

Behavioral Cost-Benefit Economics: Toward a New Normative

Approach to Policy

Nathan Berg*

University of Texas at Dallas

Key Words: Cost-Benefit Analysis, Behavioral Economics,
Status-Quo Effect, Loss Aversion, Overconfidence, Hyperbolic Discounting

*The author thanks Lloyd Jeff Dumas, Simon Fass, and the participants at the Business and Economics Society International 2002 Conference for helpful discussions on the connection between behavioral economics and real-world policy problems. The author is also grateful for the feedback of an anonymous referee whose suggestions contributed significantly.

Berg: Cecil and Ida Green Assistant Professor of Economics, School of Social Sciences, University of Texas at Dallas, GR 31 211300, Box 830688, Richardson, Texas 75083-0688. Phone 1-972-883-2088, Fax 1-972-883-2735, Email nberg@utdallas.edu

ABSTRACT

This paper addresses the question of whether the findings of behavioral economics imply that techniques used in cost-benefit analysis should be modified. The findings of behavioral economics considered include the status-quo effect, loss-aversion, overconfidence and hyperbolic discounting. These behavioral phenomena do indeed imply that concepts from cost-benefit analysis such as consumer surplus, the Kaldor-Hicks criterion, shadow-price valuation, and time discounting, need to be modified. The most important modifications follow from the status-quo effect, which provides a new reason to reject policy proposals that yield only small percentage benefits relative to costs.

I. INTRODUCTION

The last 10 years have seen behavioral economics gain increasing attention among academic and popular audiences (Schwartz 1998). Defined as the application of psychology (and other related behavioral sciences) to economics, research in behavioral economics has led to a series of discoveries showing that typical consumers systematically deviate from the rationality assumptions which are commonly invoked in traditional economic analysis. Although these “behavioral anomalies” are by now widely accepted descriptions of how consumers actually behave (Laibson and Zeckhauser 1998), the key findings of behavioral economics have yet to influence the way many policy-oriented economists arrive at policy prescriptions (Knetsch 1995).

The purpose of this paper is to reconsider the workhorse of applied normative economics — cost-benefit analysis — and ask whether the findings of behavioral economics point toward a new framework for normative analysis. In particular, four well-documented behavioral anomalies are considered: the status-quo effect, loss-aversion, overconfidence, and hyperbolic discounting. In each case, this paper analyzes the consequences of these descriptive behavioral phenomena for the fundamental normative elements of cost-benefit analysis: consumer surplus, the Kaldor-Hicks criterion, time discounting, and the shadow-price approach to valuation.

II. THE STATUS-QUO EFFECT

The *status-quo effect* refers to any instance where subjective valuations of an object depend on whether the object is being acquired or given up. Perhaps the most famous example of the status-quo effect is an experiment by Kahneman, Knetsch, and Thaler (1990), where experimental subjects were found to demand roughly twice as much money in order to give up a coffee mug (which they had been given earlier) than they were willing to pay to acquire it. The status-quo effect contradicts the idea that willingness to pay is a good indication of how much compensation should be given for taking away a good or service which is currently being consumed. For example, when the status-quo effect is present, the average amount residents of a particular neighborhood are willing to pay in order to acquire a new park (where there is currently no park) will generally be less than the amount those residents need to be compensated in order to accept the removal of a park already in existence.

This section argues that the status-quo effect, when taken seriously as a usual aspect of human preferences rather than disparaged as an instance of “irrationality,” implies that fewer policy changes are socially desirable than otherwise would be the case. For instance, if a policy change causes winners to receive \$1,000 and losers to lose \$800 (with equal numbers of each), even a \$900 transfer from winners to losers may not be enough to offset the disutility associated with moving away from the status quo. Those who lose an object or service with a market price of \$800 may require compensation of significantly more than \$800 in order to be “made whole.” If one accepts that this aspect of preferences is widespread and ought to be acknowledged in policy analysis, then the status-quo effect leads to increased conservatism with respect to the status quo, whereby projects with only slightly positive net benefits will not be socially desirable. Because of the status-quo effect, the psychic loss from change itself must be explicitly accounted for or, equivalently, the behavioral cost-benefit decision rule should become, “Adopt the proposed policy only if net benefits exceed a prespecified threshold (strictly greater than zero).”

