Chapter 10

Strategic Behaviour

> You know my methods [Watson]. Apply them. – Sherlock Holmes
> (Sir Arthur Conan Doyle: *The Sign of Four.*)

10.1 Introduction

In this chapter we focus on the conflict and cooperation that are fundamental to microeconomic problems. The principles of economic analysis that we will develop will provide a basis for the discussion of chapters 11 and 12 and provide essential tools for the wider study of microeconomics. Why a change in the direction of analysis?

Our analysis of strategic behaviour in economics focuses on the theory of games. Game theory is an important subject in its own right and it is impossible to do it justice within a chapter or so. Here we use it as a further powerful analytical tool. The methodology that we will introduce in this chapter offers new insights on concepts and techniques we have discussed earlier including the specification of the optimisation process and the nature of equilibrium. The logical processes may require some mental adjustment in order to grasp the methods involved. But, having mastered the methods, one can apply them – Sherlock Holmes style – to a wide variety of models and problems.

The chapter covers the topics in strategic behaviour by grouping them into three broad areas as follows:

- *The essential building blocks.* In sections 10.2 and 10.3 we review some of the ideas that were taken for granted in the case of perfect markets (chapters 2–7) and rethink the notion of equilibrium. Section 10.4 applies these concepts to industrial organisation.

- *Time.* In section 10.5 we examine how the sequencing of decisions in strategic interactions will affect notions of rationality and equilibrium. Section 10.6 examines these principles in the context of market structure.

- *Uncertainty*. In section 10.7 we introduce some of the issues raised in chapter 8 to the context of strategic interaction. The resulting models are quite rich and the analysis here is continued into chapter 11.

10.2 Games – basic concepts

Many of the concepts and methods of game theory are quite intuitive but, in order to avoid ambiguity, let us set run through a preliminary list of its constituent parts and note those that will require fuller treatment.

10.2.1 Players, rules and payoffs

The literature offers several alternative thumbnail sketches of the elementary ingredients of a game. The following four-part summary has claim to be a consensus approach:

Players

The "players" are the individual entities that are involved in the economic problem represented by the game. We will take these to be economic agents such as firms, households or the government. But occasionally one needs to extend the set of players in games that involve an element of exogenous uncertainty. It can be convenient to treat the random elements of the game as the actions of an extra player known as "Nature," a kind of invisible bogeyman rolling the dice behind the scenes.

Rules of play

The rules of the game focus on *moves* or *actions* of the players. The concept of "action" is a wide-ranging idea covering, for example, the consumption choices made by households, the output decisions of firms, level of taxes...

In a parlour game it is clearly specified what moves each player can legally make at each stage of the game. For a well-specified game in microeconomics this must obviously be done too. But more is involved: in both parlour games and economic problems: the *information* that is available at the point of each move can be crucial to the specification of the game. To illustrate, there is a variant of chess known as Kriegsspiel, in which the players can see their own pieces, but not those of their opponent; kings, queens, pawns and so on all work in the same way, but the rules of the game obviously become fundamentally different from ordinary chess in the light of this difference in information.

Determination of the outcome

For each set of actions or moves (including moves by "Nature" to cover the rôle of uncertainty) there is a specific outcome that is then determined almost mechanistically. The outcome could be defined in terms in terms of lists of outputs, baskets of goods or other economic quantities. It could be something as simple as the answer to the question "who wins?" It is given economic meaning by evaluation in terms of payoffs.

Payoffs

The players' objectives (utility, profits,...) are just as we have introduced them in earlier chapters. As previously we have to be careful to distinguish cases where the payoffs can be treated as purely ordinal concepts (utility in chapter 4) from those where they have cardinal significance (profits in chapter 2 or "felicity" in chapter 8).

These basic ingredients collectively permit a description of what the game is about, but not how it is to be played. To see what more is involved we have to examine some of the game's ingredients more closely: we particularly need to consider the rôle of information.

10.2.2 Information and Beliefs

Uncertainty and progressively changing information can greatly influence the possible outcomes of a game: simply turning over cards in an elementary two-person card game or in solitaire is enough to convince one of that. However more is involved. Take the Kriegsspiel versus ordinary chess example again: without being a chess expert oneself, one can see that it would be useful for player A to try to discover ways of moving his own pieces that will force player B to reveal information about the disposition of B's concealed pieces. What players think that they know is going to affect the way that they play the game and will, in turn, influence the way that information develops through time.

Because information plays such a central rôle in the way a game can unfold it is important to incorporate a precise representation of this within the microeconomic model. The key concept in characterising the situation for an individual agent at any point in the game is the agent's *information set*: this is a full description of the exact state of what is known to the agent at a particular point in the game and will usually (although not necessarily) embody complete recall about everything that has happened previously in the game. Obviously the same individual will usually have a different information set at different stages of the game. We will be able to make the definition of the information set precise once we have considered how to represent the game precisely – in 10.2.4 below.

A central idea in the discussion of "who knows what?" is the concept of "common knowledge." An appeal to common knowledge is frequently a feature of the reasoning required to analyse strategic problems and clearly has much intuitive appeal. However, the term has a precise interpretation in the context of games and microeconomics: a piece of information is common knowledge if it is known by all agents and all agents know that the other agents know it... and so on, recursively.

For cases not covered by the comforting quasi-certainty of "common knowledge" we need to introduce some concept of individual *beliefs* about the way the game works. Of course in some very special cases beliefs are almost irrelevant to the modelling of a game. But usually the use of available information in the modelling of beliefs is an important extension to the concept of rationality that we have employed in earlier chapters. If the individual agent were not making a maximising choice subject to the reasonable beliefs that he has we could say that the individual is irrational. Of course this begs the question of what constitutes "reasonable" beliefs. It also leaves open the issue of how the beliefs could or should be updated in the light of hard information that becomes available during the playing of the game, a point to which we return in section 10.7.

The explicit treatment of uncertainty in models of strategic behaviour and the unfolding of information with the passage of time are important features of microeconomic models and are considered in further detail below.

10.2.3 Strategy

The essence of the game-theoretic approach – and the reason for the title of this chapter – is the focus on *strategy.* A player's needs to be clearly distinguished from the idea of an action. Simply stated, player h's strategy is a complete contingent plan of action for all possible situations that could conceivably arise in the course of a game. It can be expressed formally as follows. Take the collection of all the information sets for agent h corresponding to reachable points within the game: a strategy s^h for agent h is a mapping from this collection to the set of actions feasible for h.

The individual's strategy is the fundamental tool that we will use to analyse the working and outcomes of games.

10.2.4 Representing a game

A game is usually a complex form of strategic interaction. To make sense of it a clear method of representation is required. There are two main forms

- The game in *extensive form* is a kind of tree diagram. The root of the tree is where the game starts and the beginning of each new branch – each *node* – characterises the situation reached at a given moment from a given sequence of actions by the players. At each terminal node (i.e. where the game ends) there is a vector of payoffs, one for each agent. For now these payoffs could be considered to be purely ordinal and need not be comparable between different agents – we will see below situations when these assumptions are no longer satisfactory.

- The game in *strategic form* (also known as *normal form*) is a kind of multidimensional spreadsheet. Each dimension (row, column, etc.) of the spreadsheet corresponds to the set of strategies for each separate player; each cell in the spreadsheet gives a list of numbers corresponding to the payoffs associated with that particular combination of strategies.

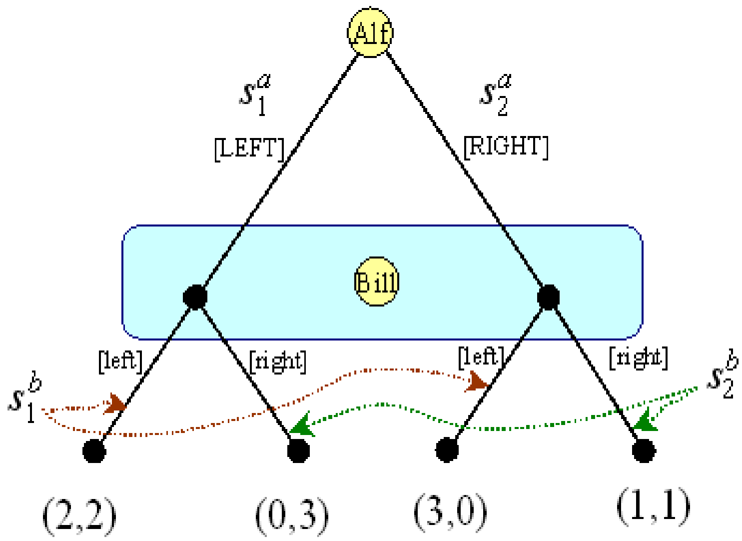

Figure 10.1: Simultaneous move, extensive form

A simple example of the two forms of representation can help here. Figure 10.1 depicts the extensive form of a game where the two players each make a move simultaneously and then the game ends. In this case the strategies for both agents are very simple – each strategy consists of exactly one action. The top of the diagram depicts Alf's choice between the two strategies s_1^a (play [LEFT]) and s_2^a (play [RIGHT]) : his choice then determines whether the left or the right hand node in the middle of the diagram is the relevant one. In the bottom part of the diagram Bill makes his choice (between the actions [left] and [right]); but in view of the simultaneous move he does not, of course, know whether the left-hand or the right-hand node is the relevant one; this lack of clarity is depicted by the shaded box around the two nodes depicting the fact that both nodes are in Bill's information set.[1] At the bottom of the figure is the list of (Alf, Bill)-payoffs resulting from each $\left(s_i^a, s_j^b\right)$-combination. Table 10.1 shows the same game in strategic form. The rows correspond to Alf's choice of strategy; the columns to Bill's choice; the contents of each cell correspond exactly to the bottom line of Figure 10.1.[2]

Note the way the concept of the information set is implicitly defined in Figure 10.1. If the agent knows for certain which node the game has reached he

[1] The game could also be one where Alf moves first but conceals the move that he has made: briefly explain why.

[2] Consider the game in the following table. Why might it be characterised as strategically trivial? Assuming that both agents are rational, what is the game's solution?

	s_1^b	s_2^b
s_1^a	$3,3$	$1,2$
s_2^a	$2,1$	$0,0$

		s_1^b [left]	s_2^b [right]
s_1^a	[LEFT]	$2,2$	$0,3$
s_2^a	[RIGHT]	$3,0$	$1,1$

Table 10.1: Simultaneous move, strategic form

has very precise information to use in making his choice, all the basket of detail associated with the knowledge of being exactly at that node; the information set contains just one point. If there is the possibility of more than one node being relevant – if the information set contains multiple points – then information is less precise. More formally we have:

Definition 10.1 *Agent h's* information set *is the set of nodes that h knows might be the actual node, but that cannot be distinguished by direct observation.*

Does it matter whether extensive form or strategic form is used? In most cases that are relevant to microeconomic modelling the choice between the two forms is largely a matter of expositional convenience, as long as the representation in each of the two forms has been properly done.[3] However it is worth noting that one particular strategic-form representation may correspond to more than one extensive form representation – it is just that the alternative extensive-form representations turn out to be economically equivalent in terms of the way the game is actually played.[4]

10.3 Equilibrium

The players – economic agents – come to the game with their strategies: what would constitute an equilibrium of the economic problem being represented by the game? To address this we can draw on the understanding of equilibrium set out in several contexts in chapters 2 to 7.

First, we introduce a concept that facilitates the definition of further concepts by re-using a term from chapter 9. A *profile* of strategies is a particular collection of strategies, one for each player in the game. Write this as

$$[s] := \left[s^1, s^2, ...\right].$$

Note that we use the same [ ] notation as for allocations in chapter 7. We also need a notation to describe the strategy being played by all those other than

[3] A couple want to decide on an evening's entertainment. He prefers to go to the West End (there's a new play); she wants to go to the East End (dog races). If they go as a couple each person gets utility level 2 if it is his/her preferred activity and 1 otherwise. However, for each person the evening would be ruined if the partner were not there to share it (utility level 0). Depict this as a game in (a) strategic form (b) extensive form.

[4] Draw another extensive-form game tree that corresponds to the strategic form given in Table 10.1.

agent h; this is of course just the profile $[s]$ with the hth component deleted, so we express this as

$$[s]^{-h} := \left[s^1, s^2, ..., s^{h-1}, s^{h+1}, s^{h+2}, ...\right]. \tag{10.1}$$

In order to evaluate the outcome of the game we will write payoffs as utilities. It makes sense to write utility as a function of strategies – in a kind of reduced form. So, for a given profile of strategies $[s]$, we write h's utility as

$$v^h\left(s^h,\ [s]^{-h}\right); \tag{10.2}$$

person h's utility is dependent on his own choice of strategy s^h and on those of everyone else in the game $[s]^{-h}$.

Let us denote the set of all feasible strategies for agent h as S^h: this gives a comprehensive description of what h can do and when he can do it. Then for a given set of agents (players) we can completely diatribe a game by just two objects, a profile of payoff functions and the corresponding list of strategy sets, as follows:

$$\left[v^1, v^2, ...\right];\ \left[S^1, S^2, ...\right] \tag{10.3}$$

These elementary building blocks allow us to introduce the essential concept to grasp in any consideration of economic strategy. This is the idea of an agent's "best response" to other agents' strategies and is defined as follows:

Definition 10.2 *The strategy $\hat{s}^h$ is h's* best response *to $[s]^{-h}$ if*

$$v^h\left(\hat{s}^h,\ [s]^{-h}\right) \geq v^h\left(s^h,\ [s]^{-h}\right) \tag{10.4}$$

for all $s^h \in S^h$, $s^h \neq \hat{s}^h$ or, equivalently, if

$$\hat{s}^h \in \arg\max_{s^h} v^h\left(s^h,\ [s]^{-h}\right) \tag{10.5}$$

The form (10.5) uses the "argmax" notation to denote the set of values of s^h that do the required maximisation job – see Appendix section A.7.5 for a formal definition. We could, of course, alter the definition to "strongly best" by replacing the "$\geq$" with "$>$" in (10.4) in which case the set on the right-hand side of (10.5) has just one element.

The best-response idea is indeed a logical extension of what we have assumed about agents in earlier chapters that focused on perfect markets. There we can see each profit-maximising firm making a "best response" in terms of inputs and outputs to a ruling set of market prices; the utility-maximising consumer makes the "best response" to the market in the light of the household budget and his or her own preferences. But now, instead of the sharp information about market conditions the individual agent has to form a view as to what the consequences will be of his own actions as they are observed and interpreted by other agents.

Contained within the concept of Definition 10.2, there is a very special case that deserves recognition in its own right. A *dominant strategy* is one that remains a best-response strategy whatever the actions of the other players in the game: there is a dominant strategy for agent h if $\hat{s}^h$ in (10.5) is actually independent of $[s]^{-h}$. Of course in many interesting cases dominant strategies just do not exist – but they are of particular interest in certain important applications as we will see in chapter 12.

The idea of the best response leads us on to the fundamental concept of equilibrium of a game.

Definition 10.3 *A* Nash equilibrium *is a profile of strategies* $[s^*]$ *such that, for all agents* h*:*

$$s^{*h} \in \arg\max_{s^h} v^h\left(s^h,\ [s^*]^{-h}\right) \tag{10.6}$$

The plain language interpretation of this is as follows. The Nash equilibrium is a situation where everyone is making the best response to everyone else. No agent has an incentive to deviate from his strategy given that all the other agents do not deviate from their policy.

Finding an equilibrium in the kind of uncomplicated games that we have used thus far can be quite easy. The method essentially follows Sherlock Holmes' dictum "when you have eliminated the impossible, whatever remains, however improbable, must be the truth." So indeed one can often find equilibrium strategies through a process of simple elimination – two examples of this are given in Exercises 10.1 and 10.2. However, in richer models the solution method can be much less straightforward.

Furthermore, although the Nash equilibrium is the main plank on which our approach to strategic behaviour is based we ought to take immediate note of three serious difficulties that are frequently encountered in applying the Nash concept to microeconomic and other problems. These difficulties are handled in 10.3.1 to 10.3.2.

