

3

Evaluating Trade-Offs: Benefit–Cost Analysis and Other Decision-Making Metrics

No sensible decision can be made any longer without taking into account not only the world as it is, but the world as it will be . . .

—Isaac Asimov, US science fiction novelist and scholar (1920–1992)

Introduction

In the last chapter we noted that economic analysis has both positive and normative dimensions. The normative dimension helps to separate the policies that make sense from those that don't. Since resources are limited, it is not possible to undertake all ventures that might appear desirable so making choices is inevitable.

Normative analysis can be useful in public policy in several different situations. It might be used, for example, to evaluate the desirability of a proposed new pollution control regulation or a proposal to preserve an area currently scheduled for development. In these cases the analysis helps to provide guidance on the desirability of a program before that program is put into place. In other contexts it might be used to evaluate how an already-implemented program has worked out in practice. Here the relevant question is: Would this be (or was this) a wise use of resources? In this chapter, we present and demonstrate the use of several decision-making metrics that can assist us in evaluating options.

Normative Criteria for Decision Making

Normative choices can arise in two different contexts. In the first context we need simply to choose among options that have been predefined, while in the second we try to find the optimal choice among all the possible choices.

Evaluating Predefined Options: Benefit–Cost Analysis

If you were asked to evaluate the desirability of some proposed action, you would probably begin by attempting to identify both the gains and the losses from that action. If the gains exceed the losses, then it seems natural to support the action.

That simple framework provides the starting point for the normative approach to evaluating policy choices in economics. Economists suggest that actions have both benefits and costs. If the benefits exceed the costs, then the action is desirable. On the other hand, if the costs exceed the benefits, then the action is not desirable.

We can formalize this in the following way. Let B be the benefits from a proposed action and C be the costs. Our decision rule would then be

If $B > C$, support the action.

Otherwise, oppose the action.¹

As long as B and C are positive, a mathematically equivalent formulation would be

If $B/C > 1$, support the action.

Otherwise, oppose the action.

So far so good, but how do we measure benefits and costs? In economics the system of measurement is anthropocentric, which simply means human centered. All benefits and costs are valued in terms of their effects (broadly defined) on humanity. As shall be pointed out later, that does *not* imply (as it might first appear) that ecosystem effects are ignored unless they *directly* affect humans. The fact that large numbers of humans contribute voluntarily to organizations that are dedicated to environmental protection provides ample evidence that humans place a value on environmental preservation that goes well beyond any direct use they might make of it. Nonetheless, the notion that humans are doing the valuing is a controversial point that will be revisited and discussed in Chapter 4 along with the specific techniques for valuing these effects.

In benefit–cost analysis, benefits are measured simply as the relevant area under the demand curve since the demand curve reflects consumers' willingness to pay. Total costs are measured by the relevant area under the marginal cost curve.

It is important to stress that environmental services have costs even though they are produced without any human input. All costs should be measured as opportunity costs. As presented in Example 3.1, the *opportunity cost* for using resources in a new or an alternative way is the net benefit lost when specific environmental services are foregone in the conversion to the new use. The notion that it is costless to convert a forest to a new use is obviously wrong if valuable ecological or human services are lost in the process.

To firm up this notion of opportunity cost, consider another example. Suppose a particular stretch of river can be used either for white-water canoeing or to generate electric power. Since the dam that generates the power would flood the rapids, the two uses are incompatible. The opportunity cost of producing power is the foregone net benefit that would have resulted from the white-water canoeing. The *marginal opportunity cost curve* defines the additional cost of producing another unit of electricity resulting from the associated incremental loss of net benefits due to reduced opportunities for white-water canoeing.

¹Actually if $B = C$, it wouldn't make any difference if the action occurs or not; the benefits and costs are a wash.

EXAMPLE
3.1

Valuing Ecological Services from Preserved Tropical Forests

As Chapter 12 makes clear, one of the main threats to tropical forests is the conversion of forested land to some other use (agriculture, residences, and so on). Whether economic incentives favor conversion of the land depends upon the magnitude of the value that would be lost through conversion. How large is that value? Is it large enough to support preservation?

A group of ecologists investigated this question for a specific set of tropical forest fragments in Costa Rica. They chose to value one specific ecological service provided by the local forest: wild bees using the nearby tropical forest as a habitat provided pollination services to aid coffee production. While this coffee (*C. arabica*) can self-pollinate, pollination from wild bees has been shown to increase coffee productivity from 15 to 50 percent.

When the authors placed an economic value on this particular ecological service, they found that the pollination services from two specific preserved forest fragments (46 and 111 hectares, respectively) were worth approximately \$60,000 per year for one large, nearby Costa Rican coffee farm. As the authors conclude:

The value of forest in providing crop pollination service alone is . . . of at least the same order [of magnitude] as major competing land uses, and infinitely greater than that recognized by most governments (i.e., zero).

These estimates only partially capture the value of this forest because they consider only a single farm and a single type of ecological service. (This forest also provides carbon storage and water purification services, for example, and these were not included in the calculation.) Despite their partial nature, however, these calculations already begin to demonstrate the economic value of preserving the forest, even when considering only a limited number of specific instrumental values.

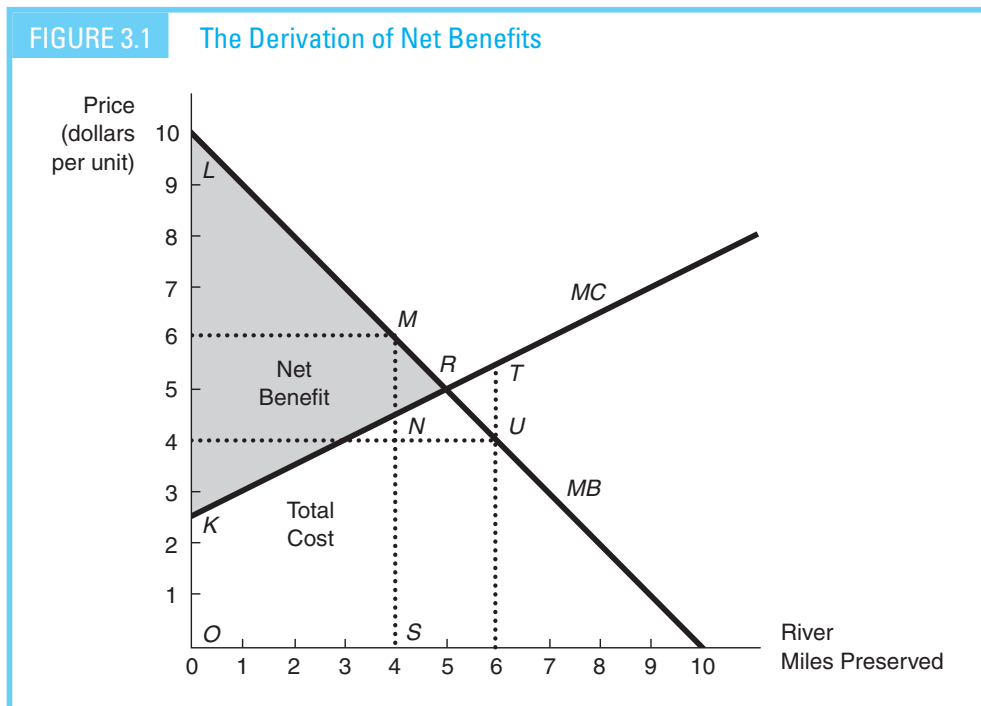
Source: Taylor H. Ricketts et al., “Economic Value of Tropical Forest to Coffee Production.” PNAS (Proceedings of the National Academy of Science), Vol. 101, No. 34, August 24, 2002, pp. 12579–12582.

Since net benefit is defined as the excess of benefits over costs, it follows that net benefit is equal to that portion of the area under the demand curve that lies above the supply curve.

