18

An Introduction to Benefit-Cost Analysis

ROBERT DORFMAN

Robert Dorfman is David A. Wells Professor of Political Economy Emeritus at Harvard University.

Chapter 5 presented the basic concepts of economic welfare theory and derived from them the conclusion that for practical purposes the best broad indicator of economic welfare is the augmented gross domestic product, AGDP. Accordingly, in reaching decisions about environmental protection policies, the effects of the alternative policies on AGDP are an important consideration.[1] Fortunately, for decision-making purposes it usually isn't necessary to estimate what the entire AGDP would be if the various alternatives (including the status quo) were adopted, but only to estimate the differential effects of the alternatives on the components of AGDP that are affected significantly.

The effect of any measure on AGDP is the resultant of its favorable effects, called benefits, and its unfavorable effects, called costs. Benefit-cost analysis (B-CA) is the task of compiling and evaluating these effects and calculating the overall change in AGDP that each alternative would produce. Though the idea is simple and commonsensical, complications, subtleties, and perplexities abound in the execution.

Most of the complications derive from the circumstance that governments deal largely in services and goods that business firms eschew because they cannot conveniently be sold, or anyway not at market-determined prices. No one can be charged for breathing air with a low concentration of sulfates, or for benefiting from many other environmen-

[1]Since any decision about environmental protection is a political decision, no simple criterion such as effects on AGDP can be completely determinative.

tal programs. As a result, no market prices for evaluating the contributions of most governmental programs to AGDP are available, which makes estimating the social values of government-provided goods and services a principal, and difficult, task of B-CA. We shall consider it in some detail below.

There are two principal broad approaches to B-CA. One, which we shall call the standard approach, is to enumerate all the ways in which a proposed environmental policy would impinge on AGDP, to estimate how great each of these effects would be and how much it would affect AGDP, and finally to aggregate all these benefits and costs into the total effect of the policy on economic welfare as measured by AGDP.

The other approach is called "contingent valuation," though "hypothetical valuation" would be a more accurate description. Its basis is simple: The effect of any project on AGDP depends on how highly people value its results and costs, and if you want to know how highly people value any project, just ask them. Accordingly, in a contingent valuation study, a properly randomized and stratified sample of the people affected is drawn, the project and its expected results are explained to each respondent, and each is asked how much she/he would be willing to pay to have the project instead of the status quo or some other basis of comparison. (Negative quotations are permitted; they indicate that the respondent prefers the status quo.) The total willingness to pay expressed by the sample, inflated to represent the whole population affected, is then an estimate of the social value of the proposed project. This approach avoids the nasty problems of assigning dollar values to nonmarket consequences, but instead asks a sample of the affected public to do the assigning. We shall have to consider below how far the general public can be relied on to perform such evaluations.

The next section will present a brief outline of the standard approach to benefit-cost analysis. Since it is not always possible to find acceptable monetary equivalents for the social values of the benefits and costs of environmental projects, complete benefit-cost evaluations are not always feasible. Accordingly, the following two sections are devoted to alternatives to B-CA that are often used when a B-CA cannot be completed.

The fourth section is an extensive discussion of nonmarket benefits—the ones for which monetary values are likely to be elusive—and includes three examples of methods frequently used to find monetary values in the absence of explicit price quotations. After that, the alternative approach to benefit-cost evaluation, the contingent valuation approach, will be considered. Finally, there is a brief discussion of two complications that tend to be neglected in B-CA, followed by an even briefer concluding section.

BENEFIT-COST ANALYSIS, THE STANDARD APPROACH

The standard approach to B-CA amounts to constructing a model of the undertaking to be evaluated, and tracing through the effects of the measure on the pertinent components of the AGDP, translating those effects into monetary terms, and aggregating the resulting benefits and costs into an estimate of the net effect of the measure.

Figure 1 depicts a benefit-cost model for a proposed limitation on some harmful discharge. The chain of effects by which a governmental initiative (shown in the topmost box) affects public health, ecological integrity, and other conditions of concern is outlined in the top five boxes. In carrying out a standard B-CA, each of the reactions in the sequence (indicated by the arrows connecting the boxes) has to be estimated. These estimating tasks are usually technical, and draw on a wide range of specialties: public administration, several branches of engineering, meteorology, hydrology, public health, and biology with emphasis on several fields of ecology, to mention just a few. Questions on, or even beyond, the frontier of knowledge are likely to arise. The resultant estimates are subject to considerable ranges of uncertainty.

The sixth box, the social valuation of the changes achieved by the governmental initiative, is the most difficult and disputatious of all. Sometimes the social value of a particular kind of benefit is estimated by a public opinion survey, as in the contingent valuation approach. Because of skepticism about the reliability of people's reports about their "willingness-to-pay" for environmental improvements (as well as other things), most analysts prefer to use concrete, cash-on-the-barrelhead indications of how highly people value the changes. A great deal of ingenuity has been devoted to inferring how much people would be willing to pay for environmental goods from decisions in related markets, and a voluminous literature has accumulated. But we must finish outlining the tasks of B-CA before tackling the problem of valuing nonmarket benefits.

The final step in a B-CA is to take costs into account. As indicated in the figure, costs are incurred at several stages in the chain of responses. First, the government will have expenses for administering, monitoring, and enforcing the program. Then the firms and/or individuals whose behavior is constrained will have costs for complying with regulations, and reporting their discharges and the measures they have taken. Finally, the public at large is likely to incur some costs, both monetary,[2] and nonmonetary in the form of reductions in convenience and amenities.

These costs must be estimated, just as the various benefits were, and the

[2]Care must be taken at this point not to include both the costs the program imposes on business firms and any price increases that firms use to pass the cost along to users of the product.