The status-quo effect implies that the consumption of a commodity at time t , denoted x_t , yields varying levels of utility depending on what was consumed previously, denoted x_{t-1} . This generates a new implication about preferences over streams of consumption through time. In particular, consider the consumption path from time 0 to time T :

$$1, 1, 1, \dots, x_t, (1 + \epsilon), (1 + \epsilon), \dots, (1 + \epsilon). \quad (1)$$

This consumption path increases through time, aside from the possible exception of x_t which is free to be bigger or less than 1. With the status-quo effect, the consumption path in (1) would be preferred over a consumption path with equal present value in which the consumer gets to consume more early on and winds up with less at the end, such as

$$(1 + \epsilon\gamma^{T-t+1}), (1 + \epsilon\gamma^{T-t+1}), (1 + \epsilon\gamma^{T-t+1}), \dots, x_t, 1, 1, \dots, 1, \quad (2)$$

where $\gamma \in (0, 1)$ is the discount factor used in computing the present value of consumption paths.

An important consequence of this for policy planning is that sustainable projects should be favored over those with mostly short-run benefits. Plans where the benefits increase over the life of the project are preferred over other consumption streams with equal present value, simply because there is less pain for consumers to bear in parting with the level of consumption to which they become accustomed. With this interpretation of habituation rather than that of giving up an actual physical object, the status-quo effect is sometimes referred to in connection with a variety of related phenomena under headings such as “reference-point dependent preferences,” “habit formation,” “the endowment effect,” and “the framing effect.”

Another key normative concept from cost-benefit analysis is the Kaldor-Hicks principle, which is often invoked when considering a policy change that affects some people adversely while benefitting others. The Kaldor-Hicks principle holds that if the winners could (at least theoretically) compensate the losers so that they feel no worse off than before, and still wind up with a net surplus, then the policy change under consideration should be adopted. The main consequence of the status-quo effect here is that the Kaldor-Hicks principle becomes more stringent. This is because there is a gap in perspectives between winners and losers, where a larger transfer of benefits is required to make losers feel indifferent rather than hostile to policy changes.

For example, one may consider the pros and cons of being sent \$600 in the mail from the government in the eyes of voters who exhibit the status-quo effect. A nearly identical policy was actually implemented by the Bush Administration in 2001 as part of a tax-cutting fiscal stimulus plan. One problem with the policy that traditional cost-benefit analysis overlooks is that the psychic benefit from being sent \$600 in the mail is significantly smaller than the psychic cost of giving up an additional \$600 of lost government benefits (or lost income in the form of additional taxes) required in order to pay for the benefit. By the Kaldor-Hicks criterion, the status-quo effect implies that fewer projects and policy changes will be worthwhile, calling for a generally conservative attitude with respect to the status quo.

A second example involving the Kaldor-Hicks principle centers on the question of whether monetary or fiscal policy is better suited for smoothing business cycles. Although the current thinking of most macroeconomists favors monetary policy (Abel and Bernanke 2001, p. 417), the status-quo effect adds at least one new point in favor of fiscal policy as a tool for stimulating national output during economic downturns. When a central bank cuts interest rates to stimulate output, this amounts to a transfer from savers to borrowers. The social benefit to borrowers is probably less than the loss of utility suffered by savers, owing to the status-quo effect. On the other hand, the dis-utility borne by future generations who are asked to pay for today’s increases in government spending may similarly outweigh the benefits of business-cycle smoothing today. The increased tax bills owed by future generations will not, however, be accompanied by the same degree of psychic pain in moving from a low to a high tax bill, because many future consumers will simply not be old enough to remember the lower tax rates. Therefore, borrowing from future generations in order to finance more government spending today may be better than using monetary policy to achieve the same goal, although the morality of such a plan may be called into question.

Next, the consequences of the status-quo effect for “consumer surplus” and “producer surplus” analysis of costs and benefits are considered. The conventional efficiency analysis using the concept of consumer surplus nearly always suggests that cutting taxes is an unambiguously “good thing.” If a tax on producers is reduced, the lost revenue is more than offset by increased producer surplus and additional consumers surplus, depicted graphically in introductory textbooks as regaining a portion of the “deadweight loss” attributable to the higher-tax status quo.

In symbols, the conventional analysis states that lost tax revenue (R) is offset by an even greater amount of surplus ($R + \text{deadweight loss}$) in the wake of a tax cut. The status-quo effect may, however, turn this result on its head. Because of the status-quo effect, it cannot be assumed that the change in utility associated with a gain in R is equal to the change in utility associated with a loss of R . In fact, when the disparity between willingness to acquire R and willingness to accept losing R is sizable, and the deadweight loss is

not so large (i.e., when the supply or demand curve is inelastic), then it is entirely possible that the utility lost in giving up R dollars worth of benefits is greater than the increase in utility associated with gaining a new benefit of $R +$ deadweight loss. In light of the status-quo effect, the normative implication regarding tax cuts is ambiguous. Although there is more aggregate surplus with a tax cut, the increase in utility due to cutting taxes may not offset the loss of utility associated with public expenditures at its status-quo level.

