3 Cardinal Welfarism

3.1 Welfarism

The welfarist postulate states that the distribution of individual welfare across the agents/citizens is the only legitimate yardstick along which the states of the world can be compared. In the cardinal version of welfarism,[1] individual welfare is measured by an index of *utility,* and comparisons of the utilities of two different agents are meaningful.

Welfarism is a reductionist model of distributive justice. It views an agent as a machine producing welfare/utility at a given state of the world, and compares two feasible states by means of the two utility profiles they generate. Welfarism is endstate justice at its best; that is to say, the process by which a particular utility profile is reached (e.g., the physical allocation of the resources of the world) is devoid of ethical content: it is the means toward achieving a particular profile of utilities. For instance, the criterion "no envy" discussed in chapters 6 and 7 is irrelevant to welfarism because it relies on interpersonal comparisons of individual allocations of resources.

The most basic concept of welfarism[2] is efficiency-fitness (Pareto optimality) of which we repeat the definition already given in section 1.3. Consider two feasible states of the world x and y, resulting for agent i in the utility $u_i(x)$ and $u_i(y)$ respectively. State y is Pareto superior to state x if no agent j strictly prefers state x to state y, that is to say, $u_j(x) \leq u_j(y)$ for all j, and moreover at least one agent i strictly prefers state y, $u_i(x) < u_i(y)$. Thus y is Pareto superior to x if the move from x to y is by unanimous consent (in the sense that everyone agrees to a change that does not decrease his or her own utility level). A state x is *Pareto optimal* (efficient) if there is no feasible state y Pareto superior to x. Thus, if the current state is x, we cannot generate a consensus to move to another state y (except in the case where everyone enjoys the same utility level in both states).

The task of cardinal welfarism is to pick, among the feasible utility profiles (lists of one utility level per agent), one of the Pareto optimal ones. In many specific allocation problems, the Pareto optimality property has much bite, in the sense that it eliminates many if not most feasible allocations of the resources. Examples 3.2, 3.3, 3.8, and 3.9. On the other hand, in the fair division problems of section 2.5, all feasible allocations are Pareto optimal because there is a single commodity and everyone prefers a bigger share to a smaller one. The only way to increase agents i's utility is to give him more resources, which in turn decreases the utility of another agent. The same property (all feasible allocations are Pareto optimal) holds true in examples 3.1, 3.4, 3.5, and 3.6.

1. The ordinal version of welfarism is called social choice; it is the subject of chapter 4.
2. Be it in its cardinal or ordinal version.

A general fact is that the property of efficiency/Pareto optimality is orthogonal to distributive concerns. Typically the allocation where all of the available resources are used to the benefit of a single agent meets the criterion, although it is the most unfair of all distributive systems.

The task of the welfarist benevolent dictator is to compare normatively any two utility profiles (u_i), (u'_i) and decide which one is best. The key idea is to insist that this comparison should follow the rationality principles of individual decision-making, namely *completeness* and *transitivity*.[3] Completeness says that any two profiles can always be compared: either (u_i) is preferred to (u'_i), denoted $(u_i) \succ (u'_i)$, or the reverse preference $(u'_i) \succ (u_i)$ holds, or they are declared indifferent $(u_i) \sim (u'_i)$. Transitivity means that $(u_i) \succsim (u'_i)$—that profile (u_i) is preferred or indifferent to profile (u'_i)—and $(u'_i) \succsim (u''_i)$ imply $(u_i) \succsim (u''_i)$.

The preference relation is called a *social welfare ordering,* and the definition and comparison of various social welfare orderings is the object of cardinal welfarism. The two most prominent instances of social welfare orderings are the *classical utilitarian,* namely $(u_i) \succsim (u'_i)$, if and only if $\sum_i u_i \geq \sum_i u'_i$, and the *egalitarian one,* namely $(u_i) \succ (u'_i)$ if and only if, upon reordering the profile by increasing coordinates as (u^*_i) and $(u^{*'}_i)$, the former is lexicographically superior to the latter. The former expresses the sumfitness principle in the welfarist world, whereas the latter conveys the compensation principle. A variety of social welfare orderings in between those two are introduced in the subsequent sections.

Before we start the general discussion of social welfare orderings in section 3.2, it is important to recall that we focus exclusively on the "micro" version of welfarism, and pay only lip service to its "macro" interpretation still an influential idea in contemporary political philosophy. We look at microallocation problems, involving a small number of commodities and where utility is tailored to the problem at hand. For instance, in the problem of locating a facility (examples 3.4 and 3.8), utility measures the (negative of the) distance between the agent in question and the facility. In example 3.11 the issue is to distribute fruits from which our agents metabolize vitamins, and utility is measured by the quantity of such vitamins. And so it goes on. Thus the context dictates the interpretation of utility and, in turn, influences the choice of the social welfare ordering (see the distinction between tastes and needs in the examples just mentioned). The central assumption that individual utilities can be objectively measured and compared across different agents can be more or less convincing. The distance from an agents' home to the facility (examples 3.4, 3.5, and 3.8) is an objective fact, as is, to a large extent, the amount of a certain vitamin or drug he needs to be healthy. But his taste for a certain piece of cake, or for art, cannot be measured along a common scale.

3. A general discussion of the rationality of choice for a single decision-maker is in section 4.1.

The microwelfarist viewpoint separates the allocation problem at stake from the rest of our agent's characteristics. It assumes that the level of my utility from the allocation I receive in the microproblem can be measured independently from the rest of my characteristics. Moreover the utility of those agents not concerned by the microallocation problem should not matter either. This crucial property of *separability* is expressed axiomatically in the next section and is the basis of the additive representation discussed there.

From this axiomatic analysis, three paramount social welfare orderings emerge. In addition to the classical utilitarian collective utility function and the egalitarian social welfare ordering already mentioned, the Nash collective utility function is simply the product of individual utilities. From the theoretical discussion of section 3.2, as well as the examples of sections 3.4 and 3.5, the Nash collective utility function emerges as a sensible compromise between the egalitarian and classical utilitarian ones. The nontechnical reader is urged to skip section 3.2 and go directly to the examples listed from section 3.3 on.

In contrast to microwelfarism, macrowelfarism is an encompassing approach to social justice, where utilities measure the overall level of happiness of a given agent/citizen (the sum of his pleasures and pains, in Bentham's words), so that the choice of a social welfare ordering amounts to an entire program of social justice.

Recall from section 1.3 the two main objections to macrowelfarism. An objective grasp of individual welfare defeats the purpose of methodological individualism. And ignoring individual responsibility in the formation of one's own welfare is morally untenable.

A popular macrowelfarist method to take into account agent's responsibility in the formation of their own welfare is to use proxy commodities (called *primary goods* by Rawls) as surrogate measurement of individual welfare. The idea is that our ability or inability to lead full and satisfying lives, to achieve high or low levels of welfare, is determined by our share of certain fundamental goods: food, shelter, health, self-respect, love, education, wealth, job, and so on. The catalog of these goods is the common denominator of human nature. The actual distribution of these primary goods tells us all that we can hope to learn about actual welfares; hence it can be used as a surrogate measurement of the distribution of welfare. If the trick of primary goods, due to Rawls, succeeds in maintaining a private sphere around each person, while offering an objective index of (access to) welfare, it pushes the difficulty back without eliminating it. The choice of a method to aggregate my holdings of primary goods into a summary "index" is tantamount (although in a less obvious way) to imposing a common value system upon all individual citizens: it imposes the same trade-offs among primary goods on all citizens. Yet an important component of my value system is my own trade-off between health and wealth (e.g., I may choose a physically dangerous, well-paid job) or between education and leisure, and so on.

*3.2 Additive Collective Utility Functions

In this section we describe the most important axiomatic results of the cardinal theory of welfarism. We state a handful of axiomatic requirements pertaining to collective rationality or fairness. Then we deduce the small family of collective utility functions (and social welfare orderings) satisfying these requirements. In this family three key social welfare orderings stand out: the classical utilitarian and Nash collective utility functions and the leximin social welfare ordering.

Two basic requirements of a social welfare ordering $\succeq$, beside the property of completeness and transitivity already mentioned, are *monotonicity* and *symmetry*. The social welfare ordering $\succeq$ is monotonic in utility u_i if an increase in agent i's utility, ceteris paribus, increases social welfare. That is, if two utility profiles $u = (u_j)$ and $u' = (u'_j)$ are such that $u_j = u'_j$ for all j, $j \neq i$, and $u_i > u'_i$, then $u \succ u'$; that is, the social welfare ordering prefers the former profile to the latter. Monotonicity has much to do with Pareto optimality: a monotonic social welfare ordering is compatible with the Pareto relation[4] in the sense that if u is Pareto superior to v, then $u \succ v$. In particular, the maximal elements of a monotonic social welfare ordering, over any feasible set of utility profiles, are Pareto optimal.

The social welfare ordering $\succeq$ is symmetric if it does not pay attention to the identity of the agents, only to their utility level. If the utility profile u obtains from v simply by permuting the index of the agents in arbitrary fashion, as (7, 2, 8, 2, 4, 4, 2) obtains from (4, 8, 4, 2, 7, 2, 2), the social welfare ordering views these two profiles as equivalent: $u \sim v$. Symmetry is equal treatment of equals, namely the basic fairness axiom discussed in section 1.1: agents can only be discriminated on the basis of their utilities, not of any other exogenous factors.

Most social welfare orderings of importance[5] are represented by a *collective utility function,* namely a real-valued function $W(u_1, u_2, \ldots, u_n)$ with the utility profile for argument and the level of collective utility for value. The function W represents the social welfare ordering $\succsim$ if $u \succsim u'$ is logically equivalent to $W(u) \geq W(u')$. The monotonicity and symmetry properties of $\succsim$ translate into the properties with the same name, for W : W is strictly monotonic in each of the variables u_i and a symmetric function of the profile.

The next property is the key argument of welfarist rationality. It says that we can ignore the unconcerned agents when choosing between two particular utility profiles u and u'. That is, if agent j receives the same utility in both profiles, $u_j = u'_j$, his utility level has no influence on the comparison of u and u'. Formally, we denote by $(u_i \mid^j a)$ the utility vector identical to (u_i) except that the jth coordinate has been replaced by a. Then the property

4. This relation is transitive but not complete; therefore it is not a social welfare ordering.

5. The leximin social welfare ordering is a notable exception, discussed in the next section.

independence of unconcerned agents reads as follows:

$$(u_i \mid^j a) \succsim (u_i' \mid^j a) \Leftrightarrow (u_i \mid^j b) \succsim (u_i' \mid^j b) \qquad \text{for all } u, u', j, a, \text{ and } b \tag{1}$$

This means that an agent who has no vested interest in the choice between u and u'—because his utility is the same in both profiles—can be ignored.

If property (1) fails, the choice between two particular states of the world will depend on the utility of some agents who are truly indifferent between these two states. This runs counter to the intuition of endstate justice. Absent the property, the set of agents whose utilities can influence the social welfare ordering (whether or not they are personally affected by the subsequent decisions) must be defined precisely for any microproblem of distributive justice. By contrast, under independence of unconcerned agents, microjustice works very well with a loose, encompassing set of potentially concerned agents. The property guarantees that the social welfare ordering (and the associated collective utility function) focuses exclusively on those agents whose utility is affected by the decisions to be made.

The collective utility function W is called *additive* if there is an increasing function g of one real variable such that

$$W(u) = \sum_i g(u_i) \qquad \text{for all } u \tag{2}$$

It should be clear that the social welfare ordering represented by an additive collective utility function meets property (1). If we restrict attention to continuous[6] social welfare orderings, the following converse property holds. If the continuous ordering $\succsim$ is independent of unconcerned agents, then it is represented by an additive collective utility (2). This important theorem gives a convenient representation of a rich family of social welfare orderings.

We introduce two additional properties of the collective utility (2), that limit the choice of the function g. The first property is one of fairness, and expresses an aversion for "pure" inequality. It is called the *Pigou-Dalton transfer principle*. Say that $u_1 < u_2$ at profile u and consider a transfer of utility from agent 2 to agent 1 where u_1' and u_2', the utilities after the transfer, are such that

$$u_1 < u_1', u_2' < u_2 \quad \text{and} \quad u_1' + u_2' = u_1 + u_2$$

Thus total utility to agents 1 and 2 is preserved, and the inequality gap is reduced (note that it could be reversed: $u_2' < u_1'$ is possible). We say that the move from u to u' (where $u_j = u_j'$ for $j \geq 3$) reduces the inequality between agents 1 and 2. The Pigou-Dalton transfer

6. The social welfare ordering $\succsim$ is continuous for all u, the sets $\{v \mid v \succsim u\}$ and $\{v \mid u \succsim v\}$, called respectively the upper and lower contour sets of u, are topologically closed.

principle requires that the social welfare ordering increases (or at least, does not decrease) in a move reducing the inequality between any two agents.

