impacts trade would have on the economy and the environment. It could not, however, provide any guidance on the question of whether trade was desirable. That judgment would have to come from normative economics.

Normative analysis can arise in several different contexts. 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. Both of these types of situations share the characteristic that the alternatives being evaluated are well defined in advance. Here the relevant question is: Should we do it or shouldn't we?

A rather different context for normative economics can arise when the possibilities are more open-ended. For example, we might ask how much should we control emissions of greenhouse gases (which contribute to global warming) and how should we achieve that degree of control? Or we might ask how much forest of various types should be preserved? Answering these questions requires us to consider the entire range of possible outcomes and to select the best or optimal one. Although that is a much more difficult question to answer than one that asks us only to compare two predefined alternatives, the basic normative analysis framework is the same in both cases.

Normative Criteria for Decision Making

Evaluating Predefined Options

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

That simple *benefit/cost analysis* framework provides the starting point for the economic approach. 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 $B > C$, then support the action.
Otherwise, oppose the action.[2]

As long as B and C are positive, an equivalent formulation would be:

If $B/C > 1$, then 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 will 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

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

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 (see Debate 2.1).

Benefits can be derived from the demand curve for the good or service provided by the action. Demand curves measure the amount of a particular good people would be willing to purchase at various prices. In a typical situation, a person will purchase less of a commodity (or environmental service) the higher is its price. In Figure 2.2, when the price is P_0, Q_0 will be purchased, but if the price rises to P_1, purchases will fall to Q_1.

DEBATE 2.1

Should Humans Place an Economic Value on the Environment?

Arne Naess, the late Norwegian philosopher, used the term "deep ecology" to refer to the view that the nonhuman environment has "intrinsic" value, a value that is independent of human interests. Intrinsic value is contrasted with "instrumental" value in which the value of the environment is derived from its usefulness in satisfying human wants.

Two issues are raised by the Naess critique: (1) What is the basis for valuing the environment? and (2) How is the valuation accomplished? The belief that the environment may have a value that goes beyond its direct usefulness to humans is in fact quite consistent with modern economic valuation techniques. As we show in Chapter 3, economic valuation techniques now include the ability to quantify a wide range of nonuse values as well as the more traditional use values.

Controversies over how the values are derived are less easily resolved. As described in this chapter, economic valuation is based firmly upon human preferences. Deep ecology, on the other hand, would argue that allowing humans to determine the value of other species would have no more moral basis than allowing other species to determine the value of humans. Rather, deep ecologists argue, humans should only use environmental resources when necessary for survival; otherwise nature should be left alone. And, because economic valuation is not helpful in determining survival necessity, deep ecologists argue that it contributes little to environmental management.

Those who oppose all economic valuation, however, face a dilemma: When humans fail to value the environment it may be assigned a default value of zero in calculations designed to guide policy. A value of zero, however derived, will tend to justify a great deal of environmental degradation that could not be justified with proper economic valuation. As a 1998 issue of "Ecological Economics" demonstrated, a number of environmental professionals now support economic valuation as a way to demonstrate just how valuable the environment is to modern society. At the very least, support seems to be growing for the proposition that economic valuation can be a very useful means of demonstrating when environmental degradation is senseless, even when judged from a limited anthropomorphic perspective (see Example 2.1).

Sources: Costanza, R. et al., "The value of ecosystem services: putting the issues in perspective" *Ecological Economics* 25 (1998) (1): 67–72 and the other articles on valuation in that issue. Daily, Gretchen and Katherine Ellison, *The New Economy of Nature: The Quest to Make Conservation Profitable* (Washington: Island Press, 2003).

