13 Optimality and Comparability

13.1 Introduction

On April 17, 1975, the Khmer Rouge seized power in Cambodia. Pol Pot began to implement his vision of Year Zero in which all inequalities—of class, money, education, and religion—would be eliminated. Driven by their desire to achieve what they perceived as the social optimum, the Khmer Rouge attempted to engineer a return to a peasant economy. In the process they slaughtered an estimated two million people, approximately one-quarter of Cambodia's population. The actions of the Khmer Rouge are an extreme example of the pursuit of equality and the willingness to accept an immense loss in order to achieve it. In normal circumstances governments impose a limit on the cost they are willing to pay for an improvement in equality.

When it comes to the efficiency/equity trade-off the Second Theorem of Welfare Economics has very strong policy implications. These were touched on in chapter 2 but were not developed in detail at that point. This was because the primary value of the theorem is what it says about issues of distribution. To fully appreciate the Second Theorem, it is necessary to view it from an equity perspective and to assess it in the light of its distributional implications.

This chapter will begin by investigating the implications of the Second Theorem for economic policy. This is undertaken on the premise that a social planner is able to make judgments between different allocations of utility. The concept of an optimal allocation is developed and the Second Theorem is employed to show how this can be achieved. Once this analysis has been accomplished, questions are raised about the applicability of lump-sum taxes and the value of Pareto-efficiency as a criterion for social decision-making. This provides a basis for re-assessing the interpretation of the First Theorem of Welfare Economics.

The major deficiency of Pareto-efficiency is identified as its inability to trade utility gains for one consumer against losses for another. This is important since most policy changes will involve some people gaining while other people lose. To proceed further, the informational basis for making welfare comparisons has to be addressed. We describe different forms of utility and different degrees of comparability of utility among consumers. These concepts are then related to Arrow's Impossibility Theorem and the potential for constructing a social welfare function.

13.2 Social Optimality

The importance of the Second Theorem of Welfare Economics for policy analysis is very easily explained. In designing economic policy, a policy maker will always aim to achieve a Pareto-efficient allocation. If an allocation that was not Pareto-efficient was selected, then it would be possible to raise the welfare of at least one consumer without harming any other. It is hard to imagine why any policy maker would want to leave such gains unexploited. If it is presumed that this argument is correct, the set of allocations from which a policy maker will choose reduces to the Pareto-efficient allocations.

Suppose that a particular Pareto-efficient allocation has been selected as the policy maker's preferred outcome. The Second Theorem shows that this allocation can be achieved by making the economy competitive and providing each consumer with the level of income needed to purchase the consumption bundle assigned to him in the chosen allocation. The consumers will then trade, and the chosen equilibrium will emerge as the competitive equilibrium. This is the process of *decentralization*. In achieving the decentralization of the allocation, only two policy tools are employed: the encouragement of competition and a set of lump-sum taxes to ensure that each consumer has the required income. If this approach could be applied in practice, then economic policy analysis would reduce to the formulation of a set of rules that guarantee competition and the calculation and redistribution of the lump-sum taxes. The subject matter of public economics, and economic policy, in general, would then be closed.

Looking at this process in detail, the first point that arises is the question of selecting the most preferred allocation. There are a number of ways to imagine this being done. An obvious one would be to consider voting, either over the alternative allocations directly or else for the election of a body (a "government"), to make the choice. Alternatively, the consumers could agree for it to be chosen at random or else they might hold unanimous views, perhaps via conceptions of fairness, about what the outcome should be. The method that is considered here is to assume that there is a social planner (which could be the elected government). This planner forms social preferences over the alternative allocations by taking into account the utility levels of the consumers. The most preferred allocation according to the social preferences is the one that is chosen.

To see how this method functions, consider the set of Pareto-efficient allocations described by the contract curve in the left-hand part of figure 13.1. Each point on the contract curve is associated with an indifference curve for consumer 1 and an indifference curve for consumer 2. These indifference curves correspond to a pair of

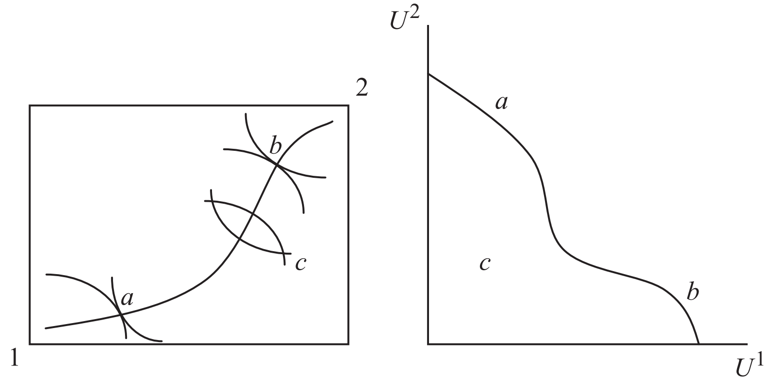

Figure 13.1
Utility possibility frontier

utility levels $\{U^1, U^2\}$ for the two consumers. As the move is made from the southwest corner of the Edgeworth box to the northeast corner, the utility of consumer 1 rises and that of 2 falls. These utility levels can be plotted by observing that each pair of utility levels on the contract curve can be represented as a point in utility space. The loci formed by these points is usually called the *utility possibility frontier*. This is shown in the right-hand panel of figure 13.1 where the utility values corresponding to the points a, b, and c are plotted. Points such as a and b lie on the frontier: they are Pareto-efficient, so it is not possible to raise both consumers' utilities simultaneously. Point c is off the contract curve and is inefficient according to the Pareto criterion. It therefore lies inside the utility possibility frontier.

The utility possibility frontier describes the Pareto-efficient options from which the social planner will choose. It is now necessary to describe how the choice is made. To do this, it is assumed that the social planner measures the welfare of society by aggregating the individual consumers' welfare levels. Given the pair of welfare levels $\{U^1, U^2\}$, the function determining the aggregate level of welfare is denoted by $W\left(U^1, U^2\right)$. This is termed a *Bergson–Samuelson social welfare function*. Basically, given individual levels of happiness, it imputes a social level of happiness. Embodied within it are the equity considerations of the planner. Two examples of social welfare functions are the *utilitarian* $W = U^1 + U^2$ and the *Rawlsian* (or maxi-min) $W = \min\{U^1, U^2\}$. The social indifference curves for these welfare functions are illustrated in figure 13.2, alongside those for an "intermediate" social welfare function. These curves show combinations of the two consumers' utilities that give a constant level of social welfare. The view on equity taken by the social planner translates into their willingness to trade off the

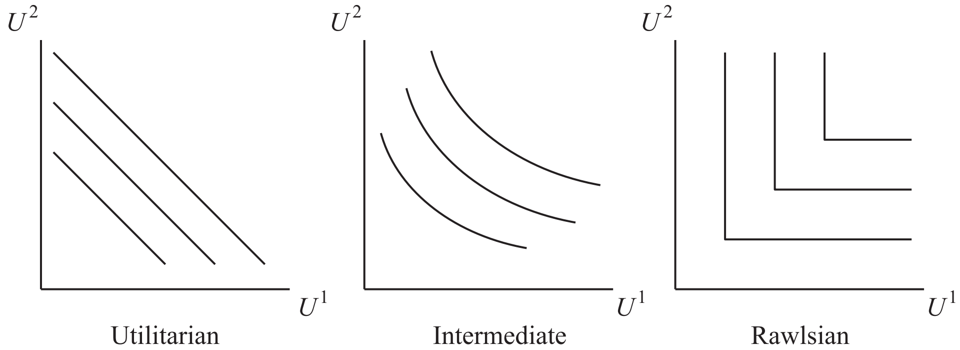

Figure 13.2
Social indifference curves

utility of one consumer against the utility of the other. This determines the shape of the indifference curves. From the shape of the indifference curves it can be seen that the utilitarian and Rawlsian social welfare functions represent two extremes. The utility of one consumer can be substituted perfectly for that of another with the utilitarian social welfare function, but no substitution is possible for the Rawlsian. The intermediate case allows imperfect substitution.

Given the welfare function, the social planner considers the attainable allocations of utility described by the contract curve and chooses the one that provides the highest level of social welfare. Indifference curves of the welfare function can be drawn as in figure 13.3. The social planner then selects the outcome that achieves the highest indifference curve. This optimal point on the utility possibility locus, denoted by point o, can then be traced back to an allocation in the Edgeworth box. This allocation represents the socially optimal division of resources for the economy given the preferences captured by the social welfare function. If these preferences were to change, so would the optimal allocation.

Having chosen the socially optimal allocation, the reasoning of the Second Theorem is applied. Lump-sum taxes are imposed to ensure that the incomes of the consumers are sufficient to allow them to purchase their allocation conforming to point o. Competitive economic trading then takes place. The chosen socially optimal allocation is achieved through trade as the equilibrium of the competitive economy. This process is called decentralization because the allocation is achieved as a consequence of individuals making optimizing decisions rather than the social planner imposing the allocation.