Consider Figure 3.1, which illustrates the net benefits from preserving a stretch of river. Let’s use this example to illustrate the use of the decision rules introduced earlier. For example, let’s suppose that we are considering preserving a four-mile stretch of river and that the benefits and costs of that action are reflected in Figure 3.1. Should that stretch be preserved? Why or why not? We will return to this example later.

Finding the Optimal Outcome

In the preceding section we examined how benefit–cost analysis can be used to evaluate the desirability of specific actions. In this section we want to examine how this approach can be used to identify “optimal” or best approaches.



In subsequent chapters, which address individual environmental problems, the normative analysis will proceed in three steps. First we will identify an optimal outcome. Second we will attempt to discern the extent to which our institutions produce optimal outcomes and, where divergences occur between actual and optimal outcomes, to attempt to uncover the behavioral sources of the problems. Finally we can use both our knowledge of the nature of the problems and their underlying behavioral causes as a basis for designing appropriate policy solutions. Although applying these three steps to each of the environmental problems must reflect the uniqueness of each situation, the overarching framework used to shape that analysis is the same.

To provide some illustrations of how this approach is used in practice, consider two examples: one drawn from natural resource economics and another from environmental economics. These are meant to be illustrative and to convey a flavor of the argument; the details are left to upcoming chapters.

Consider the rising number of depleted ocean fisheries. Depleted fisheries, which involve fish populations that have fallen so low as to threaten their viability as commercial fisheries, not only jeopardize oceanic biodiversity, but also pose a threat to both the individuals who make their living from the sea and the communities that depend on fishing to support their local economies.

How would an economist attempt to understand and resolve this problem? The first step would involve defining the optimal stock or the optimal rate of harvest of the fishery. The second step would compare this level with the actual stock and harvest levels. Once this economic framework is applied, not only does it become clear that stocks are much lower than optimal for many fisheries, but also the reason

for excessive exploitation becomes clear. Understanding the nature of the problem has led quite naturally to some solutions. Once implemented, these policies have allowed some fisheries to begin the process of renewal. The details of this analysis and the policy implications that flow from it are covered in Chapter 13.

Another problem involves solid waste. As local communities run out of room for landfills in the face of an increasing generation of waste, what can be done?

Economists start by thinking about how one would define the optimal amount of waste. The definition necessarily incorporates waste reduction and recycling as aspects of the optimal outcome. The analysis not only reveals that current waste levels are excessive, but also suggests some specific behavioral sources of the problem. Based upon this understanding, specific economic solutions have been identified and implemented. Communities that have adopted these measures have generally experienced lower levels of waste and higher levels of recycling. The details are spelled out in Chapter 8.

In the rest of the book, similar analysis is applied to population, energy, minerals, agriculture, air and water pollution, and a host of other topics. In each case the economic analysis helps to point the way toward solutions. To initiate that process we must begin by defining “optimal.”

Relating Optimality to Efficiency

According to the normative choice criterion introduced earlier in this chapter, desirable outcomes are those where the benefits exceed the costs. It is therefore a logical next step to suggest that optimal policies are those that maximize net benefits (benefits–costs). The concept of *static efficiency*, or merely *efficiency*, was introduced in Chapter 2. An allocation of resources is said to satisfy the static efficiency criterion if the economic surplus from the use of those resources is maximized by that allocation. Notice that the net benefits area to be maximized in an “optimal outcome” for public policy is identical to the “economic surplus” that is maximized in an efficient allocation. Hence efficient outcomes are also optimal outcomes.

Let’s take a moment to show how this concept can be applied. Previously we asked whether an action that preserved four miles of river was worth doing (Figure 3.1). The answer was yes because the net benefits from that action were positive.

Static efficiency, however, requires us to ask a rather different question, namely, what is the optimal (or efficient) number of miles to be preserved? We know from the definition that the optimal amount of preservation would maximize net benefits. Does preserving four miles maximize net benefits? Is it the efficient outcome?

We can answer that question by establishing whether it is possible to increase the net benefit by preserving more or less of the river. If the net benefit can be increased by preserving more miles, clearly, preserving four miles could not have maximized the net benefit and, therefore, could not have been efficient.

Consider what would happen if society were to choose to preserve five miles instead of four. Refer back to Figure 3.1. What happens to the net benefit? It increases by area *MNR*. Since we can find another allocation with greater net benefit, four miles of preservation could not have been efficient. Could five? Yes. Let’s see why.

We know that five miles of preservation convey more net benefits than four. If this allocation is efficient, then it must also be true that the net benefit is smaller for levels of preservation higher than five. Notice that the additional cost of preserving the sixth unit (the area under the marginal cost curve) is larger than the additional benefit received from preserving it (the corresponding area under the demand curve). Therefore, the triangle *RTU* represents the reduction in net benefit that occurs if six miles are preserved rather than five.

Since the net benefit is reduced, both by preserving less than five and by preserving more than five, we conclude that five units is the preservation level that maximizes net benefit (the shaded area). Therefore, from our definition, preserving five miles constitutes an efficient or optimal allocation.²

One implication of this example, which will be very useful in succeeding chapters, is what we shall call the “first equimarginal principle”:

First Equimarginal Principle (the “Efficiency Equimarginal Principle”): Social net benefits are maximized when the social marginal benefits from an allocation equal the social marginal costs.

The social marginal benefit is the increase in social benefits received from supplying one more unit of the good or service, while social marginal cost is the increase in cost incurred from supplying one more unit of the good or service.

This criterion helps to minimize wasted resources, but is it fair? The ethical basis for this criterion is derived from a concept called *Pareto optimality*, named after the Italian-born Swiss economist Vilfredo Pareto, who first proposed it around the turn of the twentieth century.

Allocations are said to be Pareto optimal if no other feasible allocation could benefit at least one person without any deleterious effects on some other person.

Allocations that do not satisfy this definition are suboptimal. Suboptimal allocations can always be rearranged so that some people can gain net benefits without the rearrangement causing anyone else to lose net benefits. Therefore, the gainers could use a portion of their gains to compensate the losers sufficiently to ensure they were at least as well off as they were prior to the reallocation.

Efficient allocations are Pareto optimal. Since net benefits are maximized by an efficient allocation, it is not possible to increase the net benefit by rearranging the allocation. Without an increase in the net benefit, it is impossible for the gainers to compensate the losers sufficiently; the gains to the gainers would necessarily be smaller than the losses to the losers.

Inefficient allocations are judged inferior because they do not maximize the size of the pie to be distributed. By failing to maximize net benefit, they are forgoing an opportunity to make some people better off without harming others.

²The monetary worth of the net benefit is the sum of two right triangles, and it equals $(1/2)(\$5)(5) + (1/2)(\$2.50)(5)$ or \$18.75. Can you see why?

Comparing Benefits and Costs Across Time

The analysis we have covered so far is very useful for thinking about actions where time is not an important factor. Yet many of the decisions made now have consequences that persist well into the future. Time is a factor. Exhaustible energy resources, once used, are gone. Biological renewable resources (such as fisheries or forests) can be overharvested, leaving smaller and possibly weaker populations for future generations. Persistent pollutants can accumulate over time. How can we make choices when the benefits and costs may occur at different points in time?

Incorporating time into the analysis requires an extension of the concepts we have already developed. This extension provides a way for thinking not only about the magnitude of benefits and costs, but also about their timing. In order to incorporate timing, the decision rule must provide a way to compare net benefits received in different time periods. The concept that allows this comparison is called *present value*. Therefore, before introducing this expanded decision rule, we must define present value.

Present value explicitly incorporates the time value of money. A dollar today invested at 10 percent interest yields \$1.10 a year from now (the return of the \$1 principal plus \$0.10 interest). The present value of \$1.10 received one year from now is therefore \$1, because given \$1 now, you can turn it into \$1.10 a year from now by investing it at 10 percent interest. We can find the present value of any amount of money (X) received one year from now by computing $X/(1 + r)$, where r is the appropriate interest rate (10 percent in our above example).

