

COST-BENEFIT ANALYSIS

Paris is well worth a Mass.

—ATTRIBUTED TO HENRI IV OF FRANCE

If you visited Boston during the last decade, you probably noticed that traffic downtown was particularly congested. The reason was the “Big Dig,” a massive \$14.6 billion public works project that involved the construction of new roads and another tunnel to Logan Airport. Many people have doubts that it was worth the money. How would one go about thinking about this issue? Infrastructure projects like the Big Dig are just one variety of the thousands of public projects that are under consideration at any given time, everything from breast cancer screening programs to space exploration. How should the government decide whether or not to pursue a particular project? The theory of welfare economics provides a framework for deciding: Evaluate the social welfare function before and after the project, and see whether social welfare increases. If it does, then do the project.

This method is correct, but not very useful. The amount of information required to specify and evaluate a social welfare function is enormous. While social welfare functions are valuable for thinking through certain conceptual problems, they are generally not much help for the day-to-day problems of project evaluation. However, welfare economics does provide the basis for **cost-benefit analysis**—a set of practical procedures for guiding public expenditure decisions.¹

Most government projects and policies result in the private sector having more of some scarce commodities and less of others. At the core of cost-benefit analysis is a set of systematic procedures for valuing these commodities, which allows policy analysts to determine whether a project is, on balance, beneficial. Cost-benefit analysis allows policymakers to attempt to do what well-functioning markets do automatically—allocate resources to a project as long as the marginal social benefit exceeds the marginal social cost.

► PRESENT VALUE

Project evaluation usually requires comparing costs and benefits from different time periods. For example, preschool education for poor children requires substantial expenditures in the present and then yields returns in the future. In this section we discuss issues that arise in comparing dollar amounts from different time periods.

¹ Boardman et al. [2006] discuss the links between welfare economics and cost-benefit analysis.

Initially, we assume that no price inflation occurs. We show later how to take inflation into account.

Projecting Present Dollars into the Future

Suppose that you take \$100 to the bank and deposit it in an account that yields 5 percent interest after taxes. At the end of one year, you will have $(1 + 0.05) \times \$100 = \105 —the \$100 initially deposited, plus \$5 in interest. Suppose further that you let the money sit in the account for another year. At the end of the second year, you will have $(1 + 0.05) \times \$105 = \110.25 . This can also be written as $(1 + 0.05) \times (1 + 0.05) \times 100 = (1 + 0.05)^2 \times 100$. Similarly, if the money is deposited for three years, it will be worth $(1 + 0.05)^3 \times \$100$ by the end of the third year. More generally, if \$ R are invested for T years at an interest rate of r , at the end of T years, it will be worth $\$R \times (1 + r)^T$. This formula shows the future value of money invested in the present.

Projecting Future Dollars into the Present

Now suppose that someone offers a contract that promises to pay you \$100 *one year from now*. The person is trustworthy, so you do not have to worry about default. (Also, remember there is no inflation.) What is the maximum amount that you should be willing to pay *today* for this promise? It is tempting to say that a promise to pay \$100 is worth \$100. But this neglects the fact that the promised \$100 is not payable for a year, and in the meantime you are forgoing the interest that could be earned on the money. Why should you pay \$100 today to receive \$100 a year from now, if you can receive \$105 a year from now simply by putting the \$100 in the bank today? Thus, the value today of \$100 payable one year from now is *less* than \$100. The **present value** of a future amount of money is the maximum amount you would be willing to pay today for the right to receive the money in the future.

To find the very most you would be willing to give up now in exchange for \$100 payable one year in the future, you must find the number that, when multiplied by $(1 + 0.05)$ just equals \$100. By definition, this is $\$100/(1 + 0.05)$ or approximately \$95.24. Thus, when the interest rate is 5 percent, the present value of \$100 payable one year from now is $\$100/(1 + 0.05)$. Note the symmetry with the familiar problem of projecting money into the future that we just discussed. To find the value of money today one year in the future, you *multiply* by one plus the interest rate; to find the value of money one year in the future today, you *divide* by one plus the interest rate.

Next consider a promise to pay \$100 *two years from now*. In this case, the calculation has to take into account the fact that if you invested \$100 yourself for two years, at the end it would be worth $\$100/(1 + 0.05)^2$. The most you would be willing to pay today for \$100 in two years is the amount that when multiplied by $(1 + 0.05)^2$ yields exactly \$100, that is, $\$100/(1 + 0.05)^2$, or about \$90.70.

In general, when the interest rate is r , the present value of a promise to pay \$ R in T years is simply $\$R/(1 + r)^T$.² Thus, even in the absence of inflation, a dollar in the future is worth less than a dollar today and must be “discounted” by an amount

² This assumes the interest rate is constant at r . Suppose that the interest rate changes over time, so in year 1 it is r_1 , in year 2, r_2 , and so on. Then the present value of a sum \$ R_T payable T years from now is $\$R_T/[(1 + r_1) \times (1 + r_2) \times \dots \times (1 + r_T)]$.

cost-benefit analysis

A set of procedures based on welfare economics for guiding public expenditure decisions.

present value

The value today of a given amount of money to be paid or received in the future.

discount rate

The rate of interest used to compute present value.

discount factor

The number by which an amount of future income must be divided to compute its present value. If the interest rate is r and the income is receivable T periods in the future, the discount factor is $(1 + r)^T$.

that depends on the interest rate and when the money is receivable. For this reason, r is often referred to as the **discount rate**. Similarly, $(1 + r)^T$ is called the **discount factor** for money T periods into the future. Note that the further into the future the promise is payable (the larger is T), the smaller is the present value. Intuitively, the longer you have to wait for a sum to be paid, the less you are willing to pay for it today, other things being the same.

Finally, consider a promise to pay $\$R_0$ today, and $\$R_1$ one year from now, and $\$R_2$ two years from now, and so on for T years. How much is this deal worth? By now, it is clear that the naive answer ($\$R_0 + \$R_1 + \dots + \$R_T$) is wrong because it assumes that a dollar in the future is exactly equivalent to a dollar in the present. Without dividing by the discount factor, adding up dollars from different points in time is like adding apples and oranges. The correct approach is to convert each year's amount to its present value and then add them.

Table 8.1 shows the present value of each year's payment. To find the present value (PV) of the income stream, $\$R_0, \$R_1, \$R_2, \dots, \R_T we simply add the figures in the last column:

$$PV = R_0 + \frac{R_1}{(1 + r)} + \frac{R_2}{(1 + r)^2} + \dots + \frac{R_T}{(1 + r)^T} \quad (8.1)$$

The importance of computing present value is hard to overestimate. Ignoring it can lead to serious errors. In particular, failure to discount makes ventures that yield returns in the future appear more valuable than they really are. For example, consider a project that yields a return of \$1 million 20 years from now. If the interest rate is 5 percent, the present value is \$376,889 [= \$1,000,000/(1.05)²⁰]. If $r = 10$ percent, the present value is only \$148,644 [= \$1,000,000/(1.10)²⁰].

Inflation

How do we modify the procedure when the price level is expected to increase in the future? To begin, consider a project that, in present prices, yields the same return each year. Call this return $\$R_0$. Now assume that inflation occurs at a rate of 7 percent per year, and the dollar value of the return increases along with all prices. Therefore, the dollar value of the return one year from now, $\$R_1$, is $(1.07) \times \$R_0$. Similarly, two years into the future, the dollar value is $\$R_2 = (1.07)^2 \times R_0$. In general, this same return has a dollar value in year T of $\$R_T = (1 + 0.07)^T \times R_0$.

Table 8.1 Calculating present value

Dollars Payable	Years in Future	Discount Factor	Present Value
R_0	0	1	R_0
R_1	1	$(1 + r)$	$R_1 / (1 + r)$
R_2	2	$(1 + r)^2$	$R_2 / (1 + r)^2$
...
R_T	T	$(1 + r)^T$	$R_T / (1 + r)^T$

In order to compute the present value of an income stream, divide each year's amount by the corresponding discount factor and then sum these terms across all years.

