




Pearson International Edition

ENVIRONMENTAL &
NATURAL RESOURCE
ECONOMICS

EIGHTH EDITION

Tom Tietenberg
Lynne Lewis

2

Valuing the Environment: Concepts

When you have eliminated the impossible, whatever remains, however improbable, must be the truth.

—Sherlock Holmes, From Sir Arthur Conan Doyle's *The Sign of Four* (1890)

Introduction

Before examining specific environmental problems and the policy responses to them, it is important that we develop and clarify the economics approach, so that we will have some sense of the forest before examining each of the trees. By having a feel for the conceptual framework, it becomes easier not only to deal with individual cases, but also, perhaps more importantly, to see how they fit into a comprehensive approach.

In this chapter we develop the general conceptual framework used in economics to approach environmental problems. We begin by examining the relationship between human actions, as manifested through the economic system, and the environmental consequences of those actions. We can then establish criteria for judging the desirability of the outcomes of this relationship. These criteria provide a basis for identifying the nature and severity of environmental problems, and a foundation for designing effective policies to deal with them.

Throughout this chapter the economic point of view is contrasted with alternative points of view. These contrasts bring the economic approach into sharper focus and stimulate deeper and more critical thinking about all possible approaches.

The Human Environment Relationship

The Environment as an Asset


In economics the environment is viewed as a composite asset that provides a variety of services. It is a very special asset, to be sure, because it provides the life-support systems that sustain our very existence, but it is an asset nonetheless. As with other assets, we wish to enhance, or at least prevent undue depreciation of, the value of this asset so that it may continue to provide aesthetic and life-sustaining services.

The environment provides the economy with raw materials, which are transformed into consumer products by the production process, and energy, which fuels this transformation. Ultimately these raw materials and energy return to the environment as waste products (see Figure 2.1).

The environment also provides services directly to consumers. The air we breathe, the nourishment we receive from food and drink, and the protection we derive from shelter and clothing are all benefits we receive either directly or indirectly from the environment. In addition, anyone who has experienced the exhilaration of white-water canoeing, the total serenity of a wilderness trek, or the breathtaking beauty of a sunset will readily recognize that the environment provides us with a variety of amenities for which no substitute exists.

If the environment is defined broadly enough, the relationship between the environment and the economic system can be considered a *closed system*. For our purposes, a closed system is one in which no inputs (energy, matter, and so on) are received from outside the system and no outputs are transferred outside the system. An *open system*, by contrast, is one in which the system imports or exports matter or energy.

FIGURE 2.1 The Economic System and the Environment



If we restrict our conception of the relationship in Figure 2.1 to our planet and the atmosphere around it, then clearly we do not have a closed system. We derive most of our energy from the sun, either directly or indirectly. We have also sent spaceships well beyond the boundaries of our atmosphere. Nonetheless, historically speaking, for *material* inputs and outputs (not including energy), this system can be treated as a closed system because the amount of exports (such as abandoned space vehicles) and imports (for example, moon rocks) are negligible. Whether the system remains closed depends on the degree to which space exploration opens up the rest of our solar system as a source of raw materials.

The treatment of our planet and its immediate environs as a closed system has an important implication that is summed up in the *first law of thermodynamics*—energy and matter can neither be created nor destroyed.¹ The law implies that the mass of materials flowing into the economic system from the environment has either to accumulate in the economic system or return to the environment as waste. When accumulation stops, the mass of materials flowing into the economic system is equal in magnitude to the mass of waste flowing into the environment.

Excessive wastes can, of course, depreciate the asset; when they exceed the absorptive capacity of nature, wastes reduce the services that the asset provides. Examples are easy to find: air pollution can cause respiratory problems; polluted drinking water can cause cancer; smog obliterates scenic vistas; climate change can lead to flooding of coastal areas.

The relationship of people to the environment is also conditioned by another physical law, the *second law of thermodynamics*. Known popularly as the *entropy law*, this law states that “entropy increases.” *Entropy* is the amount of energy unavailable for work. Applied to energy processes, this law implies that no conversion from one form of energy to another is completely efficient and that the consumption of energy is an irreversible process. Some energy is always lost during conversion, and the rest, once used, is no longer available for further work. The second law also implies that in the absence of new energy inputs, any closed system must eventually use up its available energy. Since energy is necessary for life, life ceases when useful energy flows cease.

