

# Conceptual Foundations of Cost-Benefit Analysis

It seems only natural to think about the alternative courses of action we face as individuals in terms of their costs and benefits. Is it appropriate to evaluate public policy alternatives in the same way? The CBA of the highway sketched in Chapter 1 identifies some of the practical difficulties analysts typically encounter in measuring costs and benefits. Yet, even if analysts can measure costs and benefits satisfactorily, evaluating alternatives solely in terms of their net benefits may not always be appropriate. An understanding of the conceptual foundations of CBA provides a basis for determining when CBA can be appropriately used as a decision rule, when it can usefully be part of a broader analysis, and when it should be avoided.

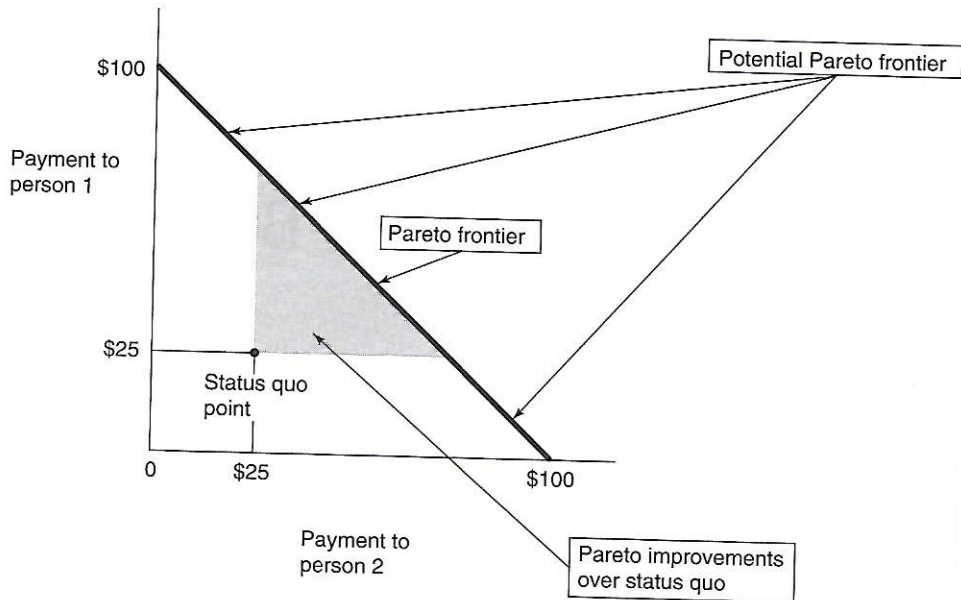
The goal of *allocative*, or *Pareto*, *efficiency* provides the conceptual basis for CBA. In this chapter we provide a nontechnical introduction to Pareto efficiency. We then explain its relationship to *potential Pareto efficiency*, which provides the practical basis for actually doing CBA. Our exploration of the roles of Pareto efficiency and potential Pareto efficiency in CBA provides a basis for distinguishing it from other analytical frameworks. It also provides a basis for understanding the various philosophical objections commonly made against the use of CBA for decision making.

## CBA AS A FRAMEWORK FOR MEASURING EFFICIENCY

CBA can be thought of as providing a framework for measuring efficiency.<sup>1</sup> Though we develop a more formal definition of efficiency in the following section, it can be thought of as a situation in which resources, such as land, labor, and capital, are deployed in their highest valued uses in terms of the goods and services they create. In situations in which analysts care only about efficiency, CBA provides a method for making direct comparisons among alternative policies. Even when goals other than efficiency are important, CBA serves as a yardstick that can be used to provide information about the relative efficiency of alternative policies. Indeed, analysts rarely encounter situations in which efficiency is not one of the relevant goals. Critical evaluation of these assertions requires a more precise definition of efficiency.

### Pareto Efficiency

A simple and intuitively appealing definition of efficiency, referred to as *Pareto efficiency*, underlies modern welfare economics and CBA. *An allocation of goods is*



**FIGURE 2-1** Pareto Efficiency

*Pareto efficient if no alternative allocation can make at least one person better off without making anyone else worse off.* An allocation of goods is inefficient, therefore, if an alternative allocation can be found that would make at least one person better off without making anyone else worse off. One would have to be malevolent not to want to achieve Pareto efficiency—why forgo gains to persons that would not inflict losses on others?


Figure 2-1 illustrates the concept of Pareto efficiency in a simple situation involving the allocation of a fixed amount of money between two persons. Imagine that the two persons will receive any total amount of money of up to \$100 if they agree on how to split it between themselves. Assume that if they do not agree, then each person receives just \$25. The vertical axis measures the amount of money received by person 1, and the horizontal axis measures the amount of money received by person 2. The point labeled \$100 on the vertical axis represents the outcome in which person 1 receives the entire \$100. Similarly, the point labeled \$100 on the horizontal axis represents the outcome in which person 2 receives the entire \$100. The line connecting these two extreme points, which we call the *potential Pareto frontier*, represents all the feasible splits between the two persons that allocate the entire \$100. Splits involving less than \$100 lie within the triangle formed by the potential Pareto frontier and the axes. The one labeled (\$25, \$25) is such a point. This point represents the status quo in the sense that it gives the amounts the two persons receive if they do not reach an agreement about splitting the \$100. The segment of the potential Pareto frontier that gives each person at least as much as the status quo is called the *Pareto frontier*.

The lightly shaded triangle formed by the lines through the status quo point and the Pareto frontier represents all the alternative allocations that would make at least one of the persons better off than the status quo without making the other worse off. The existence of these points, which are feasible alternatives to the status quo that make at least one person better off without making the other worse off, means that the status quo is not Pareto efficient. Movement to any one of these points is called a *Pareto improvement*. Any Pareto improvement that does not lie on the potential Pareto frontier would leave open the possibility of further Pareto improvements and thus not provide a Pareto-efficient allocation. Only on the potential Pareto frontier is it impossible to make a feasible reallocation that makes one person better off without making the other person worse off.

It should be clear that the segment of the potential Pareto frontier that guarantees at least \$25 to each person represents all the Pareto efficient allocations relative to the status quo. Each of these points makes a Pareto improvement over the status quo and leaves no opportunity for further improvements. The segment of the potential Pareto frontier that represents actual Pareto improvements depends upon the status quo. In other words, implicit in the concept of Pareto efficiency are the initial starting positions of the members of society. We return later to the significance of the difference between the potential and actual Pareto frontiers in our discussion of criticisms of CBA.

### Net Benefits and Pareto Efficiency

The link between positive net social benefits (henceforth, net benefits) and Pareto efficiency is straightforward: *if a policy has positive net benefits, then it is possible to find a set of transfers, or “side payments,” that makes at least one person better off without making anyone else worse off.* A full understanding of this link requires some reflection on how one measures benefits and costs in CBA. In particular, as illustrated in Figure 2-2, it requires one to consider willingness to pay (WTP) as the method for valuing the outputs of a policy and opportunity cost as the method for valuing the resources required to implement the policy. Though we develop these important concepts more fully in the next three chapters in the context of market exchange, the simple introductions that follow provide the basis for understanding the link between net benefits and Pareto efficiency.