What could your dollar earn in two years at r percent interest? Because of compound interest, the amount would be $\$1(1 + r)(1 + r) = \$1(1 + r)^2$. It follows then that the present value of X received two years from now is $X/(1 + r)^2$.

By now the pattern should be clear. The present value of a *one-time* net benefit received n years from now is

$$PV[B_n] = \frac{B_n}{(1 + r)^n}$$

The present value of a stream of net benefits $\{B_0, \dots, B_n\}$ received over a period of n years is computed as

$$PV[B_0, \dots, B_n] = \sum_{i=0}^n \frac{B_i}{(1 + r)^i}$$

where r is the appropriate interest rate and B_0 is the amount of net benefits received immediately. The process of calculating the present value is called *discounting*, and the rate r is referred to as the discount rate.

The number resulting from a present-value calculation has a straightforward interpretation. Suppose you were investigating an allocation that would yield the following pattern of net benefits on the last day of each of the next five years: \$3,000,

Year	1	2	3	4	5	Sum
Annual Amounts	\$3,000	\$5,000	\$6,000	\$10,000	\$12,000	\$36,000
Present Value ($r = 0.06$)	\$2,830.19	\$4,449.98	\$5,037.72	\$7,920.94	\$8,967.10	\$29,205.92

Year	1	2	3	4	5	6
Balance at Beginning of Year	\$29,205.92	\$27,958.28	\$24,635.77	\$20,113.92	\$11,320.75	\$0.00
Year-End Fund Balance before Payment ($r = 0.06$)	\$30,958.28	\$29,635.77	\$26,113.92	\$21,320.75	\$12,000.00	
Payment	\$3,000	\$5,000	\$6,000	\$10,000	\$12,000	

\$5,000, \$6,000, \$10,000, and \$12,000. If you use an interest rate of 6 percent ($r = 0.06$) and the above formula, you will discover that this stream has a present value of \$29,205.92 (see Table 3.1). Notice how each amount is discounted back the appropriate number of years to the present and then these discounted values are summed.

What does that number mean? If you put \$29,205.92 in a savings account earning 6 percent interest and wrote yourself checks, respectively, for \$3,000, \$5,000, \$6,000, \$10,000, and \$12,000 on the last day of each of the next five years, your last check would just restore the account to a \$0 balance (see Table 3.2). Thus, you should be indifferent about receiving \$29,205.92 now or in the specific five-year stream of benefits totaling \$36,000; given one, you can get the other. Hence, the method is called present value because it translates everything back to its current worth.

It is now possible to show how this analysis can be used to evaluate actions. Calculate the present value of net benefits from the action. If the present value is greater than zero, the action should be supported. Otherwise it should not.

Dynamic Efficiency

The static efficiency criterion is very useful for comparing resource allocations when time is not an important factor. How can we think about optimal choices when the benefits and costs occur at different points in time?

The traditional criterion used to find an optimal allocation when time is involved is called *dynamic efficiency*, a generalization of the static efficiency concept already developed. In this generalization, the present-value criterion provides a way for comparing the net benefits received in one period with the net benefits received in another.

An allocation of resources across n time periods satisfies the dynamic efficiency criterion if it maximizes the present value of net benefits that could be received from all the possible ways of allocating those resources over the n periods.

Applying the Concepts

Having now spent some time developing the concepts we need, let's take a moment to examine some actual studies in which they have been used.

Pollution Control

Benefit–cost analysis has been used to assess the desirability of efforts to control pollution. Pollution control certainly confers many benefits, but it also has costs. Do the benefits justify the costs? That was a question the U.S. Congress wanted answered, so in Section 812 of the Clean Air Act Amendments of 1990 it required the U.S. Environmental Protection Agency (EPA) to evaluate the benefits and costs of the U.S. air pollution control policy initially over the 1970–1990 period and subsequently over the 1990–2020 time period (see Example 3.2).

EXAMPLE 3.2

Does Reducing Pollution Make Economic Sense? Evidence from the Clean Air Act

In its 1997 report to Congress, the EPA presented the results of its attempt to discover whether the Clean Air Act had produced positive net benefits over the period 1970–1990. The results suggested that the present value of benefits (using a discount rate of 5 percent) was \$22.2 trillion, while the costs were \$0.523 trillion. Performing the necessary subtraction reveals that the net benefits were therefore equal to \$21.7 trillion. According to this study, U.S. air pollution control policy during this period made very good economic sense.

Soon after the period covered by this analysis, substantive changes were made in the Clean Air Act Amendments of 1990 (the details of those changes are covered in later chapters). Did those additions also make economic sense?

In August of 2010, the U.S. EPA issued a report of the benefits and costs of the Clean Air Act from 1990 to 2020. This report suggests that the costs of meeting the 1990 Clean Air Act Amendment requirements are expected to rise to approximately \$65 billion per year by 2020 (2006 dollars). Almost half of the compliance costs (\$28 billion) arise from pollution controls placed on cars, trucks, and buses, while another \$10 billion arises from reducing air pollution from electric utilities.

These actions are estimated to cause benefits (from reduced pollution damage) to rise from roughly \$800 billion in 2000 to almost \$1.3 trillion in 2010, ultimately reaching approximately \$2 trillion per year (2006 dollars) by 2020! For persons living in the United States, a cost of approximately \$200 per person by 2020 produces approximately a \$6,000 gain in benefits from the improvement in air quality. Many of the estimated benefits come from reduced risk of early mortality due to exposure to fine particulate matter. Table 3.3 provides a summary of the costs and benefits and includes a calculation of the benefit/cost ratio.

TABLE 3.3 Summary Comparison of Benefits and Costs from the Clean Air Act-1990–2020 (Estimates in Million 2006\$)

	Annual Estimates			Present Value Estimate
	2000	2010	2020	1990–2020
Monetized Direct Costs:				
Low ¹				
Central	\$20,000	\$53,000	\$65,000	\$380,000
High ¹				
Monetized Direct Benefits:				
Low ²				
Central	\$90,000	\$160,000	\$250,000	\$1,400,000
Central	\$770,000	\$1,300,000	\$2,000,000	\$12,000,000
High ²	\$2,300,000	\$3,800,000	\$5,700,000	\$35,000,000
Net Benefits:				
Low				
Central	\$70,000	\$110,000	\$190,000	\$1,000,000
Central	\$750,000	\$1,200,000	\$1,900,000	\$12,000,000
High	\$2,300,000	\$3,700,000	\$5,600,000	\$35,000,000
Benefit/Cost Ratio:				
Low ³				
Central	5/1	3/1	4/1	4/1
Central	39/1	25/1	31/1	32/1
High ³	115/1	72/1	88/1	92/1

¹The cost estimates for this analysis are based on assumptions about future changes in factors such as consumption patterns, input costs, and technological innovation. We recognize that these assumptions introduce significant uncertainty into the cost results; however, the degree of uncertainty or bias associated with many of the key factors cannot be reliably quantified. Thus, we are unable to present specific low and high cost estimates.

²Low and high benefit estimates are based on primary results and correspond to 5th and 95th percentile results from statistical uncertainty analysis, incorporating uncertainties in physical effects and valuation steps of benefits analysis. Other significant sources of uncertainty not reflected include the value of unquantified or unmonetized benefits that are not captured in the primary estimates and uncertainties in emissions and air quality modeling.

³The low benefit/cost ratio reflects the ratio of the low benefits estimate to the central costs estimate, while the high ratio reflects the ratio of the high benefits estimate to the central costs estimate. Because we were unable to reliably quantify the uncertainty in cost estimates, we present the low estimate as “less than X,” and the high estimate as “more than Y,” where X and Y are the low and high benefit/cost ratios, respectively.

