WARNING CONCERNING COPYRIGHT RESTRICTIONS

The copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproduction of copyrighted material.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If electronic transmission of reserve material is used for purposes in excess of what constitutes "fair use", that user may be liable for copyright infringement.

Citation: Rosen, Harvey S. Public Finance. Chicago: Irwin, 1995, pp. 238-269 & 507-516.

©: Richard D. Irwin, Inc., 1985, 1988, 1992, and 1995.

CHAPTER 12

Cost-Benefit Analysis

Paris is well worth a Mass.

attributed to Henri IV of France

Suppose that the government is presented with a project—building a road, starting a job-training program, or subsidizing mass transit—and must determine whether to pursue it. The theory of welfare economics provides a framework for dealing with such problems: Evaluate the social welfare function before and after the project, and see whether social welfare increases. If it does, then do the project.

This methodology 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 cost exceeds the marginal social benefit (see Chapter 4). However, cost-benefit analysis is not a panacea that provides a definitive "scientific" answer to every question. For example, when it comes to assigning dollar values to such

[1] Dreze and Stern [1987] discuss the links between theoretical welfare economics and cost-benefit analysis.

intangibles as national security or environmental purity, the tools of cost-benefit analysis provide no easy solutions. Nevertheless, using a cost-benefit analysis helps to ensure consistent decision making that focuses on the right issues.

Present Value

Project evaluation usually requires comparing costs and benefits from different time periods. For example, a program of preschool education for poor children requires substantial expenditures in the present and then yields returns in the future. In this section we discuss problems that arise in comparing dollar amounts from different time periods. Initially, we assume that no price inflation occurs. We show later how to take inflation into account.

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 $(1 + .05) \times \$100 = \$105$—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 $(1 + .05) \times \$105 = \$110.25$. This can also be written as $(1 + .05) \times (1 + .05) \times 100 = (1 + .05)^2 \times 100$. Similarly, if the money is deposited for three years, it will be worth $(1 + .05)^3 \times \$100$ 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.

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 $(1 + .05)$, just equals \$100. By definition, this is $\$100/(1 + .05)$, or approximately \$95.24. Thus, when the interest rate is 5 percent, the present value of \$100 payable one year from now is $\$100/(1 + .05)$. Note the symmetry with the familiar problem of projecting money into the future that we discussed already. 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 $\$100 \times (1 + .05)^2$. The most you would be willing to pay today for \$100 in two years is the amount that when multiplied by $(1 + .05)^2$ yields exactly \$100—$\$100/(1 + .05)^2$, 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 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 $(\$R_0 + \$R_1 + \cdots + \$R_T)$ 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 12.1 shows the present value of each year's payment. To find the present value *(PV)* of the income stream $\$R_0, \$R_1, \$R_2, \ldots, \$R_T$, we simply add the figures in the last column:

$$PV = R_0 + \frac{R_1}{(1 + r)} + \frac{R_2}{(1 + r)^2} + \cdots + \frac{R_T}{(1 + r)^T}. \qquad (12.1)$$

The importance of computing present values is hard to overestimate. Serious errors can be made if it is ignored. 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 $r = 10\%$, the present value is only \$148,644 [= $\$1{,}000{,}000/(1.10)^{20}$].

Inflation

We now consider how to modify the procedure when the price level is expected to increase in the future. To begin, consider a project that, in

[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 $\$R_T/[(1 + r_1) \times (1 + r_2) \times \cdots \times (1 + r_T)]$.

Table 12.1 **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	$(1 + r)^T$	$R_T/(1 + r)^T$

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, $\$\tilde{R}_1$, is $(1.07) \times \$R_0$. Similarly, two years into the future, the dollar value is $\$\tilde{R}_2 = (1.07)^2R_0$. In general, this same return has a dollar value in year T of $\$\tilde{R}_T = (1 + .07)^TR_0$.

The dollar values $\$\tilde{R}_0, \$\tilde{R}_1, \$\tilde{R}_2, \ldots, \$\tilde{R}_T$ are referred to as **nominal amounts.** Nominal amounts are valued according to the level of prices in the year the return occurs. It is possible to 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 $\$R_0, \$R_1, \$R_2, \ldots, \$R_T$, and inflation occurs at a rate of π per year, then the nominal returns are: $\$R_0, \$R_1 \times (1 + \pi), \$R_2 \times (1 + \pi)^2, \ldots, \$R_T \times (1 + \pi)^T$.

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 may be expected to increase 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 are increased. When expressed in *nominal* terms, the present value of the income stream is thus

[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 of the nominal rate. There are circumstances under which nominal interest rates may fail to rise by exactly the rate of inflation. See Chapter 16 under "Taxes and Inflation."

$$PV = R_0 + \frac{(1+\pi)R_1}{(1+\pi)(1+r)} + \frac{(1+\pi)^2 R_2}{(1+\pi)^2(1+r)^2} + \cdots + \frac{(1+\pi)^T R_T}{(1+\pi)^T(1+r)^T}. \tag{12.2}$$

A glance at Equation (12.2) indicates that it is equivalent to Equation (12.1) because all the terms involving $(1 + \pi)$ cancel out. The moral of the story is that the *same answer* is obtained 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 Rs, the discount rate must also be measured in real terms—the market rate of interest *minus* the expected inflation rate. Alternatively, if the market rate of interest is used for discounting, 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, 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, the benefits exceed the costs. If both projects are admissible, the firm 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):

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

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

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

where r represents the real rate of return that is generally available to the firm, the opportunity cost of its funds.

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

$$PV^Y = B_0^Y - C_0^Y + \frac{B_1^Y - C_1^Y}{1 + r} + \frac{B_2^Y - C_2^Y}{(1 + r)^2} + \cdots + \frac{B_{T'}^Y - C_{T'}^Y}{(1 + r)^{T'}}.$$

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 12.2, a research and development program (R&D) and drilling a new oil well. 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 well, on the other hand, has a single large payoff of $1,200 in three years.

The calculations show the discount rate chosen is important. For low values of r, the oil well is preferred to R&D. However, higher discount rates weigh against the oil well (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 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.

Table 12.2 **Comparing the present values of two projects**

	Annual Net Return			PV	
Year	*R&D*	*Oil Well*	*r =*	*R&D*	*Oil Well*
0	−$1,000	−$1,000	0	$150	$200
1	600	–0–	.01	128	165
2	–0–	–0–	.03	86	98
3	550	1,200	.05	46	37
			.07	10	−21

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 an advertising campaign and reaps a benefit of \$1.04 million in increased profits a year from now. If you were asked to compute the advertising campaign'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

$$B_0 - C_0 + \frac{B_1 - C_1}{1+\rho} + \frac{B_2 - C_2}{(1+\rho)^2} + \cdots + \frac{B_T - C_T}{(1+\rho)^T} = 0. \tag{12.3}$$

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 done. 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 \$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. 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 $-1{,}000 + 1080/1.06 = 18.87$. The present value criterion says that Y is preferable, as it should.

Sometimes it is impossible even to find a unique value of the internal rate of return. Consider the following three-period project:

Period	$B - C$
0	100
1	−260
2	165

Using Equation (12.3), we find this project's internal rate of return by solving

$$100 - \frac{260}{1 + \rho} + \frac{165}{(1 + \rho)^2} = 0.$$

This is a quadratic equation in ρ, and has *two* roots, $\rho = .1$ and $\rho = .5$.[4] The two values have quite different implications, and there is no obvious way to choose between them.