We should remember that our planet is not even approximately a closed system with respect to energy; we gain energy from the sun. The entropy law does remind us, however, that the flow of solar energy establishes an upper limit on the flow of available energy that can be sustained. Once the stocks of stored energy (such as fossil fuels and nuclear energy) are gone, the amount of energy available for useful work will be determined solely by the solar flow and by the amount that can be stored (dams, trees, and so on). Thus, over the very long run, the growth process will be limited by the availability of solar energy and our ability to put it to work.

The Economic Approach

Two different types of economic analysis can be applied to increase our understanding of the relationship between the economic system and the environment: *Positive*

¹We know, however, from Einstein's famous equation ($E = mc^2$) that matter can be transformed into energy. This transformation is the source of energy in nuclear power.

economics attempts to describe *what is*, *what was*, or *what will be*. *Normative* economics, by contrast, deals with what *ought to be*. Disagreements within positive economics can usually be resolved by an appeal to the facts. Normative disagreements, however, involve value judgments.

Both branches are useful. Suppose, for example, we want to investigate the relationship between trade and the environment. Positive economics could be used to describe the kinds of 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.

The fact that positive analysis does not, by itself, determine the desirability of some policy action does not mean that it is not useful in the policy process. Example 2.1 provides one example of the kinds of economic impact analyses that are used in the policy process.

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

Economic Impacts of Reducing Hazardous Pollutant Emissions from Iron and Steel Foundries

EXAMPLE 2.1

The U.S. Environmental Protection Agency (EPA) was tasked with developing a "maximum achievable control technology standard" to reduce emissions of hazardous air pollutants from iron and steel foundries. As part of the rule-making process, EPA conducted an *ex ante* economic impact analysis to assess the potential economic impacts of the proposed rule.

If implemented the rule would require some iron and steel foundries to implement pollution control methods that would increase the production costs at affected facilities. The interesting question addressed by the analysis is how large those impacts would be.

The impact analysis estimated annual costs for existing sources to be \$21.73 million. These cost increases were projected to result in small increases in output prices. Specifically prices were projected to increase by only 0.1 percent for iron castings and 0.05 percent for steel castings. The impacts of these price increases were expected to be experienced largely by iron foundries using cupola furnaces as well as consumers of iron foundry products. Unaffected domestic foundries and foreign producers of coke were actually projected to earn slightly higher profits as a result of the rule.

This analysis helped in two ways. First, by showing that the impacts fell under the \$100 million threshold that mandates review by the Office of Management and Budget, the analysis eliminated the need for a much more time and resource consuming analysis. Second, by showing how small the expected impacts would be, it served to lower the opposition that might have arisen from unfounded fears of much more severe impacts.

Source: Office of Air Quality Planning and Standards, United States Environmental Protection Agency. Economic Impact Analysis of Proposed Iron and Steel Foundries NESHAP Final Report. November 2002; National Emissions Standards for Hazardous Air Pollutants for Iron and Steel Foundries, Proposed Rule, Federal Register Vol. 72, No. 73 (April 17, 2007): 19150-19164.

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 have done it) or not?

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 climate change) 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 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 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 $B > C$, support the action.

Otherwise, oppose the action.²

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

If $B/C > 1$, 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

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

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 cost. 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 .

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 the valuing of 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. Proponents of 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 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 the enormous value of the environment 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.

Sources: R. Costanza et al. "The Value of Ecosystem Services: Putting the Issues in Perspective," *Ecological Economics* Vol. 25, No. 1 (1998): 67–72 and the other articles on valuation in that issue; Gretchen Daily and Katherine Ellison. *The New Economy of Nature: The Quest to Make Conservation Profitable* (Washington, DC: Island Press, 2003).

DEBATE

1.1






The meaning of these demand curves can be illustrated with this 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

³From 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 $1/2 \times \text{base} \times \text{height}$. Therefore, in our example this area is $1/2 \times \$5 \times 5 + \$5 \times 5 = \$37.50$.



Total willingness to pay is the concept we shall use to define *total benefits*. Thus, total benefits are equal to the area under the market demand curve from the origin to the allocation of interest.

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 presented in Example 2.2, 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 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.