Applying this principle to the additive collective utility (2) we find that

$$\{u_1 < u'_1, u'_2 < u_2 \text{ and } u_1 + u_2 = u'_1 + u'_2\} \Rightarrow \{g(u_1) + g(u_2) \leq g(u'_1) + g(u'_2)\}$$

which is equivalent to the concavity of the function g, namely its derivative is nonincreasing.

The next property is one of invariance. It is called *independence of common scale* (ICS). The property requires us to restrict attention to positive utilities, a feature that is automatically satisfied in most of our examples where the zero of utilities corresponds to the minimal feasible level; see examples 3.1, 3.2, and 3.5. The ICS property states that a simultaneous rescaling of every individual utility function does not affect the underlying social welfare ordering; it yields the same binary comparisons of utility profiles:

$$u \succsim u' \Leftrightarrow \lambda u \succsim \lambda u' \tag{3}$$

where the two profiles u, u' as well as the scaling constant λ are positive and otherwise arbitrary. For instance, if utilities represent money = willingness to pay (as in chapter 5), it does not matter if we compare cents, dozens of dollars, or thousands of dollars: the order of magnitude of the utility levels under comparison does not matter.

For an additive collective utility taking the form (2), the ICS property holds true only for a very specific family of power functions. To see this, apply (3) to the function (2), which yields

$$\sum_i (g(u_i) - g(u'_i)) \geq 0 \Leftrightarrow \sum_i (g(\lambda u_i) - g(\lambda u'_i)) \geq 0$$

It can further be shown that the only (increasing, continuous) functions g satisfying this property are (up to a multiplicative constant) of exactly three types:

$$g(z) = z^p \qquad \text{for a positive } p$$

$$g(z) = \log(z)$$

$$g(z) = -z^{-q} \qquad \text{for a positive } q$$

The corresponding collective utility W take the form

$$\begin{aligned} W_p(u) &= \sum_i u_i^p, \qquad \text{with } p > 0 \text{ and fixed} \\ W_0(u) &= \sum_i \log u_i \\ W^q(u) &= -\sum_i \frac{1}{u_i^q} \qquad \text{with } q > 0 \text{ and fixed} \end{aligned} \tag{4}$$

The family (4) has many interesting features. First, the particular collective utility function $\sum_i \log u_i$ is the limit of the other two families when p or q, respectively, approach zero.[7] It is called the *Nash collective utility* function, and is usually written in the equivalent multiplicative form $W_N(u) = \Pi_i u_i$. Of course, the function W_N is not additively decomposed as in (2), but it represents the same social welfare ordering $\succsim$ as the additive collective utility W_0.

Another collective utility function of interest within the family (4) is the classical utilitarian $W_1(u) = \sum_i u_i$, corresponding to $p = 1$. It embodies the idea of sum-fitness, and its implications are discussed in a variety of examples in the subsequent sections.

Finally we examine the impact of the Pigou-Dalton transfer principle on the family of utility functions (4). Consider the quadratic $W_2(u) = \sum_i u_i^2$. Far from seeking to reduce inequality, this collective utility function is actively promoting it. For instance, the following mathematical fact

$$u_1^2 + u_2^2 < (u_1 + u_2)^2 + (0)^2$$

implies that under W_2, transferring all the utility to one agent is desirable. Such a preference runs counter to the basic distributive fairness conveyed by the Pigou-Dalton principle.

As noted earlier, an additive collective utility (2) meets the Pigou-Dalton principle if and only if g is a concave function. Within the family (4), this eliminates all the functions W_p with $1 < p < +\infty$ and only those.

We are ready to sum up the results of our axiomatic discussion. Starting with a continuous collective utility function W representing the social welfare ordering $\succsim$, we imposed successively three requirements: independence of unconcerned individuals (1), independence of common scale (3), and the Pigou-Dalton transfer principle. Together, these properties leave us with a one-dimensional family of collective utility functions, namely

$$\begin{aligned} W_p(u) &= \sum u_i^p, 0 < p \leq 1 \\ W_0(u) &= \sum_i \log u_i \\ W^q(u) &= -\sum_i u_i^{-q}, 0 < q < +\infty \end{aligned} \tag{5}$$

Notice the striking similarity of the formula above with the family of equal sacrifice methods presented in section 2.4.

7. To see this, use the approximation $z^p = e^{p \log z} \simeq 1 + p \log z$, valid when p is close to zero.

Although it is defined in three pieces, the family (5) is actually continuous in the sense that W_0 is the limit of W_p and of W^q as p or q goes to zero. Two outstanding elements are the classical utilitarian W_1 and the Nash utility function W_0 (often written in multiplicative form as W_N). The third remarkable point of the family is the limit of (the social welfare ordering represented by) W^q as q goes to infinity: this is the important leximin social welfare ordering, defined and illustrated in the next section.

The three social orderings, classical utilitarian, Nash, and leximin are the three most important objects of cardinal welfarism. They are systematically compared in sections 3.4 and 3.5.

3.3 Egalitarianism and the Leximin Social Welfare Ordering

We focus in this section on the welfarist formulation of the compensation principle as the equalization of individual utilities. Full equalization is often impossible within the set of feasible outcomes—as in examples 3.1, 3.2, 3.4, and 3.5; in other cases it is feasible but incompatible with Pareto optimality—see example 3.3. The leximin social welfare ordering then selects the most egalitarian among the Pareto optimal utility distributions.

Example 3.1 Pure Lifeboat Problem As in example 2.1, some but not all agents can be allowed on the boat, and the arbitrator must choose which subset will be saved. She can pick from a given list of subsets. Suppose that five agents are labeled {1, 2, 3, 4, 5} and that the feasible subsets are

{1, 2}{1, 3}{1, 4}{2, 3, 5}{3, 4, 5}{2, 4, 5}

Thus agents 1 and 5 cannot both be in the lifeboat; we can have one of agents 2, 3, or 4 along with 1, or two of these three along with 5.

A less dramatic story is the purchase of a software program that will be available to our five agents: there are six programs to choose from, and each program is only compatible with the machines of a certain subset of agents. Or we must choose the musical background in the office space occupied by our five agents; they are six programs to choose from and a given agent likes certain programs and dislikes others: only agents 1 and 2 like the first program, and so on.[8] Note that each one of the six feasible outcomes is Pareto optimal: there is no unanimous agreement to dismiss any one of the six outcomes.

Suppose first that for each agent the utility of staying on the boat is 10 and that of swimming is 1. Then the classical utilitarian utility recommends choosing one (any one) of

8. In example 3.1 the arbitrator must choose one of the six subsets, with no possibility of compromise by randomization or timesharing. The latter is the subject of example 3.6b. See also exercise 3.6b, question c.

the three largest subsets (each one with three agents). The egalitarian arbitrator makes the same choice, based on comparing the increasing profile of utilities from lowest to highest. If a two-person subset stays on the boat, this profile is (1, 1, 1, 10, 10) and if a three-person subset stays, it is (1, 1, 10, 10, 10) which is lexicographically superior because the third ranked utility level is higher in the latter profile.

The point of the example is much sharper when we assume that the individual utility for being among the chosen ones varies across agents—such as in the radio program interpretation, some agents are more partial to "good" versus "bad" music than others:

Agent	1	2	3	4	5
Utility for good outcome	10	6	6	5	3
For bad outcome	0	1	1	1	0

Now the classical utilitarian arbitrator prefers to choose $\{1, 2\}$ or $\{1, 3\}$, for a total utility of 16, over any other subset; the second best is $\{2, 4, 5\}$ yielding total utility 15. His ranking of the six outcomes is as follows:

$$\{1, 2\} \sim \{1, 3\} \succ \{1, 4\} \sim \{2, 3, 5\} \succ \{2, 4, 5\} \sim \{3, 4, 5\}$$

The egalitarian arbitrator, by contrast, prefers any three-person subset over any two-person one; his ranking is as follows:

		$\{2, 3, 5\}$	with	utility	profile	(0, 1, 3, 6, 6)
$\{2, 4, 5\}$	$\sim$	$\{3, 4, 5\}$	with	utility	profile	(0, 1, 3, 5, 6)
$\{1, 2\}$	$\sim$	$\{1, 3\}$	with	utility	profile	(0, 1, 1, 6, 10)
		$\{1, 4\}$	with	utility	profile	(0, 1, 1, 5, 10)

Exercise 3.1 contrasts the classical utilitarian and egalitarian choices in example 3.1 for arbitrary utility functions.

Example 3.2 Fair Division with Identical Preference We must divide six indivisible objects among three agents, and each lot must contain two objects. Individual preferences over the different lots are identical: given any two lots, everyone agrees on which one is the better lot, or everyone is indifferent between the two lots.

The leximin social welfare ordering compares all feasible allocations, and does not require one to attach a common cardinal utility to each lot. For instance, assume that the common ordering of the fifteen lots from the objects a, b, c, d, e, and f is as follows:

$$\{a, b\} > \{b, f\} \sim \{b, e\} > \{c, d\} > \{a, c\} > \{d, e\} \sim \{b, c\} > \{c, f\} > \{a, d\} > \{a, e\}$$
$$\sim \{c, e\} \sim \{e, f\} > \{b, d\} \sim \{a, f\} \sim \{d, f\}$$

There are fifteen ways to split the six objects in three lots of two objects, and the leximin ordering allows us to compare all fifteen. For instance, {a, b}{c, d}{e, f} is ranked above {a, f}{d, e}{b, c}, because in the latter, one lot yields the worst welfare, whereas *all* lots yield a higher welfare in the former. Next compare {a, e}{b, f}{c, d} and {a, d}{b, f}{c, e}: both yield the same lowest welfare level (at {a, e} and {a, d} respectively); the latter gives this level to two agents and the former to only one agent: therefore the former division is better. Here the leximin social welfare ordering picks {a, c}{b, f}{d, e} as the unambiguous best split, followed by {a, b}{c, f}{d, e}: both guarantee the same lowest welfare level (for whomever gets {d, e}), but the latter gives that level to two of the three agents, against only one in the former. The fact that one agent gets the absolute best lot {a, b} in the latter split is unimportant.

We give now the general definition of the leximin social welfare ordering, also called the egalitarian social welfare ordering, and sometimes “practical egalitarianism.” Given two feasible utility profiles u and u', we rearrange them first in increasing order, from the lowest to the highest utility, and denote the new profiles u^* and u'^* : $u_1^* \leq u_2^* \leq \cdots \leq u_n^*$ and $u_1'^* \leq u_2'^* \leq \cdots \leq u_n'^*$. The leximin social welfare ordering compares u^* and u'^* lexicographically. Thus $u \succ u'$ holds if $u_1^* \succ u_1'^*$ ($u_1'^* \succ u_1^*$ implies similarly that $u' \succ u$): if the lowest utility is higher in one profile than in the other, this is enough to declare it a better profile. If $u_1^* = u_1'^*$, the leximin ordering compares the second lowest utilities u_2^* and $u_2'^*$; if they differ, the profile with the higher one is preferred. Thus $\{u_1^* = u_1'^*$ and $u_2^* > u_2'^*\}$ implies that $u \succ u'$. And so on: if the k lowest utility levels coincide in both profiles ($u_1^* = u_1'^*$ for $i = 1, \ldots, k$) and the $(k+1)$ lowest differ, the latter determines the preferred profile.

