

Behavioral Economics and the Conduct of Benefit-Cost Analysis: Towards Principles and Standards

Revised Draft

October 2010

Lisa A. Robinson, Independent Consultant
(Lisa.A.Robinson@comcast.net)

James K. Hammitt, Harvard University and Toulouse School of Economics
(jkh@harvard.edu)

Acknowledgements: This article builds on work initiated to support *Toward Principles and Standards in the Use of Benefit-Cost Analysis: A Summary of Work and a Starting Place* (2010), by Richard O. Zerbe, Jr., Tyler Blake Davis, Tyler Scott, and Nancy Garland, with funding from the John D. and Catherine T. MacArthur Foundation. Dr. Hammitt also received financial support from the French national institute for agricultural research (INRA) and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement no. 230589. We thank Daniel Herz-Roiphe, Brigitte Madrian, and Kerry Smith for their very helpful reviews of previous drafts, as well as participants in the June 2009 "New Ideas for Risk Regulation" conference who influenced the development of this paper. The opinions expressed are, of course, our own.

TABLE OF CONTENTS

ABSTRACT	3
1. INTRODUCTION	4
1.1 BEHAVIORAL VS. NEOCLASSICAL ECONOMICS.....	5
1.2 BEHAVIOR AND BENEFIT-COST ANALYSIS.....	5
2. VALUING PSYCHOLOGICAL ATTRIBUTES	9
2.1 WILLINGNESS TO PAY VS. WILLINGNESS TO ACCEPT COMPENSATION.....	10
2.2 PSYCHOLOGICAL RESPONSES TO RISK	12
2.3 CONCLUSIONS AND IMPLICATIONS	14
3. ESTIMATING TIME PREFERENCES	15
3.1 EXPONENTIAL DISCOUNTING	17
3.2 HYPERBOLIC DISCOUNTING	18
3.3 CONCLUSIONS AND IMPLICATIONS	19
4. SEPARATING PRIVATE FROM SOCIAL PREFERENCES	21
4.1 SOCIAL PREFERENCES WITHIN THE TRADITIONAL FRAMEWORK	21
4.2 BEHAVIORAL RESEARCH ON SOCIAL PREFERENCES.....	23
4.3 CONCLUSIONS AND IMPLICATIONS	24
5. IMPROVING VALUATION STUDIES	25
5.1 STATED-PREFERENCE RESEARCH.....	25
5.2 REVEALED-PREFERENCE RESEARCH	26
5.3 CONCLUSIONS AND IMPLICATIONS	28
6. DETERMINING THE ROLE OF BENEFIT-COST ANALYSIS	28
6.1 SUMMARY OF CONCLUSIONS AND RECOMMENDATIONS	29
6.2 BENEFIT-COST ANALYSIS AS APPLIED PRAGMATISM.....	31
REFERENCES	34

ABSTRACT

As traditionally conducted, benefit-cost analysis is rooted in neoclassical welfare economics, which assumes that individuals act rationally and are primarily motivated by self-interest, making decisions that maximize their own well-being. Its conduct is now evolving to reflect recent work in behavioral economics, which integrates psychological aspects of decisionmaking. We consider several implications for analyses of social programs. First, benefit-cost analysis often involves valuing nonmarket outcomes such as reductions in health and environmental risks. Behavioral research emphasizes the need to recognize that these values are affected by psychological as well as physical attributes. Second, benefit-cost analysis traditionally uses exponential discounting to reflect time preferences, while behavioral research suggests that individuals' discounting may be hyperbolic. However, steep near-term rates may largely reflect impulsive behavior and self-control problems. Third, behavioral research emphasizes the influence of social preferences on valuation. In addition to acting altruistically, individuals may act reciprocally to reward or punish others, or use the status of others as the baseline against which to assess their own well-being. Fourth, behavioral economics identifies factors that can help develop valuation studies that provide well-informed, thoughtful preferences. Finally, while behavioral research has led some to argue for a more paternalistic approach to policy analysis, an alternative is to continue to focus on describing the preferences of those affected by the policy options while working to ensure that these preferences are based on knowledge and careful reflection. Benefit-cost analysis can be best viewed as a pragmatic framework for collecting, organizing, and evaluating relevant information.

1. INTRODUCTION

Policymakers face difficult choices in determining how to best allocate scarce resources across social programs and other desirable goods and services. Benefit-cost analysis provides useful information for these decisions, by indicating the extent to which the values that individuals place on program outcomes are likely to exceed program costs. Determining these values has always been challenging, however. Most social programs lead at least in part to outcomes for which no market value exists, such as improved health and longevity or environmental quality. Instead, these values must be estimated from market behavior for related goods or by asking individuals about their willingness to pay. Recent research in behavioral economics adds to the complexity of this task, documenting ways in which individuals at times appear to act irrationally or contrary to their own interests.¹

This paper supports the development of principles and standards for benefit-cost analysis of social programs, focusing on the implications of behavioral economics. We review traditional practices, discuss findings from behavioral research, and recommend ways in which these findings might be integrated into benefit-cost analyses. We take the perspective of an analyst who has been asked to evaluate a given set of policy options and is curious about how behavioral economics might influence the assessment. We concentrate largely on empirical research results and their practical application, while recognizing that theory provides a useful, simplified model of reality that can help promote rigorous thinking about these issues.

Behavioral economics is a large and rapidly growing field, and is just beginning to move towards an unified theory that provides a cohesive alternative or supplement to the standard economic model.² Its status has several implications for our discussion. First, the dividing line between behavioral economics and conventional neoclassical economics is often murky. Both address behavior, and many findings identified as behavioral economics can be accommodated within the standard model. Second, while researchers have found an increasing number of behavioral deviations from neoclassical economic assumptions, it is not yet clear how significant many of these deviations are in terms of magnitude or pervasiveness. Much of the research involves small-scale laboratory experiments and needs to be supplemented by additional fieldwork to explore the real-world importance of the findings. Third, these deviations are often dependent on the context, and more work is needed to determine whether the deviations found in the contexts frequently studied (e.g., financial decisions) are equally important in the context of the policy decisions we consider. Finally, the rapid growth in related research means that it is not possible to be comprehensive. We select key topics for detailed consideration rather than attempting to review the entire literature. Because we are drafting this paper as behavioral economics evolves in significant ways, we often raise questions or pose alternatives without attempting to resolve the underlying issues, describing concerns that researchers and analysts may wish to consider.

Below, we briefly summarize the distinction between behavioral and neoclassical economics, then introduce the features of the benefit-cost analysis framework that are the focus of this paper. The remaining sections then discuss each topic in turn.

¹ Examples of seminal work and comprehensive reviews include Thaler (1992), Rabin (1998), Kahneman and Tversky (2000), Camerer, Loewenstein, and Rabin (2004), Fudenberg (2006), and DellaVigna (2009).

² Examples of alternative framing include: Sugden (2004, 2005a, 2009), Bernheim and Rangel (2007, 2009), Green and Hojman (2009), and Smith and Moore (2010).

1.1 Behavioral vs. Neoclassical Economics

“Behavioral” economics is somewhat of a misnomer, because all economics is concerned with how people behave in economic contexts. Behavioral economists often distinguish their work by noting that, in its simplest form, the standard economic model generally assumes that people behave self-interestedly and rationally (as “econs” or “homo economicus”), while they consider how human behavior may deviate from this model. In determining how to best conduct benefit-cost analysis, the issue becomes identifying where the findings from behavioral economics might lead to analytic approaches that differ from typical practices.

Thaler and Mullainathan (2008) argue that the standard economic model assumes three unrealistic traits: unbounded rationality, unbounded willpower, and unbounded selfishness. Bounded rationality recognizes that humans have limited capacity to process information, so often do not solve problems optimally. Instead, they may at times come to conclusions based on heuristics or simple decision rules.

Bounded willpower reflects humans’ incomplete self-control. We may engage in unhealthy behavior (such as eating or drinking too much, saving too little, or smoking) while at the same time recognizing that such behavior is damaging. Bounded selfishness refers to the fact that we may act selflessly. While altruistic behavior is recognized in traditional economics, neoclassical theory stresses self-interest as the primary motivator. Behavioral economists have documented ways in which individuals may instead act out of concern for others or based on conceptions of fairness, reflecting vengeful as well as altruistic motives.