Benefit-Cost Ratio

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

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

and the present value of the costs, C, is

$$C = C_0 + \frac{C_1}{1 + r} + \frac{C_2}{(1 + r)^2} + \cdots + \frac{C_T}{(1 + r)^T}. \tag{12.4}$$

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 equivalent to the present value criterion for admissibility.

As a basis for comparing admissible projects, however, the benefit-cost ratio is virtually useless. Consider a community 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

[4] Multiplying the equation through by $(1 + \rho)^2$ yields $100(1 + \rho)^2 - 260(1 + \rho) + 165 = 0$, or $\rho^2 - 0.6\rho + 0.05 = 0$. Applying the quadratic formula, $\rho = (.6 \pm \sqrt{.36 - .2})/2 = (.6 \pm .4)/2 = .1$ or $.5$.

of 2. The town'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 (12.1) 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*.

- 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.

Public Sector Discount Rate

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. In the next, we 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 it may be difficult to pinpoint this rate, from a conceptual point of view there is agreement that the opportunity cost of funds to the firm 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, the 16 percent should be used as the discount rate. It is irrelevant whether or not this return is taxed.

[5] See Tresch [1981, Chapter 24] for further discussion of the alternative views.

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.

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. As just argued, for those funds coming out of investment, the before-tax rate of return is the opportunity cost and therefore the appropriate discount rate. But this is not so for funds that come at the expense of consumption. Consider Nelson, who is deciding how much to consume and how much to save this year. For each dollar Nelson 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 Nelson 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 Nelson is 16 percent, but he must pay 50 percent of the return to the government in the form of taxes. All that Nelson 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.

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 [see Harberger, 1974a]. In the preceding example, if one-quarter of the funds came at the expense of investment and three-quarters at the expense of consumption, then the public sector discount rate would be 10 percent (¼ × 16 percent + ¾ × 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. The funds are collected from a variety of taxes, each of which has a different effect on consumption and investment. 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

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.[6] But why should society's

[6] Complications arise in implementing this approach when public funds come at the expense of private investment. See Lind [1982].

view of the opportunity cost of forgoing consumption differ from the opportunity cost revealed in market rates of return? Several reasons have been suggested for believing that the social discount rate is lower.

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 a degree of omniscience and benevolence that is unrealistic. Moreover, even totally selfish individuals often find it in their personal interest to 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.[7]

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. Pigou [1932, Chapter 2] described this problem as a "defective telescopic faculty." 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 public 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, a process called **learning by doing** [Arrow, 1962]. In a sense, then, investment creates positive externalities, and by the usual kinds of arguments, investment is underprovided by private markets (see Chapter 6 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

[7] Why should people invest in a project whose returns may not be realized until after they are dead? 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.

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 6).

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.

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 *social* benefits and costs may not be reflected in market prices. There are several possibilities for measuring the benefits and costs of public sector projects.

Market Prices

As noted in Chapter 4, in a properly functioning competitive economy, the price of a good simultaneously reflects marginal social costs of its 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. In such a world, prices do not necessarily reflect marginal social costs and benefits. As McKean [1977, p. 123] has suggested, "There are enough things wrong with observed market prices to make one's hair stand on end."

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 cost.[8] The **shadow price** of such a commodity is its

[8] This section is based on Layard [1977, pp. 18–22] and Gramlich [1990].

underlying social 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. Suppose a public project uses a monopolistically produced input. In contrast to perfect competition, under which price is equal to marginal cost, a monopolist's price is above marginal cost (see Chapter 4). Should the government value the input 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 the input 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 of the input will be produced, the government's use comes at the expense of private consumers, whose value of the input is measured by the 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 purchaser's price should be used. For a combination of responses, a weighted average is required.

Unemployment. Like most microeconomic tools, cost-benefit analysis generally assumes that all resources are fully employed. Nevertheless, a project may involve hiring workers who are currently involuntarily unemployed. Because hiring an unemployed worker does not lower output elsewhere in the economy, the wage the worker is paid 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, she will not necessar-

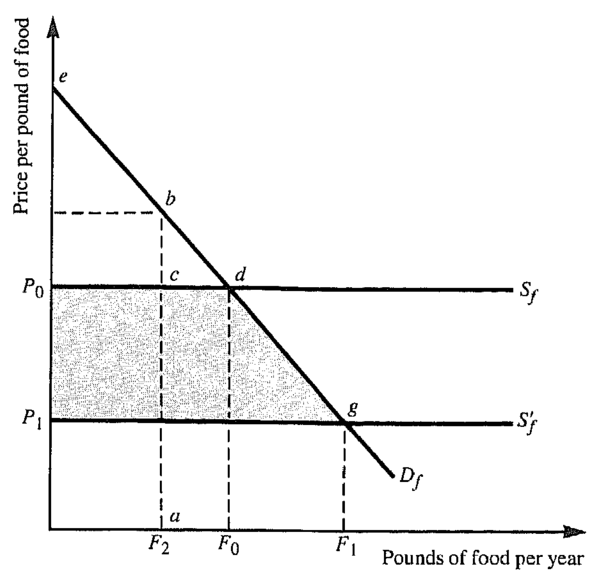

Figure 12.1
Measuring the change in consumer surplus

ily continue to be so during its entire duration. But forecasting an individual's future employment prospects is difficult. In light of the current lack of 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

In many cases, private firms are small relative to the economy. Therefore, they do not have to be concerned that changes in the amount they produce will affect the market price of their product. In contrast, public sector projects can be so large that they induce changes in market prices. 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?

This situation is depicted in Figure 12.1. Pounds of food are measured on the horizontal axis, the price per pound is measured on the vertical, and D_f is the demand schedule for food. Before the irrigation project, the supply curve is labeled S_f, and market price and quantity are P_0 and F_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 food shifts to S'_f. At the new equilibrium, the price falls to P_1, and food consumption increases to F_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 F_1 pounds of food at price P_1 rather than F_0 pounds at price P_0?"

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 Chapter 4, consumer surplus is measured by the area under the demand curve and above a horizontal line at the market price. Thus, when price is P_0, consumer surplus is P_0ed.

When the price of food falls to P_1 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 P_1, the relevant area is P_1eg. Consumer surplus has increased by the difference between areas P_1eg and P_0ed—area P_1P_0dg. Thus, the area behind the demand curve between the two prices measures the value to consumers of being able to purchase food at the lower price. Provided the planner can estimate the shape of the demand curve, the benefit of the project can be measured.

Inferences from Economic Behavior

We have so far 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. Suppose that if the government builds a road, it will save each traveler half an hour a day. 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 but common 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

[9] For further details, see Chapter 17 under "Labor Supply."

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 is more expensive. By seeing how much extra 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 similar to those described in Chapter 3 can be used to take these other variables into account. On the basis of several such studies, a reasonable estimate of the effective cost of traveling time is about 60 percent of the before-tax wage rate. [See Small, 1983.]

The value of life. A newspaper reporter once asked the head of the Federal Occupational Safety and Health Administration: "How do you place dollar values on human life or health?" [Shabecoff, 1981, p. E9]. It's a hard question. Our religious and cultural values suggest that life is priceless. Consider the events that transpired when a construction crane fell on Brigitte Gerney as she walked along a Manhattan street several years ago:

> Hundreds of police officers rerouted traffic . . . Two cranes were brought from other boroughs to lift the one that had fallen. Doctors . . . set up a mobile hospital at the construction site. Emergency Service rescue workers risked their own lives to save hers. Once she was freed, the police halted traffic for 30 blocks . . . to speed her trip to the emergency room.
>
> No city official questioned how much the rescue effort cost the city, or whether saving Mrs. Gerney's life was worth the price. To do so would have been unthinkable. "There's no point where you say that's too expensive," said [a spokesman] for the New York City Police Department. [Greer, 1985]

Similarly, if a person were asked to value his or her own life, it would not be surprising if the response indicated that 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?