**FIGURE 2-2** Categorization of Net Benefits of Projects

**Willingness to Pay.** Consider a proposed policy that would produce outputs of relevance to three people. Assume that these people make honest revelations of their assessments of the values of the outputs. Through a series of questions, we elicit the payments that each person would have to make or to receive under the policy so that he would be indifferent between the status quo and the policy with the payments. So, for example, imagine that person 1 honestly reveals that she would be indifferent between the status quo and paying \$100 to have the policy implemented. Similarly, person 2 might say that he is indifferent between the status quo and paying \$200 to have the policy implemented. These values are the WTP of persons 1 and 2 for the policy. Unlike persons 1 and 2, assume that person 3 does not like the impacts of the proposed policy and would have to receive a payment of \$250 if the policy were implemented to feel just as well off as he did under the status quo; this \$250 is the amount that would have to be given to the person in conjunction with the proposed policy so that he is indifferent between it and the status quo. The negative of this amount ( $-\$250$ ) would be the WTP of person 3 for the policy. As the policy in effect takes something away from person 3, the amount is called the person's *willingness to accept*. In this stylized example, the distinction between willingness to pay and willingness to accept is purely semantic. However, the distinction has substantive importance when one actually seeks to measure WTP through survey methods, because, as discussed in Chapters 15 and 18, people often demand larger payments to accept small decrements in some good than they are WTP to obtain small increments of exactly the same size. For now, we use WTP inclusively, assuming that its elicitation appropriately takes account of whether people perceive policies as giving them something or as taking away something.

The algebraic sum of these willingness-to-pay values is the appropriate measure of the net benefits of the impacts of the policy. In this example, the willingness-to-pay amounts can be divided into \$300 of benefits ( $\$100 + \$200$ ) accruing to persons 1 and 2 and \$250 of costs ( $-\$250$ ) accruing to person 3. The net benefits are thus positive and equal to \$50. If these were the only three persons affected by the policy, and if the policy required no resources to implement, then the \$50 would be the appropriate measure of net benefits from the perspective of CBA. Simple implementation of the policy would not be Pareto efficient because person 3 would be made worse off with respect to the status quo.

Yet, we can easily imagine altering the policy so that it would be Pareto efficient. For example, imagine that person 3 receives \$75 from person 1 and \$175 from person 2 as part of the policy. Now person 1 is better off than the status quo (\$100 of benefits minus \$75 given to person 3), person 2 is better off (\$200 of benefits minus \$175 given to person 3), and person 3 is no worse off (\$250 of costs from the policy minus \$250 of benefits in the form of compensation from persons 1 and 2).

The key point is that if, and only if, the aggregate net benefits of the policy as measured by the WTP of all affected individuals are positive, then there exist sets of contributions and payments that would make the policy a Pareto improvement over the status quo.

**Opportunity Cost.** The implementation of policies almost always requires the use of some inputs that could be used to produce other things of value. For example, implementing a policy to build a bridge across a river would require the use of labor, steel, concrete, construction machinery, and land that could be used to produce other things

of value to people. The concept of opportunity cost is used in CBA to place a dollar value on the inputs required to implement policies. *The opportunity cost of using an input to implement a policy is its value in its best alternative use.* Opportunity cost measures the value of what society must forgo to use the input to implement the policy.

Return to the example of the three persons whose aggregate WTP for the policy was \$50. Imagine that the policy requires inputs that have an opportunity cost of \$75. That is, if the policy were implemented, then some other members of society would have to give up goods valued at \$75. In this case, the policy does not generate enough net benefits to the three persons to allow them to compensate those who must forgo the \$75 of goods—the net benefits to society as a whole are negative \$25 (\$50 of net benefits to the three persons minus \$75 in opportunity costs to the rest of society). Thus, the policy could not be made Pareto efficient because it does not produce enough benefits to permit all those who bear costs to be compensated fully. If the opportunity cost were only \$20 instead of \$75, then net benefits to society would be \$30 and it would be possible to compensate all those who bear costs so that no one is made worse off, and some people are made better off, by the policy. In general, if the net benefits of a policy are positive, then it is potentially Pareto improving.

## USING CBA FOR DECISION MAKING

The connection between net benefits and Pareto efficiency should now be clear. *As long as analysts value all impacts in terms of willingness to pay and value all required inputs in terms of opportunity costs, then the sign of the net benefits indicates whether it would be possible to compensate those who bear costs sufficiently so that no one is made worse off and at least one person is better off.* Positive net benefits indicate the potential for compensation to make the policy Pareto efficient; negative net benefits indicate the absence of this potential.

One could imagine the following decision rule for CBA: adopt only policies that are actually Pareto efficient. In other words, only policies that yield positive benefits after providing full compensation to all those who bear costs would be adopted so that there would be no losers, only winners. Although conceptually this is appealing, such a rule would be extremely difficult to apply in practice for a number of reasons. First, it would place great informational burdens on analysts not just to measure aggregate costs and benefits, which can often be inferred from observing prices and quantities in markets, but also to measure costs and benefits for each person, a task that would generally render CBA too costly to use. Second, once the distribution of costs and benefits at the individual level were known, the administrative costs of actually making specific transfers for each government policy would almost certainly be high. Third, it is difficult to operate a practical system of compensation payments that does not distort the investment and work behavior of households. Fourth, the requirement that everyone be fully compensated would create a strong incentive for people to find ways to overstate the costs and understate the benefits that they expect to receive from policies, complicating the already difficult task of inferring how much each person is willing to pay for the outputs produced by the policy. The “actual Pareto efficiency” principle in practice would thus result in society forgoing

many policies that offer positive net benefits and the diversion of much effort toward the seeking of unjustified compensation.

### **Potential Pareto Efficiency**

CBA utilizes an alternative decision rule with somewhat less conceptual appeal, but much greater feasibility, than the actual Pareto efficiency rule. It is based on what is known as the *Kaldor-Hicks criterion*: a policy should be adopted if and only if those who will gain could fully compensate those who will lose and still be better off.<sup>2</sup> The Kaldor-Hicks criterion provides the basis for the *potential Pareto efficiency rule*, or, more commonly, the *net benefits criterion*: *adopt only policies that have positive net benefits*. As long as net benefits are positive, it is at least possible that losers could be compensated so that the policy potentially could be Pareto improving. In terms of Figure 2-1, any point on the potential Pareto frontier would pass the potential Pareto efficiency rule, while only those points on the potential Pareto frontier that guarantee at least \$25 to each person (the heavily shaded segment of the potential Pareto frontier) pass the actual Pareto efficiency rule.

In practice, the assessment of whether particular policies increase efficiency depends on whether they represent *potential Pareto improvements*. That is, do the policies provide sufficient net gains so that all losers could be compensated? Potential Pareto efficiency is achieved only when all potential Pareto improvements have been exhausted.

Several justifications, aside from feasibility, are commonly offered in defense of the potential Pareto efficiency rule. First, by always choosing policies with positive net benefits, society maximizes aggregate wealth. This indirectly helps those who are worse off in society because richer societies have greater capability for helping their poorest members and, if redistribution is a normal good (that is, other things being equal, people want more of it as their wealth increases), members of society have a greater willingness to help.<sup>3</sup> Second, it is likely that different policies will have different sets of winners and losers. Thus, if the rule is consistently applied to government activity, then costs and benefits will tend to average out across people so that each person is likely to realize positive net benefits from the full collection of policies. Third, as we discuss later in this chapter, the rule stands in contrast to the incentives in representative political systems to give too much weight to costs and benefits that accrue to organized groups and too little weight to costs and benefits that accrue to unorganized interests. Its use in public discourse may thereby reduce the chances that Pareto-inefficient policies will be adopted. Fourth, if a more equal distribution of wealth or income is an important goal, then it is possible to address it directly through transfers after a large number of efficiency-enhancing policies have been adopted. In other words, redistribution, at least in theory, can be done “wholesale” with a single redistribution program rather than “retail” in each particular program.