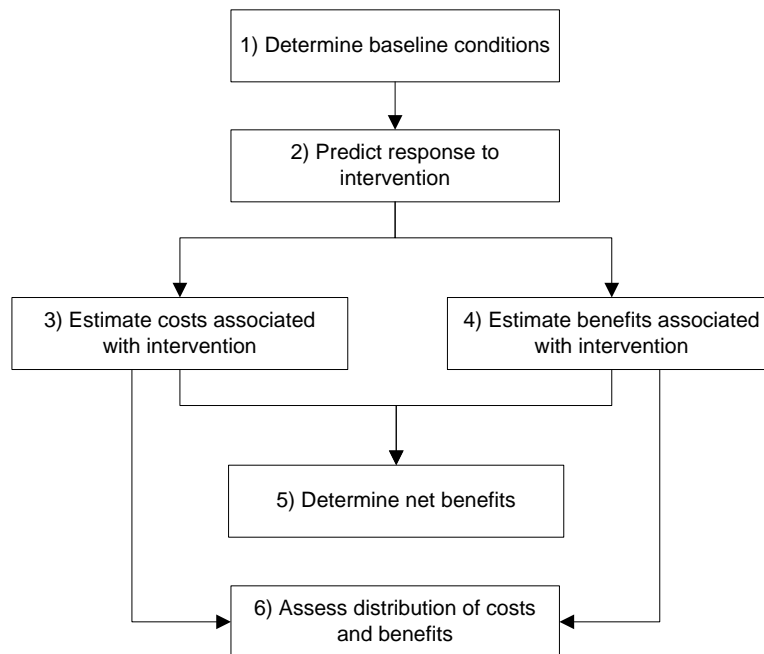
In this paper, we are primarily concerned with the use of behavioral data to determine preferences for, or the value of, the outcomes of social programs. Bernheim and Rangel (2007) frame the issues in a way that is particularly useful in this context. They note that “[p]ublic economics has positive and normative objectives; it aims both to describe the effects of public policies and to evaluate them. This agenda requires us to formulate models of human decision-making with two components – one describing choices, and the other describing well-being. Using the first component, we can forecast the effects of policy reforms on individuals’ actions, as well as on prices and allocations. Using the second component, we can determine whether these changes benefit consumers or harm them” (p. 7). This distinction between behavior and welfare is not necessary under the traditional economic model because that model assumes that individuals choose what they want; i.e., that their preferences are revealed through their behavior.

While there is no consensus on how to evaluate welfare given this distinction, Bernheim and Rangel describe two schools of thought. One is to continue to rely on revealed preferences (integrating behavior and welfare) but to expand how we think about preferences to include deviations from the standard model. The second is to modify, or even reject, the reliance on revealed preferences, suggesting that preferences should be used to measure welfare only in those cases where they appear consistent with the individual’s self-interest. While potentially protecting against errors, this more paternalistic approach runs into the possibility of abuse if individual preferences can be overridden without adequate, evidence-based justification. This tension between unquestioning acceptance of individual choices and acceptance of only those that are judged to be rational and welfare-enhancing is at the heart of many of the implications of behavioral economics for the conduct of benefit-cost analysis.

1.2 Behavior and Benefit-Cost Analysis

Implementing benefit-cost analysis involves several iterative and intertwined steps, that involve characterizing the affected universe with and without the policy intervention and assessing the social costs and benefits of each option. We provide a simplified overview of this process in Figure 1.³

Figure 1. Simplified Overview of Benefit-Cost Analysis



Typically, costs are defined as the opportunity costs of the real resources expended to develop, implement, and operate a program or to comply with regulatory or other requirements, including any market impacts. Benefits typically include the monetary value of the outcomes that are the goal of the policy: improved education, increased safety, greater employment, enhanced housing, and so forth. Ideally, any significant side effects (cost-savings or ancillary benefits) are included, and the implications of nonquantified effects and uncertainty are carefully assessed. While these analyses provide important and useful information, policy decisions are rarely, if ever, based solely on their results. Decisionmakers often seek additional information that cannot be easily captured in an economic analysis.

Behavior, and the implications of behavioral economics research, permeate each step of a benefit-cost analysis, and also influence how policy decisions are made and how the public perceives the impacts. Some analytic steps directly involve predicting future behavior, while others use behavior more indirectly to value nonmarket outcomes. We focus on largely on the latter issues in this paper, because they raise more difficult issues for the analyst.

In particular, we do not discuss the implications of behavioral economics for determining current and potential future baseline conditions in the absence of intervention, nor for estimating

³ See Robinson (2004) and Robinson (2008a) for more information on current practices for regulatory analysis; practices will vary in other policy contexts.

responses to different policies. For example, behavioral economics is helpful in predicting how individuals are likely to respond to information provision (e.g., on the caloric content of food or on energy efficiency), including information on the trade-off between short-term costs and longer-term savings.⁴ A well-conducted analysis attempts to be as realistic as possible; it seems self-evident that analysts should use whatever information appears likely to improve their predictions, regardless of whether the behavior appears rational or welfare-enhancing. The implications of behavioral economics for these types of predictions are diverse, vary significantly across different contexts, and are addressed in a large and rapidly growing literature.⁵

We also do not discuss the role of behavioral economics in estimating social costs. The focus on real resource expenditures means that market data often can be used to estimate monetary values, and such data will already reflect the effects of behavioral influences. In the case of a small program (such as a local addiction treatment center), the analyst is likely to be concerned primarily with direct expenditures on wages and benefits, space rental, equipment, supplies, and so forth, which can be directly estimated from market prices or technical (engineering) information. Implementation of the program is not likely to significantly change these prices. However, the cost analysis will be closely linked to predicting responses, as discussed above. Individual decisions (and any associated behavioral anomalies) may determine program size (e.g., the number of addicts served and the efficacy of the treatment), in turn affecting those costs that vary with participation rates.⁶

When a program is large enough to noticeably affect market conditions, accounting for behavioral factors may become more important. To predict the effects of the program on supply and/or demand conditions in particular sectors or throughout the economy, analysts may be able to rely on market data that already reflects any associated behavioral anomalies (e.g., observed supply or demand elasticities). In other cases, analysts may need to predict impacts outside the range of the available data, taking behavioral factors into account.

When market outcomes are affected by behavioral factors, the estimated resource costs may differ from those that would be estimated if the markets functioned in accordance with the standard model. As an example, if consumers undervalue future energy savings from adopting compact fluorescent light bulbs or more fuel-efficient motor vehicles (because of hyperbolic discounting, as discussed in Section 3), the market prices, and the estimated costs of requiring, these more efficient products will be higher than if consumers evaluated future cost savings using exponential discounting.⁷

Once we begin to consider the monetary valuation of nonmarket outcomes, the potential distinction between rational and irrational choices, and choices and welfare, increases in importance. Social programs are generally targeted on improving individual and societal welfare. If individuals' behavior suggests preferences that appear irrational, unstable, or contrary to their

⁴ In its simplest form, neoclassical theory assumes perfect information. However, an entire field has developed within the standard model around the economics of information, including topics such as decision making under uncertainty, insurance, the optimal acquisition of costly information, and the effects of information asymmetries (e.g., adverse selection and moral hazard).

⁵ See, for example, Diamond and Vartiainen (2007), Thaler and Sunstein (2008), and Congdon, Kling, and Mullainathan (2011).

⁶ See, for example, Bernheim and Rangel (2007) for discussion of the implications of behavioral research for addiction programs.

⁷ See, for example, Allcott and Mullainathan (2010) for discussion of the implications of behavioral research for energy policy.

self-interest, we face difficult decisions. We could use the resulting values (or a range of values) in our analysis regardless, we could explore the effects of providing more education or experience, or we could substitute expert judgment. Each of these options raises thorny practical and philosophical issues.