Economists have considered two methods for assigning finite values to human life. The first requires measuring the individual's lost earnings; the second requires estimating the value the individual puts on changes in the probability of death.

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, the president of the Amalgamated Clothing and Textiles Workers Union said, "It seems incredible that . . . anyone could seriously suggest trading lives for dollars" [Finley, 1981, p. A23]. Unfortunately, in a world of scarce resources, we have no choice in the matter. The only question is whether or not sensible ways for making the trade are used.

Valuing Intangibles

No matter how ingenious the investigator, some benefits and costs seem impossible to value: One of the benefits of the space program is increased national prestige. 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.

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? (By the way, the channel tunnel 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 sometimes be used to force planners to reveal limits on how they value intangibles. Suppose a space program'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.[12] Such information may reveal that the intangible is not valuable enough to merit doing the project. If $(C - B)$ for the space program 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 it is impossible to measure certain benefits, there may be alternative methods of attaining them. If so, systematic study of the costs of the various alternatives should be done to find the cheapest way possible. This is sometimes called **cost-effectiveness analysis.** Thus, while

[12] A similar strategy can be applied to intangible costs.

one cannot put a dollar value on national security, it still may be feasible to subject the costs of alternative weapons systems to scrutiny.

Some Pitfalls

In addition to the problems we have already discussed, Tresch [1981] has noted a number of common errors in cost-benefit analysis.

The Chain-Reaction Game

The idea is to make a proposal 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 probably 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 must be secondary losses.

The Labor Game

In 1991, the Congress debated and eventually passed a $151 billion bill to improve the quality of the nation's highways and mass transit systems. Although some of the debate concerned the benefits that would follow from improved highways, proponents of the bill emphasized the project would employ a lot of labor. As one congressman said, "This is a jobs bill as well as a transportation bill" [Cushman, 1991, p. B8].

This 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 [1981, p. 561] 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

So far, we have not considered the distributional effects of public projects. 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 have argued 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,** thus bases project selection on *potential* gains in social welfare. [See Hicks, 1940, and Kaldor, 1939.] 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 someone else.

Others believe that because the goal of government is to maximize social welfare (not profit), 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.[13] The government works continually in the background to ensure

[13] *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 economic behavior (see Chapter 14).

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.[14] (See Chapter 8.)

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 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 or her value judgments and understand their implications.

A potential hazard of introducing distributional considerations is that political concerns 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 noted in Chapter 8 under "Expenditure Incidence," it is difficult to assess the distributional implications of government fiscal activities.

Uncertainty

In 1981, the Federal Railroad Administration had to decide whether to install sensors on railroad bridges that could detect damage done to the bridges when boats accidentally crashed into their supports. The agency concluded the benefits were far less than the costs. In 1993, 44 people were killed near Mobile, Alabama, when a railroad bridge that had been damaged by a tugboat collapsed [Applebome, 1993, p. 6]. This incident is a dramatic reminder of the fact that the outcomes of public decisions cannot be predicted with certainty. 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 train-

[14] 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.

ees? Will a high-tech weapons system function properly under combat conditions?

Suppose that two projects are being considered. They have identical costs, and both affect only one citizen, Smith. 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 Smith prefer?

Note that on average, the benefit from Y is equal to that from X. This is because the expected benefit from Y is $\frac{1}{2} \times \$0 + \frac{1}{2} \times \$2,000 = \$1,000$. Nevertheless, if Smith is risk averse, she prefers X to Y. This is because project Y subjects Smith to risk, while X is a sure thing. In other words, if Smith is risk averse, she would be willing to trade project Y for a *certain* amount of money less than $1,000—she 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.

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 given weapons system will deter foreign aggression? 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

The White Cloud Peaks lie about 30 miles north of Sun Valley, Idaho.[15] This public area occupies about 20,000 acres, supports a large variety of wildlife, and with some development could provide excellent opportunities for hiking, camping, fishing, and hunting. The White Cloud Peaks also appear to have substantial deposits of the valuable metal molybdenum, and in the late 1960s, a company requested permission to mine these

[15] This section is based on Krutilla and Fisher [1975, chap. 7].

Lost earnings. The value of life is the present value of the individual's net earnings over a lifetime.[10] 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 many 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 very often 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. Garen [1988] estimated that increasing by one the number of fatalities per 100,000 workers in an occupation increases the yearly earnings in the occupation by about 0.55 percent.[11]

An appealing aspect of this approach 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. Broome [1978] has 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

[10] See Klarman [1965] for an example using this approach.

[11] See Viscusi [1993] for further discussion of such estimates. The cost-benefit analyst should also consider the psychological cost of bereavement to families and friends and the changes in financial status of relatives.

deposits. The land cannot be used for mining and for recreation simultaneously. Should the White Cloud Peaks be developed for recreational purposes? Should the area be mined instead? Or should it be left alone?

Krutilla and Fisher [1975] (K&F) employed the tools of cost-benefit analysis to answer these questions. We discuss here only one component of K&F's study—an examination of whether the incremental costs of developing the White Cloud Peaks for recreational use would exceed the incremental benefits. The analysis illustrates several of the issues raised in this chapter.

Estimating the costs and benefits of developing the White Cloud Peaks for recreation requires specifying just what *kind* of recreation. Constructing roads for Jeeps has very different costs and benefits than extending hiking trails. K&F assumed that the land would be used in such a way that its basic ecology would be preserved. Specifically, they assumed development would take the form of creating four additional trails, which would allow more access for visitors.

Earlier sections of this chapter indicated that cost-benefit analysis entails selecting a discount rate and specifying the costs and benefits for each year. We now discuss these in turn.

Discount Rate

Theoretical considerations do not pin down one particular discount rate, so K&F followed the sensible practice of selecting several and seeing whether the substantive results are sensitive to the differences. Next, we report their results for discount rates of 7 and 10 percent.[16]

Costs

Creating additional trails requires immediate outlays for equipment and labor, estimated to be $23,000 in 1971. This figure is recorded in row (1) of Table 12.3. Because these capital outlays are made immediately, their present value is also $23,000, regardless of the discount rate.

After the trails are put in, yearly expenditures must be made for their maintenance. The estimated annual cost is $12,092. Because these costs are incurred over time, however, they must be discounted. Row (2) of Table 12.3 shows the discounted present value of these maintenance expenditures over a 50-year period for both r = 7 percent and r = 10 percent. Note how the present value of maintenance is smaller with the higher discount rate. The total cost of extending the trails, given in row (3), is the present value of the capital and maintenance outlays.

If the goal is to expand recreational use *while maintaining the quality of the environment,* appropriate sanitation facilities are required. K&F estimated that initially nine toilets, at a cost of $750 per toilet, are needed. The total expenditure, $6,750, is recorded in row (4). The toilets have to

[16] K&F do not indicate whether these are intended to be real or nominal discount rates. Given their magnitudes, they appear to be nominal. However, as noted later, the benefits and costs are in real terms. As an exercise, students might want to explain what problems this creates.