Finally, the shadow-price approach to valuing a public good, another of the key normative elements of cost-benefit analysis, is re-considered here in light of the status-quo effect. If one models a policy-maker's decision process as a maximization problem (where the objective function depends on a list of variables under the decision-maker's control and another list of factors which are not), then the shadow price of a resource is defined as the increase in the maximized value of the objective function that results from increasing the availability of that resource by one unit. For example, if a consumer maximizes the utility function $u(x, y)$ subject to the budget constraint $p_x x + p_y y = W$ (where W is the consumer's wealth, and p_x and p_y are the respective prices of x and y), then the shadow price of money is the change in maximized utility that results from an extra dollar of wealth. This approach is very general and widely used for valuing a public good or a finite resource, assuming that all other resources are efficiently allocated. That is, changes in social welfare represent comparisons of two maximized values.

The key point with regard to shadow prices is to ask whether the meaning attached to them requires modification because of the new stylized facts of behavioral economics that have emerged in the last decade and a half. In the case of the status-quo effect, new difficulties in using shadow price valuation do in fact arise. For the purpose of illustrating these difficulties, let $V(x; z)$ denote the social welfare function, where z denotes the policy-maker's control variable and x denotes the decision of a representative consumer. Supposing that the policy question under consideration is whether to increase z or not, the standard cost-benefit approach dictates that the policy-maker should compare the cost of increasing z one unit with the marginal benefit $\frac{\partial \max_x V(x; z)}{\partial z}$.

The difficulty that the status-quo effect causes relates to whether this derivative is the same on the right as on the left. Although it may seem at first glance to be a dull technical detail, there is a substantive question about human behavior at the heart of the issue of left- and right-hand sided derivatives as they concern the status-quo effect. If the lost utility from losing 1 unit of z is bigger than the utility gained by increasing z one unit, even for arbitrarily small units, then consumer preferences exhibit the status-quo effect, and the social welfare function is discontinuous at the threshold point. In contrast, when V is a smooth function of z , it can be arbitrarily well approximated by a linear function, meaning that "change to the left" and "change to the right" are equal when those changes are small. The status-quo effect, however, implies that the left-hand-side derivative does not equal the right-hand-side derivative at the status quo z_0 . Formally, the status-quo effect can be defined in this context by the inequality

$$\frac{V(x^*; z_0 + \epsilon) - V(x^*; z_0)}{\epsilon} < \frac{V(x^*; z_0) - V(x^*; z_0 - \epsilon)}{\epsilon}, \quad (3)$$

where the maximizer of V in x is denoted x^* . In words, this equation says that the social gain of moving from the status quo z_0 to $z_0 + \epsilon$ is less than the gain achieved by staying with the status quo instead of reducing z_0 to $z_0 - \epsilon$. As before, the social cost of reducing benefits (which is the right-hand side of (3)) exceeds the social gain of increasing benefits (on the left-hand side), making it necessary to distinguish between gains and losses when evaluating the net benefits of policy proposals.

An anonymous referee pointed out that any enlightened cost-benefit analysis would already account for the costs of moving away from the status quo without needing behavioral economics to make this point. I agree that reasonable observers earnestly considering the impact of a policy proposal would think about the pain of change itself, without ever consulting the behavioral economics literature. However, the use of comparative statics, perhaps the dominant technique for evaluating the desirability of changes in policy parameters, is founded on the idea that adjustment costs are negligible. The rationale is that, if the pain of adjusting to a new equilibrium (after exogenous parameters change) is bounded, then it still makes sense to switch to a new equilibrium in which the economic pie is epsilon bigger, because the benefits at the new equilibrium will accumulate indefinitely into the future.

The contribution of behavioral economics is to give demonstrable evidence that preferences depend on the status quo. Of course it remains unclear how policy makers ought to weight the before- and after- preferences of an individual whose preferences change because of a policy intervention. What is clear, however, is that

traditional comparative statics give 100 percent weighting to the post-intervention preferences. If one places weight on pre-intervention preferences instead, then fewer new policy proposals will be seen as desirable.