FIGURE 2.2 The Individual Demand Curve

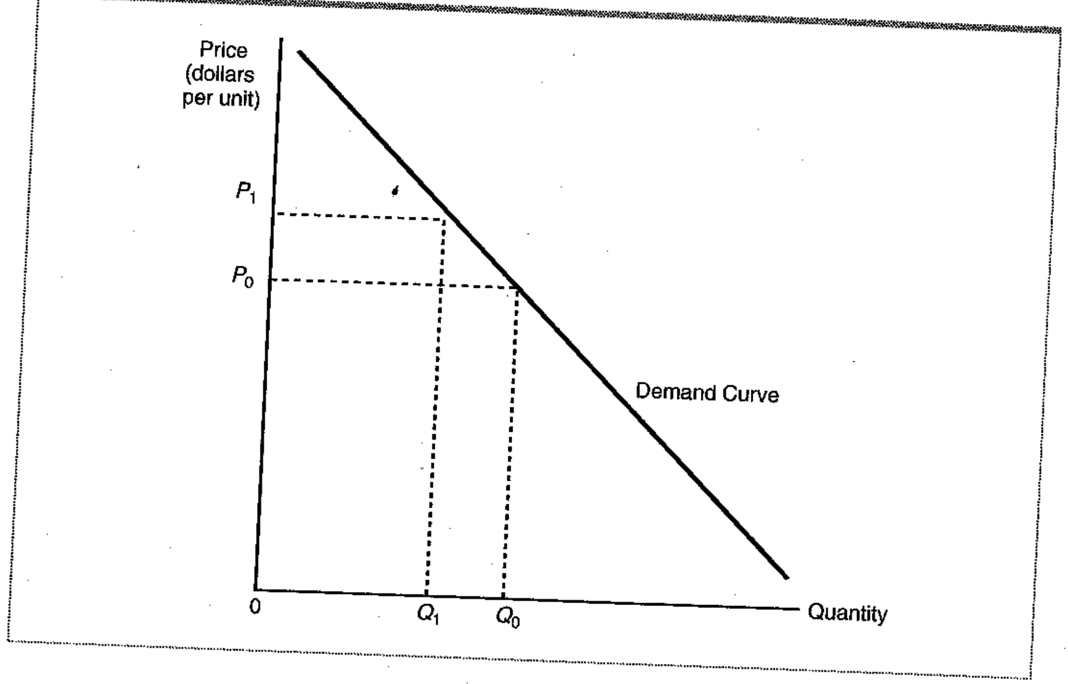

The meaning of these demand curves can be illustrated with the following hypothetical experiment. Suppose you were asked: At a price of *X* dollars, how much commodity *Y* would you buy? Your answer could be recorded as a point on a diagram, as shown in Figure 2.2. By repeating the question many times for different prices, we could trace out a locus of points. Connecting these points would yield an individual *demand curve*. Adding up all of the individual amounts demanded by all individuals at some stipulated price yields one point on the market demand curve. Connecting the points for various prices reveals the market demand curve.

For each quantity purchased, the corresponding point on the market demand curve represents the amount of money some person is willing to pay for the last unit of the good. The *total willingness to pay* for some quantity of this good—say, three units—is the sum of the willingness to pay for each of the three units. Thus, the total willingness to pay for three units would be measured by the sum of the willingness to pay for the first, second, and third units, respectively. It is now a simple extension to determine that the total willingness to pay is the area under the continuous market demand curve to the left of the allocation in question. For example, in Figure 2.3 the total willingness to pay for five units of the commodity is the shaded area.[3] Total willingness to pay is the concept we will use to define *total benefits*, specifically total benefits is equal to total willingness to pay.

[3]From simple geometry it can be noticed that for linear demand curves this area is the sum of the areas of the triangle on top plus the rectangle on the bottom. The area of a right triangle is $^1/_2 \times$ base $\times$ height. Therefore, in our example this area is $^1/_2 \times \$5 \times 5 + \$5 \times 5 = \$37.50$.