10.3.1 Multiple equilibria

In many interesting economic cases there is more than one Nash equilibrium. For example, in Table 10.2 both $\left[s_1^a, s_1^b\right]$ and $\left[s_2^a, s_2^b\right]$ are equilibria. Clearly the former generates outcomes that Pareto-dominate the latter but, as far as the Nash concept is concerned, each is equally valid as an equilibrium outcome of the game. The second example, in Table 10.3, appears more problematic: the strategy profiles $\left[s_1^a, s_2^b\right]$ and $\left[s_1^a, s_2^b\right]$ (yielding payoffs $(3, 1)$ and $(1, 3)$ respectively) are both Nash equilibria: in contrast to the previous example they are the (only) unequal outcomes of the game – either Alf is exalted and Bill ends in near despair, or vice versa.

So, in each game there are two equilibria: how to choose between them? In some cases the economic context will provide an answer (more on this below); but the Nash concept by itself is of no help.

	s_1^b	s_2^b
s_1^a	3, 3	1, 0
s_2^a	0, 1	2, 2

Table 10.2: Multiple equilibria 1

	s_1^b	s_2^b
s_1^a	2, 2	1, 3
s_2^a	3, 1	0, 0

Table 10.3: Multiple equilibria 2

10.3.2 Efficiency

The terminology "best response" that was used to underpin the Nash equilibrium concept should be treated with caution – "best" in what sense? If we are tempted to reply "best in the sense that a rational agent makes the choice that maximises his own payoff, given the environment that he is in," then we should be aware that rationality needs careful interpretation here. This can be illustrated by the example just considered in Table 10.2 – only one of the two equilibria is efficient, but both equilibria are characterised by "best responses."

The point comes out even more forcefully in the next example. To set the scene let us pose an important question about games in general – what is the worst that can happen to a rational economic agent? Formally we could write this as the *minimax* payoff for agent h:

$$\underline{v}^h := \min_{[s]^{-h}} \left[\max_{s^h} v^h \left(s^h,\ [s]^{-h} \right) \right]; \tag{10.7}$$

Checking back to definition 10.2 we see that expression enclosed in [ ] of (10.7) means that h is making the best response to everyone else's strategy; the "min" operator in (10.7) means that everyone else is trying to punish him within the rules of the game. This minimax value plays the rôle of reservation utility and provides a useful reference point in judging the outcomes of games in terms of their payoffs.

Now for the example: this is the game introduced in Figure 10.1 and Table 10.1 – a game form known as the *Prisoner's Dilemma*.[5] Note first that there

[5] Recreate the Prisoner's Dilemma from the following. Two bad guys have been arrested and are held in separate locations. The problem for the authorities is to prove that they are bad guys: evidence is only likely to come from the individuals themselves. So the authorities announce to each bad guy that if he confesses and implicates the other he will get off with a token sentence of 1 year while the other will go down for 20 years; if they both confess then they each get 10 years. Both of them know, however, that if they both stay schtumm the authorities can only get them for bad driving during the police chase: this will incur a sentence of 2 years each.

Write the game in strategic form and show that there is a dominant strategy for each aof the two bad guys. Find the Nash equilibrium payoffs and explain why it appears inefficient from the bad guys' point of view.

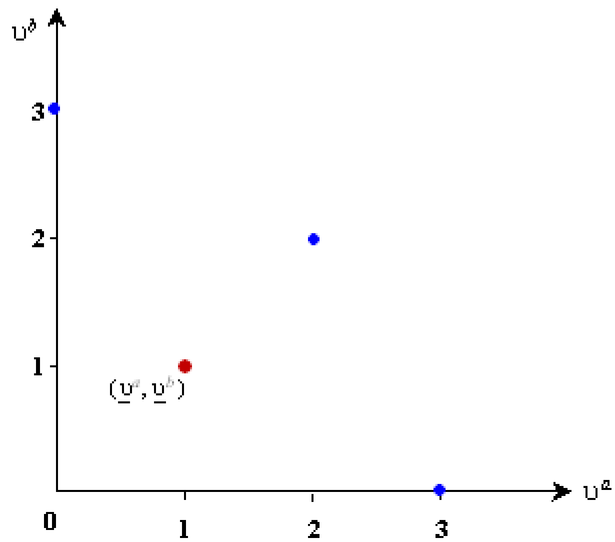

Figure 10.2: Utility possibilities: Prisoner's Dilemma

is a single Nash equilibrium at $[s_2^a, s_2^b]$; note second that it is inefficient: the strategy profile $[s_1^a, s_1^b]$ would yield higher payoffs for both agents! This is illustrated in Figure 10.2 where the utility possibilities representing the payoffs from the game consist of just the four dots.[6] The Nash equilibrium yields in fact the minimax outcome shown in the figure as the utility pair $(\underline{\upsilon}^a, \underline{\upsilon}^b)$. The equilibrium is myopically and individualistically rational, by definition. However, it is arguable that the Pareto-efficient outcome of $(3, 3)$ is where some sense of group rationality ought to lead us.[7]

This is not just a bizarre example carefully selected in order to make a recondite theoretical point. The Prisoner's Dilemma issue lies at the heart of many economic questions where group interests and narrowly defined individual interests do not coincide: we will discuss one important example from the field of industrial organisation in 10.4 below; another important area is introduced in chapter 12.

Suppose that the bad guys get the opportunity to communicate and are then put back into their separate cells: will this make a difference to the outcome of the game?

[6] Draw the same kind of diagram for the games depicted in footnote 3 ("Battle of the Sexes") and in Table 10.3 ("Chicken").

[7] Suppose all of Alf's payoffs are subjected to a given monotonically increasing transformation; and that Bill's payoffs are subjected to another monotonically increasing transformation. Show that the outcome of the game is unaffected.

	s_1^b	s_2^b
s_1^a	$2,2$	$0,3$
s_2^a	$0,1$	$1,0$

Table 10.4: No equilibrium in pure strategies

10.3.3 Existence

There may be no Nash Equilibrium at all. To see this consider the problem depicted in strategic form in Table 10.4 (more on this in exercise 10.3). Again it is set up so that strategies coincide with actions. In this case if Alf (agent a) were to select strategy s_1^a then Bill's best response is to select strategy s_2^b; but if Bill selects strategy s_2^b then Alf's best response is to go for strategy s_2^a; ... and so on round the cycle. There is no strategy profile where each agent is simultaneously making the best response to the other. What is at the bottom of the problem and can one find a way round it?

A suggested solution

Consider the best response for agent a as a function of agent b's strategy, and vice versa: it is clear that they are discontinuous. We may recall from our previous discussion of agents in perfect markets that where the response function was discontinuous it might be that there was, strictly speaking, no market equilibrium (see, for example, pages 53ff. in chapter 3); we may also recall that there is a common-sense argument to "rescue" the equilibrium concept in conventional cases. The query might come to mind whether a similar issue arises with strategic models like those depicted in Table 10.4: is lack of equilibrium in some way attributable to the discontinuity of response in this case? And is there a similar "rescue" argument? In the case of the firm and the market it made sense to appeal to a large numbers argument – on average the supply function is continuous and then we know that there is a price-taking equilibrium. But the large numbers device may not be appropriate here – perhaps there really are only two players. However there is an approach that has a similar flavour. This involves introducing an explicit probabilistic device that allows an agent to enlarge the set of available strategies. We will see how this works in the particular case of the game in Table 10.4 and then examine the issues that are involved in the extra step that apparently offers us the solution.

Suppose that Alf announces that he will adopt strategy s_1^a with probability π^a and strategy s_2^a with probability $1 - \pi^a$. Likewise Bill announces that he will adopt strategies $\left(s_1^b, s_2^b\right)$ with probabilities $\left(\pi^b, 1 - \pi^b\right)$ respectively. Furthermore, let us take the criterion for each of the agents as being their expected payoff (in utility terms). Then, from Table 10.4, if Alf takes π^b as given and chooses probability π^a his expected utility is[8]

$$\left[3\pi^b - 1\right] \pi^a + 1 - \pi^b, \tag{10.8}$$

[8] Use Table 10.4 to derive (10.8) and (10.9) .

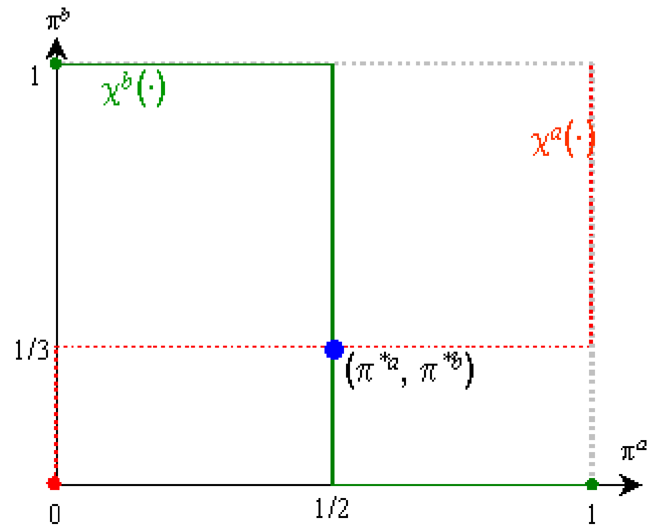

Figure 10.3: Equilibrium in mixed strategy

and if Bill takes π^a as given and chooses probability π^b, then his expected utility is

$$[1 - 2\pi^a]\,\pi^b + 3\pi^a \tag{10.9}$$

We can use (10.8) to derive Alf's choice of π^a as a best response to Bill's choice of π^b. Clearly if $\pi^b = \frac{1}{3}$ the value of π^a has no impact on Alf's expected payoff; but if $\pi^b > \frac{1}{3}$ then (10.8) is increasing in π^a and it would pay Alf to push π^a as high as it will go ($\pi^a = 1$) – i.e. he would then adopt strategy s_1^a with certainty; if $\pi^b < \frac{1}{3}$ the converse happens – (10.8) is then decreasing in π^a and Alf would adopt strategy s_2^a with certainty. Alf's best-response behaviour is summarised by the correspondence $\chi^a(\cdot)$ in Figure 10.3 (we are being picky here: χ^a is a correspondence rather than a function because it is multivalued at the point $\pi^b = \frac{1}{3}$). The expression $\chi^a(\pi^b)$ will give the set of values of π^a that constitute Alf's best response to an announced π^b.

Now think about Bill's best response to Alf's chosen probability. From (10.9) we see that his expected payoff is increasing or decreasing in π^b as $\pi^a < \frac{1}{2}$ or $\pi^a > \frac{1}{2}$, respectively. So, by similar reasoning to the Alf case, Bill's best-response correspondence $\chi^b(\cdot)$ is as depicted in Figure 10.3: for low values of π^a Bill uses strategy s_1^b with certainty and for high values of π^a he adopts s_2^b with certainty.

But now we can see an apparent solution staring at us from Figure 10.3. Put the question, "is there a probability pair such that $\pi^a \in \chi^a(\pi^b)$ and $\pi^b \in \chi^b(\pi^a)$ simultaneously?" and it is clear that the pair $(\pi^{*a}, \pi^{*b}) = \left(\frac{1}{2}, \frac{1}{3}\right)$ does the job exactly. If Alf and Bill respectively select exactly these probabilities when randomising between their two strategies then each is making a best response

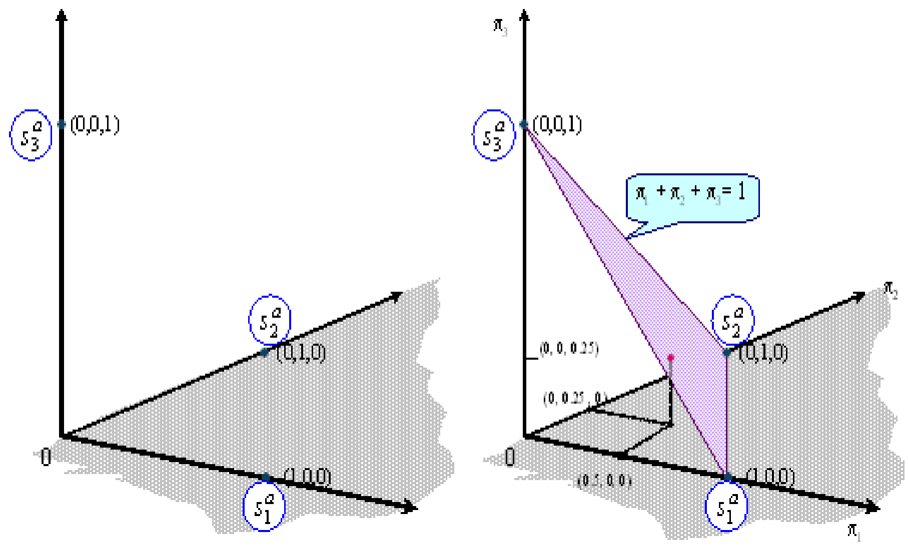

Figure 10.4: Alf's pure and mixed strategies

to the other. Again we seem to have an equilibrium in the Nash sense.

To summarise the suggested resolution of the problem, we see that each agent...

- invents his own lottery that affects the other agent's payoffs;
- knows and believes the probability with which the other agent will adopt any particular strategy;
- formulates a best-response policy by maximising expected utility in the light of that belief.

However, to make clear what is happening with this methodological development we need to re-examine the basic concepts and their meaning.

"Mixed" strategies

First let us refine the description of strategies. We ought to refer to those that have been discussed so far as *pure strategies*. If S^a, the set of pure strategies for agent a, is finite we can imagine each pure strategy as a separate radio button that agent a can press. If in a particular game there were just three pure strategies (three buttons) then we could depict the situation as on the left-hand side of Figure 10.4: each of the agent's three "buttons" is labelled both with the strategy name (s_i^a) and with what looks like the binary code for the button – (0,0,1) and so on.

By introducing randomisation we can change the whole the idea of strategies at a stroke. The picture on the right-hand side of Figure 10.4 is borrowed directly from Figure 8.18 in chapter 8. It depicts the set of lotteries amongst the three pure strategies – the shaded triangle with vertices at $(0,0,1)$, $(0,1,0)$ and $(1,0,0)$. Conventionally each such lottery is known as a *mixed strategy* and the dot in the centre of the picture denotes a mixed strategy where agent a adopts s_1^a, s_2^a, s_3^a with probabilities $0.5, 0.25, 0.25$ respectively. Obviously the idea extends readily to any situation where the number of pure strategies is finite:

Definition 10.4 *Given a finite set S^h of pure strategies for agent h, a* mixed strategy *is a probability distribution over the elements of S^h.*

We can represent the mixed strategy by writing out the elements of S^h in vector form $\left(s_1^h, s_2^h, ...\right)$ and representing the probability distribution by $\boldsymbol{\pi}^h := \left(\pi_1^h, \pi_2^h, ...\right)$ such that π_i^h is the probability that s_i^h is the strategy that is actually adopted by h.[9]

Expected utility

The extension to a mixed-strategy equilibrium also requires a new view of payoffs. In previous examples of games and strategic behaviour we were able to assume that payoffs were purely ordinal. However, by assuming that expected utility is an appropriate criterion, we now have to impose much more structure on individual agents' evaluation of outcomes. In the light of the discussion of chapter 8 (see, for example, page 188) this is not something that we should automatically assume is appropriate.

Two results

The advantage of the extended example based on Table 10.4 is that it conveniently introduces a powerful result lying at the heart of the game-theoretic approach to strategic behaviour:

Theorem 10.1 (Nash equilibrium in mixed strategies) *Every game with a finite number of pure strategies has an equilibrium in mixed strategies.*

The equilibrium in mixed strategies can include degenerate cases where $\boldsymbol{\pi}^h = (0,0,...,1,...)$ (by a linguistic paradox, of course, these "degenerate" cases involve *pure* strategies only...!). It is not hard to see where the result in Theorem 10.1 comes from in view of the result on competitive equilibrium in chapter 7. There a mapping from a convex compact set into itself was used to establish the existence of a general competitive equilibrium using a "fixed point" result (see the discussion in Appendix section C.5.2); the mapping was induced by price

[9] Introducing the possibility of mixed strategies will not change the outcome in the case of the prisoner's dilemma game form. Show this using the same reasoning as for equations (10.8) and (10.9) in the case of the game in Table 10.4.

adjustments using the excess demand function; and the set in question was the set of all normalised prices. Here we have a very similar story: the mapping is the best-response correspondence; the set is the set of mixed strategies, which has exactly the same form as in the general-equilibrium problem – compare Figure 10.4 with Figure B.21 (page 547).