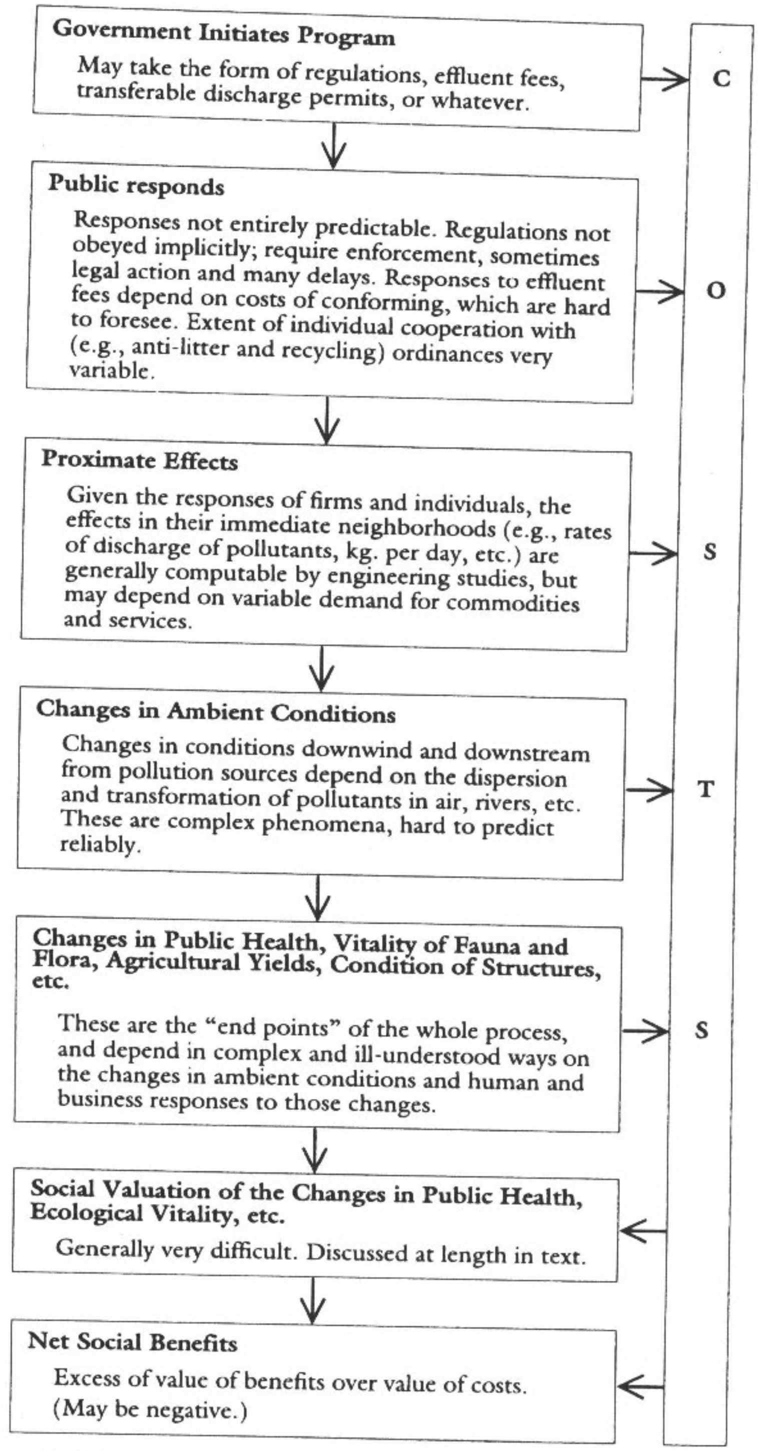

FIGURE 1. Chain of Effects of a Pollution Control Measure.

resultant total cost subtracted from the total value of the benefits to obtain the net social benefit.

In principle at least, the tasks just sketched have to be performed for each year during the expected duration of the program, since governmental programs, especially in the environmental field, generally extend over substantial periods of time, and it is not legitimate to assume that all years are the same. Usually, indeed, the costs will exceed the benefits during the early years when facilities are being constructed and installed, and the government, the firms affected, and the public are becoming acclimated to the program. Normally, it is only gradually, after a period of installation and running in, that benefits begin to exceed costs in individual years. Table 1 illustrates the kind of information that emerges for a typical year. It is, we hope, self-explanatory.

The next operation is to combine the net benefits for the individual years into an overall net benefit for the entire program. This phase raises

TABLE 1. Benefits and Costs of Hypothetical Atmospheric Antipollution Measure, Year 19xx

(A) TYPE	(B) UNIT	(C) QUANTITY	(D) UNIT VALUE ($)	(E) TOTAL VALUE (C) × (D) ($)
Benefits				
Reduced days of illness	person-days	500,000	40	$20,000,000
Reduced medical expense	$	1,250,000	1	1,250,000
Reduced crop losses	bushels	. . .	. . .	. . .
Increased days of high visibility	days	. . .	. . .	. . .
etc.				
Total Benefits				$ . . .
Costs				
Municipal:				
Capital investment	$	. . .	1	. . .
Operating & maintenance	$	. . .	1	. . .
Business:				
Capital investment	$	. . .	1	. . .
Operating & maintenance	$	. . .	1	. . .
Monitoring and enforcement	$	. . .	1	. . .
etc.				
Total costs				$ . . .
Net benefits, 19xx				$ . . .

one of the most vehemently debated aspects of the whole procedure: the question of time-discounting. The question is: Can the social worth of a program or project be obtained by simply summing its net values year by year, or must consequences that emerge in future years be reduced by discounting, and if so, at what rate?

This is a moral question, an empirical one, and an economic one, all simultaneously. The moral aspect is beyond the scope of this essay. Empirically, it is clear that people do discount future events and experiences, i.e., that there is "a perspective diminution of the future."[3] Finally, economic reasoning indicates that future events should be discounted, and that the reasonable rate of discount is the marginal productivity of capital. For, consider any desirable event, say averting a premature death. If it costs $\$D$ at present to achieve the event (by cost-effective means), then achieving X repetitions at present would cost about $\$DX$ if X is not unduly large. Now, if r is the marginal productivity of capital, $\$DX$ could be invested to yield $\$(1+r)DX$ next year, which would be enough to achieve $(1+r)X$ repetitions of the desired event. So, if society now spends some resources on achieving the result this year and simultaneously saves some to be used next year, as is often the case, it will follow that one occurrence this year is deemed neither more nor less important socially than $(1+r)$ occurrences next year; i.e., that the social discount rate for that event, whatever it is, is r per annum, the same as the marginal productivity of capital. Conclusion: The current social values of future events, even matters of life and death, fall at the rate of r per annum, where r is the marginal productivity of capital. This is the justification for discounting delayed benefits and costs in B-CA. We find this argument persuasive; not everybody does.[4]

Once the yearly net benefits have been calculated, finding the overall net benefit is straightforward. Suppose, to be specific, that the project or program is intended to be effective over a period of 50 years and that a discount rate of $100r$ percent per year is appropriate. Suppose also that the yearly net benefits have been estimated and that they are, successively, B_1, $B_2, \ldots, B_{50}$. Then the overall net benefits are given by the formula:

$$\text{Net Benefit} = \frac{B_1}{(1+r)} + \frac{B_2}{(1+r)^2} + \cdots + \frac{B_{50}}{(1+r)^{50}}.$$

In practical execution there are many variations—both shortcuts and

[3]The phrase is from Eugen von Böhm-Bawerk *Kapital und Kapitalzins, Vol. 2* (1889). English translation, *Capital and Interest,* Vol. 2, translated by G. D. Huncke and H. F. Sennholz (1959).

[4]Notice that this argument does not contend with the philosophic aspect of the question, and that it does not deal with the empirical or psychological aspect very profoundly since the existence of a positive marginal productivity of capital remains unexplained.

complications—of the procedure just described, but all benefit-cost evaluations fit this same general format.

ALTERNATIVES TO BENEFIT-COST ANALYSIS

Frequently it is not possible to estimate how much people would be willing to pay for some of the kinds of benefits that a program yields, or to avoid some kinds of costs, without resorting to unacceptable assumptions. In such cases, the B-CA analysis cannot be completed; the best that can be done is to estimate the net value of the benefits for which monetary equivalents can be found and to note the magnitudes of the remaining consequences in the most meaningful units available. We shall describe below two alternatives to B-CA that do not make such severe demands for data in monetary form.

