ECONOMICS of the PUBLIC SECTOR

THIRD EDITION

JOSEPH E. STIGLITZ

W.W. NORTON & COMPANY
NEW YORK / LONDON

5 Efficiency and Equity

FOCUS QUESTIONS

1 How do economists think systematically about how to make social choices when there are trade-offs, that is, when after finding all possible Pareto improvements, gains to the welfare of one individual must come at the expense of the welfare of others? What is the social welfare function, and why do economists find this concept useful?

2 How do economists think systematically about the trade-offs between efficiency and inequality? How do they measure poverty or inequality? How do they measure efficiency?

3 As a practical matter, how do governments translate these general principles into a form which can actually be used in decision making?

Chapter 3 took up Pareto efficiency, the condition in which no one can be made better off without making someone else worse off. It showed that in the absence of market failures, a free market would be Pareto efficient. But even if the competitive economy is efficient, the distribution of income to which it gives rise may be viewed as undesirable. One of the main consequences, and main objectives, of government activity is to alter the distribution of income.

The evaluation of a public program often entails balancing its consequences for economic efficiency and for the distribution of income. A cen-

tral objective of welfare economics is to provide a framework within which these evaluations can be performed systematically. This chapter shows how economists conceptualize the trade-offs between efficiency and equity.

EFFICIENCY AND DISTRIBUTION TRADE-OFFS

Consider again a simple economy with two individuals, Robinson Crusoe and Friday. Assume initially that Robinson Crusoe has ten oranges, while Friday has only two. This seems inequitable. Assume, therefore, that we play the role of government and attempt to transfer four oranges from Robinson Crusoe to Friday, but in the process one orange gets lost. Hence Robinson Crusoe ends up with six oranges, and Friday with five. We have eliminated most of the inequity, but in the process the total number of oranges available has been diminished. There is a **trade-off** between efficiency—the total number of oranges available—and equity.

The trade-off between equity and efficiency is at the heart of many discussions of public policy. Two questions are debated. First, there is disagreement about the nature of the trade-off. To reduce inequality, how much efficiency do we have to give up? Will one orange or two be lost in the process of transferring oranges from Crusoe to Friday?

Second, there is disagreement on the relative value to be assigned to a decrease in inequality compared to a decrease in efficiency. Some people claim that inequality is the central problem of society, and society should simply minimize the extent of inequality, regardless of the consequences to efficiency. Others claim that efficiency is the central issue. They argue that even if one wishes to help the poor, in the long run, the best way to do that is not to worry about how the pie is to be divided but to increase the size of the pie, to make it grow as rapidly as possible, so that there are more goods for everyone. These disagreements relate to social choices between equity and efficiency. We now take a closer look at these choices.

ANALYZING SOCIAL CHOICES

When economists analyze consumer choice, the opportunity set is defined by the consumer's budget constraint, and the consumer's preferences are described by indifference curves (see Chapter 3, pp. 94–96). The individual chooses the point on the budget constraint which is tangent to an indifference curve—this puts him on the highest indifference curve feasible, given the budget constraint.

Economists have tried to use the same framework for analyzing social choices. The utility possibilities curve, introduced in Chapter 3, describes the opportunity set. It gives the highest level of utility (or welfare) attainable by one individual, given the levels of utility attained by others. An economy is Pareto efficient if and only if it is operating along the utility possibilities schedule. The first fundamental theorem of welfare economics says that competitive economies are always on the utility possibilities schedule. The

second fundamental theorem of welfare economics says that every point on the utility possibilities schedule can be attained through a competitive market process if the government redistributes initial endowments accordingly. How does society select a point along the utility possibilities curve? Just as indifference curves for individuals describe how they make trade-offs between different goods, **social indifference curves** describe how society might make trade-offs between utility levels of different individuals. A social indifference curve gives those combinations of utility of, say, Crusoe and Friday, between which society is indifferent.

The two central questions of welfare economics can now be restated in terms of this social choice framework. Assume the current competitive market equilibrium is represented by the point A on the utility possibilities schedule depicted in Figure 5.1. Suppose society decides to move, say, from point A to point B along the utility possibilities schedule, representing an increase in Friday's utility and a reduction in Crusoe's utility. The first question is: What is the trade-off? The utility possibilities curve gives the answer by showing the increase in Friday's utility from U^F_0 to U^F_1 and the decrease

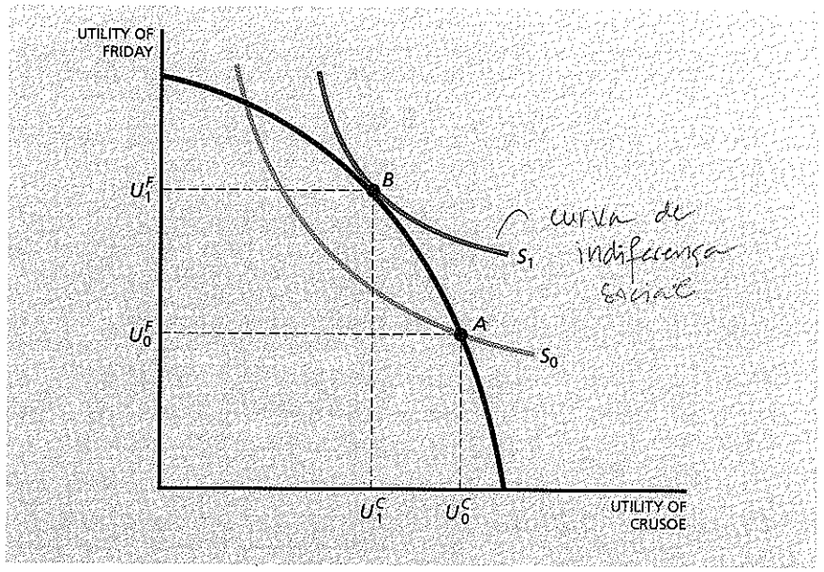

FIGURE 5.1 **Social Indifference Curves** The social indifference curves describe how society evaluates trade-offs between Friday and Crusoe; it gives the combinations of utilities between which society is indifferent. Society is better off on a higher social indifference curve, just as an individual is better off on a higher individual indifference curve. And just as the individual chooses the point on the budget constraint at which the indifference curve is tangent to the budget constraint, society's preferred point on the utility possibilities curve is the point at which the social indifference curve is tangent to the utility possibilities curve.

in Crusoe's utility from U^C_0 to U^C_1. The second question concerns social preferences: How does society evaluate the trade-off? The slope of the social indifference curves gives the trade-offs for which society is indifferent. Point B is on the social indifference curve S_1, which is tangent to the utility possibilities curve, and lies on a higher indifference curve than S_0. Point B is therefore preferred by society.

The next two sections take a closer look at each of these questions regarding trade-offs and the economist's framework for analyzing social choice.

DETERMINING THE TRADE-OFFS

As we saw in Figure 5.1, the utility possibilities schedule shows us the trade-offs of transferring utility from Crusoe to Friday. The shape of the utility possibilities schedule tells us something more about those trade-offs. Consider the utility possibilities schedule shown in Figure 5.2. Assume that the economy lies at point A, where Crusoe enjoys much more utility than Friday. Moving up and to the left along the schedule increases Friday's utility and decreases Crusoe's. Suppose we transfer oranges from Crusoe to Friday by moving in two steps, from point A to B to C. Clearly, this makes Crusoe

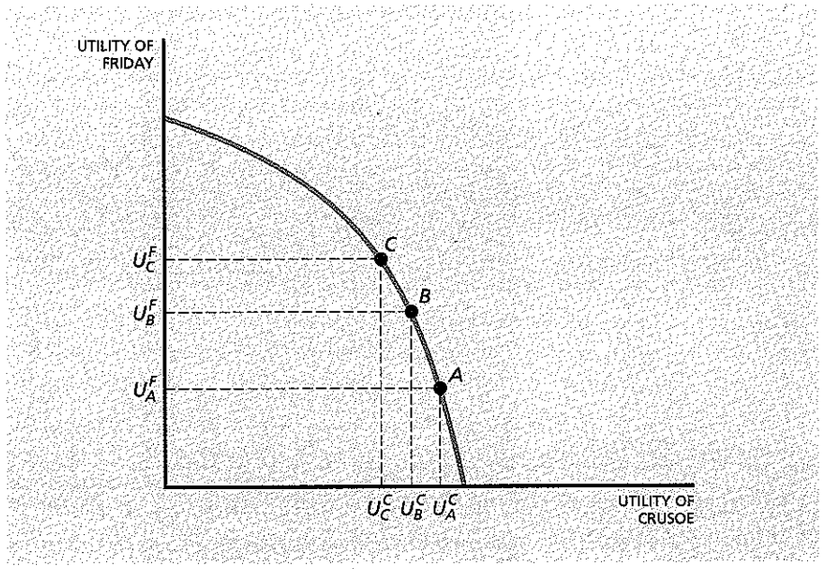

utilidade marginal decrescente

FIGURE 5.2 **Crusoe's and Friday's Utility Possibilities Curve** As oranges are transferred from Crusoe to Friday, Crusoe's utility is decreased and Friday's increased. In moving from point A to B the gain in Friday's utility appears much greater than the loss in Crusoe's utility. That is because Friday is so much worse off than Crusoe. In moving from B to C, the gain in Friday's utility is still larger than the loss in Crusoe's utility, but the trade-off has changed so that Friday's gain is smaller than the gain from A to B.

worse off. As depicted in the figure, the decreases in Crusoe's utility are small in comparison to the increases in Friday's utility.