However, in reviewing why this result works the thought might occur whether there is some other way of obtaining an existence result without using the mixed-strategy device – perhaps by appealing to the same fixed-point argument but in a transformed problem. Indeed there is, and for a class of problems that is especially relevant to microeconomic applications. Suppose, in contrast to Theorem 10.1 and the examples used so far, the set of pure strategies is infinite: for example a firm might select an output level anywhere between 0 and $\overline{q}$. Then, in many cases we can use the following:

Theorem 10.2 (Nash equilibrium with infinite strategy sets) *If the game is such that, for all agents h, the strategy sets S^h are convex, compact subsets of $\mathbb{R}^n$ and the payoff functions v^h in (10.2) are continuous and quasiconcave then the game has a Nash equilibrium in pure strategies.*

Mixed strategies: assessment

A mixed strategy can be seen as a theoretical artifice that closes up an otherwise awkward hole in the Nash-equilibrium approach to strategic behaviour. Whether it is an appropriate device depends on specific context of the microeconomic model in which it is employed and the degree to which one finds it plausible that economic actors observe and understand the use of randomisation devices as strategic tools.

This is not the last occasion on which we will find it necessary to refine the concept of equilibrium as new features and subtleties are introduced into the model of strategic behaviour. We will need to keep picking away at the concept of equilibrium as the concept of the game becomes more sophisticated and more interesting.

10.4 Application: duopoly

It is time to put the analysis to work. One of the most obvious gaps in the discussion of chapter 3 was the idea that each firm in a market might have to operate without having a given, determinate demand function. The classic instance of this is oligopoly – competition amongst the few. Each firm has to condition its behaviour not on the parameters of a determinate market environment on the conjectured behaviour of the competition.

We are going to treat this by taking a very simple version of the strategic problem. The rules of the game limit the players to exactly two – *duopoly* as a special case of oligopoly. How the game is to be played will depend on whether decisions about prices or decisions about quantities are to be treated as actions by the firms; it will also depend on whether the firms have to make their move simultaneously (more on this below).

10.4.1 Competition in quantities

We will first examine the classic version of the *Cournot model* and then interpret it in terms of the principles of strategic behaviour that we have set out earlier in this chapter. The Cournot model assumes that firms make decisions over output quantities – the market price will be determined mechanically by market demand – and they make their decisions simultaneously. As a reminder, in this simple world we can treat these quantity decisions, the actions, as strategies.

Model specification

There are two firms simultaneously making decisions on the production of the same homogeneous good. So total market output of the good is given by

$$q = q^1 + q^2 \tag{10.10}$$

where q^f is the output of firm $f = 1, 2$. There is a known market-demand curve for this single good that can be characterised by $p(\cdot)$, the *inverse demand function* for the market: this is just a way of saying that there is a known market price for any given total market output q, thus:

$$p = p(q).$$

Each firm f has a known cost function C^f that is a function just of its own output. So the profits for firm f are:

$$p(q)\, q^f - C^f\left(q^f\right). \tag{10.11}$$

Optimisation

Firm 1 assumes that q^2, the output of firm 2, is a number that is exogenously given. So, using the case $f = 1$ in (10.11), we can see that it is attempting to maximise

$$\Pi^1\left(q^1; q^2\right) := p\left(q^1 + q^2\right) q^1 - C^1\left(q^1\right) \tag{10.12}$$

on the assumption that q^2 is a constant. This is illustrated in Figure 10.5 where firm 1's objectives are represented by a family of isoprofit contours: each contour is in the form of an inverted U and profits for firm 1 are increasing in the direction of the arrow.[10] To find firm 1's optimum given the particular assumption that firm 2's output is constant at q_0^2 just draw a horizontal line at the level q_0^2; this can be repeated for any other given value of firm 1's output conditioned on a particular value of q^2. The graph of these points is conventionally known as firm 1's *reaction function*, which is a slight misnomer. The reaction function

[10] Give a one-line verbal explanation for each of these two assertions.

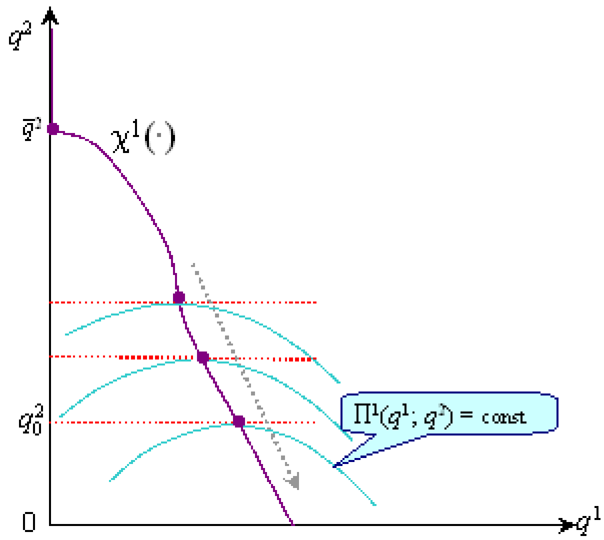

Figure 10.5: Cournot – the reaction function

might be thought of as what firm would do if it were to know of a change in the other firm's action – in simultaneous move games of course this changing about cannot actually happen.

Formally, differentiating (10.12), we have the FOC:

$$\begin{aligned} \frac{\partial \Pi^1\left(q^1; q^2\right)}{\partial q^1} &= p_q\left(q^1+q^2\right) q^1 + p\left(q^1+q^2\right) - C_q^1\left(q^1\right) \leq 0 \\ &= 0 \text{ if } q^1 > 0. \end{aligned} \tag{10.13}$$

We find q^1 as a function of q^2:

$$q^1 = \chi^1\left(q^2\right) \tag{10.14}$$

where $\chi^1(\cdot)$ is a function satisfying (10.13): this is also illustrated in Figure 10.5.[11]

Likewise for firm 2 we get a relationship χ^2 giving q^2 as a function of some arbitrary value q^1 of the output of firm 1:

$$q^2 = \chi^2\left(q^1\right). \tag{10.15}$$

Equilibrium and efficiency

Treating χ^1 and χ^2 as characterising the firms' best responses and combining them, the *Cournot-Nash* solution is then evident – see the point labelled $\left(q_C^1, q_C^2\right)$

[11] Give a brief interpretation of the straight segment of the reaction function for $q^2 >.\bar{q}^2$

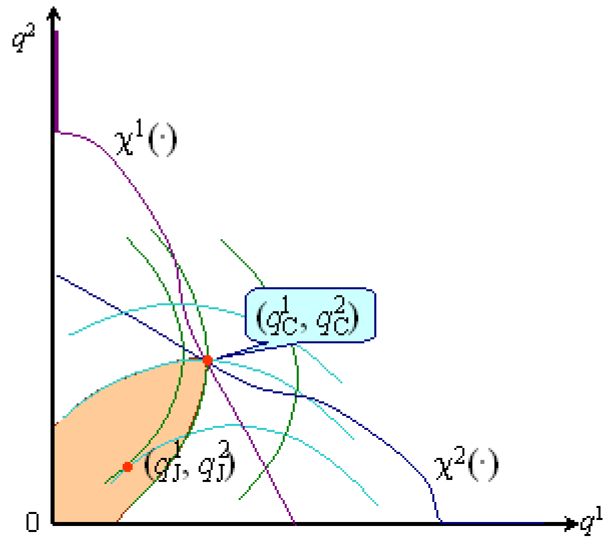

Figure 10.6: Cournot-Nash equilibrium

in Figure 10.6.[12]

Closer inspection of Figure 10.6 reveals a problem, however. Check the two sets of isoprofit contours for the two firms (firm 2's contours are those that run across the diagram in the form of a reverse C-shape): we know that any point lying below firm 1's contour that passes through the Cournot-Nash equilibrium would yield higher profits for firm 1; by the same reasoning, any point to the left of firm 2's contour through the Cournot-Nash outputs means higher profits for firm 2; so any point in the shaded area would mean higher profits for *both* firms. Both firms would benefit if they were able to restrict output and move away from the Cournot-Nash point into this zone. Clearly the Cournot-Nash equilibrium is dominated.

Collusion

Let us tackle the problem from a different direction. Suppose the two firms were able to join forces and pursue their common interest in profit: they form a *cartel*. In the context of the simple model just developed we consider the possibility that the two firms maximise joint profits and split the result between them in some agreed fashion – in effect we are treating the two firms as though they were a single monopoly with two separate plants.

In general the profits for this two-plant monopoly would be

$$p\,(q)\,q - C^1\left(q^1\right) - C^2\left(q^2\right) \tag{10.16}$$

[12] Using theorem 10.2 explain under what conditions we can be sure that the Cournot-Nash equilibrium will exist.

where q is given by (10.10). Differentiating (10.16) with respect to q^f we get:

$$p_q(q)\,q + p(q) - C_q^f\left(q^f\right) = 0 \tag{10.17}$$

$f = 1, 2$. So joint profit maximisation occurs where

$$\begin{aligned} C_q^1\left(q^1\right) &= p_q(q)\,q + p(q) && (10.18)\\ C_q^2\left(q^2\right) &= p_q(q)\,q + p(q) && (10.19)\end{aligned}$$

– marginal cost for each "plant" (each firm) equals overall marginal revenue. From this pair of equations we get the joint-profit maximising outputs $\left(q_{\mathrm{J}}^1, q_{\mathrm{J}}^2\right)$ illustrated in Figure 10.6.[13]

It is clear that the overall profits associated with $\left(q_{\mathrm{J}}^1, q_{\mathrm{J}}^2\right)$ are going to be higher than they would have been at $\left(q_{\mathrm{C}}^1, q_{\mathrm{C}}^2\right)$.

Defection

However, if the joint-profit maximising solution is to survive the two firms would each need an iron resolution and a sharp eye. Each would be tempted by a possibility that is easily demonstrated in Figure 10.6. Draw a line horizontally from $\left(q_{\mathrm{J}}^1, q_{\mathrm{J}}^2\right)$ to the right: it is clear that along this line profits for firm 1 will increase for a while as one moves rightwards. What this means is that, if firm 1 believes that firm 2 is too slow-witted to observe what is happening, then firm 1 might try to "chisel": increase its own output and profits while 2's output stays fixed.[14] Of course firm 2 may have the same temptation, with the rôles reversed (look what happens to its profits on a straight line upwards from the joint-profit maximising solution).

By now we can see the familiar form of the Prisoner's Dilemma emerging. Take a stylised version of the problem we have been discussing: the two firms have identical cost structures and, instead of being able to choose output freely. must select just one of two output levels: either low output or high. We can then reconstruct Table 10.1 as Table 10.5. If both firms choose strategy 1 [low], then each get the joint-profit maximising payoff Π_{J}, but if they both choose strategy 2 [high] then they get only the Cournot-Nash payoffs $\Pi_{\mathrm{C}} < \Pi_{\mathrm{J}}$; if they play different strategies then the one choosing [high] gets $\overline{\Pi} > \Pi_{\mathrm{J}}$ while the one playing [low], gets 0 (this is just for simplicity it could be some positive value less than Π_{C}). Likewise we can reinterpret Figure 10.1 as the extensive form of the Cournot game in Figure 10.7.[15]

[13] The point $\left(q_{\mathrm{J}}^1, q_{\mathrm{J}}^2\right)$ lies on the tangency of the two iso-profit curves such that the tangent passes through the origin. Show why this is so.

[14] What will be happening to firm 2's profits? Why?

[15] There is a possibility here that was not present when we discussed the Prisoner's Dilemma before. The payoffs can be transferred between players – contrast this with Figure 10.2 where the payoffs were in utility (that may or may not be transferable) or footnote question 5 where the payoffs were in length of prison sentence (not transferable). So firms in a cartel could agree on arbitrary divisions of total profits or on side-payments. Draw the set of possible payoffs in the Cournot game. Show that the transferability of the payoff makes no difference to the strategic outcome.

		s_1^2 [low]	s_2^2 [high]
s_1^1	[low]	$\Pi_{\rm J}, \Pi_{\rm J}$	$0, \overline{\Pi}$
s_2^1	[high]	$\overline{\Pi}, 0$	$\Pi_{\rm C}, \Pi_{\rm C}$

Table 10.5: Cournot model as Prisoner's Dilemma

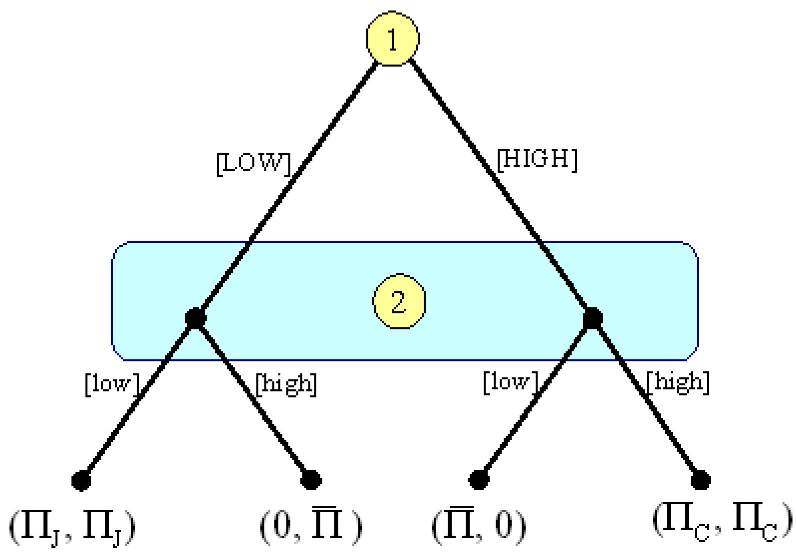

Figure 10.7: Simplified one-shot Cournot game

10.4.2 Competition in prices

Suppose we change the rules of the game for the duopoly: firms play by setting prices rather than quantities: total market output is determined by the market demand curve once the price is known. This the classic *Bertrand model*, adapted slightly here to facilitate comparison with other models.

Model specification

There is a market for a single good with a known market-demand curve. We will assume a straight-line form of this curve so that the quantity sold in this market at price p is given by:

$$q = \frac{\beta_0 - p}{\beta} \tag{10.20}$$

where β_0 and β are positive parameters. If there were a single firm with constant marginal cost c operating in this market then it would announce the following monopoly price[16]

$$p_{\mathrm{M}} = \frac{\beta_0 + c}{2}$$

However, suppose two firms supply the market: each has zero fixed cost and constant marginal cost c. They compete on price as follows. Firm 1 announces price p^1 and firm 2 announces p^2; in the light of this announcement there are three possibilities:

1. If $p^1 < p^2$ firm 1 sells $\frac{\beta_0 - p^1}{\beta}$; firm 2 sells nothing.
2. If $p^1 > p^2$ the reverse happens
3. If $p^1 = p^2 = p$ each firm sells $\frac{\beta_0 - p}{2\beta}$.

Equilibrium

How will the firms set the price? Consider the following steps of an argument:

- Clearly if one firm charges a price above the monopoly price p_{M}, the other can capture the whole market by charging exactly p_{M}.
- If one firm charges a price p above c and at or below p_{M} then the other could charge a price $p - \epsilon$ (where ϵ is a small number) and again capture the whole market.
- If one firm charges a price c then the other firm would not charge a price below this (it would make a loss were it to do that); but it could exactly match the price c, in which case we assume that the market is equally split between the firms.