The mathematical definition of the leximin ordering is slightly more involved than that of any additive collective utility in the family (5). In fact this ordering cannot be represented by *any* collective utility function. On the other hand, leximin belongs to the family (5) in a limit sense: as q goes to infinity, the social welfare ordering represented by the collective utility function W^q converges to the leximin one. Moreover the leximin ordering is independent of unconcerned individuals, independent of the common scale of utility and satisfies the Pigou-Dalton transfer principle.[9]

In many examples, such as examples 3.3 and 3.4, finding the maximum of the leximin ordering reduces to maximizing the first component $u_1^* = \min_i u_i$ of the utility profile rearranged in increasing order. In such cases we are simply maximizing the egalitarian collective utility function $W_e(u) = \min_i u_i$, meaning that we maximize the utility of the

9. Exercise 3.12 states formally the limit property and discusses these properties of the leximin.

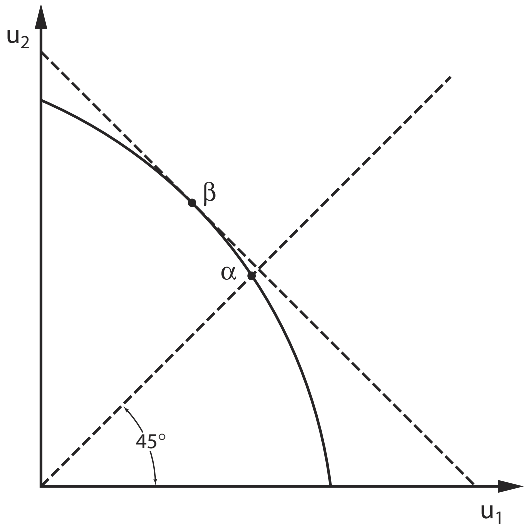

Figure 3.1
No equality/efficiency trade-off

worst-off individual. Of course W_e is not a proper representation of the leximin ordering, because $u \succsim u'$ implies that $W_e(u) \geq W_e(u')$, but the converse implication does not hold.

Example 3.3 The Equality/Efficiency Trade-off with Two Agents Consider the feasible utility sets in figures 3.1 and 3.2. We do not specify from what allocation problem these utility profiles come from. Under the welfarist postulate (section 3.1) this does not matter.

In figure 3.1 there is a fully egalitarian and efficient utility profile α. This profile is the unique maximizer of the egalitarian collective utility W_e, hence the leximin optimum as well. There is no conflict between equality and efficiency.

In figure 3.2, by contrast, the profile α is the highest feasible equal utility one, but it is not efficient: both agents enjoy a higher utility at profile α^* that maximizes the egalitarian utility W_e (and the leximin ordering). Here we have a trade-off between equality and efficiency. The egalitarian collective utility function justifies the inequality at α^* to augment the utility of the worst off agent.

The configurations in figures 3.3 and 3.4 are similar with an equality/efficiency trade-off in the latter (successfully resolved by the maximum of W_e) but not in the former.

We conclude this section with a crucial—indeed a characteristic—property of the leximin ordering. Recall that in example 3.2 all we need to define the egalitarian division in lots is the ability to rank any two lots. More generally, consider two utility profiles u and u'.

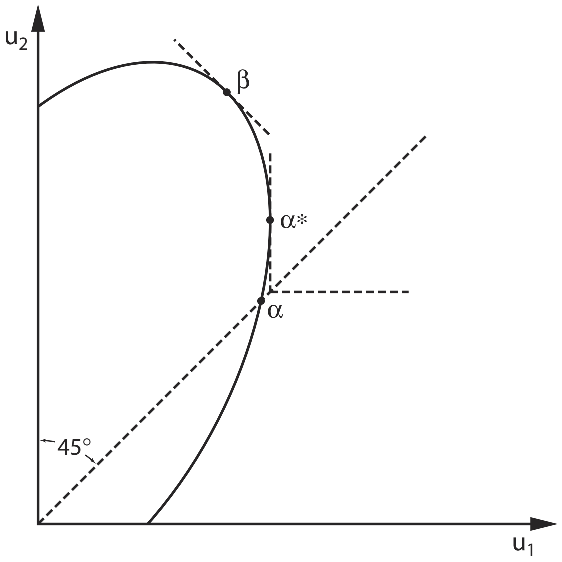

Figure 3.2
Equality/efficiency trade-off

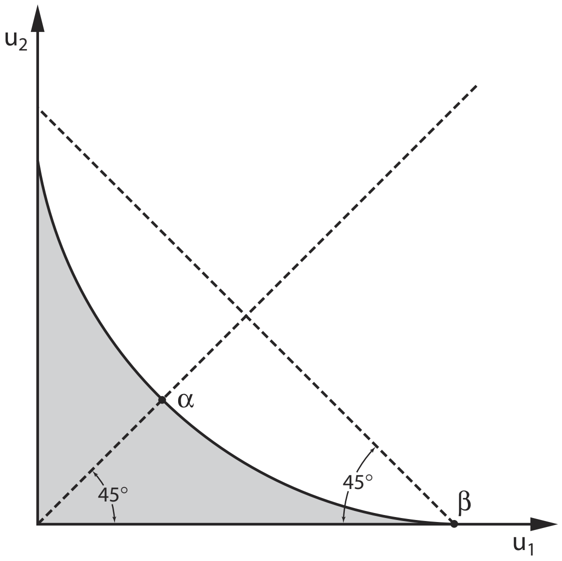

Figure 3.3
No equality/efficiency trade-off

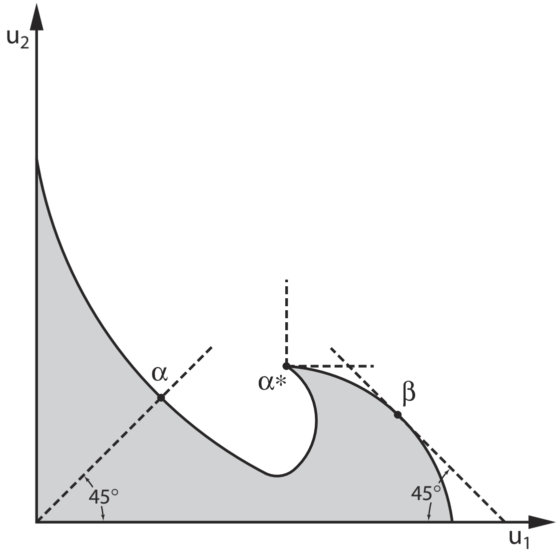

Figure 3.4
Equality/efficiency trade-off

Suppose that for any one of the n^2 pairs $u_i, u'_j, i, j = 1, \ldots, n$, we know their relative ranking, meaning that we know which one of $u_i > u'_j, u_i < u'_j$, or $u_i = u'_j$ holds. This is enough to deduce the ranking of u versus u' in the leximin ordering. For instance, let us take four agents and the following pattern:

	u'_1	u'_2	u'_3	u'_4
u_1	$>$	$<$	$>$	$<$
u_2	$<$	$<$	$=$	$<$
u_3	$>$	$>$	$>$	$>$
u_4	$=$	$<$	$>$	$<$

where an entry $>$ reads: the row u_i is greater than the column u'_j.
From this pattern we deduce the rankings of u_i, u'_j:

$$u_2 = u'_3 < u_4 = u'_1 < u_1 < u'_4 < u'_2 < u_3$$

Hence

$$u_1^* = u_1'^*, u_2^* = u_2'^* \quad \text{and} \quad u_3^* = u_1 < u_3'^* = u_4$$

We conclude that u' is preferred to u by the leximin ordering.

The invariance property underlying the example above is that the leximin ordering is preserved under a common arbitrary (in particular, nonlinear) rescaling of the utilities. Thus the comparison of u versus u' is the same as that of $v = (u_i^2)$ versus $v' = (u_i'^2)$, or of $(e^{u_i + \sqrt{u_i}})$ versus $(e^{u'_i + \sqrt{u'_i}})$, and so on. This property is called *independence of the common utility pace.*

Leximin is not the only social welfare ordering independent of the common utility pace. But it is the only one that also respects the Pigou-Dalton transfer principle.[10] This remarkable characterization result explains why this social welfare ordering occupies such a central place in cardinal welfarism.

In the next two sections, we compare systematically the three basic social orderings—leximin, classical utilitarian, and Nash—and explain along the way the axiomatic characterizations of the latter two.

3.4 Comparing Classical Utilitarianism, Nash, and Leximin

In section 3.2 we identified three outstanding welfarist solutions as the classical utilitarian and Nash collective utility functions, and the leximin social welfare ordering. We now compare them in a series of examples, where we stress their relation to the compensation and sum-fitness principles. Thus this section and the next pursue a discussion initiated in section 2.5, by testing our three welfarist solutions in more general problems of resource allocations.

The central tension between the classical utilitarian and egalitarian welfarist objectives was already uncovered in section 2.5. They are advocating different kinds of sacrifices. Under the former, the welfare of a single agent may be sacrificed for the sake of improving total welfare (the slavery of the talented—example 3.9—is a striking case in point). Under the latter, large amounts of joint welfare may be forfeited in order to improve the lot of the worst off individual (e.g., examples 3.4 and 3.6).

Example 3.4. Location of a Facility A desirable facility (examples are given below) must be located somewhere in the interval [0, 1], representing a "linear" city. Each agent i lives at a specific location x_i in [0, 1]; if the facility is located at y, agent i's *dis*utility is the distance $|y - x_i|$. The agents are spread arbitrarily along the interval [0, 1], and the problem is to find a fair compromise location.

The egalitarian solution is the easiest to compute. Suppose that there are some agents living at 0 and some living at 1. Then the egalitarian collective disutility function W_e equals

10. Exercise 3.12 offers more discussion of these two facts.

$\frac{1}{2}$ when the facility is located at $y_e = \frac{1}{2}$. For any other location y, the distance from y to 0 or to 1 is larger than $\frac{1}{2}$. Thus the unique egalitarian optimum is the *midpoint* of the range of our agents.

Classical utilitarianism chooses the *median* of the distribution of agents, namely the point y_u such that at most half of the agents live strictly to the left of y_u and at most half of them live strictly to its right. To see why this is so, observe that a move of ε to the right of y_u increases by ε the disutility of at least one-half of the agents (i.e., those located at y_u or to its left) while reducing by ε that of at most one-half. A similar argument shows that a move to the left cannot reduce the sum total of individual disutilities.[11]

The interpretation of the facility has much to do with the choice between the two solutions. If the facility is a swimming pool, or an information booth, then the utilitarian choice is more appealing because it minimizes overall transportation costs, and we accept to sacrifice the isolated rural resident: disutility is interpreted as distaste and the isolated agents have chosen freely to be so. On the other hand, if the facility meets a basic need, such as a post office or a police station, the egalitarian compromise has more appeal, because equal access to the facility is tantamount to meeting this need equally. Some cases are more ambiguous: if we are locating a fire station, the goals of equal access and of maximizing the expected return (i.e., the expected reduction of property losses) are both valid, but they pull us toward the midpoint and the median respectively.

The Nash collective utility function is not easy to use in this example because the natural zero of individual utilities is when the facility is located precisely where the agent in question lives, say x_i: then we set $u_i(y) = -\,|y - x_i|$ if the facility is located at y. The Nash utility is not defined when some utilities are negative; therefore we must adjust the zero of each agent so as to ensure she gets nonnegative utility for any choice of y. One way to do so is to set the zero of an agent's utility where the distance from his location to the facility is 1:

$$u_i^1(y) = 1 - |y - x_i|$$

Another way is to set agent i's zero where y is as far as can be from his location, namely at $y = 0$ if $x_i \geq \frac{1}{2}$ and at $y = 1$ if $x_i \leq \frac{1}{2}$. This yields the following utility:

$$u_i^2(y) = x_i - |y - x_i| \quad \text{if } i \text{ is such that } x_i \geq \tfrac{1}{2}$$

$$u_j^2(y) = 1 - x_j - |y - x_j| \quad \text{if } j \text{ is such that } x_j \leq \tfrac{1}{2}$$

11. In section 4.3 we give an alternative interpretation of the median location y as the Condorcet winner outcome: for any other location y, more than half of the citizens prefer y_u to y.

Clearly, the choice of one or the other of the two normalizations is of great consequence on the optimal location for the Nash collective utility.[12]

The great advantage of the classical utilitarian utility is to be *independent of individual zeros of utilities*. If we replace utility $u_i = -|y - x_i|$ by u_i^1 or u_i^2 above for any number of agents, the optimal utilitarian location remains the median of the distribution and the preference ranking between any two locations does not change. This independence property uniquely characterizes the classical utilitarian among all collective utility functions.[13]

Notice that the egalitarian location remains at $y_e = \frac{1}{2}$ with all utility functions normalized as u_i^1, or all as u_i^2 (exercise: Why?), but this is not a general feature. For instance, if no one lives near the location 0, the optimal egalitarian location under u_i^1 moves above $\frac{1}{2}$ but remains at $\frac{1}{2}$ under u_i^2.

Example 3.5 Location of a Noxious Facility Now the facility is a toxic waste disposal, a jail, or any other operation from which everyone wants to live as far away as possible. The distance from x_i, where agent i lives, to the facility at y measures her *utility,* instead of her *dis*utility in the previous example.