FIGURE 2.3 The Relationship of Demand to Willingness to Pay

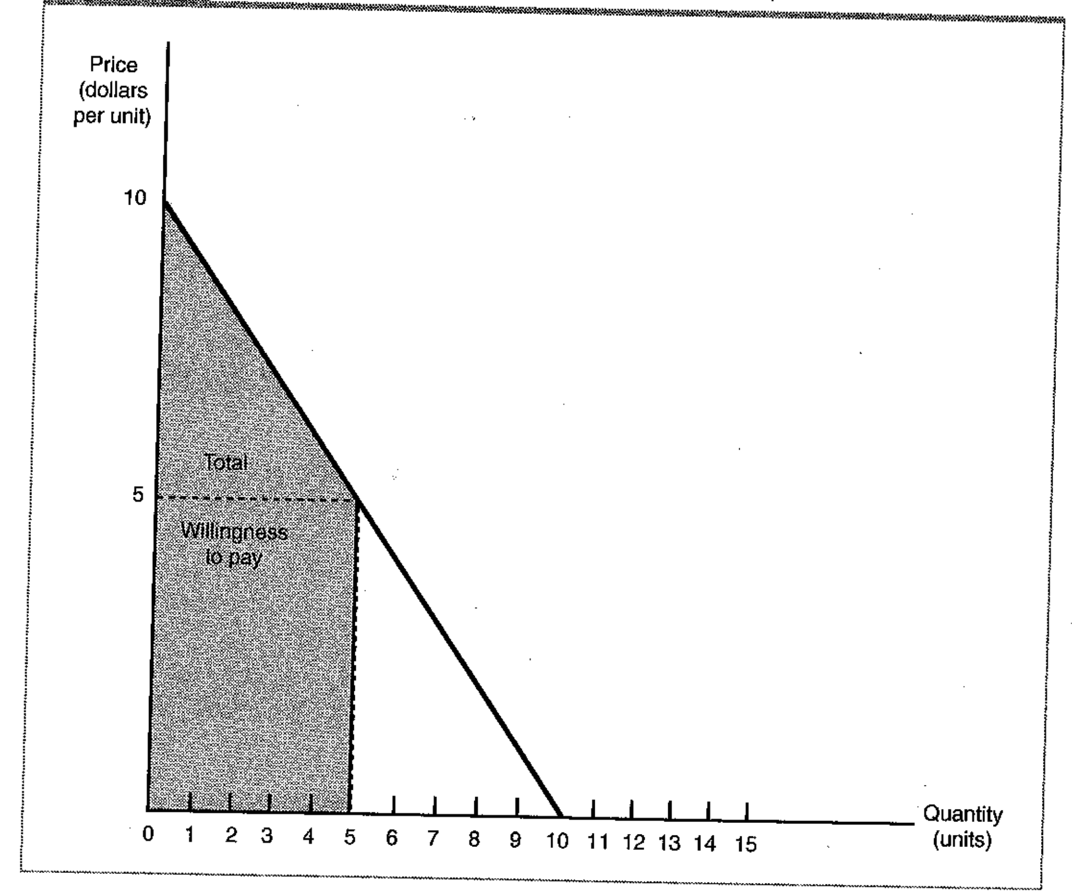

Measuring total costs on the same set of axes involves logic similar to measuring total benefits. It is important to stress that environmental services have costs even though they are produced without any human input. All costs should be measured as opportunity costs.

As illustrated in Example 2.1, the *opportunity cost* for using resources 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 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 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 forgone net benefit that would have resulted from the white-water canoeing.

In graphing costs, we will use the marginal opportunity cost curve to correspond to the *marginal willingness to pay* function used above to graph benefits. The *marginal opportunity cost*

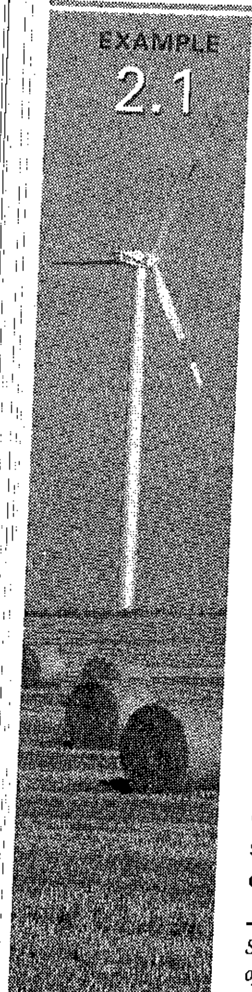

EXAMPLE 2.1

Valuing Ecological Services from Preserved Tropical Forests

As Chapter 11 explains, one of the main threats to tropical forests comes when the forested land is converted 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 decided to tackle 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 15–50 percent.

When the ecologists 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 begin to demonstrate the considerable economic value of preserving the forest, even when considering only specific instrumental values.

Source: Ricketts, Taylor H. et al., "Economic Value of Tropical Forest to Coffee Production" *PNAS (Proceedings of the National Academy of Science)* 101 August 24, 2002 (34): 12579–12582.

curve defines the additional cost of producing the last unit. In purely competitive markets, the marginal opportunity cost curve is identical to the supply curve.

Total cost is simply the sum of the *marginal costs*.[4] The total cost of producing three units is equal to the cost of producing the first unit plus the cost of producing the second unit plus the cost of producing the third unit. As with total willingness to pay, the geometric representation of the sum of the individual elements of a continuous marginal cost curve is the area under the marginal cost curve, as illustrated in Figure 2.4 by the shaded area *FGIJK*.[5]

[4]Strictly speaking, the sum of the marginal costs is equal to total variable cost. In the short run, this is smaller than total cost by the amount of the fixed cost. For our purposes this distinction is not important.