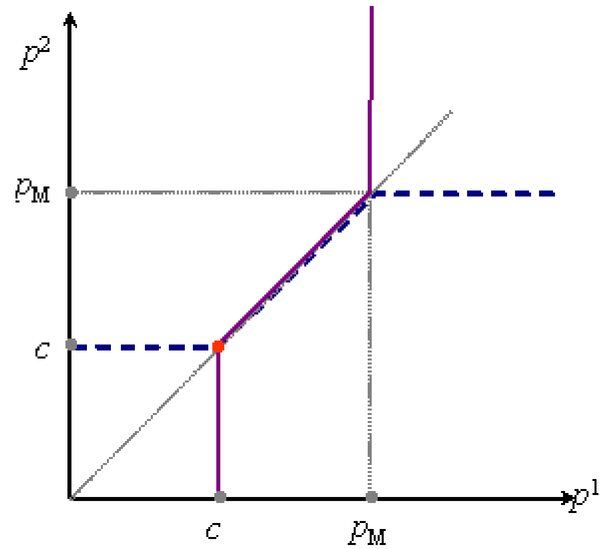

Figure 10.8: Bertrand model

This gives a complete characterisation of a function $\chi^f(\cdot)$ for each firm that would enable us to conclude how it would set its own price given the price that it anticipates would be set by the rival. In the case of firm 1 we have

$$\chi^1\left(p^2\right) = \begin{cases} p^2 - \epsilon & \text{if } p^2 > c \\ c & \text{if } p^2 \leq c \end{cases} \tag{10.21}$$

It is clear from (10.21) that there is a Nash equilibrium at (c, c).[17]

Taken at face value the result seems really remarkable. It appears that there is, effectively, a competitive outcome with just two firms. Contrast this with the case of monopoly (analysed in chapter 3) where the firm sets a price strictly greater than marginal cost with a consequent loss of efficiency. However, it is important to recognise that the rules of the game here are rather restrictive: there are constant marginal costs and no capacity constraints; the product of the two firms is perceived as identical by the customers; the game is played out simultaneously and once only – there is no idea of a true price war. Relaxing any of these assumptions would generate a much richer model; but we can think of the Bertrand model and its solution as an instructive limiting case.

[16] Derive the monopolist's optimum price in this model.

[17] In this case, strictly speaking, χ^f is not a "best-response" function: why? Take a modified version of this model where for administrative reasons it is only possible to set prices as integer values (payment is by coins in a slot machine). Marginal cost is c, an integer, and $p_M = 4c$. Illustrate the game in strategic form; explain why, in this modified model, there *is* a well-defined best-response function for each firm and confirm that the Nash equilibrium outcome is as above.

s^h	strategy for agent h
S^h	strategy set for agent h
$[s]^{-h}$	strategies for all agents other than h
v^h	payoff function for agent h
χ	best-response correspondence
$\boldsymbol{\pi}^h$	randomisation vector for agent h
τ^h	type of agent h

Table 10.6: Strategic behaviour: notation

10.5 Time

Until now, there has been a significant omission in the analysis of strategic behaviour: the lack of an explicit treatment of time. However "time" here has a significance different from that where it has popped up in earlier chapters. We have seen time in its rôle as the scope for economic flexibility (see the discussion of the short run in section 2.4) and time as a characteristic of an economic good (see the discussion of savings in section 5.3.2). Now we focus on time in its rôle of sequencing – the *ordering* of decision-making.

Taking this step means that much more becomes possible within a strategic microeconomic model. Several intuitive concepts in analysing games just make no sense without the introduction of time into the model. One cannot speak about reactions, an equilibrium path, or even threats without modelling the sequence of decision making and carefully consideration of the rôle of information in that sequence.

With this temporal dimension of the strategic problem we will find it important to extend the use and application of the tools introduced in sections 10.2 and 10.3. The distinction between strategies and actions will emerge with greater clarity and we will also need to refine the equilibrium concept. This can be illustrated by re-examining the standard game introduced in Figure 10.1. Suppose the two players now move in sequence – Alf, then Bill. The new situation is represented in extensive form in Figure 10.9. Representing the game in strategic form is a bit more complex and less transparent; but it is done in Table 10.7. There is one small development in notation here; since Bill moves second he has to condition his strategy on what Alf does when making the first move; so we will write, for example, [left-right] for the strategy which states "move left if Alf has chosen [LEFT] and move right if Alf has chosen [RIGHT]." Although at each stage of the game there are exactly two possible actions that a player can take (move left or move right) as far as Bill is concerned there are now *four* strategies $s_1^b, ..., s_4^b$ as shown in the columns of the table.

Will sequencing the play in this way alter the likely outcome of the game? In the case of this particular game the outcome is much the same[18] but in others

[18] Explain why. [hint: put yourself in Bill's position and ask "what would I do if Alf had played [LEFT]? What would I do if he had played [RIGHT]?" Then put yourself in the role of Alf and think about what is going to happen after you have made your move.]

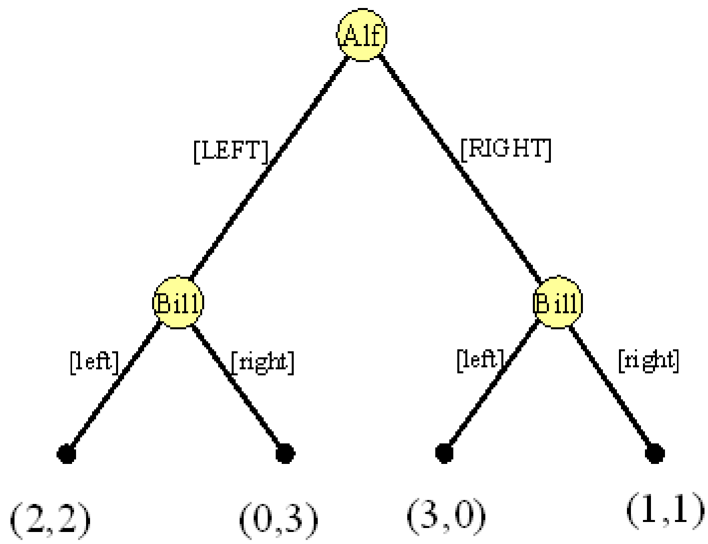

Figure 10.9: Sequential move – extensive form

		s^b_1 [left-left]	s^b_2 [left-right]	s^b_3 [right-left]	s^b_4 [right-right]
s^a_1	[LEFT]	$2,2$	$2,2$	$0,3$	$0,3$
s^a_2	[RIGHT]	$3,0$	$1,1$	$3,0$	$1,1$

Table 10.7: Simultaneous move, strategic form

there can be a drastic change.[19] However, before we treat the solution to this properly we need to consider how the explicit introduction of time allows for more elaborate and illuminating game structures. In doing so we will assume that there is *perfect information* in that everyone knows exactly what happened at earlier stages of the game (this assumption about information will be dropped in section 10.7).

10.5.1 Games and subgames

Let us begin by extending the kind of extensive-form diagram depicted in Figures 10.1 and 10.1. In Figure 10.10 there is a further stage of the game, in other words a further level of decision making with additional nodes; the payoffs after the final stage of the game are given by the payoff profiles $[\upsilon_1], ..., [\upsilon_8]$ where $[\upsilon_i] := \left(\upsilon_i^a, \upsilon_i^b\right)$ gives the payoffs to Alf and Bill in terminal node i.

There is an obvious and useful way of referring to the position of nodes in the structure: take, for example, the nodes highlighted at the bottom of the diagram are those that can be reached from node labelled *: we can think of these as *successor nodes* to *. This enables to make precise an important new concept. A glance at the figure suggests that by deleting part of the tree we can again end up with another viable game tree starting from *. Indeed it is often true that some subsets of the extensive form game can themselves be considered as games and it is these that are of special economic interest:

Definition 10.5 *A* subgame *of a game in extensive form is a subset of the game such that*

1. *It begins at a single node;*
2. *it contains all the successor nodes;*
3. *If the game contains an information set with multiple nodes then either all of these nodes are in the subset or none of them are.*

With reference to Figure 10.11 it is clear that the successor nodes to the node marked * form a subgame as do the successor nodes to the node marked **. But suppose we consider a modified structure as in Figure 10.11: here Alf's choice of actions at the start of the game has been expanded (there is a the [MID] option); furthermore there is an information set with multiple nodes (indicated by the shaded area). Again * marks the beginning of a subgame; but the successor nodes to the node marked # do not form a subgame.[20] The advantage of this new concept is that it permits a naturally intuitive description of the way a game unfolds through time. Think again about the chess analogy used earlier. Even if you are not a chess player you may have seen the kind of chess puzzles that appear in newspapers: typically you are given the position

[19] Take the model of footnote 3. What happens if the players move sequentially? What if they have to move simultaneously?
[20] Explain why.

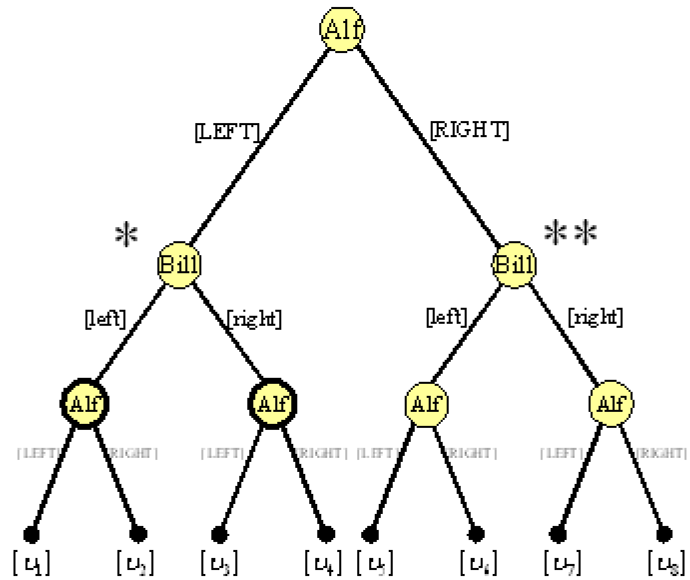

Figure 10.10: Game and subgame (1)

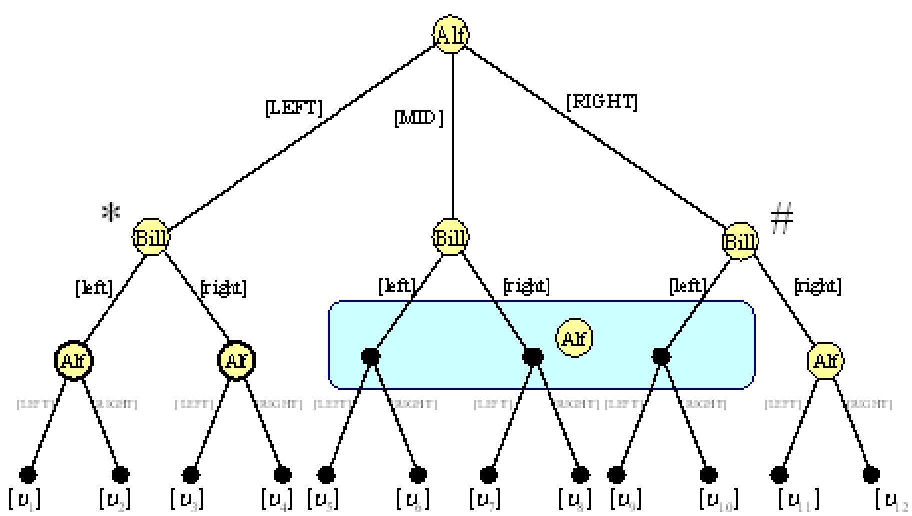

Figure 10.11: Game and subgame (2)

that a game has reached after many moves; then you are asked to finish off the game. Given that the position shown in the puzzle can be reached by a sequence of legal chess moves the puzzle is a subgame of the original game.

In the same way it is interesting to examine the "endgame" of situations of strategic economic interaction. By analysing the endgame one gets a better understanding of the whole of the game: this leads us naturally on to a further discussion of solution concepts.

10.5.2 Equilibrium: more on concept and method

In the light of the multi-period nature of games we need not only to re-examine the way in which a solution is derived but also what is meant by a satisfactory solution. The reason for this is that, as we will see, some Nash equilibria can appear as unattractive when examined from the point of view of each subgame.

So, how to solve for an equilibrium in this case? We will start with some useful intuition and then move on to a more formal concept. Again we can use another good principle from Sherlock Holmes (from *A Study in Scarlet*): "In solving a problem of this sort, the grand thing is to be able to reason backwards. That is a very useful accomplishment, and a very easy one, but people do not practise it much." This intuition is exactly what is required: start at the end of the game and work back through the stages of the game – a process usually known as backwards induction.

To see how this works let us apply the method to solve the game in the case of Figure 10.10. Suppose it is true that $\upsilon_1^a > \upsilon_2^a$ and $\upsilon_3^a > \upsilon_4^a$ and so on (we could easily retell the story if the inequalities were different). Then, if the game had reached the lower left-hand node where it is Alf's turn to play, obviously Alf would choose 1; so the value of reaching this node is effectively $[\upsilon_1] = (\upsilon_1^a, \upsilon_1^b)$; reasoning in this way we can see that the value associated with reaching each of the other nodes on the same level of this diagram is $[\upsilon_3]$, $[\upsilon_5]$, $[\upsilon_7]$ respectively. We have effectively reduced a three-stage game to a two-stage game with payoffs $[\upsilon_1]$, $[\upsilon_3]$, $[\upsilon_5]$, $[\upsilon_7]$. We can then solve the two-stage game using the same method – see footnote 18 above.

Associated with the backward-induction method we can now introduce a refined concept of equilibrium in a multi-stage game:

Definition 10.6 *A profile of strategies is a* subgame-perfect equilibrium *for a game if*

1. *It is a Nash equilibrium*
2. *It induces actions that are consistent with a Nash equilibrium in every subgame.*

Two key points about this concept and the associated backwards-induction algorithm should be noted right away:

- Definition 10.6 is quite demanding because it says something about *all* the subgames, even if one might have thought that some individual subgames are not particularly interesting and are unlikely to be actually reached in practice.

- The straightforward backward-induction method is not going to be suitable for all games with richer information sets. We will come back to this point in section 10.7.4 below.

Now for the reason why the concept of equilibrium needs to be refined in this way when we take into account the temporal sequence of a game: some Nash equilibria involve strategies that lack credibility. What we mean by this is as follows. Imagine reaching the final stage of a game at a position where a specific move by player h may well damage the opponent(s) but would cause serious damage to player h himself. Taking the subgame starting from this position as a game in its own right it is clear that h would not rationally make the move; so, in the context of the overall game, threatening to make this move should the position be reached is unlikely to be impressive. Yet there may well be Nash equilibria of the whole game that imply the use of such empty threats: clearly there is a good case for discarding such strategy combinations as candidates for equilibria and focusing just on those that satisfy subgame perfection (definition 10.6).

This point is illustrated in 10.12. Alf gets to play first; Bill knows that if Alf plays [RIGHT] then Bill gets a payoff of 2; but if they play the sequence [LEFT],[right] then the situation would be disastrous for Bill – he would get a payoff of no more than 1. Can Bill dissuade Alf from playing [LEFT] by threatening to play [left] as well, so reducing Alf's payoff to 0?

On checking the strategic form in Table 10.8 we can see that there are four Nash equilibria $\left[s_2^a, s_1^b\right]$, $\left[s_2^a, s_2^b\right]$, $\left[s_1^a, s_3^b\right]$ and $\left[s_1^a, s_4^b\right]$: the first two of these are equivalent in their outcomes; likewise the third and fourth equilibria are equivalent. So it appears that the case where Alf's strategy is to play [RIGHT] and Bill's strategy is to play [left] whatever Alf does $\left[s_2^a, s_1^b\right]$ is a valid equilibrium outcome of the game. But it is a bit odd. Put the case that on Monday Alf plays [LEFT] anyway and then says to Bill (who plays on Tuesday) "what are you going to do about that?" Presented with this *fait accompli* one could imagine Bill thinking on Monday night that maybe he ought to make the best of a bad job and play [right]: the reasoning is that on Monday night we are at the node marked * and, viewed from this standpoint, Bill would do better to play [right] on Tuesday in order to secure a payoff of 1 rather than 0. Knowing that this is how a rational opponent would reason on Monday night, Alf is unlikely to be impressed by a threat from Bill on Sunday of "I'll play [left] whatever happens." So, although $\left[s_2^a, s_1^b\right]$ is a Nash equilibrium, it is not subgame perfect.[21]

[21] Back in the 1960s nuclear strategists (seriously?) discussed the idea of a "Doomsday machine." This was to be a gizmo that would automatically launch world-devastating nuclear strike if (a) any nuclear missile landed on its home territory or (b) any attempt was made to disarm it. Could a similar device assist Bill?