### **Application of the Decision Rule in Practice**

Two policies can be thought of as independent if the adoption of one does not influence the costs and benefits of the other. When all relevant projects are independent, the CBA decision rule is simple: *adopt all policies that have positive net benefits*. A more

**TABLE 2-1** Choosing Efficient Projects and the Use of Net Benefits versus Benefit-Cost Ratios

	<b>Costs</b> (millions of dollars)	<b>Benefits</b> (millions of dollars)	<b>Net Benefits</b> (millions of dollars)	<b>Benefits/Costs</b>
No project	0	0	0	—
Project A	1	10	9	10
Project B	10	30	20	3
Project C	4	8	4	2
Project D	3	5	2	1.7
Projects C and D	7	21	14	3
Project E	10	8	−2	0.8

(1) *No constraints: Choose A, B, and combination C and D (net benefits equal \$43 million).*

(2) *All projects mutually exclusive: Choose B (net benefits equal \$20 million).*

(3) *Total costs cannot exceed \$10 million: Choose A and combination C and D (net benefits equal \$23 million).*

Source: Adapted from David L. Weimer and Aidan R. Vining, *Policy Analysis: Concepts and Practice*, 5th ed. (Upper Saddle River, NJ: Pearson Prentice Hall, 2011), Figure 16.2.

general version of the rule applies in situations involving multiple policies that may enhance or interfere with each other: *choose the combination of policies that maximizes net benefits.* Physical, budgetary, and other constraints may limit the combinations of policies that are feasible.

Consider the list of projects in Table 2-1. Interpret the costs and benefits as being expressed in terms of present values, so that they can be directly compared with dollars of current consumption. Note that projects C and D are shown as synergistic. That is, the net benefits from adopting both together exceed the sum of the net benefits from adopting each one independently. Such might be the case if project C were a dam that created a reservoir that could be used for recreation as well as hydroelectric power and D were a road that increased access to the reservoir. Of course, projects can also interfere with each other; for instance, the dam might reduce the benefits of a downstream recreation project. The important point is that care must be taken to determine interactions among projects so that the combinations of projects providing the greatest net benefits in aggregate can be readily identified.

Suppose we could choose any combination of projects; then we should simply choose all those with positive net benefits—namely, projects A, B, and combination C and D.

Suppose now the policies are mutually exclusive. For example, we cannot drain a swamp to create agricultural land and simultaneously preserve it as a wildlife refuge. When all the available policies are mutually exclusive, efficiency is maximized by choosing the one with the largest net positive benefits—project B, with net benefits of \$20 million. Assume, however, that all projects are mutually exclusive, except C and D, which can be built together to obtain synergistic gains. By taking the combination of C and D to be a separate project, we can consider all the projects on the list to be mutually exclusive. Looking down the column labeled “Net Benefits,” we see that project B still offers the largest net benefits and therefore should be the one selected, but the combination of C and D offers the next highest net benefits.

Analysts often compare programs in terms of *benefit-cost ratios*. Note that project B, which offers the largest net benefits, does not have the largest ratio of benefits to

costs. Project A has a benefit-cost ratio of 10, while project B has a benefit-cost ratio of only 3. Nevertheless, project B should be selected because it offers larger net benefits than project A. This comparison shows how the benefit-cost ratio can sometimes confuse the choice process when the projects under consideration are of different scale (that is, project B involves substantially higher costs than project A). Furthermore, the benefit-cost ratio is sensitive to whether negative WTP (willingness to accept) amounts are subtracted from benefits or added to costs. For example, imagine that the cost of \$10 million for project B was opportunity costs and the benefits of \$30 million consisted of \$40 million for one group and  $-\$10$  million for another. Treating the negative WTP as a cost rather than as a negative benefit would leave the net benefits unchanged but lower the benefit-cost ratio from 3 to 2. Thus, benefit-cost ratios are subject to manipulation. For these reasons, *we recommend that analysts avoid using benefit-cost ratios and rely instead on net benefits to rank policies.*

Return to Table 2-1 and interpret the listed costs as public expenditures exactly equal to opportunity costs and the listed benefits as the willingness-to-pay values for all project effects. Now assume that, while none of the projects are mutually exclusive in a physical sense, total public expenditures (costs) cannot exceed \$10 million because of a budget constraint that is binding for political reasons. If project B is selected, then the budget constraint is met, and net benefits of \$20 million result. If project A and the combination of projects C and D are selected instead, then the budget constraint is also met, but net benefits of \$23 million result. No other feasible combination offers larger net benefits. Thus, under the budget constraint, net benefits are maximized by choosing projects A and the combination of C and D.

## FUNDAMENTAL ISSUES RELATED TO WILLINGNESS TO PAY

Three sets of fundamental issues arise with respect to the interpretation of WTP as a measure of benefits in the assessment of the efficiency of policies. First, a theoretical limitation in the aggregation of willingness-to-pay amounts across individuals opens the possibility that the net benefits criterion will not lead to fully satisfactory rankings of policies. Second, normative issues arise because of the dependence of WTP on the distribution of wealth in society. Third, normative issues also arise with respect to the issue of *standing*, which concerns whose WTP counts in the aggregation of benefits.

### **The Theoretical Limitation of Willingness to Pay as Basis for Social Orderings**

Although using net benefits as a basis for choosing efficient public policies is intuitively appealing, its implementation through the aggregation of the willingness-to-pay amounts of the members of society confronts a fundamental theoretical limitation: ranking policies in terms of net benefits does not guarantee a transitive social ordering of the policies.

A *transitive* ordering requires that if  $X$  is preferred to  $Y$ , and  $Y$  is preferred to  $Z$ , then  $X$  is preferred to  $Z$ . The logic of transitivity seems so clear that it is usually taken as an axiom of rationality in the preferences of individuals. We would certainly be

**TABLE 2-2** Cyclical Social Preferences under Pairwise Majority Rule Voting

<i>Preference Ordering</i>	<i>Voter 1</i>	<i>Voter 2</i>	<i>Voter 3</i>
First Choice	<i>X</i>	<i>Z</i>	<i>Y</i>
Second Choice	<i>Y</i>	<i>X</i>	<i>Z</i>
Third Choice	<i>Z</i>	<i>Y</i>	<i>X</i>

(1) *Pairwise Voting Outcomes: X versus Y, X wins; Y versus Z, Y wins; X versus Z, Z wins.*

(2) *Implied Social Ordering: X is preferred to Y, Y is preferred to Z, but Z is preferred to X!*

skeptical about the mental state of someone who tells us she prefers apples to oranges, and she prefers oranges to peaches, but she prefers peaches to apples. This violation of transitivity implies a cyclical, and therefore ambiguous, ordering of the alternatives. Clearly, transitivity is a desirable property of any preference ordering.

If every member of a society has transitive preferences, does it follow that reasonable procedures for aggregating their preferences will always produce a transitive social ordering? An example makes clear that the answer is no. Consider a common aggregation procedure: majority rule voting over pairs of alternatives. Imagine that society consists of three voters who have preferences over three alternatives, *X*, *Y*, and *Z*, as displayed in Table 2-2. Specifically, voter 1 prefers *X* to *Y* to *Z*, voter 2 prefers *Z* to *X* to *Y*, and voter 3 prefers *Y* to *Z* to *X*. If the voters express their sincere preferences in each round of voting, then we would find that given the choice between *X* and *Y*, voters 1 and 2 (a majority) would vote for *X* because they each prefer it to *Y*. Similarly, given the choice between *Y* and *Z*, a majority would vote for *Y*. Yet in a choice between *X* and *Z*, a majority would vote for *Z*. Thus, the implied social ordering is intransitive because *X* is preferred to *Y*, *Y* is preferred to *Z*, but *Z* is preferred to *X*!