[5]Notice again that this area is the sum of a right triangle and a rectangle. In Figure 2.4 the total variable cost of producing five units is $18.75. Why?

FIGURE 2.4 The Relationship of Marginal Cost and Total Cost

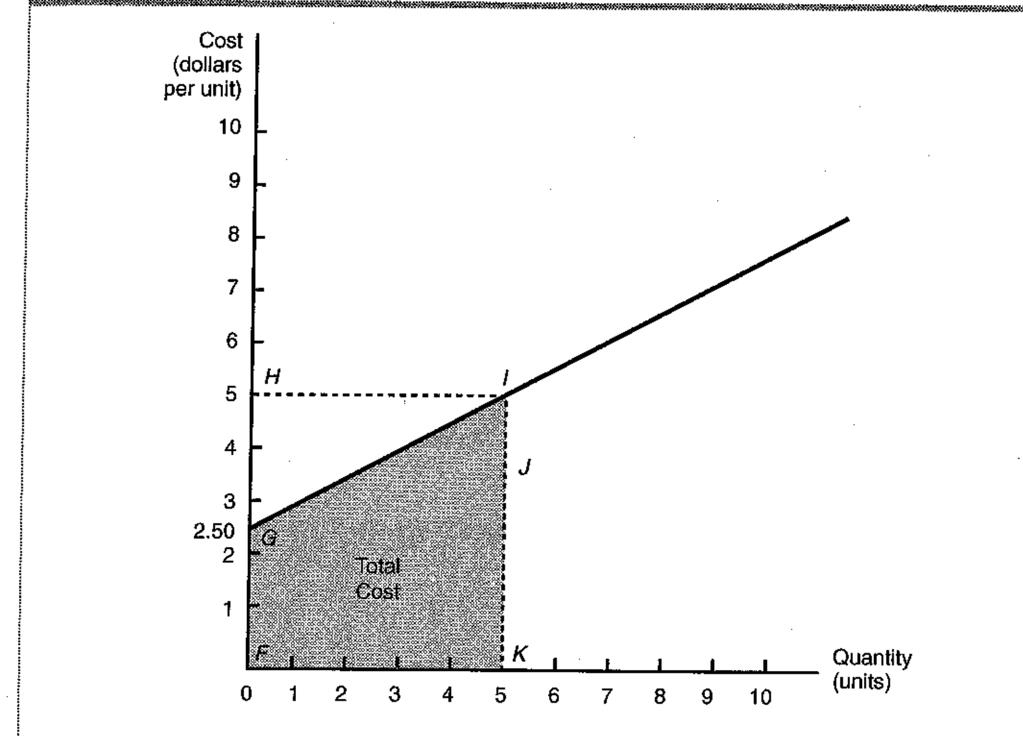

Since net benefit is defined as the excess of benefits over costs, it follows that net benefit is equal to that portion of the area under the demand curve that lies above the supply curve. Consider Figure 2.5, which combines the information in Figures 2.3 and 2.4

Let's now use this apparatus to illustrate the use of the decision rules introduced earlier. Let's suppose for example that we are considering preserving a four-mile stretch of river and that the benefits and costs of that action are reflected in Figure 2.5. Should that stretch be preserved?

The answer is clearly yes. The total benefits (the area under the marginal willingness to pay curve) are clearly larger than the total costs (the area under the marginal cost curve). Since *net benefits* are positive, it makes sense to preserve four miles of the river.

Incorporating the Timing of Benefits and Costs: The Present Value Criterion

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 occur at different times?