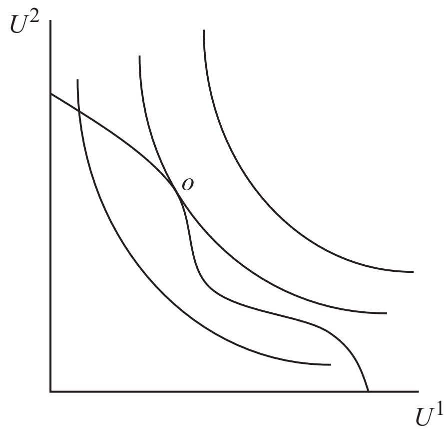

Figure 13.3
Social optimality

The decentralization argument shows that the use of the Second Theorem allows the economy to achieve the outcome most preferred by its social planner. Given the economy's limited initial stock of resources, the socially optimal allocation reaches the best trade-off between efficiency and equity as measured by the social welfare function. In this way the application of the Second Theorem can be said to solve the economic problem, since the issues of both efficiency and equity are resolved to the greatest extent possible and there is no better outcome attainable. Clearly, if this reasoning is applicable, all that a policy maker has to do is choose the allocation, implement the required lump-sum taxes, and ensure that the economy is competitive. No further policy or action is required. Once the incomes are set, the economy will take itself to the optimal outcome.

13.3 Lump-Sum Taxes

The role of lump-sum taxes has been made very explicit in describing the application of the Second Theorem. In the economic environment envisaged, lump-sum taxes are the only tool of policy that is required beyond an active competition policy. To justify the use of policies other than lump-sum taxes, it must be established that such taxes are either not feasible or else are restricted in the way in which they can be employed. This is the purpose of the next two sections. The results described are important in their own right, but they also provide important insights into the design of other forms of taxation.

In order for a tax to be lump sum, the consumer on whom the tax is levied must not be able to affect the size of the tax by changing their behavior. Most tax instruments encountered in practice are not lump sum. Income taxes cannot be lump sum by this definition because a consumer can work more or less hard and vary income in response to the tax. Similarly commodity taxes cannot be lump sum because consumption patterns can be changed. Estate duties are lump sum at the point at which they are levied (since, by definition, the person on which they are levied is dead and unable to choose any other action) but can be affected by changes in behavior prior to death (e.g., by making gifts earlier in life).

There are some taxes, though, that are close to being lump sum. For example, taxing every consumer some fixed amount imposes a lump-sum tax. Setting aside minor details, this was effectively the case of the UK Poll Tax levied in the late 1980s as a source of finance for local government. This tax was unsuccessful for two reasons. First, taxpayers could avoid paying the tax by ensuring that their names did not appear on any official registers. Usually this was achieved by moving house and not making any official declaration of the new address. It appears large numbers of taxpayers did this (unofficial figures put the number as high as 1 million). This "disappearance" is a change in behavior that reduces the tax burden. Second, the theoretical efficiency of lump-sum taxes rests partly on the fact that their imposition is costless, though this was far from the case with the Poll Tax. As it turned out, the difficulty of actually collecting and maintaining information on the residential addresses of all households made the imposition of a uniform lump-sum tax prohibitively expensive. The mobility of taxpayers proved to be much greater than had been expected. Therefore, although the structure of lump-sum taxes makes them appear deceptively simple to collect, this may not be the case in practice, since the tax base (people) is highly mobile and keen to evade. Consequently, in practice, even a uniform lump-sum tax has proved difficult and costly to administer.

However, the costs of collection are only part of the issue. The primary policy concern is the possibility of employing *optimal* lump-sum taxes. Optimal here means a tax that is chosen, via application of the Second Theorem, to achieve the income distribution necessary to decentralize the chosen allocation of the social planner. The optimal lump-sum tax system is unlikely to be a uniform tax on each consumer. This is because the role of the lump-sum taxes is fundamentally redistributive, so the taxes will be highly differentiated across consumers. Since even uniform lump-sum taxes are implemented with difficulty, the use of differentiated taxes presents even greater problems.

The extent of these problems can be seen by considering the information needed to calculate the taxes. First, the social planner must be able to construct the contract curve of Pareto-efficient allocations so that the social optimum can be selected. Second, the planner needs to predict the equilibrium that will emerge for all possible income levels so that the incomes needed to decentralize the chosen allocation can be determined. Both of these steps require knowledge of the consumers' preferences. Finally the social planner must also know the value of each consumer's endowment in order to calculate their incomes before taxes and hence the lump-sum taxes that must be imposed. The fundamental difficulty is that these economic characteristics, preferences and endowments, are *private information*. As such they are known only to the individual consumers and are not directly observed by the social planner. The characteristics may be partly revealed through market choices, but these choices can be changed if the consumers perceive any link with taxation. The fact that lump-sum taxes are levied on private information is the fundamental difficulty that hinders their use.

Some characteristics of the consumers are public information, or at least can be directly observed. Lump-sum taxes can then be levied on these characteristics. For example, it may be possible to differentiate lump-sum taxes according to characteristics of the consumers such as sex, age, or eye color. However, these characteristics are not those that are directly economically relevant as they convey neither preference information nor relate to the value of the endowment. Although we could differentiate taxes on this basis, there is no reason why we should want to do so.

This returns us to the problem of private information. Since the relevant characteristics such as ability are not observable, the social planner must either rely on consumers honestly reporting their characteristics or infer them from the observed economic choices of consumers. If the planner relies on the observation of choices, there is invariably scope for consumers to change their market behavior, which then implies that the taxes cannot be lump sum. When reports are the sole source of information, unobserved characteristics cannot form a basis for taxation unless the tax scheme is such that individuals are faced with incentives to report truthfully.

As an example of the interaction between taxes and reporting, consider the following. Let the quality of a consumer's endowment of labor be determined by their IQ level. Given a competitive market for labor, the value of the endowment is then related to IQ. Assume that there are no economically relevant variables other than IQ, so that any set of optimal lump-sum taxes must be levied on IQ. If the level of lump-sum tax was inversely related to IQ and if all households had to complete IQ tests, then the tax system would not be cheated because the incentive would always be to maximize the score on the test. In this case the lump-sum taxes are said to be *incentive compatible*, meaning that they give incentives to behave honestly. In contrast, if the taxes were positively related to IQ, a testing procedure could easily be manipulated by the high-IQ consumers who would intentionally choose to perform poorly. If such a system were put into place, the mean level of tested IQ would be expected to fall considerably. This indicates the potential for misrevelation of characteristics, and the system would not be incentive compatible. Clearly, if a high-IQ results in higher earnings and, ultimately, greater utility, a redistributive policy would require the use of lump-sum taxes that increased with IQ. The tax policy would not be incentive compatible. As the next section shows, such problems will always be present in any attempt to base lump-sum taxes on unobservable characteristics.

13.4 Impossibility of Lump-Sum Taxes

Imagine that each individual in a society can be described by a list of personal attributes upon which the society wishes to condition taxes and transfers (e.g. tastes, needs, talents, and endowments). Individuals are also identified by their names and possibly other publicly observable attributes (e.g., eye color), which are not judged to be relevant attributes for taxation. The list of personal attributes associated to every agent is not publicly known but is the private information of each individual. This implies that the lump-sum taxes the government would like to implement must rely on information about personal attributes which individuals must either report or reveal indirectly through their actions.

Lump-sum taxes are *incentive incompatible* when at least one individual, who understands how the information that is reported will be used, chooses to report falsely. We have already argued that there can be incentive problems in implementing optimal lump-sum taxes. What we now wish to demonstrate is that these problems are fundamental ones and will always afflict any attempt to implement optimal lump-sum taxes. In brief, the argument will show that optimal lump-sum taxes cannot be incentive compatible. This does not mean that lump-sum taxes cannot be used—for instance, all individuals could be taxed the same amount—but only that the existence of private information places limits on the extent to which taxes can be differentiated before incentives for the false revelation of information come into play. These issues are first illustrated for a particular example and then a general result is provided.

Before describing the general result a good illustration of the failure of incentive compatibility is provided in the following example due to Mirrlees. Assume that individuals can have one of two levels of ability: either low or high. The low ability level is denoted by s_l and the high ability level by s_h with $s_l < s_h$. For simplicity, suppose the number with high ability is equal to the number with low. The two types have the same preferences over consumption, x, and labor, ℓ as represented by the utility function $U(x, \ell) = u(x) - v(\ell)$. It is assumed that the marginal utility of consumption is decreasing in x and the marginal disutility of labor is increasing in ℓ.

To determine the optimal lump-sum taxes, suppose that the government can observe the ability of each individual and impose taxes that are conditioned upon ability. Let the tax on an individual of ability level i be $T_i > 0$ (or a subsidy if $T_i < 0$). The budget constraint of a type i is

$$x_i = s_i \ell_i - T_i, \tag{13.1}$$

where earnings are $s_i \ell_i$. Given the lump-sum taxes, each type chooses labor supply to maximize utility subject to this budget constraint. The choice of labor supply equates the marginal utility of additional consumption to the disutility of labor

$$s_i \frac{\partial u}{\partial x_i} - \frac{\partial v}{\partial \ell_i} = 0. \tag{13.2}$$

This provides a labor-supply function $\ell_i = \ell_i(T_i)$.

