

## FOURTEEN

### Why Is Cost-Benefit Analysis So Controversial?

Robert H. Frank\*

Robert H. Frank (1945– ) holds a B.S. in math from Georgia Institute of Technology and an M.A. in statistics and a Ph.D. in economics from University of California, Berkeley. Since serving as a Peace Corps volunteer in Nepal, he has been an economics professor at Cornell University. His “Economic Scene” column appears monthly in the *New York Times*. His work at the intersection of economics and ethics has focused on moral emotions and concerns about relative position.

#### Overview

The cost-benefit principle says we should take those actions, and only those actions, whose benefits exceed their costs. For many, this principle’s commonsensical ring makes it hard to imagine how anyone could disagree. Yet critics of cost-benefit analysis are both numerous and outspoken. Many of them argue that cost-benefit analysis is unacceptable as a matter of principle. I begin by noting why many find this argument largely unpersuasive. I then examine several conventions adopted by cost-benefit analysts that do appear to yield misleading prescriptions. Finally, I consider the possibility that the cost-benefit principle may itself suggest why we might not always want to employ cost-benefit analysis as the explicit rationale for our actions.

#### The Incommensurability Problem

The cost-benefit principle says we should install a guardrail on a dangerous stretch of mountain road if the dollar cost of doing so is less than the implicit

---

\* Henrietta Johnson Louis Professor of Management and Professor of Economics, Cornell University, Johnson Graduate School of Management. This paper was prepared for presentation at the University of Chicago Law School conference Cost-Benefit Analysis, September 17–18, 1999. I thank William Schulze for helpful discussions.

*Journal of Legal Studies*, vol. 29 (2000): 913–30. Copyright © 2000 by the University of Chicago. Reproduced by permission of the University of Chicago Press.

dollar value of the injuries, deaths, and property damage thus prevented. Many critics respond that placing a dollar value on human life and suffering is morally illegitimate.<sup>1</sup>

The apparent implication is that we should install the guardrail no matter how much it costs or no matter how little it affects the risk of death and injury.

Given that we live in a world of scarcity, however, this position is difficult to defend. After all, money spent on a guardrail could be used to purchase other things we value, including things that enhance health and safety in other domains. Since we have only so much to spend, why should we install a guardrail if the same money spent on, say, better weather forecasting would prevent even more deaths and injuries?

More generally, critics object to the cost-benefit framework's use of a monetary metric to place the pros and cons of an action on a common footing. They complain, for example, that when a power plant pollutes the air, our gains from the cheap power thus obtained simply cannot be compared with the pristine view of the Grand Canyon we sacrifice.

Even the most ardent proponents of cost-benefit analysis concede that comparing disparate categories is extremely difficult in practice. But many critics insist that such comparisons cannot be made even in principle. In their view, the problem is not that we do not know how big a reduction in energy costs would be required to compensate for a given reduction in air quality. Rather, it is that the two categories are simply incommensurable.

This view has troubling implications. In the eyes of the cost-benefit analyst, any action – even one whose costs and benefits are hard to compare – becomes irresistibly attractive if its benefits are sufficiently large and its costs are sufficiently small. Indeed, few people would oppose a new technology that would reduce the cost of power by half if its only negative effect were to degrade our view of the Grand Canyon for just one 15-second interval each decade.<sup>2</sup> By the same token, no one would favor adoption of a technology that produced only a negligible reduction in the cost of power at the expense of a dark cloud that continuously shielded North America from the rays of the Sun. We live in a continuous world. If the first technology is clearly acceptable, and the second clearly unacceptable, some intermediate technology is neither better nor worse than the status quo. And we should count any technology that is better than that one as an improvement.

Scarcity is a simple fact of the human condition. To have more of one good thing, we must settle for less of another. Claiming that different values are incommensurable simply hinders clear thinking about difficult tradeoffs.

Notwithstanding their public pronouncements about incommensurability, even the fiercest critics of cost-benefit analysis cannot escape such

tradeoffs. For example, they do not vacuum their houses several times a day, nor do they get their brakes checked every morning. The reason, presumably, is not that clean air and auto safety do not matter, but that they have more pressing uses of their time. Like the rest of us, they are forced to make the best accommodations they can between competing values.

### **General Reservations about Consequentialist Ethics**

Many critics of cost-benefit analysis fault it for being rooted in utilitarianism or some closely related form of consequentialist ethical theory.<sup>3</sup> Consequentialist theories hold that the right course of action is the one that leads to the best consequences, where “consequences” under the utilitarian variant means “highest total utility.” Critics often attack consequentialism by citing examples in which its purported conclusions clash with the reader’s ethical intuitions. One popular example invokes the “utility monster,” someone who transforms resources into utility far more efficiently than anyone else. Critics argue that since utilitarianism says the best outcome is to give all resources to the utility monster, and since we know this to be an absurd conclusion, we must reject the ethical theory upon which cost-benefit analysis rests.