The dollar values $\$R_0, \$R_1, \$R_2, \dots, \R_T are referred to as **nominal amounts**. Nominal amounts are valued according to the level of prices in the year the return occurs. One can measure these returns in terms of the prices that exist in a single year. These are called **real amounts** because they do not reflect changes that are due merely to alterations in the price level. In our example, the real amount was assumed to be a constant $\$R_0$ measured in present prices. More generally, if the real returns in present year prices are $\$R_0, \$R_1, \$R_2, \dots, \R_T , and inflation occurs at a rate of π per year, then the nominal returns are $\$R_0, \$R_1 \times (1 + \pi), \$R_2 \times (1 + \pi)^2, \dots, \$R_T \times (1 + \pi)^T$.

But this is not the end of the story. When prices are expected to rise, lenders are no longer willing to make loans at the interest rate r that prevailed when prices were stable. Lenders realize they are going to be paid back in depreciated dollars, and to keep even in real terms, their first year's payment must also be inflated by $(1 + \pi)$. Similarly, the second year's payment must be inflated by $(1 + \pi)^2$. In other words, the market interest rate increases by an amount approximately equal to the expected rate of inflation, from r percent to $r + \pi$ percent.³

We see, then, that when inflation is anticipated, both the stream of returns and the discount rate increase. When expressed in *nominal* terms, the present value of the income stream is thus

$$PV = R_0 + \frac{(1 + \pi)R_1}{(1 + \pi)(1 + r)} + \frac{(1 + \pi)^2 R_2}{(1 + \pi)^2 (1 + r)^2} + \dots + \frac{(1 + \pi)^T R_T}{(1 + \pi)^T (1 + r)^T} \quad (8.2)$$

A glance at Equation (8.2) indicates that it is equivalent to Equation (8.1) because all the terms involving $(1 + \pi)$ cancel out. The moral of the story is that we obtain the *same* answer whether real or nominal magnitudes are used. It is crucial, however, that dollar magnitudes and discount rates be measured consistently. If real values are used for the R s, the discount rate must also be measured in real terms—the market rate of interest *minus* the expected inflation rate. Alternatively, if we discount by the market rate of interest, returns should be measured in nominal terms.

► PRIVATE SECTOR PROJECT EVALUATION

As we noted at the beginning of the chapter, the central problem in cost-benefit analysis is valuing the inputs and outputs of government projects. A useful starting point is to consider the same problem from a private firm's point of view.

Suppose a firm is considering two mutually exclusive projects, X and Y . The real benefits and costs of project X are B^X and C^X , respectively; and those for project Y are B^Y and C^Y . For both projects, the benefits and costs are realized immediately. The firm must answer two questions: First, should either project be done at all; are the projects *admissible*? (The firm has the option of doing neither project.) Second,

³ The product of $(1 + r)$ and $(1 + \pi)$ is $1 + r + \pi + r\pi$. Thus, the nominal rate actually exceeds the real rate by $\pi + r\pi$. However, for numbers of reasonable magnitude, $r\pi$ is negligible in size, so $r + \pi$ is a good approximation. Under some circumstances, nominal interest rates may fail to rise by exactly the rate of inflation. See Chapter 17 under "Taxes and Inflation."

nominal amounts

Amounts of money that are valued according to the price levels that exist in the years that the amounts are received.

real amounts

Amounts of money adjusted for changes in the general price level.

if both projects are admissible, which is *preferable*? Because both benefits and costs occur immediately, answering these questions is simple. Compute the net return to project X , $B^X - C^X$, and compare it to the net return to Y , $B^Y - C^Y$. A project is admissible only if its net return is positive, that is, if the benefits exceed the costs. If both projects are admissible and the firm can only adopt one of them, it should choose the project with the higher net return.

In reality, most projects involve a stream of real benefits and returns that occur over time rather than instantaneously. Suppose that the initial benefits and costs of project X are B_0^X and C_0^X , those at the end of the first year are B_1^X and C_1^X , and those at the end of the last year are B_T^X and C_T^X . We can characterize project X as a stream of net returns (some of which may be negative):

$$(B_0^X - C_0^X), (B_1^X - C_1^X), (B_2^X - C_2^X), \dots, (B_T^X - C_T^X)$$

The present value of this income stream (PV^X) is

$$PV^X = B_0^X - C_0^X + \frac{B_1^X - C_1^X}{(1+r)} + \frac{B_2^X - C_2^X}{(1+r)^2} + \dots + \frac{B_T^X - C_T^X}{(1+r)^T}$$

where r is the discount rate that is appropriate for a private sector project. (Selection of a discount rate is discussed shortly.)

Similarly, suppose that project Y generates streams of costs and benefits B^Y and C^Y over a period of T^Y years. (There is no reason for T and T^Y to be the same.) Project Y 's present value is

$$PV^Y = B_0^Y - C_0^Y + \frac{B_1^Y - C_1^Y}{(1+r)} + \frac{B_2^Y - C_2^Y}{(1+r)^2} + \dots + \frac{B_{T^Y}^Y - C_{T^Y}^Y}{(1+r)^{T^Y}}$$

Since both projects are now evaluated in present value terms, we can use the same rules that were applied to the instantaneous project described earlier. The **present value criteria** for project evaluation are that:

- A project is admissible only if its present value is positive.
- When two projects are mutually exclusive, the preferred project is the one with the higher present value.

The discount rate plays a key role in the analysis. Different values of r can lead to very different conclusions concerning the admissibility and comparability of projects.

Consider the two projects shown in Table 8.2, a research and development program (R&D) and an advertising campaign. Both require an initial outlay of \$1,000. The R&D program produces a return of \$600 at the end of the first year and \$550 at the end of the third year. The advertising campaign, on the other hand, has a single large payoff of \$1,200 in three years.

The calculations show that the discount rate chosen is important. For low values of r , the advertising is preferred to R&D. However, higher discount rates weigh against the advertising (where the returns are concentrated further into the future) and may even make the project inadmissible.

Thus, one must take considerable care that the value of r represents as closely as possible the firm's actual opportunity cost of funds. If the discount rate chosen is too high, it tends to discriminate against projects with returns that come in the relatively distant future and vice versa. The firm's tax situation is relevant in this context. If the going market rate of return is 10 percent, but the firm's tax rate

present value criteria

Rules for evaluating projects stating that (1) only projects with positive net present value should be carried out; and (2) of two mutually exclusive projects, the preferred project is the one with the higher net present value.

Table 8.2 Comparing the present value of two projects

Year	Annual Net Return		$r =$	PV	
	R&D	Advertising		R&D	Advertising
0	-\$1,000	-\$1,000	0	\$150	\$200
1	600	0	0.01	128	165
2	0	0	0.03	86	98
3	550	1,200	0.05	46	37
			0.07	10	-21

The choice of the discount rate can affect which of two projects yields higher present value. In this example, a lower discount rate makes the advertising project relatively more attractive, while a higher discount rate makes the R&D project relatively more attractive.

is 25 percent, its after-tax return is only 7.5 percent. Because the after-tax return represents the firm's opportunity cost, it should be used for r .

Several criteria other than present value are often used for project evaluation. As we will see, they can sometimes give misleading answers, and therefore, the present value criteria are preferable. However, these other methods are popular, so it is necessary to understand them and to be aware of their problems.

Internal Rate of Return

A firm is considering the following project: It spends \$1 million today on a new computer network and reaps a benefit of \$1.04 million in increased profits a year from now. If you were asked to compute the computer network's "rate of return," you would probably respond, "4 percent." Implicitly, you calculated that figure by finding the value of ρ that solves the following equation:

$$-\$1,000,000 + \frac{\$1,040,000}{(1+\rho)} = 0$$

We can generalize this procedure as follows: If a project yields a stream of benefits (B) and costs (C) over T periods, the **internal rate of return** (ρ) is defined as the ρ that solves the equation

$$B_0 - C_0 + \frac{B_1 - C_1}{(1+\rho)} + \frac{B_2 - C_2}{(1+\rho)^2} + \dots + \frac{B_T - C_T}{(1+\rho)^T} = 0 \quad (8.3)$$

The internal rate of return is the discount rate that would make the present value of the project just equal to zero.