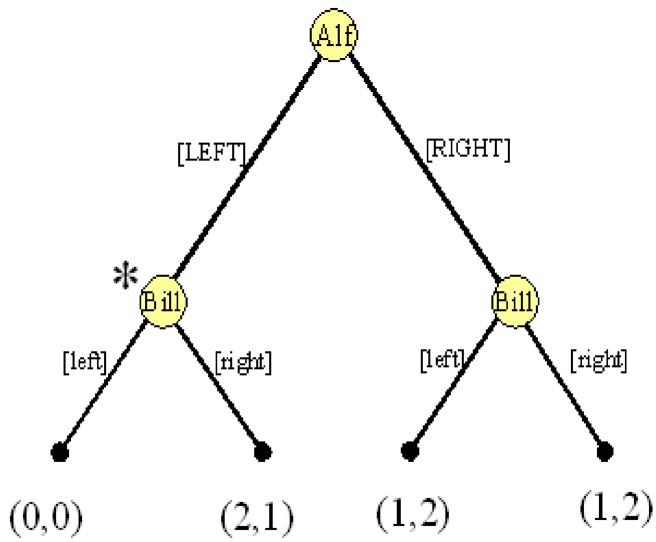

Figure 10.12: An incredible threat

		s_1^b [left-left]	s_2^b [left-right]	s_3^b [right-left]	s_4^b [right-right]
s_1^a	[LEFT]	$0, 0$	$0, 0$	$\mathbf{2}, \mathbf{1}$	$\mathbf{2}, \mathbf{1}$
s_2^a	[RIGHT]	$\mathbf{1}, \mathbf{2}$	$\mathbf{1}, \mathbf{2}$	$1, 2$	$1, 2$

Table 10.8: Incredible threat: strategic view

By restricting attention to equilibria that satisfy subgame perfection we are insisting on an important aspect of consistency in economic behaviour. In doing this we have to consider what a player would do in positions that are not actually played out.

10.5.3 Repeated interactions

For some purposes it is useful to jump from the case of comparatively few stages to the case of arbitrarily many. The principles that can be learned from this apparently arbitrary exercise have some profound implications. They can illuminate the possibilities for long-term cooperative outcomes that may appear absent from a myopic analysis of a simple model of strategic interaction.

The basic idea of a *repeated game* is simple. One joins together multiple instances of an atemporal game: the analysis models a repeated encounter between the players in apparently the same simple situation of economic conflict. Figure 10.13 shows an outline of the setup for the Prisoner's Dilemma game: the same players face the same outcomes from their actions that they may choose in periods $1, 2, ..., t, ...$ The example of the Prisoner's Dilemma is particularly instructive given its importance in microeconomics and, as noted earlier (page 280), the somewhat pessimistic outcome of an isolated implementation of the game.

What makes the repeated game different from a collection of unrelated games of identical structure with identical players? The key point is history. One typically assumes that everyone can know all the information about actual play that has accumulated at any particular stage of the game – the perfect-information assumption again. Individual strategies can then be conditioned on this information and may be used to support equilibrium outcomes that could not have arisen from play by rational economic agents of an isolated single encounter.

The stage game

The basic building block of repeated-interactions analysis is the *stage game*. This is just an instance of one of the simultaneous-play atemporal games that were considered in section 10.3: in particular we can see that each stage in Figure 10.13 is just a copy of Figure 10.1. It is important to distinguish between what goes on in a single play of the stage game and strategy in the game as a whole. If an instance of the stage game were to be played in isolation, of course, we can take strategies as being equivalent to actions; but if the stage game is taken as a component of the repeated game then the individual strategies refer to planned choices over the entire sequence of play: the actions at stage $t+1$ will have been conditioned by the sequence of behaviour up to t.

It is also important to understand the relationship between the payoffs that emerge from an isolated instance of the stage game and those that might be obtainable from a repeated version of the game in which strategies can be conditioned on history. In Figure 10.14, based on Figure 10.2, we have introduced the set of all payoffs that could be reached by mixing the payoffs from the pure

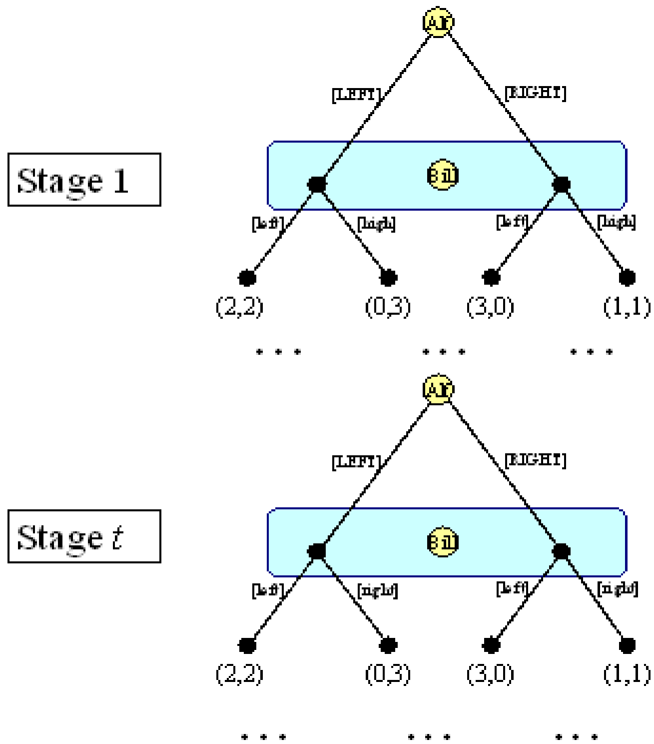

Figure 10.13: Repeated Prisoner's Dilemma

strategy combinations in the basic Prisoner's Dilemma game: these are represented by the heavily shaded lozenge shape. The mixes could be achieved by agreeing on a coordinated randomisation plan or by taking it in turns to use different strategy combinations, for example. Note the following features of this figure:

- The "south-west" corner of the shaded set represents the minimax outcomes for the two players – the worst that can happen to player h in a particular instance of the stage game; as we know it is also the Nash-equilibrium outcome of the stage game.

- The set represented by lightly shaded area north-east of this point consists of all the payoffs that would be Pareto improvements over the Nash-equilibrium outcome.

- The set $\mathbb{U}^*$, as the intersection of these two sets, consists of payoffs that are an improvement on the Nash outcome and that can be represented as mixtures of payoffs in a one-shot stage game.

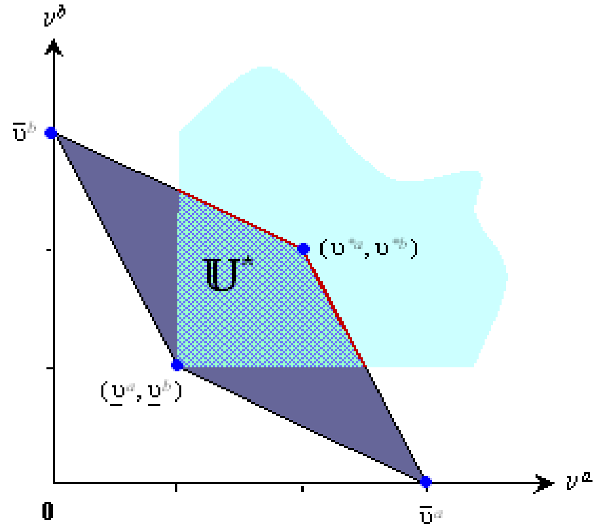

Figure 10.14: Utility possibilities: Prisoner's Dilemma with "mixing"

- The points on the north-east boundary of $\mathbb{U}^*$ correspond to the Pareto-efficient outcomes.

The issue is, can one achieve a Pareto-efficient outcome in $\mathbb{U}^*$ or, indeed, anything other than the minimax value at $\left(\underline{\upsilon}^a, \underline{\upsilon}^b\right)$? As we know (check footnote 9) the use of mixed strategies in an isolated play of Prisoner's Dilemma does nothing to alter the single Nash-equilibrium outcome at $\left(\underline{\upsilon}^a, \underline{\upsilon}^b\right)$; however, it may be that through the structure of repetitive play other points $\mathbb{U}^*$ are implementable as equilibrium outcomes.

The repeated game

To investigate this possibility we need a model of payoffs in an infinite-horizon world. Obviously this is based on the model of payoffs in a typical stage game: but we also need a method of aggregating payoffs across the stages. The aggregation method is a generalisation of the intertemporal utility function in equation (5.13). Specifically, let $\upsilon^h(t)$ denote the payoff for agent h in period t and introduce the possibility of pure time preference in the form of a discount factor δ lying between zero and one inclusive; then the value of payoff stream $\left(\upsilon^h(1), \upsilon^h(2), ..., \upsilon^h(t) ...\right)$ is given by

$$[1-\delta]\sum_{t=1}^{\infty}\delta^{t-1}\,\upsilon^h(t) \tag{10.22}$$

Note two technical points about the specification of (10.22). First, the term $[1-\delta]$ performs a normalisation rôle : if the payoff in the stage game were

Alf's action in $0, ..., t$	**Bill's action at** $t+1$
[LEFT][LEFT].... [LEFT]	[left]
Anything else	[right]

Table 10.9: Bill's trigger strategy s_{T}^{b}

constant throughout all time, so that $\upsilon^h(t) = \upsilon_0$, then the overall payoff is itself υ_0. Second, if we allow $\delta \to 1$ then the overall payoff becomes a simple average with current utility components being given equal weight with those in the indefinite future.

Why an infinite number of periods? The short answer is that this ensures that there is always a tomorrow. In many situations if there were to be a known Last Day then the game would "unravel": you just have to imagine yourself at the Last Day and then apply the Sherlock-Holmes working-backwards method that we outlined in section 10.5.2 above.

How could rational players use the information from a history of play in a repeated game? We can illustrate a method in an argument by example on Figure 10.14. Suppose Alf and Bill collectively recognise that it would be in their interests if they could maintain actions in each period that would guarantee them the Pareto-efficient payoffs $\left(\upsilon^{*a}, \upsilon^{*b}\right)$ in each period; to do this they need to play [LEFT],[left] every period. The problem is they cannot trust each other, nor indeed themselves: Alf has the temptation to jump at the possibility of getting the payoff $\overline{\upsilon}^a$ by being antisocial and playing [RIGHT]; Bill has a similar temptation. To forestall this suppose that they each adopt a strategy that (1) rewards the other party's cooperative behaviour by responding with the action [left] and (2) punishes antisocial behaviour with the action [right], thus generating the minimax payoffs $\left(\underline{\upsilon}^a, \underline{\upsilon}^b\right)$. What gives the strategy bite is that the punishment action applies to *every* period after the one where the antisocial action occurred: the offender is cast into outer darkness and minimaxed for ever. This is known as a *trigger strategy*.

Consider the trigger strategy for Bill, s_{T}^{b}, set out in detail in Table 10.9: would it persuade Alf to behave cooperatively? The gain to Alf from behaving antisocially in period t is $\overline{\upsilon}^a - \upsilon^{*a}$. The consequence for Alf in every period from $t+1$ onwards is a difference in utility given by $\upsilon^{*a} - \underline{\upsilon}^a$ per period; so Alf would not find it worth while to behave antisocially if[22]

$$\overline{\upsilon}^a - \upsilon^{*a} \leq \frac{\delta}{1-\delta}\left[\upsilon^{*a} - \underline{\upsilon}^a\right] \tag{10.23}$$

The trigger strategy for Alf follows the same reasoning – just interchange the a and b labels.

Now let us examine whether the strategy pair $\left[s_{\mathrm{T}}^{a}, s_{\mathrm{T}}^{b}\right]$ constitutes an equilibrium that would support the Pareto-efficient payoffs. Note first that if there were antisocial behaviour at t then the sequence of actions prescribed by Table 10.9 and its counterpart for s_{T}^{a} together constitute a Nash equilibrium for the

[22] Explain why.

subgame that would then start at $t+1$: Alf could not increase his payoff by switching from [RIGHT] to [LEFT] given that Bill is playing [left]; likewise for Bill. The same conclusion follows for any subgame starting after $t+1$. Note second that if δ is large enough[23] and [LEFT],[left] has been played in every period up till t then it is clear from (10.23) that Alf would not wish to switch to [RIGHT]; again a similar statement follows for Bill. So $\left[s_{\mathrm{T}}^{a}, s_{\mathrm{T}}^{b}\right]$ is a subgame-perfect perfect equilibrium that will implement $\left(\upsilon^{*a}, \upsilon^{*b}\right)$.[24]

It is important to recognise that this reasoning is not specific to an isolated example, as the following key result shows:

Theorem 10.3 (The Folk Theorem) *In a two-person infinitely repeated game any combination of actions observed in any finite number of stages is the outcome of a subgame-perfect equilibrium if the discount factor is sufficiently close to 1.*

Theorem 10.3 – known as the Folk theorem because informal versions of it were around well before it was formally stated and proved – tells us that *any* point in $\mathbb{U}^*$ can be supported as a subgame-perfect equilibrium, given a condition on the utility function in (10.22). However, this does not mean that the result turns just on a quirk of individuals' intertemporal preferences. We can consider the discount factor to be a product of a factor derived from a person's impatience – a pure preference parameter – and the probability that the person will be around to enjoy utility in the next period. (Check out the reasoning in Exercise 8.9 to convince yourself of this). So, in this case we can imagine that although in principle the game could go on forever, there is a probability that it will end in finite time. Then Theorem 10.3 requires both that this probability be "sufficiently low" and that the individual agents be "sufficiently patient."

Although I have tagged Theorem 10.3 as *The* Folk Theorem there is actually a family of results that deal with this type of issue in the field of repeated games: the version stated here is somewhat conservative. Some results focus only on Nash equilibria (which, perhaps, rather misses the point since credibility is important), some deal with more than two agents (but ensuring subgame-perfection then gets a bit tricky) and some discuss repeated games of finite length. However, in assessing the contribution of the Folk Theorem(s) it is important to be clear about the main message of the result.

The implication of Theorem 10.3 is that there is a wide range of possible equilibria in infinitely repeated games: it does not predict that rational behaviour will generate one specific outcome. Should it seem troubling that there are so many equilibrium outcomes for the repeated game? Perhaps not: we can think of Theorem 10.3 as a kind of possibility result demonstrating that strategic problems that do not have "sensible" solutions in the short run may yet be susceptible of sensible solution in the long run through induced cooperation.

[23] We need to have $\underline{\delta} \leq \delta \leq 1$. What is the value of $\underline{\delta}$?

[24] In the answer to footnote question 7 it is shown that a monotonic transformation of utilities does not change the outcome of the Prisoner's Dilemma one-shot game. Could such a transformation affect the repeated game?

10.6 Application: market structure

The temporal sequence on which we have focused plays an important rôle in the analysis of industrial organisation. We will illustrate its contribution by considering three applications.

10.6.1 Market leadership

First, let us revisit the simple competition-in-quantities version of the duopoly model Explicit recognition of the time sequence within the game structure permits the strategic modelling of an important economic phenomenon, market leadership.

Assume that social customs or institutional rules (of what sort, or from where, we do not enquire) ensure that firm 1 gets the chance to move first in deciding output – it is the *leader*. The follower (firm 2) observes the leader's output choice q^1 and then announces its output q^2. What would we expect as a solution?