To appreciate how demanding B-CA's data requirements are, consider the task of estimating the monetary value of the health benefits of a regulation intended to reduce emissions of sulphur oxide into the atmosphere. Among the data required are:

1. The amount by which power plants and other emitting sources will actually reduce their discharges in response to the regulation,
2. Given the amount of reduction at the sources, the amount by which the concentrations of sulphur oxide and its chemical products will be reduced at various places in the city,
3. Given the amount of reduction at different places in the city, the numbers of people who are exposed to various concentrations for various lengths of time,
4. Given the numbers of people exposed to different concentrations and durations, the reductions in the number of days of illness (perhaps distinguished by severity) and in medical expenses,
5. Given the reduction in days of illness, the social value to attach to it.

All this for one item in the table. Particularly difficult are steps 2 (because the process of the diffusion of pollutants in the atmosphere is complicated and not well understood), 4 (because the health effects of exposure to airborne and waterborne pollutants, technically called the "dosage-response curves," are known only very roughly), and 5 (because of the difficulty of attaching monetary values to nonmonetary consequences).

Such perplexities abound in the evaluation of every environmental protection measure or program. These difficulties, along with some methods for contending with them, are explored further in the last few selections in the volume. It is clear that precision is not attainable and that large ranges of uncertainty should be attached to all estimates of benefits, costs, and their net difference. It would be good practice for evaluators to indicate the ranges of uncertainty that they believe appropriate, but they rarely do so.

The most difficult estimate of all is the social evaluation of changes in health, environmental amenities, and the quality of life generally. Many benefit-cost evaluations do not even attempt it. In those evaluations, the nonmonetary consequences are simply omitted from the tables, and the totals of benefits and costs shown include the economically measurable consequences only. There is discretion in declining to attempt the nearly impossible, but there is also danger since effects omitted from the benefit-cost calculation tend to be given insufficient weight in making decisions based upon the calculation. Perhaps because of this danger, the usual practice is to attempt to make the estimates called for by column (D), however tenuous they may be, and to include all the anticipated effects of the proposal in the calculation.

On the other hand, for many practical purposes it is not necessary to assign dollar values to such consequences as reducing mortality, increasing the clarity of waters used for recreation, or preserving scenic areas or endangered species. Two methods for getting around the problem will be sketched briefly. They are most effective when there are only one or two types of benefit that defy evaluation in monetary terms, as is frequently the case.

The simpler method is to find the lowest unit value of the problematic type of benefit (or cost) that would lead to a total of benefits greater than the total of costs. In terms of Table 1, the only major category that resists dollar valuation is reduction in days of illness. If it should turn out that a valuation as low as $10 a day would be enough to make total benefits exceed total costs, almost any public body would find the proposed measure worthwhile. Further, if the net benefits of the proposed measure exceeded those of any alternative whenever illness-days were valued at $10 or more, this measure would be preferred to any of the alternatives. The effect of this device is to avoid estimating the AGNP contribution of the measure, but, instead, to find a lower bound for it.

The second method to be discussed is to construct a "trade-off diagram" or table. One is illustrated in Figure 2. In this figure four alternative proposals for reducing atmospheric pollution in an area are compared. It doesn't matter what the proposals are, so we shall not specify them. The monetary benefits, those measurable in dollar terms, are plotted horizontally. Since they may be positive but more usually are negative (because atmospheric improvements cost money—i.e., resources or economic output), the zero point for economic benefits is in the middle of the diagram. The nonmonetary benefits, in this case days of illness avoided, are plotted vertically. Proposal O is the status quo; it neither costs anything nor reduces illness, but serves as a basis for comparison. Proposal I is clearly

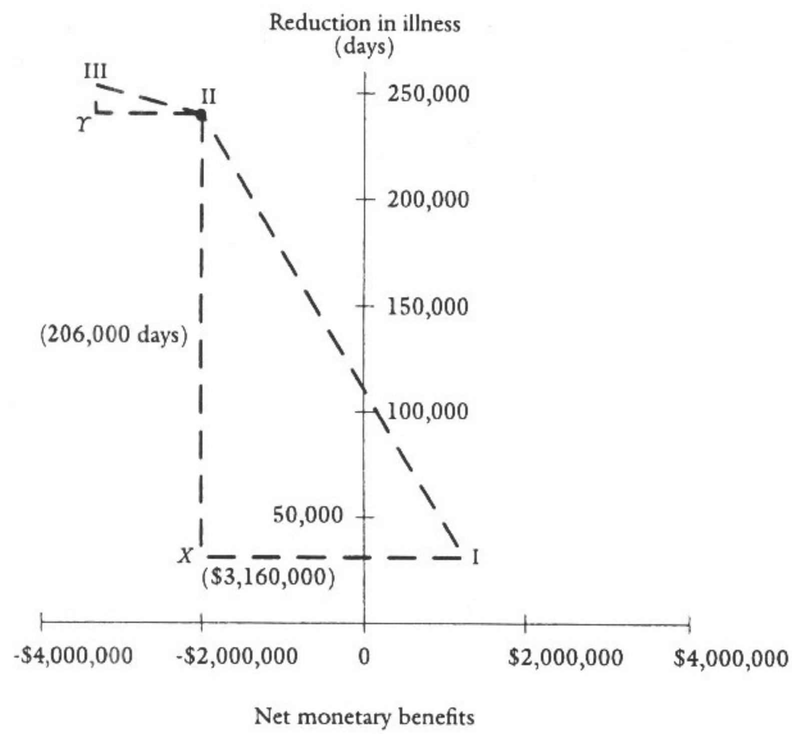

FIGURE 2. Trade-off Diagram between Monetary Costs and Reduction in Illness.

superior to Proposal O. It achieves some reduction in illness-days and, as a bonus, actually yields economic benefits in excess of its costs.

The other two proposals both have negative net economic benefits, that is, they entail economic sacrifices. But, to compensate, they lead to significant reductions in days of illness.

Such a diagram displays the major consequences of available alternatives without evaluating them in better-or-worse terms. Instead, it makes vivid the terms on which one kind of benefit or cost can be traded for the other; hence its name. The dashed lines in the diagram illustrate the comparisons among Proposals I, II, and III. In comparing Proposals I and II, the length of the line *II-X* (206,000 days) indicates the additional days of illness saved by adopting Proposal II rather than Proposal I, and the length of *I-X* ($3,160,000) represents the extra cost of Proposal II. Therefore the ratio of lengths, *I-X/II-X*, represents the cost per day of illness saved ($15.30) when the reduction is achieved by adopting Proposal II instead of Proposal I. The diagram doesn't tell whether reducing illness by this means is a bargain or exorbitant. That depends on the value that the community places on a day's reduction in illness, which is precisely the number that is virtually impossible to ascertain. Similarly, the diagram shows that Proposal III saves 12,500 more days of illness than Proposal II, at an additional cost of $1,333,000. Favoring Proposal III over Proposal II implies willingness to spend about $107 to avoid one day of illness in the community. Again the diagram says nothing about whether the additional cost is worthwhile.

Ultimately the choice among such proposals is a political decision dependent on social values, rather than one that is essentially dictated by a compelling measure of efficiency. A trade-off diagram can help make the decision an informed one.

These two devices (and there are others) show that even when definite prices cannot be assigned to nonmarket consequences, rational analysis and comparison of environmental programs is possible. A third device, called "cost-effectiveness analysis" (C-EA), is probably the most effective way to avoid resorting to unpersuasive estimates of the monetary value of nonmarket benefits.