An obvious admissibility criterion is to accept a project if ρ exceeds the firm's opportunity cost of funds, r . For example, if the project earns 4 percent while the firm can obtain 3 percent on other investments, the project should be undertaken. The corresponding comparability criterion is that if two mutually exclusive projects are both admissible, choose the one with the higher value of ρ .

Project selection using the internal rate of return can, however, lead to bad decisions. Consider project X that requires the expenditure of \$100 today and yields \$110 a year from now, so that its internal rate of return is 10 percent. Project Y requires

internal rate of return

The discount rate that would make a project's net present value zero.

\$1,000 today and yields \$1,080 in a year, generating an internal rate of return of 8 percent. (Neither project can be duplicated.) Assume that the firm can borrow and lend freely at a 6 percent rate of interest.

On the basis of internal rate of return, X is clearly preferred to Y . However, the firm makes only \$4 profit on X (\$10 minus \$6 in interest costs), while it makes a \$20 profit on Y (\$80 minus \$60 in interest costs). Contrary to the conclusion implied by the internal rate of return, the firm should prefer Y , the project with the higher profit. In short, when projects differ in size, the internal rate of return can give poor guidance.⁴ In contrast, the present value rule gives correct answers even when the projects differ in scale. The present value of X is $-100 + 110/1.06 = 3.77$, while that of Y is $-1,000 + 1,080/1.06 = 18.87$. The present value criterion says that Y is preferable, as it should.

Benefit-Cost Ratio

Suppose that a project yields a stream of benefits $B_0, B_1, B_2, \dots, B_T$, and a stream of costs $C_0, C_1, C_2, \dots, C_T$. Then the present value of the benefits, B , is

$$B = B_0 + \frac{B_1}{(1+r)} + \frac{B_2}{(1+r)^2} + \dots + \frac{B_T}{(1+r)^T}$$

and the present value of the costs, C , is

$$C = C_0 + \frac{C_1}{(1+r)} + \frac{C_2}{(1+r)^2} + \dots + \frac{C_T}{(1+r)^T} \quad (8.4)$$

The **benefit-cost ratio** is defined as B/C .

Admissibility requires that a project's benefit-cost ratio exceed one. Application of this rule always gives correct guidance. To see why, note simply that $B/C > 1$ implies that $B - C > 0$, which is just the present value criterion for admissibility.

As a basis for comparing admissible projects, however, the benefit-cost ratio is virtually useless. Consider a state that is studying two methods for disposing of toxic wastes. Method I is a toxic waste dump with $B = \$250$ million, $C = \$100$ million, and therefore a benefit-cost ratio of 2.5. Method II involves sending the wastes in a rocket to Saturn, which has $B = \$200$ million, $C = \$100$ million, and therefore a benefit-cost ratio of 2. The state's leaders choose the dump because it has the higher value of B/C . Now suppose that in their analysis of the dump, the analysts inadvertently neglected to take into account seepage-induced crop damage of \$40 million. If the \$40 million is viewed as a reduction in the dump's benefits, its B/C becomes $\$210/\$100 = 2.1$, and the dump is still preferred to the rocket. However, the \$40 million can just as well be viewed as an increase in costs, in which case $B/C = \$250/\$140 = 1.79$. Now the rocket looks better than the dump!

We have illustrated that there is an inherent ambiguity in computing benefit-cost ratios because benefits can always be counted as "negative costs" and vice versa. Thus, by judicious classification of benefits and costs, any admissible project's benefit-cost ratio can be made arbitrarily high. In contrast, a glance at Equation (8.1) indicates that such shenanigans have no effect whatsoever on the present value criterion because it is based on the *difference* between benefits and costs rather than their *ratio*.

⁴ This result rests on the assumption that neither project can be duplicated. Otherwise, duplicating project X 10 times would yield a \$100 profit, which is greater than the \$80 profit of project Y .

We conclude that the internal rate of return and the benefit-cost ratio can lead to incorrect inferences. The present value criterion is the most reliable guide.

► DISCOUNT RATE FOR GOVERNMENT PROJECTS

Sensible decision making by the government also requires present value calculations. However, the public sector should compute costs, benefits, and discount rates differently than the private sector. This section discusses problems in the selection of a public sector discount rate. We then turn to problems in evaluating costs and benefits.

As suggested previously, the discount rate chosen by private individuals should reflect the rate of return available on alternative investments. Although in practice pinpointing this rate may be difficult, from a conceptual point of view the firm's opportunity cost of funds gives the correct value of r .

There is less consensus on the conceptually appropriate discount rate for government projects. We now discuss several possibilities.⁵

Rates Based on Returns in the Private Sector

Suppose the last \$1,000 of private investment in the economy yields an annual rate of return of 16 percent. If the government extracts \$1,000 from the private sector for a project, and the \$1,000 is entirely at the expense of private sector investment, society loses the \$160 that would have been generated by the private sector project. Thus, the opportunity cost of the government project is the 16 percent rate of return in the private sector. Because it measures the opportunity cost, 16 percent is the appropriate discount rate. It is irrelevant whether or not this return is taxed. Whether it all stays with the investor or part goes to the government, the before-tax rate of return measures the value of output that the funds would have generated for society.

In practice, funds for a given project are collected from a variety of taxes, each of which has a different effect on consumption and investment. Hence, contrary to the assumption made earlier, it is likely that some of the funds for the government project would come at the expense of consumption as well as investment. What is the opportunity cost of funds that come at the expense of consumption? Consider Kenny, who is deciding how much to consume and how much to save this year. For each dollar Kenny consumes this year, he gives up one dollar of consumption next year *plus* the rate of return he would have earned on the dollar saved. Hence, the opportunity cost to Kenny of a dollar of consumption now is measured by the rate of return he would have received if he had saved the dollar. Suppose the before-tax yield on an investment opportunity available to Kenny is 16 percent, but he must pay 50 percent of the return to the government in the form of taxes. All that Kenny gives up when he consumes an additional dollar today is the *after-tax* rate of return of 8 percent. Because the after-tax rate of return measures what an *individual* loses when consumption is reduced, dollars that come at the expense of consumption should be discounted by the after-tax rate of return.

⁵ See Tresch [2002, Chapter 24] for further discussion of the alternative views.

benefit-cost ratio

The ratio of the present value of a stream of benefits to the present value of a stream of costs for a project.

Because funds for the public sector reduce both private sector consumption and investment, a natural solution is to use a weighted average of the before- and after-tax rates of return, with the weight on the before-tax rate equal to the proportion of funds that comes from investment, and that on the after-tax rate the proportion that comes from consumption. In the preceding example, if one-quarter of the funds come at the expense of investment and three-quarters at the expense of consumption, then the public sector discount rate is 10 percent ($\frac{1}{4} \times 16$ percent + $\frac{3}{4} \times 8$ percent). Unfortunately, in practice it is hard to determine what the proportions of sacrificed consumption and investment actually are for a given government project. And even with information on the impact of each tax on consumption and investment, it is difficult in practice to determine which tax is used to finance which project. The inability to determine reliably a set of weights lessens the usefulness of this approach as a practical guide to determining discount rates.

Social Discount Rate

social rate of discount

The rate at which society is willing to trade off present consumption for future consumption.

An alternative view is that public expenditure evaluation should involve a **social rate of discount**, which measures the valuation *society* places on consumption that is sacrificed in the present. But why should society's view of the opportunity cost of forgoing consumption differ from the opportunity cost revealed in market rates of return? The social discount rate may be lower for several reasons.

