, Berlin.
- Avriel, M. (1976). *Nonlinear Programming: Analysis and Methods*. Prentice-Hall, Englewood Cliffs, NJ.
- Azencott, R. and Ruget, G. (1977). Mélanges d'équations différentielles et grands écarts à la loi des grandes nombres. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 38:1–54.
- Balder, E. J. (2002). A unifying pair of Cournot–Nash equilibrium existence results. *Journal of Economic Theory*, 102:437–470.
- Balkenborg, D., Kuzmics, C., and Hofbauer, J. (2008). Refined best-response correspondence and dynamics. Unpublished manuscript, University of Exeter, Northwestern University, and University of Vienna.
- Balkenborg, D. and Schlag, K. H. (2001). Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595.
- Balkenborg, D. and Schlag, K. H. (2007). On the evolutionary selection of sets of Nash equilibria. *Journal of Economic Theory*, 133:295–315.

- Banerjee, A. and Weibull, J. W. (2000). Neutrally stable outcomes in cheap-talk coordination games. *Games and Economic Behavior*, 32:1–24.
- Basu, K. and Weibull, J. W. (1991). Strategy sets closed under rational behavior. *Economics Letters*, 36:141–146.
- Beckmann, M., McGuire, C. B., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven.
- Beggs, A. W. (2005). Waiting times and equilibrium selection. *Economic Theory*, 25:599–628.
- Ben-Shoham, A., Serrano, R., and Volij, O. (2004). The evolution of exchange. *Journal of Economic Theory*, 114:310–328.
- Benaïm, M. (1998). Recursive algorithms, urn processes, and the chaining number of chain recurrent sets. *Ergodic Theory and Dynamical Systems*, 18:53–87.
- Benaïm, M. (1999). Dynamics of stochastic approximation algorithms. In Azéma, J. et al., editors, *Séminaire de Probabilités XXXIII*, pages 1–68. Springer, Berlin.
- Benaïm, M. and Hirsch, M. W. (1999a). Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior*, 29:36–72.
- Benaïm, M. and Hirsch, M. W. (1999b). On stochastic approximation algorithms with constant step size whose average is cooperative. *Annals of Applied Probability*, 30:850–869.
- Benaïm, M., Hofbauer, J., and Hopkins, E. (2009). Learning in games with unstable equilibria. *Journal of Economic Theory*, 144:1694–1709.
- Benaïm, M., Hofbauer, J., and Sandholm, W. H. (2008). Robust permanence and impermanence for stochastic replicator dynamics. *Journal of Biological Dynamics*, 2:180–195.
- Benaïm, M., Hofbauer, J., and Sorin, S. (2005). Stochastic approximation and differential inclusions. *SIAM Journal on Control and Optimization*, 44:328–348.
- Benaïm, M., Hofbauer, J., and Sorin, S. (2006). Stochastic approximation and differential inclusions II: Applications. *Mathematics of Operations Research*, 31:673–695.
- Benaïm, M. and Sandholm, W. H. (2007). Logit evolution in potential games: Reversibility, rates of convergence, large deviations, and equilibrium selection. Unpublished manuscript, Université de Neuchâtel and University of Wisconsin.
- Benaïm, M. and Weibull, J. W. (2000). Deterministic approximation of stochastic evolution in games. Working paper 534, Research Institute of Industrial Economics.
- Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.

- Benaïm, M. and Weibull, J. W. (2009). Mean-field approximation of stochastic population processes in games. Unpublished manuscript, Université de Neuchâtel and Stockholm School of Economics.
- Berger, U. (2005). Fictitious play in $2 \times n$ games. *Journal of Economic Theory*, 120:139–154.
- Berger, U. (2007a). Brown's original fictitious play. *Journal of Economic Theory*, 135:572–578.
- Berger, U. (2007b). Two more classes of games with the continuous-time fictitious play property. *Games and Economic Behavior*, 60:247–261.
- Berger, U. (2008). Learning in games with strategic complementarities revisited. *Journal of Economic Theory*, 143:292–301.
- Berger, U. and Hofbauer, J. (2006). Irrational behavior in the Brown-von Neumann-Nash dynamics. *Games and Economic Behavior*, 56:1–6.
- Bergin, J. and Lipman, B. L. (1996). Evolution with state-dependent mutations. *Econometrica*, 64:943–956.
- Bhaskar, V. (1998). Noisy communication and the evolution of cooperation. *Journal of Economic Theory*, 82:110–131.
- Bhatia, N. P. and Szegő, G. P. (1970). *Stability Theory of Dynamical Systems*. Springer, Berlin.
- Billingsley, P. (1995). *Probability and Measure*. Wiley, New York, third edition.
- Binmore, K., Gale, J., and Samuelson, L. (1995a). Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, 8:56–90.
- Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265.
- Binmore, K. and Samuelson, L. (1999). Evolutionary drift and equilibrium selection. *Review of Economic Studies*, 66:363–393.
- Binmore, K., Samuelson, L., and Vaughan, R. (1995b). Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35.
- Binmore, K., Samuelson, L., and Young, H. P. (2003). Equilibrium selection in bargaining models. *Games and Economic Behavior*, 45:296–328.
- Bishop, D. T. and Cannings, C. (1978). A generalised war of attrition. *Journal of Theoretical Biology*, 70:85–124.
- Bisin, A. and Verdier, T. (2001). The economics of cultural transmission and the dynamics of preferences. *Journal of Economic Theory*, 97:298–319.

- Björnerstedt, J. and Weibull, J. W. (1996). Nash equilibrium and evolution by imitation. In Arrow, K. J. et al., editors, *The Rational Foundations of Economic Behavior*, pages 155–181. St. Martin's Press, New York.
- Blume, A., Kim, Y.-G., and Sobel, J. (1993). Evolutionary stability in games of communication. *Games and Economic Behavior*, 5:547–575.
- Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424.
- Blume, L. E. (1995). The statistical mechanics of best response strategy revision. *Games and Economic Behavior*, 11:111–145.
- Blume, L. E. (1997). Population games. In Arthur, W. B., Durlauf, S. N., and Lane, D. A., editors, *The Economy as an Evolving Complex System II*, pages 425–460. Addison-Wesley, Reading, MA.
- Blume, L. E. (2003). How noise matters. *Games and Economic Behavior*, 44:251–271.
- Bomze, I. M. (1986). Non-cooperative two-person games in biology: A classification. *International Journal of Game Theory*, 15:31–57.
- Bomze, I. M. (1990). Dynamical aspects of evolutionary stability. *Monatshefte für Mathematik*, 110:189–206.
- Bomze, I. M. (1991). Cross entropy minimization in uninvadable states of complex populations. *Journal of Mathematical Biology*, 30:73–87.
- Bomze, I. M. (2002). Regularity versus degeneracy in dynamics, games, and optimization: A unified approach to different aspects. *SIAM Review*, 44:394–414.
- Bomze, I. M. and Pötscher, B. M. (1989). *Game Theoretical Foundations of Evolutionary Stability*. Springer, Berlin.
- Bomze, I. M. and Weibull, J. W. (1995). Does neutral stability imply Lyapunov stability? *Games and Economic Behavior*, 11:173–192.
- Border, K. C. (2001). Comparing probability distributions. Unpublished manuscript, Caltech.
- Börgers, T. and Sarin, R. (1997). Learning through reinforcement and the replicator dynamics. *Journal of Economic Theory*, 77:1–14.
- Boyce, D. E., Mahmassani, H. S., and Nagurney, A. (2005). A retrospective on Beckmann, McGuire and Winsten's *Studies in the Economics of Transportation*. *Papers in Regional Science*, 84:85–103.
- Boylan, R. T. (1995). Continuous approximation of dynamical systems with randomly matched individuals. *Journal of Economic Theory*, 66:615–625.