Utility theory helps explain this outcome. Economists use the term **utility function** to describe the relationship between the number of oranges and Friday's level of utility; the extra utility Friday gets from an extra orange is called his **marginal utility**. These are shown in panels A and B of Figure 5.3. At each point, marginal utility is the slope of the utility function—the

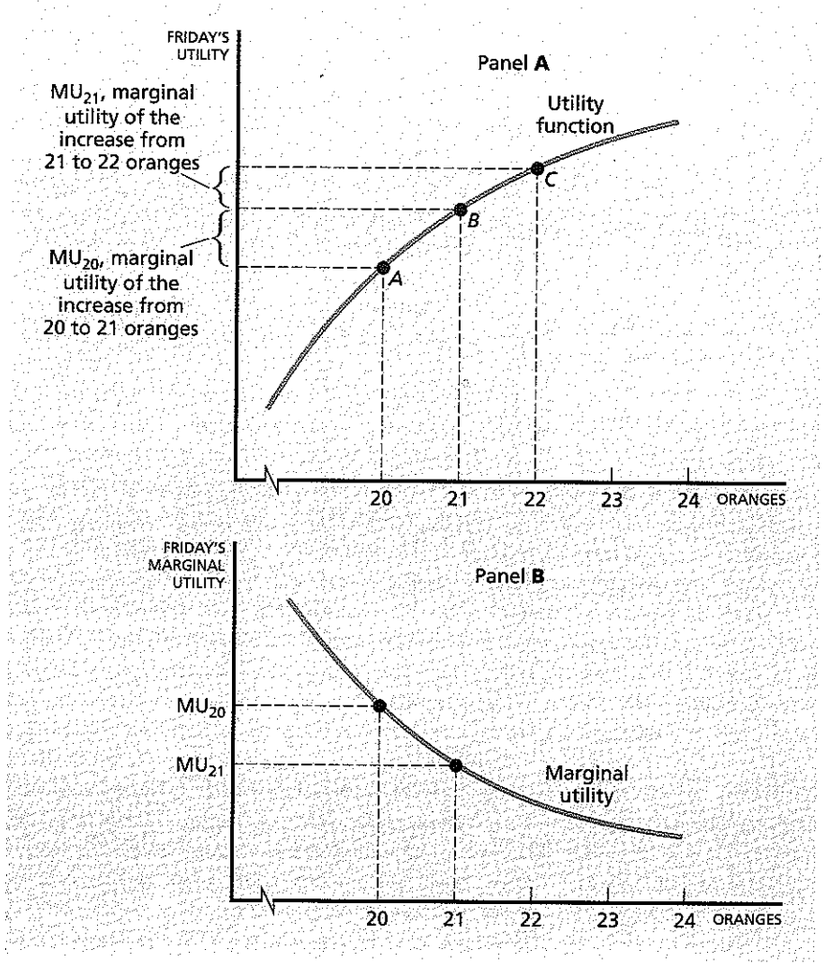

FIGURE 5.3 **The Utility Function and Marginal Utility** Panel A shows the utility function. As we give Friday more oranges, his utility increases, but each additional orange leaves him less extra utility. Panel B shows marginal utility: the extra utility Friday gets from an extra orange decreases as the number of oranges increases, corresponding to the decreasing slope of the utility function.

change in utility from a unit change in orange consumption. Notice that as more oranges are consumed, utility rises more slowly, and marginal utility falls. (Thus, the slope of the utility function at point *C* is less than the slope at *A* or *B*.) This is because Friday enjoys the first orange very much, the next one a little less, and additional oranges still less. Finally, he becomes satiated and derives very little additional enjoyment from an additional orange. As an individual consumes more of any good, the extra gain from having one extra unit of that good becomes smaller. This phenomenon is referred to as **diminishing marginal utility.**[1]

By the same token, as we take away oranges from Crusoe, his utility decreases; and as we take away more and more oranges, the extra utility he loses from each additional loss of an orange increases. That is why with diminishing marginal utility, the utility possibilities schedule has the shape depicted in Figures 5.1 and 5.2. This shape says that when Friday has very little income (few oranges), we can increase his utility a great deal with a small decrease in Crusoe's utility, but when Friday is much better off, we can increase his utility only a little with even a large decrease in Crusoe's utility.

There is a second important determinant of the shape of the utility possibilities schedule—the efficiency with which we can transfer resources from one individual to another. In our society, the way we transfer resources from one group (say, the rich) to another (say, the poor) is by taxing the rich and subsidizing the poor. The way we do that normally interferes with economic efficiency. The rich may work less hard than they would otherwise, because they reap only a fraction of the returns to their effort; while the poor may work less hard because by working harder, they may lose eligibility for benefits. The magnitude of these disincentives—a subject of considerable controversy—affects the entire shape of the utility possibilities schedule. In Figure 5.4, the blue line represents the utility possibilities schedule assuming that it is costless to transfer resources. The black line, lying far below the previous locus, except at the point *C*—the point which occurs without any redistribution—represents the schedule when transfers are very costly.

EVALUATING THE TRADE-OFFS

The second basic concept used in analyzing social choices is the social indifference curve. As described in Chapter 3, an indifference curve gives those combinations of goods which give the individual the same level of utility. Just as individuals derive utility from the goods they consume, we can think of society as deriving its welfare from the utility received by its members. The **social welfare function** gives the level of social welfare corresponding to a particular set of levels of utility attained by members of society. The *social indifference curve* is defined as the set of combinations of utility of different individuals (or groups of individuals) that yields equal levels of welfare

[1] We write the utility function as $U = U(C_1, C_2, \ldots, C_n)$, where $C_1, C_2, \ldots, C_n$ represent the quantities of consumption of the various goods. Marginal utility of, say, C_1, is then simply the increase in U (utility) from an increase in consumption of C_1. Diminishing marginal utility implies that successive increments in C_1 yield successively smaller increments to U.

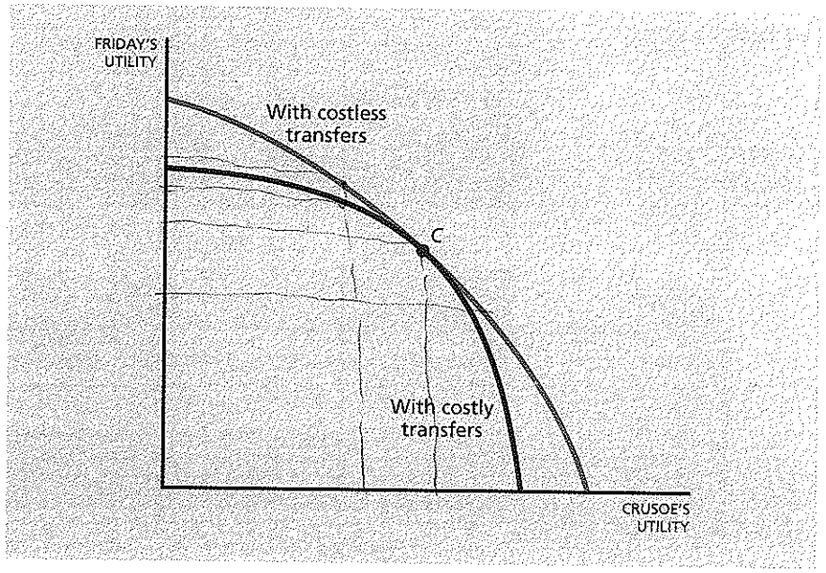

FIGURE 5.4 **Utility Possibilities Schedule with Costly Transfers** The set of points we can achieve through redistribution, when transfers are costly, lies within the utility possibilities curve, given costless transfers.

to society—for which, in other words, the social welfare function has the same value.

The social welfare function provides a basis for ranking any allocation of resources: we choose those allocations which yield higher levels of social welfare. The **Pareto principle** says that we should prefer those allocations in which at least some individuals are better off and no one is worse off. It says that if some individuals' utility is increased and no one else's utility is decreased, social welfare increases. Thus, in Figure 5.5 those combinations to the northeast of *A* make everyone better off, and hence satisfy the Pareto principle.

Unfortunately, most choices involve trade-offs, with some individuals being made better off and others worse off. At point *B* the second group is better off than at *A*, but the first group is worse off. We thus need a stronger criterion, and this is what the social welfare function provides. The social indifference curves provide a convenient diagrammatic way of thinking about the kinds of trade-offs society faces in these situations. Thus, in Figure 5.5 all combinations of the utilities of Groups 1 and 2 that are on the social indifference curve labeled W_2 yield a higher level of social welfare than those combinations on the curve labeled W_1. This shows that *B* is preferred to *A*.