Is the possibility of obtaining an intransitive social ordering peculiar to the use of pairwise majority rule voting to produce rankings alternatives? Surprisingly, it can result from any rule for creating a social ordering that satisfies certain minimal requirements. We cannot expect any rule for creating a social ranking of policy alternatives to be fully satisfactory. In 1951, Kenneth Arrow proved that any *social choice rule* that satisfies a basic set of fairness conditions can produce an intransitive social ordering.<sup>4</sup> *Arrow's theorem* applies to any rule for ranking alternatives in which two or more persons must rank three or more alternatives. It requires any such scheme to satisfy at least the following conditions to be considered fair: First, each person is allowed to have any transitive preferences over the possible policy alternatives (*axiom of unrestricted domain*). Second, if one alternative is unanimously preferred to a second, then the rule for choice will not select the second (*axiom of Pareto choice*). Third, the ranking of any two alternatives should not depend on what other alternatives are available (*axiom of independence*). Fourth, the rule must not allow any one person dictatorial power to impose his or her preferences as the social ordering (*axiom of nondictatorship*). Arrow's theorem states that any fair rule for choice (one that satisfies the four previous axioms) will not guarantee a transitive *social ordering* of policy alternatives. That is, it is possible that individual preferences are such that the social ordering will be

cyclical. Thus, unless the net benefit rule, which is a social choice rule, violates one of the axioms, it cannot guarantee a transitive social ordering of policies.

In order to ensure that the use of WTP in the implementation of the net benefit rule will produce a transitive social ordering of policies, some restrictions, violating the axiom of unrestricted domain, must be placed on the preferences that individuals are allowed to hold.<sup>5</sup> Economic models commonly assume that individual preferences are represented by utility functions (numerical representations of preference orderings) that exhibit positive but declining marginal utility; that is, other things equal, incremental consumption of any good increases utility but not by as much as the previous incremental unit. Unfortunately, this relatively weak restriction of the domain of preferences (it rules out preferences that cannot be represented by such utility functions) is not enough to guarantee that the net benefit rule based on WTP will *always* produce a transitive social ordering. Two additional restrictions are required: (1) The utility functions of individuals must be such that the individual demand curves that they imply can be aggregated into a market demand curve with the sum of individual incomes as an argument, and (2) all individuals must face the same set of prices.<sup>6</sup> The first restriction is quite strong in that it requires each individual's demand for each good to increase linearly with increasing income and have the same rate of increase for each individual. The second restriction, generally satisfied when all goods are traded in markets, may be violated when policies allocate quantities of goods to individuals who cannot resell them in markets.

The necessity of restricting the allowed preferences of individuals to guarantee a transitive social ordering from the use of WTP in the implementation of the net benefits criterion makes clear that it is an imperfect criterion for assessing the relative efficiency of alternative policies.<sup>7</sup> Of course, analysts can avoid this theoretical problem by assuming that the preferences of individual consumers conform to restrictive assumptions consistent with the existence of an appropriate aggregate demand function. Alternatively, analysts can avoid it by assuming that policies affect the price of only a single good. Indeed, as discussed in the next three chapters, analysts seeking to estimate WTP typically work with an aggregate, or market, demand schedule for a single good, implicitly assuming away price effects in the markets for other goods.

Despite its theoretical imperfection as a measure of efficiency, WTP is an intuitively appealing and practical concept for guiding the implementation of the net benefits criterion. As discussed next, however, its dependence on the distribution of wealth raises a serious normative concern about its use.

### **Dependence of Willingness to Pay on the Distribution of Wealth**

The willingness of a person to pay to obtain a desired policy impact will tend to be higher the greater the wealth that she or he has available. Consequently, the sum of the willingness of persons to pay, the benefit measure in CBA, depends on their levels of wealth. If the distribution of wealth in society were to be changed, then it would be likely that the sum of individuals' willingness-to-pay amounts would change as well, perhaps altering the ranking of alternative policies in terms of their net benefits.

The dependence of net benefits on the distribution of wealth would not pose a conceptual problem if losers from adopted policies were *actually* compensated so that

the adopted policies would produce actual, rather than potential, Pareto improvements. From a utilitarian perspective, Pareto improvement guarantees that the sum of utilities of individuals in society increases. In application of the potential Pareto principle, however, it is possible that an adopted policy could actually lower the sum of utilities if people with different levels of wealth had different *marginal utilities of money*.<sup>8</sup> As an illustration, consider a policy that gives \$10 of benefits to a person with high wealth and inflicts \$9 of costs on a person with low wealth. If the low-wealth person's marginal utility of money is higher than that of the high-wealth person, then it is possible that the utility loss of the low-wealth person could outweigh the utility gain of the high-wealth person. Thus, while the Pareto principle allows us to avoid interpersonal utility comparisons by guaranteeing increases in aggregate utility for policies with positive net benefits, the potential Pareto principle does not do so.

The implication of the dependence of WTP on wealth is that the justification for the potential Pareto principle weakens for policies that concentrate costs and benefits on different wealth groups. Policies with positive net benefits that concentrate costs on low-wealth groups may not increase aggregate utility; moreover, policies with negative net benefits that concentrate benefits on low-wealth groups may not decrease aggregate utility. However, if the potential Pareto principle is consistently applied and adopted policies do not produce consistent losers or winners, then the overall effects of the policies taken together will tend to make everyone better off. Hence, concerns about reductions in aggregate utility would be unfounded.

Critics of CBA sometimes question the validity of the concept of Pareto efficiency itself because it depends on the status quo distribution of wealth. Returning to Figure 2-1, note that the location of the Pareto frontier would change if the location of the status quo point were changed. Some have advocated the formulation of a social welfare function that maps the utility, wealth, or consumption of all individuals in society into an index that ranks alternative distributions of goods.<sup>9</sup> In this broader framework incorporating distributional values, an efficient policy is one that maximizes the value of the social welfare function. But how does society determine the social welfare function? Unfortunately, Arrow's theorem, as well as practical difficulties in obtaining needed information, precludes the formulation of a social welfare function through any fair collective choice procedure.<sup>10</sup> In practice, it must therefore be provided subjectively by the analyst.<sup>11</sup> We believe that it is usually better to keep the subjective distributional values of analysts explicit by comparing policies both in terms of efficiency and the selected distributional criteria, as illustrated in the discussion of multigoal analysis and distributionally weighted CBA later in this chapter. As an alternative, analysts can report net benefits by wealth or income group as well as for society as a whole.

### **Dependence of Net Benefits on Assumptions About Standing**

The question of whose WTP should count in the aggregation of net benefits has come to be known as the issue of standing.<sup>12</sup> It has immediate practical importance in at least three contexts: the jurisdictional definition of society, the exclusion of socially unacceptable preferences, and the inclusion of the preferences of future generations. A recognition of social constraints, rights, and duties often helps answer the question of standing.

**Jurisdictional Definition of Society.** The most inclusive definition of society encompasses all people, no matter where they live or to which government they owe allegiance. Analysts working for the United Nations or some other international organization might very well adopt such a universalistic, or global, perspective. Yet for purposes of CBA, most analysts define society at the national level. The basis for this restriction in jurisdiction is the notion that the citizens of a country share a common constitution, formal or informal, that sets out fundamental values and rules for making collective choices. In a sense, they consent to being a society. Furthermore, they accept that the citizens of other countries have their own constitutions that make them distinct polities.

The distinction between universal and national jurisdiction becomes relevant in the evaluation of policies whose impacts spill over national borders. For example, if U.S. analysts adopt the national-level jurisdiction as defining society, then they would not attempt to measure the willingness of Canadian residents to pay to avoid pollution originating in the United States that exacerbates acid rain in Canada. Of course, the willingness of U.S. citizens to pay to reduce acid rain in Canada should be included in the CBA, though in practice, it would be very difficult to measure.

As in the highway example discussed in Chapter 1, a similar issue arises with respect to subnational units of government. As an illustration, consider a city that is deciding whether to build a bike path. Assume that a CBA from the national perspective (giving standing to everyone in the country) predicts that the project will generate \$1 million in benefits (which all accrue to city residents) and \$2 million in costs (which are also borne by city residents), thereby resulting in negative \$1 million in net benefits (or \$1 million in net costs). Also assume, however, that through an intergovernmental grants program the national government will repay the city's \$2 million of costs resulting from this particular project. The grant appears to the city residents as a \$2 million benefit offsetting \$2 million in local costs. Thus, from the perspective of the city, the bike path generates \$1 million in net benefits rather than \$1 million in net costs.