Now suppose that the government is utilitarian and chooses the lump-sum taxes to maximize the sum of utilities. Then the optimal lump-sum taxes solve

$$\max_{\{T_l, T_h\}} \sum_{l,h} u(s_i \ell_i (T_i) - T_i) - v(\ell_i(T_i)), \tag{13.3}$$

subject to government budget balance, which requires

$$T_h + T_l = 0, \tag{13.4}$$

since there are an equal number of the two types. This budget constraint can be used to substitute for T_l in (13.3). Differentiating the resulting expression with respect to the tax T_h and using the first-order condition (13.2) for the choice of labor supply, the optimal lump-sum taxes are characterized by the condition

$$\frac{\partial u}{\partial x_h} = \frac{\partial u}{\partial x_l}. \tag{13.5}$$

Since the marginal utility of consumption is decreasing in x_i, the optimality condition (13.5) implies that there is equality of consumption for the two types, $x_h = x_l$. When this conclusion is combined with (13.2) and the fact that $s_l < s_h$, it follows that

$$\frac{\partial v}{\partial \ell_l} = s_l \frac{\partial u}{\partial x_l} < s_h \frac{\partial u}{\partial x_h} = \frac{\partial v}{\partial \ell_h}. \tag{13.6}$$

Under the assumption of an increasing marginal disutility of labor, this inequality shows that the optimal lump-sum taxes should induce the outcome $\ell_h > \ell_l$, so the more able work harder than the less able. The motivation for this outcome is that working the high-ability type harder is the most efficient way to raise the level of total income for the society which can then be redistributed using the lump-sum taxes. Thus the high-ability type works harder than the low-ability type but only gets to consume the same. Therefore, the high-ability type is left with a lower utility level than the low-ability type after redistribution.

Now suppose that the government can observe incomes but cannot observe the ability of each individual. Assume that it still attempts to implement the optimal lump-sum taxes. The taxes are obviously not incentive compatible because, if the high–ability type understand the outcome, they can always choose to earn as little as the low–ability type. Doing so then qualifies the high–ability type for the redistribution aimed at the low-ability type. This will provide them with a higher utility level than if they did not act strategically. The optimal lump-sum taxes cannot then be implemented with private information.

Who would work hard if the government stood ready to tax away the resulting income? Optimal (utilitarian) lump-sum redistribution makes the more able individuals worse off because it requires them to work harder but does not reward them with additional consumption. In this context it is profitable for the more able individuals to make themselves seem incapable. Many people believe there is something unfair about inequality that arises from the fact that some people are born with superior innate ability or similar advantage over others. But many people also think it morally right that one should be able to keep some of the fruits of one's own effort. This example may have been simple but its message is far-reaching. The Soviet Union and other communist economies have shown us that it is impossible to generate wealth without simultaneously offering adequate material incentives. Incentive constraints inevitably limit the scope for redistribution.

This example is now shown to reflect a general principle concerning the incentive compatibility of optimal lump-sum taxes. We state the formal version of this result for a "large economy," which is an economy where the actions of an individual are insignificant relative to the economy as a whole. In other words, there is a continuum of different agents, which is the mathematical form of the idealized competitive economy with a very large number of small agents with no market power. The theorem shows that optimal lump-sum taxation is never incentive compatible.

Theorem 13.1 (Hammond) In a large economy, redistribution through optimal lump-sum taxes is always incentive incompatible.

The logic behind this theorem is surprisingly simple. A system of optimal lump-sum taxes is used to engineer a distribution of endowments that will decentralize the first-best allocation. The endowments after redistribution must be based on the agents' characteristics (recall that in the analysis of the Second Theorem the taxes were based on knowledge of endowments and preferences), so assume the endowment of an agent with characteristics θ_i is given by $e^i = e(\theta_i)$. For those characteristics that are not publicly observable, the government must rely on an announcement of the values by the agents. Assume, for simplicity, that none of the characteristics can be observed. Then the incentive exists for each agent to announce the set of characteristics that maximize the value of the endowment at the equilibrium prices p. This is illustrated in figure 13.4 where θ_1 and θ_2 are two potential announcements, with related endowments $e(\theta_1)$ and $e(\theta_2)$, and θ^* is the announcement that maximizes $pe(\theta)$. The announcement of θ^* leads to the highest budget constraint from among the set of possible announcements and, by giving the agent maximum choice, allows the highest level of utility to be attained. Consequently all agents will announce θ^* and the optimal lump-sum taxes are not incentive compatible.

The main points of the argument can now be summarized. To implement the Second Theorem as a practical policy tool, it is necessary to employ optimal lump-sum taxes. Such taxes are unlikely to be available in practice or to satisfy all the criteria required

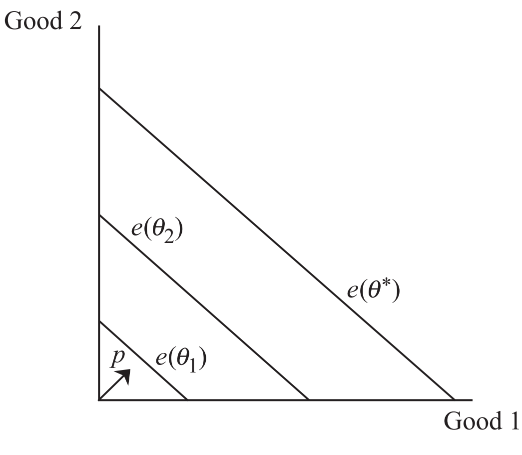

Figure 13.4
Optimal lump-sum taxes and incentive compatibility

of them. The taxes may be costly to collect and the characteristics on which they need to be based may not be observable. When characteristics are not observable, the relationship between taxes and characteristics can give consumers the incentive to make false revelations. It is therefore best to treat the Second Theorem as being of considerable theoretical interest but of very limited practical relevance. The theorem shows us what could be possible, not what is possible.

Lump-sum taxes can achieve the optimal allocation of resources provided all information is public. If some of the characteristics that are relevant for taxation are private information, then the optimal lump-sum taxes are not incentive compatible. Information limitations therefore place a limit on the extent to which redistribution can be undertaken using lump-sum taxation. It is the impracticality of lump-sum taxation that provides the motive for studying the properties of other tax instruments. The income taxes and commodity taxes that are analyzed in chapters 16 and 15 are second-best solutions and are used because the first-best solution, lump-sum taxation, is not available. Lump-sum taxes are used as a benchmark from which to judge the relative success of these alternative instruments. Lump-sum taxes also help clarify what it is that we are really trying to tax.

13.5 Redistribution In-Kind

The lump-sum taxes we have been discussing are a very immediate form of redistribution. In practice, there are numerous widely used methods of redistribution that do not directly involve taxation. Governments frequently provide goods such as education or health services at less than their cost, which may be viewed as a redistributional policy. One may expect that a cash transfer of the same value would have more redistributional power than such in-kind transfer programs. This is mistaken. There are three reasons why transfers in-kind may be superior to the cash transfers achieved through standard tax-transfer programs.

One reason is *political*. Political considerations dictate that many governments ensure that the provision of programs like education, pension, and basic health insurance is universal. Without this feature the programs would not have the political support required to be adopted or continued. For instance, public pensions and health care would be far more vulnerable politically if they were targeted to the poor and not available to others. Redistribution through cash would be even more vulnerable. It should be noted that because a government program is universal, it does not follow that there is no redistribution. First, if the program is financed by proportional income taxation, the rich will contribute more to its finance than the poor. Second, even if everyone contributes the same to the program, it is possible that the rich will not use the publicly provided good to the same extent as the poor. Consider, for example, a program of public provision of basic health care that is available to everyone for free and financed by a uniform tax on all households. Assume that there exists a private health care alternative with higher quality than the public system but only available at a cost. Since the rich can afford the higher quality, they will use the private health care, even though free public health care is available. These rich households still pay their contribution to the public program, and thus the poor households derive a net benefit from this cross-subsidization.

Another reason for preferring in-kind redistribution is *self-selection*. What ultimately limits redistribution is that it will eventually become advantageous for higher ability people to earn lower incomes by expending less effort and thereby paying the level of taxes (or receiving the transfers) intended for the lower ability groups. The self-selection argument is that anything that makes it less attractive for people to mimic those with lesser ability will extend the limit to redistribution. The use of in-kind transfers can obtain a given degree of redistribution more efficiently because of differences in preferences among different income groups. Consider two individuals who differ not only in their ability but also in their health status. Suppose that lesser ability means also poorer health, so the less able spend relatively more on health. Then both income and health expenditures act as a signal of ability. It follows that the limits to redistribution can be relaxed if transfers are made partly in the form of provision of health care (or equivalently with full subsidization of health expenditures). The reason is simply that the more able individual (with less tendency to become ill) is less likely to claim in-kind benefits in the form of health care provision than he would be to claim cash benefits. To take another example, suppose that the government is considering redistribution either in cash or in the form of low-quality housing. All households, needy or not, would like the cash transfer. However, few non-needy households would want to live in low-quality housing as they can afford better housing. Thus self-selection occurs, and the non-needy drop out of the housing program, which is taken up only by the needy. In short, transfers in-kind invite people to self-select in a way that reveals their neediness. When need is correlated with income-earning ability, then in-kind transfers can relax incentive and selection constraints, thereby improving the government's ability to redistribute income.