In this paper, we take the perspective that analysts should avoid making judgments about whether values are “rational” or “irrational,” but should make every effort to ensure that studies are designed to elicit well-informed, thoughtful preferences. We use the terms “mistakes” or “errors” to reflect choices that diverge from how an individual would define his or her own preferences given perfect information, full reflection, and self-control, absent the biases that may result from cognitive or emotional challenges. While such perfection in decisionmaking may be impossible in reality, our hope is to at least attempt to avoid paternalistic views of what individuals “should” prefer, deferring instead to the preferences individuals express when provided opportunities for contemplation and learning. In other words, we maintain the traditional reliance on individuals’ own definition of their welfare, while recognizing that education and experience may be needed to aid them in developing this definition.

Data limitations play an important role in benefit-cost analyses, particularly in valuing nonmarket effects. The research base rarely includes studies that address an outcome that is identical in all respects to the outcome of a particular social policy. Time and resource constraints mean that analysts are generally unable to conduct new primary valuation research. Instead, they follow the benefit-transfer framework, which involves taking values developed in one context (a primary research study or group of studies) and applying them in a somewhat dissimilar context (the policy analysis) based on careful review of the literature. While values may be adjusted to reflect differences in the characteristics of the outcomes or the affected population, the available research is often insufficient to support quantitative adjustment. These differences are frequently addressed in more qualitative terms. Analysts can explore the resulting uncertainty using sensitivity or probabilistic analysis, or by examining breakeven values; i.e., the values at which the benefits of a policy no longer exceed its costs, or at which the ranking of the policy options changes. Because using multiple values and/or discussing these issues qualitatively can complicate presentation of the results, analysts need to carefully summarize the implications for busy decisionmakers.

In the sections that follow, we begin by discussing several issues related to the implications of behavioral research for defining what it is that we are trying to measure, regardless of whether we are conducting primary research or reviewing studies for benefit transfer. In Section 2, we focus specifically on benefits, considering how behavioral research might affect the attributes included in nonmarket valuation studies. In Section 3, we turn to the discounting of future costs and benefits to reflect their timing. Section 4 discusses the distinction between private and social preferences. In Section 5, we turn to issues related to the quality of valuation studies, considering how behavioral research might affect their design and implementation.

Section 6 concludes by looking at the more general implications of behavioral economics. Practitioners often advocate benefit-cost analysis as a positive, descriptive exercise, that investigates what individuals would prefer (given the current distribution of income and other baseline conditions) and summarizes the results. Substituting expert judgment for observed choices is a slippery slope: there is often no clear dividing line between irrational and rational decisions or stable and unstable preferences. Behavior that appears irrational or unstable on the

surface may in fact reflect an underlying rationality; the problem may be that the investigator simply has not discovered or does not understand the rationale.⁸

Benefit-cost analysis can be viewed as a normative exercise even within this traditional framework. For example, Adler and Posner (2006) describe BCA as “applied moral philosophy” (p. 4) and Just et al. (2004) argue that welfare economics is normative because it “is concerned with what ‘ought’ to be” (p. 3). Regardless, substituting experts’ or others’ beliefs about what individual preferences “should” be firmly moves benefit-cost analysis into the normative realm (see Hammitt 2009). However, as Smith (2007) notes, experts make mistakes too, and there is no consensus on whose expert judgment should prevail. As discussed in Section 6, we advocate a pragmatic view of these analyses as attempting to reflect the thoughtful, well-informed preferences of those affected to the greatest extent possible.

2. VALUING PSYCHOLOGICAL ATTRIBUTES

The value of the benefits of social programs – such as those targeted on reducing crime, increasing education, improving housing, or decreasing environmental, health, or safety risks – frequently cannot be fully captured by directly referencing market behavior. Instead, these values are estimated through revealed- and stated-preference research. Revealed-preference studies use data from market transactions or observed behavior to estimate the value of related nonmarket goods, while stated-preference studies involve asking respondents how they would behave in a hypothetical market. Regardless of whether they are conducting new primary research or transferring values from existing studies, analysts need to start with a clear definition of the key attributes of the outcome of interest. In this section, we discuss how behavioral research might affect the attributes that analysts consider important.

Behavioral economics highlights ways in which psychological responses may lead to values that appear irrational or inconsistent with the standard economic model. The seminal work on these anomalies is encapsulated by Prospect Theory (proposed in Kahneman and Tversky 1979), which suggests that preferences depend on: (1) the reference point from which they are measured; (2) whether the change is a loss or a gain (with losses valued more than gains); (3) the distance from the reference point (with changes valued less if farther away); and (4) whether the probability is large or small (with small probabilities overweighted).⁹

The importance of psychological concerns also has been long recognized in traditional research. For example, numerous studies have explored how risk perception affects the valuation of mortality risks, leading individuals to value risks of the same outcome and magnitude differently depending on their cause. Much of this work builds on research summarized in Slovic (1987), which suggests that individuals are more likely to want to see a risk reduced if it is more dreaded (i.e., perceived as more uncontrollable, catastrophic, likely to be fatal, inequitable, risky to future generations, difficult to reduce, risk increasing (rather than decreasing), and/or involuntary). Individuals also have a greater desire for addressing risks that are unknown or unfamiliar (i.e., that are unobservable, unknown to those exposed, new, unknown to science, or have delayed effects).

⁸ Smith and Moore (2010) note that: “[d]ecisions that appear incoherent or contradictory may simply reflect the analysts’ failure to fully specify the constraints to choice.” (p. 226) These constraints include factors such cognitive capacity, the energy available for decisionmaking, and physical dexterity, as well as budget constraints.

⁹ More generally, related scholarship indicates that changes in probability near zero or one are overweighted and changes in intermediate probabilities are underweighted.

The effects of psychological responses on valuation, and the types of nonmarket outcomes associated with social programs, are too diverse for us to be comprehensive here.¹⁰ Instead, we begin by addressing the distinction between willingness to pay (WTP) and willingness to accept (WTA) compensation. This distinction is fundamental to valuation in almost every context and has been identified by several scholars as an area where behavioral economics has significant implications. We then consider the valuation of mortality risk reductions in light of behavioral findings related to psychological attributes and the interpretation of probabilities. We use mortality risk reductions as an example both because they are an important outcome of many social programs and because their value has been relatively well-studied.

2.1 Willingness to Pay vs. Willingness to Accept Compensation

WTP and WTA can be used to value beneficial or harmful changes. For a beneficial outcome, WTP represents the maximum amount of money an individual would be willing to give up in exchange for the amenity, while WTA represents the minimum amount he would need to be paid to forego, rather than gain, the amenity. For a harmful outcome, WTP is the maximum an individual would pay to avoid the harm, and WTA is the minimum he would require to accept the harm. Under conventional assumptions, Willig (1976) demonstrates that these values should be similar as long as income effects are negligible; i.e., if purchases of the good represent a small proportion of income and if changes in income do not lead to large changes in demand. Willig's analysis pertains to private goods, where the individual chooses the quantity to purchase. For public goods, where the individual cannot choose the quantity, Hanemann (1991) finds that WTP and WTA may diverge significantly when there are no private goods that are good substitutes for the public good.

A number of empirical studies have found substantial differences between WTP and WTA (see, for example, Horowitz and McConnell 2002).¹¹ Behavioral economists argue that these differences cannot be fully explained by income and substitution effects. Kahneman, Knetsch, and Thaler (1991) and several subsequent studies highlight the implications of Prospect Theory for this divergence, identifying it as a major concern for benefit-cost analysis.

In particular, Knetsch (2005, 2010) notes that behavioral studies indicate that WTP and WTA differ due to the combined effects of two factors. First, values depend on whether individuals view the reference state as their present status or as their status after the change. Second, they will value the change more highly if it is viewed as a loss from this reference state rather than as a gain. This means that, if the reference state is one's current status, WTP for a gain will be smaller than WTA for a loss of the same magnitude. If the reference point is instead one's future status, Knetsch argues that WTP and WTA will be between the values found if current status is used as the reference state.

¹⁰ We do not discuss research on happiness (or life satisfaction) in detail. Several scholars have suggested that these approaches should replace the use of WTP or WTA measures in BCA (see, for example, Kahneman and Sugden 2005 and Layard 2010 as well as the August 2008 special issue of the *Journal of Public Economics*). However, more work is needed to improve related research methods, collect empirical data, and assess the implications for policy analysis. For a thoughtful critique of related issues, see Smith (2008).