COST-EFFECTIVENESS ANALYSIS

A reasonable question to ask of any proposed project is: Is it the cheapest way to do the job? The task of cost-effectiveness analysis is to answer that question. More formally stated, a project or a program is said to be "cost-effective" if it attains some specified physical or social goals and has net monetary benefits at least as great as those of any other project that attains them. The application of this test is cost-effectiveness analysis (C-EA). In terms of Figure 2, Project X is not cost-effective. It reduces illness the same amount as Project I but costs $3,160,000 more.[5] In general, a cost-effective project is necessarily on the production possibility frontier; if not, there would be some project that provided the same nonmonetary benefits and also either provided more benefits measurable in terms of money or cost less.

In operation, C-EA is much like B-CA. A reasonable range of alternatives that attain the designated goals have to be inspected, and for each the monetary values of benefits and costs have to be estimated, year by year, for all kinds of benefits and costs except the ones for which minimum levels have been prescribed. The advantage of C-EA is that it avoids the baffling need to find monetary equivalents for every type of benefit. But with that advantage comes a serious drawback, namely, the need to specify target levels for the critical kinds of benefit in advance. Usually there is not much justification for the target levels chosen, and, strictly speaking, the evaluation is valid for only those targets.

[5]That is just another way of saying that its nonmonetary benefits are that much lower.

In effect, one baffling task, specifying target levels, has been substituted for another, estimating monetary values of marginal changes in the levels of the benefits. It is often easier to specify reasonable target levels (e.g., the minimum level of dissolved oxygen in a stream) than to estimate the corresponding monetary value of the estimate (e.g., the social value of a marginal increase in dissolved oxygen). On the other hand, a C-EA generally provides less useful information than a B-CA unless there is some strong justification for the choice of the target level.

This limitation of C-EA can usually be relaxed by estimating the marginal social cost of changes in the prescribed targets. For example, in terms of Table 1, one might estimate the net monetary social cost of abating polluting emissions enough to reduce illness-days by 500,000 per year and also enough to reduce illness-days by 550,000. If the cost turns out to be very low, say $10 per illness-day, the more stringent abatement level would probably be judged worthwhile. Indeed, it might seem advisable to test the cost of still more stringent abatement goals until an abatement level is found at which a further reduction in illness-days is no longer an obvious bargain. In this way, practical judgments can be reached about abatement levels without expressing the social values of the results in terms of money. Of course, groping in this way to find a socially desirable critical level may prove very expensive.

VALUING NONMARKET BENEFITS

We have already noted that the most perplexing obstacle to the performance of B-CA is the need to find monetary values for nonmarket benefits, and that C-EA ameliorates this difficulty, though only partially. Most of the papers that follow in this section discuss different methods for valuing nonmarket benefits. In Chapter 20, Knetsch and Davis compare two methods, the "travel cost method" and contingent valuation. (Both will be discussed below.) Randall, Ives, and Eastman's paper on "bidding games," Chapter 21, is a classic and sober exposition of the contingent valuation approach. Chapter 22, by Landefeld and Seskin, summarizes the results of a large and varied set of efforts to place a value on measures that avert premature deaths; and Schelling's Chapter 23 discusses some of the philosophic issues that such efforts raise. In Chapter 25, Krupnick and Portney illustrate how the resulting estimates of values of nonmarket benefits are applied to appraising and evaluating a major, and very ambitious, environmental protection program. In Chapter 19, Arrow treats the specialized but important question of choosing the social rate of discount to be used in evaluating deferred benefits and costs. Finally, Wilson and Crouch (Chapter 24) discuss the problem of allowing for risk and uncertainty in benefit-cost evaluation, as well as elsewhere.

The ingenuity and effort devoted to establishing monetary equivalents to the social values of nonmarket benefits and costs are impressive. We now face the ungrateful task of appraising the results of that effort, of asking how valid and useful the estimates derived by the various approaches are. As a preliminary, we have to be clear about what the estimates are intended to represent and how they are supposed to be used.

The best starting point is probably the notion of consumers' sovereignty, which serves so well as a basis for estimating the social value of private goods. The skeleton of the argument on which consumers' sovereignty rests is: (1) the social value of any good lies in its contribution to the welfares of individual citizens, (2) every citizen (aside from children and a few other exceptions) is the authoritative judge of how much any good contributes to her welfare, (3) for private goods this judgment is conveyed by the amount the citizen is willing to pay for a marginal unit of the good or its services, (4) firms in competitive markets are responsive to the amounts citizens are willing to pay for private goods, and (5) the firms' responses result in an efficient allocation of production, i.e., one such that it is impossible to satisfy any citizen more without reducing some other citizen's level of welfare.

Now, if this doctrine does so well for private goods, shouldn't it be applied to public goods and the services of common property resources also? The major obstacle seems to be to discover the data conveyed in step 3, namely, how much citizens are willing to pay for a marginal unit of each public good or common property service. This is precisely the datum needed to fill the gap in a B-CA. We must therefore ask of each method for obtaining monetary expressions for the social value of public goods how well its results are likely to represent a public goods analog to the competitive market price of a private good.

REVEALED PREFERENCE METHODS

The revealed preference methods are a collection of more or less ad hoc devices connected only by the fact that they all infer consumers' valuations of environmental improvements from their choices in markets affected by them. We shall illustrate this diverse collection by discussing three frequently used methods.

"VALUE OF A LIFE." Since time immemorial, public health, as reflected in death rates and illness rates, has been a major concern of governments in general, and of environmental protection in particular. All such programs are expensive and, since budgets are always limited, judgments have to be made about which programs justify their draft on the budget. For this, as well as for other purposes, economists have searched assiduously for indications of the amounts that people are willing to pay to reduce the frequency of premature deaths and illnesses. The estimates cited by Landefeld and Seskin in Chapter 22 attest to how wide and thorough this search has been.

The "labor market" method for estimating this willingness-to-pay is very likely the prevalent one. It rests on Adam Smith's assertion that in equilibrium the net balance of pecuniary and nonpecuniary advantages and disadvantages must be the same in all occupations. On this ground, then, other things being equal, the more dangerous an occupation, the greater should be its wage. More exactly, we should be able to estimate how highly workers esteem their personal safety by analyzing how much they have to be paid to be willing to engage in occupations with different degrees of hazard. The principal information required for making such estimates are data on wage rates in several occupations with differing levels of occupational hazard, and the probabilities of fatal accidents (or illnesses) in those industries. As can be seen in Landefeld and Seskin's Table 2 (p 383), the results of those studies vary widely.[6] A more recent, and in some respects more sophisticated, estimate by Moore and Viscusi (1988) found that workers were paid an average of $3,400 per year (in terms of 1977 dollars) for an increase of 1/1,000 in the probability of suffering a fatal accident during a year.[7]

This very ingenious method for inferring the monetary value that people place on risks of death shares with the other methods to be discussed below a heavy reliance on some strong, and perhaps inapplicable, assumptions of economic theory. It assumes that the labor market is a perfect market, in which the workers are well informed about the risks and rewards of different occupations and in which any qualified worker can choose without impediment the job that suits her best. It assumes that there are no frictions, that when the wages or risks of an occupation change workers can quickly and at small expense enter or leave it. And it presumes that they have accurate psychological perceptions of very small risks, of the order of 1 in 10,000.