- Braess, D. (1968). Über ein Paradoxen der Verkehrsplanung. *Unternehmensforschung*, 12:258–268.
- Brémaud, P. (1999). *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer, New York.
- Brown, G. W. (1949). Some notes on computation of games solutions. Report P-78, The Rand Corporation.
- Brown, G. W. (1951). Iterative solutions of games by fictitious play. In Koopmans, T. C. et al., editors, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, New York.
- Brown, G. W. and von Neumann, J. (1950). Solutions of games by differential equations. In Kuhn, H. W. and Tucker, A. W., editors, *Contributions to the Theory of Games I*, volume 24 of *Annals of Mathematics Studies*, pages 73–79. Princeton University Press, Princeton.
- Bulow, J. and Klemperer, P. (1999). The generalized war of attrition. *American Economic Review*, 89:175–189.
- Burke, M. A. and Young, H. P. (2001). Competition and custom in economic contracts: A case study of Illinois agriculture. *American Economic Review*, 91:559–573.
- Cabrales, A. (1999). Adaptive dynamics and the implementation problem with complete information. *Journal of Economic Theory*, 86:159–184.
- Cabrales, A. (2000). Stochastic replicator dynamics. *International Economic Review*, 41:451–481.
- Cabrales, A. and Ponti, G. (2000). Implementation, elimination of weakly dominated strategies, and evolutionary dynamics. *Review of Economic Dynamics*, 3:247–282.
- Cabrales, A. and Serrano, R. (2007). Implementation in adaptive better-response dynamics. Unpublished manuscript, Universidad Carlos III and Brown University.
- Carmona, G. and Podczeck, K. (2009). On the existence of pure-strategy equilibria in large games. *Journal of Economic Theory*, 144:1300–1319.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press, Cambridge.
- Chamberland, M. and Cressman, R. (2000). An example of dynamic (in)consistency in symmetric extensive form evolutionary games. *Games and Economic Behavior*, 30:319–326.
- Conley, C. (1978). *Isolated Invariant Sets and the Morse Index*. American Mathematical Society, Providence, RI.

- Cooper, R. W. (1999). *Coordination Games: Complementarities and Macroeconomics*. Cambridge University Press, Cambridge.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2004). Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29:961–976.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2008). A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64:457–469.
- Cournot, A. A. (1838). *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. L. Hachette, Paris. Translated by N. T. Bacon (1898), Macmillan, London.
- Cowan, S. G. (1992). *Dynamical Systems Arising from Game Theory*. PhD thesis, University of California, Berkeley.
- Cressman, R. (1992). *The Stability Concept of Evolutionary Game Theory: A Dynamic Approach*. Springer, Berlin.
- Cressman, R. (1995). Evolutionary game theory with two groups of individuals. *Games and Economic Behavior*, 11:237–253.
- Cressman, R. (1996a). Evolutionary stability in the finitely repeated prisoner’s dilemma game. *Journal of Economic Theory*, 68:234–248.
- Cressman, R. (1996b). Frequency-dependent stability for two-species interactions. *Theoretical Population Biology*, 49:189–210.
- Cressman, R. (1997). Local stability of smooth selection dynamics for normal form games. *Mathematical Social Sciences*, 34:1–19.
- Cressman, R. (2000). Subgame monotonicity in extensive form evolutionary games. *Games and Economic Behavior*, 32:183–205.
- Cressman, R. (2003). *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge.
- Cressman, R. (2005). Stability of the replicator equation with continuous strategy space. *Mathematical Social Sciences*, 50:127–147.
- Cressman, R. (2006). Uninvadability in N -species frequency models for resident-mutant systems with discrete or continuous time. *Theoretical Population Biology*, 69:253–262.
- Cressman, R. (2009). Continuously stable strategies, neighborhood superiority and two-player games with continuous strategy space. *International Journal of Game Theory*, 38:221–247.
- Cressman, R., Garay, J., and Hofbauer, J. (2001). Evolutionary stability concepts for N -species frequency-dependent interactions. *Journal of Theoretical Biology*, 211:1–10.

- Cressman, R. and Hofbauer, J. (2005). Measure dynamics on a one-dimensional continuous trait space: Theoretical foundations for adaptive dynamics. *Theoretical Population Biology*, 67:47–59.
- Cressman, R., Hofbauer, J., and Riedel, F. (2006). Stability of the replicator equation for a single-species with a multi-dimensional continuous trait space. *Journal of Theoretical Biology*, 239:273–288.
- Cressman, R. and Schlag, K. H. (1998). On the dynamic (in)stability of backwards induction. *Journal of Economic Theory*, 83:260–285.
- Crouzeix, J.-P. (1998). Characterizations of generalized convexity and generalized monotonicity: A survey. In Crouzeix, J.-P. et al., editors, *Generalized Convexity, Generalized Monotonicity: Recent Results*, pages 237–256. Kluwer, Dordrecht.
- Crow, J. F. and Kimura, M. (1970). *An Introduction to Population Genetics Theory*. Harper and Row, New York.
- Dafermos, S. and Sparrow, F. T. (1969). The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards B*, 73:91–118.
- Dawid, H. and MacLeod, W. B. (2008). Hold-up and the evolution of investment and bargaining norms. *Games and Economic Behavior*, 62:26–52.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, Oxford.
- Dawkins, R. (1982). *The Extended Phenotype*. Oxford University Press, Oxford.
- Dekel, E., Ely, J. C., and Yilankaya, O. (2007). Evolution of preferences. *Review of Economic Studies*, 74:685–704.
- Dekel, E. and Gul, F. (1997). Rationality and knowledge in game theory. In Kreps, D. M. and Wallis, K. F., editors, *Advances in Economics and Econometrics: Theory and Applications*, volume 1, pages 87–172. Cambridge University Press, Cambridge.
- Dembo, A. and Zeitouni, O. (1998). *Large Deviations Techniques and Applications*. Springer, New York, second edition.
- Demichelis, S. and Germano, F. (2000). On the indices of zeros of Nash fields. *Journal of Economic Theory*, 94:192–217.
- Demichelis, S. and Germano, F. (2002). On (un)knots and dynamics in games. *Games and Economic Behavior*, 41:46–60.
- Demichelis, S. and Ritzberger, K. (2003). From evolutionary to strategic stability. *Journal of Economic Theory*, 113:51–75.
- Demichelis, S. and Weibull, J. W. (2008). Language, meaning and games: A model of communication, coordination and evolution. *American Economic Review*, 98:1292–1311.