Sources: U.S. Environmental Protection Agency, THE BENEFITS AND COSTS OF THE CLEAN AIR ACT, 1970 to 1990 (Washington, DC: Environmental Protection Agency, 1997), Table 18, p. 56; and the U.S. Environmental Protection Agency Office of Air and Radiation, THE BENEFITS AND COSTS OF THE CLEAN AIR ACT, 1990 to 2020 – Summary Report, 8/16/2010 and Full Report available at <http://www.epa.gov/oar/sect812/prospective2.html> (accessed on 12/31/2010).

In responding to this congressional mandate, the EPA set out to quantify and monetize the benefits and costs of achieving the emissions reductions required by U.S. policy. Benefits quantified by this study included reduced death rates and lower incidences of chronic bronchitis, lead poisoning, strokes, respiratory diseases, and heart disease as well as the benefits of better visibility, reduced structural damages, and improved agricultural productivity.

We shall return to this study later in the book for a deeper look at how these estimates were derived, but a couple of comments are relevant now. First, despite the fact that this study did not attempt to value all pollution damage to ecosystems that was avoided by this policy, the net benefits are still strongly positive. While presumably the case for controlling pollution would have been even stronger had all such avoided damage been included, the desirability of this form of control is evident even with only a partial consideration of benefits. An inability to monetize everything does not necessarily jeopardize the ability to reach sound policy conclusions.

Although these results justify the conclusion that pollution control made economic sense, they do not justify the stronger conclusion that the policy was efficient. To justify that conclusion, the study would have had to show that the present value of net benefits was maximized, not merely positive. In fact, this study did not attempt to calculate the maximum net benefits outcome and if it had, it would have almost certainly discovered that the policy during this period was not optimal. As we shall see in Chapters 15 and 16, the costs of the chosen policy approach were higher than necessary to achieve the desired emissions reductions. With an optimal policy mix, the net benefits would have been even higher.

Preservation versus Development

One of the most basic conflicts faced by environmental policy occurs when a currently underdeveloped but ecologically significant piece of land becomes a candidate for development. If developed, the land may not only provide jobs for workers, wealth for owners, and goods for consumers, but also it may degrade the ecosystem, possibly irreversibly. Wildlife habitat may be eliminated, wetlands may be paved over, and recreational opportunities may be gone forever. On the other hand, if the land were preserved, the specific ecosystem damages caused by development could be prevented, but the opportunity for increased income and employment provided by development would have been lost. These conflicts become intensified if unemployment rates in the area are high and the local ecology is rather unique.

One such conflict arose in Australia from a proposal to mine a piece of land in an area known as the Kakadu Conservation Zone (KCZ). Decision makers at that time had to decide whether it should be mined or preserved. One way to examine that question is to use the techniques above to examine the net benefits of the two alternatives (see Example 3.3).

Choosing between Preservation and Development in Australia

EXAMPLE

3.3

The Kakadu Conservation Zone, a 50-square-kilometer area lying entirely within the Kakadu National Park (KNP), was initially set aside by the government as part of a grazing lease. The current issue was whether it should be mined (it was believed to contain significant deposits of gold, platinum, and palladium) or added to the KNP, one of Australia's major parks. In recognition of its unique ecosystem and extensive wildlife as well as its aboriginal archeological sites, much of the park has been placed on the U.N. World Heritage List.

Mining would produce income and employment, but it could also cause the ecosystems in both the KCZ and KNP to experience irreversible damage. What value was to be placed on those risks? Would those risks outweigh the employment and income effects from mining?

To provide answers to these crucial questions, economists conducted a benefit–cost analysis using a technique known as contingent valuation. (We shall go into some detail about how this technique works in Chapter 4, but for now it can suffice to note that this is a technique for eliciting “willingness-to-pay” information.) The value of preserving the site was estimated to be A\$435 million, while the present value of mining the site was estimated to be A\$102 million.

According to this analysis, preservation was the preferred option and it was the option chosen by the government.

Source: Richard T. Carson, Leanne Wilks, and David Imber, “Valuing the Preservation of Australia’s Kakadu Conservation Zone.” OXFORD ECONOMIC PAPERS, Vol. 46, Supplement (1994), pp. 727–749.

Issues in Benefit Estimation

The analyst charged with the responsibility for performing a benefit–cost analysis encounters many decision points requiring judgment. If we are to understand benefit–cost analysis, the nature of these judgments must be clear in our minds.

Primary versus Secondary Effects. Environmental projects usually trigger both primary and secondary consequences. For example, the primary effect of cleaning a lake will be an increase in recreational uses of the lake. This primary effect will cause a further ripple effect on services provided to the increased number of users of the lake. Are these secondary benefits to be counted?

The answer depends upon the employment conditions in the surrounding area. If this increase in demand results in employment of previously unused resources, such as labor, the value of the increased employment should be counted. If, on the other hand, the increase in demand is met by a shift in previously employed resources from one use to another, it is a different story. In general, secondary employment benefits should be counted in high unemployment areas or when the particular skills demanded are underemployed at the time the project is commenced.

This should not be counted when the project simply results in a rearrangement of productively employed resources.

Accounting Stance. The accounting stance refers to the geographic scale at which the benefits are measured. Who benefits? If a proposed project is funded by a national government, but benefits a local or regional area, a benefit–cost analysis will look quite different depending on whether the analysis is done at the regional or national scale.

With and Without Principle. The “with and without” principle states that only those benefits that would result from the project should be counted, ignoring those that would have accrued anyway. Mistakenly including benefits that would have accrued anyway would overstate the benefits of the program.

Tangible versus Intangible Benefits. *Tangible benefits* are those that can reasonably be assigned a monetary value. *Intangible benefits* are those that cannot be assigned a monetary value, either because data are not available or reliable enough or because it is not clear how to measure the value even with data.³ Quantification of intangible benefits is the primary topic of the next chapter.

How are intangible benefits to be handled? One answer is perfectly clear: They should not be ignored. To ignore intangible benefits is to bias the results. That benefits are intangible does not mean they are unimportant.

Intangible benefits should be quantified to the fullest extent possible. One frequently used technique is to conduct a sensitivity analysis of the estimated benefit values derived from less than perfectly reliable data. We can determine, for example, whether or not the outcome is sensitive, within wide ranges, to the value of this benefit. If not, then very little time has to be spent on the problem. If the outcome is sensitive, the person or persons making the decision bear the ultimate responsibility for weighing the importance of that benefit.

Approaches to Cost Estimation

Estimating costs is generally easier than estimating benefits, but it is not easy. One major problem for both derives from the fact that benefit–cost analysis is forward-looking and thus requires an estimate of what a particular strategy *will* cost, which is much more difficult than tracking down what an existing strategy *does* cost.

Two approaches have been developed to estimate these costs.

The Survey Approach. One way to discover the costs associated with a policy is to ask those who bear the costs, and presumably know the most about them, to reveal the magnitude of the costs to policy-makers. Polluters, for example, could be

³The division between tangible and intangible benefits changes as our techniques improve. Recreation benefits were, until the advent of the travel-cost model, treated as intangible. The travel cost model will be discussed in the next chapter.

asked to provide control-cost estimates to regulatory bodies. The problem with this approach is the strong incentive not to be truthful. An overestimate of the costs can trigger less stringent regulation; therefore, it is financially advantageous to provide overinflated estimates.

The Engineering Approach. The engineering approach bypasses the source being regulated by using general engineering information to catalog the possible technologies that could be used to meet the objective and to estimate the costs of purchasing and using those technologies. The final step in the engineering approach is to assume that the sources would use technologies that minimize cost. This produces a cost estimate for a “typical,” well-informed firm.

The engineering approach has its own problems. These estimates may not approximate the actual cost of any particular firm. Unique circumstances may cause the costs of that firm to be higher, or lower, than estimated; the firm, in short, may not be typical.

The Combined Approach. To circumvent these problems, analysts frequently use a combination of survey and engineering approaches. The survey approach collects information on possible technologies, as well as special circumstances facing the firm. Engineering approaches are used to derive the actual costs of those technologies, given the special circumstances. This combined approach attempts to balance information best supplied by the source with that best derived independently.