III. LOSS AVERSION

Loss aversion, in contrast to risk aversion, characterizes preferences that are risk-loving over losses, and risk-averse over gains. This is not the only definition in circulation, however. Some authors describe as “loss averse” any expected utility function that is “kinked” at the status quo, with the slope to the right-hand side (over gains) strictly less than the slope to the left-hand side. This kinked utility function that weights losses heavier than gains is equivalent to the status-quo effect discussed in the previous section. Therefore, this section focuses on the first definition, i.e., risk-aversion over gains and attraction to risk when evaluating uncertain losses. See Kahneman, Knetsch and Thaler (1991) for evidence, definitions, and background information relating loss aversion to other behavioral “anomalies.”

Focusing on the risk-loving/risk-averting definition, loss aversion can be characterized in terms of concavity and convexity of the expected utility function. Of course, risk-averse preferences (subject to regularity conditions described in Mas-Colell, Whinston and Green 1995 pp. 167-173) can be represented by a globally concave expected utility function. Loss-averse preferences, on the other hand, are represented by an expected utility function that is concave over gains (relative to a threshold level of payoffs which may, for convenience, be taken to be zero) and convex over losses.

The idea behind loss aversion is that consumers, on average, are risk averse over gains but would prefer to gamble when dealing with losses (Kahneman, Knetsch and Thaler 1991). For example, a loss-averse consumer might prefer \$500 with probability 1 over flipping a fair coin generating payoffs of \$1000 or \$0, reflecting risk aversion over positive payoffs. But when asked to decide between giving up \$500 with probability 1 or flipping the fair coin generating *losses* of either \$1000 or \$0, that same consumer, if she is loss averse, prefers to gamble.

Loss aversion raises potentially serious issues for policy-makers, because the collective will of voters may support action that takes big gambles with extremely bad outcomes. Under loss-aversion, once things are bad, i.e., once payoffs are sufficiently negative, the additional loss in utility from losing more is negligible. For instance, in contemplating a nuclear accident where 10,000 people die and a costly safety measure that might help reduce deaths to 1,000, loss-averse policy-makers and voters may not place much value on avoiding 9,000 deaths. The reason is that 1,000 deaths is already so terrible that an additional loss of 9,000 lives causes far less dis-utility than nine times the dis-utility of 1,000 deaths. Loss averse preferences may lead decision-makers to develop projects that might cause catastrophes, i.e., policies which are objectively bad from a pure frequency point of view (negative expected payoffs) but which seem reasonable to most people because “very bad” outcomes are penalized only slightly more than merely “bad” outcomes are.

Given this propensity to not place much weight on extremely bad outcomes, the public may fail to demand enough compensation to cover the expected losses associated with a new risk accompanying a policy change. This causes a problem for those relying on the Kaldor-Hicks principle in justifying their policy ideas, because *ex ante* subjective benefits are so far out of line with *ex post* reality. The following numerical example illustrates how a loss-averse voter may accept compensation that is far less than the expected loss associated with taking on a new risk.

Suppose the expected utility function

$$u(x) = x^{0.5} \text{ if } x > 0 \quad \text{and} \quad u(x) = -(-x)^{0.5} \text{ if } x < 0 \tag{4}$$

represents a loss averse voter’s preferences. Then this voter will accept being subjected to the gamble “negative \$10,000 with probability 0.5 and \$0 with probability 0.5” in exchange for \$50, even though the expected loss from this policy change is \$5,000.

This possibility that society will unwittingly accept risks with negative net expected payoffs implies that loss aversion should be explicitly considered in analyzing policy issues where there is a small probability of catastrophe. (See Kunreuther 1996 for more applications of behavioral theory to the applied problem of disaster insurance.) Nuclear arms, genetically modified agriculture, and unilateral military action are leading examples of such issues.

Loss aversion has other surprising consequences for the key normative elements of cost-benefit analysis. For one thing, demand curves may slope upward as a result of loss aversion. To understand why, one may

think of a two-state payoff problem where there is an “up” state with a positive payoff and a “down” state with negative payoff. In order to increase the down-state payoff, the up-state payoff must be reduced, because payoffs are scarce. Optimality for an interior choice of the down-state payoff requires equalizing marginal benefits across both states. When the price of down-state benefit (in terms of up-state benefit forgone) goes up, the up-state payoff falls and marginal utility in the up-state increases. To equate marginal benefits across both states, the marginal benefit in the down-state must be brought up. But because the expected utility function is convex in the region of negative payoffs, this means that the down-state payoff must be increased. Thus, the demand for down-state payoffs is an increasing function of own price, at least over a certain range.