If the agents are spread all over the interval [0, 1], the egalitarian collective utility is zero everywhere: $W_e(y) = 0$ for all y because there is always someone living at y. The leximin social welfare ordering, on the other hand, wants to locate the noxious facility at a point where the *density* of agents is lowest. If there are several such points, it breaks ties in favor of a location where the second derivative of the density is lowest (exercise: Why?).

The utilitarian collective utility $W_1(y)$ is now minimal at the median y_u and maximal at one of the two endpoints 0 or 1.[14] Thus it is enough to compare $W_1(0)$ and $W_1(1)$. Denoting by $f(x)\,dx$ the population density at x, we compute

$$W_1(0) < W_1(1) \Leftrightarrow \int_0^1 xf(x)\,dx < \int_0^1 (1-x)f(x)\,dx$$

$$\Leftrightarrow Ex = \int_0^1 xf(x)\,dx < \tfrac{1}{2}$$

We conclude that the utilitarian location is the endpoint farthest away from the mean location Ef.

12. And in each case this optimum is neither easy to compute nor to interpret.

13. Even those that are not independent of unconcerned agents (see section 3.2).

14. Each utility function is convex in the variable y. Therefore so is W_1; a convex function reaches its maximum over an interval at one of the endpoints.

We turn to a simple, yet important model where the problem is to pick a fair compromise between several pure public goods. The model brings to the fore the contrasting distributive policies of our three basic collective utility functions.

Example 3.6a Time-Sharing The n agents work in a common space (e.g., fitness room) where the radio must be turned on one of five available stations (one of them may be the "off" station). As their tastes differ greatly, they ask the manager to share the time fairly between the five stations.

Each agent likes some stations, and dislike some; if we set her utility at 0 or 1 for a station she dislikes or likes, respectively, we have a pure lifeboat problem as in example 3.1. The difference is that we allow mixing between the five decisions: the manager chooses a list of timeshares x_k, $k = 1, \ldots, 5$ such that $x_k \geq 0$ and $x_1 + x_2 + \cdots + x_5 = 1$.

In this example we assume a simple preference pattern that makes it easy to compute and contrast the solution chosen by our three basic collective utility functions. Each agent likes exactly one station and dislikes the other four, with corresponding utilities at 1 and 0. There are n_k fans of station k, with $n_1 + \cdots + n_5 = n$.

The classical utilitarian manager chooses the "tyranny of the majority": the station with the largest support is on all the time (and if there are several such stations, any mixing between them is optimal as well). The egalitarian manager does exactly the opposite, namely it pays not attention to the size of support and plays each station $\frac{1}{5}$th of the time (provided that each station has at least one fan) so that everyone is happy 20 percent of the time.

The Nash collective utility picks an appealing compromise between the two extremist solutions above. The relative sizes of n_k matter *and* everyone is guaranteed some share of her favorite station. The optimal times shares x_k for the Nash utility maximize $\sum_k n_k \log x_k$ under the constraint $\sum_k x_k = 1$; therefore $x_k^* = n_k/n$, namely the time share of each station is proportional to the number of its fans. This can be interpreted as random dictatorship: each agent gets to choose the station he likes for $1/n$th of the time.

The proportional time shares make good sense in the radio-sharing story, because we interpret utilities as subjective tastes for one type of music or the other. Alternative interpretations of utilities yield a very different intuition.

Consider a pure lifeboat decision, where we must choose, literally, whom to save: say five boats are about to sink, and we can only help one of them. Who would hesitate to give all his help to the most populated boat, as utilitarianism recommends? Flipping a fair coin to decide which boat to help allows the rescuer to give an equal chance of survival to every person, but our claim is that he won't and that utilitarianism is compelling here.

In some other contexts utilities may measure the satisfaction of a need: our agents are away from home, and the radio broadcasts news from their hometown station. They come from five different towns and station k gives news from town k only. Now the egalitarian solution makes a lot of sense!

Example 3.6b Time-Sharing Five agents share a radio as in example 3.6a, and the preferences of three of them (agents 3, 4, and 5) are somewhat flexible, in the sense that they like two of the five stations according to the following pattern:

		Station				
		a	b	c	d	e
	1	1	0	0	0	0
	2	0	1	0	0	0
Agent	3	0	0	1	1	0
	4	0	0	0	1	1
	5	0	0	1	0	1

The utilitarian manager shares the time between the three stations c, d, and e but never plays stations a or b. The egalitarian manager selects $x_a = x_b = \frac{2}{7}$, $x_c = x_d = x_e = \frac{1}{7}$, so that everyone listens to a program he or she likes 28.6 percent of the time. The utilitarian solution is too harsh on agents 1, 2, while the egalitarian solution appears too soft on these two agents: agents 3, 4, 5 should be somewhat rewarded for the flexibility of their preferences.

The Nash collective utility function recommends a sensible compromise between utilitarianism and egalitarianism: it plays each station with equal probability $\frac{1}{5}$. To check this, we note that outcomes a, b play a symmetrical role, hence are allocated the same time share x; similarly each one of c, d, and e receives the same time share y. The Nash maximization problem is now

$$\text{maximize} \quad x^2(2y)^3 \quad \text{under} \quad x, y \geq 0, 2x + 3y = 1$$

A straightforward computation gives the optimal solution $x^* = y^* = \frac{1}{5}$.

To conclude this section, we illustrate the great advantage of the Nash collective utility function in the variant of example 3.6a where individual utilities for listening to the "right" kind of music differ across agents: a supporter i of station k enjoys utility u_i if k is on and 0 otherwise. Both the classical utilitarian and egalitarian collective utility functions pay a great deal of attention to the relative intensities of these utilities. For instance, the egalitarian arbitrator computes for each k the smallest individual utility a_k among the fans of station k and allocates to k a time share proportional $1/a_k$ (exercise: prove this claim). And classical utilitarianism may end up broadcasting exclusively a station with a handful of very vocal supporters.

The Nash utility function, by contrast, is *independent of individual scale of utilities*. In our example, this means that the intensity u_i of agent i's musical pleasure is irrelevant to the choice of a fair time-sharing. Indeed, for a profile of time shares $(x_1, \ldots, x_5)$ the collective utility W_N is computed as follows:

$$W_N = \sum_{k=1}^{5} \sum_{i \in N_k} \log(u_i . x_k)$$
$$= \left(\sum_i \log\ u_i \right) + \sum_{k=1}^{5} n_k \log\ x_k$$

where N_k is the set of k fans. The maximization of this collective utility is independent of the numbers u_i; hence the Nash arbitrator still recommends giving station k a time share proportional to n_k.

Independence of individual scales of utilities eliminates the possibility of influencing the arbitrator's choice by distorting the intensity of one's utility for good music. If the arbitrator is classical utilitarian, it is clearly to agent's i advantage to increase u_i (which may make his favorite station look best); if he is egalitarian, it is advantageous to decrease u_i (exercise: prove these claims). The Nash collective utility function is immune to both kinds of distortions, which is a considerable advantage in a context where utility measures subjective tastes. Notice that the profitability of increasing one's utility scale under classical utilitarianism (resp. to decrease it under egalitarianism) is fully general (not limited to our simple example): exercise 3.14 explains this important property. Similarly, under an egalitarian or a Nash arbitrator, it is always profitable to increase one's zero of utility.

The Nash collective utility function is uniquely characterized among all collective utility functions,[15] by the property independence of individual scales. Thus each of the three basic social welfare orderings is characterized by a specific independence property: the independence of common utility pace picks the leximin social welfare ordering,[16] and the independence of individual zeros captures the classical utilitarian collective utility.[17]

3.5 Failures of Monotonicity

Some paradoxical features of welfarism affect our three basic collective utility functions. The main issue is how the optimal solution reacts when the resources of the economy change and was already discussed in section 2.5. There we noticed that the property may fail under classical utilitarianism if some individual utility functions are not concave in the

15. Even those that fail independence of unconcerned agents.

16. See the discussion at the end of section 3.3.

17. As discussed after example 3.4.

amount of resources. Below we give examples displaying the same failure with concave utility functions both for the classical utilitarian and Nash arbitrators.

For the egalitarian solution, resource monotonicity is always satisfied if the optimal allocation gives equal utilities to all agents—namely if there is no equality/efficiency trade-off (example 3.3). This important fact is obvious: if my utility and yours remain equal whenever an arbitrary parameter affecting the allocation of resources changes, we will both benefit or we will both suffer from the change, or we will both be unaffected. Shocks do not break the egalitarian harmony.

However, practical egalitarianism (i.e., maximization of the leximin ordering) may lead to the paradoxical feature. An elementary example involves indivisible goods.

Ann, Bob, and Chris want to play tennis, but they do not have enough racquets (balls and courts are not scarce). Ann enjoys playing against the wall as much as against Bob or Chris. Bob and Chris hate to play the wall and only enjoy playing against each other or against Ann. If only one racquet is available, the leximin ordering tells us to give it to Ann (indeed, this is the only efficient allocation of the resources). If two racquets are available, the egalitarian solution is an equitable time-sharing arrangement where the three pairs take turns on the court and everyone plays $\frac{2}{3}$ of the time: thus Ann is worse off after the resource increase.

Our next fair division example has divisible goods and concave utility functions, therefore no equality/efficiency trade-off. Thus the egalitarian optimum has equal utilities for both agents, and resource monotonicity is automatically verified. On the other hand, maximizing the Nash or classical utilitarian collective utility leads to monotonicity failures.

Example 3.7 Dividing Complementary Goods Jones and Smith both use a different mix of two goods labeled A and B. To produce one unit of utility, Jones needs one unit of A for two of B while Smith needs two units of A for one of B.

Jones: $u_1(a_1, b_1) = \min\{2a_1, b_1\}$

Smith: $u_2(a_2, b_2) = \min\{a_2, 2b_2\}$

A "serious" example is the mix of labor and capital to produce a certain service: good A is labor and good B a certain machine, both measured in hours; Jones's technology is less labor intensive than Smith's. A "frivolous" one involves a cocktail of two liquors that they mix in different proportions.

Suppose first that 12 units of each good are available. The set of feasible utility profiles is depicted on figure 3.5a. Any collective utility function that does not like inequality (i.e., meeting the Pigou-Dalton transfer principle in section 3.2) chooses the equal utility profile $(u_1, u_2) = (8, 8)$ coming from the allocation $(a_1, b_1) = (4, 8)$, $(a_2, b_2) = (8, 4)$.

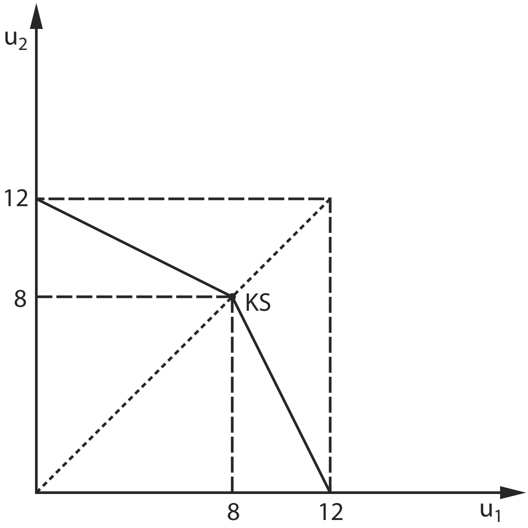

Figure 3.5a
Feasible utility profiles in example 3.7

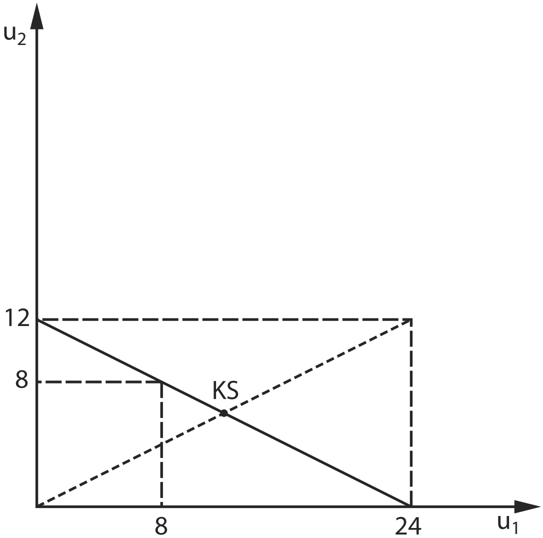

Figure 3.5b
Modified feasible utility set

Next suppose that 12 more units of good B are available, for a total of 12 units of A and 24 units of B. The new feasible set of utility profiles is depicted on figure 3.5b: it has increased to the triangle of all pairs (u_1, u_2) such that $u_1 + 2u_2 = 24$ and $u_i \geq 0$: for any x, $0 \leq x \leq 12$, the utility profile $(u_1, u_2) = (24 - 2x, x)$ comes from $(a_1, b_1) = (12 - x, 24 - 2x)$, $(a_2, b_2) = (x, 2x)$. The egalitarian arbitrator still picks the utility profile (8, 8): the additional resources are simply discarded. The classical utilitarian solution gives all the resources to agent 1 so that $(u_1, u_2) = (24, 0)$. Agent 2's utility loss is less severe under the Nash solution: the corresponding optimal utility profile is $(u_1, u_2) = (12, 6)$, namely the solution of the following problem:

$$\max \log u_1 + \log u_2 \quad \text{under constraint} \quad u_1 + 2u_2 = 24, u_i \geq 0, \quad i = 1, 2$$

The failure of resource monotonicity in the example above is generalized in example 7.12 to many more solutions than the two above.