Social welfare functions can be thought of as a tool economists use to summarize assumptions about society's attitudes toward different distributions of income and welfare. If society is very concerned about inequality, it might not care that Crusoe has to give up seventy oranges for Friday to get

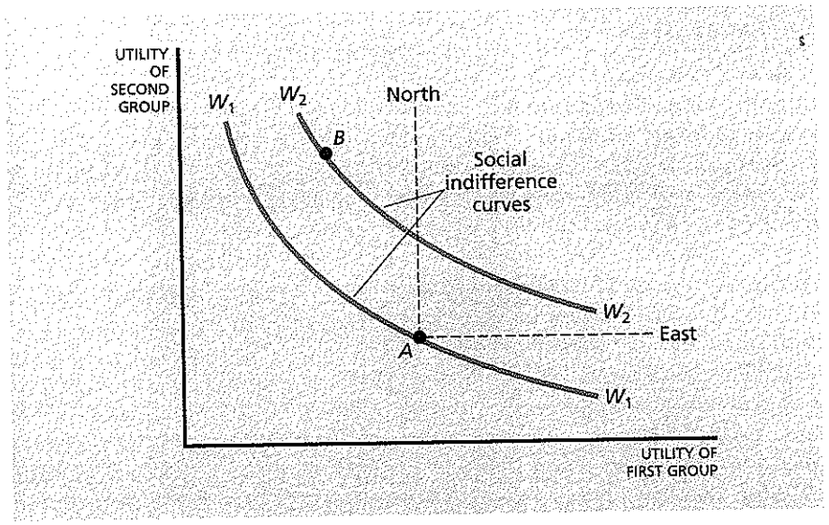

FIGURE 5.5 **Social Indifference Curves** Society is willing to trade off some decrease in one group's utility for an increase in another's. A social indifference curve gives those combinations of utilities of Group 1 and Group 2 between which society is indifferent. Points on the social indifference curve labeled W_2 yield a higher level of social welfare than do points on the social indifference curve labeled W_1.

one orange, "since Crusoe has so many to begin with." So long as Friday is poorer than Crusoe, any sacrifice on Crusoe's part that makes Friday better off would be justified. On the other hand, society might not care at all about inequality; it could value an orange in the hands of Friday exactly the same as an orange in the hands of Crusoe, even though Friday is much poorer. In that case it would focus only on efficiency, on the number of oranges available. No redistribution of oranges from Crusoe to Friday would be justified if, in the process, a single orange was lost.

UTILITARIANISM Social welfare functions—and the associated social indifference curves—can take a variety of shapes, as illustrated in Figure 5.6. Panels A and B illustrate two different cases. In panel A the social indifference curve is a straight line, implying that no matter what the level of utility of Friday and Crusoe, society is willing to trade off one "unit" of Friday's utility against one unit of Crusoe's. The view represented by this social indifference curve has a long historical tradition. Jeremy Bentham was the leader of a group, called **utilitarians,** which argued that society should maximize the sum of the utilities of its members; in our simple example with two individuals, the social welfare function is

$$W = U_1 + U_2.$$

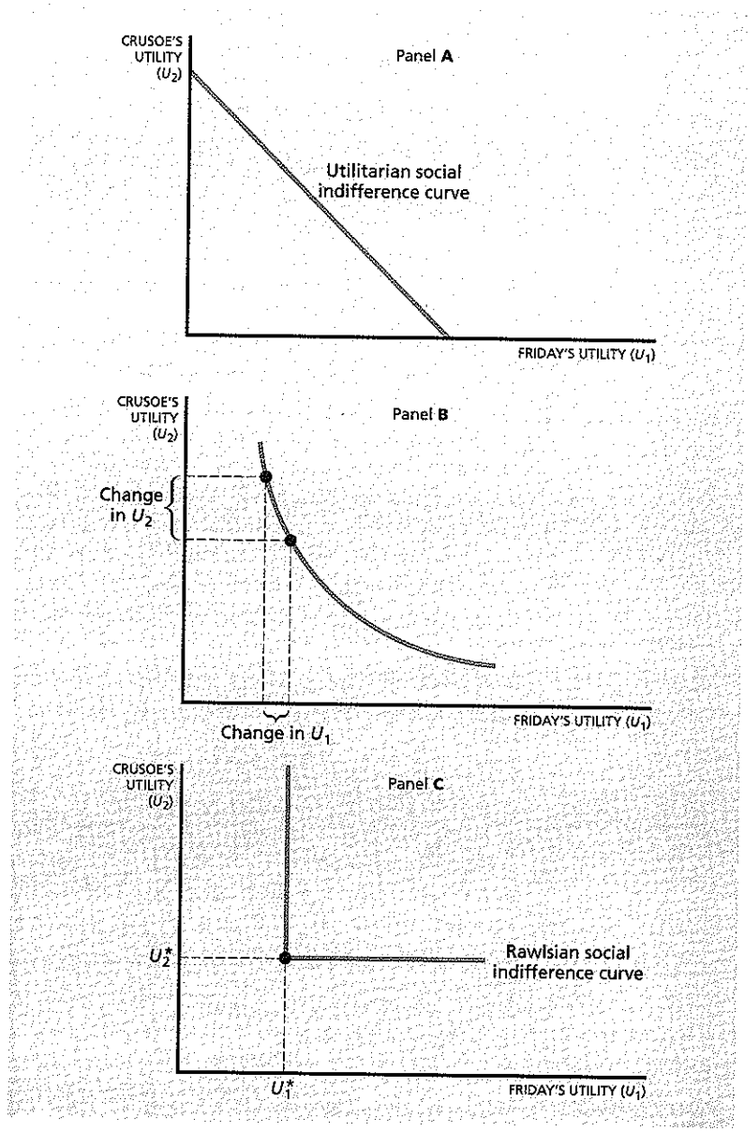

FIGURE 5.6 **Alternative Shapes of Social Indifference Curves** (A) A utilitarian is willing to give up some utility for Crusoe so long as Friday gains at least an equal amount of utility. The social indifference curves are straight lines. (B) Some argue that society requires more than an equal increase in the utility (U_2) of a rich individual to compensate for a decrease in the utility (U_1) of a poor individual. (C) Rawls maintains that no amount of increase in the welfare of the rich can compensate for a decrease in the welfare of the poor. This implies that the social indifference curves are L-shaped.

It is clear with this social welfare function, the social indifference curve has the shape depicted in panel A.

It is important to emphasize that with a utilitarian social welfare function, society is not indifferent to an increase of one orange (or one dollar of income) for Individual 1 and a decrease of one orange (or one dollar of income) for Individual 2. If Individual 1 has a lower level of income (fewer oranges) than Individual 2, then the increase in utility of Individual 1 from one more orange (one more dollar) will be greater than the decrease in utility for Individual 2. What the utilitarian social welfare function says is that the *utility* of any individual should be weighted equally to the utility of any other individual.

Many would argue that when one individual is worse off than another, society is not indifferent to a decrease in the utility of the poorer (Individual 1) matched by an equal increase in the utility of the richer (Individual 2). Society should be willing to accept a decrease in the utility of the poor only if there is a much larger increase in the utility of the rich. The social indifference curve reflecting these values is drawn in Figure 5.6B, where it appears not as a straight line but as a curved one; as the poorer individual becomes worse and worse off, the increment in utility of the richer individual that makes society indifferent must be larger and larger (i.e., the slope of the social indifference curve becomes steeper and steeper).

RAWLSIANISM The extreme position of this debate has been taken by John Rawls, a professor of philosophy at Harvard University. Rawls argues that the welfare of society only depends on the welfare of the worst-off individual; society is better off if you improve his welfare but gains nothing from improving the welfare of others. There is, in his view, no trade-off. If Friday is worse off than Crusoe, then anything that increases Friday's welfare increases social welfare. As oranges are transferred from Crusoe to Friday, it makes no difference how many are lost in the process (how inefficient the transfer process is), so long as Friday gets *something*. To put it another way, no amount of increase in the welfare of the better-off individual could compensate society for a decrease in the welfare of the worst-off individual. Diagrammatically, this is represented by an L-shaped social indifference curve, as in Figure 5.6C[2].

TWO CAVEATS

While many public sector economists have made extensive use of the concepts of social welfare functions and the utility possibilities curve, these concepts have also been extensively criticized, on several grounds.

INTERPERSONAL COMPARISONS We assume that when an individual consumes more, her utility rises. But we cannot measure the level of utility or the

[2] The social welfare function is written:

$W = \min \{U_1, \ldots U_n\}$.

Social welfare reflects only the utility of the worst-off member of society.

SOCIAL CHOICE IN THEORY

1 *Construct the opportunity set.* The utility possibilities schedule describes how much one person's utility has to be decreased when another's is increased.

2 *Define preferences.* Social indifference curves describe how much society is willing to decrease one person's utility to increase another's by a given amount.

3 *Adopt programs that increase social welfare.* Find the programs that put society on the highest social indifference curve.

change in utility. Social welfare functions seem to assume not only that there is a meaningful way of measuring an individual's utility,[3] but that there is a meaningful way of comparing the utility of different individuals. For example, with the utilitarian social welfare function we add up the utility of the different members of society. Because we add Crusoe's and Friday's utility together, we are, in effect, assuming that somehow we can compare *in a meaningful numerical way* their level of utility. But when we transfer an orange from Robinson to Friday, how can we compare in an objective way the value of Friday's gain and Robinson's loss?

The same problem arises with a Rawlsian social welfare function, where we are told to maximize the welfare of the worst-off member of society. To judge who is worst off, we must somehow compare utilities.

Many economists believe these **interpersonal comparisons** cannot be made in any meaningful way. I may claim that although I have a much higher income than my brother, I am unhappier; not only that, I may claim that I know how to spend income so much better that the extra increment in my utility from a dollar given to me is much greater than the extra increment in utility that he would get from receiving an extra dollar. How could anyone prove that I was wrong (or right)? Because there is no way of answering this question, economists argue that there can be no scientific basis for making welfare comparisons.

Since there is no "scientific" basis for making such welfare comparisons, many economists believe economists should limit themselves to describing the consequences of different policies, pointing out who are the gainers and who are the losers, and that should be the end of their analysis. They believe the only circumstances in which economists should make welfare

[3] In some situations, it may be possible to use the amount of money an individual would be willing to pay for an object as a *measure* of the utility of that object. But this does not resolve the problem of comparing utilities across individuals.