A third reason is the idea of *time consistency* that we introduced in chapter 3. Here the argument for in-kind transfers relies on the inability of government to commit to its future actions. Unlike the argument of Strotz (1956) on government time inconsistency, this does not arise from a change in government objective over time (e.g., because of elections) nor from the fact that the government is not welfaristic or rational. The time-consistency problem arises from a perfectly rational government that fully respects individual preferences but that does not have the power to commit to its policy in the long run. The time-consistency problem is obvious with regard to pensions. To the extent that households expect governments to provide some basic pension to those with too little savings, their incentive to save for retirement consumption and provide for themselves is reduced. Anticipating this, the government may prefer to provide public pensions. A related time-consistency problem can explain why transfer programs, such as social security, education, and job training are in-kind. If a welfaristic government cannot commit not to come to the rescue of those in need in the future, potential recipients will have little reason to invest in their education or to undertake job training because the government will help them out anyway. Again, the government can improve both economic efficiency and redistribution by making education and job training available at less than their cost, rather than making cash transfers of equivalent value.

13.6 Aspects of Pareto-Efficiency

The analysis of lump-sum taxation has raised questions about the practical value of the Second Theorem of Welfare Economics. Although the theorem shows how an optimal allocation can be decentralized, the means to achieve the decentralization may be absent. If the use of lump-sum taxes is restricted, the government must resort to alternative policy instruments. All alternative instruments will be distortionary and will not achieve the first-best.

These criticisms do not extend to the First Theorem of Welfare Economics, which states only that a competitive equilibrium is Pareto-efficient. Consequently the First Theorem implies no policy intervention, so it is safe from the restrictions on lump-sum taxes. However, at the heart of the First Theorem is the use of Pareto-efficiency as a method for judging the success of an economic allocation. The value of the First Theorem can only be judged once a deeper understanding of Pareto-efficiency has been developed.

The Pareto criterion was introduced into economics by the Italian economist Vilfredo Pareto at the beginning of the twentieth century. This was a period of reassessment in economics during which the concept of utility as a measurable entity was rejected. Alongside this rejection of measurability, the ability to compare utility levels between consumers also had to be rejected. Pareto-efficiency was therefore constructed explicitly to allow comparisons of allocations without the need to make any interpersonal comparisons of utility. As will be seen, this avoidance of interpersonal comparisons is both its strength and its main weakness.

To assess Pareto-efficiency, it is helpful to develop the concept in three stages. The first stage defines the idea of making a *Pareto improvement* when moving from one allocation to another. From this can be constructed the *Pareto preference* order that judges whether one allocation is preferred to another. The final stage is to use Pareto preference to find the most preferred states, which are then defined as *Pareto-efficient*. Reviewing each of these steps allows us to assess the meaning and value of the concept.

Consider a move from economic state s_1 to state s_2. This is defined as a Pareto improvement if it makes some consumers strictly better off and none worse off. If there are H consumers, this definition can be stated formally by saying a Pareto improvement is made in going from s_1 to s_2 if

$$U^h(s_2) > U^h(s_1) \quad \text{for at least one consumer, } h, \tag{13.7}$$

and

$$U^h(s_2) \geq U^h(s_1) \quad \text{for all consumers } h = 1, \ldots, H. \tag{13.8}$$

The idea of a Pareto improvement can be used to construct a preference order over economic states. If a Pareto improvement is made in moving from s_1 to s_2, then state s_2 is defined as being *Pareto-preferred* to state s_1. This concept of Pareto preference defines one state as preferred to another if all consumers are at least as well off in that state and some are strictly better off. It is important to note that this stage of the construction has converted the set of individual preferences of the consumers into social preferences over the states.

The final stage is to define Pareto-efficiency. The earlier definition can be re-phrased as saying that an economic state is *Pareto-efficient* if there is no state that is Pareto-preferred to it. That is, no move can be made from that state to another that achieves a Pareto improvement. From this perspective, we can view Pareto-efficient states as being the “best” relative to the Pareto preference order. The discussion now turns to assessing the usefulness of Pareto preference in selecting an optimal state from a set of alternatives. By analyzing a number of examples, several deficiencies of the concepts will become apparent.

The simplest allocation problem is to divide a fixed quantity of a single commodity between two consumers. Let the commodity be a cake, and assume that both consumers prefer more cake to less. The first observation is that no cake should be wasted—it is always a Pareto improvement to move from a state where some is wasted to one with the wasted cake given to one, or both, of the consumers. The second observation is that any allocation in which no cake is wasted is Pareto-efficient. To see this, start with any division of the cake between the two consumers. Any alternative allocation must give more to one consumer and less to the other; therefore, since one must lose some cake, no change can be a Pareto improvement.

From this simple example two deficiencies of Pareto-efficiency can be inferred. First, since no improvement can be made on an allocation where none is wasted, extreme allocations such as giving all of the cake to one consumer are Pareto-efficient. This shows that even though an allocation is Pareto-efficient, there is no implication that it need be good in terms of equity. This illustrates quite clearly that Pareto-efficiency is not concerned with equity. The cake example also illustrates a second point: there can be a multiplicity of Pareto-efficient allocations. This was shown in the cake example by the fact that every nonwasteful allocation is Pareto-efficient. This multiplicity of efficient allocations limits the value of Pareto-efficiency as a tool for making allocative decisions. For the cake example, Pareto-efficiency gives no guidance whatsoever in deciding how the cake should be shared, other than showing that none should be thrown away. In brief, Pareto-efficiency fails to solve even this simplest of allocation problems.

The points made in the cake division example are also relevant to allocations within a two-consumer exchange economy. The contract curve in figure 13.5 shows the set of Pareto-efficient allocations, and there is generally an infinite number of these. Once again the Pareto preference ordering does not select a unique optimal outcome. In addition the competitive equilibrium may be as the one illustrated in the bottom left corner of the box. This has the property of being Pareto-efficient, but it is highly inequitable and may not find much favor using other criteria for judging optimality.

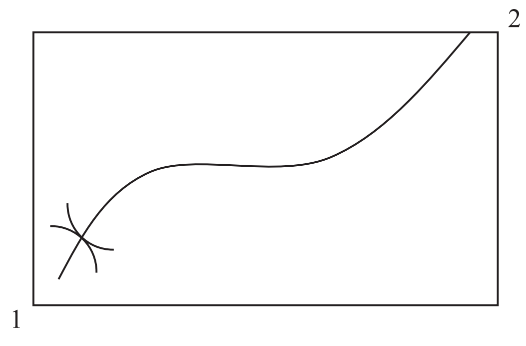

Figure 13.5
Efficiency and inequity

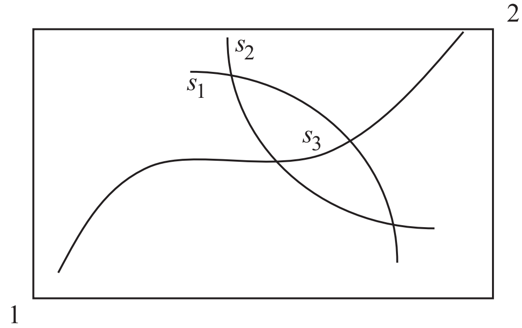

Figure 13.6
Incompleteness of Pareto ranking

Another failing of the Pareto preference ordering is that it is not always able to compare alternative states. In formal terms, it does not provide a *complete ordering* of states. This is illustrated in figure 13.6 where the allocations s_1 and s_2 cannot be compared, although both can be compared to s_3 (s_3 is Pareto-preferred to both s_1 and s_2). When faced with a choice between s_1 and s_2, the Pareto preference order is silent about which should be chosen. It should be noted that this incomparability is not the same as indifference. If the preference order were indifferent between two states, then they are judged as equally good. Incomparability means the pair of states simply cannot be ranked.

The basic mechanism at work behind this example is that the Pareto preference order can only rank alternative states if there are only gainers or only losers as the move is made between the states. If some gain and some lose, as in the choice between s_1 and s_2 in figure 13.6, then the preference order is of no value. Such gains and losses are invariably a feature of policy choices and much of policy analysis consists of weighing up the gains and losses. In this respect Pareto-efficiency is insufficient as a basis for policy choice.

To summarize these arguments, Pareto-efficiency does not embody any concept of justice, and highly inequitable allocations can be efficient under the criterion. In many situations there are very many Pareto-efficient allocations, in which case the criterion provides little guidance for policy choice. Finally Pareto-efficiency may not provide a complete ordering of states, so some states will be incomparable under the criterion. The source of all these failing is that the Pareto criterion avoids weighing gains against losses, but it is just such judgments that have to be made in most allocation decisions. To make a choice of allocation, the evaluation of the gains and losses has to be faced directly.

13.7 Social Welfare Functions

The social welfare function was employed in section 13.2 to introduce the concept of a socially optimal allocation. At that point it was simply described as a means by which different allocations of utility between consumers could be socially ranked. What was not done was to provide a convincing description of where such a ranking could come from or of how it could be constructed. Three alternative interpretations will now be given, each of which provides a different perspective on the social welfare function.

The first possibility is that the social welfare function captures the distributive preferences of a central planner or dictator. Under this interpretation there can be two meanings of the individual utilities that enter the function. One is that they are the planner's perception of the utility achieved by each consumer at their level of consumption. This provides a consistent interpretation of the social welfare function, but problems arise in its relation to the underlying model. To see why this is so, recall that the Edgeworth box and the contract curve within it were based on the actual preferences of the consumers. There is then a potential inconsistency between this construction and the evaluation using the planner's preferences. For example, what is Pareto-efficient under the true preferences may not be one under the planner's (it need not even be an equilibrium).