Table 12.3 **Incremental costs for expanded recreational use of White Cloud Peaks**

	Present Value	
Trail Extensions	r = *7%*	r = *10%*
(1) Initial outlays: $23,000	$ 23,000	$ 23,000
(2) Maintenance outlays per year: $12,092	170,981	119,835
(3) Total present value of trail extensions costs (sum of rows *1* and *2*)	193,981	142,835
Sanitary Facilities		
Outlays for equipment:		
(4) In year 1, nine toilets at a cost of $750 per toilet	$ 6,750	$ 6,750
(5) In year 30, nine toilets	887	387
(6) In year 60, nine toilets	114	20
(7) Maintenance and operation outlays per year: $100	1,414	990
(8) Total present value of sanitation costs (sum of rows *4, 5, 6,* and *7*)	$ 9,165	$ 8,147
(9) Present value of all costs (sum of rows *3* and *8*)	$203,146	$150,982

SOURCE: Computations based on Krutilla and Fisher [1975, p. 163].

be replaced at 30-year intervals. Rows (5) and (6) show the present value of the same $6,750 toilet expenditure made 30 and 60 years in the future, respectively. In addition, the sanitation facilities have to be maintained at an annual cost of $100. The present value of this stream of expenditures at the two discount rates is recorded in row (7). Row (8) gives the sum of all the sanitation costs, and row (9) adds these to the costs of extending the trails. In our earlier notation [Equation (12.4)], the figures in row (9) represent C, the present value of the project's costs, at each discount rate.

Benefits

A key piece of information is required: How much would people be willing to pay for the privilege of using the new trails? To answer this question, K&F took advantage of previous econometric analysis of recreational demand for wilderness areas. The dependent variable was the willingness to pay for a wilderness trip as expressed by individuals in a survey. The right-hand side variables included the length of the trip, congestion at the camp site, and various socioeconomic characteristics of the respondents. On the basis of this regression, K&F estimated that in the initial year (1971), a typical user would be willing to pay $10 per day. They further estimated that the area would be used 4,600 recreation days that year, making the first year's benefit $46,000.

The analysis required estimates of benefits in future years as well. K&F assumed use of the area would grow initially by 10 percent per year, and willingness to pay would increase by 4 percent per year.[17] Of course, an uninterrupted growth rate of 10 percent a year over a long time would result in the area being overrun by visitors. On the basis of discussions with district rangers and other experts, K&F postulated a maximum number of possible users consistent with maintaining environmental quality. After this figure is reached, the annual number of visits is not allowed to grow.

On the basis of K&F's willingness-to-pay calculations, the benefits flowing from the increased expenditures for recreational use are as shown in Table 12.4. In our earlier notation, Table 12.4 gives the value of B, the present value of the benefits, for each discount rate.

Computation of the net present value of the project is now straightforward. For each discount rate, take the benefit value from Table 12.4, and subtract from it the cost figure from row (9) of Table 12.3. This computation reveals that when $r = 7$ percent, benefits minus costs are $787,854, and when $r = 10$ percent, benefits minus costs are $239,018. Thus, for either choice of discount rate, $(B - C)$ exceeds zero, and by the present value criterion, the project is admissible. This does not complete the analysis, however, because we still have to find out whether the net present value exceeds that of the alternative project, molybdenum mining. However, discussion of that part of the study would take us too far afield. Interested readers should consult Krutilla and Fisher [1975] for further details.

Comments

This White Cloud Peaks analysis illustrates some important aspects of practical cost-benefit studies:

1. The analysis is often interdisciplinary because economists alone do not have the technical expertise to evaluate all costs and benefits. Thus, engineers were required to predict construction

Table 12.4 **Incremental benefits for expanded recreational use of White Cloud Peaks**

Present Value	
$r = 7\%$	$r = 10\%$
$991,000	$390,000

SOURCE: Krutilla and Fisher [1975, p. 171].

[17] As use increases, some congestion occurs, which tends to lower people's willingness to pay. (A visit to a crowded wilderness area isn't as pleasant as a visit to a quiet one.) K&F estimated the dampening effect of increased congestion costs and factored it into their willingness-to-pay calculations.

costs of the trails, wilderness use experts to estimate the maximum number of campers the area can support, and so forth.

2. Evaluation of costs and benefits, especially those arising in the future, is likely to require ad hoc assumptions. As just shown, it is difficult enough to estimate *current* willingness to pay for a commodity that is not priced on the market; projecting such figures decades into the future is even tougher. In most cases, all the analyst can do is assume that willingness to pay follows some pattern of growth, and see whether the substantive results change when alternative reasonable patterns are postulated. Moreover, in situations characterized by so much uncertainty, it may overburden the analysis to include distributional considerations. An investigator who can barely predict the total number of users two decades in the future can hardly be expected to estimate the distribution of their incomes as well.
3. For all its limitations, cost-benefit analysis is a remarkably useful way of summarizing information. It also forces analysts to make explicit their assumptions so that the reasons for their ultimate recommendation are clear. Indeed, any other way of making decisions is implicitly just a cost-benefit analysis without its assumptions explicitly stated.

Use by Government

Much effort has been devoted to refining the techniques of cost-benefit analysis. Have these methods been put to work by the government? Stipulations that certain kinds of federal projects be subjected to cost-benefit analysis began appearing in the late 1930s.[18] However, not until the mid-1960s did such analysis receive major public interest. At that time, President Lyndon Johnson's Great Society (a catchall term for a number of social programs) was leading to a tremendous expansion in the scope and magnitude of government programs. With so many new projects being proposed, some kind of systematic evaluation procedure was needed.

The Johnson administration therefore mandated that each program's costs and benefits be computed, and that they be compared to the returns of relevant alternatives. However, this mandate had little impact on the style of government project selection. This was partly due to the many practical difficulties in implementing cost-benefit analysis, particularly when there is no consensus as to what the government's objectives are. In addition, many bureaucrats lacked either the ability or the temperament to perform the analysis—particularly when it came to their own

[18] Gramlich [1981, pp. 7–11] provides a concise and useful discussion of the history of cost-benefit analysis in the United States up to the early 1980s.

programs. And neither were politicians particularly eager to see their pet projects subjected to scrutiny.

The use of cost-benefit analysis received some encouragement in 1981, when Ronald Reagan issued an executive order requiring that all new regulations had to pass a cost-benefit test. George Bush issued a similar order in 1992. As a consequence, the preparation of analyses of the costs and benefits of regulation has become standard operating procedure, particularly in the environmental area. Although there have been defects in some of these analyses, the policy of requiring them seems to be a desirable step.[19]

On the other hand, in certain vital areas not only has cost-benefit analysis failed to take hold, but it also has been expressly forbidden:

- The Clean Air Act prohibits costs from being considered when air quality standards are being set.
- 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 might be 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.

Such laws are unfortunate. 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.
- Cost-benefit analyses consider net benefits received over time. To make these comparable, the present value of all future net benefits 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,

[19] See the essays in Smith [1984].

and the social discount rate. Choosing among them depends on the type of private activity displaced—investment or consumption—and the degree to which private markets are believed to reflect society's preferences.

- 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. However, forecasting unemployment is difficult.
  If large government projects change equilibrium prices, consumer surplus can be used to measure the effect on individuals.
  For nonmarketed commodities, no prices are available. In some instances, the values of these commodities can be inferred by observing people's behavior. Two examples are computing the benefits of saving time and the benefits of a reduced probability of death.
- Certain intangible benefits and costs simply cannot be measured. The safest approach is to exclude these in a cost-benefit analysis and then calculate how large intangibles 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 mistakenly counted twice.
- Whether distributional considerations belong 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. Because of the potential for political manipulation, distributional weights must be introduced explicitly.
- In uncertain situations, individuals favor less risky projects, other things being the same. In general, the costs and benefits of uncertain projects must be adjusted to reflect this.