One can make an argument that the city should treat its residents as the relevant society and, hence, should not give standing to nonresidents. The city government has a charter to promote the welfare of its residents. The city by itself can do relatively little to affect national policy—even if it does not take advantage of all the opportunities offered by the national government, other cities probably will. Furthermore, analysts who do not adopt the city's perspective, but instead employ only the broader national perspective, risk losing influence, a possibility of special concern to analysts who earn their living by giving advice to the city.

Adopting the subnational perspective, however, makes CBA a less valuable decision rule for public policy. We believe that *analysts should ideally conduct CBA from the national perspective*. They may, of course, also conduct a parallel CBA from the subnational perspective as a response to the narrower interests of their clients. If major impacts spill over national borders, then the CBA should be done from the global as well as the national perspective.

**Jurisdictional Membership.** Deciding the jurisdictional definition of society leaves open a number of questions about who should be counted as members of the jurisdiction. For example, almost all analysts agree that citizens of their country, whether living

domestically or abroad, should have standing. With respect to noncitizens in their country, most analysts would probably give standing to those who were in the country legally. Less consensus exists with respect to the standing of other categories of people: should illegal aliens have standing? What about the children of illegal aliens?

One source of guidance for answering these types of questions is the system of legally defined rights.<sup>13</sup> For example, a ruling by the courts that the children of illegal aliens are entitled to access publicly funded education might encourage the analyst to give these children standing in CBA. Reliance on legally defined rights to determine standing, however, is not always morally acceptable. It would not have been right to deny standing in CBA to slaves in the antebellum United States, nonwhites in apartheid South Africa, or Jews in Nazi Germany simply because they lacked legal rights. Therefore, legal rights alone cannot fully resolve the issue of standing in CBA. They provide a presumption, but one that analysts may sometimes have an ethical responsibility to challenge. Democratic regimes usually provide mechanisms for challenging such presumptions, but often with personal cost to individual analysts.

One other issue of membership deserves brief mention. CBA is anthropocentric. *Only the willingness of people to pay counts.* Neither flora nor fauna have standing. That is not to say that their “interests” have no representation. Many people are willing to pay to preserve a species, and some are even willing to pay to preserve individual animals or plants. As discussed in Chapter 9, it is conceptually correct within the CBA framework to take account of these willingness-to-pay amounts, though doing so effectively is very often beyond our analytical reach.

**Exclusion of Socially Unacceptable Preferences.** People sometimes hold preferences that society seeks to suppress through widely supported legal sanctions. For instance, although some people would be willing to pay for the opportunity to have sexual relations with children, most countries attempt to thwart the expression of such preferences through strict criminal penalties. Should such socially unacceptable preferences be given standing in CBA?

One approach to answering this question adds duties and prohibitions to legal rights as sources of guidance about social values. Together they can be thought of as social constraints that should be taken into account in CBA just as the analyst takes into account physical and budgetary constraints.<sup>14</sup> Clear and widely accepted legal sanctions may help identify preferences that should not have standing.

An important application arises in estimating the net benefits of policies that are intended to reduce the amount of criminal behavior in society. Some analysts count reductions in the monetary returns to crime as a cost borne by criminals, offsetting the benefits of reduced criminal activity enjoyed by their victims.<sup>15</sup> As the returns from crime are illegal and widely viewed as wrong, however, the social constraint perspective argues against treating them in this manner.

The issue of the standing of preferences can be especially difficult for analysts to resolve when they are dealing with foreign cultures. Consider, for instance, the CBA of a program to bring water to poor communities in Haiti.<sup>16</sup> Analysts found that husbands had negative willingness to pay amounts for the time that their wives saved from easier access to water. By contemporary standards in most urban settings, people would generally regard these preferences as unworthy. Yet in the cultural context of rural Haiti at the time,

they were consistent with prevailing norms. Should these preferences of husbands have standing? In practice, lack of data to estimate willingness-to-pay amounts for this sort of impact usually spares analysts from having to answer such difficult questions.

**Inclusion of the Preferences of Future Generations.** Some policies adopted today, such as the disposal of nuclear wastes or the restoration of wilderness areas, may have impacts on people not yet born. Though we believe that these people should have standing in CBA, there is no way to measure their WTP directly because they are not yet here to express it.<sup>17</sup> How serious a problem does this pose for CBA?

The absence of direct measures of the willingness of future generations to pay for policy impacts generally poses few problems for two reasons. First, because few policies involve impacts that appear only in the far future, the WTP of people alive today for the effects during their lifetimes can be used to some extent to predict how future generations will value them. Second, as most people alive today care about the well-being of their children, grandchildren, and great-grandchildren, whether born or yet to be born, they are likely to include the interests of these generations to some extent in their own valuations of impacts. Indeed, because people cannot predict with certainty the place that their future offspring will hold in society, they are likely to take a very broad view of future impacts.

In Chapters 9 and 10, we return to the question of the standing of future generations when we discuss existence value and the social discount rate.

## CONCERNS ABOUT THE ROLE OF CBA IN THE POLITICAL PROCESS

The most vocal critics of CBA fear that it subverts democratic values. Some see the monetizing of impacts as a profane attempt to place a price on everything. Others see CBA as undermining democracy. Though these fears are largely unfounded, they deserve explicit consideration by advocates of CBA.

### Does CBA Debase the Terms of Public Discourse?

A number of objections have been raised to the effort made in CBA to value all policy impacts in terms of dollars: Pricing goods not normally traded in markets—for example, life itself—decreases their perceived value by implying that they can be compared to goods that are traded in markets; pricing such goods reduces their perceived value by weakening the claim that they should not be for sale in any circumstance; and pricing all goods undercuts the claim that some goods are “priceless.”<sup>18</sup> The language and conceptual frameworks that people use almost certainly affect the nature of debate to some extent. It is not clear, however, how influential the technical concepts of economics are in actually shaping public discourse. In any event, the correct interpretation of how nonmarket goods are monetized largely undercuts the charge that CBA debases public discourse.

Consider the issue of the monetization of the value of life. On the surface it may appear that economists are implying that a price can be put on someone’s life. A closer look, which we provide in Chapters 14 and 16, indicates that the value of life estimated by economists is a measure of how much people are willing to pay to reduce

## EXHIBIT 2-1

Does wealth produce happiness? Surveys conducted within countries consistently find that rich people (say those in the top quarter of the income distribution) on average report being happier than poorer people (say those in the bottom quarter of the income distribution). Yet, if one looks at either of these groups over time, one discovers that its absolute level of happiness is roughly constant despite the fact that economic growth has made it richer. Similarly, comparing the happiness of the rich (or poor) across countries generally shows similar levels of happiness despite substantial differences in the overall levels of wealth between the countries. What explains this puzzle? Richard Layard suggests two psychological effects that move up the norm to which people

compare their own circumstances as societies become wealthier: habituation and rivalry. Habituation involves getting used to things we have—an initial feeling of happiness from acquisition tends to evaporate as we get used to having the good. Rivalry involves comparing one's situation to those in a reference group—happiness depends on relative position.

These phenomena raise concerns about interpreting changes in social surplus as changes in aggregate happiness. A policy that increased everyone's income would certainly pass the net benefits test. Yet extreme habituation might quickly return everyone to their initial levels of utility, or extreme rivalry would result in no utility gains at all because no one's relative position changes!