The alternative meaning of the utilities is that they are the actual utilities of the consumers. This leads directly into the central difficulty faced in the concept of social welfare. In order to evaluate all allocations of utility it must be possible to determine the social value of an increase in one consumer's utility against the loss in another's. This is only possible if the utilities are comparable across the consumers. More will be said about this below.

The second interpretation of the social welfare function is that it captures some ethical objective that society should be pursuing. Here the social welfare function is determined by what is viewed as the just objective of society. There are two major examples of this. The *utilitarian philosophy* of aiming to achieve the greatest good for society as a whole translates into a social welfare function that is the sum of individual utilities. In this formulation only the total sum of utilities counts, so it does not matter how utility is distributed among consumers in the society. Alternatively, the *Rawlsian philosophy* of caring only for the worst-off member of society leads to a level of social welfare determined entirely by the minimum level of utility in that society. With this objective the distribution of utility is of paramount importance. Gains in utility achieved by anyone other than the worst-off consumer do not improve social welfare.

Although this approach to the social welfare function is internally consistent, it is still not entirely satisfactory. The utilitarian approach requires that the utilities of the consumers be added in order to arrive at the total sum of social welfare. The Rawlsian approach necessitates the utility levels being compared in order to find the lowest. The nature of the utility comparability is different for the two approaches (being able to add utilities is different to being able to compare), but both rely on some form of comparability. This again leads directly into the issue of utility comparisons.

The final view that can be taken of the social welfare function is that it takes the preferences of the individual consumers (represented by their utilities) and aggregates these into a social preference. This aggregation process would be expected to obey certain rules; for instance, if all consumers prefer one state to another, it should be the case that the social preference also prefers the same state. The structure of the social welfare function then emerges as a consequence of the rules the aggregation must obey.

Although this arrives at the same outcome as the other two interpretations, it does so by a distinctly different process. In this case it is the set of rules for aggregation that are foremost rather than the form of social welfare. That is, the philosophy here would be that if the aggregation rules are judged as satisfactory, then society should accept the social welfare function that emerges from their application, whatever its form. An example of aggregating preferences is the rule of majority voting (despite the failings already identified in chapter 11), since the minority must accept what the majority chooses.

The consequences of constructing a social welfare function by following this argument are of fundamental importance in the theory of welfare economics. In fact doing so leads straight back into Arrow's Impossibility Theorem, which was described in chapter 11. The next section is dedicated to interpreting the theorem and its implications in this new setting.

13.8 Arrow's Theorem

Although they appear very distinct in nature, both majority voting and the Pareto criterion are examples of procedures for aggregating individual preferences into a social preference. It has been shown that neither is perfect. The Pareto preference order can be incomplete and unable to rank some of the alternatives. Majority voting always leads to a complete social preference order, but this may not be transitive. What Arrow's Impossibility Theorem has shown is that such failings are not specific to these aggregation procedures. All methods of aggregation will fail to meet one or more of its conditions, so the Impossibility Theorem identifies a fundamental problem at the heart of generating social preferences from individual preferences.

The conditions of Arrow's theorem were stated in terms of the rankings induced by individual preferences. However, since individual preferences can usually be represented by a utility function, the theorem also applies to the aggregation of individual utility functions into a social welfare function. The implication behind applying the theorem is that a social welfare function does not exist that can aggregate individual utilities without conflicting with one, or more, of the conditions $I.N.P.U.T$. This means that whatever social welfare function is proposed, there will be some set of utility functions for which it conflicts with at least one of the conditions. In other words, no ideal social welfare function can be found. No matter how sophisticated the aggregation mechanism is, it cannot overcome this theorem.

Since the publication of Arrow's theorem there has been a great deal of research attempting to find a way out of the dead end into which it leads. One approach that has been tried is to consider alternative sets of aggregation rules. For instance, transitivity of the social preference ordering can be relaxed to quasi-transitivity (only strict preference is transitive) or weaker versions of condition I and condition P can be used. Most such changes just lead to further impossibility theorems for these different sets of rules. Modifying the rules does not therefore really seem to be the way forward out of the impossibility.

What is at the heart of the impossibility is the limited information contained in individual utility functions. Effectively all that is known is the individuals' rankings of the alternatives—which is best, which is worst, and how they line up in between. What the rankings do not give is any strength of feeling either between alternatives for a given individual or across individuals for a given option. Such strength of feeling is an essential art of any attempt to make social decisions. Consider, for instance, a group of people choosing where to dine. In this situation a strong preference in one direction ("I really don't want to eat fish") usually counts for more than a mild preference ("I don't really mind, but I would prefer fish"). Arrow's theorem rules out any information of this kind.

Using information on how strongly individuals feel about the alternatives can be successful in choosing where to dine. It is interesting that the strength of preference comparisons can be used in informal situations, but this does not demonstrate that it can be incorporated within a scientific theory of social preferences. This issue is now addressed in detail.

13.9 Interpersonal Comparability

Earlier in this chapter it was noted that Pareto-efficiency was originally proposed because it provided a means by which it was possible to compare alternative allocations without requiring interpersonal comparisons of welfare. It is also from this avoidance of comparability that the failures of Pareto-efficiency emerge. This point is at the core of the Impossibility Theorem. To proceed further, this section first reviews the development of utility theory in order to provide a context and then describes alternative degrees of utility comparability.

Nineteenth-century economists viewed utility, the level of happiness of an individual, as something that was potentially measurable. Advances in psychology were expected to deliver the machinery for conducting the actual measurement. If utility were measurable, it follows naturally that it would be comparable among individuals. This ability to measure utility, combined with the philosophy that society should aim for the greatest good, came to provide the underpinnings of utilitarianism. The measurability of utility permitted social welfare to be expressed by the sum of individual utilities. Ranking states by the value of this sum then gave a means of aggregating individual preferences that satisfied all of the conditions of the impossibility theorem except for the information content. If the envisaged degree of measurability could be achieved, then the restrictions of the impossibility theorem are overcome.

This concept of measurable and comparable utility began to be dispelled in the early twentieth century. There were two grounds for this rejection. First, no means of measuring utility had been discovered, and it was becoming clear that the earlier hopes would not be realized. Second, advances in economic theory showed that there was no need to have measurable utility in order to construct a coherent theory of consumer choice. In fact the entire theory of the consumer could be derived by specifying only the consumer's preference ordering. The role of utility then became strictly secondary—it could be invoked to give a convenient function to represent preferences if necessary but was otherwise redundant. Since utility had no deeper meaning attached to it, any increasing monotonic transformation of a utility function representing a set of preferences would also be an equally valid utility function. Utility was simply an ordinal concept, with no natural zero or units of measurement. By the very construction of utility, comparability between different consumers' utilities was a meaningless concept. This situation therefore left no scientific basis on which to justify the comparability of different consumer's utility levels.

This perspective on utility, and the consequent elimination of utility comparisons among consumers, created the need to develop concepts for social comparisons, such as Pareto-efficiency, that were free of interpersonal comparisons. However, the weaknesses of these criteria soon became obvious. The analytical trend since the 1960s has been to explore the consequences of re-admitting interpersonal comparability into the analysis. The procedure adopted is basically to assume that comparisons are possible. This permits the derivation of results from which interpretations can be obtained. These are hoped to provide some general insights into policy that can be applied, even though utility is not actually comparable in the way assumed.

There are even some economists who would argue that comparisons are possible. One basis for this is the claim that all consumers have very similar underlying preference orderings. All prefer to have more income to less, and consumers with equal incomes make very similar divisions of expenditures between alternative groups of commodities. For example, expenditure on food is similar, even though the actual foodstuffs purchased may be very different. In modeling such consumers, it is possible to assert that they all have the same utility function guiding their choices. This makes their utilities directly comparable.

So far comparability has been used as a catch-all phrase for being able to draw some contrast between the utility levels of consumers. In fact many different degrees of comparability can be envisaged. For instance, the claim that one household has a higher level of utility than another requires rather less comparability than claiming it has 15 percent more utility. Different degrees of comparability have implications for the way in which individual utilities can be aggregated into a social preference ordering.

The starting point for discussing comparability is to define the two major forms of utility. The first is *ordinal utility*, which is the familiar concept from consumer theory. Essentially an ordinal utility function is no more than just a numbering of a consumer's indifference curves, with the numbering chosen so that higher indifference curves have higher utility numbers. These numbers can be subjected to any form of transformation without altering their meaning, provided that the transformation leaves the ranking of the numbers unchanged—higher indifference curves must still have larger utility numbers attached. Because they can be so freely transformed, there is no meaning to differences in utility levels between two situations for a single consumer except which of the two provides the higher utility.