Consequentialist moral philosophers have attempted to show that their theories, properly construed, do not imply the conclusions suggested by such examples.<sup>4</sup> But even if these disputes are never fully resolved, we may note that the theories favored by the rival camps reach remarkably similar decisions regarding a broad range of ethical questions. As a practical matter, then, the mere fact that cost-benefit analysis is closely identified with consequentialist ethical theories would not seem to imply that its prescriptions are systematically misleading.

### **Discounting the Future**

As traditionally implemented, cost-benefit analysis attempts to put all relevant costs and benefits on a common temporal footing. A discount rate is chosen, which is then used to compute all relevant future costs and benefits in present-value terms. Most commonly, the discount rate used for present-value calculations is an interest rate taken from financial markets.

Though some critics complain about this practice, use of a market interest rate to discount future monetary costs and benefits commands broad approval. After all, if the annual interest rate on financial deposits is 7 percent, one can cover a \$1,000 cost 10 years hence by depositing only \$500 today.

There is less widespread agreement about using a market interest rate to discount future subjective utility. As Stanley Jevons argued, for example, “To secure a maximum benefit in life, all future pleasures or pains, should act upon us with the same force as if they were present, allowance being made for their uncertainty . . . But no human mind is constituted in this perfect way: a future feeling is always less influential than a present one.”<sup>5</sup>

On this view, if failure to adopt more stringent air quality standards today means that respiratory illnesses will be more common a generation from now, those illnesses should receive roughly the same weight as if they were to occur today. Having been born later should not mean that one’s enjoyment and suffering receive less weight in important policy decisions. Of course, a complete cost-benefit calculation would also want to make allowance for possible improvements in medical technology that would make the consequences of a given illness less severe in the future.

Whatever the ultimate merits of this position, it does not argue against the use of cost-benefit analysis as a matter of principle. If analysts agree that future experiences should receive roughly the same weight as current ones, the costs and benefits associated with any policy change can simply be calculated on that basis.

### Distributional Issues

Distributional issues have long been a favorite target of critics of cost-benefit analysis. Their objection, in a nutshell, is that because willingness to pay is based on income, cost-benefit analysis assigns unjustifiably large decision weight to high-income persons. Implicit in this objection is the view that everyone’s preferences regarding policy decisions should receive the same weight, irrespective of income.

Critics presumably have the interests of the poor in mind when they press this objection. Yet it is not clear that the poor themselves would want policy decisions to be made on some basis other than willingness to pay. Consider, for example, a community consisting of three voters – one rich, the other two poor. Up for decision is a proposal to switch the local public radio station from an all-music format to an all-talk format. The rich voter would be willing to pay \$1,000 to see this change enacted, while the poor voters would be willing to pay \$100 each to prevent it. If each voter’s interests are weighted equally, the switch will not be adopted. Yet, in cost-benefit terms, failure to switch results in a net loss of \$800.

Under the circumstances, little ingenuity is required to design a proposal that would command unanimous support. The switch could be made

conditional, for example, on the rich voter making an additional \$500 contribution to the public treasury, which could then be used to reduce the taxes of each poor voter by \$250.

Critics may respond that although such transfers would be fine in principle, the poor lack the political muscle to assure they are carried out. In an imperfect world, they argue, we get better results by resolving such issues on a one-person, one-vote basis.

But this response simply will not do. If the poor lack the political power to bargain for compensation in return for supporting a policy that harms them, what gives them the power to block that policy in the first place? But if they have that power, they necessarily have the power to bargain for compensation. After all, any policy that passes the cost-benefit test but creates net losses for the poor can be transformed into a Pareto improvement by simply making the tax system more progressive.

Critics of cost-benefit analysis are correct that using unweighted willingness-to-pay measures virtually assures a mix of public programs that are slanted in favor of the preferences of high-income persons. But rather than abandon cost-benefit analysis, we have a better alternative. We can employ unweighted willingness-to-pay measures without apology, and use the welfare and tax system to compensate low-income families *ex ante* for the resulting injury. The compensation need not – indeed cannot – occur on a case-by-case basis. Rather, low-income persons could simply be granted the welfare and tax breaks required by distributive justice, plus additional concessions reflecting their expected loss from the implementation of cost-benefit analysis using unweighted willingness-to-pay measures.

My point in offering this defense of standard cost-benefit analysis is not that granting additional political power to the poor would be a bad idea. Rather, it is that abandoning cost-benefit analysis is a gratuitously wasteful way of trying to achieve that goal. Rich and poor alike have an interest in making the economic pie as large as possible. Any policy that passes the cost-benefit test makes the economic pie larger. And when the pie is larger, everyone can have a larger slice.