*Source:* Adapted from Richard Layard, "Happiness: Has Social Science a Clue?" Lionel Robbins Memorial Lectures, London School of Economics, Lecture 1: Income and Happiness: Rethinking Economic Policy, March 3, 4, and 5, 2003.

their risk of death; in other words, it is the value of a *statistical life*, the WTP to avoid risks that will result in one less death in a population. Although it may not be appropriate to place a dollar value on the life of any particular person, it is appropriate to use the value of a statistical life in assessing proposed policies that change the risk of death that people face.

Every day people voluntarily make trade-offs between changes in the risk of death and other values: driving faster to save time increases the risk of being involved in a fatal traffic accident; eating fatty foods is pleasurable but increases the risk of fatal heart disease; skiing is exhilarating but risks fatal injury. Is it inappropriate to take account of these preferences in valuing the impacts of public policies? Most economists would answer no. Indeed, valuing statistical lives seems less problematic than attempting to place a dollar value on a specific person by estimating the person's forgone future earnings, the procedure employed by courts in cases of wrongful death.

### Does CBA Undermine Democracy?

Some critics of CBA charge that it undermines democracy by imposing a single goal, efficiency, in the assessment of public policies. Their charge would be justified if the appropriate comparison were between a world in which public policy is determined solely through democratic processes that gave equal weight to all interests and a world in which public policy is determined strictly through the application of CBA. But this is an inappropriate comparison for two reasons. First, actual governmental processes fall far short of "ideal democracy." Second, at most, CBA has modest influence in public policy making.<sup>19</sup>

The interests of vocal constituencies, often those who can organize themselves in anticipation of obtaining concentrated benefits or avoiding concentrated costs, typically receive great attention from those in representative governments who wish to be reelected or advance to higher office. Less vocal constituencies usually have their interests represented less well. The interests of many of these less vocal constituencies are often better reflected in CBA. For example, CBA takes account of the individually small, but in aggregate large, costs borne by consumers because of government price-support programs that raise prices to the benefit of a small number of well-organized agricultural producers. But CBA rarely serves as the decisive decision rule for public policy. Indeed, it is difficult to identify important public policies selected *solely* on the basis of CBA.

A realistic assessment of representative democracy and the current influence of CBA should allay concerns that the latter is subverting the former. To the extent it is influential, CBA probably contributes to more democratic public policy by paying attention to diffuse interests typically underrepresented in a representative democracy. It would have to become much more influential before it could possibly be viewed as undermining democratic processes. Despite our hopes that the readers of this book will help make the use of CBA more prevalent, we have no concerns about it being too influential in the near future.

## LIMITATIONS OF CBA: OTHER ANALYTICAL APPROACHES

It is important for analysts to realize the limitations of CBA. Two types of circumstances make the net benefits criterion an inappropriate decision rule for public policy. First, technical limitations may make it impossible to quantify and then monetize all relevant impacts as costs and benefits. Second, goals other than efficiency are relevant to the policy. For example, some policies are intended to affect the equality of outcomes or opportunity. Nevertheless, even when the net benefits criterion is not appropriate as a decision rule, CBA usually provides a useful yardstick for comparing alternative policies in terms of efficiency along with other goals.

### Technical Limitations to CBA

CBA in its pure form requires that all impacts relevant to efficiency be quantified and made commensurate through monetization. Only when all the costs and benefits are expressed in dollars can the potential Pareto principle be applied through the calculation of net benefits. Limitations in theory, data, or analytical resources, however, may make it impossible for the analyst to measure and value all impacts of a policy as commensurate costs and benefits. Nonetheless, it may still be desirable to do a qualitative cost-benefit analysis or, if all but one important effect can be monetized, to switch from CBA to cost-effectiveness analysis. A brief description of each of these alternative approaches follows.

**Qualitative CBA.** The advice given by Benjamin Franklin at the beginning of Chapter 1 can be thought of as a prescription for qualitative CBA. In conducting qualitative CBA, the analyst typically monetizes as many of the impacts as possible and

then makes qualitative estimates of the relative importance of the remaining costs and benefits. Consider, for instance, a program to plant trees along an urban highway. The cost of the program, which consists only of the expenditures that must be made to hire a contractor to plant and to maintain the trees, can be directly monetized. The benefits, however, include a number of effects that are likely to be difficult to monetize: the visual pleasure the trees give to motorists, the reduction of noise in adjoining neighborhoods, and the filtering of pollutants from the air. With sufficient resources, the analyst would be able to monetize these benefits through a variety of techniques such as surveys of motorists and comparisons with the effects of other noise-reduction programs on property values. But because the program involves relatively small costs, it is unlikely that such efforts would be justified. Instead, a reasonable approach would be to list these benefits with rough estimates of their order of magnitude.

Analysts who lack the time, data, or other resources needed to value all relevant impacts directly may be able to make use of estimates found in other cost-benefit analyses or economic research. For example, most analysts doing CBA do not directly estimate people's WTP for reductions in mortality risk. Instead, as discussed in Chapters 13, 14, and 16, they rely on econometric studies investigating how people trade such things as changes in wages for changes in levels of risk.

When possible, analysts should quantify the impacts of the policy; that is, they should estimate the numeric values of the nonmonetized impacts. For example, consider analysis of a proposed regulation to restrict commercial fishing practices so that fewer dolphins will be killed per ton of tuna harvested. The regulation produces a benefit because some people have a positive willingness to pay for dolphin deaths avoided. Actually monetizing, that is, measuring the WTP, is a difficult task that might not be feasible for the analyst conducting the CBA. Even if monetization is infeasible, however, it is useful to attempt to predict the number of dolphins saved by the regulation. Doing so increases the usefulness of the qualitative CBA for others by conveying the magnitude of the impact of the regulation. Additionally, the client or other users of the analysis may be able to provide estimates of the willingness of people to pay for each dolphin saved so that (fully monetized) CBA becomes feasible.

Analysts often face a more complicated choice than simply whether to quantify a category of costs or benefits. Empirical measures can have varying degrees of accuracy, ranging from precise estimates in which we have great confidence to imprecise estimates in which we have little confidence. The decision to quantify, and with what degree of effort, should reflect the value of the increased precision that can be obtained and the costs of obtaining it. In other words, we should make such decisions within a CBA framework!

**Cost-Effectiveness Analysis.** Analysts can often quantify impacts but not monetize them all. If the analysts are unable or unwilling to monetize the major benefit, then cost-effectiveness analysis may be appropriate. Because not all of the impacts can be monetized, it is not possible to estimate net benefits. The analysts can, however, construct a ratio involving the quantitative, but nonmonetized, benefit and the total dollar costs. A comparison allows the analyst to rank policies in terms of the cost-effectiveness criterion. However, unlike the net benefits criterion of CBA, it does not directly allow the analysts to conclude that the highest-ranked policy contributes to greater efficiency.

Return to the qualitative CBA of the fishing regulation discussed earlier. Suppose that, except for the benefit from avoided dolphin deaths, all the impacts could be monetized to a net cost of  $c$  dollars. If the number of avoided dolphin deaths were  $n_d$ , then the analyst could construct an effectiveness-cost ratio for the regulation,  $n_d/c$ , which can be interpreted as the average number of dolphins saved per dollar of cost borne. (Alternatively, the analyst could construct the cost-effectiveness ratio as  $c/n_d$ , which can be interpreted as the average dollar cost per dolphin saved.) Now imagine a number of alternative regulations, each of which involves different net costs and a different number of dolphins saved. A cost-effectiveness ratio can be calculated for each of these programs to facilitate comparison across alternative regulations.

Using cost-effectiveness analysis for decision making usually requires that some additional information be brought to bear. If the objective is to save as many dolphins as possible at a net cost of no more than  $c^*$ , then the analyst should select the most effective regulation from among those with net costs of less than  $c^*$ . Alternatively, if the objective is to save at least  $n_d^*$  dolphins, then the analyst should select the regulation with the lowest cost from among those regulations saving at least  $n_d^*$ . This is not necessarily the alternative with the best cost-effectiveness ratio. For example, if  $n_d^* = 1,000$ , one would choose a regulation that saved 1,000 dolphins at a cost of \$1 million (\$1,000 per dolphin saved) over an alternative regulation that saved 500 dolphins at a cost of \$50 thousand (\$100 per dolphin saved).