Discussion Questions

1. The Army Corps of Engineers frequently employs the cost-benefit ratio for evaluating public projects. In Senate hearings on a breakwater project, one senator was impressed that the project had "an amazingly high 17 to 1 benefit-cost ratio" [US Congress, 1973, p. 1]. If you were on the senator's staff, what would you have said?
2. The city of Sundown is considering a plan to build a convention center to revitalize the downtown area. The project requires hiring workers in the construction union. Because of unionization, the workers' wage rate is above the value that would be determined by supply and demand. However, city analysts have econometrically estimated the wage that would prevail if the labor market were competitive.
   a. If the convention center increases the total number of construction workers employed, what is the appropriate cost of labor?
   b. Suppose instead that total employment is unaffected; workers leave private projects to work on the convention center. Does your answer to (*a*) change? Why?
3. A project yields an annual benefit of \$25 a year, starting next year and continuing forever. What is the present value of the benefits if the interest rate is 10 percent? [Hint: The infinite sum $x + x^2 + x^3 + \ldots$ is equal to $x/(1 - x)$, where x is a number less

than 1.] Generalize your answer to show that if the perpetual annual benefit is B and the interest rate is r, then the present value is B/r.

4. An outlay of $1,000 today yields an annual benefit of $80 beginning next year and continuing forever. There is no inflation and the market interest rate is 10 percent before taxes and 5 percent after taxes.
   a. What is the internal rate of return?
   b. Taxes levied to fund the project come entirely from consumer spending. Is the project admissible? Why? Suppose instead that taxes are collected by reducing private firms' investments. Is the project admissible in this case? Finally, suppose consumers spend 60 cents of their last dollar and save 40 cents. Is the project admissible now? Explain your calculations.
   c. Suppose the social discount rate is 4 percent. What is the present value of the project?
   d. Now suppose 10 percent annual inflation is anticipated over the next 10 years. How are your answers to (*a*), (*b*), and (*c*) affected?

5. Bill rides the subway at a cost of 75 cents per trip, but would switch if the price were any higher. His only alternative is a bus that takes five minutes longer, but costs only 50 cents. He makes 10 trips per year. The city is considering renovations of the subway system that would reduce the trip by 10 minutes, but fares would rise by 40 cents per trip to cover the costs. The fare increase and reduced travel time both take effect in one yea and last forever. The interest rate is 25 percent.
   a. As far as Bill is concerned, what are the present values of the project's benefits and costs?
   b. The city's population consists of 55,000 middle-class people, all of whom are identical to Bill, and 5,000 poor people. Poor people are either unemployed or have jobs close to their homes, so they do not use any form of public transportation. What are the total benefits and costs of the project for the city as a whole? What is the net present value of the project?
   c. Some members of the city council propose an alternative project that consists of an immediate tax of $1.25 per middle-class person to provide "free" legal services for the poor in both of the following two years. The legal services are valued by the poor at a total of $62,500 per year. (Assume this amount is received at the end of each of the two years.) What is the present value of the project?
   d. If the city must choose between the subway project and the legal services project, which should it select?
   **e.* What is the "distributional weight" of each dollar received by a poor person that would make the present values of the two projects just equal? That is, how much must each dollar of income to a poor person be weighted relative to that of a middle-class person? Interpret your answer.

* Difficult.

APPENDIX

Calculating the Certainty Equivalent Value

This appendix shows how to calculate the certainty equivalent value of an uncertain project.

Consider Jones, whose current earnings are E dollars. He enters a job-training program with an unpredictable effect on his future earnings. The program will leave his annual earnings unchanged with a probability of ½, or it will increase his earnings by y dollars, also with a probability of ½.[20] The benefit of the program is the amount that Jones would be willing to pay for it, so the key problem here is to determine that amount. A natural answer is $y/2$ dollars, the expected increase in his earnings.[21] However, this value is too high, because it neglects the fact that the outcome is uncertain and therefore subjects Jones to risk. As long as Jones does not like risk, he would give up some income in return for gaining some security. When the benefits or costs of a project are risky, they must be converted into **certainty equivalents,** the amounts of *certain* income that the individual would be willing to trade for the set of uncertain outcomes generated by the project.

The notion of certainty equivalence is illustrated in Figure 12.A. The horizontal axis measures Jones's income, and the vertical axis indicates the amount of his utility. Schedule OU is Jones's utility function, which shows the total amount of utility associated with each income level. Algebraically, the amount of utility associated with a given income level, I, is $U(I)$. The shape of the schedule indicates that as income increases, utility also increases, but at a declining rate—there is diminishing marginal utility of income.[22]

To find the amount of utility associated with any income level, simply go from the horizontal axis up to OU, and then off to the vertical axis. For example, if the training project yields no return so that Jones's income is E, then his utility is $U(E)$, as indicated on the vertical axis. Similarly, if the project succeeds so that Jones's income increases by y, his total income is $(E + y)$, and his utility is $U(E + y)$.

[20] Probabilities of ½ are used for simplicity. The general results hold regardless of the probabilities chosen.

[21] Expected earnings are found simply by multiplying each possible outcome by the associated probability and then adding: $(½ \times 0) + (½ \times y) = y/2$.

[22] If marginal utility were increasing, all the derived results would be reversed. The assumption of diminishing marginal utility is more plausible.

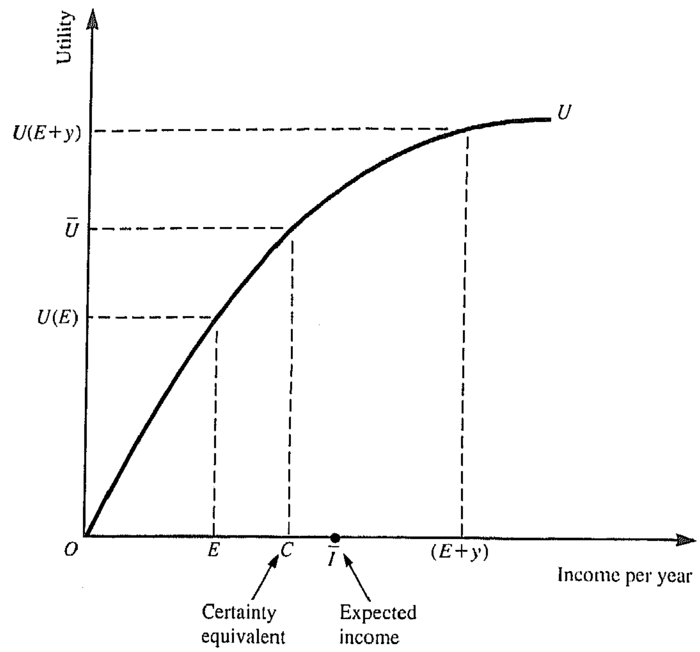

Figure 12.A
Computing the certainty equivalent of a risky project

Because each outcome occurs with a probability of ½, Jones's average or expected *income* is $E + y/2$, which lies halfway between E and $(E + y)$ and is denoted $\bar{I}$. However, what Jones really cares about is not expected income, but expected *utility*.[23] Expected utility is just the average of the utilities of the two outcomes, or $\frac{1}{2}U(E) + \frac{1}{2}U(E + y)$. Geometrically, expected utility is halfway between $U(E)$ and $U(E + y)$, and is denoted by $\bar{U}$.

We are now in a position to find out exactly how much certain money the job-training program is worth to Jones. All we have to do is find the amount of income that corresponds to utility level $\bar{U}$. This is shown on the horizontal axis as C, which is by definition the certainty equivalent. It is crucial to note that C is less than $\bar{I}$—the certainty equivalent of the job training program is *less* than the expected income. This is consistent with the intuition developed earlier. Jones is willing to pay a premium of $(\bar{I} - C)$ in exchange for the security of a sure thing.