Concern for Future Generations It is the duty of public sector decision makers to care about the welfare not only of the current generation of citizens but of future generations as well. The private sector, on the other hand, is concerned only with its own welfare. Hence, from a social point of view, the private sector devotes too few resources to saving—it applies too high a discount rate to future returns. However, the idea of government as the unselfish guardian of the interests of future generations assumes an unrealistic degree of omniscience and benevolence. Moreover, even totally selfish individuals often engage in projects that benefit future generations. If future generations are expected to benefit from some project, the anticipated profitability is high, which encourages investment today. Private firms plant trees today in return for profits on wood sales that may not be realized for many years.⁶

Paternalism Even from the point of view of their own narrow self-interest, people may not be farsighted enough to weigh adequately benefits in the future; they therefore discount such benefits at too high a rate. The government should use the discount rate that individuals *would* use if they knew their own good. This is a paternalistic argument—government forces citizens to consume less in the present, and in return, they have more in the future, at which time they presumably thank the government for its foresight. Like all paternalistic arguments, it raises the fundamental philosophical question of when the government's preferences should be imposed on individuals.

Market Inefficiency When a firm undertakes an investment, it generates knowledge and technological know-how that can benefit other firms. In a sense,

⁶ Why should people invest in a project whose returns may not be realized until after they are dead? Because investors can always sell the rights to future profits to members of the younger generation and hence consume their share of the anticipated profits during their lifetimes.

then, investment creates positive externalities, and by the usual kinds of arguments, investment is underprovided by private markets (see Chapter 5 under "Positive Externalities"). By applying a discount rate lower than the market's, the government can correct this inefficiency. The enormous practical problem here is measuring the actual size of the externality. Moreover, the theory of externalities suggests that a more appropriate remedy would be to determine the size of the marginal external benefit at the optimum and grant a subsidy of that amount (see again Chapter 5).

It appears, then, that none of the arguments against using market rates provides much specific guidance with respect to the choice of a public sector discount rate. Where does this leave us? It would be difficult to argue very strongly against any public rate of discount in a range between the before- and after-tax rates of return in the private sector. One practical procedure is to evaluate the present value of a project over a range of discount rates and see whether or not the present value stays positive for all reasonable values of r . If it does, the analyst can feel some confidence that the conclusion is not sensitive to the discount rate. *Sensitivity analysis* is the process of conducting a cost-benefit analysis under a set of alternative reasonable assumptions and seeing whether the substantive results change.

Government Discounting in Practice

Historically, the federal government has used a variety of discount rates, depending on the agency and the type of project. According to recent rules issued by the US Office of Management and Budget (OMB), federal agencies are now required to conduct two separate analyses when evaluating their projects: one using a real discount rate of 7 percent and another using a real discount rate of 3 percent. This convention is very much in line with the economic reasoning discussed earlier in this chapter. Seven percent is an estimate of the private return on investment, so it is the appropriate discount rate for projects that extract resources from private investment. Three percent is an estimate of the rate at which society discounts future consumption, so it is the appropriate discount rate for projects that primarily extract resources from private consumption. Because it is usually difficult to know whether a government project is taking resources from private investment or private consumption, OMB's recommendation of using both discount rates allows one to see whether the substantive results are sensitive to the difference. Further, for government projects that affect future generations, OMB recommends an additional sensitivity analysis using discount rates of 1 to 3 percent. This is consistent with the notion, also discussed earlier, that the social discount rate may be lower than the market rate of return.

In the context of federal budget planning, there are major inconsistencies in the conventions used for discounting. When a new tax or expenditure program is introduced, its effects over a five-year period must be reported to determine whether or not they will put the budget out of balance.⁷ For these purposes, all that matters are the sums of the relevant taxes or expenditures; future flows are discounted at a rate of zero. Thus, for example, a policy that increased spending by a billion dollars today and was financed by a tax of a billion dollars five years from now would be viewed as having no effect on the deficit, while in present value terms, the package would lose money.

⁷ For some purposes, the Senate requires flows over a 10-year period.

Beyond the five-year window, the fiscal consequences of fiscal proposals are ignored; in effect, they are discounted at a rate of infinity! Consider a policy that raises \$5 billion within the first five years, but after 10 years loses \$20 billion. Under current budgetary rules, such a policy is scored as creating a surplus, while with any reasonable discount rate, its long-run effect is to lose money for the government. There is, in fact, some evidence that this peculiar fashion of discounting has biased government decision making in favor of policies that increase revenue in the short term but reduce it in the long term [Bazelon and Smetters, 1999].

► VALUING PUBLIC BENEFITS AND COSTS

The next step in project evaluation is computing benefits and costs. From a private firm's point of view, their computation is relatively straightforward. The benefits from a project are the revenues received; the costs are the firm's payments for inputs; and both are measured by market prices. The evaluation problem is more complicated for the government because market prices may not reflect *social* benefits and costs. Consider, for example, a highway expansion that might do some damage to the environment. One can imagine both the private and public sectors undertaking this project, but the private and public cost-benefit analyses would be rather different, because the public sector should take into account social costs, which include externalities.

We now discuss several ways for measuring the benefits and costs of public sector projects.

Market Prices

As noted in Chapter 3, in a properly functioning competitive economy, the price of a good simultaneously reflects its marginal social cost of production and its marginal value to consumers. It would appear that if the government uses inputs and/or produces outputs that are traded in private markets, then market prices should be used for valuation.

The problem is that real-world markets have many imperfections, such as monopoly, externalities, and so on. Therefore, prices do not necessarily reflect marginal social costs and benefits. The relevant question, however, is not whether market prices are perfect, but whether they are likely to be superior to alternative measures of value. Such measures would either have to be made up or derived from highly complicated—and questionable—models of the economy. And, whatever their problems, market prices provide plenty of information at a low cost. Most economists believe that in the absence of any glaring imperfections, market prices should be used to compute public benefits and costs.

Adjusted Market Prices

The prices of goods traded in imperfect markets generally do not reflect their marginal social costs.⁸ The **shadow price** of such a commodity is its underlying social

⁸ For further details, see Boardman et al. [2006].

marginal cost. Although market prices of goods in imperfect markets diverge from shadow prices, in some cases the market prices can be used to *estimate* the shadow prices. We discuss the relevant circumstances next. In each case, the key insight is that the shadow price depends on how the economy responds to the government intervention.

Monopoly In the nation of South Africa, the production of beer is monopolized by the company South African Breweries, Ltd. Imagine that the Education Ministry is contemplating the purchase of some beer for a controlled experiment to determine the impact of beer consumption on the performance of college students. How should the project's cost-benefit analysis take into account the fact that this input is monopolistically produced?

In contrast to perfect competition, under which price is equal to marginal cost, a monopolist's price is above marginal cost (see Chapter 3). Should the government value the beer at its market price (which measures its value to consumers) or at its marginal production cost (which measures the incremental value of the resources used in its production)?

The answer depends on the impact of the government purchase on the market. If production of beer is expected to increase by the exact amount used by the project, the social opportunity cost is the value of the resources used in the extra production—the marginal production cost. On the other hand, if no more beer will be produced, the government's use comes at the expense of private consumers, who value the beer at its demand price. If some combination of the two responses is expected, a weighted average of price and marginal cost is appropriate. (Note the similarity to the previous discount rate problem.)