- Dindoš, M. and Mezzetti, C. (2006). Better-reply dynamics and global convergence to Nash equilibrium in aggregative games. *Games and Economic Behavior*, 54:261–292.
- Dokumacı E. and Sandholm, W. H. (2007a). Schelling redux: An evolutionary model of residential segregation. Unpublished manuscript, University of Wisconsin.
- Dokumacı E. and Sandholm, W. H. (2007b). Stochastic evolution with perturbed payoffs and rapid play. Unpublished manuscript, University of Wisconsin.
- Droste, E., Hommes, C., and Tuinstra, J. (2002). Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games and Economic Behavior*, 40:232–269.
- Dugatkin, L. A. and Reeve, H. K., editors (1998). *Game Theory and Animal Behavior*. Oxford University Press, Oxford.
- Dupuis, P. (1988). Large deviations analysis of some recursive algorithms with state dependent noise. *Annals of Probability*, 16:1509–1536.
- Dupuis, P. and Nagurney, A. (1993). Dynamical systems and variational inequalities. *Annals of Operations Research*, 44:9–42.
- Durrett, R. (2005). *Probability: Theory and Examples*. Brooks-Cole, Belmont, CA, third edition.
- Eigen, M. and Schuster, P. (1979). *The Hypercycle: A Principle of Natural Self-Organization*. Springer, Berlin.
- Ellingsen, T. and Robles, J. (2002). Does evolution solve the hold-up problem? *Games and Economic Behavior*, 39:28–53.
- Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61:1047–1071.
- Ellison, G. (2000). Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45.
- Ellison, G. and Fudenberg, D. (2000). Learning purified mixed equilibria. *Journal of Economic Theory*, 90:84–115.
- Ely, J. C. (2002). Local conventions. *Advances in Economic Theory*, 2:1 (30).
- Ely, J. C. and Sandholm, W. H. (2005). Evolution in Bayesian games I: Theory. *Games and Economic Behavior*, 53:83–109.
- Ely, J. C. and Yilankaya, O. (2001). Nash equilibrium and the evolution of preferences. *Journal of Economic Theory*, 97:255–272.
- Eshel, I. (1983). Evolutionary and continuous stability. *Journal of Theoretical Biology*, 103:99–111.

- Eshel, I., Samuelson, L., and Shaked, A. (1998). Altruists, egoists, and hooligans in a local interaction model. *American Economic Review*, 88:157–179.
- Eshel, I. and Sansone, E. (2003). Evolutionary and dynamic stability in continuous population games. *Journal of Mathematical Biology*, 46:445–459.
- Ethier, S. N. and Kurtz, T. G. (1986). *Markov Processes: Characterization and Convergence*. Wiley, New York.
- Fisher, R. A. (1930). *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford.
- Folland, G. B. (1999). *Real Analysis: Modern Techniques and Their Applications*. Wiley, New York, second edition.
- Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232.
- Foster, D. P. and Young, H. P. (1998). On the nonconvergence of fictitious play in coordination games. *Games and Economic Behavior*, 25:79–96.
- Freidlin, M. I. (1978). The averaging principle and theorems on large deviations. *Russian Mathematical Surveys*, 33:117–176.
- Freidlin, M. I. and Wentzell, A. D. (1998). *Random Perturbations of Dynamical Systems*. Springer, New York, second edition.
- Friedberg, S. H., Insel, A. J., and Spence, L. E. (1989). *Linear Algebra*. Prentice-Hall, Englewood Cliffs, NJ, second edition.
- Friedman, D. (1991). Evolutionary games in economics. *Econometrica*, 59:637–666.
- Friedman, D. and Ostrov, D. N. (2008). Conspicuous consumption dynamics. *Games and Economic Behavior*, 64:121–145.
- Friedman, D. and Yellin, J. (1997). Evolving landscapes for population games. Unpublished manuscript, UC Santa Cruz.
- Friedman, J. W. and Mezzetti, C. (2001). Learning in games by random sampling. *Journal of Economic Theory*, 98:55–84.
- Friesz, T. L., Bernstein, D., Mehta, N. J., Tobin, R. L., and Ganjalizadeh, S. (1994). Day-to-day dynamic network disequilibria and idealized traveler information systems. *Operations Research*, 42:1120–1136.
- Fudenberg, D. and Harris, C. (1992). Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57:420–441.
- Fudenberg, D. and Imhof, L. A. (2006). Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262.

- Fudenberg, D. and Imhof, L. A. (2008). Monotone imitation dynamics in large populations. *Journal of Economic Theory*, 140:229–245.
- Fudenberg, D. and Kreps, D. M. (1993). Learning mixed equilibria. *Games and Economic Behavior*, 5:320–367.
- Fudenberg, D. and Levine, D. K. (1998). *Theory of Learning in Games*. MIT Press, Cambridge.
- Fudenberg, D., Nowak, M. A., Taylor, C., and Imhof, L. A. (2006). Evolutionary game dynamics in finite populations with strong selection and weak mutation. *Theoretical Population Biology*, 70:352–363.
- Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press, Cambridge.
- Garay, B. M. and Hofbauer, J. (2003). Robust permanence for ecological differential equations, minimax, and discretizations. *SIAM Journal on Mathematical Analysis*, 34:1007–1039.
- Gaunersdorfer, A. and Hofbauer, J. (1995). Fictitious play, Shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11:279–303.
- Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.
- Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2008). Quantal response equilibrium. In Blume, L. E. and Durlauf, S. N., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, Basingstoke, U.K., second edition.
- Golman, R. and Page, S. E. (2009). Basins of attraction and equilibrium selection under different learning rules. *Journal of Evolutionary Economics*, forthcoming.
- Gordon, W. B. (1972). On the diffeomorphisms of Euclidean space. *American Mathematical Monthly*, 79:755–759.
- Gorodeisky, Z. (2008). Stochastic approximation of discontinuous dynamics. Unpublished manuscript, Hebrew University.
- Gorodeisky, Z. (2009). Deterministic approximation of best-response dynamics for the Matching Pennies game. *Games and Economic Behavior*, 66:191–201.
- Goyal, S. (2007). *Connections: An Introduction to the Economics of Networks*. Princeton University Press, Princeton.
- Goyal, S. and Janssen, M. C. W. (1997). Non-exclusive conventions and social coordination. *Journal of Economic Theory*, 77:34–57.
- Guckenheimer, J. and Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer, Berlin.