Analysts often encounter situations in which they or their clients are unable or unwilling to monetize impacts such as human lives saved, injuries avoided, and the acres of old-growth forest preserved. Because cost-effectiveness analysis may be useful in these situations, we consider it in greater depth in Chapter 18.

### The Relevance of CBA when Goals Other than Efficiency Matter

One goal, efficiency, underlies CBA. The general public, politicians, and even economists, however, very often consider goals reflecting other values to be relevant to the evaluation of public policies proposed to solve social problems. Although efficiency almost always is one of the relevant goals in policy analysis, other goals, such as equality of opportunity, equality of outcome, expenditure constraints, political feasibility, and national security, for instance, may be as, or even more, important. Indeed, the spenders and guardians we met in Chapter 1 behave as if they are responding to goals other than efficiency. When goals in addition to efficiency are relevant, as well as when efficiency is the only goal, but relevant impacts cannot be confidently monetized, *multigoal analysis* provides the appropriate framework. In the special case in which efficiency and equality of outcome are the only relevant goals, *distributionally weighted CBA* may be an appropriate technique.

**Multigoal Analysis.** The most general analytical framework is multigoal analysis. At the heart of multigoal analysis lies the notion that all policy alternatives should be compared in terms of all the relevant goals. Though multigoal analysis can be prescribed as a number of distinct steps,<sup>20</sup> three of its aspects are especially important. First, the analyst must move from relevant social values to general goals to specific impact categories that can be used as yardsticks for evaluating alternative policies. For example, the value

**TABLE 2-3** Evaluation Matrix Worksheet for Alternative Family Aid Policies

<i>Goals</i>	<i>Impact Categories</i>	<i>Policy Alternatives</i>		
		<i>Policy A (status quo)</i>	<i>Policy B</i>	<i>Policy C</i>
Efficiency	Labor earnings Investment in human capital Administrative costs			
Quality of life of poorest families	Number of fami- lies below poverty line Number of one-parent families Educational achievement of family members			
Political feasibility	Probability of adoption of required legislation			

of human dignity may imply a goal of improving equality of opportunity, which might be expressed as quantifiable impacts such as increasing participation in higher education and expanding workforce participation. Second, the analyst must evaluate each alternative policy, including the status quo, with respect to each of the impacts. Third, as no policy alternative is likely to dominate the others in terms of improvement in all the goals, the analyst usually can only make a recommendation to adopt one of the alternatives by carefully considering and making a subjective judgment concerning the trade-offs in the achievement of goals it offers relative to the other alternatives.

As a simple example, consider a multigoal analysis of alternative income transfer policies intended to help poor families. The analyst might construct the worksheet shown in Table 2-3 as a checklist for keeping track of the relevant goals. Increasing efficiency and improving the quality of life of poor families are both appropriate substantive goals. The former captures the aggregate gain or loss to society from transfers; the latter captures that portion of the gain or loss accruing to the poorest families. The goal of achieving political feasibility might be added to take account of the fact that a consensus on the relative importance of the substantive goals among politicians is unlikely. In this example, it can be thought of as an instrumental goal that is valuable not for its own sake but because it helps achieve the substantive goals. The major efficiency impacts are likely to be work disincentives for the recipients of aid and the real resource costs of administering the aid policy. If both of these impacts could be monetized, then the criterion for measuring efficiency would simply be the sum of the net benefits of these two impacts as measured in CBA. If either one of them could not be monetized, however, then efficiency would be stated in terms of the two impacts. The goal of improving the quality of life of poor families would probably be expressed in terms of such impacts as reducing the number of families below the poverty line, reducing the number of one-parent families, and increasing the educational achievement of family members.

The impact associated with the additional goal of political feasibility might be the change in the probability of passage of legislation required to implement the policy.

Before selecting among the alternative policies, the analyst should fill in all the cells of a matrix like the one shown in Table 2-3. Each cell would contain a prediction of the effect of a particular policy in terms of a particular impact category. By filling in all the cells, the analyst seeks to gain a comprehensive comparison of the alternatives across all the impact categories, and hence across their related goals.

Note that one can think of CBA, qualitative CBA, and cost-effectiveness analysis as special cases of multigoal analysis. In the case of CBA, there is one goal (efficiency) with one criterion (net benefits) so that the evaluation matrix has only one row, and the choice among alternatives is trivial (simply select the policy with the largest net benefits). In the case of qualitative CBA, there is also one goal, but because all relevant impacts cannot be monetized, it corresponds to several criteria, one for each impact. In the case of cost-effectiveness analysis, the goal of efficiency is often combined with some other goal such as satisfying a constraint on monetary costs or achieving some target level of reduction in the quantified but nonmonetized impact.

**Distributionally Weighted CBA.** If both efficiency and equality of income are relevant goals and their relative importance can be quantified, then distributionally weighted CBA provides an alternative decision rule to the maximization of net benefits. Instead of considering aggregate net benefits as in standard CBA, net benefits are calculated for each of several relevant groups distinguished by income, wealth, or some similar characteristic of relevance to a distributional concern. As discussed in more detail in Chapter 19, the net benefits of each group are multiplied by a weighting factor, selected by the analyst to reflect some distributional goal, and then summed to arrive at a number that can be used to rank alternative policies.

The major problem analysts encounter in doing distributionally weighted CBA is arriving at an appropriate and acceptable set of weights. One general approach, which takes as a desirable social goal increasing equality of wealth, involves making the weights inversely proportional to wealth (or income) to favor policies that tend to equalize wealth (or income) in the population.<sup>21</sup> Another general approach, which takes as a desirable social goal the raising of the position of the least advantaged in society, involves placing a higher weight on the net benefits of those with incomes or wealth below some threshold levels than on those with incomes or wealth above the threshold. As reasonable arguments can be made in support of each of these approaches, the absence of a consensus about appropriate weights is not surprising.<sup>22</sup>

Obviously, developing weights that allow a single quantitative criterion for ranking alternative policies makes the choice among policy alternatives easier. Yet this ease is achieved only by making an assumption that forces efficiency and equality of outcomes to be fully commensurate. Dissatisfaction with the strong assumption required to do this has led a number of analysts to suggest that distributionally weighted CBA should always be done in conjunction with standard CBA to make clearer the efficiency implications of the selected weights.<sup>23</sup> In doing so, the study becomes, in effect, a multigoal analysis, raising the question of whether an explicit treatment of efficiency and equality as separate goals might not be a more appropriate framework when both efficiency and distributional concerns are important. Cost-effectiveness analysis

## EXHIBIT 2-2

A study by the Congressional Budget Office to assess three alternatives for reducing U.S. consumption of gasoline listed the following criteria:

This study weighs the relative merits of tightening CAFE standards, raising the federal gasoline tax, and creating a cap-and-trade program against several major criteria:

*Cost-Effectiveness.* Reducing gasoline consumption would impose costs (both monetary and nonmonetary) on various producers and consumers. A cost-effective policy would keep those costs to a minimum.

*Predictability of Gasoline Savings.* How reliably would the policy bring about the desired reduction in gasoline consumption?

*Effects on Safety.* How would the policy alter the number and severity of traffic accidents?

*Effects on Other External Costs Related to Driving.* Reducing gasoline consumption would affect not only the United States' energy security and carbon emissions but other driving-related

external costs (ones whose full weight is borne by society at large rather than by an individual). Those external costs include traffic congestion, the need for highway construction and maintenance, and emissions of air pollutants besides carbon dioxide.