¹¹ This issue arises primarily in stated-preference studies. As discussed in Smith and Moore (2010), hedonic models (such as those examining wage-risk trade-offs to estimate the value of mortality risks) result in point estimates of marginal WTP based on equilibrium conditions. Information on discrete changes is needed to detect disparities between WTP and WTA.

This divergence raises the question of which measure is most appropriate for policy analysis. Traditionally, many economists argue that, because WTP and WTA differ in their starting points, the choice of measure should be based on property rights.¹² If the property right is associated with the status quo, then WTP to obtain an improvement is the correct measure. If the property right is instead associated with the change, then WTA to forgo the improvement is the correct measure. However, particularly for nonmarket goods, these property rights are often not well-defined. In addition, the legal definitions of rights may not correspond with how individuals identify the starting point when assessing changes in their own welfare. In particular, they may view the status quo as the more intuitive basis for valuing the changes associated with implementation of new social programs.

Once a reference state is established, differing values could be applied depending on whether a policy results in a gain or a loss from that state. However, as discussed in Guria et al. (2005) and elsewhere, this approach could lead to recommendations that vary depending on the perspective: an *ex ante* evaluation of a proposed program could support a differing conclusion than an *ex post* evaluation of an existing program; a proposal to introduce a program could be evaluated differently than a proposal to abolish a program. When the individuals who benefit from a policy change are the ones to bear the costs, it seems illogical for the evaluation of the policy to depend on the reference point. In addition, if a program involves both losses and gains, the use of different values could lead to counterintuitive results. For example, if a policy reduces mortality risks for some and increases them for others, the use of differing values could lead to rejection (on efficiency grounds) of a proposal that saves more lives, or could favor a proposal that saves fewer lives but has higher net benefits, depending on how the risks are allocated across the different groups (Hammit and Treich 2007). The extent to which these sorts of problems arise will depend on the extent to which the values vary when viewed from these differing perspectives; for many outcomes these differences have not yet been well-studied.

Much of the work identifying the behavioral anomalies that may underlie the divergence between WTP and WTA measures has been conducted in a laboratory setting with students as subjects, so often lacks the types of feedback mechanisms associated with real world exchanges. In addition, laboratory experiments usually involve exchanges of relatively simple goods (money, coffee mugs) that lack the multiple attributes associated with the outcomes of social programs. For more complex (and often less familiar) nonmarket goods, differences between WTP and WTA measures may be driven at least in part by attributes of the outcomes themselves for which sufficient controls are not included in the data analysis and by challenges related to how the values are elicited or measured, as well as by loss aversion and other behavioral anomalies. Thus more study is needed to understand why WTP and WTA estimates diverge for particular types of outcomes and to estimate the size of the difference.

Policy analysts often focus on WTP estimates for more practical reasons, including skepticism about WTA estimates that appear implausibly large.¹³ The extent to which these differences stem from problems with the data and methods used for nonmarket valuation, rather than from the sorts of concerns noted above, remains unclear. One particular methodological

¹² More precisely, these measures are based on the concepts of compensating and equivalent variation (or compensating and equivalent surplus for public goods). The two measures differ in their starting points: for a beneficial outcome, compensating variation references the level of utility without the improvement, while equivalent variation references the level with the improvement. For more discussion of the relationship to property rights, see Freeman (2003).

¹³ This preference for WTP estimates due to perceived WTA measurement problems is reflected in current guidelines, such as OMB (2003).

challenge arises from the relationship of these measures to income. Because WTP cannot exceed an individual's ability to pay (e.g., income or wealth), stated-preference researchers can remind respondents to consider their budget constraints and can identify values that appear unrealistic given reported income levels. In contrast, WTA amounts are unconstrained, and may lead respondents to overstate what they would in fact accept when hypothetical surveys are used to elicit values. Large values also may be reported as protest bids when respondents do not accept the scenario presented by the researchers. Estimates of WTA that are consistent with an underlying utility function can be obtained instead using a preference-calibration approach (Smith et al. 2006).

Thus more research is needed to examine the extent to which WTP and WTA estimates are likely to diverge for the particular outcomes of concern in benefit-cost analyses of social programs. In the interim, analysts can test the sensitivity of their findings to variation in these values. For example, if the benefits analysis relies on WTP estimates but consideration of the reference state and loss aversion suggests that WTA may be more appropriate, then the analyst may wish to test the impact of larger values on the results.

2.2 Psychological Responses to Risk

Behavioral research has many other potentially significant implications for the attributes considered in nonmarket valuation, including how risks are perceived and valued. Some of these issues relate to the cognitive processing of risk information, including the misinterpretation of probabilities and the tendency to rely on simple heuristics or decision rules. In this section, we focus more directly on underlying preferences, while recognizing that it can be very difficult to distinguish “real” from “mistaken” choices or values. We illustrate these issues using research on the value of small mortality risk reductions.

Typically, the value for mortality risks is expressed as the value per statistical life (VSL). The VSL represents the value of small risk changes (e.g., of 1 in 10,000) in a defined time period, expressed as a “statistical” life for convenience – it is not the value of saving an individual's life with certainty.¹⁴ The VSL has been estimated in over 60 revealed-preference studies (Viscusi and Aldy 2003) and over 70 stated-preference studies (Lindhjem et al. 2010); the implications of these studies have been assessed in numerous literature reviews and meta-analyses as well as in guidance for regulatory analyses.¹⁵

There is substantial evidence that both personal characteristics and risk characteristics affect the VSL; its variation is not limited to the potential differences between WTP and WTA discussed above. The influence of personal characteristics (such as income or age) can be described at least in part by standard economic theory, such as the lifecycle consumption model (see for example, Hammitt 2007, Hammitt and Robinson 2010). Risk characteristics include both physical attributes (such as whether the risk is latent or involves significant morbidity prior to death) and psychological attributes (such as whether the risk is perceived as voluntarily incurred or under an individual's control). Research on the effects of these perceptions suggests that individuals may value risks of the same expected magnitude (e.g., 1 in 10,000) and same outcome (e.g., immediate death) differently if they stem from causes that are viewed differently.

¹⁴ In other words, the VSL is equal to individual WTP for a small risk change in a defined time period, divided by the risk change; it is the local slope of an indifference curve between risk and wealth (see Hammitt 2000).

¹⁵ See, for example, the meta-analyses in Viscusi and Aldy (2003) and Lindhjem et al. (2010), the reviews in Robinson (2008b) and Robinson and Hammitt (2009), and the guidance in EPA (2000) and OMB (2003).

More work is needed to better understand how these perceptions affect valuation; however, the available empirical evidence suggests that risks viewed as less controllable, voluntary, or familiar, or as more feared or ambiguous, may be valued up to twice as much as other risks (Robinson et al. 2010).

This variation is not necessarily inconsistent with the standard economic model: changes in attributes are expected to lead to changes in value. However, some research findings illustrate the influence of the types of anomalies emphasized by behavioral economists. First, when faced with uncertain risk information, individuals respond differently than when faced with a point estimate equivalent to the expected value of the range, reflecting ambiguity aversion.¹⁶ For example, work by Viscusi et al. (1991), Shogren (2005), and Riddel and Shaw (2006) suggests that WTP for fatal or nonfatal risk reductions increases as risk ambiguity increases. Second, research shows that individuals tend to overweight small risks (particularly when they are viewed as fearsome), consistent with Prospect Theory (see, for example, the review in Johannson-Stenman 2008). Third, individuals are often insensitive to changes in small risks, reporting the same or similar values for risks that differ in magnitude (see, for example, Hammitt and Graham 1999, Corso et al. 2001, Robinson and Hammitt 2009). Although these findings may in part reflect misunderstanding or miscommunication of the risks being valued, they may also reflect stable or “real” underlying preferences.

These anomalies do not necessarily create problems for the analyst. Ideally, the values used in benefit-cost analysis would reflect all of the attributes of the risk, including the ambiguity and fear associated specifically with that risk (see, for example, Robinson et al. 2010). However, in some cases, this may mean that the analyst would need to test a wide range of values either because of deficiencies in the research base or because of uncertainty regarding the applicability of the results.