Upward sloping demand curves change the standard consumer surplus analysis significantly. When market demand is increasing in price, taxing producers can cause consumer surplus to rise, although there is still a deadweight loss to consider in terms of aggregate surplus. It may only be a theoretical curiosity, but when demand slopes upward, it is actually possible to increase aggregate surplus by taxing consumers, although this involves a sizable transfer of surplus away from consumers. If the distribution of surplus among consumers and producers is a concern, then the consumer-welfare-improving tax on producers may be a relevant alternative for policy-makers to explore.

Finally, the implication of loss aversion for the shadow-price approach is considered. Whether there is anything new in this regard depends on whether the expected utility function is kinked at the threshold point (usually taken to be the zero-payoff point) or not. The reason for this is entirely similar to the issues at stake regarding discontinuity caused by the status-quo effect. In short, loss aversion can lead to discontinuities (although it does not follow necessarily) and, in such cases, shadow prices are poorly defined. If our attitudes jump discretely when considering even small losses compared to small gains, then the expected utility function is kinked and shadow values at that value will not be equal when evaluated as left-hand-side or right-hand-side derivatives. The meaning of this relates directly to the discrete difference in marginal benefit when considering losses as compared to gains.

IV. OVERCONFIDENCE

Overconfidence has been defined in more than one way in the economics literature. The experimental evidence regarding overconfidence is reviewed in Daniel, Hirshleifer, and Subrahmanyam (1998). For purposes of re-considering cost-benefit analysis, it is sufficient to define overconfidence as a systematic distortion in consumer beliefs about the probability distributions of random variables in the economy that leads to an artificially high level of demand. It is therefore meaningful to distinguish between an undistorted demand curve (which would prevail if consumers’ beliefs matched objective probability distributions) and a distorted or “overconfident” demand curve that lies to the right of the undistorted demand curve (in Marshallian price-quantity space).

If a policy-maker suspects that consumers are overconfident and demand too much of a particular product, e.g., smokers whose beliefs about the health consequences of their actions are systematically biased and therefore demand too many cigarettes, then a tax (either on consumers or on producers) can restore the market to the undistorted level of output. The partial-equilibrium surplus arguments against taxes do not apply here, because excessive expenditure in one market implies under-expenditure in others, meaning that a gap between prices and marginal social benefits prevails. Thus, overconfidence and other behavioral phenomena involving distorted beliefs about probability distributions imply that consumer surplus arguments are inappropriate.

As for the Kaldor-Hicks principle, the willingness of those who lose under a new policy to accept compensation in exchange for agreeing to the policy change will be distorted relative to the objective value of benefits and transfers. Systematically skewed perceptions must be brought in explicitly when invoking the Kaldor-Hicks principle. On a case-by-case basis, the divergence between current perceptions about the value of benefits and transfers and objective values, if measurable, need to be considered in order to maximize social welfare.

Similarly, shadow-price valuations can be systematically skewed because of overconfidence. “Distorted demand” is nearly equivalent to “distorted marginal benefit,” which ultimately leads to a distorted social welfare function as long as consumer perceptions are taken at face value. This raises the question of whether policy-makers should replace subjective benefits of consumers which they feel are distorted with “corrected” valuations. Policy-makers who distinguish between an objective demand curve and an overconfident one will subject themselves to the charge of paternalism. However, when there is compelling empirical evidence that

consumers make systematic mistakes, there is little ethical ground for asserting that policy-makers should, in the name of consumer sovereignty, maximize social welfare as mis-perceived by consumers. It can be concluded that, at this point, the normative implications of overconfidence are trickier to deal with than those of the other behavioral phenomena discussed above.

V. HYPERBOLIC DISCOUNTING

If a consumer is indifferent between \$100 in seven days and \$105 dollars in eight days, yet prefers \$100 today over \$105 tomorrow, then that consumer's attitudes toward time may be characterized as reflecting *hyperbolic discounting*. When a consumer's preferences exhibit hyperbolic discounting, it means that consumer is impatient in the short-run and relatively patient in the long-run. Hyperbolic discounting has been put forward as a leading explanation of phenomena such as procrastination in filing income tax returns, the will-power problems faced by dieters, the dilemmas faced by anyone planning for their own retirement, and the policy problems surrounding social security programs (Rabin 1996).