COMPARING INDIVIDUAL AND SOCIAL CHOICES

INDIVIDUAL CHOICES	SOCIAL CHOICES
STEP ONE: DEFINE OPPORTUNITY SET	
Budget constraint	Utility possibilities curve
STEP TWO: DEFINE PREFERENCES	
Individual indifference curve	Social indifference curve
STEP THREE: CHOOSE PREFERRED POINT	
Tangency between individual indifference curve and budget constraint	Tangency between social indifference curve and utility possibilities curve

judgments is when the policy change is a Pareto improvement. Unfortunately, as we have said, few policy changes are Pareto improvements, and hence without making interpersonal comparisons of welfare, economists have little to say regarding policy.

WHENCE SOCIAL WELFARE FUNCTIONS? The second set of objections concerns the very nature of social welfare functions. Individuals have preferences; they can decide whether they prefer some combination of apples and oranges to another combination. Society consists of many individuals; but society itself does not have preferences. We can describe the preferences of each individual, but whose preferences does the social welfare function represent? If there were a dictator, the answer to that question would be easy: the social welfare function would reflect the preferences of the dictator. But in a democratic society, there is no easy answer to the question. Some individuals—particularly the rich—may care little for redistribution, while others—particularly the poor—may argue that greater weight should be placed on redistribution.

As a descriptive matter—as part of a positive analysis—societies seldom exhibit consistency. One of the results to be described in Chapter 7 explains why this is not unexpected. Most economists think of the concepts we have described—as part of a normative analysis—as tools that help us think systematically about the trade-offs society constantly must face. And, as we noted earlier, the systematic analysis of these trade-offs actually constitutes an important part of the process by which decisions get made.

SOCIAL CHOICES IN PRACTICE

In practice, government officials do not derive utility possibilities schedules, nor do they write down social welfare functions. But their approach to deciding whether, say, to undertake any particular project does reflect the concepts we have introduced.

First, they attempt to identify and measure the net benefits (benefits minus costs) received by different groups. Second, they ascertain whether the project is a Pareto improvement, that is, whether everyone is better off. If so, clearly the project should be undertaken (this is the Pareto principle).

If the project is not a Pareto improvement, matters are more difficult. Some gain, some lose. The government needs to make an overall judgment. One commonly used approach looks at two summary statistics, describing "efficiency" and "equity" effects. *Efficiency* is measured by simply summing the gains or losses for each individual (which are calculated in a manner to be described shortly). *Equity* is measured by looking at some overall measure of inequality in society. If a project has net positive gains (positive efficiency effects) and reduces measured inequality, it should be undertaken. If a project has net positive losses and increases measured inequality, it should not be undertaken. If the efficiency measure shows gains but the equality measure shows losses (or vice versa), there is a trade-off, which is evaluated using a social welfare function: how much extra inequality is society willing to accept for an increase in efficiency?

There are numerous examples where choices between equality and efficiency have to made. For instance, in general, the more a tax system redistributes income, the greater the inefficiencies it introduces. There is a trade-off between equality and efficiency. There are, of course, important instances of poorly designed tax systems; such tax systems put the economy below its utility possibilities schedule. In such cases, it may be possible to increase both equality and efficiency.

We now take a closer look at how economists measure efficiency and inequality.

MEASURING BENEFITS

The first problem is how to measure the benefits of some program or project to particular individuals. In the discussion of utility theory above, we described how giving Friday more oranges increased his utility. But how do we measure this?

The standard way this is done is in terms of *willingness to pay*. We ask how much an individual would be willing to pay to be in one situation rather than another. For example, if Joe likes chocolate ice cream more than vanilla, it stands to reason that he would be willing to pay more for a scoop of chocolate ice cream than for a scoop of vanilla. Or if Diane would rather live in California than in New Jersey, it stands to reason that she would be willing to pay more for the West Coast location.

Notice that how much a person is willing to pay is different from how much he *has* to pay. Just because Joe is willing to pay more for chocolate ice cream than for vanilla does not mean he will have to pay more. What he has to pay depends on market prices; what he is willing to pay reflects his preferences.

Using willingness to pay as our measure of utility, we can construct a diagram like panel A of Figure 5.7, which shows the level of utility Mary receives from sweatshirts as the number of sweatshirts she buys increases. This

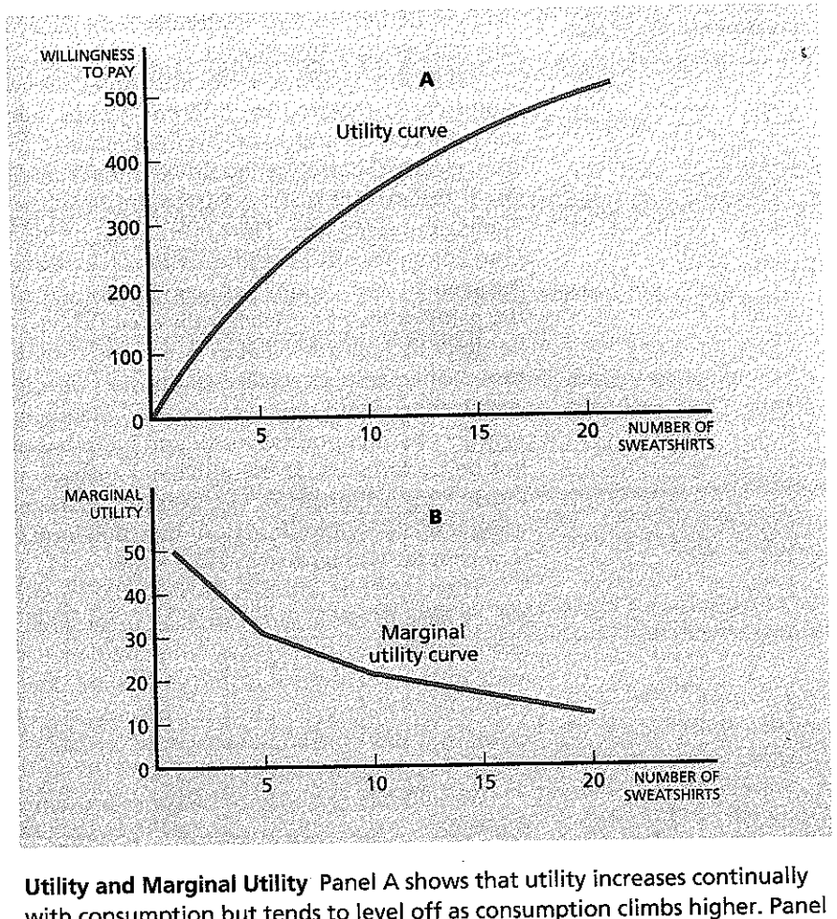

FIGURE 5.7 **Utility and Marginal Utility** Panel A shows that utility increases continually with consumption but tends to level off as consumption climbs higher. Panel B explicitly shows marginal utility; notice that it declines as consumption increases.

information is also given in Table 5.1. Here we assume that Mary is willing to pay $200 for five sweatshirts, $228 for six sweatshirts, $254 for seven sweatshirts, and so on. Thus, five sweatshirts give her a utility of 200, six a utility of 228, and seven sweatshirts a utility of 254. Mary's willingness to pay increases with the number of sweatshirts, reflecting the fact that additional sweatshirts give her additional utility. The extra utility of an additional sweatshirt measured here by the additional amount she is willing to pay, is the marginal utility. The numbers in the third column of Table 5.1 give the marginal (or extra) utility she received from her last sweatshirt. When Mary owns five sweatshirts, an additional sweatshirt yields her an additional or marginal utility of 28 (228 − 200); when she owns six sweatshirts, an additional one

TABLE 5.1 **Utility and Marginal Utility**

NUMBER OF SWEATSHIRTS	MARY'S WILLINGNESS TO PAY (UTILITY)	MARGINAL UTILITY
0	0	
1	50	50
2	95	45
3	135	40
4	170	35
5	200	30
6	228	28
7	254	26
8	278	24
9	301	23
10	323	22
11	344	21
12	364	20
13	383	19
14	401	18
15	418	17
16	434	16
17	449	15
18	463	14
19	476	13
20	488	12

gives her a marginal utility of only 26 (254 − 228). Panel B traces the marginal utilities of each of these increments.[4]

ORDINARY AND COMPENSATED DEMAND CURVES

We can use the concept of willingness to pay to construct a demand curve. We have already asked how much Mary is willing to pay for each additional sweatshirt. If the price of sweatshirts is $29, then she will buy five sweatshirts. She would have been willing to pay $30 for the fifth sweatshirt, so clearly, the marginal benefit of the fifth sweatshirt exceeds its cost; but she is only willing to pay $28 for the sixth sweatshirt, so the marginal benefit is less than the cost. Thus, the marginal utility curve drawn in panel B in Figure 5.7 can also be thought of as the demand curve.

But it is a special demand curve, called the **compensated demand curve,** which differs slightly from the **ordinary demand curve.** Recall that we constructed the compensated demand curve by asking how much Mary would

[4] Since marginal utility is the extra utility from an extra unit of consumption, it is measured by the slope of the utility curve in panel A.

be willing to pay for each additional sweatshirt; thus, as we give her more sweatshirts, we are always keeping her at exactly the same level of utility.