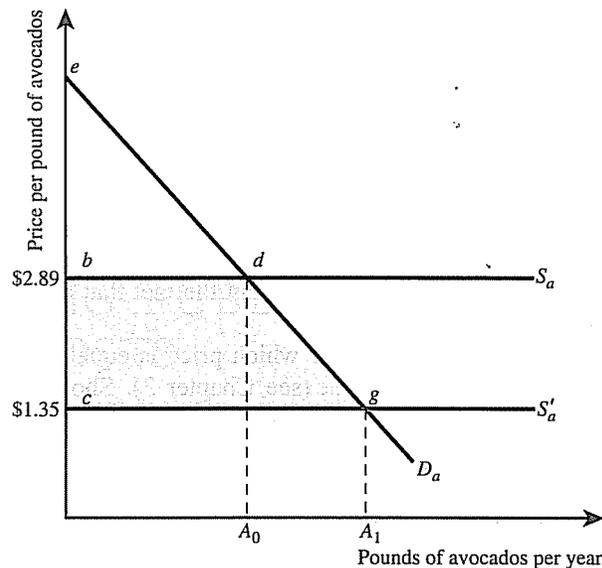
Taxes If an input is subject to a sales tax, the price received by the producer of the input is less than the price paid by the purchaser. This is because some portion of the purchase price goes to the tax collector. When the government purchases an input subject to sales tax, should the producer's or purchaser's price be used in the cost calculations? The basic principle is the same as that for the monopoly case. If production is expected to expand, then the producer's supply price is appropriate. If production is expected to stay constant, the consumer's price should be used. A combination of responses requires a weighted average.

Unemployment If a worker for a public sector project is hired away from a private job, then society's opportunity cost is the worker's wage rate in the private sector, because it reflects the value of the lost output that the worker had been producing. Things get trickier when the project employs someone who is currently involuntarily unemployed. Because hiring an unemployed worker does not lower output elsewhere in the economy, the wage the worker is paid by the government does not represent an opportunity cost. All that is forgone when the worker is hired is the leisure he or she was consuming, the value of which is presumably low if the unemployment is involuntary. There are two complications, however: (1) If the government is running its stabilization policy to maintain a constant rate of employment, hiring an unemployed worker may mean reducing employment and output elsewhere in the economy. In this case, the social cost of the worker is his or her wage. (2) Even if the worker is involuntarily unemployed when the project begins, he or she may not necessarily be so during its entire duration. But forecasting an individual's future employment prospects is difficult. In light of the current lack of

shadow price

The underlying social marginal cost of a good.

Figure 8.1
Measuring the change in consumer surplus
A government irrigation project reduces the cost of producing avocados, thus shifting the supply curve to S'_a . The reduction in the price increases consumer surplus by $bcdg$.



consensus on the causes and nature of unemployment, the pricing of unemployed resources remains a problem with no agreed-on solution. In the absence of a major depression, valuation of unemployed labor at the going wage is probably a good approximation for practical purposes.

Consumer Surplus

A private firm is generally small relative to the economy, so changes in its output do not affect the market price of its product. In contrast, public sector projects can be so large that they change market prices, and this affects the way in which benefits should be calculated. For example, a government irrigation project could lower the marginal cost of agricultural production so much that the market price of food falls. But if the market price changes, how should the additional amount of food be valued—at its original price, at its price after the project, or at some price in between?

The situation for a hypothetical avocado-growing region is depicted in Figure 8.1. Pounds of avocados are measured on the horizontal axis, the price per pound is measured on the vertical, and D_a is the demand schedule for avocados. Before the irrigation project, the supply curve is labeled S_a , and market price and quantity are \$2.89 and A_0 , respectively. (The supply curve is drawn horizontally for convenience. The main points would still hold even if it sloped upward.)

Suppose that after more land is brought into production by the irrigation project, the supply curve for avocados shifts to S'_a . At the new equilibrium, the price falls to \$1.35, and avocado consumption increases to A_1 . How much better off are consumers? Another way of stating this question is, "How much would consumers be willing to pay for the privilege of consuming A_1 pounds of avocados at price \$1.35 rather than A_0 pounds at price \$2.89?"

The economic tool for answering this question is **consumer surplus**—the amount by which the sum that individuals would have been willing to pay exceeds the sum they actually have to pay. As shown in the appendix to this book, consumer surplus is measured by the area under the demand curve and above a horizontal line at the market price. Thus, when the price is \$2.89, consumer surplus is ebd .

When the price of avocados falls to \$1.35 because of the irrigation project, consumer surplus is still the area under the demand curve and above a horizontal line at the going price, but because the price is now \$1.35, the relevant area is ecg . Consumer surplus has increased by the difference between areas ecg and ebd —area $bcdg$. Thus, the area behind the demand curve between the two prices measures the value to consumers of being able to purchase avocados at the lower price. Provided the planner can estimate the shape of the demand curve, the project's benefit can be measured.

If the supply curve of the commodity under consideration is upward sloping, then changes in producer surplus (also explained in the appendix at the end of the book) can be brought into play. For example, in the cost-benefit analysis of rent controls, the change in landlords' surplus could be estimated given information on the shape of the supply curve of rental housing.

Inferences from Economic Behavior

So far we have been dealing with cases in which market data can serve as a starting point for valuing social costs and benefits. Sometimes the good in question is not explicitly traded, so no market price exists. We discuss two examples of how people's willingness to pay for such commodities can be estimated.

The Value of Time One important component of Boston's Big Dig project mentioned at the beginning of this chapter was a 3.5-mile stretch of highway that cost \$6.5 billion. It was estimated that with the new highway in place, the ride from downtown to the airport would be reduced from 45 minutes to 8 minutes. Was this a good deal? While it is true that "time is money," to do cost-benefit analysis we need to know *how much* money. A common way to estimate the value of time is to take advantage of the theory of leisure-income choice. People who have control over the amount they work do so up to the point where the subjective value of leisure is equal to the income they gain from one more hour of work—the after-tax wage rate. Thus, the after-tax wage can be used to value the time that is saved.⁹

Although this approach is useful, it has two major problems: (1) Some people cannot choose their hours of work. Involuntary unemployment represents an extreme case. (2) Not all uses of time away from the job are equivalent. For example, to avoid spending time on the road, a person who hated driving might be willing to pay at a rate exceeding his wage. On the other hand, a person who used the road for pleasure drives on weekends might not care very much about the opportunity cost of time, particularly if she could not work on weekends anyway.

Several investigators have estimated the value of time by looking at people's choices between modes of transportation that involve different traveling times. Suppose that in a given community people can commute to work either by bus or by train. The train takes less time, but it is more expensive. By seeing how much extra

consumer surplus

The amount by which consumers' willingness to pay for a commodity exceeds the sum they actually have to pay.

⁹For further details, see Chapter 18 under "Labor Supply."

money people are willing to pay for the train, we can infer how much they are willing to pay to reduce their commuting time, and hence how they value that time. Of course, other characteristics of people, such as their incomes, affect their choice of travel mode. Statistical techniques like those described in Chapter 2 can be used to take these variables into account. On the basis of several such studies, a reasonable estimate of the effective cost of traveling time is about 50 percent of the after-tax wage rate (see von Wartburg and Waters [2004]).

The Value of Life In the wake of September 11, a fund was set up to compensate the families of the victims. A journalist asked the man in charge of the fund, "What is a life worth?" He responded, "You'd have to be a rabbi or a priest to try to answer that" [Henriques, 2001, p. WK10]. Indeed, our religious and cultural values suggest that life is priceless. Consider the events that transpired a few years ago when a 22-month-old boy fell into an abandoned well. A rescue team with special training worked for 13 hours through the night to dig a separate hole to reach the trapped baby. A camera was dropped into the well to monitor the boy, and paramedics and a doctor were also on hand through the night to give medical advice. In the news accounts of this story, not a single person questioned whether saving the child's life was worth the cost. Arguing that any price was too high for saving his life would have been unthinkable. Similarly, if you were asked to value your own life, it would not be surprising if only the sky was the limit.

Such a position presents obvious difficulties for cost-benefit analysis. If the value of life is infinite, any project that leads even to a single life being saved has an infinitely high present value. *This leaves no sensible way to determine the admissibility of projects.* If every road in America were a divided four-lane highway, traffic fatalities would doubtless decrease. Would this be a good project? Similarly, any project that cost even one life would have an infinitely low value. In this context, consider the fact that to meet government mandated fuel efficiency standards, automobile manufacturers produce lighter cars than would otherwise be the case. But lighter cars are associated with higher fatality rates in accidents. Do fuel standards therefore automatically fail cost-benefit tests?

Economists have considered two methods for assigning finite values to human life, one based on lost earnings and the other on the probability of death.