A more formal characterization of hyperbolic discounting requires one to distinguish between consumption at three points in time: x_0, x_1 , and x_2 . At time 0, a hyperbolic discounter discounts the utility of consuming $x_2, u(x_2)$, by the discount factor $\gamma^2\alpha$, and discounts the utility of $x_1, u(x_1)$, by the discount factor $\gamma\alpha$. The relative discount rate of x_2 in terms of x_1 at $t = 0$ is simply $\frac{\gamma^2\alpha}{\gamma\alpha} = \gamma$. The additional term α is a second discount factor reflecting the myopia of hyperbolic discounters. To see this, consider that the discount rate moving from time 0 to time 1 is $\gamma\alpha$, which is smaller than the discount rate moving from time 1 to time 2 (from the perspective of time 0). Thus, the hyperbolic discounter may be described as choosing x_0, x_1, x_2 to maximize the objective

$$u(x_0) + \alpha\gamma u(x_1) + \alpha\gamma^2 u(x_2). \quad (5)$$

The essence of hyperbolic discounting is that the demand for x_2 relative to x_1 is greater at time 0 than it is at time 1. At time 1, the myopic nature of the hyperbolic discounter dominates and x_2 is seen as relatively less valuable. But at time 0 when x_1 and x_2 are both in the future, x_2 is a better substitute for x_1 . This means that the demand curve for x_2 as a function of the relative price of x_2 in terms of x_1 actually shifts inward when moving from time 0 to time 1.

Consumer surplus is therefore sensitive to the period in which it is measured. And shadow prices are different at different points in time. Perhaps most problematic of all, projects with benefits far in the future are under-valued from the perspective of tomorrow, while those with costs far in the future are over-valued in the present.

Realizing that voters and policy-makers alike may exhibit hyperbolic discounting preferences is important. When consumers or voters are thought to be hyperbolic discounters, there is a gap between how they value relative benefits in the present versus the future. Ethicists may debate which among a single individual's multiple time-dated preferences ought to prevail for the purposes of designing policy. But even without waiting for a concrete answer to that difficult problem, those who conduct cost-benefit analysis will want to consider the time-specific nature of valuations, and acknowledge the possibility that individual voters or consumers may have a tendency to contradict themselves through time in their opinions on whether a project is worthwhile or not.

In general, cost-benefit decisions with immediate benefits and deferred costs are the most likely to be reversed through time, as voters who have hyperbolic discounting preferences move from assigning net positive valuations to assigning net negative valuations for the same project. To counter this tendency, a cost-benefit analyst may want to consider using a set of discount factors that discount at a decelerating rate (in contrast to the accelerating rate characterizing hyperbolic discounters). The rationale for this would be to offset subjective valuations that tend to be too present-oriented because of hyperbolic discounting. For example, if cost-benefit analysis calls for benefits (and costs) one year from now to be discounted by a factor of γ , then benefits and costs two years from now could be discounted by less than γ^2 to control for the tendency of subjective valuations to place too much weight on the present. One such sequence of weights is

$$1 + \gamma^T - \gamma^{(T-t)} \quad (6)$$

for $t = 0, 1, \dots, T$.

VI. SATISFICING

Herbert Simon (1957) popularized the notion that, due to cognitive limitations and computational complexity, decision-makers often fail to maximize and instead search their choice set until a sufficiently high level of satisfaction is achieved — a decision process referred to as *satisficing*. The important point is that decision-makers do not typically maximize. Rather than maximize, those who satisfice experiment until an action that is “good enough” is achieved. Therefore, in departing from the maximization framework which is dominant within economics, satisficing represents a radical break from tradition. In the context of cost-benefit analysis, satisficing implies a number of further breaks with tradition. It is not straightforward, however, to see how cost-benefit analysis ought to be modified in light of what is known about satisficing.

When consumers do not maximize, their demand curves can no longer be interpreted as “marginal benefit curves,” where marginal benefits are measured along the price axis. In this case, the rationale for using changes in consumer and producer surplus to analyze economic efficiency evaporates. Similarly, if consumers do not maximize their own utility, one may reasonably wonder how a policy-maker can hope to do so. What this means is that shadow prices are probably not well defined and, even if they were, they would be exceedingly difficult to measure.