We give now a pure public good example where both the egalitarian and utilitarian solutions fail resource monotonicity.

Example 3.8 Location of a Facility (continued) In this variant of example 3.4, the road network is depicted on figure 3.6. The agents live on the roads *AB*, *BC*, *CD*, and *DE* where

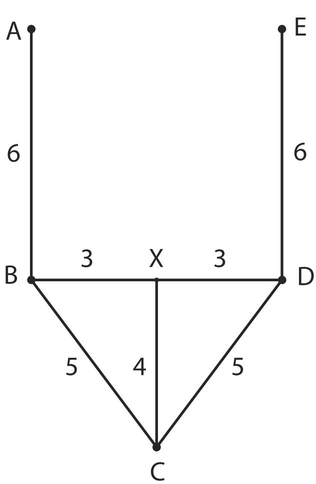

Figure 3.6
Road network in example 3.8

their density is constant and equal to one. Nobody lives on *BX*, *XC*, and *XD*. As before, the distance of the facility to one's location is the disutility (negative of utility).

Suppose first that the facility can only be located wherever the agents live (thus the inner roads *BX*, *XC*, and *XD* are not feasible locations). Then by the symmetry of our problem, the egalitarian and the classical utilitarian collective utilities are maximized at C. Actually C is the optimal location for *any* collective utility function, by virtue of symmetry and monotonicity (exercise: prove this claim).

Now suppose every point of the road network is a feasible location of the facility. This is unambiguously an increase of the available resources. The optimal location for the egalitarian collective utility is now at X, from where the distance to any agent is at most 9 miles, whereas with the facility at C, agents living at A or F are 11 miles away. Thus the agents living at C and near C see their utility decrease as the resources improve.[18]

The optimal location for the classical utilitarian is also X (exercise: Why?). Therefore resource monotonicity fails for this solution as well.

Our last paradoxical example is a famous one. It shows the utilitarian and Nash solutions penalizing the more productive agent in a distinctly unpalatable fashion.

Example 3.9 Slavery of the Talented In this simple production economy, two agents can use their labor to produce corn, at a constant productivity s_i, $i = 1, 2$: one unit of agent i's labor produces s_i units of corn. Agents consume corn and leisure, and these two goods are perfect complements: if agent i consumes z_i units of corn and y_i of leisure, her final utility is $\min\{z_i, y_i\}$. Finally each agent can split 20 hours of time between x_i units of labor and y_i of leisure: $x_i + y_i = 20$.

Consider first the benchmark case where both agents are equally productive at $s_1 = s_2 = 1$. The efficient production plan treating the two agents equally (i.e., respecting the symmetry of the problem) is the "decentralized" outcome where each agent keeps the corn he produces; hence $x_i = y_i = z_i = 10$, $i = 1, 2$, and each agent gets 10 utils, $u_i = 10$. The egalitarian, Nash, or any collective utility function that is strictly averse to inequality[19] picks this allocation uniquely. The utilitarian function is an exception, as it is indifferent to inequality.[20]

Next suppose that agent 1's productivity raises to $s_1' = 2$, while agent 2's productivity remains $s_2 = 1$. Efficiency commands to give agent i exactly the same amount of leisure

18. Those living no more than 4 miles away from C.

19. That is to say, a collective utility that increases from a Pigou-Dalton transfer (section 3.2).

20. Here all efficient allocations maximize the utilitarian function. These allocations are parametrized by the amount of labor $x_1 = \lambda$ supplied by agent 1, an arbitrary number between 0 and 20. Then $u_1 = y_1 = z_1 = 20 - \lambda$, $x_2 = 20 - \lambda$, $u_2 = y_2 = z_2 = \lambda$.

and of corn, i.e., $z_i = 20 - x_i$. The feasibility constraint is now

$$20 - x_1 + 20 - x_2 = 2x_1 + x_2 \Leftrightarrow 3x_1 + 2x_2 = 40, \qquad x_1, x_2 \geq 0$$

Under the egalitarian collective utility, the increase in agent 1's talent benefits both agents equally, who end up with $u_1 = u_2 = 12$ utils. The corresponding allocation is $x_i = 8$, $u_i = z_i = y_i = 12$, $i = 1, 2$: both agents work less hard and 4 units of the corn produced by agent 1 are transferred to agent 2.

Under the classical utilitarian and the Nash solutions, the fate of agent 1 is less enviable. Say that agent i works x_i hours, $i = 1, 2$, so that total output is $z = 2x_1 + x_2$. Efficiency commands to give exactly $z_i = 20 - x_i$ units of corn to agent i. Thus the classical utilitarian solution maximizes $(20 - x_1) + (20 - x_2)$ under the above feasibility constraint.

The optimal solution is $x_1^* = 13.3$, $x_2^* = 0$, resulting in the allocation $u_1 = z_1 = y_1 = 6.7$, $u_2 = z_2 = y_2 = 20$. This is slavery because the talented agent 1 works 33 percent harder and consumes 33 percent less than before acquiring his special talent. He is also frustrated to see agent 2 reap all the benefits and get a "free ride," whereas he (agent 1) experiences a sharp decrease in his utility!

The utilitarian argument is that one more unit of leisure for agent 2 has a lower opportunity cost, in terms of lost production, than one unit for agent 1, and this argument holds until agent 2 is totally exonerated from work.

The Nash collective utility function yields a milder slavery of the talented, but slavery nevertheless. The Nash arbitrator solves the program

$$\max \log(20 - x_1) + \log(20 - x_2) \qquad \text{under } 3x_1 + 2x_2 = 40$$

The solution is $x_1^0 = 10$, $x_2^0 = 5$ hence the allocation $u_1 = z_1 = y_1 = 10$, $u_2 = z_2 = y_2 = 15$. Here agent 1 does not suffer anymore from his productivity boost, but he does not benefit either; all the benefit goes to the untalented agent who ends up working less hard and consuming more corn (and leisure) than the talented one.

In section 6.2 we propose a different solution of the above production problem based on the Lockean theory of entitlements rather than welfarism. This solution rules out any externality in productivity; hence it eliminates slavery entirely.

3.6 Bargaining Compromise

The bargaining compromise places bounds on individual utilities that depend on the physical outcomes of the allocation problem; thus it moves a step away from the strict postulate of welfarism (section 3.1).

Example 3.10 Priorities Ann, Bob, and Charles work in the same company. Each needs a computer repair job, and their respective repair jobs are not equally long: Ann's repair can be done in 1 hour; Bob's takes 4 hours, and Charles's takes 5 hours.

There is a single repairman in the company. Since an agent must stay idle until the completion of his or her repair, the total waiting time until the repair is completed measures his or her disutility.

The classical utilitarian solution minimizes total waiting time, and this is achieved by serving the shortest repair job first. So the respective disutilities are

Ann: 1, *Bob*: 5, *Charles*: 10

Total disutility is 16, and it is a simple matter to check that any other ordering of the jobs yields a larger total waiting time. This is hard on Charles, especially so if we reduce the difference between Bob's and Charles' repair jobs. If Charles's job is one minute longer, he still has to wait 4 hours more.

If the only available choices are the six deterministic orderings of Ann, Bob, and Charles, the leximin ordering selects the same ordering Ann < Bob < Charles, as the reader can easily verify. However, in this example we allow randomization over the six orderings in order to achieve equitable compromises where two agents with nearly identical characteristics (job length) have nearly identical expected waiting times.

Classical utilitarianism refuses to compromise, because the ordering above uniquely minimizes total waiting time. The contrast with the egalitarian solution could not be sharper, as the latter gives to each participant precisely the same expected disutility. Note that the smallest waiting time u that can be guaranteed to all three participants is $u = 7.1$,[21] which obtains for instance by randomizing as follows over three different schedulings of the three jobs:

Scheduling	u_A	u_B	u_C	Probability
B, A, C	5	4	10	0.4
A, C, B	1	10	6	0.1
C, B, A	10	9	5	0.5
	7.1	7.1	7.1	Expected utility

Now we see that Ann, whose job is shortest by far, is served first only 10 percent of the time, whereas Bob is first 40 percent of the time and Charles 50 percent.[22] The solution ignores the differences between the delay externalities caused by jobs of different length.

21. To check this, observe that for *any* scheduling of the three jobs, the utility of the three agents satisfy $u_A + 4u_B + 5u_C = 71$. This observation is generalized in exercise 3.10.

22. Note that the probability of agent i being served first is proportional to the length of her job. This is a general property: see exercise 3.10.

The bargaining compromise here equalizes "priorities" instead of utilities: each agent has an equal right to be served first, or second, or third. If we randomize over all six orderings, with equal probability to each ordering, the expected waiting time of our three agents are

$$u_A = \frac{1}{3}1 + \frac{1}{6}5 + \frac{1}{6}6 + \frac{1}{3}10 = 5.5$$

$$u_B = \frac{1}{3}4 + \frac{1}{6}5 + \frac{1}{6}9 + \frac{1}{3}10 = 7$$

$$u_C = \frac{1}{3}5 + \frac{1}{6}6 + \frac{1}{6}9 + \frac{1}{3}10 = 7.5$$

This outcome is an egalitarian compromise in the following *relative* sense: everyone ends up half-way between his or her worst wait (i.e., 10) and his or her best wait (i.e., 1 for Ann, 4 for Bob, 5 for Charles). These two bounds of the best and worst wait are very natural, but their meaning goes beyond the mere description of welfare: they depend on the set of feasible outcomes in the particular allocation problem.

The choice of the zero and/or the scale of individual utilities is crucial whenever a social welfare ordering picks the solution: with the exception of the classical utilitarian (independent of individual zeros but not scales) and the Nash collective utility function (independent of individual scales but not zero), all other social welfare orderings depend on both the individual zeros and scales.

The bargaining version of welfarism incorporates an objective definition of the zero of individual utilities, which corresponds to the worst outcome deemed acceptable from the point of view of a certain agent. In some cases this outcome is interpreted as the *disagreement* outcome because each agent has the strategic option to "walk away" from the arbitration table, so the arbitrator must take the corresponding utility as a hard lower bound. In other cases, like examples 3.10 and 3.11, zero utility simply comes from the worst feasible outcome in the allocation problem; hence we call it the *minimal utility*.

The bargaining approach then applies a scale invariant solution to the zero normalized problem, which in turns ensures that the solution is independent of both individual zeros and scales of utilities.

The two prominent bargaining methods are the Nash bargaining and Kalai-Smorodinsky solutions introduced in our next example.

Example 3.11 Ann and Bob represent two companies selling related yet different products, and share a retail outlet. They can set up the outlet in three different modes denoted a, b, and c that bring the following volumes of sales (in thousands of dollars):

	a	b	c
Ann	60	50	30
Bob	80	110	150

(6)

Both managers are only interested in maximizing the volume of their own sales (which may not be the same thing as maximizing profit) and accounting rules prohibit any cash transfers. Thus the only tool for compromises is time-sharing among the three modes: over a yearly season, they can mix them in arbitrary proportions such that x, y, z such that $x + y + z = 1$.

Applying any one of our three basic welfarist solutions to the raw utilities given in (6) make little sense. For instance, the egalitarian collective utility picks outcome a where Ann's utility is highest. But the fact that Ann's business always yields a smaller volume of sales should not matter: the issue is to find a compromise between three feasible outcomes over which the agents have opposite preferences; the relative size of Ann's business to Bob's business is irrelevant.