In the cases described so far, the costs are relatively easy to quantify and the problem is simply finding a way to acquire the best information. This is not always the case, however. Some costs are not easy to quantify, although economists have developed some ingenious ways to secure monetary estimates even for those costs.

Take, for example, a policy designed to conserve energy by forcing more people to carpool. If the effect of this is simply to increase the average time of travel, how is this cost to be measured?

For some time, transportation analysts have recognized that people value their time, and quite a literature has now evolved to provide estimates of how valuable time savings or time increases would be. The basis for this valuation is opportunity cost—how the time might be used if it weren’t being consumed in travel. Although the results of these studies depend on the amount of time involved, individuals seem to value their travel time at a rate not more than half their wage rates.

The Treatment of Risk

For many environmental problems, it is not possible to state with certainty what consequences a particular policy will have, because scientific estimates themselves often are imprecise. Determining the efficient exposure to potentially toxic substances requires obtaining results at high doses and extrapolating to low doses, as well as extrapolating from animal studies to humans. It also requires relying upon epidemiological studies that infer a pollution-induced adverse human health impact from correlations between indicators of health in human populations and recorded pollution levels.

For example, consider the potential damages from climate change. While most scientists now agree on the potential impacts of climate change, such as sea level rise and species losses, the timing and extent of those losses are not certain.

The treatment of risk in the policy process involves two major dimensions: (1) identifying and quantifying the risks; and (2) deciding how much risk is acceptable. The former is primarily scientific and descriptive, while the latter is more evaluative or normative.

Benefit–cost analysis grapples with the evaluation of risk in several ways. Suppose we have a range of policy options A, B, C, D and a range of possible outcomes E, F, G for each of these policies depending on how the economy evolves over the future. These outcomes, for example, might depend on whether the demand growth for the resource is low, medium, or high. Thus, if we choose policy A , we might end up with outcomes AE, AF , or AG . Each of the other policies has three possible outcomes as well, yielding a total of 12 possible outcomes.

We could conduct a separate benefit–cost analysis for each of the 12 possible outcomes. Unfortunately, the policy that maximizes net benefits for E may be different from that which maximizes net benefits for F or G . Thus, if we only knew which outcome would prevail, we could select the policy that maximized net benefits; the problem is that we do not. Furthermore, choosing the policy that is best if outcome E prevails may be disastrous if G results instead.

When a dominant policy emerges, this problem is avoided. A *dominant policy* is one that confers higher net benefits for every outcome. In this case, the existence of risk concerning the future is not relevant for the policy choice. This fortuitous circumstance is exceptional rather than common, but it can occur.

Other options exist even when dominant solutions do not emerge. Suppose, for example, that we were able to assess the likelihood that each of the three possible outcomes would occur. Thus, we might expect outcome E to occur with probability 0.5, F with probability 0.3, and G with probability 0.2. Armed with this information, we can estimate the expected present value of net benefits. The *expected present value of net benefits* for a particular policy is defined as the sum over outcomes of the present value of net benefits for that policy where each outcome is weighted by its probability of occurrence. Symbolically this is expressed as

$$EPVNB_j = \sum_{i=0}^I P_i PVNB_{ij}, \quad j = 1, \dots, \mathcal{J}, \quad (3.1)$$

where

$EPVNB_j$ = expected present value of net benefits for policy j

P_i = probability of the i th outcome occurring

$PVNB_{ij}$ = present value of net benefits for policy j if outcome i prevails

\mathcal{J} = number of policies being considered

I = number of outcomes being considered

The final step is to select the policy with the highest expected present value of net benefits.

This approach has the substantial virtue that it weighs higher probability outcomes more heavily. It also, however, makes a specific assumption about society's preference for risk. This approach is appropriate if society is risk-neutral. *Risk-neutrality* can be defined most easily by the use of an example. Suppose you were allowed to choose between being given a definite \$50 or entering a lottery in which you had a 50 percent chance of winning \$100 and a 50 percent chance of winning nothing. (Notice that the expected value of this lottery is $\$50 = 0.5(\$100) + 0.5(\$0)$.) You would be said to be risk-neutral if you would be indifferent between these two choices. If you view the lottery as more attractive, you would be exhibiting *risk-loving* behavior, while a preference for the definite \$50 would suggest *risk-averse* behavior. Using the expected present value of net benefits approach implies that society is risk-neutral.

Is that a valid assumption? The evidence is mixed. The existence of gambling suggests that at least some members of society are risk-loving, while the existence of insurance suggests that, at least for some risks, others are risk-averse. Since the same people may gamble and own insurance policies, it is likely that the type of risk may be important.

Even if individuals were demonstrably risk-averse, this would not be a sufficient condition for the government to forsake risk-neutrality in evaluating public investments. One famous article (Arrow and Lind, 1970) argues that risk-neutrality is appropriate since "when the risks of a public investment are publicly borne, the total cost of risk-bearing is insignificant and, therefore, the government should ignore uncertainty in evaluating public investments." The logic behind this result suggests that as the number of risk bearers (and the degree of diversification of risks) increases, the amount of risk borne by any individual diminishes to zero.

When the decision is irreversible, as demonstrated by Arrow and Fisher (1974), considerably more caution is appropriate. Irreversible decisions may subsequently be regretted, but the option to change course will be lost forever. Extra caution also affords an opportunity to learn more about alternatives to this decision and its consequences before acting. Isn't it comforting to know that occasionally procrastination can be optimal?

There is a movement in national policy in both the courts and the legislature to search for imaginative ways to define acceptable risk. In general, the policy approaches reflect a case-by-case method. We shall see that current policy reflects a high degree of risk aversion toward a number of environmental problems.

Distribution of Benefits and Costs

Many agencies are now required to consider the distributional impacts of costs and benefits as part of any economic analysis. For example, the U.S. EPA provides guidelines on distributional issues in its "Guidelines for Preparing Economic Analysis." According to the EPA, distributional analysis "assesses changes in social welfare by examining the effects of a regulation across different sub-populations and entities." Distributional analysis can take two forms: *economic impact analysis* and *equity analysis*. Economic impact analysis focuses on a broad characterization of who gains and who loses from a given policy. Equity analysis examines impacts

on disadvantaged groups or sub-populations. The latter delves into the normative issue of equity or fairness in the distribution of costs and benefits. The issue of environmental justice will be considered further in Chapter 19.⁴ Loomis (2011) outlines several approaches for incorporating distribution and equity into benefit–cost analysis.

Choosing the Discount Rate

The discount rate can be defined conceptually as the social opportunity cost of capital. This cost of capital can be divided further into two components: (1) the riskless cost of capital and (2) the risk premium. The choice of the discount rate can influence policy decisions. Recall that discounting allows us to compare all costs and benefits in current dollars, regardless of when the benefits accrue or costs are charged. Suppose, a project will impose an immediate cost of \$4,000,000 (today's dollars), but the \$5,500,000 benefits will not be earned until 5 years out. Is this project a good idea? On the surface it might seem like it is, but recall that \$5,500,000 in 5 years is not the same as \$5,500,000 today. At a discount rate of 5 percent, the present value of benefits minus the present value of costs is positive. However, at a 10 percent discount rate, this same calculation yields a negative value, since the present value of costs exceeds the benefits. Can you reproduce the calculations that yield these conclusions?

As Example 3.4 indicates, this has been, and continues to be, an important issue. When the public sector uses a discount rate lower than that in the private sector, the public sector will find more projects with longer payoff periods worthy of authorization. And, as we have already seen, the discount rate is a major determinant of the allocation of resources among generations as well.

Traditionally, economists have used long-term interest rates on government bonds as one measure of the cost of capital, adjusted by a risk premium that would depend on the riskiness of the project considered. Unfortunately, the choice of how large an adjustment to make has been left to the discretion of the analysts. This ability to affect the desirability of a particular project or policy by the choice of discount rate led to a situation in which government agencies were using a variety of discount rates to justify programs or projects they supported. One set of hearings conducted by Congress during the 1960s discovered that, at one time, agencies were using discount rates ranging from 0 to 20 percent.