The second form of utility is *cardinal utility*. Cardinal utility imposes restrictions beyond those of ordinal utility. With cardinal utility one can only transform utility numbers by multiplying by a constant and then adding a constant, so an initial utility function U becomes the transformed utility $\tilde{U} = a + bU$, where a and b are the constants. Any other form of transformation will affect the meaning of a cardinal utility function. The typical place where cardinal utility is found is in the economics of uncertainty, since an expected utility function is cardinal. This cardinality is a consequence of the fact that an expected utility function must provide a consistent ranking for different probability distributions of the outcomes. (A noneconomic example of a cardinal scale is temperature. It is possible to convert Celsius to Fahrenheit by multiplying by $\frac{9}{5}$ and adding 32. The converse transformation from Fahrenheit to Celsius is to multiply by $\frac{5}{9}$ and subtract 32.) With these definitions it now becomes possible to talk in detail about comparability and noncomparability.

Noncomparability can arise with both ordinal and cardinal utility. What noncomparability means is that we can apply different transformations to different consumers' utilities. To express this in formal terms, let U^1 be the utility function of consumer 1 and U^2 the utility function of consumer 2. Then noncomparability arises if the transformation f^1 can be applied to U^1 and a different transformation f^2 to U^2, with no relationship between f^1 and f^2. Why is this noncomparable? The reasoning is that by suitably choosing f^1 and f^2, it is always possible to start with one ranking of the initial utilities and to arrive at a different ranking of the transformed utilities. The utility information therefore does not provide sufficient information to make a comparison of the two utility levels.

Comparability exists when the transformations that can be applied to the utility functions are restricted. With ordinal utility there is only one possible degree of comparability. This occurs when the ordinal utilities for different consumers can be subjected only to the same transformation. The implication of this is that the transformation preserves the ranking of utilities among different consumers. So, if one consumer has a higher utility than another before the transformation, the same consumer will have a higher utility after the transformation. Letting this transformation be denoted by f, then if $U^1 \geq U^2$, it must be the case that $f(U^1) \geq f(U^2)$. This form of comparability is called *ordinal level comparability*.

If the underlying utility functions are cardinal, there are two forms of comparability that are worth discussing. The first form of comparability is to assume that the constant multiplying of utility in the transformation must be the same for all consumers, but the constant that is added can differ. Hence for two consumers the transformed utilities are $\tilde{U}^1 = a^1 + bU^1$ and $\tilde{U}^2 = a^2 + bU^2$, so the constant b is the same for both. This is called *cardinal unit comparability*. The implication of this transformation is that it now becomes meaningful to talk about the effect of changes in utility, meaning that gains to one consumer can be measured against losses to another—and whether the gain exceeds the loss is not affected by the transformation. The second degree of comparability for cardinal utility is to further restrict the constant a in the transformation to be the same for both consumers. For all consumers the transformed utility becomes $\widetilde{U}^h = a + bU^h$. It is now possible for both changes in utility and levels of utility to be compared. This form of comparability is called *cardinal full comparability*.

The next step is to explore the implications of these comparabilities for the construction of social welfare functions. It will be shown that each form of comparability implies different permissible social welfare functions.

13.10 Comparability and Social Welfare

The discussion of Arrow's Impossibility Theorem showed that the failure to successfully generate a social preference ordering from a set of individual preference orderings was the result of limited information. The information content of an individual's preference order involves nothing more than knowing how they rank the alternatives. A preference order does not convey any information on the strength of preferences or allow comparison of utility levels across consumers. When more information is available, it becomes possible to find social preference orderings that satisfy the conditions I, N, P, U, T. Such information can be introduced by building social preferences on individual utility functions that allow for comparability.

What this section shows is that for each form of comparability there is a specification of social welfare function that is consistent with the information content of the comparable utilities. To explain what is meant by consistent, recall that comparability is described by a set of permissible transformations of utility. A social welfare function is *consistent* if it ranks the set of alternative social states in the same way for all permissible transformations of the utility functions. Since increasing the degree of comparability reduces the number of permissible transformations, it has the effect of increasing the set of consistent social welfare functions.

Let the utility obtained by consumer h from allocation s be $U^h(s)$. A transformation of this basic utility function is denoted by $\tilde{U}^h(s) = f^h\big(U^h(s)\big)$. The value of social welfare at allocation s using the basic utilities is $W(s) = W\big(U^1(s), \ldots, U^H(s)\big)$, and that from using the transformed utilities is $\tilde{W}(s) = W\Big(\tilde{U}^1(s), \ldots, \tilde{U}^H(s)\Big)$. Given alternative allocations A and B, the social welfare function is consistent with the transformation (and hence the form of comparability) if $W(A) \geq W(B)$ implies $\tilde{W}(A) \geq \tilde{W}(B)$. In words, if A generates higher social welfare than B for the basic utilities, it will also do so for the transformed utilities.

Table 13.1
Allocations and utility

	x^1	y^1	U^1	x^2	y^2	U^2
A	4	9	6	3	2	5
B	16	1	4	2	5	7

To demonstrate these points, assume there are two consumers with the basic utility functions $U^1 = [x]^{1/2}[y]^{1/2}$ and $U^2 = x + y$, where x and y are the consumption levels of the two goods. Further assume that there are two allocations A and B with the consumption levels, and the resulting utilities, as shown in table 13.1.

The first point to establish is that it is possible to find a social welfare function that is consistent with ordinal level comparability but none that is consistent with ordinal noncomparability. What level comparability allows is the ranking of consumers by utility level (think of placing the consumers in a line with the lowest utility level first). A position in this line (e.g., the first, or the tenth, or the nth) can be chosen, and the level of utility of the consumer in that position used as the measure of social welfare. This process generates a *positional* social welfare function. The best known example is the Rawlsian social welfare function, $W = \min\{U^h\}$, which judges social welfare by the minimum level of utility in the population. An alternative that shows other positions can be employed (though not one that is often used) is to measure social welfare measure by the maximum level of utility, $W = \max\{U^h\}$.

That such positional welfare functions are consistent with ordinal level comparability but not with ordinal noncomparability is shown in table 13.2 using the allocations A and B introduced above. For the social welfare function $W = \min\{U^h\}$, the welfare level in allocation A is 5 and that in allocation B is 4. Therefore allocation A is judged superior using the basic utilities. An example of a pair of transformations that satisfy ordinal noncomparability are $\tilde{U}^1 = f^1(U^1) = 3U^1$ and $\tilde{U}^2 = f^2(U^2) = 2U^2$. The levels of utility and resulting social welfare are displayed in the upper part of table 13.2. The table shows that the preferred allocation is now B, so the transformation has changed the preferred social outcome. With ordinal level comparability, the transformations $f^1(U^1)$ and $f^2(U^2)$ must be the same. For example, let the transformation be given by $\tilde{U}^h = f(U^h) = (U^h)^2$. The values of the transformed utilities in the lower part of the table confirm that allocation A is preferred—as it was with the basic utilities. The positional social welfare function is therefore consistent with ordinal level comparability.

Table 13.2
Noncomparability and level comparability

Noncomparability	A	B
$\tilde{U}^1 = f^1(U^1) = 3U^1$	18	12
$\tilde{U}^2 = f^2(U^2) = 2U^2$	10	14
$W = \min\{\tilde{U}^h\}$	10	12
Level comparability	A	B
$\tilde{U}^1 = f(U^1) = (U^1)^2$	36	16
$\tilde{U}^2 = f(U^2) = (U^2)^2$	25	49
$W = \min\{\tilde{U}^h\}$	25	16

Although cardinal utility is often viewed as stronger concept than ordinal utility, cardinality alone does not permit the construction of a consistent social welfare function. Recalling that transformations of the form $f^h = a^h + b^h U^h$ can be applied with noncomparability, it can be seen that even positional welfare functions will not be consistent, since a^h can always be chosen to change the social ranking generated by the transformed utilities compared to that generated by the basic utilities. In contrast, if utility satisfies cardinal unit comparability, it is possible to use social welfare functions of the form

$$W = \sum_{h=1}^{H} \alpha^h U^h, \tag{13.9}$$

where the α^h are constants. To demonstrate this, and to show that social welfare function is not consistent with cardinal noncomparability, assume that $\alpha^1 = 2$ and $\alpha^2 = 1$. Then, under the basic utility functions, the social welfare levels in the two allocations are $W(A) = 2 \times 6 + 5 = 17$ and $W(B) = 2 \times 4 + 7 = 15$, so allocation A is preferred. The upper part of table 13.3 displays two transformations satisfying noncomparability and the implied value of social welfare. This shows that allocation B will be preferred with the transformed utility. Therefore the social welfare function is not consistent with the transformations. With cardinal unit comparability, the transformations are restricted to have a common value for b^h, so $\tilde{U}^h = a^h + bU^h$. Two such transformations are selected, and the resulting utility levels are given in the lower part of the table. Calculation of the social welfare shows the preferred allocation to be A as it was with the basic utilities. Therefore with cardinal level comparability, social welfare functions

Table 13.3
Cardinal utility

Noncomparability	A	B
$\tilde{U}^1 = f^1(U^1) = 2 + 2U^1$	14	10
$\tilde{U}^2 = f^2(U^2) = 5 + 6U^2$	35	47
$W = 2\tilde{U}^1 + \tilde{U}^2$	63	67
Level comparability	A	B
$\tilde{U}^1 = f^1(U^1) = 2 + 3U^1$	20	14
$\tilde{U}^2 = f^2(U^2) = 5 + 3U^2$	20	26
$W = 2\tilde{U}^1 + \tilde{U}^2$	60	54

of the form (13.9) are consistent and provide a social ranking that is invariant for the permissible transformations.