To construct the ordinary demand curve we need to know how many units of the commodity Mary would buy at each price. As the price is lowered, Mary not only demands more, but is made better off. As prices are lowered, individuals substitute the cheaper good for others goods. If the price of sweatshirts is lowered, Mary will substitute sweatshirts for sweaters. This is called the **substitution effect.** Because of the lower price, Mary is better off; if she bought exactly the same amount of goods that she did before, she would have money left over. She spreads this money around. Some of it is spent on buying sweatshirts. The increase in demand for sweatshirts as a result of the fact that Mary is better off—it is *as if* she had more income—is called the **income effect.** If we take away this extra money, we have the compensated demand curve; we eliminate the income effect. Thus, the compensated demand curve reflects only the substitution effect. In most cases, the differences between the two are negligible. If Mary spends one-tenth of 1 percent of her income on sweatshirts, taking away the extra income has almost no effect on her demand for sweatshirts, or any other commodity. Thus, Figure 5.8 shows the ordinary and compensated demand curves as being almost the same, with the ordinary demand curve being slightly flatter (lowering the price from its current level, P_0, results in a slightly greater increase in the quantity demanded, and raising the price from its current level results in a slightly greater decrease in quantity demanded).

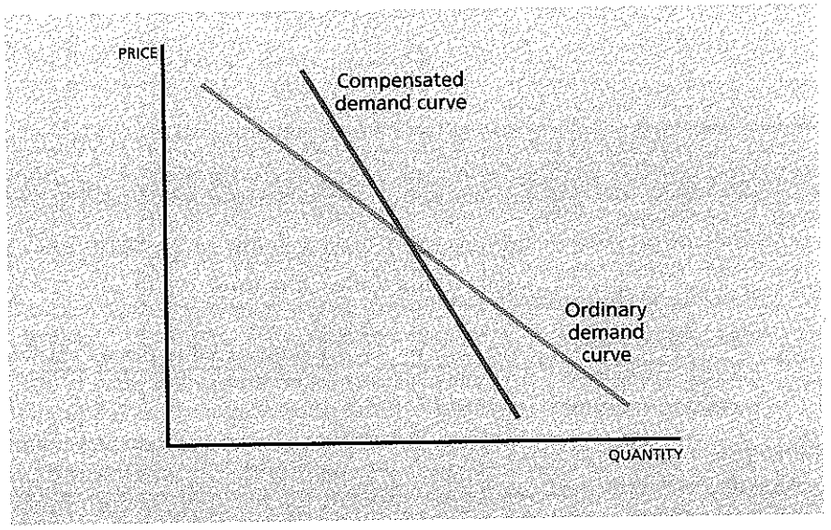

FIGURE 5.8 **Compensated versus Uncompensated Demand Curves** The compensated demand curve gives the demand for a good assuming, as price is changed, that money is taken away or given to the individual to leave him just as well off as he was before the price change. It thus measures only the substitution effect associated with the price changes. Because as price is lowered individuals are better off, and as a result buy slightly more of (normal) commodities, the ordinary demand curve is slightly flatter than the compensated demand curve.

CONSUMER SURPLUS

The difference between what an individual is willing to pay and what he has to pay is called his **consumer surplus.** Mary would have been willing to pay \$50 for the first sweatshirt, \$45 for the second, \$40 for the third, and so on. But if the market price is \$29, that is all she has to pay for each sweatshirt. Thus, on the first sweatshirt she gets a surplus of \$21 (\$50, what she was willing to pay, minus \$29, what she actually pays); on the second sweatshirt she gets a surplus of \$16; on the third sweatshirt she gets a surplus of \$11, and so on. The total consumer surplus is thus the sum: \$21 + \$16 + \$11 + \$6 + \$1 = \$55.

Diagrammatically, the consumer surplus is depicted in Figure 5.9 as the shaded area under the compensated demand curve and above the price line. Of course, since the compensated and uncompensated demand curves are almost the same, typically, we calculate the consumer surplus simply by looking at the area under the ordinary demand curve above the price line.

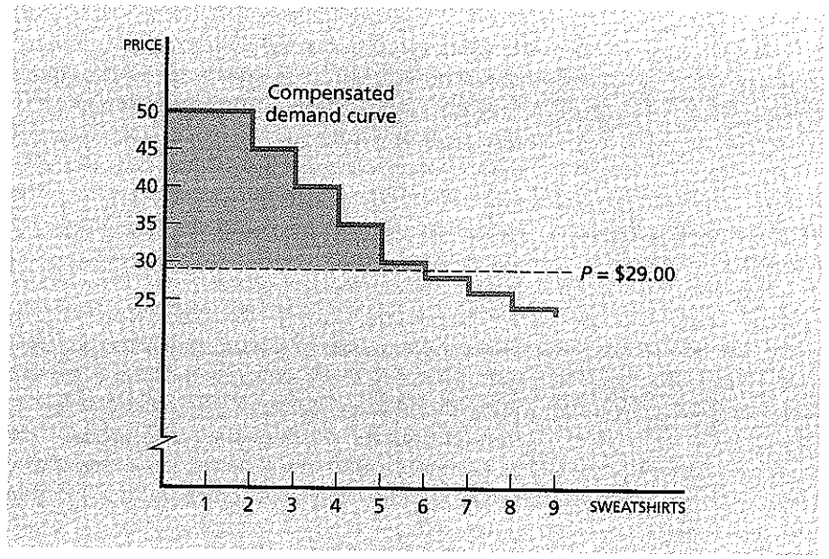

FIGURE 5.9 **Graphical Representation of Consumer Surplus** An individual's surplus is the difference between what he is willing to pay (represented by the area beneath the demand curve) and what he actually pays (the area under the price line). The consumer surplus here is indicated by the shaded region.

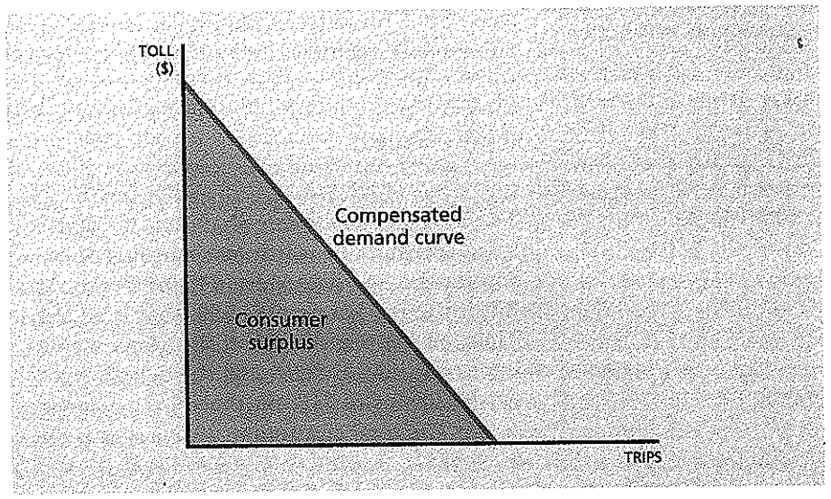

FIGURE 5.10 **Measuring the Benefits of a Government Project: Building a Bridge** The benefits of a bridge for which no tolls will be charged can be measured by the total area under the demand curve—the total consumer surplus.

USING CONSUMER SURPLUS TO CALCULATE THE BENEFITS OF A GOVERNMENT PROJECT The compensated demand curve can be useful for measuring the benefits of government projects. For instance, constructing a bridge upon which no toll will be charged can be thought of as lowering the price from "infinity" (one simply can't buy trips across a nonexistent bridge) to zero. The welfare gain is just the total consumer surplus, the area under the demand curve in Figure 5.10. This measures the maximum individuals could pay and still be as well off with the bridge as they were without it. Clearly, if the consumer surplus is less than the cost of the bridge, it does not pay to construct it, while if the consumer surplus is greater than the cost of the bridge, it does pay to build it.

There are several ways that economists go about trying to measure consumer surplus and willingness to pay. For many goods, there are data with which economists can construct the demand curve (the quantity that individuals are willing to purchase at each price) and the compensated demand curve.[5] In that case, willingness to pay can be calculated simply as the area under the compensated demand curve. For some goods, such as the Grand Canyon, there is no market demand curve. Yet the government still might want to know how much citizens are willing to pay to preserve it in its pristine condition. Economists have designed elaborate survey techniques to elicit meaningful answers from individuals concerning their willingness to pay. These methods are discussed at greater length in Chapter 11.

[5] As was noted previously, for most goods the compensated and uncompensated demand curves are very similar. If the *income elasticity* (the percentage increase in the demand for the good when income increases by 1 percent) is known, one can calculate the compensated demand curve from the uncompensated demand curve.

MEASURING AGGREGATE SOCIAL BENEFITS

We have now described how we can measure the benefits that an individual receives. Social benefits are typically measured by adding up the benefits received by all individuals. The numbers obtained represent the total willingness to pay of all individuals in society. The difference between the total willingness to pay and the total costs of a project can be thought of as the net "efficiency" effect of the project. It is a dollar value of the net benefits.

MEASURING INEFFICIENCY

In assessing alternative policies, economists have put particular emphasis on economic efficiency. Taxes are criticized for discouraging work effort, monopolies for restricting production and driving up prices. To measure the dollar value of an inefficiency, economists use exactly the same methodology they use to measure the dollar value of a new project. There, we calculated the consumer surplus associated with the project. Here, we calculate the consumer surplus associated with the elimination of the inefficiency. That is, economists ask: "How much would an individual be willing to give up to have the inefficiency eliminated?" Consider the inefficiency caused by a tax on cigarettes. We ask each individual how much he would be willing to pay to have the tax on cigarettes eliminated. Say his answer is $100. Thus eliminating the cigarette tax and imposing in its place a $100 **lump-sum tax** (that is, a tax that the individual would have to pay regardless of what he did) leaves his welfare unchanged. The difference between the revenue raised by the cigarette tax (say, $80) and the lump-sum tax that the individual would be willing to pay is called the **deadweight loss** or **excess burden** of the tax. It is the measure of the inefficiency of the tax. Taxes, other than lump-sum taxes, give rise to a deadweight loss because they cause individuals to forgo more-preferred consumption in favor of less-preferred consumption in order to avoid payments of the tax. Thus, even a tax that raises no government revenue—because individuals completely avoid purchasing the taxed commodity—can have a substantial excess burden.