Insensitivity to risk changes may be particularly troubling. Economic theory suggests that WTP should increase almost proportionately to the size of the risk change, as long as the change is small, which means that the VSL would be nearly constant. If instead, for example, individuals report that they are willing to pay \$700 per year for both a 5 in 10,000 and a 1 in 10,000 annual risk reduction, this translates into a VSL ranging from \$1.4 million to \$7 million – a wide enough range to potentially affect whether a policy appears cost-beneficial. One possible explanation may be that individuals are indifferent between risk changes of these magnitudes. However, as discussed in Corso et al. (2001), it appears more likely that this result reflects misinterpretation of the probabilities, which can be reduced by using visual aids that more effectively communicate the size of the change. While many studies that use these aids find increased sensitivity to risk magnitude, some continue to find a degree of insensitivity (e.g., Alberini et al. 2004). However, it may be preferable to use this sensitivity as a criteria for evaluating the quality of the studies for benefit-transfer, rather than assuming that it reflects indifference between risk changes of these magnitudes. More generally, differences between the risks studied and the risks associated with various social policies will lead to uncertainty in the estimates, the implications of which will need to be addressed in the analysis.

Our focus on mortality risks skirts some difficult questions, which we believe cannot be answered definitively. For example, there is substantial evidence that individuals’ evaluations

¹⁶ In his seminal work on this topic, Ellsberg (1961) notes that ambiguity depends on the amount, type, reliability, and unanimity of information on probabilities and the resulting degree of confidence one has in the data. He indicates that many people prefer payments that they will receive with certainty to those that are uncertain, even when the uncertain outcome has a higher expected value.

differ before, during, and after an experience (e.g., Kahneman 2000a, 2000b). This means that, for example, the values for reducing the risk of a particular injury or illness may depend on whether the individual has experienced that health effect. One could argue it is the prospective, *ex ante* perspective of the inexperienced, healthy individual that matters in policy analysis, assuming that the goal of the policy is to help the currently healthy person avoid transitioning into the less healthy state.¹⁷ Alternatively, one could argue that the experienced, ill individual is better informed and hence his or her values should be used. A third choice could involve somehow integrating values across inexperienced and experienced health states over time. These sorts of debates again argue for using a range of values in benefit-cost analysis, to determine how the differences affect the results.

2.3 Conclusions and Implications

The first step in valuing any outcome involves defining it, including both its physical and psychological characteristics. Behavioral economics suggests a number of attributes that may be worth investigating when valuing the nonmarket benefits of social programs. Assuming that our goal is to provide information on the amount of money that affected individuals would be willing to trade for the outcome of concern, it seems appropriate to examine the effects of these attributes regardless of whether the results appear consistent with the standard economic model; the findings may provide important insights. While inconsistencies in policy recommendations could potentially arise when such context-dependent values are used in benefit-cost analysis, it is difficult to determine the importance of these inconsistencies without more information on the extent to which values might vary.

As we discuss in more detail later, addressing these issues in part requires ensuring that valuation studies are well-designed, to help separate values that reflect misinformation or misunderstanding from values that reflect more stable and carefully-considered preferences. However, we also need more research that tests the effects of these attributes on valuation, so that we are better able to sort out what is and is not important in different contexts, as well as determine the extent to which they are outweighed or counterbalanced by other considerations. For example, it is unclear whether differences between WTP and WTA are more or less important than whether a mortality risk is particularly feared or ambiguous. Regardless of whether we view these issues from a traditional or behavioral perspective, our understanding of how nonmarket values vary in different contexts is incomplete.

Two types of research would be helpful. First, because the effects of different attributes are interrelated, studies are needed that consider them jointly – holistically assessing the specific outcome of concern. Second, studies that consider the effects of varying attributes one-by-one are also useful, because they provide insights that can be used to make adjustments (or to calibrate results) when benefit transfers are conducted. These studies should investigate the

¹⁷ When assessing the cost-effectiveness of health policies and medical interventions, recommendations for best practices suggest that, while descriptions of health states should be based on information from patients or others familiar with the condition, preferences for different health states should be based on the community or societal perspective. One way of describing the desirability of this perspective “...is to imagine that we are looking at the world before we are born, or at least before we encounter any serious health problems, and to ask what kind of world we would like it to be. In that “ex ante” position we would not yet know what health problems we were destined to develop – only that there was some chance that we might develop any of them” (Gold et al. 2006, p. 7). Similar recommendations have been developed for cost-effectiveness analysis of environmental, health, and safety regulations (Institute of Medicine 2006).

extent of heterogeneity in the population as well, to determine whether changes in attributes are perceived and valued similarly by different individuals.

Such studies can take years to complete, and funding for nonmarket valuation work is very limited. In the absence of conclusive empirical research, analysts will need to carefully describe the potential implications of differences between the study and policy outcome when interpreting benefit-cost analysis results, including both factors that are, and are not, consistent with the standard economic model. We expect that analysts will continue to use the benefit-transfer framework to explore the effects of context differences both quantitatively and qualitatively; using sensitivity, breakeven, or probabilistic analysis where appropriate to test the implications of related uncertainties.

In conclusion, we recommend that analysts avoid making *a priori* judgments about whether values appear “rational” or “stable,” and instead consider the following.

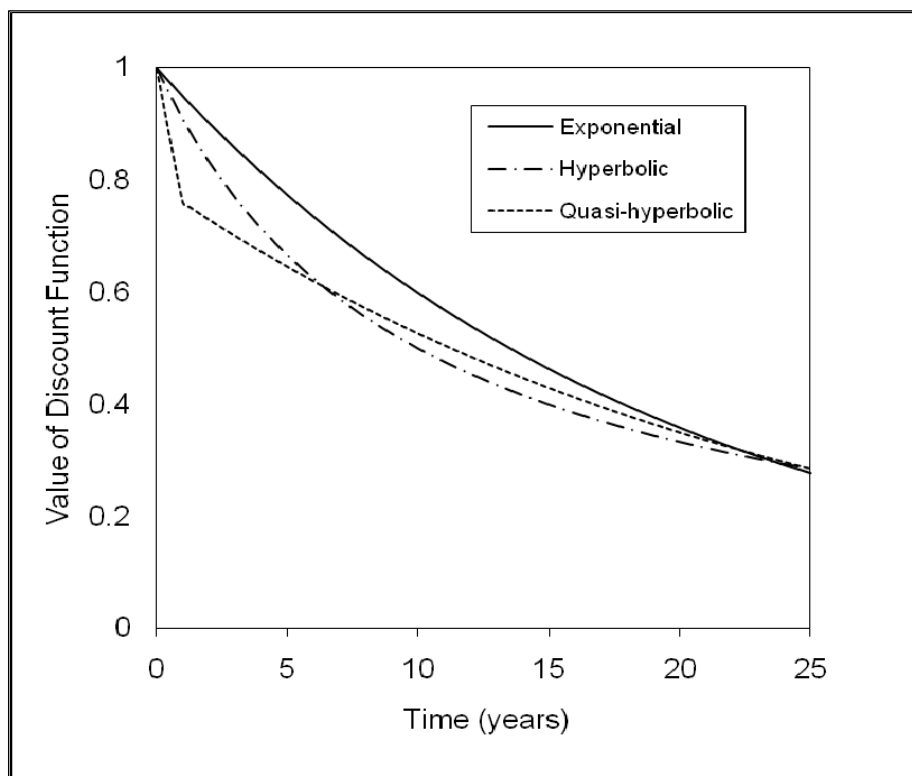
- 1) Studies should be designed to test the effects of the psychologically-salient attributes found in behavioral research on benefit values, as well as the effects of other physical and psychological attributes.
- 2) Such studies should consider both the holistic effect of the full range of attributes relevant to a particular context, and the effects of varying the attributes one-by-one to develop adjustments for transferring values to other contexts.
- 3) Values should not be rejected unless the study does not meet basic criteria for quality or adhere to generally-accepted principles for best practices. Rejected values, and the basis for rejecting them, should be clearly documented.
- 4) When values are uncertain, sensitivity, probabilistic, or breakeven analysis should be used to test the effects of this uncertainty on the results.
- 5) Where quantitative estimates are not available, the potential effects of both psychologically-salient and physical attributes should be discussed qualitatively.