There is also a technical statistical problem, called the "specification problem," in this and the other methods to be discussed. Clearly riskiness is not the only characteristic that accounts for differences in wage rates

[6]Each study has its own technical peculiarities. Some relate to the risk of fatal injury only, others relate to all occupational accidents; only Thaler and Rosen include deaths from occupational diseases, and that inaccurately. In some studies, the workers covered are classified by industry, in others by occupation. The different studies also are based on samples from somewhat different working populations. The investigators should not be blamed for these discrepancies; they all had to rely on data gathered by government agencies and insurance companies for administrative rather than research purposes.

[7]This is equivalent to a "value of life" of $3.4 million. It should be noted that all these "value of life" estimates are really based on estimates of compensation paid for small increases in small probabilities of death.

among occupations. The effects of other factors, such as skill level, regularity of employment, pleasantness or unpleasantness of the tasks, section of the country, and many more have to be allowed for by statistical manipulation in order to isolate the effect of riskiness. The different studies accomplish this differently, depending on the data that were available and the judgment of the investigator.

The problem of specifying the relationships assumed in statistical analyses afflicts all empirical research in the social sciences. The best way to handle it is to try a number of plausible specifications in search of the one that fits the data best and to test whether the conclusions of the study are "robust," i.e., essentially the same under all plausible specifications. In each case the author has to choose the most satisfactory assumptions that her data permit, and the reader has to rely on her own judgment to decide whether to be satisfied with the specification chosen.[8]

VALUING IMPROVED URBAN ENVIRONMENT. A frequent problem in evaluating environmental programs is estimating how much the public is willing to pay for the improvement in environmental conditions that such programs achieve. One strategy is to analyze the results of the improvement into components such as reduced mortality rates, increased useful lives of structures exposed to the weather, reduced frequency of smog episodes, etc., and to estimate the public's willingness to pay for each of these. An alternative strategy, the one we shall explain here, is to cut through all those details and try to estimate directly the value that the public attaches to the improved environmental conditions.

The best indicator we have of the value the public places on environmental conditions is the observable relationship between them and property values. Table 2, which is taken from an important paper on methods for valuing nonmarket benefits,[9] gives some idea of what is involved in detecting and measuring the effect of environmental conditions on home values and rentals. There are two columns, each of which presents the coefficients of fourteen variables in a linear equation that describes the effects of those variables on the selling price of homes.[10] Thirteen of the variables are "nuisance variables," whose effects on the value of the house have to be eliminated in order for the effect of the environmental variables—the atmospheric concentration of nitrogen dioxide on the left and of suspended particulates on the right—to stand out clearly. In somewhat less elliptical notation, the equation for nitrogen dioxide concentration says:

[8]If you've ever wondered why so much of social science is controversial, here is a large part of the answer.

[9]See D. S. Brookshire, M. A. Thayer, W. D. Schulze, and R. C. d'Arge (1982).

[10]These equations are "least squares fits" to data recorded for the sales of 634 single-family dwellings in the Los Angeles metropolitan area in 1977–78.

TABLE 2. Estimated Hedonic Price Equations. Dependent Variable = log (Home Sale Price).

INDEPENDENT VARIABLE	NO_2 EQUATION	TSP EQUATION
Housing Structure Variables		
Sale date	.018591	.018654
Age	−.018171	−.021411
Living area	.00017568	.00017507
Bathrooms	.15602	.15703
Pool	.058063	.058397
Fireplaces	.099577	.099927
Neighborhood Variables		
Log (Crime)	−.08381	−.10401
School quality	.0019826	.001771
Ethnic composition	.027031	.043472
Housing density	−.000066926	−.000067613
Public safety expenditures	.00026192	.00026143
Accessibility Variables		
Distance to beach	−.011586	−.011612
Distance to employment	−.28514	−.26232
Air Pollution Variables		
log (TSP)		−.22183
log (NO_2)	−.22407	
Constant	5.4566	4.0527
R_2	.89	.89
Degrees of Freedom	619	619

Source: Adapted from Brookshire, Thayer, Schulze, and d'Arge (1982), with the kind permission of the authors.

Log(selling price) = 5.4566 + .01859*(Sale date, months since Dec. 1976) − .01817*(Age in yrs. at date of sale) + . . . − .22407*log (average annual concentration of NO_2 at nearest air-monitoring station, parts per billion).

The only number of interest in all of this is the coefficient −.22407, which asserts that on the average the value of a house fell by ⅔ percent for every 1 percent increase in the concentration of NO_2 in its neighborhood. Nevertheless, the data on all fifteen variables had to be compiled, and the coefficients estimated.

Before moving on to interpretation, it should be noted that the two equations shown are virtually the same, coefficient by coefficient, with only one or two exceptions. This is no coincidence; it indicates that the two pollution variables, NO_2 concentration and TSP concentration, are indistinguishable statistically, a frequent condition technically called "colinearity." It occurs whenever two or more variables tend to vary together, as the two pollution variables do in this instance; the NO_2 concentration tends to be high in areas where the TSP concentration is high, and low where it is low. Thus the two equations measure the same thing: the effect of changes in air pollution in general on housing prices.

But, of course, these studies are not undertaken out of interest in the determination of housing prices, but rather as a step toward evaluating people's willingness to pay for decreases in atmospheric pollution. At first glance, it might appear that all that is necessary after determining how much an average family is willing to pay for a house with improved atmosphere is to multiply by the number of houses. Not so. In the first place, improving a city's atmosphere affects its shopping districts, workplaces, recreation areas, etc., in addition to its residential areas. Cleaning up the air in the central business district may increase housing values throughout the city, but it will not be reflected adequately in the values of houses exposed to different levels of pollutant concentrations. The sum of the induced increases in housing values therefore omits these components of the social value of a reduction in urban air pollution.

But there is a more subtle effect, which is likely to work in the other direction. Consider a city which, like most, has its more and its less polluted neighborhoods, and suppose that the houses in the less polluted neighborhoods are more expensive for that reason and perhaps others. If an environmental regulation were to reduce pollution chiefly in the most polluted residential areas, house prices would rise there, but house prices in the less polluted areas might fall because the price differential would be likely to fall if the pollution differential did. The total value of real estate in the city might remain the same, or even decline. What, then, could be said about the social value of the atmospheric improvement?[11]

The step from the induced change in house values to the monetary measure of the change in social welfare is thus difficult. But the formula for the induced change in house values is itself suspect. Just as with the formulas for reductions in the probability of fatal accidents, it depends on stringent and dubious economic assumptions. The house buyers must be well informed about the degree of atmospheric pollution near the houses they consider and its effects on health and maintenance costs. The costs and other obstacles in the way of changing houses when prices or environmental conditions change must be moderate. All householders must have the same utility or preference function. The equation describing consumers' preferences must be correctly specified and the variables in it must be measured accurately. There must be a single housing market, so that all prospective buyers can choose among all available houses, and it must

[11]For the answer, see Strotz (1968) or Lind (1973).

behave at least approximately like a competitive market. And so on. These strict requirements led K-G. Mäler (1977, p. 368) to conclude: "Together these difficulties show conclusively that there is no real possibility of estimating willingness to pay for environmental quality from property value studies."