### Measurement Problems

To discover whether an action satisfies the cost-benefit test, we must come up with concrete measures of its costs and benefits. Notwithstanding the logical difficulties raised by claims of incommensurability, this much is clear: constructing plausible measures of the costs and benefits of specific actions is often very difficult. In practice, analysts try to estimate costs and benefits

either by using survey methods or by drawing inferences from market behavior. Both approaches, however, are fraught with difficulty.

### Survey Methods

How much is the preservation of a virgin redwood forest worth? Proponents of the contingent-valuation method generate estimates by asking people how much they would be willing to pay to see the forest preserved. Responses in such surveys are problematic for several reasons.

One difficulty is that the valuations are often implausibly large. For example, if the amount someone would pay to prevent a specific stretch of coastline from being fouled by an oil spill were applied to all coastlines worldwide, the resulting sum would typically far exceed his total wealth.<sup>6</sup> Responses in contingent-valuation surveys are also highly sensitive to how questions are phrased and to the format provided for responses.<sup>7</sup>

But perhaps the most troubling feature of contingent-valuation surveys is that respondents are often willing to pay more, by several orders of magnitude, to prevent a harmful effect than to undo a harmful effect that has already occurred. Richard Thaler coined the term “loss aversion” to describe this tendency.<sup>8</sup> Loss aversion means not just that the pain of losing a given amount is larger, for most of us, than the pleasure from gaining that same amount. It is much larger.

Thaler illustrates the asymmetry by asking students to consider the following hypothetical questions:

1. By attending class today, you have been exposed to a rare, fatal disease. The probability that you have the disease is one in a thousand. If you have the disease you will die a quick and painless death in one week. There is a cure for the disease that always works, but it has to be taken now. We do not know how much it will cost. You must say now the most you would be willing to pay for this cure. If the cure ends up costing more you won't get it. If it costs less, you will pay the stated price, not the maximum you stated. How much will you pay?
2. We are conducting experiments on the same disease for which we need subjects. A subject will just have to expose him or herself to the disease and risk a one-in-a-thousand chance of death. What is the minimum fee you would accept to become such a subject?<sup>9</sup>

In each scenario, respondents are asked, in effect, how much they value a one in 1,000 reduction in the probability of death. But whereas the first scenario asks how much they would pay to eliminate a risk of death to which

they have already been exposed, the second asks them how much they would have to be paid before exposing themselves to a similar risk voluntarily. The median responses were approximately \$800 for the first question and \$100,000 for the second.<sup>10</sup> Similar disparities between willingness to pay and willingness to accept are observed in contingent-valuation surveys that pose environmental questions.<sup>11</sup> Disparities in other domains are typically smaller, though few surveys find willingness-to-pay values that are more than half as large as the corresponding values for willingness to accept.<sup>12</sup> These disparities, needless to say, pose formidable hurdles for analysts who employ contingent-valuation methods.

### Hedonic Methods

These and other problems inherent in survey methods have led many analysts to favor hedonic pricing models, which attempt to infer valuations from observable market behavior. In typical applications, analysts estimate the value of noise reduction by examining how residential housing prices vary with ambient noise levels, or the value of safety by examining how wages vary with workplace injury levels.<sup>13</sup>

Hedonic pricing models assume that the wage-safety gradient tells us how much workers value safety. Is this a tenable assumption? The argument in support of it is a simple application of invisible-hand theory. If an amenity – say, a guardrail on a lathe – costs \$50 per month to install and maintain, and if workers value it at \$100 per month, then firms that do not install one risk losing valued employees to a competitor who does. After all, if a competitor were to pay a worker \$60 per month less than he earns from his current employer, it could cover the cost of the safety device with \$10 to spare, while providing an overall compensation package that is \$40 per month more attractive than his current employer's.

To this argument, critics respond that labor markets are not workably competitive in practice. Incomplete information, worker immobility, and other imperfections force workers to accept whatever conditions employers offer. But even if a firm were the only employer in a labor market, it would still have a clear incentive to install a \$50 safety device that is worth \$100 to the worker. Failure to do so would leave cash on the table.

Other critics suggest that workers often do not know about the safety devices they lack. But this claim is also troubling because firms would have strong incentives to call these devices to workers' attention. After all, both the firm and its workers come out ahead when a cost-effective safety device is adopted.

With respect to the charge that labor markets are not effectively competitive, critics of hedonic pricing models have failed to meet the burden of proof. Worker mobility between firms is high, as is entry by new firms into existing markets, and cartel agreements have always been notoriously unstable. Information is never perfect, but if a new employer in town is offering a better deal, word sooner or later gets around.