Lost Earnings Under the lost earnings method, the value of life is the present value of the individual's net earnings over a lifetime. If an individual dies as a consequence of a given project, the cost to society is just the expected present value of the output that person would have produced. This approach is often used in law courts to determine how much compensation the relatives of accident fatalities should receive. However, taken literally, this approach means that society would suffer no loss if the aged, infirm, or severely handicapped were summarily executed. This implication is sufficiently bizarre that the method is rejected by most economists.

Probability of Death A second approach has as its starting point the notion that most projects do not actually affect with *certainty* a given individual's prospects for living. Rather, it is more typical for a change in the *probability* of a person's death to be involved. For example, you do not know that cancer research will save *your* life. All that can be determined is that it may reduce the *probability* of your death. The reason this distinction is so important is that even if people view their lives as having

infinite value, they continually accept increases in the probability of death for finite amounts of money. An individual driving a light car is subject to a greater probability of death in an auto accident than someone in a heavy car, other things being the same. People are willing to accept the increased risk of death because of the money they save by purchasing lighter cars.

Another way that people reveal their risk preferences is by their occupational choices. Some jobs involve a higher probability of death than others. Suppose we compare two workers who have identical job qualifications (education, experience, etc.), but one has a riskier job than the other. The individual in the riskier job is expected to have a higher wage to compensate for the higher probability of death. The difference between the two wages provides an estimate of the value that people place on a decreased probability of death.¹⁰

In the same spirit, there have been many studies of the amounts that people are willing to pay for safety devices, such as smoke alarms, that reduce the probability of death by a given amount. Different studies come up with quite different results, but a rough guess on the basis of such research is that the value of a life is between \$4 million and \$10 million [Viscusi, 2006]. Now, you might think that this range is so great as to be useless. However, these estimates can be very useful in weeding out senseless projects. For example, the regulations relating to the emergency floor lights on commercial planes cost about \$900,000 per life saved. These regulations clearly pass the admissibility criterion. On the other hand, governmental asbestos removal rules cost more than \$100 million per life saved.

An appealing aspect of this approach to valuing life is that it puts the analysis on the same willingness-to-pay basis that is so fruitful in other contexts. It remains highly controversial, however. Critics have argued that the probabilistic approach is irrelevant once it is conceded that *some* people's lives are *certainly* going to be at stake. The fact that we happen to be ignorant of just who will die is beside the point. This position leads us back to where we started, with no way to value projects that involve human life.

This academic controversy has become a matter of public concern because of various proposals to subject government safety and environmental regulations to cost-benefit analysis. In an attack on one proposal, an environmental lobbyist stated that one simply could not "put a price on saving lives" [Wetstone, 1995]. Unfortunately, in a world of scarce resources, we have no choice in the matter. The only question is whether or not sensible ways for setting the price are used.

Valuing Intangibles

No matter how ingenious the investigator, some benefits and costs are impossible to value: One of the benefits of the space shuttle program is increased national prestige. Indeed, President George W. Bush argued that space exploration "is a desire written in the human heart." Creating national parks gives people the thrill of enjoying beautiful scenery. The mind boggles at putting a dollar value on these "commodities." Three points must be kept in mind when intangible items might be important.

¹⁰ See Viscusi and Aldy [2003] for further discussion of such estimates.

First, intangibles can subvert the entire cost-benefit exercise. By claiming that they are large enough, *any* project can be made admissible. A journalist commenting on Britain's deliberations about whether to construct a tunnel below the English channel gave this advice: "Build it, not because dreary cost-benefit analysis says it will pay but because Britain needs a big project to arouse it" [Will, 1985]. However, presumably anyone who favors a particular project can make a case on the basis of its ability to "arouse." How does one then choose among projects? (The channel tunnel, of course, was ultimately built. When completed in 1994, it cost \$15 billion, more than twice the original estimate.)

Second, the tools of cost-benefit analysis can be used to force planners to reveal limits on how they value intangibles. Suppose the space shuttle's measurable costs and benefits are C and B , respectively, and its intangible benefits, such as national prestige, are an unknown amount X . Then if the measured costs are greater than measured benefits, X must exceed $(C - B)$ for the program to be admissible. Such information may reveal that the intangible is not valuable enough to merit doing the project. If $(C - B)$ for the space shuttle were \$10 million per year, people might agree that its contribution to national prestige was worth it. But if the figure were \$10 billion, a different conclusion might emerge.

Finally, even if measuring certain benefits is impossible, there may be alternative methods of attaining them. Systematic study of the costs of various alternatives should be done to find the cheapest way possible. This is sometimes called **cost-effectiveness analysis**. Thus, while one cannot put a dollar value on national security, it still may be feasible to subject the costs of alternative weapons systems to scrutiny.

cost-effectiveness
analysis

Comparing the costs of
the various alternatives
that attain similar
benefits to determine
which one is the
cheapest.

► GAMES COST-BENEFIT ANALYSTS PLAY

In addition to the problems we have already discussed, Tresch [2002] has noted a number of common errors in cost-benefit analysis.

The Chain-Reaction Game

An advocate for a proposal can make it look especially attractive by counting secondary profits arising from it as part of the benefits. If the government builds a road, the primary benefits are the reductions in transportation costs for individuals and firms. At the same time, though, profits of local restaurants, motels, and gas stations increase. This leads to increased profits in the local food, bed-linen, and gasoline-production industries. If enough secondary effects are added to the benefit side, eventually a positive present value can be obtained for practically any project.

This procedure ignores the fact that the project may induce losses as well as profits. After the road is built, the profits of train operators decrease as some of their customers turn to cars for transportation. Increased auto use may bid up the price of gasoline, decreasing the welfare of many gasoline consumers.

In short, the problem with the chain-reaction game is that it counts as benefits changes that are merely transfers. The increase in the price of gasoline, for example, transfers income from gasoline consumers to gasoline producers, but it does not represent a net benefit of the project. As noted later, distributional considerations

may indeed be relevant to the decision maker. But if so, consistency requires that if secondary benefits are counted, so should secondary losses.

The Labor Game

In 2005, Congress debated bills that authorized billions of dollars for a variety of transportation and energy projects. Opponents of the bills argued that their costs were excessive. In response, Representative Roy Blunt of Missouri declared, "This is about jobs, jobs, jobs, jobs, jobs" [Hulse, 2005, p. A20].

The congressman's statement is a typical example of the argument that some project should be implemented because of all the employment it "creates." Essentially, the wages of the workers employed are viewed as *benefits* of the project. This is absurd, because wages belong on the cost, not the benefit, side of the calculation. Of course, as already suggested, it is true that if workers are involuntarily unemployed, their social cost is less than their wage. Even in an area with high unemployment, it is unlikely that all the labor used in the project would have been unemployed, or that all those who were unemployed would have remained so for a long time.

The Double-Counting Game

Suppose that the government is considering irrigating some land that currently cannot be cultivated. It counts as the project's benefits the sum of (1) the increase in value of the land, *and* (2) the present value of the stream of net income obtained from farming it. The problem here is that a farmer can *either* farm the land and take as gains the net income stream *or* sell the land to someone else. Under competition, the sale price of the land just equals the present value of the net income from farming it. Because the farmer cannot do both simultaneously, counting both (1) and (2) represents a doubling of the true benefits.

This error may seem so silly that no one would ever commit it. However, Tresch [2002, p. 825] points out that at one time double counting was the official policy of the Bureau of Reclamation within the US Department of the Interior. The bureau's instructions for cost-benefit analysts stipulated that the benefits of land irrigation be computed as the *sum* of the increase in land value and the present value of the net income from farming it.