FIGURE 2.5 The Derivation of Net Benefits

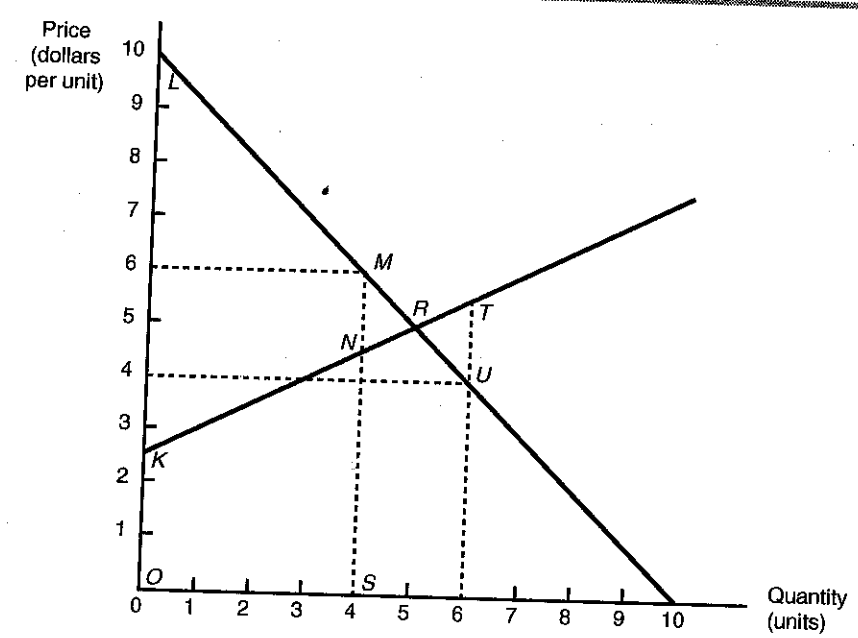

Incorporating time into the analysis requires an extension of the concepts we have already developed. This extension provides a way of thinking not only about the magnitude of benefits and costs, but also about timing. In order to incorporate timing, the decision rule must provide a way to compare the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

What could your dollar earn in two years at r percent interest? Because of compound interest, the amount would be $\$1(1 + r)(l + r) = \$1(1 + r)^2$. 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 net present value of a *one-time* net benefit received n years from now is

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

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

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

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*.[6]

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, $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 2.1).

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 zero balance (see Table 2.2). 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.

TABLE 2.1 Demonstrating Present Value Calculations

Year	1	2	3	4	5	Sum
Annual Amounts	3,000	5,000	6,000	10,000	12,000	36,000
Present Value (r = .06)	$2,830.19	$4,449.98	$5,037.72	$7,920.94	$8,967.10	$29,205.92

TABLE 2.2 Interpreting Present Value Calculations

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 = .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	

[6]The discount rate should equal the social opportunity cost of capital. In Chapter 4 we examine the questions of whether private firms can be expected to use the socially correct discount rate. In Chapter 3 we discuss how the discount rate is chosen by the government.

It is now possible to show how this analysis can be used to evaluate actions. First, 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 be rejected.

Finding the Optimal Outcome

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

In subsequent chapters that address individual environmental problems, the normative analysis will proceed in three steps: (1) We will identify an optimal outcome; (2) we will attempt to discern the extent to which our institutions produce optimal outcomes, and where divergences occur between actual and optimal outcomes, to uncover the behavioral sources of the problems; and (3) 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. How these three steps are applied to each of the environmental problems will reflect the uniqueness of each situation, but the overarching framework used to shape that analysis will be the same.

To provide some concreteness of this approach, 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 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 was actually applied, not only did it reveal that stocks are much lower than optimal for many fisheries, but it also identified the reason for excessive exploitation. 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 actual analysis not only revealed that current waste levels are excessive, but also suggested some specific behavioral sources of the problem. Based upon this understanding, targeted 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.

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 what is meant by optimal.

Static Efficiency

The chief normative economic criterion for choosing among various allocations occurring at the same point in time is called *static efficiency*, or merely *efficiency*. An allocation of resources is said to satisfy the static efficiency criterion if the net benefit from the use of those resources is maximized by that allocation.

Let's show how this concept can be applied by returning to Figure 2.5. Previously we asked whether an action that preserved four miles of river was worth doing. The answer was yes because the net benefits from that action were positive.

Static efficiency, however, requires us to ask a rather different question: What is the efficient number of miles to be preserved? We know from the definition that the efficient amount of preservation would maximize net benefits. Do four units maximize net benefits?

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. 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. Are 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. Therefore, from our definition, preserving five miles constitutes an efficient allocation.[7]

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

> *First Equimarginal Principle (the Efficiency Equimarginal Principle): Net benefits are maximized when the marginal benefits from an allocation equal the marginal costs.*

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 some people without any deleterious effects on at least one other person.*

Allocations that do not satisfy this definition are suboptimal. Suboptimal allocations can always be rearranged so that some people are better off and no one is hurt by the rearrangement. 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

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

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, there is no way the gainers could sufficiently compensate the losers; 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 net benefit. By failing to maximize net benefit, they are forgoing an opportunity to make some people better off without harming others.

For a different lens on determining efficiency consider the following numerical example:

Suppose the demand (marginal willingness to pay) for some environmental service can be expressed as $MWP = \$80 - 2q$ and the marginal cost curve as $MC = \$10$.