3. ESTIMATING TIME PREFERENCES

Evaluating the benefits and costs of social programs often involves comparing impacts that occur in different time periods, using discounting to reflect time preferences. The main distinction between the traditional approach to discounting and the results of behavioral research is whether outcomes are discounted at a constant exponential rate or hyperbolically. This distinction reflects simplifying assumptions rather than theoretical differences. Exponential discounting can be traced to Samuelson’s (1937) work on discounted utility, which was intended to highlight certain theoretical relationships. Samuelson himself recognized that his assumptions were a simplification, neither necessarily predictive of actual behavior nor associated with a normative view of welfare. More recent work (Weitzman 2001) suggests that hyperbolic discounting is appropriate under the standard economic model when the discount rate is uncertain.¹⁸ Gollier and Zeckhauser (2005) also discuss theoretical reasons why aggregating individual discount rates will lead to a rate that decreases over time. Thus the distinction that we explore in this section relates more to practical application than to theory. Figure 2 provides an example of exponential, hyperbolic, and quasi-hyperbolic functions.

¹⁸ This results because the discount factor $[1/(1+r)^t]$ is a nonlinear function of the discount rate r over time t , so the expected value of the discount factor corresponds to smaller values of r as t increases.

Figure 2. Exponential, Hyperbolic, and Quasi-Hyperbolic Discounting Functions (based on Berns et al. (2007) Figure 1.)



Discounting is another area where behavioral economics as well as other research emphasizes the need to clearly define what it is that we are trying to measure. If discounting is intended to reflect real resource expenditures, then relying on market rates seems appropriate. These rates measure what an investment might otherwise earn on average; i.e., the opportunity costs of diverting funds to a particular social program. However, if market imperfections lead to significant divergence between these rates and aggregated or average individual preferences, then further examination of these preferences may be desirable.

One problem that arises in examining this possible divergence is the difficulty of disentangling the effect of timing from the effects of other characteristics of an outcome, such as the reference dependence and loss aversion discussed earlier. Frederick et al. (2002) note that, for a given time delay, the empirical evidence suggests that rates vary depending on the context: “(1) gains are discounted more than losses; (2) small amounts are discounted more than large amounts; (3) greater discounting is shown to avoid delay of a good than to expedite its receipt; (4) in choices over sequences of outcomes, improving sequences are often preferred to declining sequences though positive time preference dictates the opposite; and (5) in choices over sequences, violations of independence are pervasive, and people seem to prefer spreading consumption over time in a way that diminishing marginal utility alone cannot explain” (p. 362). As they note, these problems relate more to how the utility function is specified than to the discount rate. To avoid redundancy, below we focus on the effects of timing alone, while recognizing that the dividing line between pure time preferences and other types of preferences is somewhat murky.

In this discussion, we address discounting of the monetary value of costs and benefits, not the underlying physical impacts. To the extent that timing affects the value of these physical outcomes, its impact is best represented by using their time-specific monetary value rather than by adjusting the discount rate. For example, if latent health effects are valued differently than effects experienced immediately, this difference should be reflected in the unit values applied to the health effects at the time they manifest, which then can be discounted at the same rate as other monetary amounts included in the analysis.

The time period over which an analysis is conducted has important implications for discounting. We concentrate here on discounting annual quantities over an intra-generational time period, consistent with the likely time frame of interest for analyses of social programs. We do not discuss the additional complications that arise when assessing programs (such as climate change or nuclear-waste storage) where inter-generational impacts are of major importance. Below, we first briefly review the traditional exponential approach, then discuss the evidence from behavioral economics and describe the implications.

3.1 Exponential Discounting

Time can influence the value of costs and benefits in a variety of different ways. Outcomes further in the future may be more uncertain or risky; new opportunities, information, or technologies may arise while current options may disappear; and an individual's preferences may change as he or she ages. In theory, discounting as traditionally implemented should not reflect any of these factors (which may be addressed separately elsewhere in the analysis); it should simply reflect the "pure" effect of timing.

Typically, benefit-cost analysis is conducted with constant discount rates (using the same rate for both costs and benefits), although rates that change over time have been used in some cases for longer-term impacts.¹⁹ Exponential (i.e., constant-rate) discounting assumes that time preferences are constant over different periods.²⁰

For government analyses of social programs, rates are established by OMB in Circular A-94 (1992) and Circular A-4 (2003); the former focuses more on government programs while the latter focuses on economically-significant regulations.²¹ Both suggest the use of exponential discounting, using a 7 percent real annual rate to approximate the average marginal pretax rate of return on private investments.²² Circular A-4 requires that analysts also report the results using a 3 percent rate to reflect consumption time preferences (or the "social rate of time preference").²³ OMB derives the 3 percent rate from the pretax rate of return on long-term government debt to

¹⁹ For example, the United Kingdom's guidance for evaluation of public programs (HM Treasury 2003) suggests that declining discount rates be used for analyses that cover periods greater than 30 years. Wietzman (2001) also suggests declining rates for assessing climate change policies.

²⁰ This means that the present value (PV) of a benefit or cost (S) at time (t) using a discount rate (r) is calculated as $PV = S_t / (1+r)^t$.

²¹ Both Circulars note that the shadow price approach would be preferable but is not recommended due to difficulties in its implementation. In addition to discussing social programs, Circular A-94 also discusses rates to be used in analyses that reflect trade-offs within the Federal budget, which are based on Treasury borrowing rates and updated annually.

²² All discount rates are reported as real rates, net of inflation.

²³ Three percent is also recommended as the base rate for cost-effectiveness analysis of health and medical interventions, with an alternative of 5 percent (for comparability with previous studies) and a range from zero to 7 percent in sensitivity analysis (Gold et al. 1996).

approximate the interest paid on savings, assuming that the savings rate represents the average by which consumers discount future consumption.

The use of two alternative rates reflects uncertainty about whether these programs primarily affect the allocation of capital or private consumption. In theory, the rates would not diverge in perfectly competitive markets, but in actuality economic distortions such as taxes lead to differences. Given this and other sources of uncertainty, OMB also requires that agencies provide a schedule that shows how the undiscounted costs and benefits will be distributed over time and discusses the use of alternative rates in sensitivity analysis. Thus the traditional approach for intra-generational discounting relies on a number of simplifying assumptions, and includes recommendations for testing these assumptions by considering the impact of alternative discount rates.

3.2 Hyperbolic Discounting

Time preferences have been one of the most active and well-developed components of behavioral economics research. As summarized in Chabris et al. (2008), numerous studies conducted over many years have found higher discount rates in the near-term than over the longer term. As a result, behavioral economists have explored hyperbolic functions with discounted values that drop steeply in the immediate future and more gently over the longer run.²⁴ These functions are often described as “present-biased.”

The exponential function is often described as the only discounting approach that yields dynamically consistent decisions.²⁵ Under hyperbolic and other functions, the preferred choice can change solely because time passes. In other words, a project that appears desirable (e.g., the present value of its net benefits is positive) in time period 1 may appear undesirable (e.g., may have a negative net present value) in time period 2. As noted by Laibson (1997) “from today’s perspective, the discount rate between two far-off periods, t and $t + 1$, is the long-term low discount rate. However, from the time t perspective, the discount rate between t and $t + 1$ is the short-term high discount rate” (p. 445-446). This inconsistency is often referred to as a “preference reversal,” although it can perhaps be better characterized as a sort of tug-of-war between these two perspectives or “as a game between a sequence of short-run impulsive selves and a long-run patient self” (Fudenberg and Levine 2006, p. 1449).

Many of the studies that find declining rates were conducted in laboratory settings. In these experiments, a small number of subjects (often students) are offered choices between receiving differing amounts of money (or another reward, often small) sooner or later. Declining rates have also been found in some field studies that address real-world behavior. Frederick et al. (2002) list 42 studies conducted between 1978 and 2002, including 34 laboratory experiments and eight field studies. The number of studies is increasing rapidly. Laibson (2010) lists almost 20 additional studies completed since 2002, including several conducted under field rather than laboratory conditions.