We can calculate the deadweight loss using compensated demand curves. Assume the cost of producing a cigarette is c_0, and the tax raises the price from c_0 to $c_0 + t$, where t is the tax per pack. We assume the individual consumes q_0 packs of cigarettes with the tax, and q_1 after the tax has been removed (but replaced by a lump-sum tax that leaves him no better or no worse off than when there was a cigarette tax). We have drawn the resulting compensated demand curve in Figure 5.11. The deadweight loss is measured by the shaded area *ABC*, the area under the compensated de-

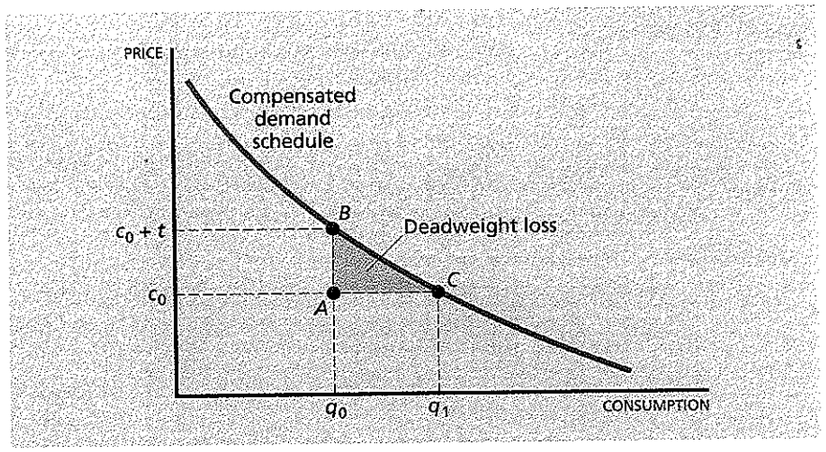

FIGURE 5.11 **Measuring Inefficiencies** The area *ABC* measures the deadweight loss, the efficiency loss as a result of a cigarette tax. A lump-sum tax that would have the same effect on the individual's welfare as the cigarette tax would raise an additional revenue of *ABC*.

mand schedule and above c_0, between the output with and without the tax.

The triangle *ABC* is sometimes called a **Harberger triangle,**[6] in honor of Chicago economist Arnold Harberger, who used such triangles not only to measure the inefficiencies associated with distortionary taxation but also to measure other inefficiencies, such as those associated with monopoly. Why does the Harberger triangle provide a measure of deadweight loss? The price tells us the value of the last unit consumed; that is, at q_0, the individual is willing to trade off $p_0 = c_0 + t$ units of "income" (with which he could have purchased other goods) for one more pack of cigarettes. Of course, when the individual has $q_0 + 1$ packs of cigarettes, he will value an additional pack of cigarettes less than when he has q_0 packs, and so the price he is willing to pay will fall.

Assume that initially consumption is 100 packs, and consumption increases by 10 packs when the tax is removed; the tax is 10 cents, and the cost of production is $1.00 per pack. (Tax revenue is 100 packs times 10 cents per pack, or $10.) The individual is willing to pay $1.10 for the first additional pack, $1.09 for the second, $1.08 for the third, and so on. If the tax

[6] See for instance, A. Harberger, "Taxation, Resource Allocation and Welfare," in *The Role of Direct and Indirect Taxes in the Federal Revenue System*, ed. J. Due (Princeton, N.J.: Princeton University Press, 1964), reprinted in A. Harberger, *Taxation and Welfare* (Chicago: University of Chicago Press, 1974).

CONSUMER SURPLUS

- Measured by the area under the (compensated) demand curve
- Used to measure the value of a government project or assess the magnitude of an inefficiency

were eliminated, and the price fell to c_0, the cost of production ($1.00 a pack), the total amount that the individual would be willing to pay would be 10 cents times 100 packs = $10 (the amount he saves on the first 100 packs he has purchased, which is equal to the tax revenue); plus 10 cents for the 101st pack (the difference between how much he values the 101st pack and what he must pay), 9 cents for the 102nd pack, etc. Remember, we are calculating how much *more* he would be willing to pay beyond the $1.00 that he will have to pay for each pack. The total he would be willing to pay is thus $10.50. Since the tax raised a revenue of $10, the deadweight loss is 50 cents. This is, of course, just the area under the compensated demand curve and above c_0, between q_0 and q_1.

QUANTIFYING DISTRIBUTIONAL EFFECTS

Assessing the distributional effects of a project or a tax is often far more complex than assessing the efficiency effects. There are many groups in a society, and each may be affected differently. Some poor individuals may be hurt, some helped; some middle-income individuals may be helped, others hurt. In some cases the rich may be helped the most, the poor helped moderately, and the middle class made only slightly worse off.

In practice, governments focus on a few summary measures of inequality. Since the poor are of particular concern, they receive special attention. The **poverty index** measures the fraction of the population whose income lies below a critical threshold; below that threshold, individuals are considered to be in poverty. In 1997, the poverty threshold for a family of four was $16,404.[7]

Another measure is the **poverty gap.** The poverty index only counts the number of individuals who are below the poverty threshold; it does not look at how far below that threshold they are. The poverty gap asks, how much income would we have to give to the poor to bring them all up to the poverty threshold?

Two other measures are briefly discussed in the appendix to this chapter.

[7] This is an estimate for 1997. U.S. Bureau of the Census, Current Population Survey, accessed at the U.S. Census Bureau web site (http://www.census.gov) on 9/4/98.

DRAWING A POVERTY LINE

The official poverty line determines how many people the government counts as poor. But what determines the poverty line itself?

In the late 1960s, an official at the Social Security Administration, Molly Orshansky, developed a method of measuring poverty from a survey of household expenditures. She found that a typical family spent one-third of its income on food. She then gathered information on minimum food budgets for families of various sizes, and multiplied that number by 3 to get an estimate of the poverty line for the different family sizes. With minor changes, Orshansky's poverty line was officially adopted in 1969 and it has been increased by the overall rate of inflation since then.

There are a number of questions one can ask about how poverty is measured: here are three.

First, the survey Orshansky relied on to find that households spent one-third of their income on food was taken in 1955. Since then, household expenditures have shifted. Households now spend a much lower percentage of income on food, perhaps one-fourth or one-fifth. If the minimum food budget were accordingly multiplied by 4 or 5, the poverty line would be much higher.

Second, the poverty line does not take in-kind benefits into account. In-kind benefits include any benefits that are not received in cash form, like Medicaid, food stamps, and subsidized school lunches. If those benefits are measured as additional income, the number of people below the poverty line falls by about 20 percent.

Finally, some critics have proposed that poverty should be thought of as a relative rather than an absolute concept. They argue that those at the bottom of society, say, the bottom 5 or 10 or 20 percent, are poor relative to everyone else. Poverty is more appropriately viewed as an extreme case of inequality.

For many, this last criticism goes too far. They fear that a relative concept of poverty could reduce the moral urgency of fighting poverty. There is broad social support for efforts to assure that people have basic levels of food, housing, clothing, and medical care, even if defining those amounts is controversial.

In 1995, a National Academy of Sciences study proposed major revisions in how we measure poverty. While there was agreement about including noncash income, the difficult problems of how best to include health care expenditures were not fully resolved. Should a sick, poor person who receives $150,000 for a kidney transplant have that added to his income, in which case he now appears to be in an upper income bracket? The study proposed an adjustment in the poverty level that went beyond just taking into account inflation, but it did not propose increasing the poverty level in proportion to increases in average income, which would have made poverty a purely relative phenomenon. But even this compromise generated a strong dissent from one of the members of the Academy's panel.

SOURCES: Joyce E. Allen and Margaret C. Simms, "Is a New Yardstick Needed to Measure Poverty?" *Focus*, February 1990, pp. 6–8; *Measuring Poverty: A New Approach* (Washington, D.C.: National Academy of Sciences, 1996).

THREE APPROACHES TO SOCIAL CHOICES

We now have the basic tools for describing social choices in those difficult cases where the project does not constitute a Pareto improvement. There are three approaches, which we shall refer to as the *compensation principle, trading-off measures,* and the *weighted benefits approach.*

THE COMPENSATION PRINCIPLE

What happens if the total willingness to pay exceeds the total costs, but the costs borne by some individuals exceed their willingness to pay? Should the project be undertaken? The **compensation principle** says that if the aggregate willingness to pay exceeds the cost, the project should be undertaken. Most economists criticize this principle, for it ignores distributional concerns. Only if the compensation is actually paid to those adversely affected can we be sure that the project is desirable, for then it is a Pareto improvement.

Because the compensation principle does not pay adequate attention to distributional concerns, economists have turned to two other approaches.

TRADE-OFFS ACROSS MEASURES

With a measure of efficiency (net benefits) and a measure of inequality, public decision-making, conceptually at least, should be easy: one simply evaluates whether the increase in efficiency is worth the increase in inequality, or vice versa.