What level of q would represent the static efficient allocation of that environmental service and how large are the net benefits? We know from the first equimarginal principle that in an efficient allocation, where net benefits are maximized, $MWP = MC$.

So $MWP = \$80 - 2q = MC = \$10$.

Solving this for q yields $q = 35$, the static efficient level of q. The associated marginal willingness to pay and marginal cost are both \$10.

We can now draw the diagram that shows the net benefits (see Figure 2.6).

To calculate the net benefits we need to subtract the total cost (the area under the marginal cost curve) from the total benefits (the area under the marginal willingness to pay curve). Geometrically speaking, the total benefits are the sum of the triangle bounded by the *MWP* curve, the vertical axis and the *MC* curve from $q = 0$ to $q = 35$ plus the rectangle formed by the horizontal axis from $q = 0$ to $q = 35$ and $MC = \$10$ over that same interval.

The area of the benefit triangle is

$$\frac{1}{2}(base)(height) = \frac{1}{2}(\$35)(\$70) = \$1,225.$$

FIGURE 2.6 Finding the Efficient Outcome

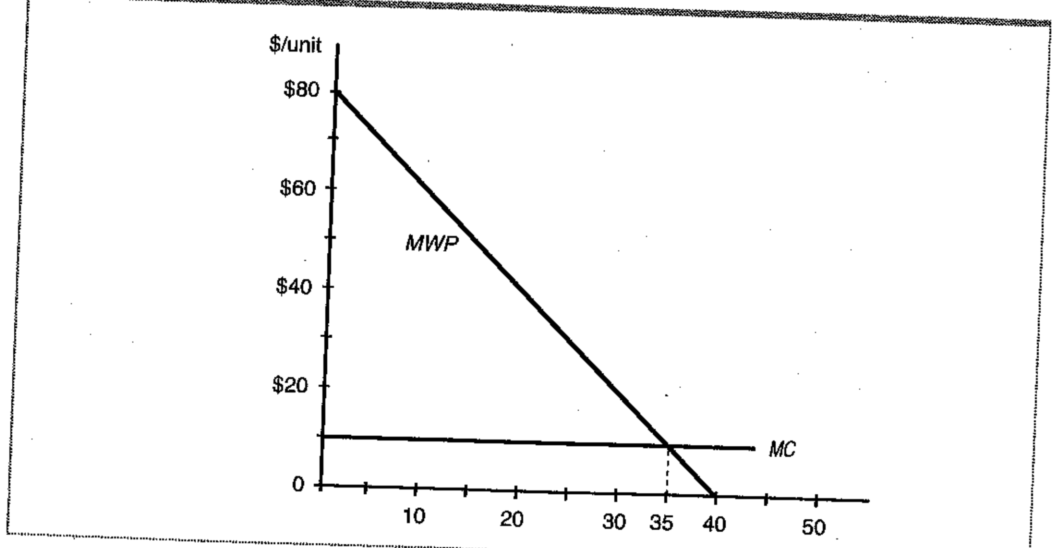

Do you see where the \$70 comes from? It is the vertical distance from the value of the *MWP* when $q = 0$ (in other words the place where it crosses the vertical axis) minus the height of the *MC* (\$10).

The area of the benefit rectangle is $(base)(height) = (\$35)(\$10) = \$350$.

Hence, total benefits are the sum or \$1,575. Total costs are the rectangle under $MC = \$10$ from $q = 0$ to $q = 35$ or $(base)(height) = (\$35)(\$10) = \$350$.

Net benefits are total benefits minus total cost or \$1,225. Notice that in this case the net benefits are simply equal to the benefit triangle. In general, that will be true when the marginal cost is constant as it is in this example.

If you want more practice, find the efficient level of q and both the *MWP* and the *MC* if the marginal cost function were $MC = \$2q$ instead of $MC = \$10$. If you obtained $q = 20$ and $MC = MWP = \$40$, you nailed it!

Dynamic Efficiency

The static efficiency criterion is very useful for comparing resource allocations when time is not an important factor. How can we make choices when the benefits and costs may occur at different times?

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 Environmental Protection Agency (EPA) to evaluate the benefits and costs of the United States air pollution control policy over the 1970–1990 period (see Example 2.2).

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 United States policy. Benefits quantified by this study included reduced death rates and lower incidences of chronic bronchitis, lead poisoning, stroke, respiratory disease, and heart disease as well as the benefits of better visibility, reduced damage to structures, and improved agricultural productivity. They were unable to quantify many suspected ecosystem effects.