- Güth, W. (1995). An evolutionary approach to explaining cooperative behavior by reciprocal incentives. *International Journal of Game Theory*, 24:323–344.
- Güth, W. and Yaari, M. E. (1992). Explaining reciprocal behavior in simple strategic games: An evolutionary approach. In Witt, U., editor, *Explaining Process and Change: Approaches to Evolutionary Economics*, pages 23–34. University of Michigan Press, Ann Arbor.
- Haigh, J. (1975). Game theory and evolution. *Advances in Applied Probability*, 7:8–11.
- Hamilton, W. D. (1967). Extraordinary sex ratios. *Science*, 156:477–488.
- Hamilton, W. D. (1996). *Narrow Roads of Gene Land*, volume 1. W. H. Freeman/Spektrum, Oxford.
- Hammerstein, P. and Selten, R. (1994). Game theory and evolutionary biology. In Aumann, R. J. and Hart, S., editors, *Handbook of Game Theory*, volume 2, chapter 28, pages 929–993. Elsevier, Amsterdam.
- Harker, P. T. and Pang, J.-S. (1990). Finite-dimensional variational inequality and non-linear complementarity problems: A survey of theory, algorithms, and applications. *Mathematical Programming*, 48:161–220.
- Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2:1–23.
- Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge.
- Hart, S. (2002). Evolutionary dynamics and backward induction. *Games and Economic Behavior*, 41:227–264.
- Hart, S. (2005). Adaptive heuristics. *Econometrica*, 73:1401–1430.
- Hart, S. and Mas-Colell, A. (2001). A general class of adaptive strategies. *Journal of Economic Theory*, 98:26–54.
- Hart, S. and Mas-Colell, A. (2003). Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93:1830–1836.
- Hartman, P. (1964). *Ordinary Differential Equations*. Wiley, New York.
- Hauert, C. (2007). VirtualLabs in evolutionary game theory. Software. <http://www.univie.ac.at/virtuallabs>.
- Heifetz, A., Shannon, C., and Spiegel, Y. (2007). What to maximize if you must. *Journal of Economic Theory*, 133:31–57.
- Henry, C. (1973). An existence theorem for a class of differential equations with multivalued right-hand side. *Journal of Mathematical Analysis and Applications*, 41:179–186.

- Herold, F. and Kuzmics, C. (2008). Evolutionary stability of discrimination under observability. *Games and Economic Behavior*, forthcoming.
- Herz, A. V. M. (1994). Collective phenomena in spatially extended evolutionary games. *Journal of Theoretical Biology*, 169:65–87.
- Hewitt, E. and Stromberg, K. (1965). *Real and Abstract Analysis*. Springer, Berlin.
- Hines, W. G. S. (1980). Three characterizations of population strategy stability. *Journal of Applied Probability*, 17:333–340.
- Hines, W. G. S. (1987). Evolutionary stable strategies: A review of basic theory. *Theoretical Population Biology*, 31:195–272.
- Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). *Fundamentals of Convex Analysis*. Springer, Berlin.
- Hirsch, M. W. (1988). Systems of differential equations that are competitive or cooperative III: Competing species. *Nonlinearity*, 1:51–71.
- Hirsch, M. W. and Smale, S. (1974). *Differential Equations, Dynamical Systems, and Linear Algebra*. Academic Press, San Diego.
- Hirsch, M. W., Smale, S., and Devaney, R. L. (2004). *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. Elsevier, Amsterdam.
- Hofbauer, J. (1981a). A general cooperation theorem for hypercycles. *Monatshefte für Mathematik*, 91:233–240.
- Hofbauer, J. (1981b). On the occurrence of limit cycles in the Volterra-Lotka equation. *Nonlinear Analysis*, 5:1003–1007.
- Hofbauer, J. (1985). The selection mutation equation. *Journal of Mathematical Biology*, 23:41–53.
- Hofbauer, J. (1995a). Imitation dynamics for games. Unpublished manuscript, University of Vienna.
- Hofbauer, J. (1995b). Stability for the best response dynamics. Unpublished manuscript, University of Vienna.
- Hofbauer, J. (1996). Evolutionary dynamics for bimatrix games: A Hamiltonian system? *Journal of Mathematical Biology*, 34:675–688.
- Hofbauer, J. (2000). From Nash and Brown to Maynard Smith: Equilibria, dynamics, and ESS. *Selection*, 1:81–88.
- Hofbauer, J. and Hopkins, E. (2005). Learning in perturbed asymmetric games. *Games and Economic Behavior*, 52:133–152.

- Hofbauer, J. and Imhof, L. A. (2009). Time averages, recurrence and transience in the stochastic replicator dynamics. *Annals of Applied Probability*, forthcoming.
- Hofbauer, J., Mallet-Paret, J., and Smith, H. L. (1991). Stable periodic solutions for the hypercycle system. *Journal of Dynamics and Differential Equations*, 3:423–436.
- Hofbauer, J., Oechssler, J., and Riedel, F. (2009). Brown-von Neumann-Nash dynamics: The continuous strategy case. *Games and Economic Behavior*, 65:406–429.
- Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294.
- Hofbauer, J. and Sandholm, W. H. (2006). Survival of dominated strategies under evolutionary dynamics. Unpublished manuscript, University of Vienna and University of Wisconsin.
- Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69.
- Hofbauer, J. and Sandholm, W. H. (2009). Stable games and their dynamics. *Journal of Economic Theory*, 144:1665–1693.
- Hofbauer, J., Schuster, P., and Sigmund, K. (1979). A note on evolutionarily stable strategies and game dynamics. *Journal of Theoretical Biology*, 81:609–612.
- Hofbauer, J., Schuster, P., and Sigmund, K. (1981). Competition and cooperation in catalytic self-replication. *Journal of Mathematical Biology*, 11:155–168.
- Hofbauer, J. and Sigmund, K. (1988). *Theory of Evolution and Dynamical Systems*. Cambridge University Press, Cambridge.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge.
- Hofbauer, J. and Sigmund, K. (2003). Evolutionary game dynamics. *Bulletin of the American Mathematical Society (New Series)*, 40:479–519.
- Hofbauer, J. and Sorger, G. (1999). Perfect foresight and equilibrium selection in symmetric potential games. *Journal of Economic Theory*, 85:1–23.
- Hofbauer, J., Sorin, S., and Viossat, Y. (2007). Time average replicator and best reply dynamics. Unpublished manuscript, University of Vienna and University of Paris VI and IX.
- Hofbauer, J. and Swinkels, J. M. (1996). A universal Shapley example. Unpublished manuscript, University of Vienna and Northwestern University.
- Hofbauer, J. and Weibull, J. W. (1996). Evolutionary selection against dominated strategies. *Journal of Economic Theory*, 71:558–573.