The Kaldor-Hicks criterion, on the other hand, may actually be easier to implement when large-scale satisficing prevails. Because consumers do not maximize, those who lose under new policies may not demand anything more if they began from a point of satisfaction that comfortably exceeded their target level of utility. The target level of utility, in turn, raises a variety of other issues in connection with responsible or optimal policy-making. For instance, one wonders if there is any rationale at all for policy intervention if those who benefit are already happier than they “need” to be from the standpoint of their target utility. The issue of whether policies designed to change the thresholds themselves are worthwhile is yet another deep question which is raised by behavioral economics without being resolved. One counter-intuitive prospect to keep in mind is that, as standards of living rise, so do most individuals’ utility targets, leading to a condition where more people fall short of their target and live in a state of subjective deprivation as a consequence. Such possibilities may complicate the meaning of “economic progress.”

VII. APPLICATION: DRINKING ON CAMPUS

At this stage, after reviewing how developments in behavioral economics force one to modify standard cost-benefit analysis, it is appropriate to ask whether the modifications, as a whole, are useful. The question is whether behavioral economics points toward a new, more empirically valid normative procedure for answering the prescriptive policy questions facing society. I assert that the answer to this question is “yes.” This positive assessment must be qualified, however, with the acknowledgment that the findings of behavioral economics do not always point clearly to specific modifications.

In order to explore how far one might progress in applying the behavioral cost-benefit ideas described above, it is helpful to take up a specific policy issue such as one which is being debated on the campus where I work: whether a new pub located on our campus ought to be allowed to sell liquor. Due to space limitations, a detailed account is not presented here. Rather, several modifications to the cost-benefit analysis of liquor on campus are outlined, each corresponding to one of the behavioral findings described above.

On the cost side, drinking on campus would require additional security, additional insurance premiums, the administrative costs of acquiring a liquor license and engaging in the political struggle with city officials (in Richardson, Texas, where retailers are currently forbidden from selling “hard” liquor), and an expected loss of certain students whose families might be opposed by the idea of on-campus drinking. On the benefit side must be counted the increased pleasure and time savings of students and faculty who enjoy an atmosphere in which drinks are sold. Additionally, students who live on campus would not need to drive home after drinking, implying that the new policy might be expected to reduce drunk driving among the student body. The sale of liquor on campus can also be expected to build a certain amount of new social capital in the form of stronger friendships, more collaboration, and a greater accumulation of shared experiences that contribute positively to socially constructive pursuits such as education, job placement, and alumni fund raising.

A basic modification reflected in the behavioral cost-benefit procedure follows from the status-quo effect and the idea that change, itself, is painful. The “change is painful” principle implies that the decision rule dictating whether liquor on campus is a good idea or not, should not be “Sell alcohol if the benefits outweigh the costs.” Rather, the decision rule should be, owing to status-quo effect, “Sell alcohol if the benefits outweigh the costs by a certain pre-specified amount.” Opponents and proponents should try to agree in advance by how much benefits must exceed costs in order for the policy proposal to be adopted.

They may want to use psychological studies and parameter estimates from other sources to quantify the extent to which status-quo effects are important. An agreeable policy rule might be: “If the ratio of benefits to costs exceeds 1.10, the policy should be implemented.” The extra 10 percent compensates for the “pain of change” which is otherwise unaccounted for.

Loss aversion is relevant to the liquor on campus issue in placing values on extremely negative outcomes. Knowing that loss-averse students and decision-makers have a propensity to gamble over extremely bad outcomes, special care should be taken that the costs associated with an increase in student deaths, for instance, are accounted for in line with objective frequencies insofar as they are available.

Overconfidence may or may not be relevant in the case of campus liquor policy. Overconfidence could lead to exaggerated projected benefits of having liquor on campus. But it could equally well lead to exaggerated projected benefits of being a liquor-free campus. There seems to be no unambiguous implication of overconfidence here.

Hyperbolic discounting could play a potentially important role in analyzing the liquor-on-campus policy question. Many of the benefits identified — those related to new social capital and better alumni fundraising prospects — are payoffs that will accumulate relatively far into the future. Many of the costs, on the other hand, e.g., the political costs and one-time start-up costs of implementing a new bar on campus, must be dealt with in the near term. If the valuations of near-term and short-term benefits are subject to decision-makers who are hyperbolic discounters, then this will tend toward defeating the liquor-on-campus proposal. In fact, it is common to hear administrators and students who are in a position to try gaining approval on-campus liquor sales to articulate the view that, “It would take so long to go through all the hassles of obtaining a liquor license that the fight isn’t worth it.” At this point, should the cost-benefit analyst take this kind of heavy discounting of the distant future seriously, accepting it as the way consumers right now actually feel? Or should some corrective weights be applied that give added weight to the future? If both sides of this policy debate agreed that everyone concerned tends to be overly focused on the near-term and that corrective measures ought to be taken, then a set of corrective weights, like those recommended above, might be considered. It should be noted, however, that this may contradict the spirit of earlier suggestions to accept the status-quo effect rather than to ignore it as an anomaly. Ultimately, policy makers must decide which behavioral phenomena deserve standing in the analysis and which ought to be considered anomalous.