First let us note that the Nash concept does not give us much leverage. In fact, using the reaction function given in (10.15), any non-negative output pair $\left(q^1, q^2\right)$ satisfying $q^2 = \chi^2\left(q^1\right)$ can be taken as the outcome of a Nash equilibrium to the sequentially played game described above; but given the sequence of decision making we know that many of these equilibria will involve incredible threats – they are not subgame perfect.[25] To find the subgame-perfect equilibrium consider first the subgame that follows firm 1's output decision; clearly this involves firm 2 choosing $\chi^2\left(q^1\right)$ as a best response to whatever q^1 has been selected; reasoning backwards firm 1 will therefore select its output so as to maximise its profits conditional on firm 2's best response.

The upshot of this argument is that the leader effectively manipulates the follower by choosing its own output appropriately. Given the reaction function (10.15), the leader's expression for profits becomes

$$p\left(q^1 + \chi^2\left(q^1\right)\right) q^1 - C^1\left(q^1\right) \tag{10.24}$$

The prerogative of being the leader is the opportunity to construct an opportunity set for oneself from the responses of one's opponents: this is illustrated in Figure 10.15 where firm 2's reaction function χ^2 marks out the boundary of firm 1's opportunity set. This is the essence of the *Stackelberg* model of duopoly.

The solution to the Stackelberg duopoly problem (10.24) is depicted by the point $\left(q_{\mathrm{S}}^1, q_{\mathrm{S}}^2\right)$ in Figure 10.15: the leader's isoprofit contour is tangent to the follower's reaction function at this point. The leader has a *first-mover advantage* in that firm 1's profits will be higher than those of firm 2 and, indeed higher than would be the case at the Cournot-Nash solution.

[25] Let q_{M} be the profit-maximising output for firm 2 if it were an monopolist and assume that $\chi^1\left(q_{\mathrm{M}}\right) = 0$ in the case of simultaneous play – see (10.14). Show that in the sequential-play game a strategy pair yielding the output combination $(0, q_{\mathrm{M}})$ is a Nash equilibrium but is not a subgame-perfect solution.

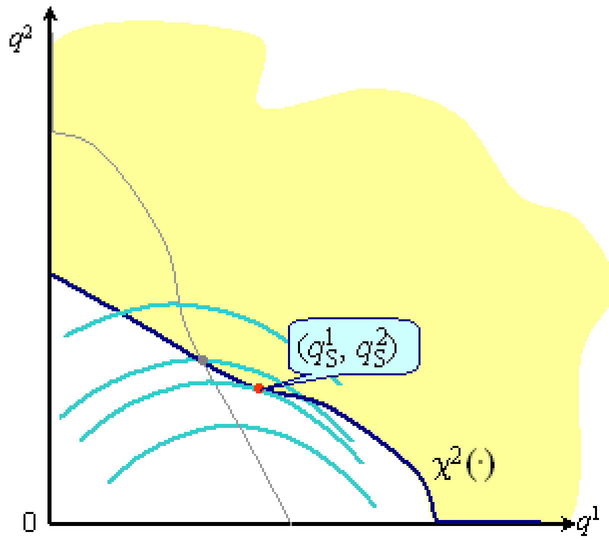

Figure 10.15: Leader-follower

However, the Stackelberg analysis leaves upon a fundamental and important question – what constitutes a credible leader? How is the leadership position maintained? There are two responses here. First, this special duopoly model establishes some important principles that are relevant for other economic applications (see chapter 12). Second, we can dig a little deeper into the issues of industrial organisation that are raised by this model; this we will handle in the topic of market entry.

10.6.2 Market entry

In chapter 3 we considered a simple mechanism of introducing new firms to a market (page 56); but the mechanism was almost mechanical and took no account of the strategic issues involved in the relationship between the incumbent firm(s) and the potential entrants that are challenging them. Here we will use the analysis of time in games as the basis for modelling a strategic model of entry.

The point of departure is the story depicted in Figure 10.12 and Table 10.8. Replace player Alf with a potential entrant firm (here [LEFT] means "enter the industry", [RIGHT] means "stay out") and Bill as the incumbent (so [left] means "fight a potential entrant", [right] means "accommodate a potential entrant"). The numbers in the example depict the case where the incumbent's position is relatively weak and so the subgame-perfect equilibrium is one where the incumbent immediately accommodates the potential entrant without a fight.[26]

[26] Suppose the first two payoff pairs in Figure 10.12 are changed from "$(0,0)$ $(2,1)$" to

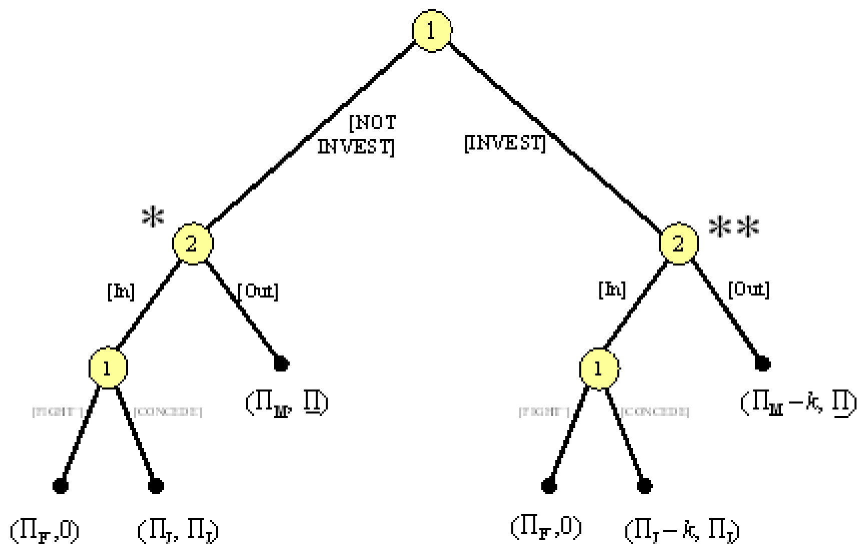

Figure 10.16: Entry deterrence

However, the model is rather naive and inflexible: the relative strength of the positions of the incumbent and the challenger are just hardwired into the payoffs and do not offer much economic insight. What if the rules of the game were altered a little? Could an incumbent make credible threats? The principal way of allowing for this possibility within the model of market structure is to introduce a "commitment device" (see footnote 21). A simple and realistic example of this is where a firm incurs *sunk costs*: this means that the firm spends money on some investment that has no resale value.[27] A simple version of the idea is depicted in Figure 10.16.

Figure 10.16 is based on Figure 10.12 but now there are now three stages of the game. Stages two and three correspond to the story that we have just described; the subgame starting at the node marked * on the left is effectively the same game as we discussed before where the incumbent conceded immediately; in the corresponding subgame on the right-hand side (starting at the node marked **) the payoffs for the incumbent have been changed so that, in this case, it will no longer be profitable to concede entry to the challenger.[28] In the first stage the incumbent makes a decision whether or not to invest an

"$(0,1)$ $(2,0)$." How will this alter the equilibrium of the game? What interpretation can be given in terms of the model of contested entry? What is the equilibrium?

[27] You set up a window cleaning business. You buy a ladder, window cleaning fluid and 1000 leaflets to publicise your business in the neighbourhood. Identify (i) variable costs, (ii) fixed costs, (iii) sunk costs.

[28] Show that the subgame starting at the left-hand node marked 2 in Figure 10.16 is essentially the same game, up to an ordinal transformation of payoffs, as the game in Figure 10.12.

amount that will cost a given amount k: this decision is publicly observable.

The decision on investment is crucial to the way the rest of the game works. The following is common knowledge.

- If the challenger stays out it makes a reservation profit level $\underline{\Pi}$ and the incumbent makes monopoly profits Π_{M} (less the cost of investment if it had been undertaken in stage 1).
- If the incumbent concedes to the challenger then they share the market and each gets Π_{J}.
- If the investment is not undertaken then the cost of fighting is Π_{F}.
- If the investment is undertaken in stage 1 then it is recouped, dollar for dollar, should a fight occur. So, if the incumbent fights, it makes profits of exactly Π_{F}, net of the investment cost.

Now consider the equilibrium. Let us focus first on the subgame that follows on from a decision by the incumbent to invest (for the case where the incumbent does not invest see Exercise 10.11). If the challenger were to enter after this then the incumbent would find that it is more profitable to fight than concede as long as

$$\Pi_{\mathrm{F}} > \Pi_{\mathrm{J}} - k. \tag{10.25}$$

Now consider the first stage of the game: is it more profitable for the incumbent to commit the investment than just to allow the no-commitment subgame to occur? Yes if the net profit to be derived from successful entry deterrence exceeds the best that the incumbent could do without committing the investment:

$$\Pi_{\mathrm{M}} - k > \Pi_{\mathrm{J}}. \tag{10.26}$$

Combining the two pieces of information in (10.25) and (10.26) we get the result that deterrence works (in the sense of having a subgame-perfect equilibrium) as long as k has been chosen such that:

$$\Pi_{\mathrm{J}} - \Pi_{\mathrm{F}} < k < \Pi_{\mathrm{M}} - \Pi_{\mathrm{J}}. \tag{10.27}$$

In the light of condition (10.27) it is clear that, for some values of Π_{F}, Π_{J} and Π_{M}, it may be impossible for the incumbent to deter entry by this method of precommitting to investment.

There is a natural connection with the Stackelberg duopoly model. Think of the investment as advance production costs: the firm is seen to build up a "war-chest" in the form of an inventory of output that can be released on to the market. If deterrence is successful, this stored output will have to be thrown away. However, should the challenger choose to enter, the incumbent can unload inventory from its warehouses without further cost. Furthermore the newcomer's optimal output will be determined by the amount of output that the incumbent will have stashed away and then released. We can then see that the overall game becomes something very close to that discussed in the leader-follower model of section 10.6.1, but with the important difference that the rôle of the leader is now determined in a natural way through a common-sense interpretation of timing in the model.

10.6.3 Another look at duopoly

In the light of the discussion of repeated games (section 10.5.3) it is useful to reconsider the duopoly model of section 10.4.1. Applying the Folk Theorem enables us to examine the logic in the custom and practice of a tacit cartel. The story is the familiar one of collusion between the firms in restricting output so as to maintain high profits; if the collusion fails then the Cournot-Nash equilibrium will establish itself.

First we will oversimplify the problem by supposing that the two firms have effectively a binary choice in each stage game – they can choose one of the two output levels as in the discussion on page 290. Again, for ease of exposition, we take the special case of identical firms and we use the values given in Table 10.5 as payoffs in the stage game:

- If they both choose [low], this gives the joint-profit maximising payoff to each firm, Π_{J}.
- If they both choose [high], gives the Cournot-Nash payoff to each firm, Π_{C}.
- If one firm defects from the collusive arrangement it can get a payoff $\overline{\Pi}$.

Using the argument for equation (10.23) (see also the answer to footnote 23) the critical value of the discount factor is

$$\underline{\delta} := \frac{\overline{\Pi} - \Pi_{\mathrm{J}}}{\overline{\Pi} - \Pi_{\mathrm{C}}}$$

So it appears that we could just carry across the argument of page 304 to the issue of cooperative behaviour in a duopoly setting. The joint-profit maximising payoff to the cartel could be implemented as the outcome of a subgame-perfect equilibrium in which the strategy would involve punishing deviation from cooperative behaviour by switching to the Cournot-Nash output levels for ever after. But it is important to make two qualifying remarks.

First, suppose the market is expanding over time. Let $\tilde{\Pi}(t)$ be a variable that can take the value $\overline{\Pi}$, Π_{J} or Π_{C} Then it is clear that the payoff in the stage game for firm f at time t can be written

$$\Pi^f(t) = \tilde{\Pi}(t)[1+g]^{t-1}$$

where g is the expected growth rate and the particular value of $\tilde{\Pi}(t)$ will depend on the actions of each of the players in the stage game. The payoff to firm f of the whole repeated game is the following present value:

$$
\begin{aligned}
& [1-\delta] \sum_{t=1}^{\infty} \delta^{t-1}\, \Pi^{f}(t) \\
= \; & [1-\delta] \sum_{t=1}^{\infty} \tilde{\delta}^{t-1}\, \tilde{\Pi}(t)
\end{aligned} \tag{10.28}
$$

where $\tilde{\delta} := \delta\,[1+g]$. So it is clear that we can reinterpret the discount factor as a product of pure time preference, the probability that the game will continue and the expected growth in the market. We can see that if the market is expected to be growing the effective discount factor will be higher and so in view of Theorem 10.3 the possibility of sustaining cooperation as a subgame-perfect equilibrium will be enhanced.

Second, it is essential to remember that the argument is based on the simple Prisoner's Dilemma where the action space for the stage game just has the two output levels. The standard Cournot model with a continuum of possible actions introduces further possibilities that we have not considered in the Prisoner's Dilemma. In particular we can see that minimax level of profit for firm f in a Cournot oligopoly is *not* the Nash-equilibrium outcome, Π_{C}. The minimax profit level is zero – the other firm(s) could set output such that the f cannot make a profit (see, for example, point $\overline{q}_2$ in Figure 10.5). However, if one were to set output so as to ensure this outcome in every period from $t+1$ to ∞, this would clearly not be a best response by any other firm to an action by firm f (it is clear from the two-firm case in Figure 10.6 that $(0, \overline{q}_2)$ is not on the graph of firm 2's reaction function); so it cannot correspond to a Nash equilibrium to the subgame that would follow a deviation by firm f. Everlasting minimax punishment is not credible in this case.[29]

10.7 Uncertainty

As we have seen, having precise information about the detail of how a game is being played out is vital in shaping a rational player's strategy – the Kriegsspiel example on page 272 is enough to convince of that. It is also valuable to have clear ideas about the opponents' characteristics a chess player might want to know whether the opponent is "strong" or "weak," the type of play that he favours and so on.

These general remarks lead us on to the nature of the uncertainty to be considered here. In principle we could imagine that the information available to a player in the game is *imperfect* in that some details about the history of the game are unknown (who moved where at which stage?) or that it is *incomplete*

[29] Draw a diagram similar to Figure B.33 to shaw the possible payoff combinations that are consistent with a Nash equilibrium in infinitely repeated subgame. Would everlasting minimax punishment be credible if the stage game involved Bertrand competition rather than Cournot competition?

in that the player does not fully know what the consequences and payoffs will be for others because he does not know what type of opponent he is facing (risk-averse or risk-loving individual? high-cost or low-cost firm?). Having created this careful distinction we can immediately destroy it by noting that the two versions of uncertainty can be made equivalent as far as the structure of the game is concerned. This is done by introducing one extra player to the game, called "Nature." Nature acts as an extra player by making a move that determines the characteristics of the players; if, as is usually the case, Nature moves first and the move that he/she/it makes is unknown and unobservable, then we can see that the problem of incomplete information (missing details about types of players) is, at a stroke, converted into one of imperfect information (missing details about history).

10.7.1 A basic model

We focus on the specific case where each economic agent h has a *type* τ^h. This type can be taken as a simple numerical parameter; for example it could be an index of risk aversion, an indicator of health status, a component of costs. The type indicator is the key to the model of uncertainty: τ^h is a random variable; each agent's type is determined at the beginning of the game but the realisation of τ^h is only observed by agent h.

Payoffs

The first thing to note is that an agent's type may affect his payoffs (if I become ill I may get lower level of utility from a given consumption bundle than if I stay healthy) and so we need to modify the notation used in (10.2) to allow for this. Accordingly, write agent h's utility as

$$V^h\left(s^h, [s]^{-h};\ \tau^h\right) \tag{10.29}$$

where the first two arguments argument consists of the list of strategies – h's strategy and everybody else's strategy as in expression (10.2) – and the last argument is the type associated with player h.