- Hopkins, E. (1999). A note on best response dynamics. *Games and Economic Behavior*, 29:138–150.
- Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica*, 70:2141–2166.
- Hopkins, E. and Seymour, R. M. (2002). The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43:1157–1190.
- Horn, R. A. and Johnson, C. R. (1985). *Matrix Analysis*. Cambridge University Press, Cambridge.
- Huck, S. and Oechssler, J. (1999). The indirect evolutionary approach to explaining fair allocations. *Games and Economic Behavior*, 28:13–24.
- Hurkens, S. (1995). Learning by forgetful players. *Games and Economic Behavior*, 11:304–329.
- Hutson, V. (1984). A theorem on average Lyapunov functions. *Monatshefte für Mathematik*, 98:267–275.
- Hutson, V. and Schmitt, K. (1992). Permanence and the dynamics of biological systems. *Mathematical Biosciences*, 111:1–71.
- Imhof, L. A. (2005). The long-run behavior of the stochastic replicator dynamics. *Annals of Applied Probability*, 15:1019–1045.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press, Princeton.
- Jacobsen, H. J., Jensen, M., and Sloth, B. (2001). Evolutionary learning in signalling games. *Games and Economic Behavior*, 34:34–63.
- Jordan, J. S. (1993). Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior*, 5:368–386.
- Josephson, J. (2008). Stochastic better reply dynamics in finite games. *Economic Theory*, 35:381–389.
- Josephson, J. and Matros, A. (2004). Stochastic imitation in finite games. *Games and Economic Behavior*, 49:244–259.
- Kandori, M. (1997). Evolutionary game theory in economics. In Kreps, D. M. and Wallis, K. F., editors, *Advances in Economics and Econometrics: Theory and Applications*, volume I, pages 243–277. Cambridge University Press, Cambridge.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.
- Kandori, M. and Rob, R. (1995). Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–414.

- Kandori, M. and Rob, R. (1998). Bandwagon effects and long run technology choice. *Games and Economic Behavior*, 22:84–120.
- Kandori, M., Serrano, R., and Volij, O. (2008). Decentralized trade, random utility and the evolution of social welfare. *Journal of Economic Theory*, 140:328–338.
- Kaniovski, Y. M. and Young, H. P. (1995). Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior*, 11:330–363.
- Khan, M. A. and Sun, Y. (2002). Non-cooperative games with many players. In Aumann, R. J. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 46, pages 1761–1808. North Holland.
- Kim, Y.-G. and Sobel, J. (1995). An evolutionary approach to pre-play communication. *Econometrica*, 63:1181–1193.
- Kimura, M. (1958). On the change of population fitness by natural selection. *Heredity*, 12:145–167.
- Koçkesen, L., Ok, E. A., and Sethi, R. (2000). The strategic advantage of negatively interdependent preferences. *Journal of Economic Theory*, 92:274–299.
- Kojima, F. (2006a). Risk-dominance and perfect foresight dynamics in n -player games. *Journal of Economic Theory*, 128:255–273.
- Kojima, F. (2006b). Stability and instability of the unbeatable strategy in dynamic processes. *International Journal of Economic Theory*, 2:41–53.
- Kojima, F. and Takahashi, S. (2007). Anti-coordination games and dynamic stability. *International Game Theory Review*, 9:667–688.
- Kojima, F. and Takahashi, S. (2008). p -dominance and perfect foresight dynamics. *Journal of Economic Behavior and Organization*, 67:689–701.
- Kosfeld, M. (2002). Stochastic strategy adjustment in coordination games. *Economic Theory*, 20:321–339.
- Koutsoupas, E. and Papadimitriou, C. H. (1999). Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS '99)*, Trier, Germany, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer.
- Krantz, S. G. and Parks, H. R. (1999). *The Geometry of Domains in Space*. Birkhäuser, Boston.
- Krishna, V. and Sjöström, T. (1998). On the convergence of fictitious play. *Mathematics of Operations Research*, 23:479–511.
- Kuhn, H. W. (2003). *Lectures on the Theory of Games*. Princeton University Press, Princeton.

- Kukushkin, N. S. (2004). Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48:94–110.
- Kuran, T. and Sandholm, W. H. (2008). Cultural integration and its discontents. *Review of Economic Studies*, 75:201–228.
- Kurtz, T. G. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58.
- Kurtz, T. G. (1976). Limit theorems and diffusion approximations for density dependent Markov chains. *Mathematical Programming Study*, 5:67–78.
- Kuzmics, C. (2004). Stochastic evolutionary stability in extensive form games of perfect information. *Games and Economic Behavior*, 48:321–336.
- Lahkar, R. (2007). The dynamic instability of dispersed price equilibria. Unpublished manuscript, University College London.
- Lahkar, R. and Sandholm, W. H. (2008). The projection dynamic and the geometry of population games. *Games and Economic Behavior*, 64:565–590.
- Lang, S. (1997). *Undergraduate Analysis*. Springer, New York, second edition.
- Lax, P. D. (2007). *Linear Algebra and Its Applications*. Wiley, Hoboken, NJ, second edition.
- Leonard, R. J. (1994). Reading Cournot, reading Nash: The creation and stabilisation of Nash equilibrium. *The Economic Journal*, 104:492–511.
- Luce, R. D. and Raiffa, H. (1957). *Games and Decisions: Introduction and Critical Survey*. Wiley, New York.
- Mailath, G. J. (1998). Do people play Nash equilibrium? Lessons from evolutionary game theory. *Journal of Economic Literature*, 36:1347–1374.
- Mailath, G. J. and Samuelson, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, Oxford.
- Marden, J. R., Arslan, G., and Shamma, J. S. (2009a). Cooperative control and potential games. *IEEE Transactions on Systems, Man and Cybernetics. Part B: Cybernetics*, forthcoming.
- Marden, J. R. and Shamma, J. S. (2008). Revisiting log-linear learning: Asynchrony, completeness and a payoff-based implementation. Unpublished manuscript, Georgia Tech and Caltech.
- Marden, J. R., Young, H. P., Arslan, G., and Shamma, J. S. (2009b). Payoff-based dynamics for multi-player weakly acyclic games. *SIAM Journal on Control and Optimization*, 48:373–396.

- Marsden, J. E. and Ratiu, T. S. (2002). *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*. Springer, Berlin, second edition.
- Maruta, T. (1997). On the relationship between risk-dominance and stochastic stability. *Games and Economic Behavior*, 19:221–234.
- Maruta, T. (2002). Binary games with state dependent stochastic choice. *Journal of Economic Theory*, 103:351–376.
- Mas-Colell, A. (1984). On a theorem of Schmeidler. *Journal of Mathematical Economics*, 13:201–206.
- Mathevet, L. (2007). Supermodular Bayesian implementation: Learning and incentive design. Unpublished manuscript, Caltech.
- Matsui, A. (1991). Cheap-talk and cooperation in a society. *Journal of Economic Theory*, 54:245–258.
- Matsui, A. (1992). Best response dynamics and socially stable strategies. *Journal of Economic Theory*, 57:343–362.
- Matsui, A. and Matsuyama, K. (1995). An approach to equilibrium selection. *Journal of Economic Theory*, 65:415–434.
- Matsui, A. and Oyama, D. (2006). Rationalizable foresight dynamics. *Games and Economic Behavior*, 56:299–322.
- Maynard Smith, J. (1972). Game theory and the evolution of fighting. In *On Evolution*, pages 8–28. Edinburgh University Press, Edinburgh.
- Maynard Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47:209–221.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.
- McFadden, D. (1981). Econometric models of probabilistic choice. In Manski, C. F. and McFadden, D., editors, *Structural Analysis of Discrete Data with Econometric Applications*, pages 198–272. MIT Press, Cambridge.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38.
- Miękisz, J. (2004). Statistical mechanics of spatial evolutionary games. *Journal of Physics A*, 37:9891–9906.

- Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58:1255–1278.
- Milgrom, P. and Roberts, J. (1991). Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3:82–100.
- Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62:157–180.
- Milgrom, P. R. and Weber, R. J. (1985). Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10:619–632.
- Milnor, J. W. (1965). *Topology from the Differentiable Viewpoint*. Princeton University Press, Princeton.
- Minty, G. J. (1967). On the generalization of a direct method of the calculus of variations. *Bulletin of the American Mathematical Society*, 73:315–321.
- Miyazawa, K. (1961). On the convergence of the learning process in 2×2 non-zero-sum two-person games. Research memorandum 33, Econometric Research Program, Princeton University.
- Möbius, M. M. (2000). The formation of ghettos as a local interaction phenomenon. Unpublished manuscript, MIT.
- Monderer, D. and Sela, A. (1997). Fictitious play and no-cycling conditions. Unpublished manuscript, The Technion.
- Monderer, D. and Shapley, L. S. (1996a). Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68:258–265.
- Monderer, D. and Shapley, L. S. (1996b). Potential games. *Games and Economic Behavior*, 14:124–143.
- Montgomery, J. (2009). Intergenerational cultural transmission as an evolutionary game. Unpublished manuscript, University of Wisconsin.
- Moran, P. A. P. (1962). *The Statistical Processes of Evolutionary Theory*. Clarendon Press, Oxford.
- Morris, S. (2000). Contagion. *Review of Economic Studies*, 67:57–78.
- Morris, S., Rob, R., and Shin, H. S. (1995). p -dominance and belief potential. *Econometrica*, 63:145–157.
- Myatt, D. P. and Wallace, C. C. (2003). A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113:286–301.

- Myatt, D. P. and Wallace, C. C. (2008a). An evolutionary analysis of the volunteer's dilemma. *Games and Economic Behavior*, 62:67–76.
- Myatt, D. P. and Wallace, C. C. (2008b). When does one bad apple spoil the barrel? An evolutionary analysis of collective action. *Review of Economic Studies*, 75:499–527.
- Myatt, D. P. and Wallace, C. C. (2009). Evolution, teamwork, and collective action: Production targets in the private provision of public good. *Economic Journal*, 119:61–90.
- Myerson, R. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge.
- Nachbar, J. H. (1990). 'Evolutionary' selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19:59–89.
- Nagurney, A. (1999). *Network Economics: A Variational Inequality Approach*. Kluwer, Dordrecht, second edition.
- Nagurney, A. and Zhang, D. (1996). *Projected Dynamical Systems and Variational Inequalities with Applications*. Kluwer, Dordrecht.
- Nagurney, A. and Zhang, D. (1997). Projected dynamical systems in the formulation, stability analysis, and computation of fixed demand traffic network equilibria. *Transportation Science*, 31:147–158.
- Nash, J. F. (1950a). Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, 36:48–49.
- Nash, J. F. (1950b). *Non-Cooperative Games*. PhD thesis, Princeton.
- Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics*, 54:287–295.
- Nemytskii, V. V. and Stepanov, V. V. (1960). *Qualitative Theory of Differential Equations*. Princeton University Press, Princeton.
- Nöldeke, G. and Samuelson, L. (1993). An evolutionary analysis of backward and forward induction. *Games and Economic Behavior*, 5:425–454.
- Nöldeke, G. and Samuelson, L. (1997). A dynamic model of equilibrium selection in signaling markets. *Journal of Economic Theory*, 73:118–156.
- Norman, T. W. L. (2008). Dynamically stable sets in infinite strategy spaces. *Games and Economic Behavior*, 62:610–627.
- Norris, J. R. (1997). *Markov Chains*. Cambridge University Press, Cambridge.
- Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap/Harvard, Cambridge.

- Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359:826–829.
- Nowak, M. A. and May, R. M. (1993). The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos*, 3:35–78.
- Nowak, M. A., Sasaki, A., Taylor, C., and Fudenberg, D. (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650.
- Oechssler, J. and Riedel, F. (2001). Evolutionary dynamics on infinite strategy spaces. *Economic Theory*, 17:141–162.
- Oechssler, J. and Riedel, F. (2002). On the dynamic foundation of evolutionary stability in continuous models. *Journal of Economic Theory*, 107:141–162.
- Ok, E. A. (2007). *Real Analysis with Economic Applications*. Princeton University Press, Princeton.
- Ok, E. A. and Vega-Redondo, F. (2001). On the evolution of individualistic preferences: An incomplete information scenario. *Journal of Economic Theory*, 97:231–254.
- Okada, D. and Tercieux, O. (2009). Log-linear dynamics and local potential. Unpublished manuscript, Rutgers University and Paris School of Economics.
- Osborne, M. (2004). *An Introduction to Game Theory*. Oxford University Press, Oxford.
- Oyama, D. (2002). p -dominance and equilibrium selection under perfect foresight dynamics. *Journal of Economic Theory*, 107:288–310.
- Oyama, D., Takahashi, S., and Hofbauer, J. (2008). Monotone methods for equilibrium selection under perfect foresight dynamics. *Theoretical Economics*, 3:155–192.
- Oyama, D. and Tercieux, O. (2009). Iterated potential and robustness of equilibria. *Journal of Economic Theory*, 144:1726–1769.
- Papadimitriou, C. H. (2001). Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC '01), Hersonissos, Greece*, pages 749–753.
- Patriksson, M. (1994). *The Traffic Assignment Problem: Models and Methods*. VSP, Utrecht.
- Pawlowitsch, C. (2008). Why evolution does not always lead to an optimal signaling system. *Games and Economic Behavior*, 63:203–226.
- Pemantle, R. (2007). A survey of random processes with reinforcement. *Probability Surveys*, 4:1–79.
- Pohley, H.-J. and Thomas, B. (1983). Non-linear ESS models and frequency dependent selection. *BioSystems*, 16:87–100.