If, despite these checks, some firms still managed to exploit their workers by paying less than a competitive wage, we should expect these firms to have relatively high profits. In fact, however, we observe just the opposite correlation. Year in and year out, the firms paying the highest wages are most profitable.<sup>14</sup>

But even if labor markets are workably competitive, the same theory of revealed preference that makes hedonic models so attractive also sounds a cautionary note. It calls our attention to a related form of behavioral evidence, namely, the laws we choose to adopt. Scholars in the law and economics movement have long argued that laws tend to evolve in ways that maximize wealth.<sup>15</sup> This characterization presumably also applies to laws regulating health and safety in the workplace, which by now have been enacted by virtually all industrial democracies. These laws pose a challenge to the hedonic pricing model's assumption that safety risks are fully reflected in compensating wage differentials. If this assumption were correct, safety regulations would entail costs that exceed their benefits and therefore should not have been enacted in the first place. But although these regulations have often been criticized on practical grounds, they appear in no imminent political danger.

Does the political success of safety regulation suggest that hedonic pricing models are misleading? I believe it does, but not for the reasons usually given. In what follows I construct an example to illustrate an alternative rationale for safety regulation, one that is independent of market power and imperfect information.

### Positional Concerns and Revealed Preference

Consider a hypothetical community with only two members, Sherwin and Gary. Each gets satisfaction from three things – from his income, from his safety on the job, and from his position on the economic ladder. Each must choose between two jobs – a safe job that pays \$300 per week and a risky job that pays \$350 per week. The value of safety to each is \$100 per week, and each evaluates relative income as follows: Having more income than his neighbor provides the equivalent of \$100 per week worth of additional



		<b>Sherwin</b>	
		Safe job @ \$300/week	Unsafe job @ \$350/week
<b>Gary</b>	Safe job @ \$300/week	\$400/week each	\$300/week for Gary \$450/week for Sherwin
	Unsafe job @ \$350/week	\$450/week for Gary \$300/week for Sherwin	\$350/week each

Figure 14.1. The effect of concerns about relative income on worker choices regarding safety.

satisfaction; having less income than his neighbor means the equivalent of a \$100 per week reduction in satisfaction; and having the same income as his neighbor means no change in the underlying level of satisfaction. Will Sherwin and Gary choose optimally between the two jobs?

If we viewed each person’s decision in isolation, the uniquely correct choice would be the safe job. Although it pays \$50 per week less than the risky job, the extra safety it provides is worth \$100 per week. So if we abstract from the issue of concern about relative income, the value of the safe job is \$400 per week (its \$300 salary plus \$100 worth of safety), which is \$50 per week more than the \$350 value of the risky job.

Once we incorporate concerns about relative income, however, the decision logic changes in a fundamental way. Now the attractiveness of each choice depends on the job chosen by the other. The four possible combinations of choices and the corresponding levels of satisfaction are shown in Figure 14.1.

Suppose, for example, that Gary chooses the safe job. If Sherwin then chooses the unsafe job, he ends up with total satisfaction worth \$450 – \$350 in salary plus \$100 from having more income than Gary. Gary, for his part, ends up with only \$300 worth of total satisfaction – \$300 in salary plus \$100 from safety minus \$100 from having lower income than Sherwin. Alternatively, suppose Gary chooses the unsafe job. Then Sherwin again does better to accept the unsafe job, for by so doing he gets \$350 worth of satisfaction rather than only \$300. Since the payoff matrix is symmetric, each player’s dominant strategy is to choose the unsafe job. Analysts equipped with the hedonic pricing model will conclude that these workers must value the extra safety at less than \$50 per week.

But this inference is clearly wrong. Note that if each chooses a safe job, each will get \$400 worth of total satisfaction – \$300 of income, \$100 worth of satisfaction from safety, and zero satisfaction from relative position. If each had instead chosen the unsafe job, each would have had \$350 of income, zero satisfaction from safety, and each would again have had the same level of income, so again zero satisfaction from relative position. If we compare the upper-left cell of Figure 14.1 to the lower-right cell, then, we can say unequivocally that Sherwin and Gary would be happier if each took a safe job at lower income than if each chose an unsafe job with more income. By assumption, the extra safety is worth more than its cost.

The discrepancy arises because the job safety choice confronts workers with a Prisoner's Dilemma. If they could choose collectively, they would pick the safe job, an outcome they prefer to what happens when they choose independently. On this interpretation, safety regulation is attractive not because it prevents exploitation, but because it mitigates the consequences of consumption externalities.