In finding dollar valuations for the increased utility of those who want alcohol on campus and the disutility of those who do not, the non-equivalence of willingness to pay and willingness to be compensated for having a policy change against one’s wishes should be taken into consideration. In this case, the amount that alcohol-opponents would be willing to pay in order to change a “wet” campus to a “dry” campus would underestimate what those opponents lose in moving from a dry campus to a wet campus, because of status-quo effects. Attention to this asymmetry of benefits with respect to the status quo should be kept in mind when determining dollar values for policy-related changes in subjective well being.

VIII. CONCLUSION

The approach outlined here begins with a short list of recently established behavioral tendencies which are generally not considered in traditional economic analysis, and attempts to pro-actively avoid systematic mistakes in public decision-making which may arise because of those behavioral tendencies. Rather than merely implying that traditional cost-benefit analysis is flawed, the new stylized facts of behavioral economics (which include the status-quo effect, loss aversion, overconfidence, and hyperbolic discounting) actually point to a constructive technique for modifying cost-benefit analysis. The modified cost-benefit analysis is referred to here as “behavioral cost-benefit analysis.”

Two criticisms are often leveled at the idea of drawing normative implications from behavioral economics. This article, in several places, considered adjusting subjective valuations in order to correct for systematic gaps between perceptions and objective outcomes that can lead to serious problems in the domain of public policy. Many critics regard the idea of such adjustments as paternalistic. A second complaint comes from those who are not convinced by the empirical record on which the behavioral phenomena discussed in this article rest. To these critics, the modifications proposed here are viewed as *ad hoc* rationalizations for tinkering with otherwise objective analysis. These charges can be at least partially refuted by turning back to the experimental and field research that led to the discovery and subsequent testing of the behavioral phenomena discussed here. More research is underway to refine our understanding of these phenomena, as researchers subject these ideas to tougher and tougher tests.

In response to the second charge of making *ad hoc* assumptions, it should be pointed out that interest in behavioral economics grew out of dissatisfaction with the *ad hoc* nature of traditional rationality assumptions and their failure to match actual economic behavior. Including behavioral phenomena in cost-benefit analysis should instead be seen as an analytical choice that takes consumer preferences even *more* seriously than before. Consequently, by insisting on beginning with a description of typical consumer behavior which is as accurate as possible, more weight is given to testing out the basic assumptions and avoiding *ad hoc* modeling choices. Ultimately, policy-making that rests on a set of descriptively false behavioral assumptions will be less apt to do what voters want. Behavioral cost-benefit analysis begins with this premise and attempts to make good on the hope that improved empirical validity with regard to consumer behavior translates into improved policy action.

REFERENCES

- Abel, Andrew B., and Ben S. Bernanke, *Macroeconomics*, 4th ed. (Boston: Addison Wesley, 2001).
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam, "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*. Vol. 53, No. 6, 1998, pp.1839–1885.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler, "The Endowment Effect, Loss Aversion, and Status Quo Bias: Anomalies," *Journal of Economic Perspectives*, Vol. 5, No. 1, 1991, pp.193-206.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler, "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, Vol. 98, No. 6, 1990, pp.1325-48.
- Knetsch, Jack L., "Assumptions, Behavioral Findings, and Policy Analysis," *Journal of Policy Analysis and Management*, Vol. 14, No.1, 1995, pp.68-78.
- Kunreuther, Howard, "Mitigating Disaster Losses Through Insurance," *Journal of Risk and Uncertainty*, Vol. 12, No. 2, 1996, pp.171-87.
- Laibson, David, and Richard Zeckhauser, "Amos Tversky and the Ascent of Behavioral Economics," *Journal of Risk and Uncertainty*, Vol. 16, No. 1, 1998, pp.7-47.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory*, (New York: Oxford University Press, 1995).
- Rabin, M., "Psychology and Economics," Working Paper, University of California at Berkeley, 1996.
- Schwartz, Hugh H., *Rationality Gone Awry? Decision Making Inconsistent with Economic and Financial Theory*, (Westport: Praeger, 1998).
- Simon, Herbert, *Models of Man*, (New York: John Wiley and Sons, 1957).