THE TRAVEL COST METHOD. The preceding two examples both depended on using data about a market assumed to be in equilibrium to infer consumers' preferences concerning some nonmarket goods. The travel cost method is based on an entirely different principle. Its basic idea was discovered in 1844 by E. J. Dupuit, a French engineer who used it to estimate the economical amount to spend on proposed bridges. Dupuit reasoned that the maximum economical expenditure for any bridge is just the greatest amount that users of the bridge would be willing to pay to have it in place, and that this willingness-to-pay is equal graphically to the area under the demand curve for the use of the bridge up to the abscissa for the amount of traffic expected. Tersely stated, this area, usually called the "consumers' surplus," is the amount that consumers would be willing to pay for the use of the bridge (or other facility) in excess of the amount they are required to pay.[12] The social value of a public service or facility of any sort is just the consumers' surplus that it generates. To estimate it, all that is needed is to estimate the demand curve and compute the area under it.

Unfortunately there are no markets for nonmarket goods, so estimating the demand curve for one is not altogether straightforward or even possible in most cases. But it is possible if two conditions are satisfied: (1) though there is no explicit market, some inconvenience or expense is needed in order to acquire the good; and (2) the amount of inconvenience or expense is observable and different for different people. Use of national and state parks and other scenic sites and recreational areas satisfies these conditions, and the travel cost method is used frequently to estimate the social values of such facilities.[13]

The first step in estimating the social value of a scenic or recreation site by the travel cost method is to derive the demand curve for visits from data on the numbers of visits from several points of origin at different distances from the site. Though this task can be complicated in practice, it can be explained adequately in terms of the simple situation where several communities use a single, isolated recreation site.

Two special assumptions are needed. The first is the usual one that consumers behave reasonably rationally. By this assumption, each consumer will use the recreation site to the extent where an additional use

[12]Consumers' surplus is discussed in virtually all texts on price theory. The explanations in Dorfman (1978), pp. 134–36, and Mansfield (1988), pp. 99–101, are probably as good as any.
[13]For a practical example, see Knetsch and Davis's Chapter 20.

would cost her more than the experience would be worth. The second assumption is more special. It is that the residents of the cities and towns that use the site have similar enough tastes so that their responses to differences in the cost of using the site can be interpreted as responses along a common demand curve.

The key data for estimating the demand curve for the use of the site by use of these assumptions are the proportion of the population of each community served by the site that actually visits it. These are found from a sample survey of visitors to the site, in which the visitors are asked their point of origin, the frequency with which they visit the site, and any desired auxiliary information. At this point, a little notation will be helpful. Let V be the total number of visitors to the site in a season, v the number included in the sample, v_i the number of visitors from community i in the sample, and P_i the population of community i. Then $(V/v)v_i$ is an estimate of the total number of visitors from community i, and $(V/v)v_i/P_i$ is the proportion of the population of the community that uses the site. U_i will denote this proportion.

The rationality assumption implies that the people in community i who visit the site will be those, and only those, for whom the value of a visit is as great as the cost. It is thus reasonable to assume that the value of a visit to the least eager visitor from any community, i.e., marginal value of visits from that community, is approximately equal to the cost of those visits.

The cost of reaching and using the site from each community that uses it can be estimated by multiplying the number of hours of travel needed (round trip) by a unit-time cost equivalent to the opportunity cost per time-unit, and adding any entrance fee, other expenses at the site, plus mileage costs and out-of-pocket travel expenses. Call the cost TC_i. Then, for each community that uses the site, the coordinate pair (U_i, TC_i) shows the relation between the cost of using the site to people in that community and the proportion of the population who do so. These coordinate pairs form the desired demand curve. One is illustrated in Figure 3. By Dupuit's argument, the shaded area represents the consumers' surplus per 1,000 population for a community from which the average cost of a visit to Yosemite is $40.

Clearly many essential complications were ignored in this exposition, including the effects of competing recreation sites; the size, recreation facilities, and other characteristics of the site studied; complementary attractions and pleasures at and en route to the site; and the size of families, incomes, and other characteristics of the communities served. All these factors shift the demand curve. In principle, they can be allowed for by introducing nuisance variables, much as in the previous two examples, to the extent that the available data permit. In practice, estimating a demand curve from travel cost data is a delicate procedure, sensitive to specification error. Introducing complications like the foregoing would

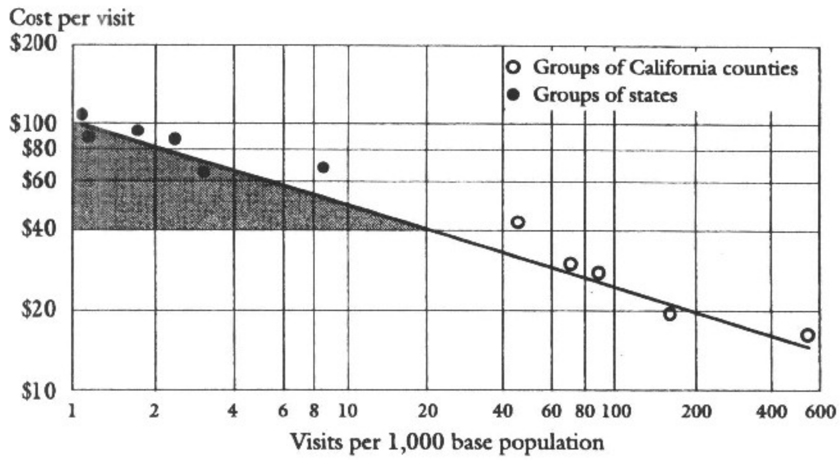

FIGURE 3. Estimated Demand Curve for Visits to Yosemite National Park, 1953. Notice the log-log scales. (Source: From Clawson and Knetsch (1966), p. 73.)

magnify the complexity of the calculations, vastly increase the requirements for data, and call for additional strong economic assumptions. In short, though the travel cost method is exceedingly ingenious, it is adapted to only very simple situations, and becomes awkward and unreliable when necessary complications are introduced.

CONTINGENT VALUATION METHODS

Since, as mentioned above, contingent valuation methods are controversial, an extensive literature is devoted to evaluating them.[14] The criticisms of contingent valuation all stem from the simple fact that a hypothetical question is not a real choice. In the literature, six types of bias that can afflict contingent valuation surveys have been distinguished and identified.

STRATEGIC BIAS occurs when a respondent chooses her answers in the hope of influencing the results of the survey. It is most likely to occur when respondents believe that the results of the survey will influence a pending political decision or have some other practical effect. Its presence can be detected by giving identical questionnaires to several samples, some of which are informed that a decision in which they are interested depended on the results. That kind of information had little or no effect on answers

[14]Probably the best general reviews are R. G. Cummings, D. S. Brookshire, and W. D. Schulze, et al. (1986), and R. C. Mitchell and R. T. Carson (1989).

to the questionnaire in several experimental tests. It appears that most people don't bother to falsify information in a survey for strategic reasons.

HYPOTHETICAL BIAS is the obverse of strategic bias. It occurs when respondents believe that their answers will not have any significant effects. It has several aspects, including: (1) people may be more generous when expressing willingness to make payments than in committing themselves actually to do so; (2) people may give offhand and ill-considered responses when they feel that nothing is at stake; and (3) especially when dealing with environmental questions, people may not visualize alternatives in a hypothetical situation as vividly as they do in a real one.