Many modern disciples of Adam Smith appear reluctant to introduce concerns about relative position into normative economic models. Yet as Smith himself recognized, such concerns are a basic component of human nature:

Consumable commodities are either necessities or luxuries. By necessities I understand not only the commodities which are indispensably necessary for the support of life, but whatever the custom of the country renders it indecent for creditable people, even of the lowest order, to be without. A linen shirt, for example, is, strictly speaking, not a necessary of life. The Greeks and Romans lived, I suppose, very comfortably though they had no linen. But in the present times, through the greater part of Europe, a creditable day-labourer would be ashamed to appear in public without a linen shirt, the want of which would be supposed to denote that disgraceful degree of poverty which, it is presumed, nobody can well fall into without extreme bad conduct. Custom, in the same manner, has rendered leather shoes a necessary of life in England. The poorest creditable person of either sex would be ashamed to appear in public without them.<sup>16</sup>

As Smith clearly understood, concerns about relative income need not entail a desire to have more or better goods than one's neighbors. People with low relative income experience not just psychological discomfort but also more tangible economic costs.<sup>17</sup> A resident of a remote Indian mountain village has no need for a car, but a resident of Los Angeles cannot meet even the most minimal demands of social existence without one. A family that wants to send its children to a good school must buy a house in a good

school district, yet such houses are often beyond reach for families with low relative income. Similarly, if only 10 percent of houses have views and everybody cares equally strongly about having a view, then only people in the top 10 percent of the income distribution will get one.

Measuring the social value of a consumption good by summing what individuals spend on it is similar to measuring the social value of military armaments by summing the amounts that individual nations spend on them. Both measurements are problematic because they ignore the influence of context on demand.

Consider a simple model in which individuals apportion their income between consumption ( $C$ ) and workplace safety ( $S$ ) and in which the representative individual's utility depends not only on her absolute levels of consumption and safety, but also on her relative consumption. For example, suppose the  $i$ th individual's utility is given by<sup>18</sup>

$$U_i = U_i[C_i, S_i, R(C_i)], \tag{14.1}$$

where  $R(C_i)$  denotes her rank in the consumption distribution,  $0 \leq R(C_i) \leq 1$ . If  $f(C)$  is the density function for the observed values of consumption in the population, then

$$R(C_i) = \int_0^{C_i} f(C) dc.$$

Let  $M_i$  denote the individual's income,  $P_c$  the price of the consumption good, and  $P_s$  the price of safety. If the individual takes  $f(C)$  as given, the first-order condition for maximum utility is given by

$$U_{i1}/U_{i2} + [U_{i3} f(C_i)C]/U_{i2} = P_c/P_s, \tag{14.2}$$

where  $U_{ij}$  denotes the first partial derivative of  $U_i$  with respect to its  $j$ th argument.

The second term on the left-hand side of equation (2) reflects the fact that when an individual buys an additional unit of the consumption good, her payoff is not just the direct utility it provides but also the utility from the implied advance in the consumption ranking. But other individuals also perceive this second reward, and when all respond to it, the resulting consumption ranking remains as before. As a result, consumers spend more on consumption and less on safety than is socially optimal.

Suppose consumers could agree collectively to ignore the effect of individual consumption changes on consumption rank – that is, suppose they

could agree to assume that  $R'(C) = f(C) = 0$ . The first-order condition in equation (2) would then simplify to

$$U_{i1}/U_{i2} = P_c/P_s, \quad (14.3)$$

which is the familiar first-order condition from models in which consumption rank does not matter. Suppressing the rank term would lead individuals to consume less and spend more on safety than before. Equation (3), not equation (2), defines the socially optimal allocation.

The driving force behind this market failure is that the utility from consumption is more context dependent than the utility from safety. If utility had been equally context dependent for each good, there would have been no distortion.

Is the extent to which satisfaction depends on context different in different domains? Sara Solnick and David Hemenway recently conducted a survey of graduate students in the public health program at Harvard University in an attempt to answer this question.<sup>19</sup> They began by asking each subject to choose between the following hypothetical worlds:

- A: You earn \$50,000 a year, others earn \$25,000;
- B: You earn \$100,000 a year, others earn \$200,000.

Fifty-six percent of subjects chose the first world. Solnick and Hemenway then asked each subject to choose between worlds in which their relative and absolute income levels were the same, but their relative and absolute vacation times differed:

- C: You have 2 weeks of vacation each year, others have 1 week;
- D: You have 4 weeks of vacation each year, others have 8 weeks.

This time only 20 percent chose the first world, less than half as many as in the first question. On its face, this suggests that satisfaction from consumption is more strongly context dependent than satisfaction from vacation time.

Other important consumption categories also appear to be less sensitive than material goods consumption to interpersonal comparisons. Consider traffic congestion, whose adverse effects on health and psychological well-being are similar to those of prolonged exposure to loud, unpredictable noise.<sup>20</sup> The effect of such noise on subjects in the laboratory occurs independently of the amount of noise to which other subjects are exposed, suggesting that the demand for goods is more context sensitive than the demand for such environmental amenities as freedom from noise and traffic congestion.