Total utility in classical utilitarian fashion—maximal at c—is similarly irrelevant. We wish to define a fair compromise that depends neither on the scale nor on the zero of both individual utilities.

For minimal utility of either player, we pick the lowest feasible volume of sales: 30K for Ann and 80K for Bob. Indeed, this level is guaranteed even by conceding to the other agent his or her favorite outcome. This yields the new utility table:

	a	b	c
Ann	30	20	0
Bob	0	30	70

(7)

The idea of a random ordering, successful in example 3.10, suggests letting Ann and Bob each have their way 50 percent of the time: this means that $x = z = \frac{1}{2}$, outcomes a and c each with a timeshare $\frac{1}{2}$. But the resulting normalized utilized vector is (15, 35), whereas the outcome $y' = 0.8$, $z' = 0.2$ (b or c with respective timeshares 0.8 and 0.2) yields the utilities (16, 38), and hence is Pareto superior.

It turns out that any combination of a and c is Pareto inferior to some combination of a and b, or of b and c: this is apparent on figure 3.7, where compromises of a and c produce the utility vectors in the segment AC.

There are now two simple ways to select the shares x, y, z without taking into account the scales of individual utilities. The first one is to maximize the Nash collective utility function:

$$\begin{aligned} &\max \log(30x + 20y) + \log(30y + 70z) \\ &\text{under} \quad x + y + z = 1, x, y, z \geq 0 \end{aligned} \tag{8}$$

The second way is the Kalai-Smorodinsky solution, equalizing the *relative* utility gains, namely the ratio of the actual gain to the maximum feasible gain. In this example the maximal feasible gains are 30 and 70 for Ann and Bob respectively. Therefore the KS solution selects the shares x, y, z so as to

$$\begin{aligned} &\text{maximize } \frac{30x + 20y}{30} = \frac{30y + 70z}{70} \\ &\text{under} \quad x + y + z = 1, x, y, z \geq 0 \end{aligned} \tag{9}$$

The resolution of programs (8) and (9) is greatly simplified by taking a look at the feasible utility set of the normalized utilities (7). Figure 3.7 reveals that the efficient compromises among a, b, and c involve either a and b only ($z = 0$: interval AB) or b and c only ($x = 0$: interval BC). On each one of these two intervals, there is only one degree of freedom, so the resolution of programs (8) and (9) becomes easy.

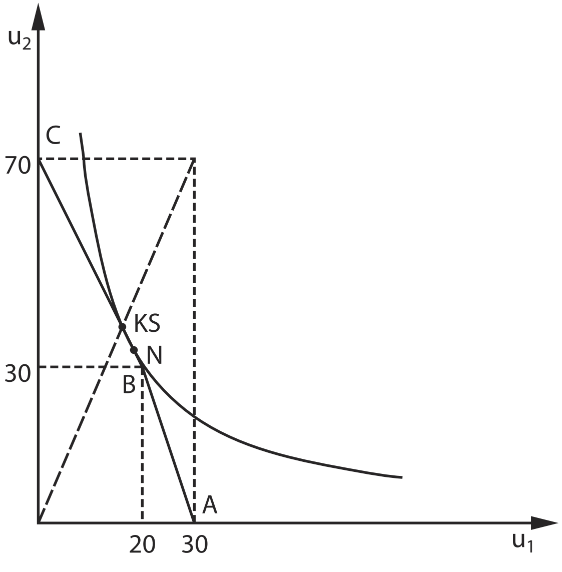

Figure 3.7
Bargaining solutions in example 3.11

Consider first the interval AB, corresponding to $z = 0$ and to the utility vectors $(10x + 20, 30 - 30x)$ for $0 \leq x \leq 1$. We see that equation (9) is impossible, namely

$$\frac{10x + 20}{30} = \frac{30 - 30x}{70} \Rightarrow x = -\frac{5}{16}$$

Therefore the KS solution lies on BC, corresponding to $x = 0, z = 1 - y$:

$$\frac{20y}{30} = \frac{70 - 40y}{70} \Rightarrow y = \frac{21}{26}, z = \frac{5}{26}$$

We turn to the resolution of program (8). We know that its optimal solution lies either on AB or on BC. We check that B is the solution of (8) on AB, but on BC we can do better, namely maximize $\log(20y) + \log(30y + 70(1 - y))$ under $0 \leq y \leq 1 \Rightarrow y = \frac{7}{8}$, hence

Nash solution: $\quad y = \frac{7}{8}, z = \frac{1}{8} \Rightarrow u_1 = 17.5, u_2 = 35$

KS solution: $\quad y = \frac{21}{26}, z = \frac{5}{26} \Rightarrow u_1 = 16.1, u_2 = 37.7$

where the utilities are normalized as in (7).

Note that both solutions are Pareto superior to the random dictator outcome $a/2 + c/2$, with associated utilities (15, 35). This is a general property, discussed below, of our two bargaining solutions.

We give now the general definition of the Nash and Kalai-Smorodinsky bargaining solutions. The data are a set U of feasible utility profiles and a distinguished minimal utility profile u^0. See figure 3.8 where an important feature is the fact that the set U is convex.[23] We set the zero of agent i's utility at u_i^0: figure 3.9.

The Nash bargaining solution maximizes the Nash utility under this normalization of individual zeros, namely $\Pi_i(u_i - u_i^0)$. Of course the maximization bears exclusively on those utility profiles in U such that $u_i \geq u_i^0$ for all i.

Next we compute the maximal utility level $u_i^{\max}$ that agent i can achieve whenever other agents receive at least their minimal disagreement utility: that is to say, $u_i^{\max}$ solves the program max u_i over all $u \in U$ such that $u \geq u^0$. The quantity $\delta_i = u_i^{\max} - u_i^0$ is the maximal feasible gain of agent i above and beyond his minimal utility. The KS solution equalizes the relative gains (fraction of maximal feasible gains) of all agents. It is the unique

23. It contains the segment joining any two of its points.

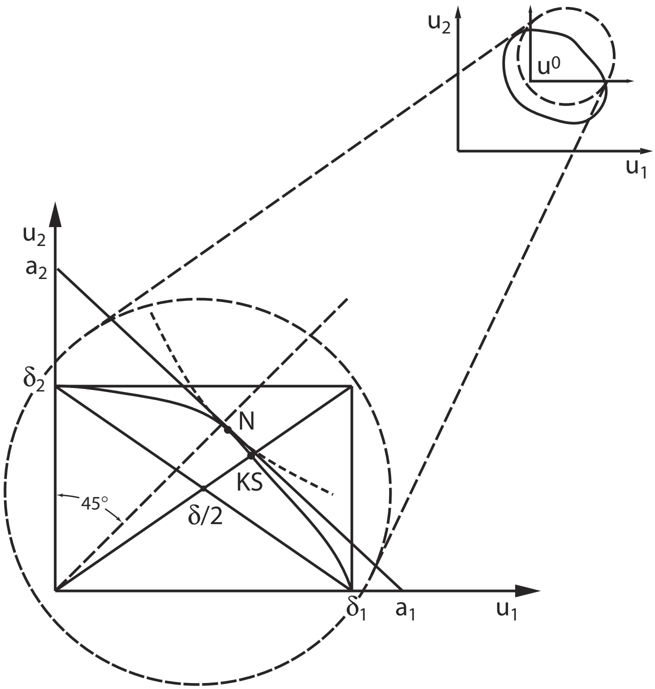

Figures 3.8 and 3.9
Nash and KS solutions

utility profile u such that

$$u \text{ is efficient and } \quad \frac{u_i - u_i^0}{\delta_i} = \frac{u_j - u_j^0}{\delta_j} \quad \text{for all } i, j.$$

Figure 3.9 shows the geometry of this construction in a two-agent example: draw first the utopian utility profile δ where each agent gets δ_i; the KS solution is at the intersection of the efficiency frontier of U with the line from the zero profile to the utopian profile.

The geometric characterization of the Nash solution is an interesting first-order condition, namely a property of the line tangent to the efficiency frontier of U at the Nash point N. Writing a_i for the intersection of this line with the u_i axis, N is simply the midpoint of a_1a_2, as shown on figure 3.9.[24]

24. Note that this property implies that the tangent line is orthogonal to the vector (u_2, u_1): hence if (du_1, du_2) is a small variation of u at N along the efficiency frontier, we have $u_2du_1 + u_1du_2 = 0$, which is the first-order condition for the maximization of u_1u_2.

Both the Nash and the KS solutions are independent of the individual scales of utilities. We know this to be true for the Nash utility function, hence for the outcome maximizing this function over U. As for the KS solution, we note that the ratios $(u_i - u_i^0)/\delta_i$ are invariant under a rescaling of utility u_i (because both numerator and denominator are multiplied by the same rescaling factor), which proves the point.

Another appealing feature shared by the Nash and KS solutions is to guarantee for agent i his minimal utility *plus* $1/n$th of his maximal feasible gain δ_i:

$$u_i \geq u_i^0 + \frac{1}{n}\left(u_i^{\max} - u_i^0\right) \tag{10}$$

In other words, we draw an agent at random (with uniform probability) and give him the entire feasible surplus while other agents merely get their minimal utility. The resulting expected utility is a lower bound of what every agent receives under the Nash or the KS solution. The proof of this claim in the two-agent case is clear on figure 3.9.[25]

Our last example emphasizes a property of the KS solution that sets it apart from the Nash bargaining solution, and from any solution maximizing a collective utility function after normalizing individual zeros in some objective fashion. The KS solution depends on the entire shape of the feasible set U: the solution is not independent of irrelevant utility profiles, meaning utility profiles that are "far" from the equitable compromise.

Example 3.12 Vitamins A bottle containing 10 grams of vitamin X and 10 grams of vitamin Y must be shared by Ann and Bob, who both need to increase their level of zygum, a certain compound that can only be metabolized from vitamin X or vitamin Y. The zero utility outcome is that no one gets any vitamin: the agents hold no claim on any of the resources, which are entirely under the control of the benevolent dictator.

We learn first that both Ann and Bob metabolize 1 unit of zygum from 1 gram of vitamin X or Y. Thus the utility (quantity of zygum) they derive from the allocation (x_i, y_i) is $u_i = x_i + y_i$. By the symmetry of the problem, the only fair utility profile is $(u_1, u_2) = (10, 10)$ (10 grams of vitamins per agent).

Now further testing reveals that Bob's metabolism is only half as efficient at producing zygum from vitamin Y than originally thought. From the allocation (x_2, y_2), Bob derives $u_2 = x_2 + (y_2/2)$. Ann's metabolism, on the other hand, still gives her $u_1 = x_1 + y_1$. The efficient allocation of vitamins now precludes giving positive amounts of vitamin Y to Bob *and* of vitamin X to Ann (for they would be able to find a mutually advantageous swap).

25. Note that U contains both points $(0, \delta_2)$ and $(\delta_1, 0)$. Therefore, because U is convex, it contains their midpoint $\delta/2$. The KS solution lies on the segment $[0, \delta]$ beyond $\delta/2$, which proves (10). Next we check $a_i \geq \delta_i$, again by the convexity of U: therefore the Nash solution $(a_1/2, a_2/2)$ in Pareto superior to $(\delta_1/2, \delta_2/2)$, establishing (10).

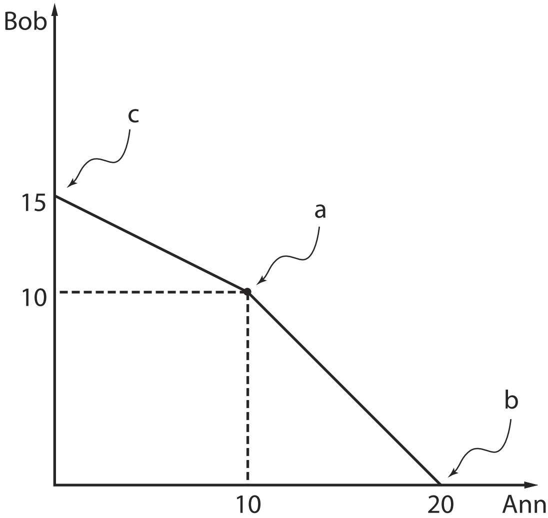

Figure 3.10
Feasible utility set in example 3.12

The efficiency frontier depicted on figure 3.10 comes from *either* giving all 10 grams of vitamin Y to Ann (segment ab) or giving all 10 grams of vitamin X to Bob (segment ac).