The results are diverse. The studies reviewed by Frederick et al. imply annual discount rates ranging from zero to thousands of percent, based on data collected for timeframes as short as one day to as long as 25 years. Frederick et al. note that this lack of agreement reflects at least

²⁴ Alternatives to the hyperbolic discounting model have been proposed by some to explain these research results, including subadditive discounting (Read 2001) and similarity relations (Rubenstein 2003).

²⁵ If preferences are characterized by date rather than distance from the present, hyperbolic discounting can be dynamically consistent (Harvey 1994).

in part the difficulties inherent in isolating pure time preferences from other (theoretically distinct) considerations.

While many studies demonstrate hyperbolic discounting, this pattern appears to apply primarily to relatively short near-term time periods. Frederick et al. report that if they exclude studies with time horizons less than one year, discount rates no longer decline over time, clustering around an average annual discount factor of 0.8. This factor implies an annual discount rate of 25 percent, however, well above market rates. Frederick et al. note that the high rates may be due to the effects of several confounding factors that tend to bias the results of these studies upwards.

Laibson et al. (2007) explore time preferences over the life-cycle more systematically, using a structural model with data on age-specific income, credit-card borrowing, marginal propensity to consume, retirement-wealth accumulations, household characteristics, mortality rates, and other factors. If they restrict their model to a single (exponential) function, they find an annual discount rate of about 16.7 percent. However, their modeling rejects this single rate hypothesis. Allowing a quasi-hyperbolic function, they find a short-term annualized discount rate of 39.5 percent and a long-term annualized rate of 4.3 percent.

Behavioral economists have explored a number of motivations behind these patterns. The high near-term rates are often described as resulting from imperfect self-control, which leads individuals to seek immediate gratification even if it diverges from their own longer-term preferences.²⁶ Examples of this behavior, such as eating dessert despite wanting to lose weight, are abundant. Individuals vary in the extent to which they correctly predict these types of problems. Researchers identify “sophisticated” consumers as those who are more fully aware of the potential for future self-control issues, while “naive” consumers are more likely to incorrectly predict their future behavior. Sophisticated consumers may implement self-control measures or use commitment devices (such as avoiding temptation or establishing penalties) to reinforce their self control.

3.3 Conclusions and Implications

Although most economists agree that discounting is needed, there is some disagreement on the appropriate rate for social programs even within the traditional approach to benefit-cost analysis. Conceptually, under this framework, the rate should reflect the opportunity cost of investing in the intervention of concern; i.e., the best alternative use of the funds. However, it is often unclear what types of investment or consumption are affected when particular social programs are implemented, leading to the widespread use of generic defaults for discount rates, ideally accompanied by sensitivity analysis.

Traditionally, benefit-cost analysis involves discounting the monetary values of future impacts at a constant exponential rate, with real rates often in the range of 3 percent to 7 percent annually. Behavioral economics suggests that discounting instead follows a hyperbolic pattern, with varying rates. These findings pose three challenges for analysts: (1) they suggest significant differences between short- and long-term discount rates; (2) they suggest that discounting should be hyperbolic rather exponential; (3) they suggest that individual behavior indicates time preferences that may differ from market rates.

²⁶ As discussed in Berns et al. (2007) and Chabris et al. (2008), behavioral economists are also exploring the potential neurological basis of the tendency to discount hyperbolically.

In the context of social programs, rates reflecting a longer-term view appear more appropriate than the steep near-term rates found in behavioral research (regardless of whether they are exponential or hyperbolic) for two reasons. First, these programs are generally intended to operate over a several year period, which means that the planning horizon is consistent with longer-term rates. Second, these programs are focused on providing lasting (rather than temporary) improvements in welfare. To the extent that short-term rates reflect impulsive behavior and self-control problems (rather than more patient and thoughtful consideration), they are inconsistent with this goal. To reflect true improvements in social welfare, it appears desirable to focus on the time preferences that result from more careful assessment of long-range well-being.²⁷

As noted earlier, one source of tension between the rates traditionally used and the results from behavioral and other research on preferences relates to clearly defining what it is that we are trying to measure. Market rates indicate real resource costs, while behavioral research reflects preferences found in laboratory experiments or in field studies limited to particular market goods. (Stated-preference studies have also been used to explore discount rates.) These studies generally consider decisions made by individuals, while aggregate or average rates are needed for assessing social costs and benefits.²⁸ These perspectives can be integrated if market rates are viewed as a reasonable approximation of preferences; i.e., if we assume that these rates reflect the aggregate of individual choices; that these choices reflect individual preferences; and that these preferences correspond to individual welfare.

However, we believe that it is premature to recommend specific changes in the rates or functional forms used for discounting in benefit-cost analysis. While the reviews cited earlier compile the results of numerous studies, they do not evaluate these studies against criteria for quality or for applicability to social programs. Systematic review is needed both to determine the extent to which discount rates might decrease over time (and whether they may be close-to-constant over the long run as suggested by Frederick et al.) and the extent to which they might differ from the 3 percent to 7 percent range now often used to represent market rates.

In the meantime, what advice can we provide our prototypical policy analyst?

- 1) Analysts should provide a schedule of undiscounted costs and benefits, in the form of a bar chart (or other graphic) or a table, so that the decisionmaker can inspect this pattern. Simply summarizing costs and benefits as a net present value does not provide adequate information given related uncertainties.
- 2) In choosing a discount rate or set of rates, as well as a functional form (exponential or hyperbolic) analysts must be clear about what, conceptually, they intend to represent for review by decisionmakers. The quality of the underlying data and related assumptions should also be described.

²⁷ This does not mean that short-term rates should be completely ignored: they play an important role in predicting behavior separate from the discounting of costs and benefits, as noted in Section 1.2.

²⁸ We thank Kerry Smith for highlighting this point.

- 3) Analysts should consider the impact of alternative rates on their results, in particular emphasizing the discount rates at which benefits do, and do not, exceed costs, and/or at which the relative ranking of different policy options change.²⁹ Comparing these switch points to rates reported in the empirical literature will provide insights into the implications of related uncertainties for decisionmaking.

4. SEPARATING PRIVATE FROM SOCIAL PREFERENCES

While the outcomes of social programs vary in the extent to which they meet the economic definition of a public good (i.e., are non-rival and non-excludable), these programs differ from markets for private goods and services along several dimensions that are affected by how individuals view others when determining their own preferences. Although in its simplest form, the standard economic model assumes that individuals are primarily self-interested, other-regarding preferences have been long-recognized within the traditional framework. As examples, we briefly discuss the treatment of altruistic motives, and note issues related to valuing a program rather than a private good and to assessing equity along with economic efficiency.

Behavioral studies reinforce the importance of these types of considerations, suggesting that individuals often act selflessly. They also identify other types of interpersonal considerations that may affect preferences, including vengeful as well as welfare-enhancing motives. Behavioral research on these issues is less well-developed than research in the areas discussed previously, but provides potentially useful insights.

4.1 Social Preferences within the Traditional Framework

While self-interested preferences are a standard assumption of the most simplified version of the traditional economic model, scholars have long recognized that individuals also care about others' welfare. Incorporating other-regarding preferences into benefit-cost analysis is difficult due to data limitations as well as conceptual challenges; however, related issues have been explored extensively within the traditional framework.

In particular, the appropriate treatment of altruism has received substantial attention. Economic theory distinguishes between two forms. Pure, or non-paternalistic, altruism means that I respect the preferences of others: I value the benefits they receive and the costs they incur exactly the same as they do. In contrast, paternalistic altruism involves ignoring other's preferences to some extent. If those affected by the costs and benefits of a policy are pure altruists, counting both their altruistic and private values in benefit-cost analysis simply scales the costs and benefits upwards without affecting the overall analytic conclusions, so that ignoring or including altruistic values leads to identical results (Jones-Lee 1991, Bergstrom 2006). If their altruism instead varies depending on the outcome (e.g., is greater for health than for other aspects of well-being) or depending on the individuals affected (e.g., is greater for poor individuals than for the wealthy), then it can affect the sign of net benefits.³⁰ In particular, there is some evidence

²⁹ In some cases, there may be more than one switch point (i.e., internal rate of return), depending on the distribution of costs and benefits over time.