Experimental tests have revealed the reality of this source of bias. For example, when a sample of hunters was asked how much they would pay for a certain type of hunting license and a matched sample was allowed to bid in an auction of identical licenses, the hypothetical bids ranged from 30 percent to 65 percent higher than the actual ones. Similar experiments have disclosed the other types of hypothetical bias, but in some tests no biases were detected. Of course, hypothetical bias does not have to be detected on every occasion in order to be considered a dangerous pitfall when contingent valuation methods are used.

INFORMATION BIAS results when respondents are asked about unfamiliar options or contingencies. There is little point in asking ordinary people how much they would be willing to pay to have the concentration of sulphur oxides in their neighborhood reduced from 0.06 ppm to 0.03. On the other hand, it would strain the patience of most respondents if they were required to listen to an explanation of the meaning and effects of such a change. In consequence, the options that are of interest in a contingent valuation survey are likely to be vaguely delineated and even more vaguely understood.[15]

Quite apart from being ill-informed, respondents are confronted with an unfamiliar task when they are asked how much they would be willing to contribute to obtain or prevent a change in environmental conditions. Most people have never faced that choice in earnest, and it is hard for them to conjecture how they would respond if they should encounter it in earnest. Schelling, in Chapter 23, points out how difficult it is to induce a respondent to examine her own preferences for small changes in the risk of so serious an event as death.

VEHICLE BIAS refers to the effect of the mode of payment when respondents are asked "willingness-to-pay" types of questions. Respondents are likely to give different answers depending on whether they are told that the

[15]For examples, just think of the questions asked by Gallup or Roper or any other public opinion polling organization.

payment will be added to their tax bill, or deducted from their paychecks, or whatever. Thus the answers to such questions do not reflect the perceived value of the project in pure form, but rather an admixture of this value with the respondent's feelings about the mode of payment.

POPULATION SURVEY BIASES. In addition to those special problems of contingent valuation surveys, they are subject to the problems of public opinion surveys in general. For example: people like to make good impressions, even on interviewers.[16] Even when they are anonymous, they do not like to reveal themselves to be indifferent to nature, or stingy, or whatnot. People also prefer to be agreeable. Other things being almost equal, we'd rather say yes than no, and we are likely to at least shade our feelings if talking to a member of a minority or a female interviewer. The framing and form of a question is likely to affect the answer nearly as much as the substance. All in all, interpreting a human being's responses is inescapably tricky.

Finally, there is **NONRESPONSE BIAS.** In mail surveys, often fewer than half the questionnaires are returned with usable answers. There is always some nonresponse, and depending on its causes and extent, the sample of usable answers will depart more or less strongly from being truly representative of the population studied. There are correctives, such as following up on nonrespondents, but they are expensive and imperfect.

* * *

So much for weaknesses of contingent valuation surveys. They have two salient virtues. First, they are immune to most of the defects of the alternative approaches, which have been described. Second, they capture some aspects of social value that elude other methods, in particular the "non-use values." Non-use values, originally pointed out by Krutilla (Chapter 12) and Weisbrod (1964), are the values that people place on environmental or cultural amenities or conditions that they do not experience personally or expect to experience. If Niagara Falls were threatened with destruction, wouldn't you contribute at least a little to a fund that could save it for other people and future generations to see and marvel at?[17]

[16] Including the people who code the questionnaires in a mail survey.

[17] Before leaving contingent valuation, we should mention the continuing discussion of willingness-to-pay (WTP) vs. willingness-to-accept (WTA), since a contingent valuation survey has to choose between them in framing its questions. WTP is code for the most that the respondent would be willing to pay to obtain some environmental improvement. It corresponds to Hicks's concept of the compensating variation in income. WTA is the smallest recompense for which the respondent would be willing to forego an environmental improvement or consent to a degradation. According to consumption theory, WTP and WTA should be about equal when they are small in proportion to income; in contingent valuation surveys, WTA invariably turns out to be greater than WTP, and often four to six times as great.

How to explain? Kahneman and Tversky (1979) confirmed long ago that marginal utility curves have a discontinuity at zero: a marginal increase in wealth has less psychological

The upshot of these considerations appears to be that while contingent valuation has some enthusiastic advocates, the bulk of the profession is holding its judgment in abeyance and evaluating projects and programs by the standard approach.

TWO UNWELCOME COMPLICATIONS

We have now reviewed the essentials of benefit-cost analysis as it is practiced, but have to mention two complications that ought to be incorporated in practice. One is allowance for considerations of equity and political feasibility in the distribution of the benefits and costs, and the other is allowance for the uncertainty and inaccuracy of estimates of benefits and costs and their components.

The need to consider equity and political feasibility arises from the circumstance that the benefits and costs of environmental improvement often accrue unequally to different segments of the population, with most of the benefits going to some segments of the population, and other segments getting most of the costs. There is no need to expand on the resentments, tensions, and feelings of injustice that are likely to result.

In 1936, when benefit-cost analysis was first formalized in the United States, the instruction from the Congress was that projects should be undertaken only "if the benefits to whomsoever they may accrue are in excess of the estimated costs." This standard, though nominally adhered to this very day, was early seen to be grossly inadequate, and has been circumvented habitually from the very outset. In politics, it won't do to be too candid about matters of distribution. Thus the admonitions of the Water Resources Council, which officially sets the standards for federal benefit-cost analyses, are widely disregarded. But we repeat them here because they are sound and because they have to be obeyed, albeit surreptitiously.

In brief, the Water Resources Council recommends that the benefit and costs in a benefit-cost analysis be disaggregated to show their incidence to meaningful population segments (not necessarily congressional districts, though that might be popular). It appears to be expedient, as well as conformable to prevalent standards of justice, to design projects to be

impact than a decrease of the same amount. I've noticed that myself. My personal explanation is that people form ego attachments to their possessions and entitlements, and psychologically regard and resist decreases as invasions of their established rights. On this ground, the observed discrepancies between WTP and WTA are not errors introduced by contingent valuation, but reflections of discontinuities in people's preferences that public decisions should take into account.

Pareto efficient insofar as possible, i.e., so that all interested population groups receive benefits that exceed their shares of the costs. It is, however, rarely possible to meet this standard completely.

The role of a benefit-cost analysis is then to exhibit the balance of benefits and costs for each significant population segment. As everyone recognizes, these, and not the grand totals, are the significant data for political decision making, and the political process will proceed more smoothly and effectively to the extent that these facts of life are recognized and agreed on by all concerned. This is good advice, but honored mostly deviously, to great social cost.

Our review of benefit-cost analysis has emphasized that the estimates on both sides of the account are uncertain and inaccurate. When estimates are checked by being reestimated by independent methods or by being audited in the light of history, it is entirely usual for discrepancies of 200 percent or 300 percent or even an order of magnitude to be disclosed. Such is the nature of the beast. Therefore, (1) the third "significant" figure should never be taken seriously and the second should generally be regarded with some skepticism, and (2) the report on a benefit-cost analysis should indicate the ranges of uncertainty of the principal estimates. Chapter 25, Krupnick and Portney's benefit-cost analysis of urban air pollution control programs, illustrates conscientious reporting of the uncertainties inherent in benefit-cost assessments. Another example of good reporting is A. M. Freeman's report (1982) on the costs and benefits of air and water pollution controls, from which Table 3 is taken. Notice that the range of uncertainty is an order of magnitude in several instances. Nevertheless, Freeman was able to reach interesting and significant conclusions. The important thing in writing a benefit-cost report is to avoid seeming to tell the reader more than you can possibly know, and in reading one, to recognize the inherent imprecision even when the writer doesn't remind you.