We have shown, then, that proper evaluation of the costs and benefits of an uncertain project requires more than finding the project's expected value. The latter must be reduced by a risk premium that depends on the shape of the individual's utility function.

[23] Those who are familiar with the theory of uncertainty will recognize the implicit assumption that individuals have "Von Neumann-Morgenstern utility functions." See Henderson and Quandt [1980].

In a way, this is a disappointing outcome, because the computation of an expected value is much simpler than that of a certainty equivalent. Fortunately, it turns out that in many cases the expected value is enough. Suppose a new bomber is being considered, and because the technology is not completely understood, analysts are unsure of its eventual cost. The cost will be either $15 per family or $25, each with probability ½. Although in the aggregate a large amount of money is at stake, on a per-*family* basis, the sums involved are quite small compared to income. In terms of Figure 12.A, the two outcomes are very close to each other on curve *OU*. As points on *OU* get closer and closer together, the expected value and certainty equivalent become virtually identical, other things being the same. Intuitively, people do not require a risk premium to accept a gamble that involves only a small amount of income.[24]

Thus, for projects that spread risk over large numbers of people, expected values can provide good measures of uncertain benefits and costs. But for cases in which risks are large relative to individuals' incomes, certainty equivalents must be computed.

---

[24] As the points on *OU* come closer together, the shape of the curve between them becomes approximately linear. When the utility function is linear, the risk premium is always zero, a fact that is easily verifiable diagrammatically. See Arrow and Lind [1970].

CHAPTER 21

Public Finance in a Federal System

A community of a higher order should not interfere in the internal life of a community of a lower order, . . . but rather should support it in case of need and help to coordinate its activity with the activities of the rest of society, always with a view to the common good.

John Paul II

In the summer of 1982, the city of Glen Cove on Long Island banned diplomats from the former Soviet Union who lived there from using its tennis courts and other recreational facilities. The stated reason for the ban was the fact that the property inhabited by the Soviets was exempt from local property taxes.[1] Glen Cove's action caused an international furor, and the State Department requested that the city desist from meddling in foreign affairs. But the mayor responded, "Unless the State Department wants to pay up all the property taxes the Soviets have never had to pay like other Glen Cove residents, then the Russians will have to stay off the tennis courts."

While the Cold War is now over, this incident highlights three enduring issues that surround the operation of the US system of public finance.

- Subfederal units of government function with considerable autonomy. Attempts from the outside to change their behavior are likely to be met with active or passive resistance. Is decentralized decision making desirable?
- Different types of public services are customarily provided by various levels of government. The reason the Glen Cove incident received so much attention—and created such amusement (diplomats excepted)—was the incongruity of a local level of govern-

[1] Note, however, that the ban originated after federal officials said that the Soviets had installed eavesdropping equipment in their mansion. The controversy was finally settled in 1984, when the Soviets agreed to make some payment for city services.

Table 21.1 **Distribution of all US government expenditure by level of government, selected years**

	Federal	*State*	*Local*
1900	34.1%	8.2%	57.7%
1910	30.1	9.0	60.9
1920	39.7	9.8	50.5
1930	32.5	16.3	51.2
1938	45.5	16.2	38.3
1950	59.3	15.2	26.5
1960	57.6	13.8	28.6
1971	48.4	18.6	33.0
1980	54.9	18.1	27.0
1988	57.1	17.4	25.6
1991	55.5	18.6	26.0

SOURCE: Werner Pommerehne, "Quantitative Aspects of Federalism: A Study of Six Countries," in *The Political Economy of Fiscal Federalism,* ed. W. Oates (Lexington, MA: D.C. Heath, 1977), p. 311, except for figures after 1980, which are computed from various editions of US Bureau of the Census, *Statistical Abstract of the United States.*

point is that if local and state government spending behavior is constrained by the central government, the centralization ratio underestimates the true extent of centralization in the system. Conversely, if states and localities effectively lobby the federal government to achieve their own ends, the centralization ratio may overestimate the degree of decentralized economic power. In fact, a substantial amount of state and local spending is dictated by the federal government. The federal government simply mandates that the subfederal government provide certain services, but without a corresponding increase in financial support. These unfunded mandates, which cover areas as diverse as handicapped rights, hazardous waste disposal, and motor vehicle safety, are estimated to cost the states and localities $2 to $5 billion annually [Zimmerman, 1992, p. 5]. If we include mandated expenditures on Medicaid on the list, the total amounts to about 20 percent of state spending.

Table 21.1 shows how the distribution of United States government expenditure by level of government has been changing. The long-run trend has been for the centralization ratio to increase, although the movement upward has not been steady.

Figure 21.1 shows the division of public spending by level of government for various government functions. The figures indicate that a number of important activities are in the hands of state and local governments. In the context of the current debate over welfare reform, it is noteworthy that more than 75 percent of public welfare expenditures are made by state and local governments. This figure gives a somewhat exaggerated view of the importance of decentralized redistribution for two reasons:

ment in effect making foreign policy. International relations "belong" to the central government. On the other hand, decisions on the quantity and type of recreational facilities "belong" to localities. How should different functions be allocated to various levels of government?

- The story illustrates the crucial role of property taxes in the finance of US local governments: no property taxes, no tennis. Are locally raised taxes a good way for the services provided by municipalities to be financed? Or should the money come from the state and federal governments?

These are important issues in the United States where, as noted earlier, there are a multitude of governmental jurisdictions. The appropriate division of power among them has been a matter of controversy since the founding of the nation. This chapter examines the normative and positive aspects of public finance in a federal system. We devote special attention to the fiscal issues that currently confront localities.

Background

Oates [1972, p. 17] provides a useful economic definition of **federal government:**

> A public sector with both centralized and decentralized levels of decision making in which choices made at each level concerning the provision of public services are determined largely by the demands for those services of the residents of (and perhaps others who carry on activities in) the respective jurisdictions.

One federal system is characterized as being more centralized than another when more of its decision-making powers are in the hands of authorities with a larger jurisdiction. The most common measure of the extent to which a system is centralized is the **centralization ratio,** the proportion of total direct government expenditures made by the central government. ("Direct" government expenditure comprises all expenditure except transfers made to other governmental units.) Centralization ratios vary widely across nations. In France, it is 77 percent; in Canada, 41 percent; and in the United States, 55.5 percent.[2]

The centralization ratio is by no means a foolproof indicator. Suppose that states make expenditures for highways, but the money comes in the form of grants from the federal government. Congress decides that no state will receive highway grants unless it mandates a 55-mile-per-hour speed limit. Virtually every state complies.[3] Who is really in charge? The

---

[2] Computed from Organization for Economic Cooperation and Development [1993], except see Table 21.1 for the United States.

[3] In 1986, the Department of Transportation moved against Arizona and Vermont for not enforcing the speed limit. In 1987, the Congress gave states the option of increasing speed limits on some roads up to 65 miles per hour.

Figure 21.1
Percentage of expenditures by selected functions and levels of government.