During the early 1970s the Office of Management and Budget published a circular that required, with some exceptions, all government agencies to use a discount rate of 10 percent in their benefit–cost analysis. A revision issued in 1992 reduced the required discount rate to 7 percent. This circular also includes guidelines for benefit–cost analysis and specifies that certain rates will change annually.⁵ This standardization reduces biases by eliminating the agency's ability to choose a discount rate that justifies a predetermined conclusion. It also allows a project to be

⁴[http://yosemite.epa.gov/ee/epa/eed.nsf/webpages/Guidelines.html/\\$file/Guidelines.pdf](http://yosemite.epa.gov/ee/epa/eed.nsf/webpages/Guidelines.html/$file/Guidelines.pdf)

⁵Annual rates can be found at <http://www.whitehouse.gov/omb/>. 2010 rates can be found at http://www.whitehouse.gov/omb/circulars_a094/a94_appx-c.

The Importance of the Discount Rate

Let's begin with an historical example. For years the United States and Canada had been discussing the possibility of constructing a tidal power project in the Passamaquoddy Bay between Maine and New Brunswick. This project would have heavy initial capital costs, but low operating costs that presumably would hold for a long time into the future. As part of their analysis of the situation, a complete inventory of costs and benefits was completed in 1959.

Using the same benefit and cost figures, Canada concluded that the project should not be built, while the United States concluded that it should. Because these conclusions were based on the same benefit–cost data, the differences can be attributed solely to the use of different discount rates. The United States used 2.5 percent while Canada used 4.125 percent. The higher discount rate makes the initial cost weigh much more heavily in the calculation, leading to the Canadian conclusion that the project would yield a negative net benefit. Since the lower discount rate weighs the lower future operating costs relatively more heavily, Americans saw the net benefit as positive.

In a more recent illustration of why the magnitude of the discount rate matters, on October 30, 2006 economist Nicholas Stern from the London School of Economics issued a report using a discount rate of 0.1 percent that concluded that the benefits of strong, early action on climate change would considerably outweigh the costs. Other economists, such as William Nordhaus of Yale University, who prefer a discount rate around 6 percent, believe that optimal economic policies to slow climate change involve only modest rates of emissions reductions in the near term, followed by sharp reductions in the medium and long term.

In this debate the desirability of strong current action is dependent (at least in part) on the size of the discount rate used in the analysis. Higher discount rates reduce the present value of future benefits from current investments in abatement, implying a smaller marginal benefit. Since the costs associated with those investments are not affected nearly as much by the choice of discount rate (remember that costs occurring in the near future are discounted less), a lower present value of marginal benefit translates into a lower optimal investment in abatement.

Far from being an esoteric subject, the choice of the discount rate is fundamentally important in defining the role of the public sector, the types of projects undertaken, and the allocation of resources across generations.

Sources: Edith Stokey and Richard Zeckhauser. *A Primer for Policy Analysis* (New York: W. W. Norton, 1978): 164–165; Raymond Mikesell. *The Rate of Discount for Evaluating Public Projects* (Washington, DC: The American Enterprise Institute for Public Policy Research, 1977): 3–5; the Stern Report: <http://webarchive.nationalarchives.gov.uk/> and http://www.hm-treasury.gov.uk/sternreview_index.htm; William Nordhaus. "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature* Vol. XLV (September 2007): 686–702

EXAMPLE 3.4

considered independently of fluctuations in the true social cost of capital due to cycles in the behavior of the economy. On the other hand, when the social opportunity cost of capital differs from this administratively determined level, the benefit–cost analysis will not, in general, define the efficient allocation.

Divergence of Social and Private Discount Rates

Earlier we concluded that producers, in their attempt to maximize producer surplus, also maximize the present value of net benefits under the “right” conditions, such as the absence of externalities, the presence of properly defined property rights, and the presence of competitive markets within which the property rights can be exchanged.

Now let’s consider one more condition. If resources are to be allocated efficiently, firms must use the same rate to discount future net benefits as is appropriate for society at large. If firms were to use a higher rate, they would extract and sell resources faster than would be efficient. Conversely, if firms were to use a lower-than-appropriate discount rate, they would be excessively conservative.

Why might private and social rates differ? The social discount rate is equal to the social opportunity cost of capital. This cost of capital can be separated into two components: the risk-free cost of capital and the risk premium. The *risk-free cost of capital* is the rate of return earned when there is absolutely no risk of earning more or less than the expected return. The *risk premium* is an additional cost of capital required to compensate the owners of this capital when the expected and actual returns may differ. Therefore, because of the risk premium, the cost of capital is higher in risky industries than in no-risk industries.

One difference between private and social discount rates may stem from a difference in social and private risk premiums. If the risk of certain private decisions is different from the risks faced by society as a whole, then the social and private risk premiums may differ. One obvious example is the risk *caused* by the government. If the firm is afraid its assets will be taken over by the government, it may choose a higher discount rate to make its profits before nationalization occurs. From the point of view of society—as represented by government—this is not a risk and, therefore, a lower discount rate is appropriate. When private rates exceed social rates, current production is higher than is desirable to maximize the net benefits to society. Both energy production and forestry have been subject to this source of inefficiency.

Another divergence in discount rates may stem from different underlying rates of time preference. Such a divergence in time preferences can cause not only a divergence between private and social discount rates (as when firms have a higher rate of time preference than the public sector), but even between otherwise similar analyses conducted in two different countries.

Time preferences would be expected to be higher, for example, in a cash-poor, developing country than in an industrialized country. Since the two benefit–cost

analyses in these two countries would be based upon two different discount rates, they might come to quite different conclusions. What is right for the developing country may not be right for the industrialized country and vice versa.

Although private and social discount rates do not always diverge, they may. When those circumstances arise, market decisions are not efficient.

A Critical Appraisal

We have seen that it is sometimes, but not always, difficult to estimate benefits and costs. When this estimation is difficult or unreliable, it limits the value of a benefit–cost analysis. This problem would be particularly disturbing if biases tended to increase or decrease net benefits systematically. Do such biases exist?

In the early 1970s, Robert Haveman (1972) conducted a major study that shed some light on this question. Focusing on Army Corps of Engineers water projects, such as flood control, navigation, and hydroelectric power generation, Haveman compared the *ex ante* (before the fact) estimate of benefits and costs with their *ex post* (after the fact) counterparts. Thus, he was able to address the issues of accuracy and bias. He concluded that

In the empirical case studies presented, ex post estimates often showed little relationship to their ex ante counterparts. On the basis of the few cases and the a priori analysis presented here, one could conclude that there is a serious bias incorporated into agency ex ante evaluation procedures, resulting in persistent overstatement of expected benefits. Similarly in the analysis of project construction costs, enormous variance was found among projects in the relationship between estimated and realized costs. Although no persistent bias in estimation was apparent, nearly 50 percent of the projects displayed realized costs that deviated by more than plus or minus 20 percent from ex ante projected costs.⁶

In the cases examined by Haveman, at least, the notion that benefit–cost analysis is purely a scientific exercise was clearly not consistent with the evidence; the biases of the analysts were merely translated into numbers.

Does their analysis mean that benefit–cost analysis is fatally flawed? Absolutely not! It does, however, highlight the importance of calculating an accurate value and of including all of the potential benefits and costs (e.g., nonmarket values). It also serves to remind us, however, that benefit–cost analysis is not a stand-alone technique. It should be used in conjunction with other available information. Economic analysis including benefit–cost analysis can provide useful information, but it should not be the only determinant for all decisions.