The previous two sections have described how we measure total efficiency and inequality. These are just statistics, numbers that help to summarize the impacts of a project or program. Such summary statistics, while useful, often submerge some of the detailed information which is important in public decision-making. Ideally, we would look at the impacts on each individual, and then use the social welfare function to add up the effects. In practice, the government doesn't attempt to identify impacts on every individual, but it does at-

SOCIAL CHOICE IN PRACTICE

- Identify Pareto improvements.
- If some individuals are better off while others are worse off, identify groups of individuals who are better off and groups that are worse off (by income, region, age), and gains and losses of each major group.
  - Ascertain whether aggregate net benefits are positive (compensation principle).
  - Look at change in measure of efficiency and measure of inequality, and evaluate trade-offs.
  - Calculate weighted net benefits, weighting gains and losses to poor more heavily than those to rich, according to the social welfare function.

tempt to ascertain the effects on each major group. For instance, it may look at the impact on individuals in different income categories—say, families with incomes below $10,000, between $10,000 and $20,000, and so forth.

WEIGHTED NET BENEFITS

This may be all the information required for policy makers to make a decision. If the aggregate net benefits (the sum of the willingnesses to pay minus costs) is positive, and if the poor are net beneficiaries and the rich are net losers, then the project increases both efficiency and equity and should be adopted. But often, matters are more complicated. For instance, the poor and the rich may be worse off, but middle-income individuals better off. How do we assess such a change? Again, we turn to our social welfare function to add up the effects. We assign weights to the net gains of different groups to summarize the impacts in a single number. The social welfare function tells us how to do that. Because of the concern for equity, effects on higher-income groups are weighted less heavily. How much less heavily may determine whether it is desirable to undertake a project. For instance, a project that helps the middle class but hurts the poor and the rich might not be undertaken if we weight the losses of the poor much more heavily than the gains to the middle class.[8]

[8] Given the importance of these weights in social decision-making, economists have looked for a rational basis for assigning weights. One way is to think about how rapidly marginal utility diminishes with increased income. Inferences about this can be made from observing individual behavior in risky situations: if marginal utility diminishes very rapidly, individuals will be very averse to undertaking large risks, and will be willing to pay large premiums to divest themselves of risk. On the basis of this evidence, most economists argue that a doubling of income will lower the marginal utility of income by a factor of between 2 and 4, so that a change in the income of a middle-class individual with an income of $30,000 should be weighted half to a quarter of the same change in income of a poor individual with an income of $15,000.

The use of weights can be thought of as based on three assumptions: first, that there is diminishing marginal utility; second, that different individuals have the same relation between utility and income; and third, that society is concerned with total utility—the sum of the utilities of all individuals (the utilitarian social welfare function). While each of these assumptions may be questioned, we can also think of these procedures as simply a convenient way to summarize data that decision makers often find helpful.

REVIEW AND PRACTICE

SUMMARY

1 Welfare economics—or normative economics—is concerned with criteria for evaluating alternative economic policies. In general, it takes into account both efficiency *and* equity.

2 The social welfare function provides a framework within which the distributional consequences of a policy may be analyzed. It specifies the increase in utility of one individual that is required to compensate for a decrease in utility of another.

3 In the utilitarian social welfare function, social welfare is equal to the sum of the utilities of the individuals in society. In the Rawlsian social welfare function, social welfare is equal to the utility of the worst-off individual in society.

4 The concept of consumer surplus—how much individuals are willing to pay for a project or program in addition to what they have to pay—is used to measure the aggregate benefits of a project or program. The concept of deadweight loss is used to measure the inefficiency of a tax; it asks how much extra revenue could have been generated by a lump-sum tax that would have left individuals just as well off as the tax that was imposed.

5 As a practical matter, in evaluating alternative proposals we do not detail the impact each proposal has on each individual in society, but rather we summarize its effects by describing its impact on some measure of inequality (or on some well-identified groups) and describing the efficiency gains or losses. Alternative proposals often present trade-offs between efficiency and distribution; to get more equality one has to give up some efficiency. Differences in views arise concerning the nature of the trade-offs (how much efficiency one needs to give up to get some increase in equality), and values (how much efficiency one should be willing to give up, at the margin, to get some increase in equality).

6 Three approaches for making social choices when there is not a Pareto improvement are:

a the compensation principle

b trade-offs across measures of efficiency and equality

c the weighted benefits approach

7 The poverty index measures the fraction of the population whose income lies below some threshold.

KEY CONCEPTS

Social indifference curves
Diminishing marginal utility
Social welfare function
Utilitarianism
Rawlsianism
Interpersonal utility comparisons
Rawlsian social welfare function
Deadweight loss
Excess burden
Poverty index
Poverty gap
Compensation principle
Weighted benefits approach

QUESTIONS AND PROBLEMS

1 Assume that Friday and Crusoe have identical utility functions described by the following table.

Utility Functions for Friday and Crusoe

NUMBER OF ORANGES	UTILITY	MARGINAL UTILITY
1	11	
2	21	
3	30	
4	38	
5	45	
6	48	
7	50	
8	51	

Draw the utility function. Fill in the marginal utility data in the table above, and draw the marginal utility function.

2 Assume that there are eight oranges to be divided between Friday and Crusoe. Take a utilitarian view—assume that social welfare is the sum of the utilities of the two individuals. Using the data from Problem 1, what is the social welfare corresponding to each possible allocation of oranges? What allocation maximizes social welfare? Show that it has the property that the marginal utility of an extra orange given to each individual is the same.

3 Now take a Rawlsian view and assume that the social welfare function is the level of utility of the individual with the lowest utility level. Using the data from Problem 1, and again assuming there are eight oranges, what is the social welfare associated with each allocation of oranges? What allocation maximizes social welfare?

REVIEW AND PRACTICE

4 Draw the utility possibilities schedule based on the data from Problem 1. Mark the points that maximize social welfare under the two alternative criteria from Problems 2 and 3.

5 Assume that Crusoe's and Friday's utility functions are described in Problem 1. But assume now that initially Crusoe has six oranges and Friday two. Assume that for every two oranges taken away from Crusoe, Friday gets only one, an orange being lost in the process. What does the utility possibilities schedule look like now? Which of the feasible allocations maximizes social welfare with a utilitarian social welfare function? With a Rawlsian social welfare function?

6 If marginal utility did not decrease at all for both Friday and Crusoe, what would the utility possibilities schedule look like?

7 Consider an accident where an individual loses his leg. Assume that it lowers his utility at each level of income but increases his marginal utility (at each level of income), though only slightly. Show diagrammatically the utility functions before and after the accident. Show that if you were a utilitarian, you would give more income to the individual after the accident, but that even after the transfer, the individual with the accident is worse off than before. Show the compensation that a Rawlsian would provide.

Is it possible for a utilitarian to give more to the individual who had experienced the accident than a Rawlsian?

Under what circumstances would a utilitarian give nothing to an individual who had experienced an accident?

8 For each of the following policy changes, explain why the change is or is not likely to be a Pareto improvement:

a Building a park, financed by an increase in the local property tax rate.

b Building a park, financed by the donation of a rich philanthropist; the city acquires the land by exercising the right of eminent domain.[9]

c Increasing medical facilities for lung cancer, financed out of general revenues.

d Increasing medical care facilities for lung cancer, financed out of an increase in the cigarette tax.

e Replacing the system of agricultural price supports with a system of income supplements for poor farmers.

f Protecting the automobile industry from cheap foreign imports by imposing quotas on the importation of foreign cars.

[9] The right of eminent domain gives public authorities the right to take property, with compensation, for public uses.

g Increasing social security benefits, financed by an increase in the payroll tax.

h Replacing the primary reliance at the local level on the property tax with state revenues obtained from an income tax.

i Eliminating rent control laws.

In each case, state who the losers (if any) are likely to be. Which of these changes might be approved under the compensation principle? Which might be approved under a Rawlsian social welfare function?

9 Assume you are shipwrecked. There are ten of you in a lifeboat; you know that it will take ten days to reach shore and that there are only rations for ten man-days. (The ration is the minimum amount needed for survival.) How would a utilitarian allocate the rations? How would a Rawlsian? Some people think that even Rawlsian criteria are not sufficiently egalitarian. What might an extreme egalitarian individual advocate? What does Pareto efficiency require? In each case, state what assumptions you need to make to make the decision.

APPENDIX ALTERNATIVE MEASURES OF INEQUALITY

In the text, we introduced the two most commonly used measures of inequality. These measures are criticized, however, for focusing exclusively on the impact on the very poor. In this appendix, we discuss two more-inclusive measures.