With cardinal full comparability the transformations must satisfy $\tilde{U}^h = a + bU^h$. One interesting example of the forms of social welfare function that are consistent with such transformations is

$$W = \bar{U} + \gamma \min\left\{U^h - \bar{U}\right\}, \quad \bar{U} = \frac{\sum_{h=1}^{H} U^h}{H}, \tag{13.10}$$

where γ is a parameter that can be chosen. This form of social welfare function is especially interesting because it is the utilitarian social welfare function when $\gamma = 0$ and Rawlsian when $\gamma = 1$. To show that this function is not consistent for cardinal unit comparability, assume $\gamma = \frac{1}{2}$. For the basic utilities it follows for allocation A that $\bar{U} = \frac{6+5}{2} = 5.5$ and for allocation B, $\bar{U} = \frac{4+7}{2} = 5.55$. The social welfare levels are then $W = 5.5 + \frac{1}{2}\min\{6 - 5.5, 5 - 5.5\} = 5.25$ for allocation A and $W = 5.5 + \frac{1}{2}\min\{4 - 5.5, 7 - 5.5\} = 4.75$ for allocation B. The social welfare function would select allocation A. The upper part of table 13.4 displays the welfare levels for two transformations that satisfy cardinal level comparability. With these transformed utilities the welfare function would select allocation B, so the social welfare function is not valid for these transformations. The lower part of the table displays a transformation that satisfies cardinal full comparability. For this transformation the social welfare function selects allocation A for both the basic and the transformed utilities. This demonstrates the consistency.

These calculations have demonstrated that if we can compare utility levels among consumers, then a consistent social welfare function can be constructed. The resulting

Table 13.4
Level comparability and full comparability

Level comparability	A	B
$\tilde{U}^1 = f^1(U^1) = 7 + 3U^1$	25	19
$\tilde{U}^2 = f^2(U^2) = 1 + 3U^2$	16	22
$W = \bar{U} + \frac{1}{2}\min\{\tilde{U}^h - \bar{U}\}$	18.25	19.75
Full comparability	A	B
$\tilde{U}^1 = f(U^1) = 1 + 3U^1$	19	13
$\tilde{U}^2 = f(U^2) = 1 + 3U^2$	16	22
$W = \bar{U} + \frac{1}{2}\min\{\tilde{U}^h - \bar{U}\}^2$	16.75	15.25

social welfare function must agree with the information content in the utilities, so each form of comparability leads to a different consistent social welfare function. As the information increases, so does the range of consistent social welfare functions. Expressed differently, for each of the cases of comparability the problem of aggregating individual preferences leads to a well-defined form of social welfare function. All these social welfare functions will generate a social preference ordering that completely ranks the alternative states. They are obviously stronger in content than majority voting or Pareto-efficiency. The drawback is that they are reliant on stronger utility information that may simply not exist.

13.11 Conclusions

This chapter has cast a critical eye over the efficiency theorems of chapter 2. Although these theorems are important for providing a basic framework in which to think about policy, they are not an end in their own right. This perspective is based on the limited practical applicability of the lump-sum transfers needed to support the decentralization in the Second Theorem and the weakness of Pareto-efficiency as a method of judging among economic states.

Although at first sight the theorems apparently have very strong policy implications, they become weakened when placed under critical scrutiny. But they are not without value. Much of the subject matter of public economics takes as its starting point the practical shortcomings of these theorems and attempts to find a way forward to something that is applicable. A knowledge of what could be achieved if the optimal lump-sum transfers were available provides a means of assessing the success of what can be achieved and shows ways in which improvements in policy can be made.

The other aspect involved in the Second Theorem is the selection of the optimal allocation to be decentralized. This choice requires a social welfare function that can be used to judge different allocations of utility among consumers. Such a social welfare function can only be constructed if the consumers' utilities are comparable. The chapter described several different forms of comparability and of the social welfare functions that are consistent with them.

Further Reading

Arrow's Impossibility Theorem was first demonstrated in:

Arrow, K. J. 1950. A difficulty in the concept of social welfare. *Journal of Political Economy* 58: 328–46.

The theorem is further elaborated in:

Arrow, K. J. 1951. *Social Choice and Individual Values*. New York: Wiley.

A comprehensive textbook treatment is given by:

Kelly, J. 1987. *Social Choice Theory: An Introduction*. Berlin: Springer-Verlag.

The concept of a social welfare function was first introduced by:

Bergson, A. 1938. A reformulation of certain aspects of welfare economics. *Quarterly Journal of Economics* 68: 233–52.

An analysis of limitations on the use of lump-sum taxation is contained in:

Mirrlees, J. A. 1986. The theory of optimal taxation. In K. J. Arrow and M. D. Intrilligator, eds., *Handbook of Mathematical Economics*, vol. 3. Amsterdam: North-Holland, 1197–1249.

An economic assessment of the UK poll tax is conducted in:

Besley, T., Preston, I., and Ridge, M. 1997. Fiscal anarchy in the UK: Modelling poll tax non-compliance. *Journal of Public Economics* 64: 137–52.

For a more complete theoretical treatment of the information constraint on redistribution see:

Guesnerie, R. 1995. *A Contribution to the Pure Theory of Taxation*. Cambridge, UK: Cambridge University Press.

Hammond, P. 1979. Straightforward incentive compatibility in large economies. *Review of Economic Studies* 46: 263–82.

Roberts, K. 1984. The theoretical limits to redistribution. *Review of Economic Studies* 51: 177–95.

Two excellent reviews of the central issues that arise with redistribution:

Boadway, R., and Keen, M. 2000. Redistribution. In A. B. Atkinson and F. Bourguignon, eds., *Handbook of Income Distribution*, vol. 1. Amsterdam: North Holland, 677–789.

Stiglitz, J. 1987. Pareto-efficient and optimal taxation and the new welfare economics. In A. Auerbach and M. Feldstein, eds., *Handbook of Public Economics*, vol. 2. Amsterdam: North Holland, 991–1042.

The self-selection argument for in-kind redistribution is in:

Akerlof, G. 1978. The economics of tagging as applied to the optimal income tax, welfare programs and manpower planning. *American Economic Review* 68: 8–19.

Besley, T., and Coate, S. 1991. Public provision of private goods and the redistribution of income. *American Economic Review* 81: 979–84.

Blackorby, C., and Donaldson, D. 1988. Cash versus kind, self-selection and efficient transfers. *American Economic Review* 78: 691–700.

Garfinkel, I. 1973. Is in-kind redistribution efficient? *Quarterly Journal of Economics* 87: 320–30.

Estimates of the incentive effects of welfare programs are in:

Eissa, N., and Liebman, J. 1996. Labour supply response to the earned income tax credit. *Quarterly Journal of Economics* 111: 605–37.

Gruber, J. 2000. Disability, insurance benefits and labor supply. *Journal of Political Economy* 108: 1162–83.

The government time-consistency problem is in:

Bruce, N., and Waldman, M. 1991. Transfers in kind: Why they can be efficient and nonpaternalistic. *American Economic Review* 81: 1345–51.

Hillier, B., and Malcomson, J. M. 1984. Dynamic inconsistency, rational expectations and optimal government policy. *Econometrica* 52: 1437–52.

Strotz, R. H. 1956. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 24: 165–80.

Comparability of utility is discussed in:

Ng, Y.-K. 2003. *Welfare Economics*. Basingstoke: Macmillan.

A discussion of the relation between social welfare functions and Arrow's theorem can be found in:

Samuelson, P. A. 1977. Reaffirming the existence of "reasonable" Bergson–Samuelson social welfare functions. *Economica* 44: 81–88.

Several of Sen's papers that discuss these issues are collected in:

Sen, A. K. 1982. *Choice, Welfare and Measurement*. Oxford: Basil Blackwell.

The application of Arrow's theorem to economic allocation problems is discussed in:

Le Breton, M. and Weymark, J.A. (2011) Arrovian social choice theory on economic domains. In K. Arrow, A. Sen, and K. Suzumura, eds., *Handbook of Social Choice and Welfare*, vol. 2. Amsterdam: North Holland, 191–299.

Exercises

13.1 Should a social planner be concerned with the distribution of income or the distribution of utility? How does the answer relate to needs and abilities?

13.2 Sketch the indifference curves of the Bergson–Samuelson social welfare function $W = U^1 + U^2$. What do these indifference curves imply about the degree of concern for equity of the social planner? Repeat for the welfare function $W = \min\{U^1, U^2\}$.

13.3 Show that an anonymous social welfare function must have indifference curves that are symmetric about the 45 degree line. Will an optimal allocation with an anonymous social welfare function and a symmetric utility possibility frontier always be equitable?

13.4 Assume that the preferences of the social planner are given by the function $W = \frac{[U^1]^\varepsilon}{\varepsilon} + \frac{[U^2]^\varepsilon}{\varepsilon}$. What effect does an increase in ε have on the curvature of a social indifference curve? Use this result to relate the value of ε to the planner's concern for equity.

13.5 There are H consumers who each have utility function $U^h = \log(M^h)$. If the social welfare function is given by $W = \sum U^h$, show that a fixed stock of income will be allocated equitably. Explain why this is so.