The point is that the utility profile (10, 10) is still feasible (all the vitamin X to Bob, all the vitamin Y to Ann), and that Bob's decrease in productivity only eliminates certain utility profiles such as (0, 20) that were unfair in the first place. Any social welfare ordering called (10, 10) optimal in the former problem still calls it optimal in the latter problem. This includes the Nash collective utility, the leximin social ordering, and any social welfare ordering that strictly improves under a Pigou-Dalton transfer.

The KS solution takes a different viewpoint. In the first problem it picks $u = (10, 10)$ but in the second it recommends $u' = (11.4, 8.6)$. To see this, check that the maximal feasible utilities are $\delta'_1 = 20$, $\delta'_2 = 15$; therefore equality of relative benefits means $u'_1/20 = u'_2/15$. An efficient allocation with associated utility vector on ab (figure 3.10) takes the form

Ann: $x_1 = z, \quad y_1 = 10 \Rightarrow u'_1 = 10 + z$

Bob: $x_2 = 10 - z, \quad y_2 = 0 \Rightarrow u'_2 = 10 - z$

thus equality of relative benefits yields $z = 1.43$, and the announced utility vector.

The decrease of Bob's maximal utility weakens his position, even though it relates to the "irrelevant" allocations where he would get a higher utility than Ann does. The KS solution is even less a welfarist solution than the Nash solution because it takes into account not only an exogenous notion of minimal utility, but also the corresponding maximal feasible surplus of each participant.

In the "vitamin" interpretation of our example, the KS solution is not very appealing. It is, however, plausible if utility represents subjective tastes instead of needs. Think of the division of ten free meals in restaurant X and ten free meals in restaurant Y. Say that restaurant Y is strictly nonsmoking, whereas restaurant X has a smoking section. Bob is a smoker who enjoys one meal where he can smoke as much as two where he can't. Is it unfair to give Ann ten meals at Y and one meal at X, whereas Bob gets nine meals at X? After all, meals at Y are to Bob a low-quality commodity, so he cannot object to Ann getting a larger number of such meals than he gets of "good" meals. This kind of argument takes us directly into the discussion of fair division with respect to heterogeneous preferences and to the no-envy property, the subject of chapter 7.

In chapter 7 we stress the systematic connection between the two bargaining solutions, Nash and KS, and the two central fair division methods, known respectively as the competitive allocation with equal incomes and the egalitarian-equivalent solution. See, in particular, examples 7.12 and 7.10, which are related to examples 3.7 and 3.12.

3.7 Introduction to the Literature

Rawls (1971, 1988) introduced the notion of primary goods; its critique as briefly discussed in section 3.1 is well articulated by Roemer (1996); see also Sen (1985).

The central result for section 3.2 is the representation of a separable social welfare orderings by additive collective utilities, also known as the Debreu-Gorman theorem; see Debreu (1960) and Gorman (1968). A systematic treatment is in Blackorby et al. (1978). The further role of invariance axioms and of the Pigou-Dalton transfer principle was developed through the seventies and is nicely summarized by Sen (1977) and Roberts (1980a, b).

The egalitarian collective utility appeared first in Kolm (1972), and the leximin preordering was axiomatized by d'Aspremont and Gevers (1977); see section 3.3. The problem of example 3.2 is the subject of Brams and Fishburn (2000).

The informal comparison of the leximin, Nash, and classical utilitarian solutions in section 3.4 is inspired by chapters 1, 2, and 3 in Moulin (1988), which provides a systematic formal presentation of the material in sections 3.2 to 3.6.

Example 3.4 on the location of a facility is related to the central model of party competition—Black (1958), whereby the "facility" is the political platform submitted to the voters whose ideal platforms are spread over the left-right spectrum represented by an interval. See the discussion of voting over single-peaked preferences in section 4.3.

Resource monotonicity plays an important role in this book; in addition to section 3.5, the concept is discussed in sections 2.5, 6.6, and 7.6. The idea appeared first in axiomatic bargaining, where it yields a simple characterization of the egalitarian solution; see Kalai (1977). A more systematic discussion of this idea in axiomatic bargaining is in Thomson (1999). Its application to resource allocation problems are reviewed in Roemer (1996) and Moulin (1995).

The two classic bargaining solutions of section 3.6 were first axiomatized by Nash (1953) and Kalai and Smorodinsky (1975). Several surveys on axiomatic bargaining are now available: Roth (1979), Peters (1992), and Thomson (1999).

Example 3.12 is inspired by Yaari and Bar-Hillel (1984), who conducted stimulating experiments on fairness in resource allocation. The slavery of the talented—example 3.8—is due to Mirrlees (1974).

Exercises to Chapter 3

Exercise 3.1 Variant of Example 3.1

We have five agents and the six feasible subsets are the same as in example 3.1.

a. Assume that the utility of agent i being saved is u_i, and zero otherwise, with $u_i > 0$. Show that the leximin ordering always picks one of the three subsets with three agents.

b. Assume from now on that the utility of being saved is u_i, and v_i otherwise, with $u_i > v_i > 0$. Show that $u_i, v_i, i = 1, \ldots, 5$ can be chosen so that $\{1, 2\}$ is the unique optimal choice of the leximin ordering.

c. Find some values of u_i, v_i such that the arbitrator ranks all three subsets of size two above those of size three, whereas the classical utilitarian arbitrator does just the opposite.

Exercise 3.2 Fair Division with Identical Preferences

We have 3 gold coins, 5 silver coins, and 8 bronze ones. As in example 3.1, all agents have identical preferences over lots. We assume that a gold coin is worth two silver ones, or three bronze ones. Thus we measure the common utility for a lot by adding 6 utils for a gold coin, 3 utils for a silver one, and 2 utils for a bronze one. In particular, total utility is 49, irrespective of the number n of agents among whom the sixteen coins must be divided.

Given n, call "n-equal division" a division of p into n integers a_i such that any two integers a_i, a_j differ by at most 1. For instance, with $n = 5$, an equal division of 49 is (10, 9, 10, 10, 10) but (10, 9, 9, 11, 10) is not.

Clearly, an n-equal division of p exists for all n and all p, and is unique up to permuting the a_i.

a. Suppose that the 16 coins can be divided in n lots such that the corresponding profile of utilities is a n-equal division of 49. Show that these allocations, and only these, maximize the leximin social welfare ordering.

b. Show that for $n = 2, 3, 4, 5, 6, 8, 9$, the 16 coins can be divided in lots in such a way that the corresponding utility profile is a n-equal division of 49. Show that this is not possible for any other choice of n.

c. Find the division of the 16 coins selected by the egalitarian arbitrator for $n = 7$ and for $n = 10$.

Exercise 3.3 Cake Division with Altruism

One unit of cake is to be distributed between Ann and Bob. The utility of each agent has two components:

The "selfish" utility increase derived from one's own consumption and measured by a function $u(x_i)$ where x_i is one's own share.

The "selfish" utility of the *other agent.*

These two components are combined in some proportion, and the proportion measures the degree of altruism of each agent. Specifically, we denote by a and b the shares of Ann and Bob, and we write their utility for a division (a, b) of the cake as follows:

Ann: $u(a) + \lambda_A u(b)$

Bob: $u(b) + \lambda_B u(a)$

Here λ_A is Ann's "degree of altruism," $0 \leq \lambda_A \leq 1$, and the interpretation of λ_B is similar. We assume that the common function u (measuring utility increase from own consumption) is increasing and concave.

The goal of the exercise is to assume that Ann is more altruistic than Bob, namely $\lambda_A > \lambda_B$, and to find out if she receives a bigger share, smaller share, or equal share of the cake:

- If the utilitarian collective utility is maximized
- If the egalitarian collective utility is maximized
- If the Nash collective utility is maximized

Answer first in the two following examples, then with maximal generality, namely without specifying u, λ_A, or λ_B:

Example 1: $\lambda_A = \frac{1}{3}$, $\lambda_B = \frac{1}{4}$, $u(x) = x$
Example 2: $\lambda_A = \frac{1}{3}$, $\lambda_B = \frac{1}{4}$, $u(x) = \sqrt{x}$

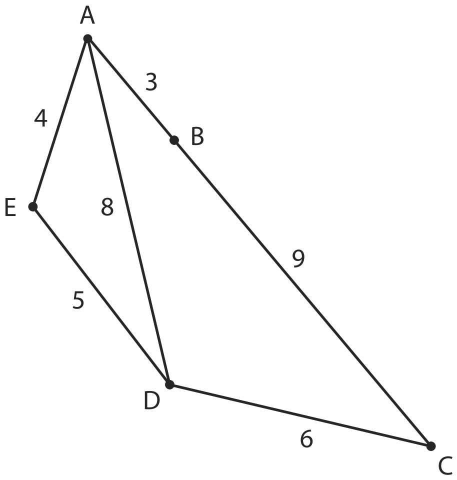

Figure 3.11
Road network for exercise 3.4a

Exercise 3.4 Location of a Facility on a Network with Loops

a. Consider the road network of figure 3.11, where each dot represents an agent (five agents in total), and numbers represent distances in miles. As in example 3.8, the distance between two points is that of the shortest path on the network connecting them. Disutility is the distance between one's location and that of the facility. The facility can be located anywhere on the road network.

Where will the classical utilitarian arbitrator locate a desirable facility? *Hint:* Check first that the utilitarian optimum must be at one of the five points where an agent resides.

What about the egalitarian arbitrator?

b. Consider the road network of figure 3.12. Two agents live at B, three live at C, and four at D (9 agents in total). There is no direct road between B and C. Find the optimum locations for the egalitarian and the classical utilitarian arbitrators.

c. Suppose now that the dotted line between B and C in figure 3.12 is a new road of length 8. Answer the same questions as in b. Which agents benefit and which are hurt by the increase in the resources?

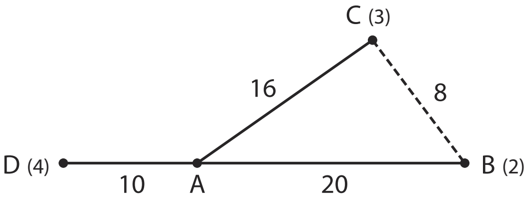

Figure 3.12
Road network for exercise 3.4b

*Exercise 3.5 Location of a Noxious Facility

a. Consider a variant of example 3.5 where the densities of the agents over the interval [0, 1] are as follows:

Density 2 on $\left[0, \frac{1}{3}\right]$
Density 1 on $\left[\frac{1}{2}, \frac{2}{3}\right]$ and on $\left[\frac{3}{4}, 1\right]$
Density 0 on $\left]\frac{1}{3}, \frac{1}{2}\right[$ and on $\left]\frac{2}{3}, \frac{3}{4}\right[$

Show that the location selected by the classical utilitarian is $y_u = 1$. Show that the egalitarian arbitrator selects $y_e = 5/12$.

b. Next consider the following densities:

Density 2 on $\left[0, \frac{1}{3}\right]$
Density 1 on $\left[\frac{1}{2}, \frac{2}{3}\right]$
Density 3 on $\left[\frac{5}{6}, 1\right]$
Density 0 on $\left]\frac{1}{3}, \frac{1}{2}\right[$ and on $\left]\frac{2}{3}, \frac{5}{6}\right[$

Find the locations selected by the classical utilitarian, egalitarian (leximin), and Nash arbitrators.

c. Consider the road network of question a in exercise 3.4 depicted in figure 3.11. Show that the classical utilitarian locates the noxious facility on the road *CD*, one mile away from C. Show that the egalitarian selects the midpoint between B and C.

d. Consider the road network of questions b and c in exercise 3.4 depicted on figure 3.12. Show that the egalitarian arbitrator picks the same location for both networks (with or without the direct road BC). Show that the utilitarian arbitrator selects B in the network of questions b, and the location 2 miles away from B on BC in the network of question c.

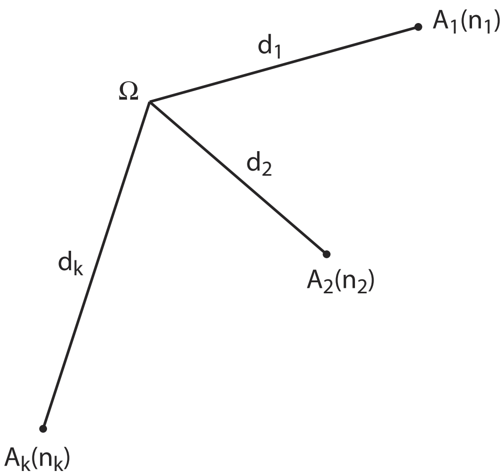

Figure 3.13
Road network for exercise 3.6

*Exercise 3.6 Location on a Star-Tree

The road network depicted on figure 3.13 is a "star-tree." The outer node A_k is connected to the center O by a direct route of length d_k, and there is no other road going through A_k. For concreteness, assume ten nodes $A_1, \ldots, A_{10}$. There are n_k agents living at A_k, and no one lives anywhere else. So $n = n_1 + \cdots + n_{10}$ is the total number of agents.

a. In this question the facility is a "good" one, as in example 3.4 and exercise 3.4. Show that the optimal egalitarian location is the midpoint between the two outer nodes A_i, A_j farthest away from the center (i.e., d_i and d_j are the two largest distances). Check that this location is unambiguous, even if there are several possible choices for A_i, A_j.