Interpersonal comparisons also appear relatively unimportant for savings, at least in the short run. Thus, whereas most of us know what kinds of houses our friends live in and what kinds of cars they drive, we are much less likely to know how large their savings accounts are. But even if everyone's savings balance were on public display, at least some important individual rewards from current consumption would still depend more on context than those from saving. Many parents, for example, might gladly settle for a diminished standard of living in retirement if by saving less they could meet the payments on a house in a better school district.<sup>21</sup> And the same incentives would lead many parents to accept less safe, more regimented, but better paying, jobs. As before, however, the positional gains enjoyed by families that make such choices are offset by the corresponding positional losses experienced by other families.

How might a cost-benefit analyst adjust conventional estimates to counteract the biases introduced by concerns about relative consumption? One simple method would make use of surveys in which subjects are periodically asked to report how much additional income a family would need to maintain a constant level of subjective well-being in the face of a rise in the incomes of others. Using data collected in several European countries, B. M. S. van Praag and Arie Kapteyn estimate an elasticity of roughly 0.6 – that is, that a family would need about a 6 percent increase in its real income to compensate for a 10 percent increase in the incomes of all others in the community.<sup>22</sup> If we take this estimate at face value for illustrative purposes, we can employ it to construct a simple multiplier for adjusting willingness-to-pay values generated by hedonic pricing models.

Suppose, for example, that a study in which wages were regressed on mortality rates in the workplace found that individual workers are willing to give up 2 percent of their incomes each year in exchange for a one in 1,000 reduction in the probability of dying in a workplace accident. This estimate tells us that a worker earning \$50,000 per year would be willing to pay \$1,000 per year for the additional safety, even though the expenditure would reduce his relative consumption by 2 percent. The Kapteyn–van Praag estimate suggests that this worker would be willing to pay roughly \$600 more for the same increment in safety if he could be assured that his relative income would be unaffected by the expenditure – as would be the case, for example, if everyone else made similar expenditures on safety.

An adjustment based on the van Praag–Kapteyn survey data would thus call an upward revision by 60 percent in the willingness-to-pay values inferred from hedonic pricing models. It would be easy to quarrel, of course, with an adjustment procedure based on survey responses like these. Other,

more objective procedures might be pursued. Elsewhere, for example, I have argued that one can infer the value of relative income by examining the relationship between wages, local rank, and productivity among groups of coworkers.<sup>23</sup> In any event, the mere fact that an adjustment procedure may be flawed clearly does not imply that it yields worse estimates than we would get by simply ignoring concerns about relative consumption.

In sum, if demands for some goods are more highly context sensitive than demands for others, then individual spending decisions cannot be aggregated to estimate social valuations for cost-benefit analysis. In general, the sum of individual valuations will be smaller than social value for goods whose demands are relatively sensitive to context and greater than social value for those whose demands are relatively insensitive to context. And because contextual forces influence demands in powerful ways,<sup>24</sup> we have ample reason to be skeptical of hedonic pricing models, even those based on perfectly competitive markets with complete information.

As before, however, the implication is not that the cost-benefit approach is invalid as a matter of principle. Rather, it is that, as currently implemented, its prescriptions may be substantially misleading. If so, the remedy is not to abandon cost-benefit analysis but to amend conventional estimating procedures.

### Impulse-Control Problems and Revealed Preference

Hedonic pricing models also assume that we can infer the values people place on future events by observing the choices they make. On this view, if a person accepts a one in 10 chance of contracting a serious illness 1 year from now in return for a payment of \$100 now, then the cost of taking that risk, expressed as a present value, cannot be more than \$100. Compelling experimental evidence, however, suggests grounds for skepticism.<sup>25</sup> Consider, for example, the pair of choices A and B:

A: \$100 tomorrow versus \$105 a week from tomorrow;

B: \$100 after 52 weeks versus \$105 after 53 weeks.

The rational choice model on which hedonic pricing models are based says that people will discount future costs and benefits exponentially at their respective rates of time preference. If so, people should always choose similarly under alternatives A and B. Since the larger payoff comes a week later in each case, the ordering of the present values of the two alternatives must be the same in both, irrespective of the rate at which people discount. When

people confront such choices in practice, however, most pick the \$100 option in A, whereas most choose the \$105 option in B.

Substantial experimental evidence suggests that individuals discount future costs and benefits not exponentially, as assumed by the rational choice model, but hyperbolically.<sup>26</sup> The psychological impact of a cost or benefit falls much more sharply with delay under hyperbolic discounting than under exponential discounting. One consequence is that preference reversals of the kind just discussed are all but inevitable under hyperbolic discounting. The classic reversal involves choosing the larger, later reward when both alternatives occur with substantial delay, then switching to the smaller, earlier reward when its delay falls below some threshold. Thus, from the pair of alternatives labeled B above, in which both rewards come only after a relatively long delay, most subjects chose the larger, later reward, whereas from the pair labeled A, most chose the earlier, smaller reward.