13.6 For a social welfare function $W = W(U^1(M^1), \ldots, U^H(M^H))$, where M^h is income, the "social marginal utility of income" is defined by $\frac{\partial W}{\partial U^h}\frac{\partial U^h}{\partial M^h}$. If $U^h = [M^h]^{1/2}$ for all h, show that the social marginal utility of income is decreasing in M^h for a utilitarian social welfare function. Use this to argue that a fixed stock of income will be distributed equally. Show that the argument extends to any anonymous and concave social welfare function when all consumers have the same utility function.

13.7 The two consumers that constitute an economy have utility functions $U^1 = x_1^1 x_2^1$ and $U^2 = x_1^2 x_2^2$.

a. Graph the indifference curves of the consumers, and show that at every Pareto-efficient allocation $\frac{x_1^1}{x_2^1} = \frac{x_1^2}{x_2^2}$.

b. Employ the feasibility conditions and the result in part a to show that Pareto-efficiency requires $\frac{x_1^2}{x_2^2} = \frac{\omega_1}{\omega_2}$, where ω_1 and ω_2 denote the endowments of the two goods.

c. Using the utility function of consumer 2, solve for x_1^2 and x_2^2 as functions of ω_1, ω_2, and U^2.

d. Using the utility function of consumer 1, express U^1 as a function of ω_1, ω_2, and U^2.

e. Assuming that $\omega_1 = 1$ and $\omega_2 = 1$, plot the utility possibility frontier.

f. Which allocation maximizes the social welfare function $W = U^1 + U^2$?

13.8 "Government intervention in markets is essential if we wish to achieve a fair allocation of resources." Is this correct?

13.9 Consider three individuals with utility indicators $U^A = M^A$, $U^B = \nu M^B$ and $U^C = \gamma M^C$.

a. Show that there are values of ν and γ that can generate any social ordering of the income allocations $a = (5, 2, 5)$, $b = (4, 6, 1)$, and $c = (3, 4, 8)$ when evaluated by the social welfare function $W = U^A + U^B + U^C$.

b. Assume instead that $U^A = \upsilon + \gamma M^A$, $U^B = \upsilon + \gamma M^B$ and $U^C = \upsilon + \gamma M^C$. Show that the evaluation via the utilitarian social welfare function is unaffected by the choices of υ and γ.

c. Now assume $U^h = [M^h]^\gamma$, where $h = A, B, C$. Show that the preferred outcome under the social welfare function $W = \min_{\{h\}}\{U^A, U^B, U^C\}$ is unaffected by choice of γ but that for the welfare function $W = U^A + U^B + U^C$ is affected.

d. Explain the answers to parts a through c in terms of the comparability of utility.

13.10 Provide an argument to establish that the optimal allocation must be Pareto-efficient. What assumptions have you placed on the social welfare function?

13.11 Consider an economy with two individuals (1 and 2). A, B, and C are three points that belong to the utility possibility frontier of the economy. The individual utilities (U^1, U^2) at the three points are as follows:

Points	U^1	U^2
A	4	21
B	8	17
C	14	6

Now consider point D.

Point	U^1	U^2
D	9	15

a. Does point D lie on the utility possibility frontier? How does the answer change if you know the utility possibility frontier is concave?

b. Are there any points on the utility possibility frontier that are Pareto-preferred to C? Justify your answer.

c. Which of the points A, B, C, or D lies on the highest indifference curve of a utilitarian social welfare function?

d. Which point is on the highest indifference curve of a Rawlsian social welfare function? Explain the answer using the solution to part a.

13.12 The most general form of a social welfare function *SWF* can be written as $W = W(U^1, \ldots, U^H)$.

a. Explain the following properties that a *SWF* may satisfy: nonpaternalism, Pareto principle, anonymity (the names of the agents do not matter), and concavity (aversion to inequality).

b. Consider two agents $h = 1, 2$ with utilities U^1 and U^2. Depict the social indifference curve of the utilitarian *SWF* in (U^1, U^2)-space. Which of the properties in part a does it satisfy?

c. Depict the social indifference curves of the maxi-min or Rawlsian *SWF*. Contrast to the utilitarian *SWF* with respect to the aversion to inequality. Which properties does the Rawlsian *SWF* satisfy?

d. The Bernoulli–Nash social welfare function is given by the product of individual utilities. Discuss the distributional properties of the Bernoulli–Nash *SWF*.

13.13 Consider the *SWF* of the form $W = \left[\sum_h [U^h]^\eta\right]^{1/\eta}$ with $-\infty < \eta \leq 1$. Show that this *SWF* reduces to the utilitarian *SWF* when $\eta = 1$, to the Bernoulli–Nash *SWF* when $\eta = 0$, and to the maxi-min Rawlsian *SWF* when $\eta \to -\infty$.

13.14 Are the following statements true or false? Provide examples to demonstrate your answer.

a. A Pareto improvement is always obtained when the economy moves from a point inside the utility possibility frontier to a point on the frontier.

b. A policy intervention will increase social welfare if and only if it is a Pareto improvement.

c. A policy intervention will increase social welfare for every Bergson–Samuelson social welfare function if and only if it is a Pareto-improvement.

13.15 A fixed amount $\overline{x}$ of a good has to be allocated between two individuals, $h = 1, 2$ with utility functions $U^h = \alpha^h x^h$ (with $\alpha^h > 0$), where x^h is the amount of the good allocated to consumer h.

a. How should $\overline{x}$ be allocated to maximize a utilitarian *SWF*? Illustrate the answer graphically. How do the optimal values of x^1 and x^2 change among the cases $\alpha^1 < \alpha^2$, $\alpha^1 = \alpha^2$, and $\alpha^1 > \alpha^2$?

b. What is the allocation maximizing the Bernoulli–Nash *SWF*? Illustrate graphically. How do the optimal values of x^1 and x^2 change with the preference parameters α^1 and α^2?

c. What is the allocation maximizing the maxi-min Rawlsian *SWF*? Illustrate graphically. How does the allocation change with preference parameters α^1 and α^2?

13.16 Show how the results of the previous exercise change if we assume a utility function of the form $U^h = \alpha^h \sqrt{x^h}$.

13.17 The 31 professors in an economics department have to vote on the location of a new coffee machine. The offices of the professors are located along one side of a corridor. Every professor would receive a utility of 20 if the coffee machine was placed outside his office. Utility is reduced by one unit for each office the professor has to pass to reach the machine.

a. Which location for the coffee machine wins a simple majority vote?

b. Which location would be chosen by a benevolent Rawlsian planner?

c. Which location would be chosen by a benevolent utilitarian planner?

d. How would the answer change if the coffee machine was an irritation, so each professor gained a unit of utility for each office he had to pass to reach the machine?

13.18 Consider a two-good exchange economy with two types of consumers. Type A have the utility function $U^A = 2\log(x_1^A) + \log(x_2^A)$ and an endowment of 3 units of good 1 and k units of good 2. Type B have the utility function $U^B = \log(x_1^B) + 2\log(x_2^B)$ and an endowment of 6 units of good 1 and $21 - k$ units of good 2.

a. Find the competitive equilibrium outcome and show that the equilibrium price $p^* = \frac{p_1}{p_2}$ of good 1 in terms of good 2 is $p^* = \frac{21+k}{15}$.

b. Find the income levels (M^A, M^B) of both types in equilibrium as a function of k.

c. Suppose that the government can make a lump-sum transfer of good 2, but it is impossible to transfer good 1. Use your answer to part b to describe the set of income distributions attainable through such transfers. Draw this in a diagram.

d. Suppose that the government can affect the initial distribution of resources by varying k. Find the optimal distribution of income if (i) the *SWF* is $W = \log(M^A) + \log(M^B)$ and (ii) $W = M^A + M^B$.

13.19 Are the following true or false? Explain your answer.

a. Cardinal utilities are always interpersonally comparable.

b. A Rawlsian social welfare function can be consistent with ordinal utility.

c. The optimal allocation with a utilitarian social welfare function is always inequitable.

13.20 The purpose of this exercise is to illustrate the potential conflict between personal liberty and the Pareto principle (first studied by Sen). Assume there is a copy of Lady Chatterley's Lover available to be read by two persons, A and B. There are three possible options: (a) A reads the book and B does not; (b) B reads the book and A does not; (c) neither reads the book. The preference ordering of A (the prude) is $c \succ_A a \succ_A b$ and the preference ordering of B (the lascivious) is $a \succ_B b \succ_B c$. Hence c is the worst option for one and the best option for the other; while both prefer a to b. Define the personal liberty rule as allowing everyone to choose freely on personal matters (like the color of one's own hair) with society as a whole accepting the choice, no matter what others think.

a. Apply the personal liberty rule to the example to derive social preferences $b \succ c$ and $c \succ a$.

b. Show that by the Pareto principle we must have a social preference cycle $a \succ b \succ c \succ$.

c. Suppose that liberalism is constrained by the requirement that the prude A decides to respect B's preferences such that A's preference for c over b is ignored. Similarly for B, only his preference for b over c is relevant but not his preference for a over c. What are the modified preference orderings of each person? Show that it leads to acyclic (transitive) social preference.

d. The second possibility to solve the paradox is to suppose that each is willing to respect the other's choice. Thus A respects B's preference for b over c and B respects A's preference for c over a. What are the modified preference orderings of each person? Show that it leads to acyclic social preference. What is then the best social outcome?