In addition to those factors, the three policy options would have other implications that policymakers may care about—such as their effects on people at different income levels and in different parts of the country and their impact on the amount of revenue collected by the federal government. (Summary, p.1)

One could imagine turning the analysis into a CBA by monetizing the effects on safety and the effects on other external costs related to driving and treating predictability of gasoline savings through sensitivity analysis. As monetizing the distributional concerns would be difficult, a multigoal analysis with the CBA assessing efficiency and a separate treatment of distributional impacts could be useful.

*Source:* Adapted from Congressional Budget Office, *Reducing Gasoline Consumption: Three Policy Options*, November 2002 ([www.cbo.gov/ftpdocs/39xx/doc3991/11-21-GasolineStudy.pdf](http://www.cbo.gov/ftpdocs/39xx/doc3991/11-21-GasolineStudy.pdf)).

might also provide a more reasonable approach than distributionally weighted CBA by posing the question in terms of achieving the most desirable redistribution possible for some fixed level of net cost.<sup>24</sup>

## CONCLUSION

CBA is a method for determining if proposed policies could potentially be Pareto improving: positive net benefits make it possible, in the sense of making available resources to compensate those who bear costs so that some people are made better off without making anyone else worse off. Willingness to pay and opportunity cost are the guiding principles for measuring costs and benefits. Much of the rest of this book deals with how to make use of the concepts in practice.

Part II sets out the conceptual foundations for CBA in more depth. It begins by reviewing supply and demand analytics (Chapter 3), and then uses these analytics to show how information revealed in markets can be used to estimate willingness to pay and opportunity cost (Chapters 4 and 5). Subsequent chapters consider how to take account of policy effects that accrue in different time periods (Chapters 6 and 10) or with

uncertainty (Chapters 7 and 8) and impacts on goods that have value beyond their direct use (Chapter 9).

Part III considers empirical approaches for predicting and valuing policy impacts. After considering general issues in prediction and valuation in Chapter 11, three chapters demonstrate estimation methods based on observed behavior. Policy demonstrations sometimes provide information of direct relevance to CBA, especially in terms of predicting policy impacts (Chapter 12). More often, however, analysts face the task of making inferences about demand schedules from data revealed in markets (Chapter 13) or through other behaviors (Chapter 14). In some circumstances, especially those arising in the context of environmental policy, the unavailability of data from observed behavior forces analysts to rely on survey responses as a basis for estimation (Chapter 15). Analysts can often make use of existing research to find values needed to monetize impacts as costs and benefits (Chapter 16), though some special care must be taken in transferring these values to CBAs of policies in developing countries (Chapter 17).

Part IV further develops two of the analytical methods introduced in this chapter that can be used as alternatives to CBA: cost-effectiveness analysis (Chapter 18) and distributionally weighted cost-benefit analysis (Chapter 19).

## EXERCISES FOR CHAPTER 2

1. Many experts claim that, although VHS came to dominate the video recorder market, Betamax was a superior technology. Assume that these experts are correct, so that, all other things equal, a world in which all video recorders were Betamax technology would be Pareto superior to a world in which all video recorders were VHS technology. Yet it seems implausible that a policy that forced a switch in technologies would be even potentially Pareto improving. Explain.
2. Let's explore the concept of willingness to pay with a thought experiment. Imagine a specific sporting, entertainment, or cultural event that you would very much like to attend—perhaps a World Cup match, the seventh game of the World Series, a Rolling Stones concert, or a Kathleen Battle performance.
  - a. What is the most you would be willing to pay for a ticket to the event?
  - b. Imagine that you won a ticket to the event in a lottery. What is the minimum amount of money that you would be willing to accept to give up the ticket?
  - c. Imagine that you had an income 50 percent higher than it is now, but that you didn't win a ticket to the event. What is the most you would be willing to pay for a ticket?
  - d. Do you know anyone who would sufficiently dislike the event that they would not use a free ticket unless they were paid to do so?
  - e. Do your answers suggest any possible generalizations about willingness to pay?
3. How closely do government expenditures measure opportunity cost for each of the following program inputs?
  - a. Time of jurors in a criminal justice program that requires more trials.
  - b. Land to be used for a nuclear waste storage facility that is owned by the government and located on a military base.
  - c. Labor for a reforestation program in a small rural community with high unemployment.
  - d. Labor of current government employees who are required to administer a new program.
  - e. Concrete that was previously poured as part of a bridge foundation.
4. Three mutually exclusive projects are being considered for a remote river valley: Project R, a recreational facility, has estimated benefits of \$10 million and costs of \$8 million; project F, a forest preserve with some recreational facilities, has estimated benefits of \$13 million and

costs of \$10 million; project W, a wilderness area with restricted public access, has estimated benefits of \$5 million and costs of \$1 million. In addition, a road could be built for a cost of \$4 million that would increase the benefits of project R by \$8 million, increase the benefits of project F by \$5 million, and reduce the benefits of project W by \$1 million. Even in the absence of any of the other projects, the road has estimated benefits of \$2 million.

- a. Calculate the benefit-cost ratio and net benefits for each possible alternative to the status quo. Note that there are seven possible alternatives to the status quo: R, F, and W, both with and without the road, and the road alone.
  - b. If only one of the seven alternatives can be selected, which should be selected according to the CBA decision rule?
5. An analyst for the U.S. Navy was asked to evaluate alternatives for forward-basing a destroyer flotilla. He decided to do the evaluation as a CBA. The major categories of costs were related to obtaining and maintaining the facilities. The major category of benefit was reduced sailing time to patrol routes. The analyst recommended the forward base with the largest net benefits. The admiral, his client, rejected the recommendation because the CBA did not include the risks to the forward bases from surprise attack and the risks of being unexpectedly ejected from the bases because of changes in political regimes of the host countries. Was the analyst's work wasted?
  6. Because of a recent wave of jewelry store robberies, a city increases police surveillance of jewelry stores. The increased surveillance costs the city an extra \$500,000 per year, but as a result, the amount of jewelry that is stolen falls. Specifically, without the increase in surveillance, jewelry with a retail value of \$1 million would have been stolen. This stolen jewelry would have been fenced by the jewelry thieves for \$600,000. What is the net social benefit resulting from the police surveillance program?
  7. (Spreadsheet recommended.) Excessive and improper use of antibiotics is contributing to the resistance of many diseases to existing antibiotics. Consider a regulatory program in the United States that would monitor antibiotic prescribing by physicians. Analysts estimate the direct costs of enforcement to be \$40 million, the time costs to doctors and health professionals to be \$220 million, and the convenience costs to patients to be \$180 million (all annually). The annual benefits of the program are estimated to be \$350 million in avoided resistance costs in the United States, \$70 million in health benefits in the United States from better compliance with prescriptions, and \$280 million in avoided resistance costs in the rest of the world. Does the program have positive net benefits from the national perspective? If not, what fraction of benefits accruing in the rest of the world would have to be counted for the program to have positive net benefits?

## NOTES

1. Unless otherwise stated, we intend efficiency to mean *allocative efficiency*, as defined in this section. A broader interpretation of efficiency, which we discuss in a later section, is the maximization of a specific social welfare function that explicitly ranks alternative allocations.
2. Nicholas Kaldor, "Welfare Propositions of Economics and Interpersonal Comparisons of Utility," *Economic Journal* 49(195) 1939, 549–552; and John R. Hicks, "The Valuation of the Social Income," *Economica* 7(26) 1940, 105–124. The principle can also be stated as suggested by Hicks: Adopt a policy if and only if it would not be in the self-interest of those who will lose to bribe those who will gain not to adopt it.
3. Those who are worse off in society may or may not have been the ones who have borne the net costs of public policies. This argument thus shifts the focus from fairness with respect to particular policies to the relative position of those in society who are worse off for whatever reason.