- Ponti, G. (2000). Cycles of learning in the Centipede game. *Games and Economic Behavior*, 30:115–141.
- Rhode, P. and Stegeman, M. (1996). A comment on “Learning, mutation, and long run equilibria in games”. *Econometrica*, 64:443–449.
- Ritzberger, K. (1994). The theory of normal form games from the differentiable viewpoint. *International Journal of Game Theory*, 23:207–236.
- Ritzberger, K. (2002). *Foundations of Noncooperative Game Theory*. Oxford University Press, Oxford.
- Ritzberger, K. and Weibull, J. W. (1995). Evolutionary selection in normal form games. *Econometrica*, 63:1371–1399.
- Roberts, A. W. and Varberg, D. E. (1973). *Convex Functions*. Academic Press, New York.
- Robinson, C. (1995). *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos*. CRC Press, Boca Raton, FL.
- Robinson, J. (1951). An iterative method of solving a game. *Annals of Mathematics*, 54:296–301.
- Robles, J. (1998). Evolution with changing mutation rates. *Journal of Economic Theory*, 79:207–223.
- Robles, J. (2008). Evolution, bargaining, and time preferences. *Economic Theory*, 35:19–36.
- Robson, A. (1990). Efficiency in evolutionary games: Darwin, Nash, and the secret handshake. *Journal of Theoretical Biology*, 144:379–396.
- Robson, A. and Vega-Redondo, F. (1996). Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory*, 70:65–92.
- Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press, Princeton.
- Rosenthal, R. W. (1973). A class of games possessing pure strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.
- Roughgarden, T. (2005). *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge.
- Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, 49:236–259.
- Roughgarden, T. and Tardos, É. (2004). Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 49:389–403.
- Samuelson, L. (1994). Stochastic stability in games with alternative best replies. *Journal of Economic Theory*, 64:35–65.

- Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge.
- Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391.
- Sandholm, W. H. (1998). Simple and clever decision rules in a model of evolution. *Economics Letters*, 61:165–170.
- Sandholm, W. H. (2001a). Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108.
- Sandholm, W. H. (2001b). Preference evolution, two-speed dynamics, and rapid social change. *Review of Economic Dynamics*, 4:637–639.
- Sandholm, W. H. (2002). Evolutionary implementation and congestion pricing. *Review of Economic Studies*, 69:81–108.
- Sandholm, W. H. (2003). Evolution and equilibrium under inexact information. *Games and Economic Behavior*, 44:343–378.
- Sandholm, W. H. (2005a). Excess payoff dynamics and other well-behaved evolutionary dynamics. *Journal of Economic Theory*, 124:149–170.
- Sandholm, W. H. (2005b). Negative externalities and evolutionary implementation. *Review of Economic Studies*, 72:885–915.
- Sandholm, W. H. (2006). A probabilistic characterization of integrability for game dynamics. Unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. (2007a). Evolution in Bayesian games II: Stability of purified equilibria. *Journal of Economic Theory*, 136:641–667.
- Sandholm, W. H. (2007b). Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132:367–382.
- Sandholm, W. H. (2007c). Simple formulas for stationary distributions and stochastically stable states. *Games and Economic Behavior*, 59:154–162.
- Sandholm, W. H. (2008a). Decompositions and potentials for normal form games. Unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. (2008b). Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. Unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. (2008c). “Strategic Learning and its Limits” by H. Peyton Young. *Games and Economic Behavior*, 63:417–420.
- Sandholm, W. H. (2009a). Evolutionary game theory. In Meyers, R. A., editor, *Encyclopedia of Complexity and Systems Science*, pages 3176–3205. Springer, Heidelberg.

- Sandholm, W. H. (2009b). Exponential protocols and large population potential games. Unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. (2009c). Large population potential games. *Journal of Economic Theory*, 144:1710–1725.
- Sandholm, W. H. (2009d). Local stability under evolutionary game dynamics. *Theoretical Economics*, forthcoming.
- Sandholm, W. H. (2009e). Orders of limits for stationary distributions, stochastic dominance, and stochastic stability. *Theoretical Economics*, forthcoming.
- Sandholm, W. H. (2009f). Stochastic imitative game dynamics with committed agents. Unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. and Dokumacı, E. (2007). Dynamo: Phase diagrams for evolutionary dynamics. Software. <http://www.ssc.wisc.edu/~whs/dynamo>.
- Sandholm, W. H., Dokumacı, E., and Lahkar, R. (2008). The projection dynamic and the replicator dynamic. *Games and Economic Behavior*, 64:666–683.
- Sandholm, W. H. and Pauzner, A. (1998). Evolution, population growth, and history dependence. *Games and Economic Behavior*, 22:84–120.
- Sato, Y., Akiyama, E., and Farmer, J. D. (2002). Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99:4748–4751.
- Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–156.
- Schmeidler, D. (1973). Equilibrium points of non-atomic games. *Journal of Statistical Physics*, 7:295–300.
- Schreiber, S. J. (2000). Criteria for C^r robust permanence. *Journal of Differential Equations*, 162:400–426.
- Schuster, P. and Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538.
- Schuster, P., Sigmund, K., Hofbauer, J., Gottlieb, R., and Merz, P. (1981a). Selfregulation of behaviour in animal societies III: Games between two populations with selfinteraction. *Biological Cybernetics*, 90:16–25.
- Schuster, P., Sigmund, K., Hofbauer, J., and Wolff, R. (1981b). Selfregulation of behaviour in animal societies I: Symmetric contests. *Biological Cybernetics*, 40:1–8.
- Schuster, P., Sigmund, K., Hofbauer, J., and Wolff, R. (1981c). Selfregulation of behaviour in animal societies II: Games between two populations without selfinteraction. *Biological Cybernetics*, 90:9–15.

- Schuster, P., Sigmund, K., and Wolff, R. (1979). Dynamical systems under constant organization. III. Cooperative and competitive behavior of hypercycles. *Journal of Differential Equations*, 32:357–368.
- Selten, R. (1980). A note on evolutionarily stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology*, 84:93–101.
- Seneta, E. (1981). *Non-negative Matrices and Markov Chains*. Springer, New York, second edition.
- Sethi, R. (1999). Evolutionary stability and media of exchange. *Journal of Economic Behavior and Organization*, 40:233–254.
- Sethi, R. and Somanathan, E. (1996). The evolution of social norms in common property resource use. *American Economic Review*, 86:766–788.
- Sethi, R. and Somanathan, E. (2001). Preference evolution and reciprocity. *Journal of Economic Theory*, 97:273–297.
- Shahshahani, S. (1979). A new mathematical framework for the study of linkage and selection. *Memoirs of the American Mathematical Society*, 211.
- Shapley, L. S. (1964). Some topics in two person games. In Drescher, M., Shapley, L. S., and Tucker, A. W., editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–28. Princeton University Press, Princeton.
- Sheffi, Y. (1985). *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall, Englewood Cliffs, NJ.
- Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge.
- Skyrms, B. (1992). Chaos in game dynamics. *Journal of Logic, Language, and Information*, 1:111–130.
- Skyrms, B. (2002). Signals, evolution, and the explanatory power of transient information. *Philosophy of Science*, 69:407–428.
- Slade, M. E. (1994). What does an oligopoly maximize? *Journal of Industrial Economics*, 42:45–51.
- Smale, S. (1967). Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73:747–817.
- Smirnov, G. V. (2002). *Introduction to the Theory of Differential Inclusions*. American Mathematical Society, Providence, RI.
- Smith, H. L. (1995). *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, Providence, RI.