The tendency to discount future costs and benefits hyperbolically gives rise to a variety of familiar impulse-control problems and, in turn, to a variety of strategies for solving them. Anticipating their temptation to overeat, people often try to limit the quantities of sweets, salted nuts, and other delicacies they keep on hand. Anticipating their temptation to spend cash in their checking accounts, people enroll in payroll deduction savings plans. Foreseeing the difficulty of putting down a good mystery novel in mid-stream, many people know better than to start one on the evening before an important meeting. Reformed smokers seek the company of nonsmokers when they first try to kick the habit and are more likely than others to favor laws that limit smoking in public places. The recovering alcoholic avoids cocktail lounges.

Effective as these bootstrap self-control techniques may often be, they are far from perfect. Many people continue to express regret about having overeaten, having drunk and smoked too much, having saved too little, having stayed up too late, having watched too much television, and so on. The exponential discounting model urges us to dismiss these expressions as sour grapes. But from the perspective of the hyperbolic discounting model, these same expressions are coherent. In each case, the actor chose an inferior option when a better one was available, and later feels genuinely sorry about it.

Hedonic pricing models use observed choices to infer discount rates, which cost-benefit analysts then use to compute present values. To the extent that many important intertemporal choices are driven by hyperbolic discounting, conventional methods will give too little weight to future costs and benefits.

### Status Quo Bias

Opposition to cost-benefit analysis may also stem from the fact that the costs of a policy change are often far easier to quantify than its benefits, especially in the domains of environmental policy and health and safety policy. In both fields, consensus about how to measure benefits has proved especially elusive. The upshot is that policy decisions in these arenas tend to be driven primarily by cost considerations, resulting in a bias in favor of the status quo. This bias may help explain why advocates of change are overrepresented among opponents of cost-benefit analysis.

The fact that benefits are more difficult to measure than costs does not provide a compelling reason to abandon cost-benefit analysis, just as the fact that costs are easier to forecast than revenues does not provide a compelling reason for firms to abandon profit maximization. In each case, we do better to act on the best information available than to act on no information at all.

### Concluding Remarks

From the preceding discussion, I draw two conclusions. One is that critics have failed to offer persuasive arguments that cost-benefit analysis is objectionable as a matter of principle. The other is that many of the methods used by cost-benefit analysts generate systematically biased prescriptions. Hedonic pricing methods overstate the value of goods and activities whose demands are relatively context sensitive. And they give too much weight to current costs and benefits, too little weight to those that occur in the future. These biases suggest an answer to the question posed in my title. Cost-benefit analysis as currently practiced may be controversial simply because it often generates misleading prescriptions.

I conclude by considering a more speculative explanation for opposition to cost-benefit analysis, one rooted in the distinction between consequentialist and deontological moral theories. The deontologists insist that immutable moral principles distinguish right conduct from wrong conduct, irrespective of costs and benefits. They insist, for example, that stealing is wrong not because it does more harm than good, but simply because it violates the victim's rights. The consequentialist resists such absolute prescriptions, confident that there could always be *some* conditions in which the gains from stealing might outweigh its costs.

Yet even the most committed consequentialists seem to recognize that statements like "Stealing is permissible whenever its benefits exceed its costs" are not rhetorically effective for teaching their children moral values. Indeed,



like the deontologists, most consequentialists teach their children that stealing is wrong as a matter of principle. Elsewhere I have argued that once we acknowledge the strategic role of moral emotions in solving commitment problems, this posture is coherent, even in purely consequentialist terms.<sup>27</sup>

Yet a potentially more worrisome aspect of the consequentialist position remains, which is that people who view their ethical choices in cost-benefit terms must also construct their own estimates of the relevant costs and benefits. The obvious concern is that their estimates will be self-serving. More than 90 percent of all drivers, for example, feel sure they are better than average.<sup>28</sup> More than 99 percent of high-school students think they are above average in terms of their ability to get along with others.<sup>29</sup> Ninety-four percent of college professors believe they are more productive than their average colleague.<sup>30</sup> The same forces that make us overestimate our skills can be expected also to distort the estimates that underlie our ethical judgments. And if these self-serving calculations lead some to disregard the common good, their example will make others more apt to do likewise.

Needless to say, people may also be prone to self-serving biases in their interpretations of deontological moral principles. In the end, which approach entails the greater risk is an empirical question. But it is at least possible that consequentialist thinking could lead to a worse outcome on balance. If this were shown to be so, consequentialists would have little choice but to endorse the deontological position (much as an atheist might support fundamentalist religious institutions on the view that threats of hell-fire and damnation are the only practical way to get people to behave themselves). They would have to view cost-benefit analysis as correct in principle yet best avoided in practice.

I hasten to add that critics of cost-benefit analysis have made no such showing. And unless they do, it seems certain that cost-benefit analysis will continue to play an important role in decision making. Under the circumstances, both friends and foes of cost-benefit analysis have a shared interest in trying to eliminate the biases that distort its prescriptions.