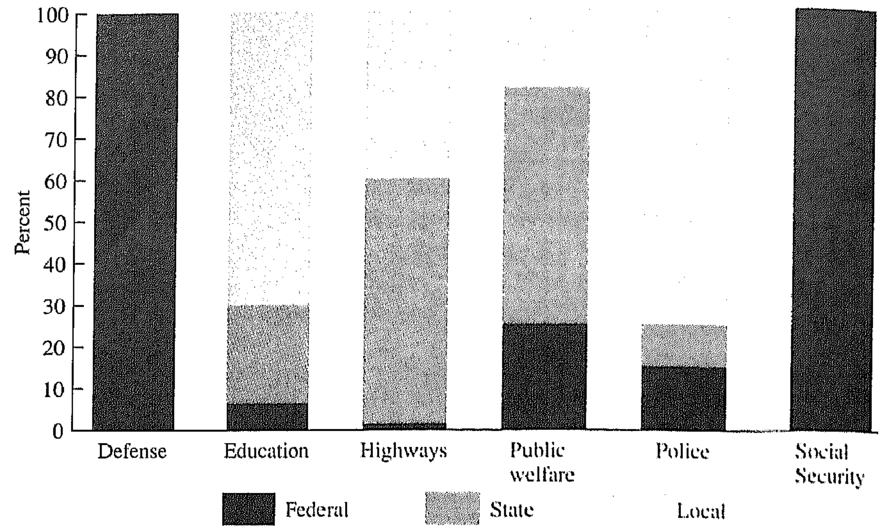

SOURCE: Computed from US Department of Commerce, Bureau of the Census, *Government Finances in 1990–1991* (Washington, DC: US Government Printing Office, 1993), p. 13. Figures are for fiscal year 1993.

(1) most of the financing (about three quarters) of these expenditures is done by the federal government; and (2) much of the spending is subject to federal standards on eligibility and benefit levels. Aid to Families with Dependent Children is an example of a program with extensive federal guidelines, even though states have some autonomy in setting benefit levels and standards.

Figure 21.1 leaves us with a critical question: Is the division of powers in the US fiscal system sensible? Before we can provide an answer, we need to discuss the special features associated with local government.

Community Formation

To understand the appropriate fiscal roles for local jurisdictions, we consider why communities are formed. In this context, it is useful to think of a community as a **club**—a voluntary association of people who band together to share some kind of benefit. This section develops a theory of clubs and uses that theory to explain how the size of a community and its provision of public goods are determined.[4]

Consider a group of people who wish to band together to purchase land for a public park. For simplicity, assume that all members of the

[4] Most club models are based on the work of Buchanan [1965]. For a survey of this area, see Sandler and Tschirhart [1980].

group have identical tastes and that they intend to share equally the use of the park and its costs. The "community" can costlessly exclude all nonmembers, and it operates with no transaction costs. Given the assumption of identical tastes, we need consider only the desires of a representative member. Two decisions must be made: how large a park to acquire and how many members to have in the community.

Assuming that it wants to maximize the welfare of its citizens, how does the community make these decisions? Consider first the relationship between the total cost per member and the number of members, *given* that a certain size park is selected. Clearly, the larger the community, the more people there are to shoulder the expense of the park, and the smaller the required contribution per member. But if the per capita cost continually decreases with membership size, why not simply invite an infinite number of people to join? The problem is that as more people join the community, the park becomes congested. The marginal congestion cost measures the dollar cost of the incremental congestion created by each new member. We assume that marginal congestion cost increases with the number of members. *The community should expand its membership until the marginal decrease in the membership fee just equals the per person marginal increase in congestion costs.*

Now turn to the flip side of the problem: For any given number of members in the community, how big should the park be? A bigger park yields greater benefits, although like most goods, we assume it is subject to diminishing marginal utility. The per member marginal cost of increased park acreage is just the price of the extra land divided by the number of members sharing its cost. *Acreage should be increased to the point where each member's marginal benefit just equals the per member marginal cost.*

We can now put together these two pieces of the picture to describe an optimal community or club. The optimal community is one in which the number of members and the level of services simultaneously satisfy the condition that the marginal cost equal the corresponding marginal benefit. Although this club model is very simple, it highlights the crucial aspects of the community-formation process. Specifically, it suggests how community size depends on the type of public goods the people want to consume, the extent to which these goods are subject to crowding, and the costs of obtaining them, among other things. However, viewing communities as clubs leaves unanswered several important questions that are relevant for understanding local public finance:

1. How are the public services to be financed? A country club can charge a membership fee, but a community normally levies taxes to pay for public goods.
2. A club can exclude nonmembers and so eliminate the free rider problem. How can communities achieve this end?

3. When people throughout the country organize themselves into many different clubs (communities), will the overall allocation of public goods prove to be equitable and efficient?

These questions are taken up in the next section.

The Tiebout Model

"Love it or leave it." When people who oppose US federal government policy are given this advice, it is generally about as constructive as telling them to "drop dead." Only in extreme cases do we expect people to leave their country because of government policy.[5] Because of the large pecuniary and psychic costs of emigrating, a more realistic option is to stay home and try to change the policy. On the other hand, most citizens are not as strongly attached to their local communities. If you dislike the policies being followed in Skokie, Illinois, the easiest thing to do may be to move a few miles away to Evanston. This section discusses the relationship among intercommunity mobility, voluntary community formation, and the efficient provision of public goods.

Chapter 5 examined the idea that markets generally fail to provide public goods efficiently. The root of the problem is that the market does not force individuals to reveal their true preferences for public goods. Everyone has an incentive to be a free rider. The usual conclusion is that some kind of government intervention is required.

In an important article, Tiebout [1956] (rhymes with "me too") argued that the ability of individuals to move among jurisdictions produces a market-like solution to the local public goods problem. As suggested in the cartoon, individuals vote with their feet and locate in the community that offers the bundle of public services and taxes they like best. Much as Jones satisfies her demand for private goods by purchasing them on the market, she satisfies her demand for public services by the appropriate selection of a community in which to live, and pays taxes for the services. In equilibrium, people distribute themselves across communities on the basis of their demands for public services. Each individual receives her desired level of public services and cannot be made better off by moving (or else she would). Hence, the equilibrium is Pareto efficient, and government action is not required to achieve efficiency.

Tiebout's Assumptions

Tiebout's provocative assertion that a quasi-market process can solve the public goods problem has stimulated a large amount of research. Much of that research has been directed toward finding a precise set of sufficient conditions under which the ability of citizens to vote with their feet leads to efficient public goods provision. Some of the conditions are as follows:[6]

[5] For example, in the 1960s, a number of young men left the country to evade military service in Vietnam.

[6] Mieszkowski and Zodrow [1989] provide more detail. Not all of these conditions were included in Tiebout's original article.

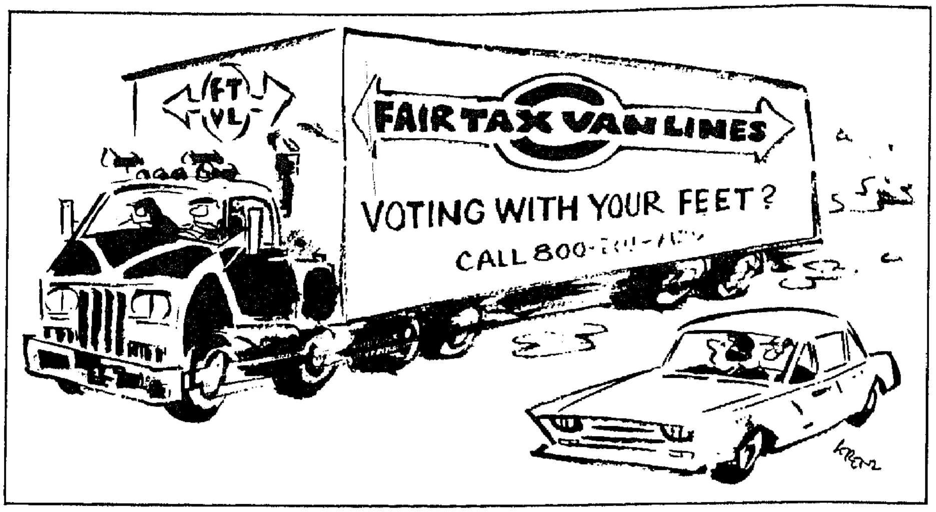

rawing by Lorenz; ©
'85 The New Yorker
agazine, Inc.