► DISTRIBUTIONAL CONSIDERATIONS

In the private sector, normally no consideration is given to the question of who receives the benefits and bears the costs of a project. A dollar is a dollar, regardless of who is involved. Some economists argue that the same view be taken in public project analysis. If the present value of a project is positive, it should be undertaken regardless of who gains and loses. This is because as long as the present value is positive, the gainers *could* compensate the losers and still enjoy a net increase in utility. This notion, sometimes called the **Hicks-Kaldor criterion**,¹¹ thus bases project selection on whether there is a *potential* Pareto improvement. The actual compensation does not have to take place. That is, it is permissible to impose costs on some members of society if that provides greater net benefits to other individuals.

Hicks-Kaldor criterion

A project should be undertaken if it has a positive net present value, regardless of the distributional consequences.

¹¹Named after the economists John Hicks and Nicholas Kaldor.

Others believe that because the goal of government is to maximize social welfare, the distributional implications of a project should be taken into account. Moreover, because it is the actual pattern of benefits and costs that really matters, the Hicks-Kaldor criterion does not provide a satisfactory escape from grappling with distributional issues.

One way to avoid the distributional problem is to assume the government can and will costlessly correct any undesirable distributional aspects of a project by making the appropriate transfers between gainers and losers.¹² The government works continually in the background to ensure that income stays optimally distributed, so the cost-benefit analyst need be concerned only with computing present values. Again, reality gets in the way. The government may have neither the power nor the ability to distribute income optimally.¹³ (See Chapter 12.)

Suppose the policymaker believes that some group in the population is especially deserving. This distributional preference can be taken into account by assuming that a dollar benefit to a member of this group is worth more than a dollar going to others in the population. This, of course, tends to bias the selection of projects in favor of those that especially benefit the preferred group. Although much of the discussion of distributional issues has focused on income as the basis for classifying people, presumably characteristics such as race, ethnicity, and gender can be used, as well.

After the analyst is given the criteria for membership in the preferred group, she must face the question of precisely how to weight benefits to members of that group relative to the rest of society. Is a dollar to a poor person counted twice as much as a dollar to a rich person, or 50 times as much? The resolution of such issues depends on value judgments. All the analyst can do is induce the policymaker to state explicitly his value judgments and understand their implications.

A potential hazard of introducing distributional considerations is that political concerns may come to dominate the cost-benefit exercise. Depending on how weights are chosen, any project can generate a positive present value, regardless of how inefficient it is. In addition, incorporating distributional considerations substantially increases the information requirements of cost-benefit analysis. The analyst needs to estimate not only benefits and costs but also how they are distributed across the population. As we discuss in Chapter 12, it is difficult to assess the distributional implications of government fiscal activities.

► UNCERTAINTY

In 2005, the levees protecting New Orleans were breached during Hurricane Katrina, leading to disastrous flooding. This catastrophe serves as a grim reminder of the fact that the outcomes of public projects are uncertain. Many important debates over project proposals center around the fact that no one knows how they will turn out. How much will a job-training program increase the earnings of welfare recipients? Will a high-tech weapons system function properly under combat conditions?

¹² *Costlessly* in this context means that the transfer system costs nothing to administer, and the transfers are done in such a way that they do not distort people's behavior (see Chapter 15).

¹³ Moreover, as the government works behind the scenes to modify the income distribution, relative prices probably change. But as relative prices change, so do the benefit and cost calculations. Hence, efficiency and equity issues cannot be separated as neatly as suggested here.

Suppose that two projects are being considered. They have identical costs, and both affect only one citizen, Kyle. Project *X* guarantees a benefit of \$1,000 with certainty. Project *Y* creates a benefit of zero dollars with a probability of one-half, and a benefit of \$2,000 with a probability of one-half. Which project does Kyle prefer?

Note that *on average*, *X* and *Y* have the same benefit. This is because the expected benefit from *Y* is $\frac{1}{2} \times \$0 + \frac{1}{2} \times \$2,000 = \$1,000$. Nevertheless, if Kyle is risk averse, he prefers *X* to *Y*. This is because project *Y* subjects Kyle to risk, while *X* is a sure thing. In other words, if Kyle is risk averse, he would be willing to trade project *Y* for a *certain* amount of money less than \$1,000—he would give up some income in return for gaining some security. The most obvious evidence that people are in fact willing to pay to avoid risk is the widespread holding of insurance policies of various kinds. (See chapter 9.) Therefore, when the benefits or costs of a project are risky, they must be converted into **certainty equivalents**—the amount of *certain* income the individual would be willing to trade for the set of uncertain outcomes generated by the project. The computation of certainty equivalents requires information on both the distribution of returns from the project and how risk averse the people involved are. The method of calculation is described in the appendix to this chapter.

The calculation of certainty equivalents presupposes that the random distribution of costs and benefits is known in advance. In some cases, this is a reasonable assumption. For example, engineering and weather data could be used to estimate how a proposed dam would reduce the probability of flood destruction. In many important cases, however, it is hard to assign probabilities to various outcomes. There is not enough experience with nuclear reactors to gauge the likelihood of various malfunctions. Similarly, how do you estimate the probability that a new AIDS vaccine will be effective? As usual, the best the analyst can do is to make explicit his or her assumptions and determine the extent to which substantive findings change when these assumptions are modified.

► AN APPLICATION: ARE REDUCTIONS IN CLASS SIZE WORTH IT?

In Chapter 7 we discussed research on the effect of class size on students' test scores. A related literature examines whether children in smaller classes have higher earnings as adults, other things being the same. In one econometric analysis of the relationship between class size and earnings, Card and Krueger [1996] estimated that a 10 percent reduction in class size is associated with future annual earnings increases of 0.4 to 1.1 percent. If it is correct, this estimate suggests that decreasing class size does produce monetary benefits.

By itself, though, this does not tell us whether implementing reductions in class size would be a sensible policy. After all, making classes smaller is costly—more teachers need to be hired, additional classrooms built, and so on. Do the benefits outweigh the costs? Peltzman [1997] employs the tools of cost-benefit analysis to address this question. His analysis illustrates several of the key issues raised in this chapter.

Cost-benefit analysis entails selecting a discount rate and specifying the costs and benefits for each year. We now discuss in turn how Peltzman deals with each of these problems.

certainty equivalent

The value of an uncertain project measured in terms of how much certain income an individual would be willing to give up for the set of uncertain outcomes generated by the project.

Discount Rate

Theoretical considerations do not pin down a particular discount rate, so Peltzman follows the sensible practice of selecting a couple and seeing whether the substantive results are sensitive to the difference. The (real) rates he chooses are 3 percent and 7 percent.

Costs

Peltzman assumes that a 10 percent reduction in class size would require 10 percent more of all inputs used in public school education—teachers, classroom space, equipment, and so on. Thus, a permanent reduction in class size of 10 percent would increase yearly costs by 10 percent. In 1994, the average cost per student in US public schools was about \$6,500, so a 10 percent increase is \$650. This cost is incurred for each of the 13 years that the student is in school. Because these costs are incurred over time, they must be discounted. Row (1) of Table 8.3 shows the present value of \$650 over a 13-year period for both $r = 3$ percent and $r = 7$ percent. In our earlier notation, these figures represent C , the present value of the project's costs (per student), at each discount rate.

This calculation of C involves a variety of simplifications; one of the most important is that the costs per year of schooling are constant. In fact, per-student costs are typically higher in high school than in elementary school. Allocating a greater proportion of the costs to future years would tend to reduce their present value.

Benefits

As noted earlier, Card and Krueger [1996] estimate that the range of returns to an increase in class size is 0.4 to 1.1 percent. Peltzman takes the midpoint of this range, 0.75 percent. He assumes that individuals go to work immediately upon leaving school, and work for the next 50 years. Hence, earnings are increased by 0.75 percent for each of the next 50 years. In 1994 median annual earnings for male workers 25 and older were \$30,000; increasing this sum by 0.75 percent implies a raise of \$225 per year over a 50-year period. Just like the costs, the benefits must be discounted. Note that the first of these \$225 flows occurs 13 years in the future; hence its present value is $\$225/(1 + r)^{13}$. The present values

Table 8.3 Costs and benefits of reducing class sizes by 10 percent

	Present Value	
	$r = 7\%$	$r = 3\%$
(1) Costs (\$650 annually for 1994 through 2006)	\$5,813	\$7,120
(2) Benefits (\$225 annually for 2007 through 2056)	\$1,379	\$4,060
(3) Benefits minus costs	-\$4,434	-\$3,060

Source: Computations based on Peltzman [1997].