Another shortcoming of benefit–cost analysis is that it does not really address the question of who reaps the benefits and who pays the cost. It is quite possible for a particular course of action to yield high net benefits, but to have the benefits

⁶A more recent assessment of costs (Harrington et al., 1999) found evidence of both overestimation and underestimation, although overestimation was more common. The authors attributed the overestimation mainly to a failure to anticipate technical innovation.

borne by one societal group and the costs borne by another. This admittedly extreme case does serve to illustrate a basic principle—ensuring that a particular policy is efficient provides an important, but not always the sole, basis for public policy. Other aspects, such as who reaps the benefit or bears the burden, are also important.

In summary, on the positive side, benefit–cost analysis is frequently a very useful part of the policy process. Even when the underlying data are not strictly reliable, the outcomes may not be sensitive to that unreliability. In other circumstances, the data may be reliable enough to give indications of the consequences of broad policy directions, even when they are not reliable enough to fine-tune those policies. Benefit–cost analysis, when done correctly, can provide a useful complement to the other influences on the political process by clarifying what choices yield the highest net benefits to society.

On the negative side, benefit–cost analysis has been attacked as seeming to promise more than can actually be delivered, particularly in the absence of solid benefit information. This concern has triggered two responses. First, regulatory processes have been developed that can be implemented with very little information and yet have desirable economic properties. The recent reforms in air pollution control, which we cover in Chapter 15, provide one powerful example.

The second approach involves techniques that supply useful information to the policy process without relying on controversial techniques to monetize environmental services that are difficult to value. The rest of this chapter deals with the two most prominent of these—cost-effectiveness analysis and impact analysis.

Even when benefits are difficult or impossible to quantify, economic analysis has much to offer. Policy-makers should know, for example, how much various policy actions will cost and what their impacts on society will be, even if the efficient policy choice cannot be identified with any certainty. Cost-effectiveness analysis and impact analysis both respond to this need, albeit in different ways.

Cost-Effectiveness Analysis

What can be done to guide policy when the requisite valuation for benefit–cost analysis is either unavailable or not sufficiently reliable? Without a good measure of benefits, making an efficient choice is no longer possible.

In such cases, frequently it is possible, however, to set a policy target on some basis other than a strict comparison of benefits and costs. One example is pollution control. What level of pollution should be established as the maximum acceptable level? In many countries, studies of the effects of a particular pollutant on human health have been used as the basis for establishing that pollutant's maximum acceptable concentration. Researchers attempt to find a threshold level below which no damage seems to occur. That threshold is then further lowered to provide a margin of safety and that becomes the pollution target.

Approaches could also be based upon expert opinion. Ecologists, for example, could be enlisted to define the critical numbers of certain species or the specific critical wetlands resources that should be preserved.

Once the policy target is specified, however, economic analysis can have a great deal to say about the cost consequences of choosing a means of achieving that objective. The cost consequences are important not only because eliminating wasteful expenditures is an appropriate goal in its own right, but also to assure that they do not trigger a political backlash.

Typically, several means of achieving the specified objective are available; some will be relatively inexpensive, while others turn out to be very expensive. The problems are frequently complicated enough that identifying the cheapest means of achieving an objective cannot be accomplished without a rather detailed analysis of the choices.

Cost-effectiveness analysis frequently involves an *optimization procedure*. An optimization procedure, in this context, is merely a systematic method for finding the lowest-cost means of accomplishing the objective. This procedure does not, in general, produce an efficient allocation because the predetermined objective may not be efficient. All efficient policies are cost-effective, but not all cost-effective policies are efficient.

In Chapter 2 we introduced the efficiency equimarginal principle. According to that principle, net benefits are maximized when the marginal benefit is equal to the marginal cost.

A similar, and equally important equimarginal principle exists for cost-effectiveness:

*Second Equimarginal Principle (the Cost-Effectiveness Equimarginal Principle):
The least-cost means of achieving an environmental target will have been achieved
when the marginal costs of all possible means of achievement are equal.*

Suppose we want to achieve a specific emissions reduction across a region, and several possible techniques exist for reducing emissions. How much of the control responsibility should each technique bear? The cost-effectiveness equimarginal principle suggests that the techniques should be used such that the desired reduction is achieved and the cost of achieving the last unit of emissions reduction (in other words, the marginal control cost) should be the same for all sources.

To demonstrate why this principle is valid, suppose that we have an allocation of control responsibility where marginal control costs are much higher for one set of techniques than for another. This cannot be the least-cost allocation since we could lower cost while retaining the same amount of emissions reduction. Costs could be lowered by allocating more control to the lower marginal cost sources and less to the high marginal cost sources. Since it is possible to find a way to lower cost, then clearly the initial allocation could not have minimized cost. Once marginal costs are equalized, it becomes impossible to find any lower-cost way of achieving the same degree of emissions reduction; therefore, that allocation must be the allocation that minimizes costs.

In our pollution control example, cost-effectiveness can be used to find the least-cost means of meeting a particular standard and its associated cost. Using this cost as a benchmark case, we can estimate how much costs could be expected to increase

EXAMPLE
3.5

NO₂ Control in Chicago: An Example of Cost-Effectiveness Analysis

In order to compare compliance costs of meeting a predetermined ambient air quality standard in Chicago, Seskin, Anderson, and Reid (1983) gathered information on the cost of control for each of 797 stationary sources of nitrogen oxide emissions in the city of Chicago, along with measured air quality at 100 different locations within the city. The relationship between ambient air quality at those receptors and emissions from the 797 sources was then modeled using mathematical equations. Once these equations were estimated, the model was calibrated to ensure that it was capable of re-creating the actual situation in Chicago. Following successful calibration, this model was used to simulate what would happen if EPA were to take various regulatory actions.

The results indicated that a cost-effective strategy would cost less than one-tenth as much as the traditional approach to control and less than one-seventh as much as a more sophisticated version of the traditional approach. In absolute terms, moving to a more cost-effective policy was estimated to save more than \$100 million annually in the Chicago area alone. In Chapters 15 and 16 we shall examine in detail the current movement toward cost-effective policies, a movement triggered in part by studies such as this one.

from this minimum level if policies that are not cost-effective are implemented. Cost-effectiveness analysis can also be used to determine how much compliance costs can be expected to change if the EPA chooses a more stringent or less stringent standard. The case study presented in Example 3.5 not only illustrates the use of cost-effectiveness analysis, but also shows that costs can be very sensitive to the regulatory approach chosen by the EPA.

Impact Analysis

What can be done when the information needed to perform a benefit–cost analysis or a cost-effectiveness analysis is not available? The analytical technique designed to deal with this problem is called *impact analysis*. An impact analysis, regardless of whether it focuses on economic impact or environmental impact or both, attempts to quantify the consequences of various actions.

In contrast to benefit–cost analysis, a pure impact analysis makes no attempt to convert all these consequences into a one-dimensional measure, such as dollars, to ensure comparability. In contrast to cost-effectiveness analysis, impact analysis does not necessarily attempt to optimize. Impact analysis places a large amount of relatively undigested information at the disposal of the policy-maker. It is up to the policy-maker to assess the importance of the various consequences and act accordingly.

On January 1, 1970, President Nixon signed the National Environmental Policy Act of 1969. This act, among other things, directed all agencies of the federal government to

include in every recommendation or report on proposals for legislation and other major Federal actions significantly affecting the quality of the human environment, a detailed statement by the responsible official on—

- i. the environmental impact of the proposed action,*
- ii. any adverse environmental effects which cannot be avoided should the proposal be implemented,*
- iii. alternatives to the proposed action,*
- iv. the relationships between local short-term uses of man's environment and the maintenance and enhancement of long-term productivity; and*
- v. any irreversible and irretrievable commitments of resources which would be involved in the proposed action should it be implemented.⁷*

This was the beginning of the environmental impact statement, which is now a familiar, if controversial, part of environmental policy-making.

Current environmental impact statements are more sophisticated than their early predecessors and may contain a benefit–cost analysis or a cost-effectiveness analysis in addition to other more traditional impact measurements. Historically, however, the tendency had been to issue huge environmental impact statements that are virtually impossible to comprehend in their entirety.