#### Notes

1. For an overview, see Robert Kuttner, *Everything for Sale* (1997).
2. The few who did object would likely invoke a variation of the “slippery-slope” argument, which holds that allowing even a single small step will lead to an inevitable slide to the bottom. Yet we move partway down slippery slopes all the time, as when we amend the laws of free speech to prohibit people from yelling “fire” in a crowded theater in which there is no fire.

3. See, for example, Steven Kelman, An Ethical Critique of Cost-Benefit Analysis, 5 Regulation 33 (1981).
4. See, for example, John Jamieson, Carswell Smart, & Bernard Williams, Utilitarianism: For and Against (1973).
5. Stanley Jevons, The Theory of Political Economy 72–73 (1941) (1871).
6. See I. Ritov & Daniel Kahneman, How People Value the Environment: Attitudes vs. Economic Values, in Psychological Approaches to Environmental and Ethical Issues in Management 33–51 (M. Bazerman *et al.* eds. 1997); and Daniel Kahneman & Jack Knetsch, Valuing Public Goods: The Purchase of Moral Satisfaction, 22 J. Envtl. Econ. & Mgmt. 57 (1992).
7. William H. Desvousges, John W. Payne, & David A. Schkade, How People Respond to Contingent Valuation Questions (EPA Grant No. R824310 Final Report, April 1998).
8. Richard Thaler, Toward a Positive Theory of Consumer Choice, 1 J. Econ. Behav. & Org. 39 (1980).
9. Richard Thaler, Precommitment and the Value of a Life, in The Value of Life and Safety 178–79 (M. W. Jones-Lee ed. 1982).
10. *Id.* at 179.
11. Ritov & Kahneman, *supra* note 6.
12. Rebecca Boyce *et al.*, An Experimental Examination of Intrinsic Values as a Source of the WTA-WTP Disparity, 82 Am. Econ. Rev. 1366 (1992).
13. See, for example, Richard Thaler & Sherwin Rosen, The Value of Saving a Life: Evidence from the Labor Market, in Household Production and Consumption 265 (N. Terleckyj ed. 1976).
14. See Lawrence Seidman, The Return of the Profit Rate to the Wage Equation, 61 Rev. Econ. & Stat. 139 (1979), and numerous studies cited therein.
15. See, for example, Richard A. Posner, Economic Analysis of Law (5th ed. 1998).
16. Adam Smith, An Inquiry into the Nature and Causes of the Wealth of Nations, bk. 5, ch. II. pt. II, art. 4 (1952) (1776).
17. On this point, see especially Amartya K. Sen, The Standard of Living (1989).
18. For a more detailed discussion of the model that follows, see Robert H. Frank, The Demand for Unobservable and Other Nonpositional Goods, 75 Am. Econ. Rev. 101 (1985).
19. Sara J. Solnick & David Hemenway, Is More Always Better? A Survey on Positional Concerns, 37 J. Econ. Behav. & Org. 373 (1998).
20. For a survey of the relevant studies, see Robert H. Frank, Luxury Fever, ch. 6 (1999).
21. Some object that a desire for high consumption rank cannot really explain low savings rates, since those who save too little now simply consign themselves to having low consumption rank in the future. Yet, as noted, having lower consumption rank in the future may be an acceptable price to pay for the ability to have high rank with respect to some forms of current consumption. What is more, to the extent that driving the right cars and wearing the right clothes function as signals of ability, and thereby help people land better jobs or more lucrative contracts, low savings now may not even entail reduced consumption rank in the future. But whereas this may be true from the perspective of each individual, it is surely not true for society as a whole. For when all of us

spend more to signal our abilities, the relative strength of each signal remains unchanged.

22. B. M. S. van Praag & Arie Kapteyn, Further Evidence on the Individual Welfare Function of Income, 4 *Eur. Econ. Rev.* 33 (1973).
23. Robert H. Frank, Are Workers Paid Their Marginal Products? 74 *Am. Econ. Rev.* 549 (1984).
24. See Robert H. Frank, *Choosing the Right Pond* (1985); and Frank, *supra* note 20.
25. See, for example, the papers in *Choice over Time* (Jon Elster & George Loewenstein eds. 1993).
26. For detailed summary of the relevant evidence, see George Ainslie, *Picoeconomics* (1992).
27. See Robert H. Frank, *Passions within Reason* (1988). For a related discussion, see Eric A. Posner, *The Strategic Basis of Unprincipled Behavior: A Critique of the Incommensurability Thesis*, 146 *U. Pa. L. Rev.* 1185 (1998).
28. See Thomas Gilovich, *How We Know What Isn't So* (1991).
29. College Board, *Student Descriptive Questionnaire* (1976–77).
30. P. Cross, *Not Can but Will College Teaching Be Improved?* *New Directions Higher Educ.*, Spring 1977, at 1.