Conditional strategies

Given that the selection of strategy involves some sort of maximisation of payoff (utility), the next point we should note is that each agent's strategy must be conditioned on his type. So a strategy is no longer a single "button" as in the discussion on page 283 but is, rather, a "button rule" that specifies a particular button to each possible value of the type τ^h. Write this rule for agent h as a function $\varsigma^h(\cdot)$ from the set of types to the set of pure strategies S^h. For example if agent h can be of exactly one of two types $\{[\text{HEALTHY}],[\text{ILL}]\}$ then agent h's button rule $\varsigma^h(\cdot)$ will generate exactly one of two pure strategies

$$s_0^h = \varsigma^h([\text{HEALTHY}])$$

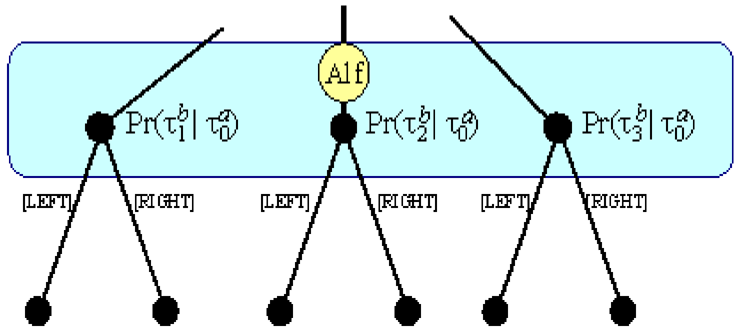

Figure 10.17: Alf's beliefs about Bill

or

$$s_1^h = \varsigma^h\left([\text{ILL}]\right)$$

according to the value of τ^h realised at the beginning of the game.

Beliefs, probabilities and expected payoffs

However, agent h does not know the types of the other agents who are players in the game. instead he has to select a strategy based on some set of *beliefs* about the others' types. These beliefs are incorporated into a simple probabilistic model: F, the joint probability distribution of types over the agents is assumed to be common knowledge. Although it is by no means essential, from now on we will simply assume that the type of each individual is just a number in $[0, 1]$.[30]

Figure 10.17 shows a stylised sketch of the idea. Here Alf, who has been revealed to be of type τ_0^a and who is about to choose [LEFT] or [RIGHT], does not know what Bill's type is at the moment of the decision. There are three possibilities, indicated by the three points in the information set. However, because Alf knows the distribution of types that Bill may possess he can at least rationally assign conditional probabilities $\Pr\left(\tau_1^b|\ \tau_0^a\right)$, $\Pr\left(\tau_2^b|\ \tau_0^a\right)$ and $\Pr\left(\tau_3^b|\ \tau_0^a\right)$ to the three members of the information set, given the type that has been realised for Alf. These probabilities are derived from the joint distribution F, conditional on Alf's own type: these are Alf's beliefs (since the probability distribution of types is common knowledge then he would be crazy to believe anything else).

Consider the way that this uncertainty affects h's payoff. Each of the other agents' strategies will be conditioned on the type which "Nature" endows them and so, in evaluating (10.29) agent h faces the situation that

$$s^h = \varsigma^h\left(\tau^h\right) \tag{10.30}$$

[30] This assumption about types is adaptable to a wide range of specific models of individual characteristics. Show how the two-case example used here, where the person is either of type [HEALTHY] or of type [ILL] can be expressed using the convention that agent h's type $\tau^h \in [0, 1]$ if the probability of agent h being healthy is π.

$$[s]^{-h} = \left[\varsigma^1\left(\tau^1\right), ..., \varsigma^{h-1}\left(\tau^{h-1}\right), \varsigma^{h+1}\left(\tau^{h+1}\right), ...\right] \tag{10.31}$$

The arguments in the functions on the right-hand side of (10.30) and (10.31) are random variables and so the things on the left-hand side of (10.30) and (10.31) are also random. Evaluating (10.29) with these random variables one then gets

$$V^h\left(\varsigma^1\left(\tau^1\right),\ \varsigma^2\left(\tau^2\right), ...;\ \tau^h\right) \tag{10.32}$$

as the (random) payoff for agent h.

In order to incorporate the random variables in (10.30)-(10.32) into a coherent objective function for agent h we need one further step. We assume the standard model of utility under uncertainty that was first introduced in chapter 8 (page 187) – the von Neumann-Morgenstern function. This means that the appropriate way of writing the payoff is in expectational terms

$$\mathcal{E}V^h\left(s^h, [s]^{-h}; \tau^h\right) \tag{10.33}$$

where s^h is given by (10.30), $[s]^{-h}$ is given by (10.31), $\mathcal{E}$ is the expectations operator and the expectation is taken over the joint distribution of types for all the agents.

Equilibrium

We need a further refinement in the definition of equilibrium that will allow for the type of uncertainty that we have just modelled. To do this note that the game can be completely described by three objects, a profile of utility functions, the corresponding list of strategy sets, and the joint probability distribution of types:

$$\left[V^1, V^2, ...\right];\ \left[S^1, S^2, ...\right];\ F \tag{10.34}$$

However, we can recast the game in a way that is familiar from the discussion of section 10.3. We could think of each agent's "button-rule" $\varsigma^h(\cdot)$ as a redefined strategy in its own right; agent h gets utility $v^h\left(\varsigma^h,\ [\varsigma]^{-h}\right)$ which exactly equals (10.33) and where v^h is just the same as in (10.2). If we use the symbol $\mathcal{S}^h$ the set of these redefined strategies or "button rules" for agent h Then (10.34) is equivalent to the game

$$\left[v^1, v^2, ...\right];\ \left[\mathcal{S}^1, \mathcal{S}^2, ...\right] \tag{10.35}$$

Comparing this with (10.3) we can see that, on this interpretation, we have a standard game with redefined strategy sets for each player.

This alternative, equivalent representation of the Bayesian game enables us to introduce the definition of equilibrium:

Definition 10.7 *A* pure strategy Bayesian Nash equilibrium *for (10.34) is a profile of rules* $[\varsigma^*]$ *that is a Nash equilibrium of the game (10.35).*

This definition means that we can just adapt (10.6) by replacing the ordinary strategies ("buttons") in the Nash equilibrium with the "button rules" $\varsigma^{*h}(\cdot)$ where

$$\varsigma^{*h}(\cdot) \in \arg\max_{\varsigma^h(\cdot)} v^h\left(\varsigma^h(\cdot),\ [\varsigma^*(\cdot)]^{-h}\right) \tag{10.36}$$

Identity

The description of this model of incomplete information may seem daunting at first reading, but there is a natural intuitive way of seeing the issues here. Recall that in chapter 8 we modelled uncertainty in competitive markets by, effectively, expanding the commodity space – n physical goods are replaced by $n\varpi$ contingent goods, where ϖ is the number of possible states-of-the-world (page 203). A similar thought experiment works here. Think of the incomplete-information case as one involving players as superheroes where the same agent can take on a number of identities. We can then visualise a Bayesian equilibrium as a Nash equilibrium of a game involving a larger number of players: if there are 2 players and 2 types we can take this setup as equivalent to a game with 4 players (Batman, Superman, Bruce Wayne and Clark Kent). Each agent in a particular identity plays so as to maximise his expected utility in that identity; expected utility is computed using the conditional attached to the each of the possible identities of the opponent(s); the probabilities are conditional on the agent's own identity. So Batman maximises Batman's expected utility having assigned particular probabilities that he is facing Superman or Clark Kent; Bruce Wayne does the same with Bruce Wayne's utility function although the probabilities that he assigns to the (Superman, Clark Kent) identities may be different.

This can be expressed in the following way. Use the notation $\mathcal{E}\left(\bullet|\ \tau_0^h\right)$ to denote conditional expectation – in this case the expectation taken over the distribution of all agents other than h, conditional on the specific type value τ_0^h for agent h – and write $[s^*]^{-h}$ for the profile of random variables in (10.31) at the optimum where $\varsigma^j = \varsigma^{*j}$, $j \neq h$. Then we have:

Theorem 10.4 *A profile of decision rules $[\varsigma^*]$ is a Bayesian Nash equilibrium for (10.34) if and only if for all h and for any τ_0^h occurring with positive probability*

$$\mathcal{E}\left(V^h\left(\varsigma^{*h}\left(\tau_0^h\right), [s^*]^{-h} \mid \tau_0^h\right)\right) \geq \mathcal{E}\left(V^h\left(s^h, [s^*]^{-h} \mid \tau_0^h\right)\right)$$

for all $s^h \in S^h$.

So the rules given in (10.36) will maximise the expected payoff of every agent, conditional on his beliefs about the other agents.

10.7.2 An application: entry again

We can illustrate the concept of a Bayesian equilibrium and outline a method of solution using an example that ties in with the earlier discussion of strategic issues in industrial organisation.

Figure 10.18 takes the story of section 10.6.2 a stage further. The new twist is that the monopolist's characteristics are not fully known by a firm trying to enter the industry. It is known that firm 1, the incumbent, has the possibility of committing to investment that might strategically deter entry: the investment would enhance the incumbent's market position. However the firm may incur either high cost or low cost in making this investment: which of the two cost levels actually applies to firm 1 is something unknown to firm 2. So the game involves first a preliminary move by "Nature" (player 0) that determines the cost type, then a simultaneous move by firm 1, choosing whether or not to invest, and firm 2, choosing whether or not to enter. Consider the following three cases concerning firm 1's circumstances and behaviour:

1. *Firm 1 does not invest.* If firm 2 enters then both firms make profits Π_{J}. But if firm 2 stays out then it just makes its reservation profit level $\underline{\Pi}$, where $0 < \underline{\Pi} < \Pi_{\mathrm{J}}$, while firm 1 makes monopoly profits Π_{M}.

2. *Firm 1 invests and is low cost.* If firm 2 enters then firm 1 makes profits $\Pi_{\mathrm{J}}^* < \Pi_{\mathrm{J}}$ but firm 2's profits are forced right down to zero. If firm 2 stays out then it again gets just reservation profits $\underline{\Pi}$ but firm 1 gets enhanced monopoly profits $\Pi_{\mathrm{M}}^* > \Pi_{\mathrm{M}}$.

3. *Firm 1 invests and is high cost.* Story is as above, but firm 1's profits are reduced by an amount k, the cost difference.

To make the model interesting we will assume that k is fairly large, in the following sense:

$$k > \max\left\{\Pi_{\mathrm{J}}^* - \Pi_{\mathrm{J}}, \Pi_{\mathrm{M}}^* - \Pi_{\mathrm{M}}\right\}.$$

In this case it is never optimal for firm 1 to invest if it has high cost (check the bottom right-hand part of Figure 10.18 to see this).

To find the equilibrium in this model we will introduce a device that we used earlier in section 10.3.3. even though we are focusing on pure (i.e. non-randomised) strategies let us suppose that firm 1 and firm 2 each consider a randomisation between the two actions that they can take. To do this, define the following:[31]

- π^0 is the probability that "Nature" endows firm 1 with low cost. This probability is common knowledge.

- π^1 is the probability that firm 1 chooses [INVEST] given that its cost is low.

- π^2 is the probability that firm 1 chooses [In].

[31] Write out the expressions for epected payoff for firm 1 and for firm 2 and verify (10.37) and (10.39).

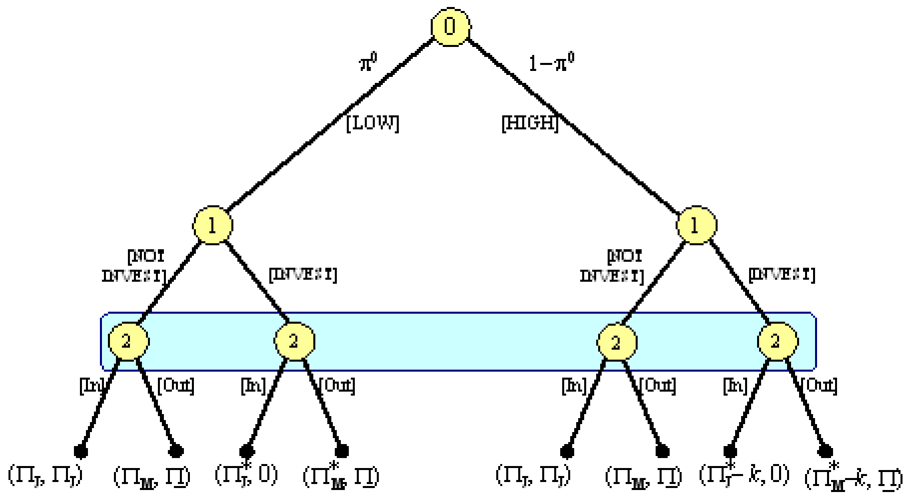

Figure 10.18: Entry with incomplete information

Then, writing out the expected payoff to firm 1, $\mathcal{E}\Pi^1$ we find that:

$$\frac{\partial \mathcal{E}\Pi^1}{\partial \pi^1} > 0 \Longleftrightarrow \pi^2 < \frac{1}{1+\gamma} \tag{10.37}$$

where

$$\gamma := \frac{\Pi_{\mathrm{J}} - \Pi_{\mathrm{J}}^*}{\Pi_{\mathrm{M}}^* - \Pi_{\mathrm{M}}} > 0. \tag{10.38}$$

Furthermore, evaluating $\mathcal{E}\Pi^2$, the expected payoff to firm 2:

$$\frac{\partial \mathcal{E}\Pi^2}{\partial \pi^2} > 0 \Longleftrightarrow \pi^1 < \frac{\Pi_{\mathrm{J}} - \underline{\Pi}}{\pi^0 \Pi_{\mathrm{J}}}. \tag{10.39}$$

The restriction on the right-hand of (10.39) only makes sense if the probability of being low-cost is large enough, that is, if

$$\pi^0 \geq 1 - \frac{\underline{\Pi}}{\Pi_{\mathrm{J}}}. \tag{10.40}$$

To find the equilibrium in pure strategies[32] check whether conditions (10.37)-(10.39) can be satisfied by probability pairs (π^1, π^2) equal to any of the values $(0,0)$, $(0,1)$, $(1,0)$ or $(1,1)$. Clearly condition (10.37) rules out $(0,0)$ and $(1,1)$. However the pair $(0,1)$ always satisfies the conditions, meaning that ([NOT INVEST],[In]) is always a pure-strategy Nash equilibrium. Likewise, if the probability of [LOW] is large enough that condition (10.40) holds, then ([INVEST],[Out]) will also be a pure-strategy Nash equilibrium.

[32] Will there also be a mixed-strategy equilibrium to this game?

The method is of interest here as much as is detail of the equilibrium solutions. It enables us to see a link with the solution concept that we introduced on page 283.

10.7.3 Mixed strategies again

One of the features that emerges from the description of Bayesian Nash equilibrium and the example in section 10.7.2 is the use of probabilities in evaluating payoffs. The way that uncertainty about the type of one's opponent is handled in the Bayesian game appears to be very similar to the resolution of the problem arising in elementary games where there is no equilibrium in pure strategies. The assumption that the distribution of types is common knowledge enables us to focus on a Nash equilibrium solution that is familiar from the discussion of mixed strategies in section 10.3.3.

In fact one can also establish that a mixed-strategy equilibrium with given players Alf, Bill, Charlie,... each of whom randomise their play, is equivalent to a Bayesian equilibrium in which there is a continuum of a-types all with Alf's preferences but slightly different types, a continuum of b-types all with Bill's preferences but with slightly different types,... and so on, all of whom play pure strategies.

The consequence of this is that there may be a response to those who see strategic arguments relying on mixed strategies as artificial and unsatisfactory (see page 285). Large numbers and variability in types appear to "rescue" the situation by showing that there is an equivalent, or closely approximating Bayesian-Nash equilibrium in pure strategies.

10.7.4 A "dynamic" approach

The discussion of uncertainty thus far has been essentially static in so far as the sequencing of the game is concerned. But it is arguable that this misses out one of the most important aspects of incomplete information in most games and situations of economic conflict. With the passage of time each player gets to learn something about the other players' characteristics through observation of the other players' actions at previous stages; this information will be taken into account in the way the game is planned and played out from then on.

In view of this it is clear that the Bayesian Nash approach outlined above only captures part of essential problem. There are two important omissions:

1. *Credibility.* We have already discussed the problem of credibility in connection with Nash equilibria of multi-stage games involving complete information (see pages 299 ff). The same issue would arise here if we considered multi-stage versions of games of incomplete information.