1. No externalities arise from local government behavior. As noted later, to the extent there are spillover effects among communities, the allocation of resources is inefficient.
2. Individuals are completely mobile. Each person can travel costlessly to a jurisdiction whose public services are best for him. The location of his place of employment puts no restriction on where he resides and does not affect his income.
3. People have perfect information with respect to each community's public services and taxes.
4. There are enough different communities so that each individual can find one with public services meeting her demands.
5. The cost per unit of public services is constant, so that if the quantity of public services doubles, the total cost also doubles. In addition, the technology of public service provision is such that if the number of residents doubles, the quantity of the public service provided must double. To see why these conditions are required for a Tiebout equilibrium to be efficient, imagine instead that the cost per unit of public services fell as the scale of provision increased. In that case, there would be scale economies of which independently operating communities might fail to take advantage.

   This assumption makes the public service essentially a publicly provided private good. "Pure" public goods (such as national defense) do not satisfy this assumption. However, many local public services such as education and garbage collection appear to fit this description to a reasonable extent.

6. Public services are financed by a proportional property tax. The tax rate can vary across communities.[7]
7. Communities can enact **exclusionary zoning laws**—statutes that prohibit certain uses of land. Specifically, they can require that all houses be of some minimum size. To see why this assumption is crucial, recall that in Tiebout equilibrium, communities are segregated on the basis of their members' demands for public goods. If income is positively correlated with the demand for public services, community segregation by income results. In high-income communities, the *level* of property values tends to be high, and, hence, the community can finance a given amount of public spending with a relatively low property tax *rate*. Low-income families have an incentive to move into rich communities and build relatively small houses. Because of the low tax rate, low-income families have relatively small tax liabilities, but nevertheless enjoy the high level of public service provision. As more low-income families get the idea and move in, the tax base per family in the community falls. Tax rates must be increased to finance the expanded level of public services required to serve the increased population.

   Since we assume perfect mobility, the rich have no reason to put up with this. They can just move to another community. But what stops the poor from following them? In the absence of constraints on mobility, nothing. Clearly, it is possible for a game of musical suburbs to develop in a Tiebout model. Exclusionary zoning prevents this phenomenon and thus maintains a stable Pareto efficient equilibrium.

Tiebout and the Real World

The Tiebout model is clearly not a perfect description of the real world. People are not perfectly mobile; there are probably not enough communities so that each family can find one with a bundle of services that suits it perfectly; and so on. Moreover, contrary to the model's implication, we observe many communities within which there are massive income differences and, hence, presumably different desired levels of public service provision. Just consider any major city.

However, we should not dismiss the Tiebout mechanism too hastily. There is a lot of mobility in the American economy. A persistent pattern is that in any given year, about 16 percent of Americans have different residences than they had the year before [US Bureau of the Census, 1993, p. 32]. Moreover, within most metropolitan areas, there is a wide range of choice with respect to type of community. As White [1975, p. 52] notes,

[7] Tiebout [1956] assumed finance by head taxes. The more realistic assumption of property taxation is from Hamilton [1975b].

"The salient fact about location choice in a large American metropolis is that households have a wide choice of places to live. Within a 20-mile radius they generally have a choice of one or more central cities and up to several hundred suburbs." Certainly, casual observation suggests that across suburbs there is considerable residential segregation by income, and that exclusionary zoning is practiced widely. In addition, it is not hard to find anecdotal evidence of classic Tiebout behavior:

> Police departments in California are in a bind: Crime is increasing, but after Proposition 13 took its toll of local property taxes, they are running out of money.[8] So some communities are turning to the 'police tax.'
>
> State law allows two-thirds of the voters of any community to override the limits of Proposition 13 and pass a supplemental assessment for police services. In affluent suburban communities, this works well. "We all gained huge savings as homeowners from Proposition 13," explains a member of the Atherton City Council. So paying up to $200 a year to help the police protect expensive homes from increasingly active burglars seems a good investment.
>
> But in larger, poorer cities, the idea did not fare well. Los Angeles and Oakland voters turned down police taxes. With lower property values, they have gained less from Proposition 13, and didn't consider a new tax so affordable. They also have less faith in police, and don't see how extra money would help.
>
> Despite the attention that it receives from politicians and the media, the fear of crime is evidently not absolute; it's just another part of life, factored in with all the others. ["Relative Crime," *New York Times*, June 22, 1981]

There have been several formal empirical tests of the Tiebout hypothesis. One type of study looks at whether the values of local public services and taxes are capitalized into local property values. The idea is that if people move in response to local packages of taxes and public services, differences in these packages should be reflected in property values.[9] A community with better public services should have higher property values, other things (including taxes) being the same. These capitalization studies are discussed later in this chapter in the context of property taxation. As noted there, capitalization does appear to be a widespread phenomenon.

Another interesting test was done by Gramlich and Rubinfeld [1982]. They analyzed responses to survey questions in which individuals were asked about their desired levels of local public expenditures. If the Tiebout mechanism is operative, we would expect to find substantial homogeneity of demands within suburbs located near many other commu-

---

[8] Proposition 13 is discussed later.

[9] There are some circumstances under which the Tiebout hypothesis does not imply that capitalization necessarily will occur [see Rubinfeld, 1987].

nities, because, in such a setting, the model suggests that those who are dissatisfied with current spending levels simply move elsewhere. On the other hand, in areas where there are few other communities nearby, it is harder to exit if you are unhappy. In such areas, people with very different demands for public goods may be lumped together in a single community. Gramlich and Rubinfeld found that compared to areas where there is little scope for choice, there are indeed relatively small differences in tastes for public goods within communities located in large metropolitan areas. This result suggests that, at least in some settings, the Tiebout model is a good depiction of reality.

Optimal Federalism

Now that we have an idea of how to characterize local governments, we return to our earlier question. What is the optimal allocation of economic responsibilities among levels of government in a federal system? The goal of the theory of optimal federalism is to determine the proper division of activities among the levels of government. Let us first briefly consider macroeconomic functions. Most economists agree that spending and taxing decisions intended to affect the levels of unemployment and inflation should be made by the central government. No state or local government is large enough to affect the overall level of economic activity. It would not make sense, for example, for each locality to issue its own money supply and pursue an independent monetary policy. Now, some macroeconomists argue that it may not be possible for even a central government to pursue effective policies to counter the business cycle. This issue is beyond the scope of this text. (See Mankiw [1990] for a discussion.) We merely note that to the extent a stabilization policy is feasible and desirable, it should be done at the national level.

With respect to the microeconomic activities of enhancing efficiency and equity, there is considerably more controversy. Posed within the framework of welfare economics, the question is whether a centralized or decentralized system is more likely to maximize social welfare. For simplicity, most of our discussion assumes just two levels of government, "central" and "local." No important insights are lost with this assumption.

Disadvantages of a Decentralized System

Consider a country composed of a group of small communities. Each community government makes decisions to maximize a social welfare function depending only on the utilities of its members—outsiders do not count.[10] How do the results compare to those that would emerge from

[10] We ignore for now the questions of how the social welfare function is determined and whether the people who run the government actually try to maximize it. (See Chapters 4 and 7.)