- Smith, M. J. (1984). The stability of a dynamic model of traffic assignment—an application of a method of Lyapunov. *Transportation Science*, 18:245–252.
- Sparrow, C., van Strien, S., and Harris, C. (2008). Fictitious play in 3×3 games: The transition between periodic and chaotic behaviour. *Games and Economic Behavior*, 63:259–291.
- Stroock, D. W. (2005). *An Introduction to Markov Processes*. Springer, Berlin.
- Swinkels, J. M. (1992). Evolutionary stability with equilibrium entrants. *Journal of Economic Theory*, 57:306–332.
- Swinkels, J. M. (1993). Adjustment dynamics and rational play in games. *Games and Economic Behavior*, 5:455–484.
- Szabó, G. and Fáth, G. (2007). Evolutionary games on graphs. *Physics Reports*, 446:97–216.
- Sznitman, A. (1991). Topics in propagation of chaos. In Hennequin, P. L., editor, *Ecole d'Eté de Probabilités de Saint-Flour XIX, 1989*, pages 167–251. Springer, Berlin.
- Takahashi, S. (2008). Perfect foresight dynamics in games with linear incentives and time symmetry. *International Journal of Game Theory*, 37:15–38.
- Tanabe, Y. (2006). The propagation of chaos for interacting individuals in a large population. *Mathematical Social Sciences*, 51:125–152.
- Taylor, C., Fudenberg, D., Sasaki, A., and Nowak, M. A. (2004). Evolutionary game dynamics in finite populations. *Bulletin of Mathematical Biology*, 66:1621–1644.
- Taylor, P. D. (1979). Evolutionarily stable strategies with two types of players. *Journal of Applied Probability*, 16:76–83.
- Taylor, P. D. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.
- Thomas, B. (1984). Evolutionary stability: States and strategies. *Theoretical Population Biology*, 26:49–67.
- Thomas, B. (1985). On evolutionarily stable sets. *Journal of Mathematical Biology*, 22:105–115.
- Topkis, D. (1979). Equilibrium points in nonzero-sum n -person submodular games. *SIAM Journal on Control and Optimization*, 17:773–787.
- Topkis, D. (1998). *Supermodularity and Complementarity*. Princeton University Press, Princeton.
- Trapa, P. E. and Nowak, M. A. (2000). Nash equilibria for an evolutionary language game. *Journal of Mathematical Biology*, 41:172–188.

- Traulsen, A. and Hauert, C. (2009). Stochastic evolutionary game dynamics. In Schuster, H. G., editor, *Reviews of Nonlinear Dynamics and Complexity*, volume 2. Wiley, New York.
- Tröger, T. (2002). Why sunk costs matter for bargaining outcomes: An evolutionary approach. *Journal of Economic Theory*, 102:28–53.
- Tsakas, E. and Voorneveld, M. (2009). The target projection dynamic. *Games and Economic Behavior*, forthcoming.
- Ui, T. (1998). Robustness of stochastic stability. Unpublished manuscript, Bank of Japan.
- Ui, T. (2000). A Shapley value representation of potential games. *Games and Economic Behavior*, 31:121–135.
- van Damme, E. (1991). *Stability and Perfection of Nash Equilibria*. Springer, Berlin, second edition.
- van Strien, S. and Sparrow, C. (2009). Fictitious play in 3×3 games: Chaos and dithering behavior. Unpublished manuscript, University of Warwick.
- Vega-Redondo, F. (1996). *Evolution, Games, and Economic Behaviour*. Oxford University Press, Oxford.
- Vega-Redondo, F. (1997). The evolution of Walrasian behavior. *Econometrica*, 65:375–384.
- Vega-Redondo, F. (2007). *Complex Social Networks*. Cambridge University Press, Cambridge.
- Vickers, G. T. and Cannings, C. (1987). On the definition of an evolutionarily stable strategy. *Journal of Theoretical Biology*, 129:349–353.
- Viossat, Y. (2007). The replicator dynamics does not lead to correlated equilibria. *Games and Economic Behavior*, 59:397–407.
- Viossat, Y. (2008). Evolutionary dynamics may eliminate all strategies used in correlated equilibria. *Mathematical Social Sciences*, 56:27–43.
- Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19:305–321.
- Vives, X. (2000). *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, Cambridge.
- Vives, X. (2005). Complementarities and games: New developments. *Journal of Economic Literature*, 43:437–479.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Prentice-Hall, Princeton.
- Voorneveld, M. (2004). Preparation. *Games and Economic Behavior*, 48:403–414.

- Voorneveld, M. (2005). Consistent retracts and preparation. *Games and Economic Behavior*, 51:228–232.
- Wardrop, J. G. (1952). Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers II*, 1:323–362.
- Wärneryd, K. (1993). Cheap talk, coordination, and evolutionary stability. *Games and Economic Behavior*, 5:532–546.
- Weibull, J. W. (1994). The ‘as if’ approach to game theory: Three positive results and four obstacles. *European Economic Review*, 38:868–881.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge.
- Weibull, J. W. (1996). The mass action interpretation. Excerpt from ‘The work of John Nash in game theory: Nobel Seminar, December 8, 1994’. *Journal of Economic Theory*, 69:165–171.
- Weibull, J. W. (2002). What have we learned from evolutionary game theory so far? Unpublished manuscript, Stockholm School of Economics.
- Young, H. P. (1993a). The evolution of conventions. *Econometrica*, 61:57–84.
- Young, H. P. (1993b). An evolutionary model of bargaining. *Journal of Economic Theory*, 59:145–168.
- Young, H. P. (1998a). Conventional contracts. *Review of Economic Studies*, 65:773–792.
- Young, H. P. (1998b). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.
- Young, H. P. (2001). The dynamics of conformity. In Durlauf, S. N. and Young, H. P., editors, *Social Dynamics*, pages 133–153. Brookings Institution Press/MIT Press, Washington/Cambridge.
- Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford University Press, Oxford.
- Zeeman, E. C. (1980). Population dynamics from game theory. In Nitecki, Z. and Robinson, C., editors, *Global Theory of Dynamical Systems (Evanston, 1979)*, number 819 in Lecture Notes in Mathematics, pages 472–497, Berlin. Springer.
- Zhang, J. (2004a). A dynamic model of residential segregation. *Journal of Mathematical Sociology*, 28:147–170.
- Zhang, J. (2004b). Residential segregation in an all-integrationist world. *Journal of Economic Behavior and Organization*, 24:533–550.