CONCLUSION

You must be aware by now that this essay is not intended to either puff or belittle benefit-cost analysis. Its goal is to portray it as it is, warts and all. There are plenty of warts. There are also plenty of strengths, enough to make benefit-cost analysis a necessary tool for any government that tries seriously to make effective use of the resources it requisitions from its community.

Benefit-cost analysis has an especially important role in the decision-making processes of a democratic community. To a large extent, the preparation of a benefit-cost analysis performs the staff work for all sides

TABLE 3. Air Pollution Control Benefits Being Enjoyed in 1978
(in billions of 1978 dollars)

CATEGORY	REALIZED BENEFITS: RANGE	REALIZED BENEFITS: MOST REASONABLE POINT ESTIMATE
1. *Health*		
Stationary Source		
Mortality	$2.8–27.8	$13.9
Morbidity	$0.3–12.4	$ 3.1
Total	$3.1–40.2	$17.0
Mobile Source	$0.0– 0.4	$ 0.0
Total Health	$3.1–40.6	$17.0
2. *Soiling and Cleaning*	$1.0– 6.0	$ 3.0
3. *Vegetation*		
Stationary Source	0	0
Mobile Source	$0.1– 0.4	$ 0.3
Total Vegetation	$0.1– 0.4	$ 0.3
4. *Materials*		
Stationary Source	$0.4– 1.1	$ 0.7
Mobile Source	$0.0– 0.3	$ 0.0
Total Materials	$0.4– 1.4	$ 0.7
5. *Property Values*		
Stationary Source	$0.9– 6.9	$ 2.3
Mobile	$0.0– 2.0	$ 0.0
Total Property Value	$0.9– 8.9	$ 2.3

Source: A. M. Freeman III (1982), p. 128.

of the inevitable debate about a proposal for environmental protection or improvement. It gathers the essential data and sets forth reasonable economic, demographic, and technical assumptions, which often serve as ground rules that help make the ensuing debate coherent and intelligible. At the very least, the analysis will rule out some of the blatant misrepresentations that frequently mar political discourse. Though the analysis may be contested, the very acts of contesting and defending it focus the attention of all disputants on relevant and comprehensible issues. In short, the sources and foci of people's disagreements are exposed.

A spectacular example of this function of benefit-cost analysis concerns a proposal to build a canal across central Florida to provide a shortcut barge route from Texas and Louisiana to the Atlantic Coast. The plan envisaged dredging and "improving" almost 50 nearly pristine miles of the Oklawaha River, probably the largest stretch of wild river through a tropical rain forest in the United States. A protracted brouhaha naturally resulted. Finally, a benefit-cost analysis was undertaken. It was carefully disaggregated according to interest groups that would be affected, such as shippers who might use the canal, hunters, fishers, environmentalists who prized the wilderness, lumbermen, and so forth. The analysis showed that scarcely any significant group would benefit greatly from the canal, and most would be net losers. After that finding, the proposal was dropped without further debate, although construction had already begun.[18]

Counterbalancing such pleasant triumphs is the notorious fact that benefit-cost analyses can be slanted, and often are. The technique is treacherously simple. If in favor of a proposal, overvalue the benefits and underestimate the costs; if opposed, do the opposite. The discussions in this paper of how benefit-cost analyses are conducted and how benefits and costs are valued should be helpful both in distorting analyses and in detecting the biases. It is the unfortunate case that truly neutral analyses are hard to find, though the cross-Florida canal episode may have been one. When dealing with benefit-cost analyses, one should never forget the maxim, *Caveat lector*.

The limitations of benefit-cost analysis have been emphasized sufficiently by now to dispel hope that it can provide decisive guidance. There is one fundamental limitation, however, that probably ought to be made more explicit. The goals and considerations that enter into any real public decision are varied, and subtle, and often left discreetly unstated. So there are almost invariably some goals and some restrictions that cannot be fitted into the benefit-cost format, but cannot be ignored either. In the end, therefore, benefit-cost analysis can be an important ingredient of the decision process, but the final decisions elude the benefit-cost accountants.

REFERENCES

Böhm-Bawerk, Eugen von (1889), *Kapital und Kapitalzins*, Band 2. Translated as *Capital and Interest*, Vol. 2, by G. D. Huncke and H. F. Sennholz (1959).

Brookshire, D. S., M. Thayer, R. Schulze, and R. d'Arge (1982), "Valuation of public goods," *Amer. Econ. Rev.*, 72:165–77.

Carter, Luther J. (1975), *The Florida Experience: Land and Water Policy in a Growth State*. Baltimore: Johns Hopkins University Press for Resources for the Future.

Clawson, Marion, and Jack L. Knetsch (1966), *Economics of Outdoor Recreation*. Baltimore: Johns Hopkins University Press for Resources for the Future.

[18]Luther Carter (1975, Chap. 9) is a vivid narrative of the first dozen years of this controversy. Unfortunately, it was written before the denouement sketched in the text was reached.

Cummings, R. G., D. S. Brookshire, W. D. Schulze, et al. (1986), *Valuing Environmental Goods: An Assessment of the Contingent Valuation Method.* Totowa, NJ: Rowman & Allenheld.

Dorfman, Robert (1978), *Prices and Markets,* 3rd edn. Englewood Cliffs, NJ: Prentice-Hall.

Dupuit, E. J. (1844), "De la mésure de l'utilité des travaux publics," *Annales des ponts et chausées, 2^me^ ser.,* vol. VIII.

Freeman, A. Myrick III (1982), *Air and Water Pollution Control: A Benefit-Cost Assessment.* New York: John Wiley.

Kahneman, Daniel, and Amos Tversky (1979), "Prospect theory: An analysis of decision under risk," *Econometrica,* 47:263–91.

Lind, Robert C. (1973), "Spatial equilibrium, the theory of rents, and the measurement of benefits from public programs," *Quart. J. Econ.,* 87:188–207.

Mäler, Karl-Göran (1977), "A note on the use of property values in estimating marginal willingness to pay for environmental quality," *J. of Environmental Econ. and Mgt.,* 4:355–69.

Mansfield, Edwin (1988), *Microeconomics, Theory and Applications,* 6th edn. New York: W. W. Norton.

Mitchell, R. C., and R. T. Carson (1989), *Using Surveys to Value Public Goods.* Washington, DC: Resources for the Future.

Moore, M. J., and W. K. Viscusi (1988). "The quantity-adjusted value of life," *Economic Inquiry,* 26:369–88.

Strotz, Robert H. (1968), "The use of land rent changes to measure the welfare benefits of land improvement," in Joseph E. Haring, ed., *The New Economics of Regulated Industries.* Los Angeles: Occidental College.

Weisbrod, Burton A. (1964), "Collective consumption services of individual consumption goods," *Quart. J. Econ.,* 77:471–77.