EXAMPLE

2.2

Valuing Ecological Services
from Preserved Tropical Forests

As Chapter 13 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): 12579–12582.

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

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.

⁴Strictly 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.

⁵Note 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





FIGURE 2.5 The Derivation of Net Benefits



Let's now use this apparatus 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 2.5. Should that stretch be preserved? Explain why or why not?

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.⁶

⁶The discount rate should equal the social opportunity cost of capital. See (Scheraga and Sussman, 1998) for details on techniques for environmental discounting. 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 the use of the discount rate for policy analysis by the government.

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 \$0 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.

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.

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.

TABLE 2.1 Demonstrating Present Value Calculations

Year	1	2	3	4	5	Sum
Annual Amounts	3000	5000	6000	10000	12000	36000
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	3000	5000	6000	10000	12000	

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 14.

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 9.

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."

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 see 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, namely, 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 miles 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 or optimal allocation.⁷

One implication of this example, which shall be very useful in succeeding chapters, is what we shall 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 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

⁷The 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?

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 net benefit. By failing to maximize net benefit, they are forgoing an opportunity to make some people better off without harming others.

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 over the 1970–1990 period (see Example 2.3).

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,

Does Reducing Pollution Make Economic Sense?

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 to 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, as shown in Table 2.3. According to this study, U.S. air pollution control policy during this period made very good economic sense.

EXAMPLE 2.3

TABLE 2.3 Monetized Benefits and Costs of the U.S. Clean Air Act, 1970–1990 (billions of 1990 dollars)

	1975	1980	1985	1990	Present Value
Benefits ^a	355	930	1,155	1,248	22,200
Costs ^b	14	21	25	26	523
Net Benefits	341	909	1,130	1,220	21,700

^aThese are the mean (average) benefits. Due to the uncertainties involved, EPA also calculated low and high estimates.

^bThese are the annualized costs. (Many investments in pollution control involve the purchase of durable equipment that lasts many years.) Rather than put all of the expense in the year of purchase, EPA distributed the costs over the useful lives of this equipment.

Source: Created by the authors from information presented in 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 on p. 56.

and improved agricultural productivity. They were unable to quantify many suspected ecosystem effects.

Two categories of costs were also quantified. The first category included the higher costs of goods and services as the costs of installing, operating, and maintaining pollution control equipment were passed on to the consumers in the form of higher prices. The second category included the costs associated with designing and implementing the regulations as well as monitoring and enforcing compliance with them.

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 the pollution damage to ecosystems that was avoided by this policy, the net benefits are strongly positive. While presumably the case for controlling pollution would have been even stronger had they been included, the desirability of this form of control is strong enough even when they are not included. 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 discovered that the policy during this period was not completely efficient. As we shall show in Chapters 15 and 16, the costs of the chosen policy approach were higher than necessary to achieve the desired emission 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 is preserved, the ecosystem benefits will be retained, but the opportunity for increased income and employment will be 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 2.4).

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.

The economics approach views the environment as a composite asset, supplying a variety of services to humanity. The intensity and composition of those services depend on the actions of humans as constrained by physical laws, such as the first and second laws of thermodynamics.

Economics has two rather different means of enhancing understanding of environmental and natural resource economics. Positive economics is useful in describing the actions of people and the impact of those actions on the environmental asset. Normative economics can provide guidance on how optimal service flows can be defined and achieved.

Normative economics uses benefit/cost analysis for judging the desirability of the level and composition of provided services. 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

Choosing Between Preservation and Development in Australia

EXAMPLE
2.4

The Kakadu Conservation Zone (KCZ), 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 3, 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): 727-749.

satisfied when the outcome maximizes the present value of net benefits from all possible uses of the resources. Future chapters examine the degree to which our social institutions yield allocations that conform to these criteria.

Discussion Question

1. It has been suggested that we should use the "net energy" criterion to make choices among various types of energy. Net energy is defined as the total energy content in the energy source minus the energy required to extract, process, and deliver it to consumers. According to this criterion, we should use those sources with the highest net energy content first. Would the dynamic efficiency criterion and the net energy criterion be expected to yield the same choice? Why or why not?

Problem

1. One convenient way to express the willingness-to-pay relationship between price and quantity is to use the inverse demand function. In an inverse