These estimates suggest that the costs of reducing class size by 10 percent outweigh the benefits, at either a 3 or 7 percent discount rate.

of the benefits per student (B) for both discount rates are recorded in row (2) of the table.

Just as was true on the cost side, the calculation of benefits involves a number of important simplifications. Men generally earn more than women, so that using median earnings for males imparts an upward bias to the estimate of the benefits. Another issue is that earnings typically increase over time instead of staying constant. Further, the analysis ignores nonmonetary returns to education, which might include a reduced likelihood to commit crime, better informed choices in elections, and so on. To the extent that such effects are present, Peltzman's estimates of the social benefits to education are too low.

The Bottom Line and Evaluation

Computation of the net present value of this project is now straightforward. For each discount rate, take the benefit figure in row (2) of Table 8.3 and subtract from it the cost in row (1). These computations, recorded in row (3), reveal that when r is 7 percent, costs exceed benefits by \$4,434, and when r is 3 percent, costs exceed benefits by \$3,060. Thus, with either discount rate, $(B - C)$ is less than zero, and reducing class size by 10 percent fails the admissibility criterion. On this basis, Peltzman concludes, tongue-in-cheek, that students would be better off if class size were raised by 10 percent, and the savings used to give each student a bond that paid the market rate of interest (p. 226).

This analysis of class-size reductions illustrates some important aspects of practical cost-benefit analysis:

- The analysis is often interdisciplinary because economists alone do not have the expertise to evaluate all costs and benefits. Thus, for example, engineering studies would be required to determine what expenditures really would be needed to expand classroom capacity by 10 percent. Similarly, if one wanted to include crime reduction in the benefits, one would want to consult sociologists who study criminal behavior.
- Evaluation of costs and benefits, especially those arising in the future, is likely to require ad hoc assumptions. We noted earlier, for example, that Peltzman's simplifying assumption that earnings are constant over time is certainly not correct. But in order to do better, one needs an alternative assumption of how earnings will rise (or fall) over time, and it is not obvious how to do that.
- In situations characterized by so much uncertainty, it may overburden the analysis to include distributional considerations. For example, an investigator who cannot predict with much precision how class size affects earnings overall can hardly be expected to estimate the distribution of the benefits by income group.
- For all its limitations, cost-benefit analysis is a remarkably useful way to summarize information. It also forces analysts to make explicit their assumptions so that the reasons for their ultimate recommendation are clear. In the case of Peltzman's examination of class size reductions, for example, because some of the assumptions are questionable, the conclusions may ultimately be proven incorrect. Nevertheless, it is an extremely valuable exercise because it establishes a rational framework within which to conduct future discussions of this important issue.

► USE (AND NONUSE) BY GOVERNMENT

This chapter clearly indicates that cost-benefit analysis is not a panacea that provides a definitive “scientific” answer to every question. Nevertheless, it helps to ensure consistent decision making that focuses on the right issues. Have these methods been put to work by the government? The federal government has been ordering that various kinds of projects be subjected to cost-benefit analysis ever since the 1930s. Presidents Reagan, Bush, and Clinton each issued executive orders requiring cost-benefit analyses for all major regulations.

That said, both Democratic and Republican administrations often ignore or fudge orders to perform cost-benefit analyses, and the Congress has not been enthusiastic about getting them done either. Indeed, Hahn et al. [2000] studied 48 major federal health, safety, and environmental regulations issued in the late 1990s, and found that agencies quantified net benefits in less than a third of them. The agencies simply do not comply with the directives that require them to perform cost-benefit analyses. Why hasn't cost-benefit analysis had more effect on the style of government decision making? Part of the answer lies in the many practical difficulties in implementing cost-benefit analysis, especially when there is no consensus as to what the government's objectives are. In addition, many bureaucrats lack either the ability or the temperament to perform the analysis—particularly when it comes to their own programs. And neither are politicians particularly interested in seeing their pet projects subjected to scrutiny.

The story gets even worse when we consider the fact that, in certain vital areas, cost-benefit analysis has actually been expressly forbidden:

- The Clean Air Act prohibits costs from being considered when air quality standards are being set. In 1997, when the President's chief environmental aide was confronted with the fact that the costs of some new environmental regulations would exceed the benefits by hundreds of billions of dollars, she replied, “It is not at all about the money. . . . These are health standards” [Cushman, 1997, p. 28]. Any other stance would have been illegal!
- The same Act requires companies to install equipment that reduces pollution as much as is feasible, regardless of how small the benefits of the incremental reduction or how large the incremental costs of the equipment.
- The Endangered Species Act requires the Fish and Wildlife Service to protect every endangered species in the United States, regardless of the cost.
- The Food, Drug, and Cosmetic Act requires the Food and Drug Administration to ban any additive to food that may induce cancer in animals or humans, regardless of how tiny the risk or how important the benefits of the substance.

A 1995 attempt by several members of Congress to change some of these laws was defeated. Moreover, in 2001 the Supreme Court upheld the constitutionality of the Clean Air Act's prohibition of cost-benefit analysis. While this may have been the right decision from a legal perspective, it was unfortunate from a policy standpoint. Although cost-benefit analysis is surely an imperfect tool, it is the only analytical framework available for making consistent decisions. Forbidding cost-benefit analysis amounts to outlawing sensible decision making.

Summary

- Cost-benefit analysis is the practical use of welfare economics to evaluate potential projects.
- To make net benefits from different years comparable, their present value must be computed.
- Other methods—internal rate of return, benefit-cost ratio—can lead to incorrect decisions.
- Choosing the discount rate is critical in cost-benefit analyses. In public sector analyses, three possible measures are the before-tax private rate of return, a weighted average of before- and after-tax private rates of return, and the social discount rate. Choosing among them depends on the type of private activity displaced—investment or consumption—and the extent to which private markets reflect society's preferences.
- In practice, the US government applies discount rates inconsistently.
- The benefits and costs of public projects may be measured in several ways:
 - Market prices serve well if there is no strong reason to believe they depart from social marginal costs.
 - Shadow prices adjust market prices for deviations from social marginal costs due to market imperfections.
 - If labor is currently unemployed and will remain so for the duration of the project, the opportunity cost is small.
 - If large government projects change equilibrium prices, consumer surplus can be used to measure benefits.

For nonmarket commodities, the values can sometimes be inferred by observing people's behavior. Two examples are computing the benefits of saving time and the benefits of reducing the probability of death.

- Certain intangible benefits and costs simply cannot be measured. The safest approach is to exclude them in a cost-benefit analysis and then calculate how large they must be to reverse the decision.
- Cost-benefit analyses sometimes fall prey to several pitfalls:
 - Chain-reaction game—secondary benefits are included to make a proposal appear more favorable, without including the corresponding secondary costs.
 - Labor game—wages are viewed as *benefits* rather than *costs* of the project.
 - Double-counting game—benefits are erroneously counted twice.
- Including distributional considerations in cost-benefit analysis is controversial. Some analysts count dollars equally for all persons, while others apply weights that favor projects for selected population groups.
- In uncertain situations, individuals favor less-risky projects, other things being the same. In general, the costs and benefits of uncertain projects must be converted to certainty equivalents.

Discussion Questions

1. “If you were running the government, would you ask whether it would be cost-effective to make children's pajamas flame-resistant, or would you just order the manufacturers to do it? Would you be moved by the pleas of crib manufacturers who told you it would cost them a bundle to move those slats closer together?”

[Herbert, 1995]. How would you respond to these questions?

2. New Jersey recently instituted an enhanced auto emissions testing system at inspection sites throughout the state. According to news reports, the new tests increased waiting times from about 15 minutes to 2 hours. How should this