The Addison-Wesley Series in Economics

Environmental &
Natural Resource
Economics

8th Edition

$2-009$

Tom Tietenberg
Colby College

Lynne Lewis
Bates College

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

$$\text{If } B > C, \text{ support the action.}$$

Otherwise, oppose the action.[2]

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

$$\text{If } B/C > 1, \text{ support the action.}$$

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

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

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

---
[2]Actually 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.

DEBATE 1.1

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).

FIGURE 2.2   The Individual Demand Curve

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]

---

[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 2.2   The Individual Demand Curve

FIGURE 2.3   The Relationship of Demand to Willingness to Pay

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.
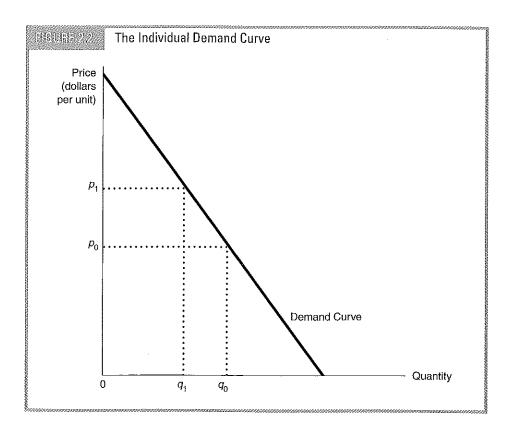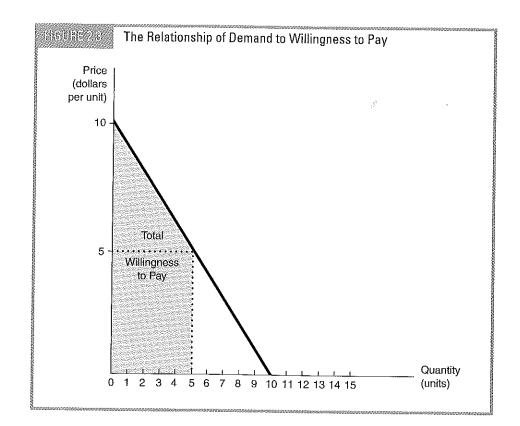
For each quantity purchased, the corresponding point on the market demand curve represents the amount of money some person is willing to pay for the last unit of the good. The *total willingness to pay* for some quantity of this good—say, three units—is the sum of the willingness to pay for each of the 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]

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.

[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 $1/2 \times \$5 \times 5 = \$5 \times 5 = \$37.50$.

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.

---

[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]Note 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?

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:
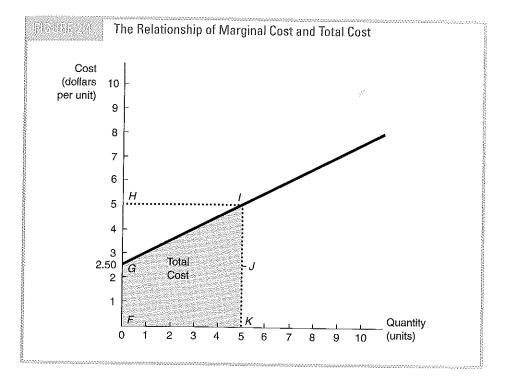
> The value of forest in providing crop pollination service alone is … of at least the same order [of magnitude] as major competing land uses, and infinitely greater than that recognized by most governments (i.e., zero).
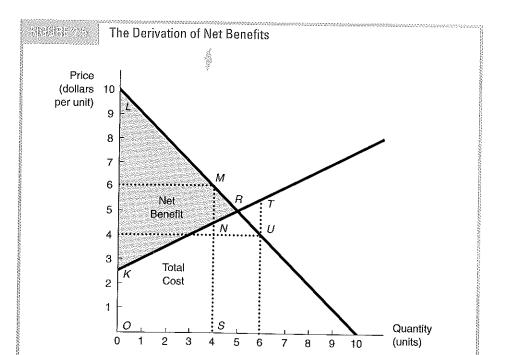
These estimates only partially capture the value of this forest because they consider only a single farm and a single type of ecological service. (This forest also provides carbon storage and water purification services, for example, and these were not included in the calculation.) Despite their partial 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.

[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]Note again that this area is the sum of a right triangle and a rectangle. In Figure 2.4 the total variable

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 $\{B_0, \ldots, B_n\}$ received over a period of n years is computed as

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

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]

---

[6]The 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.