Show that the unique classical utilitarian optimum is at the center if no outer location contains more than one-half of the agents: $n_k < n/2$ for $k = 1, \ldots, 10$. What happens in the remaining case?

b. Now the facility is noxious, as in example 3.5 and exercise 3.5. Show that the egalitarian optimum is the midpoint between the outernode farthest away from the center and the one closest to the center. That is, if $d_{i^*} = \max_i d_i$ and $d_{j^*} = \min_i d_i$, the midpoint of $A_{i^*}A_{j^*}$ maximizes the egalitarian collective utility. Can we use the leximin ordering to break ties?

Show that the optimal locations for the classical utilitarian are all the nodes A_{k^*} maximizing the product $(n - n_k)d_k$. Comment on the trade-off leading the choice of the utilitarian arbitrator.

Exercise 3.7 Time-Sharing The exercise proposes variants of examples 3.6a and 3.6b.

a. Consider example 3.6a. Compute the optimal time-sharing for the collective utility functions $W_p, 0 < p \leq 1$ and $W^q, 0 < q$, as is defined by (5) in section 3.2. Check that when p or q go to zero (resp. q goes to infinity), the solution converges toward the Nash (resp. the egalitarian) optimum.

b. In the following four examples we have one agent per row and three decisions. The arbitrator can mix between the three decisions. Find the classical utilitarian, egalitarian, and Nash solutions:

	a	b	c
Ann	1	1	0
Bob	0	0	1
Chris	0	1	0
Dave	1	0	0

	a	b	c
Ann	1	0	0
Bob	0	1	0
Chris	0	1	1
Dave	1	1	0

	a	b	c
Ann	1	0	1
Bob	0	1	0
Chris	1	0	0
Dave	0	1	1

	a	b	c
Ann	1	0	0
Bob	1	1	0
Chris	1	0	1
Dave	0	1	1

Hint: For the first and fourth example, use the symmetry between two of the three decisions; for the fourth example, the egalitarian collective utility is not enough, and you must invoke the leximin ordering.

c. Now we have four types of agents and three decisions a, b, c. The total number of agents is $n = 2m + 2p$.

	a	b	c
m	1	0	0
m	0	1	0
p	1	0	1
p	0	1	1

Find the three usual solutions, distinguishing the cases $p \geq m$ and $p < m$.

d. We have here four decisions and four agents:

	a	b	c	d
Ann	1	0	1	1
Bob	1	1	0	0
Chris	0	0	1	0
Dave	0	1	0	1

Compute the egalitarian and Nash solutions.

e. Consider example 3.1 where we have six deterministic choices and five agents. Each choice corresponds to a subset of agents who receive a utility of one, so that the six choices are

$a = \{1, 2\}, b = \{1, 3\}, c = \{1, 4\}, d = \{2, 4, 5\}, e = \{2, 3, 5\}, f = \{3, 4, 5\}$

Find the time-sharing recommended by the three usual solutions.

Exercise 3.8 Slavery of the Talented

This is a variant of example 3.9 where the only difference is the common utility function of the two agents: $u_i(z_i, y_i) = \sqrt{z_i y_i}$.

a. Assume that $s_1 = 2, s_2 = 1$. Show that if an allocation is efficient and z_i, y_i, x_i are all positive, we must have

$$z_1 = 2y_1, \quad z_2 = y_2, \quad 2y_1 + y_2 = 30, \qquad 5 \leq y_1 \leq 15,\ 0 \leq y_2 \leq 20$$

b. Show that the utilitarian solution is full slavery of the talented: when productivities are $s_1 = 2, s_2 = 1$, agent 1 works full time and consumes no corn.

c. Show that there is no failure of monotonicity under the Nash solution: both agents benefit when agent 1's productivity increases from 1 to 2; yet agent 1's gain is smaller than agent 2's.

d. Compare the two solutions above with the egalitarian solution.

Exercise 3.9 Bargaining Compromises

We consider three variants of example 3.11. Compute in each case the Nash and KS bargaining solutions

a.

	a	b	c
Ann	70	50	20
Bob	80	90	110

As in example 3.11 we normalize utilities at the worst utility between a, b, and c:

	a	b	c
Ann	40	20	10
Bob	20	30	70

b. Two agents and four outcomes. Minimal utility is at outcome a (thus the data are "pre-normalized"):

	a	b	c	d
Ann	0	1	5	6
Bob	0	11	6	3

c. Variant of question b: with have two agents with utility identical to Ann's, and three with utility identical to Bob's.

Exercise 3.10 Generalization of Example 3.10

The job of agent i requires a_i units of time, and we assume that $a_1 < a_2 < \cdots < a_n$. Disutility is total waiting time until completion of one's job.

a. Assume that the server must choose a deterministic priority ordering. Show that it chooses the ordering $\{1, 2, \ldots, n\}$ under any social welfare ordering (monotonic and symmetric).

b. Now the server can mix over all $n!$ orderings σ of $\{1, 2, \ldots, n\}$, with arbitrary probability $\pi_\sigma \geq 0, \sum_\sigma \pi_\sigma = 1$. Show that the utilitarian server chooses the ordering $\{1, 2, \ldots, n\}$ with probability 1, as in question a.

c. Show that for the profile of utilities u^σ resulting from an ordering σ of $\{1, 2, \ldots, n\}$, the sum $\sum a_i u_i^\sigma$ is independent of σ. Deduce that any convex combination of the profiles π^σ is a Pareto optimal and feasible utility profile.

d. Compute the expected utility profile when the priority ordering σ is selected at random with uniform probability on all orderings σ. Check that it is the Kalai-Smorodinsky solution if the maximal disutility (minimal utility) is $a_N = \sum_j a_j$ for every agent, and agent $i's$ minimal disutility is a_i.

e. Consider the n orderings $\sigma^k, k = 1, \ldots, n$, obtained by successive applications of the circular permutation $i \to i + 1$.

$$\sigma^1 = \{1, 2, \ldots, n\};\ \sigma^2 = \{2, 3, \ldots, n-1, 1\}; \ldots; \sigma^k = \{k, k+1, \ldots, k-1\}; \ldots$$

Check that if we choose the priority ordering σ^k with probability a_k/a_N, for $k = 1, \ldots, n$, the resulting utility profile is egalitarian.

Exercise 3.11 Sharing One Commodity

a. We must divide \$100 between two agents with the following utilities for money:

$$u_1(x) = \sqrt{x}, \ u_2(x) = 2\sqrt{x}$$

Compute the classical utilitarian, egalitarian, Nash, and KS solutions (for the latter two, take the minimal utility to be zero).

b. Answer the same questions with the utility functions:

$$u_1(x) = x^{2/3}; \quad u_2(x) = x^{1/3}$$

c. Answer the same questions with the utility functions:

$$u_1(x) = x; \quad u_2(x) = \frac{100x}{100 + x}$$

d. What happens in the problems of questions a, b, and c when the cash prize increases? Which agent gets a bigger share of the increment according to what solution?

*Exercise 3.12 Leximin and Leximax

Given a utility profile $u = (u_i)$ in $\mathbf{R}^n$, we denote by u^* (resp. $^*u$) the vector obtained by rearranging the coordinates of u increasingly (resp. decreasingly). The leximin ordering compares two profiles u, v by comparing u^* and v^* for the lexicographic ordering:

$$u \succeq v \overset{\text{def}}{\iff} \{u_1^* > v_1^*\} \quad \text{or} \quad \{u_1^* = v_1^* \text{ and } u_2^* > v_2^*\}$$

$$\text{or} \quad \{u_1^* = v_1^*, u_2^* = v_2^* \text{ and } u_3^* > v_3^*\} \ldots$$

$$\text{or} \quad \{u^* = v^*\}.$$

The leximax ordering compares u and v as the lexicographic ordering compares $^*u$ and $^*v$.

a. Show that both orderings, leximin and leximax, are independent of unconcerned agents (property 1).

b. Show that they are both independent of the common utility pace (discussed at the end of section 3.3).

c. Show that leximin meets, but leximax fails—the Pigou-Dalton transfer principle.

d. Show that leximin is the limit as q goes to infinity of the social welfare ordering W^q defined in (4). Show that leximax is the limit of W_p (also defined in 4) as p goes to infinity. The convergence statement is defined as follows. Suppose that the two profiles u, v in $\mathbf{R}^n$ are such that $W^q(u) \geq W^q(v)$ for all q large enough. Then $u \succsim v$ for the leximin ordering.

The other convergence is defined similarly.

***Exercise 3.13 Independence of Common Zero of Utilities**

Consider the family of collective utility functions:

$$V_p(u) = \sum_i e^{pu_i} \quad \text{for some fixed} \quad p, p > 0$$

$$V^q(u) = -\sum_i e^{-qu_i} \quad \text{for some fixed} \quad q, q > 0$$

a. Show that each collective utility function V_p and V^q is independent of unconcerned agents. Show that it is independent of the common zero of utilities, a property similar to (3) bearing on a simultaneous shift of all individual zeros of utilities.

b. What is the limit of the social welfare ordering V_p (resp. V^q) as p (resp. q) goes to infinity? As p (resp. q) goes to zero?

***Exercise 3.14 Distortion of Individual Zeros and Scales**

a. Consider the binary choice between the two four-person utility profiles

$$u(a) = (1, 4, 4, 0), \quad u(b) = (0, 4, 4, 2)$$

The arbitrator uses the leximin ordering to select a or b with no possibility of mixing. Now agent 1 inflates the scale of his utility by a factor of 3, so his new utility is $u'_1(a) = 3$; $u'_1(b) = 0$. Show that this distortion is profitable even if it is a lie (agent 1's true utility remains u_1).

Next suppose that the arbitrator can mix a and b, and still is an ardent egalitarian. Show that the (untrue) distortion by agent 1 ends up hurting him.

b. From now on we restrict attention to strictly positive individual utilities. We fix a social welfare ordering $\succsim$ and define the property: "increasing strategically the scale of one's utility can't hurt." For all profile u, all agent i, and all λ, $\lambda > 1$, we write $u' = (u \mid^i \lambda u_i)$ for the profile $u'_j = u_j$ if $j \neq i$ and $u'_i = \lambda u_i$.

The property above is now defined:

$$\text{for all } i, \lambda > 1, u \text{ and } v: \quad \{v \succ u \text{ and } u' \succ v'\} \Rightarrow \{u_i > v_i\}$$

Interpret this definition and explain its name. We say "increasing one's scale is profitable" if, in addition to the property above, we have

$$\text{for all } i, \text{ and all } u, v: \quad \{v \succ u \text{ and } u_i > v_i\} \Rightarrow \exists \lambda > 1 : (u \mid^i \lambda u_i) \succ (v \mid^i \lambda v_i)$$

Consider the collective utility W_p, $p > 0$ defined in (4). Show that increasing one's scale is profitable for W_p, hence for its limit as well, the leximax social ordering (defined in exercise 3.12). What about W_0, namely the Nash collective utility?

Define similarly the properties "decreasing one's scale cannot hurt" and "decreasing one's scale is profitable." Show that for the collective utility W^q, $q > 0$, in (4), decreasing one's scale is profitable. Therefore the same holds true for its limit, the leximin social welfare ordering.

c. We still assume strictly positive utilities only. We define "increasing the zero of one's utility can't hurt":

for all $i, \lambda > 0, u$ and v: $\{v \succ u\}$ and $(u \mid^i u_i - \lambda) \succ (v \mid^i v_i - \lambda) \Rightarrow \{u_i > v_i\}$

Consider an additive collective utility W as in (2). Show that its ordering meets the above property if and only if g is concave. Show that decreasing one's zero cannot hurt if and only if g is convex.

Deduce that increasing one's zero can't hurt if we are using the collective utility W^q in (4) or W_p for $0 \leq p \leq 1$. In particular, this holds true for the leximin social welfare ordering.