THE LORENZ CURVE

Economists often represent the degree of inequality in an economy by a diagram called the **Lorenz curve,** shown in Figure 5.12. The Lorenz curve shows the cumulative fraction of the country's total income earned by the poorest 5 percent, the poorest 10 percent, the poorest 15 percent, and so on. If there were complete equality, then 20 percent of the income would accrue to the lowest 20 percent of the population, 40 percent to the lowest 40 percent. The Lorenz curve would be a straight line, as depicted in panel A. On the other hand, if incomes were very concentrated, then the lowest 80 percent might receive almost nothing, and the top 5 percent might receive 80 percent of total income; in this case, the Lorenz curve would be bowed, as illustrated in panel B. When there is a great deal of inequality, the shaded area between the 45-degree line in panel B and the Lorenz curve is large. When there is complete equality, as in panel A, this area is zero. Twice the area between the 45-degree line and

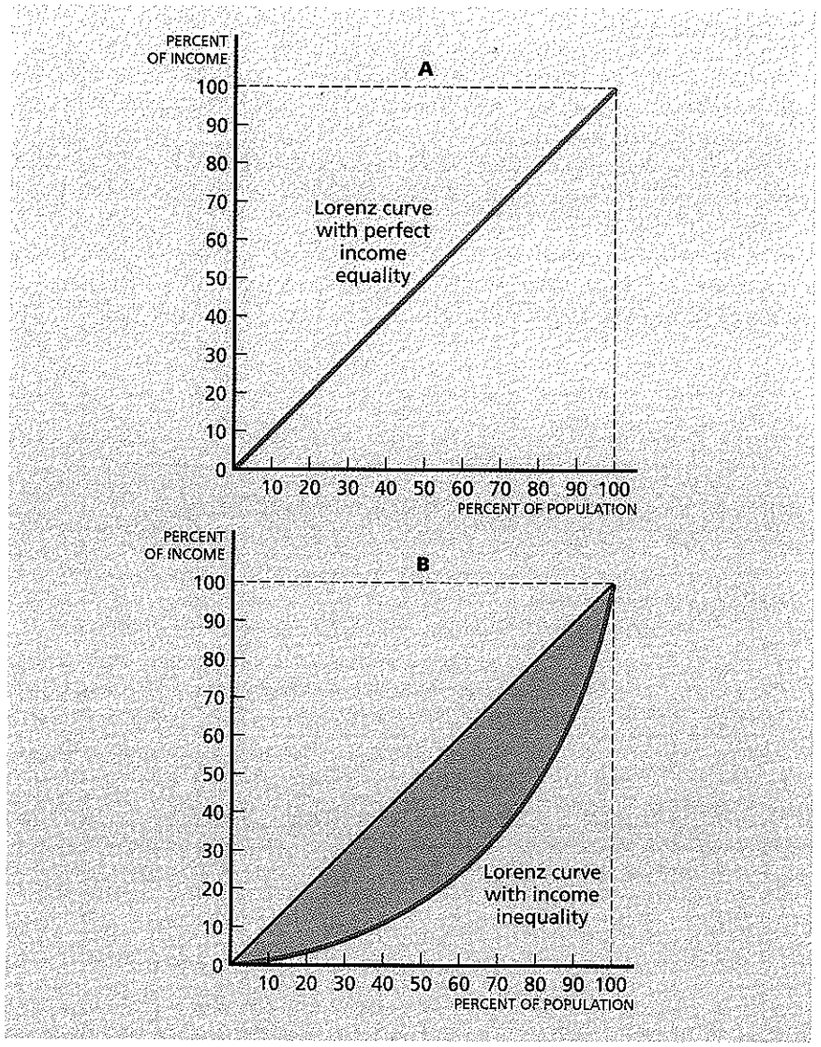

FIGURE 5.12 **The Lorenz Curve** Panel A shows a Lorenz curve for an economy in which income is evenly distributed. The bottom 20 percent of the economy has 20 percent of income, the bottom 40 percent has 40 percent of income, and so on. Panel B depicts a Lorenz curve for an economy where income is unequally distributed. The curvature of the line indicates that now the bottom 20 percent has less than 20 percent of income, the bottom 40 percent has less than 40 percent of income, and so on.

the Lorenz curve is a commonly employed measure of inequality, called the **Gini coefficient.**

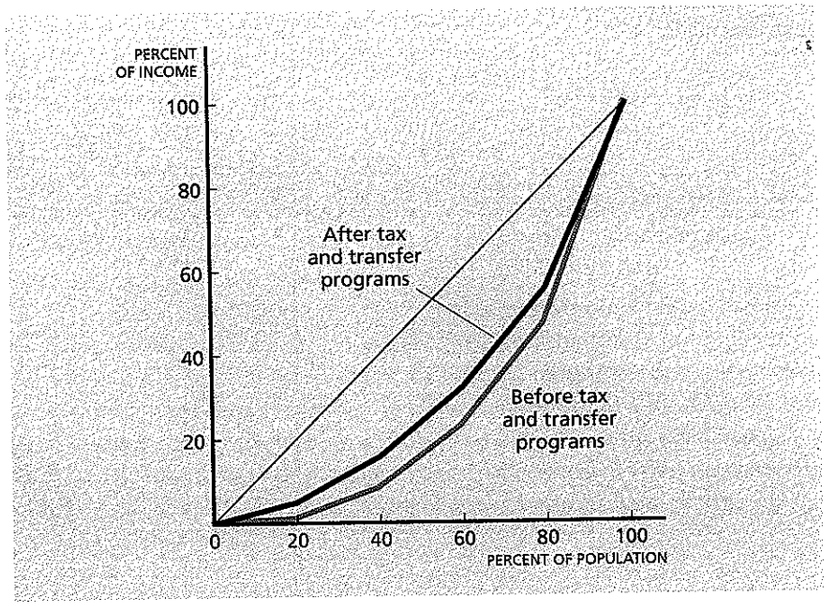

FIGURE 5.13 **Inequality Measures** Taxes and subsidies affect the distribution of income. The figure shows two Lorenz curves for the United States in 1995, one for income before taxes have been levied and government transfers have been received, and the other after. Clearly, some redistribution does take place through these mechanisms, as they move the Lorenz curve toward greater equality.

SOURCE: U.S. Census Bureau, *Current Population Survey*, March 1996, Table E.

Figure 5.13 shows Lorenz curves for the United States, both before and after the government tax and transfer programs have had their effect. The after-tax curve is decidedly inside the pre-tax, indicating that the combined effect of government redistribution programs is to make incomes more equal than the market would have made them. Thus, while the efficiency costs are less clear-cut, the redistributive gains are undeniable.

THE DALTON-ATKINSON MEASURE

There is another measure, first introduced by Sir Hugh Dalton, a professor of public finance at the London School of Economics who went on to become the Chancellor of the Exchequer for the United Kingdom. This measure was based on the premise that societies prefer more egalitarian distributions. Figure 5.14 shows two distributions of income. In distribution *B*,

more of the income is concentrated at the center, and for societies that value equality, this is the preferred distribution. We can ask, if society could move from its current distribution to a situation where income was completely equally distributed, what fraction of its total income would it be willing to give up? This fraction is called the Dalton-Atkinson measure of inequality. Of course, different individuals might have different views on the amounts that society should be *willing* to give up (this says nothing about how much they would *have* to give up to accomplish the redistribution). The amount society would be willing to give up depends on its social welfare function. With a Rawlsian social welfare function, the amount would be much larger than with a utilitarian social welfare function. An-

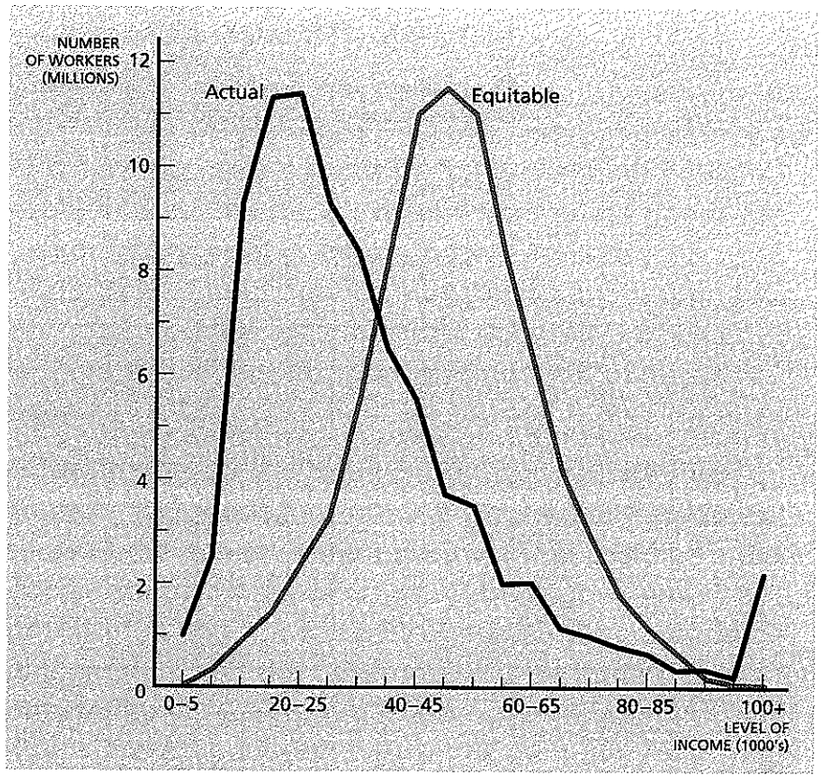

FIGURE 5.14 **Toward a More Egalitarian Income Distribution** The dark line *A* represents the actual income distribution of full-time, permanent U.S. workers in 1995. The shaded line *B* represents what the income distribution might look like under a more egalitarian setting.

SOURCE: U.S. Bureau of Labor Statistics and Bureau of the Census, *Annual Demographic Survey, March 1997 Supplement.*

thony Atkinson, of Nuffield College, Oxford, argued that the amount was significant, often between a quarter and a third of total income in more developed economies. Changes in the Dalton-Atkinson measure can be used to assess the impact on inequality of any proposed government program.[10]

[10] Formally, the Dalton-Atkinson measure can be defined as follows. Assume a utilitarian social welfare function

$$W = U(Y_1) + U(Y_2) + U(Y_3) + \ldots$$

and let Y be the average income. Then the Dalton-Atkinson measure D is given by

$$U((1 - D)Y) = U(Y_1) + U(Y_2) + U(Y_3) + \ldots$$

The measure clearly depends on the utility function. Atkinson, in his analysis, focused on constant-elasticity utility functions, which have the form

$$U = \frac{Y^{1-\alpha}}{1 - \alpha}$$

He used values of α between 1 and 2.