2. *Updating.* As information is generated by the actions of players this can be used to update the probabilities used by the players in evaluating expected utility. This is typically done by using Bayes' rule (see Appendix A, page 518).

So in order to put right the limitations of the uncertainty model one would expect to combine the "perfection" involved in the analysis of subgames with the logic of the Bayesian approach to handling uncertainty. This is exactly what is done in the following further refinement of equilibrium

Definition 10.8 *A* perfect Bayesian equilibrium *in a multi-stage game is a collection of strategies of beliefs at each node of the game such that:*

1. *the strategies form a subgame-perfect equilibrium, given the beliefs;*
2. *the beliefs are updated from prior beliefs using Bayes' rule at every node of the game that is reached with positive probability using the equilibrium strategies.*

The two parts of the definition show a nice symbiosis: the subgame-perfect strategies at every "relevant" node make use of the set of beliefs that is the natural one to use at that point of the game; the beliefs are revised the light of the information that is revealed by playing out the strategies.

However, note that the definition is limited in its scope. It remains silent about what is supposed to happen to beliefs out of equilibrium – but this issue raises complex questions and takes us beyond the scope of the present book. Note too, that in some cases the updating may be simple and drastic so that the problem of incomplete information is resolved after one stage of the game. However, despite these qualifications, the issue of strategic interactions that incorporate learning is so important and so multifaceted that we shall be devoting all of chapter 11 to it.

10.8 Summary

Strategic behaviour is not just a new microeconomic topic but a new method and a fresh way of looking at economic analysis. Game theory permits the construction of an abstract framework that enables us to think through the way economic models work in cases where the simplified structure of price-taking is inapplicable or inappropriate.

But how much should one expect from game theory? It clearly provides a collection of important general principles for microeconomics. It also offers some truly striking results, for example the demonstration that cooperative outcomes can be induced from selfish agents by the design of credible strategies that involve future punishment for "antisocial" behaviour (the folk theorem). On the other hand game theory perhaps warrants an enthusiasm that is tempered by considerations of practicality. Game theoretic approaches do not always give clear-cut answers but may rather point to a multiplicity of solutions and, where they do give clear-cut answers in principle, these answers may be almost impossible to work out in practice. To illustrate: finding all the outcomes in chess is a computable problem, but where is the computer that could do the job?

To summarise the ways in which this chapter has illustrated the contribution of the game-theoretic approach to economic principles and to point the forward to later chapters let us focus on three key aspects:

- *The nature of equilibrium.* In moving to an economic environment in which strategic issues are crucial we have had to introduce several new definitions of equilibrium; in the formal literature on this subject there are even more intellectual constructions that are candidates for equilibrium concepts. Do the subtle differences between the various definitions matter? Each can be defended as the correct way of modelling coherence of agents' behaviour in a carefully specified strategic setting. Each incorporates a notion of rationality consistent with this setting. However, as the model structure is made richer, the accompanying structure of beliefs and interlocking behaviour can appear to be impossibly sophisticated and complex. The difficulty for the economic modeller is, perhaps, to find an appropriate location on the spectrum from total naivety to hyper-rationality (more on this in chapter 12).

- *Time.* The sequencing of decisions and actions is a crucial feature of many situations of potential economic conflict because it will often affect the way the underlying game is played and even the viability of the solution concept. A modest extension of fairly simple games to more than one period enables one to develop models that incorporate the issues of power, induced cooperation.

- *Uncertainty* In chapter 8 uncertainty and risk appeared in economic decision making in the rôle of mechanistic chance. Here, the mechanistic chance can be a player in the game and clear-cut results carry over from the complete-information case, although they rest on quite strong assumptions about individual beliefs and understanding of the uncertain universe. However, we can go further. The Bayesian model opens the possibility of using the acquisition of information strategically and has implications for how we model the economics of information. This is developed in chapter 11.

10.9 Reading notes

A good introduction is provided by Dixit and Skeath (2004), Gardner (2003), Osborne (2004) or Rasmusen (2001); the older Gibbons (1992) still provides an excellent and thorough overview of the main issues; for a more advanced treatment Vega-Redondo (2003) is useful. The Nash equilibrium concept first appeared in Nash (1951); on the appropriateness of using it as a solution concept see Kreps (1990). The rationale of mixed-strategy equilibria is discussed in Harsanyi (1973) and the argument for treating "nature" as a player in the game is developed in Harsanyi (1967). For the history and precursors of the concept of Nash equilibrium see Myerson (1999); on Nash equilibrium and behaviour see

Mailath (1998) and Samuelson (2002). Subgame-perfection as an equilibrium concept is attributable to Selten (1965, 1975).

The folk theorem and variants on repeated games form a substantial literature. For an early statement in the context of oligopoly see Friedman (1971). A key result establishing sub-game perfection in repeated games is proved in Fudenberg and Maskin (1979).

The standard reference on industrial organisation is the thorough treatment by Tirole (1988); the original classic contributions whose logic underlies so much modern work are to be found in Bertrand (1883), Cournot (1838) and von Stackelberg (1934).

10.10 Exercises

10.1 *Table 10.10 is the strategic form representation of a simultaneous move game in which strategies are actions.*

	s_1^b	s_2^b	s_3^b
s_1^a	$0, 2$	$3, 1$	$4, 3$
s_2^a	$2, 4$	$0, 3$	$3, 2$
s_3^a	$1, 1$	$2, 0$	$2, 1$

Table 10.10: Elimination and equilibrium

1. *Is there a dominant strategy for either of the two agents?*
2. *Which strategies can always be eliminated as individually irrational?*
3. *Which strategies can be eliminated if it is common knowledge that both players are rational?*
4. *What are the Nash equilibria in pure strategies?*

10.2 *Table 10.11 again represents a simultaneous move game in which strategies are actions.*

	s_1^b	s_2^b	s_3^b
s_1^a	$0, 2$	$2, 0$	$3, 1$
s_2^a	$2, 0$	$0, 2$	$3, 1$
s_3^a	$1, 3$	$1, 3$	$4, 4$

Table 10.11: Pure-strategy Nash equilibria

1. *Identify the best responses for each of the players a, b.*
2. *What are the Nash equilibria in pure strategies?*

10.3 *A taxpayer has income y that should be reported in full to the tax authority. There is a flat (proportional) tax rate γ on income. The reporting technology means that that taxpayer must report income in full or zero income. The tax authority can choose whether or not to audit the taxpayer. Each audit costs an amount φ and if the audit uncovers under-reporting then the taxpayer is required to pay the full amount of tax owed plus a fine F.*

1. *Set the problem out as a game in strategic form where each agent (taxpayer, tax-authority) has two pure strategies.*
2. *Explain why there is no simultaneous-move equilibrium in pure strategies.*
3. *Find the mixed-strategy equilibrium. How will the equilibrium respond to changes in the parameters γ, φ and F?*

10.4 *Take the "battle-of-the-sexes" game of footnote 3 (the strategic form is given in Table B.1 on page 562).*

1. *Show that, in addition to the pure strategy, Nash equilibria there is also a mixed strategy equilibrium.*
2. *Construct the payoff-possibility frontier (as in Figure B.33 on page 566). Why is the interpretation of this frontier in the battle-of-the-sexes context rather unusual in comparison with the Cournot-oligopoly case?*
3. *Show that the mixed-strategy equilibrium lies strictly inside the frontier.*
4. *Suppose the two players adopt **the same** randomisation device, observable by both of them: they know that the specified random variable takes the value 1 with probability π and 2 with probability $1-\pi$; they agree to play $\left[s_1^a, s_1^b\right]$ with probability π and $\left[s_2^a, s_2^b\right]$ with probability $1-\pi$; show that this **correlated mixed strategy** always produces a payoff on the frontier.*

10.5 *Rework Exercise 10.4 for the case of the game in Table 10.3 (this is commonly known as the Chicken game).*

10.6 *Consider the three-person game depicted in Figure 10.19 where strategies are actions. For each strategy combination, the column of figures in parentheses denotes the payoffs to Alf, Bill and Charlie, respectively. (Fudenberg and Tirole 1991, page 55)*

1. *For the simultaneous-move game shown in Figure 10.19 show that there is a unique pure-strategy Nash equilibrium.*
2. *Suppose the game is changed. Alf and Bill agree to coordinate their actions by tossing a coin and playing [LEFT],[left] if heads comes up and [RIGHT],[right] if tails comes up. Charlie is not told the outcome of the spin of the coin before making his move. What is Charlie's best response? Compare your answer to part 1.*

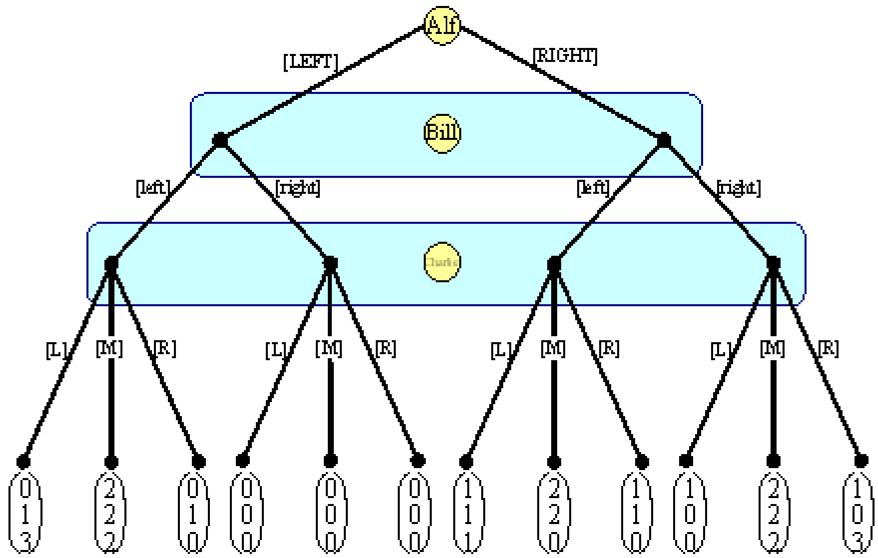

Figure 10.19: Benefits of restricting information

3. *Now take the version of part 2 but suppose that Charlie knows the outcome of the coin toss before making his choice. What is his best response? Compare your answer to parts 1 and 2. Does this mean that restricting information can be socially beneficial?*

10.7 *Consider a duopoly with identical firms. The cost function for firm f is*

$$C_0 + cq^f, f = 1, 2.$$

The inverse demand function is

$$\beta_0 - \beta q$$

where C_0, c, β_0 and β are all positive numbers and total output is given by $q = q^1 + q^2$.

1. *Find the isoprofit contour and the reaction function for firm 2.*
2. *Find the Cournot-Nash equilibrium for the industry and illustrate it in (q^1, q^2)-space.*
3. *Find the joint-profit maximising solution for the industry and illustrate it on the same diagram.*
4. *If firm 1 acts as leader and firm 2 as a follower find the Stackelberg solution.*

5. *Draw the set of payoff possibilities and plot the payoffs for cases 2-4 and for the case where there is a monopoly.*

10.8 *An oligopoly contains N identical firms. The cost function is convex in output. Show that if the firms act as Cournot competitors then as N increases the market price will approach the competitive price.*

10.9 *Two identical firms consider entering a new market; setting up in the new market incurs a once-for-all cost $k > 0$; production involves constant marginal cost c. If both firms enter the market Bertrand competition then takes place afterwards. If the firms make their entry decision sequentially, what is the equilibrium?*

10.10 *There is a cake of size 1 to be divided between Alf and Bill. In period $t = 1$ Alf offers player Bill a share: Bill may accept now (in which case the game ends), or reject. If Bill rejects then, in period $t = 2$ Alf again makes an offer, which Bill can accept (game ends) or reject. If Bill rejects, the game ends one period later with exogenously fixed payoffs of γ to Alf and $1 - \gamma$ to Bill. Assume that Alf and Bill's payoffs are linear in cake and that both persons have the same, time-invariant discount factor $\delta < 1$.*

1. *What is the backwards induction outcome in the two-period model?*
2. *How does the answer change if the time horizon increases but is finite?*
3. *What would happen if the horizon were infinite? (Rubinstein 1982, Ståhl 1972, Sutton 1986)*

10.11 *Take the game that begins at the node marked "*" in Figure 10.16 (page 307).*

1. *Show that if $\Pi_{\mathrm{M}} > \Pi_{\mathrm{J}} > \Pi_{\mathrm{F}}$ then the incumbent firm will always concede to a challenger.*
2. *Now suppose that the incumbent operates a chain of N stores, each in a separate location. It faces a challenge to each of the N stores: in each location there is a firm that would like to enter the local market). The challenges take place sequentially, location by location; at each point the potential entrant knows the outcomes of all previous challenges. The payoffs in each location are as in part 1 and the incumbent's overall payoff is the undiscounted sum of the payoffs over all locations. Show that, however large N is, all the challengers will enter and the incumbent never fights (Selten 1978).*

10.12 *In a monopolistic industry firm 1, the incumbent, is considering whether to install extra capacity in order to deter the potential entry of firm 2. Marginal capacity installation costs, and marginal production costs (for production in excess of capacity) are equal and constant. Excess capacity cannot be sold. The* *potential entrant incurs a fixed cost k in the event of entry.(Dixit 1980, Spence 1977)*

1. *Let $\underline{q}^1$ be the incumbent's output level for which the potential entrant's best response yields zero profits for the entrant. Suppose $\underline{q}^1 \neq q_{\mathrm{M}}$, where q_{M} is firm 1's output if its monopolistic position is unassailable (i.e. if entry-deterrence is inevitable). Show that this implies that market demand must be nonlinear.*

2. *In the case where entry deterrence is possible but not inevitable, show that if $q_{\mathrm{S}}^1 > \underline{q}^1$, then it is more profitable for firm 1 to deter entry than to accommodate the challenger, where q_{S}^1 is firm 1's output level at the Stackelberg solution..*

10.13 *Two firms in a duopolistic industry have constant and equal marginal costs c and face market demand schedule given by $p = k - q$ where $k > c$ and q is total output..*

1. *What would be the solution to the Bertrand price setting game?*

2. *Compute the joint-profit maximising solution for this industry.*

3. *Consider an infinitely repeated game based on the Bertrand stage game when both firms have the discount factor $\delta < 1$. What trigger strategy, based on punishment levels $p = c$, will generate the outcome in part 2? For what values of δ do these trigger strategies constitute a subgame perfect Nash equilibrium?*

10.14 *Consider a market with a very large number of consumers in which a firm faces a fixed cost of entry F. In period 0, N firms enter and in period 1 each firm chooses the quality of its product to be* HIGH*, which costs $c > 0$, or* LOW*, which costs 0. Consumers choose which firms to buy from, choosing randomly if they are indifferent. Only after purchasing the commodity can consumers observe the quality. In subsequent time periods the stage game just described is repeated indefinitely. The market demand function is given by*

$$q = \begin{cases} \varphi(p) & \textit{if quality is believed to be}\textrm{HIGH} \\ \\ 0 & \textit{otherwise} \end{cases}$$

where $\varphi(\cdot)$ is a strictly decreasing function and p is the price of the commodity. The discount rate is zero.

1. *Specify a trigger strategy for consumers which induces firms always to choose high quality. Hence determine the subgame-perfect equilibrium. What price will be charged in equilibrium?*

2. *What is the equilibrium number of firms, and each firm's output level in a long-run equilibrium with free entry and exit?*

3. What would happen if $F = 0$?

10.15 *In a duopoly both firms have constant marginal cost. It is common knowledge that this is* 1 *for firm 1 and that for firm 2 it is either* $\frac{3}{4}$ *or* $1\frac{1}{4}$*. It is common knowledge that firm 1 believes that firm 2 is low cost with probability* $\frac{1}{2}$*. The inverse demand function is*

$$2 - q$$

where q is total output. The firms choose output simultaneously. What is the equilibrium in pure strategies?