In response, the Council on Environmental Quality, which, by law, administers the environmental impact statement process, has set content standards that are now resulting in shorter, more concise statements. To the extent that they merely quantify consequences, statements can avoid the problem of “hidden value judgments” that sometimes plague benefit–cost analysis, but they do so only by bombarding the policy-makers with masses of noncomparable information.

Summary

Finding a balance in the relationship between humanity and the environment requires many choices. Some basis for making rational choices is absolutely necessary. If not made by design, decisions will be made by default.

Normative economics uses benefit–cost analysis for judging the desirability of the level and composition of provided services. Cost-effectiveness analysis and impact analysis offer alternatives to benefit–cost analysis. All of these techniques offer valuable information for decision making and all have shortcomings.

⁷83 Stat. 853.

A static efficient allocation is one that maximizes the net benefit over all possible uses of those resources. The dynamic efficiency criterion, which is appropriate when time is an important consideration, is satisfied when the outcome maximizes the present value of net benefits from all possible uses of the resources. Later chapters examine the degree to which our social institutions yield allocations that conform to these criteria.

Because benefit–cost analysis is both very powerful and very controversial, in 1996 a group of economists of quite different political persuasions got together to attempt to reach some consensus on its proper role in environmental decision making. Their conclusion is worth reproducing in its entirety:

Benefit-cost analysis can play an important role in legislative and regulatory policy debates on protecting and improving health, safety, and the natural environment. Although formal benefit-cost analysis should not be viewed as either necessary or sufficient for designing sensible policy, it can provide an exceptionally useful framework for consistently organizing disparate information, and in this way, it can greatly improve the process and, hence, the outcome of policy analysis. If properly done, benefit-cost analysis can be of great help to agencies participating in the development of environmental, health and safety regulations, and it can likewise be useful in evaluating agency decision-making and in shaping statutes.⁸

Even when benefits are difficult to calculate, however, economic analysis in the form of cost–effectiveness can be valuable. This technique can establish the least expensive ways to accomplish predetermined policy goals and to assess the extra costs involved when policies other than the least-cost policy are chosen. What it cannot do is answer the question of whether those predetermined policy goals are efficient.

At the other end of the spectrum is impact analysis, which merely identifies and quantifies the impacts of particular policies without any pretense of optimality or even comparability of the information generated. Impact analysis does not guarantee an efficient outcome.

All three of the techniques discussed in this chapter are useful, but none of them can stake a claim as being universally the “best” approach. The nature of the information that is available and its reliability make a difference.

Discussion Questions

1. Is risk-neutrality an appropriate assumption for benefit–cost analysis? Why or why not? Does it seem more appropriate for some environmental problems than others? If so, which ones? If you were evaluating the desirability of locating a hazardous waste incinerator in a particular town, would the Arrow-Lind rationale for risk-neutrality be appropriate? Why or why not?

⁸From Kenneth Arrow et al. “Is There a Role for Benefit-Cost Analysis in Environmental, Health and Safety Regulation?” *Science* Vol. 272 (April 12, 1996): 221–222. Reprinted with Permission from AAAS.

2. Was the executive order issued by President Bush mandating a heavier use of benefit–cost analysis in regulatory rule making a step toward establishing a more rational regulatory structure, or was it a subversion of the environmental policy process? Why?

Self-Test Exercises

1. Suppose a proposed public policy could result in three possible outcomes: (1) present value of net benefits of \$4,000,000, (2) present value of net benefits of \$1,000,000, or (3) present value of net benefits of $-\$10,000,000$ (i.e., a loss). Suppose society is risk-neutral and the probability of occurrence of each of these three outcomes are, respectively, 0.85, 0.10, and 0.05, should this policy be pursued or trashed? Why?
2.
 - a. Suppose you want to remove ten fish of an exotic species that have illegally been introduced to a lake. You have three possible removal methods. Assume that q_1 , q_2 , and q_3 are, respectively, the amount of fish removed by each method that you choose to use so that the goal will be accomplished by any combination of methods such that $q_1 + q_2 + q_3 = 10$. If the marginal costs of each removal method are, respectively, $\$10q_1$, $\$5q_2$, and $\$2.5q_3$, how much of each method should you use to achieve the removal cost-effectively?
 - b. Why isn't an exclusive use of method 3 cost-effective?
 - c. Suppose that the three marginal costs were constant (not increasing as in the previous case) such that $MC_1 = \$10$, $MC_2 = \$5$, and $MC_3 = \$2.5$. What is the cost-effective outcome in that case?
3. Consider the role of discount rates in problems involving long time horizons such as climate change. Suppose that a particular emissions abatement strategy would result in a \$500 billion reduction in damages 50 years into the future. How would the maximum amount spent now to eliminate those damages change if the discount rate is 2 percent, rather than 10 percent?

Further Reading

- Freeman, A. Myrick III. *The Measurement of Environmental and Resource Values*, 2nd ed. (Washington, DC: Resources for the Future, Inc., 2003). A comprehensive and analytically rigorous survey of the concepts and methods for environmental valuation.
- Hanley, Nick, and Clive L. Spash. *Cost-Benefit Analysis and the Environment* (Brookfield, VT: Edward Elgar Publishing Company, 1994). An up-to-date account of the theory and practice of this form of analysis applied to environmental problems. Contains numerous specific case studies.
- Norton, Bryan, and Ben A. Minteer. "From Environmental Ethics to Environmental Public Philosophy: Ethicists and Economists: 1973–Future," in T. Tietenberg and H. Folmer,

eds. *The International Yearbook of Environmental and Resource Economics: 2002/2003* (Cheltenham, UK: Edward Elgar, 2002): 373–407. A review of the interaction between environmental ethics and economic valuation.

Scheraga, Joel D., and Frances G. Sussman. “Discounting and Environmental Management,” in T. Tietenberg and H. Folmer, eds. *The International Yearbook of Environmental and Resource Economics 1998–1999* (Cheltenham, UK: Edward Elgar, 1998): 1–32. A summary of the “state of the art” for the use of discounting in environmental management.

Additional References and Historically Significant References are available on this book’s Companion Website: <http://www.pearsonhighered.com/tietenberg/>

Appendix

The Simple Mathematics of Dynamic Efficiency*

Assume that the demand curve for a depletable resource is linear and stable over time. Thus, the inverse demand curve in year t can be written as

$$P_t = a - bq_t \quad (1)$$

The total benefits from extracting an amount q_t in year t are then the integral of this function (the area under the inverse demand curve):

$$\begin{aligned} (\text{Total benefits})_t &= \int_0^{q_t} (a - bq) dq \\ &= aq_t - \frac{b}{2} q_t^2 \end{aligned} \quad (2)$$

Further assume that the marginal cost of extracting that resource is a constant c and therefore the total cost of extracting any amount q_t in year t can be given by

$$(\text{Total cost})_t = cq_t \quad (3)$$

If the total available amount of this resource is \bar{Q} , then the dynamic allocation of a resource over n years is the one that satisfies the maximization problem:

$$\text{Max}_q \sum_{i=1}^n \frac{aq_i - bq_i^2/2 - cq_i}{(1+r)^{i-1}} + \lambda \left[\bar{Q} - \sum_{i=1}^n q_i \right] \quad (4)$$

Assuming that \bar{Q} is less than would normally be demanded, the dynamic efficient allocation must satisfy

$$\frac{a - bq_i - c}{(1+r)^{i-1}} - \lambda = 0, \quad i = 1, \dots, n \quad (5)$$

$$\bar{Q} - \sum_{i=1}^n q_i = 0 \quad (6)$$

An implication of Equation 5 is that $(P - MC)$ increases over time at rate r . This difference, which is known as the marginal user cost, will play a key role in our thinking about allocating depletable resources over time. A fuller understanding of this concept is provided in Chapter 5.

*Greater detail on the mathematics of constrained optimization can be found in any standard mathematical economics text.