³⁰ We follow the usual assumption that the values to be counted in the benefit-cost analyses are those held by the individuals who bear its costs and/or receive its benefits. Decisionmakers may, of course, have altruistic motives or other interests that differ from those of the affected individuals.

that health and longevity are viewed more paternalistically than most other goods (e.g., Jacobsson et al. 2007).

Paternalistic motives also appear more plausible when inequalities exist: I may wish to help those in impaired health by funding a program that exceeds their willingness to pay for it (and decreasing my own consumption).^{31,32} Alternatively, selective paternalistic altruism could involve caring about the costs incurred by others rather than the risks they incur. For example, I may be willing to pay \$75 for my own risk reductions plus \$50 to offset the costs to poorer individuals, allowing them to use their funds for other purposes. In this case, the subsidy is a means to re-distribute the costs of the program without changing its total costs nor the value of the benefits that accrue. In these cases, the value may not be truly altruistic (i.e., reflect how I value these particular costs or benefits for others); it may instead reflect how I value the distribution of impacts. If this is the case, analysts will need to carefully distinguish general preferences for redistribution from preferences related to the outcome of concern.

In more complex, real-world situations, determining whether and how the values of various outcomes are differentially affected by paternalistic altruism is a difficult task. For example, Viscusi, Magat, and Forrest (1988) demonstrate that individuals are willing to pay additional amounts to reduce the risks of insecticide poisoning to others. However, they note that extrapolating their results to other contexts is not appropriate because the analysis was exploratory, the contributions were hypothetical, and the values are likely to vary across contexts. Due to these sorts of challenges, altruistic values are typically not quantified when valuing nonmarket goods.³³

Another issue that is receiving increased attention is the difference between individual WTP for a “unit” of a particular good or service for private use versus WTP for a program that benefits the entire community of which the individual is a part. With purely self-regarding preferences, these values would be equal, if the outcomes are otherwise identical. However, while several stated-preference studies consider public programs rather than, or in addition to, private goods, they have not yet been reviewed systematically to provide insights into the extent to which the values vary. Because it is difficult to design a believable scenario that is identical in all respects other than whether the good is provided privately or to a community, determining the extent to which such values diverge is challenging. For example, to be otherwise comparable, a risk-reducing program would need to deliver the same expected risk reduction as the private option, and to be described in a way that avoids strategic bias (e.g., free-riding, protest bids, or overstatement due to the hypothetical nature of the payment).³⁴

Finally, as traditionally practiced, benefit-cost analysis focuses on economic efficiency, supplemented at times by separate assessment of distributional equity. Advocates of this

³¹ We default to the usual convention of referring to WTP rather than both WTP and WTA for ease of presentation. See Section 2 for discussion of these measures.

³² It is worth considering, in this case, whether those receiving the excess health benefits (e.g., poorer members of the population) would benefit more if the altruists were willing to provide an additional \$50 in cash rather than spending it on risk reduction.

³³ One area where paternalistic values have been increasingly recognized is in valuing health risks to children; see, for example, Dockins et al. (2002).

³⁴ A related issue is whether adult WTP reflects only personal values or also incorporates values associated with the well-being of other household members regardless of their age. See Munro (2009) for discussion of related issues in the context of stated-preference studies. Similar issues may arise in revealed-preference studies. For example, researchers generally assume that VSL estimates from wage-risk studies represent individual values, but decisions to trade income for risk changes may represent a household decision and include consideration of other’s well-being.

separation between equity and efficiency suggest that programs designed primarily to achieve outcomes other than income redistribution should focus on maximizing net social welfare; taxes and similar strategies can then be used to more effectively achieve distributional goals.³⁵ However, if WTP for a particular outcome in part reflects altruistic motives (intentionally or inadvertently) that result from concerns about disparities, and this WTP is used in the analysis of efficiency, the results of the benefit-cost analysis will, at least in part, also reflect equity concerns.

Social programs can exacerbate or ameliorate current inequalities both directly (depending on who receives the benefits and who bears the costs) and indirectly (e.g., because improved health can lead to greater productivity and income). Proposals to weight costs and benefits to reflect equity concerns have generally not been accepted in the U.S. due to the lack of agreement on the appropriate weights as well as concerns about transparency.³⁶ However, several scholars have proposed approaches for more rigorously assessing equity (e.g., Sunstein 2007, Adler 2008, Graham 2008, Farrow 2009, Loomis 2009, Zerbe 2009, Johansson-Stenman and Konow 2010). The feasibility and usefulness of these approaches have not yet been carefully tested. However, they suggest that equity effects could be addressed in more detail within the traditional framework, in part by taking advantage of the existing research on social welfare functions, optimal taxation, the marginal utility of income, and related topics.

4.2 Behavioral Research on Social Preferences

The strong interest in altruism, in valuing public programs as well as private goods, and in developing proposals to better incorporate equity into benefit-cost analysis, are consistent to a large extent with the emerging behavioral research on social preferences. Both laboratory and field studies suggest that individuals frequently consider others' welfare in decisionmaking. While this behavior is often altruistic, it may also involve acting reciprocally to reward or punish others. In addition, individuals may use others' status as the baseline against which to assess their own well-being.

Charness and Rabin (2002) discuss three types of models that could explain behaviors found in laboratory settings: “ ‘[d]ifference aversion models’ assume that players are motivated to reduce differences between theirs and others’ payoffs; ‘social-welfare models’ assume that people like to increase social surplus, caring especially about helping those (themselves or others) with low payoffs; reciprocity models assume that the desire to raise or lower others’ payoffs depends on how fairly those others are behaving” (p. 817-818). They conclude that their experimental results are consistent with the social welfare and reciprocity models, but suggest that difference aversion is less important, and indicate the need for further research.

The evidence from experimental research may be limited, however, in the extent to which it predicts behavior in more complex situations. As summarized in DellaVigna (2009), the role of social preferences has been studied in some field settings, particularly charitable giving and employer-employee relationships. The results suggest varied motives and do not necessarily match the experimental findings. DellaVigna notes that it can be difficult to separate social preferences from strategic decisions such as responses to social pressures, or from the “warm

³⁵ In reality, such redistribution may be difficult to achieve due to administrative costs, political constraints, and other factors.

³⁶ Other countries appear more willing to implement equity weighting; for example, the United Kingdom suggests weighting by income level in policy analyses (HM Treasury 2003).

glow” associated with giving rather than with a particular outcome (see, for example, Kahneman and Knetsch 1992, Andreoni et al. 2008). Thus while behavioral research provides support for the presence of altruistic (as well as malevolent) preferences, it does not provide values that can be directly used in benefit-cost analysis.

Finally, research also suggests that preferences are formed in part by interpersonal comparisons, indicating that private WTP will depend not only on one’s current endowment of income, health, and so forth, but also on how this endowment compares to that of others (see, for example, Luttmer 2005, Frank 2005, and Solnick and Hemenway 2005).³⁷ This research suggests that positional concerns are stronger in some areas, and can lead to a sort of “arms race” involving cyclical increases as each individual responds to positional changes made by others, ultimately leading to an oversupply of positional goods relative to other goods.

Substantially more research is needed to determine how these positional concerns affect the valuation of outcomes from social programs. The results of the few studies to-date are inconsistent. For example, Frank and Sunstein (2001) argue that ignoring the effects of relative income leads to benefit measures that are far too low, and suggest that the VSL estimates used in regulatory analysis may be significantly understated. However, Kneisner and Viscusi’s (2005) analysis of wage-risk trade-offs suggests that the opposite is true: relative position appears to have little effect on compensating wage differentials for mortality risks, and may decrease the VSL.

4.3 Conclusions and Implications

As indicated by the above discussion, practitioners have long struggled with how to represent social preferences (such as paternalistic altruism) in benefit-cost analysis, due to both research challenges and theoretical considerations. Behavioral research suggests additional ways in which concerns for others might affect valuation, such as reciprocating “good” and “bad” behavior or improving one’s relative position.

The research on social preferences reinforces the discussion in Section 2 on the potential importance of psychological attributes in valuation. For example, reciprocity may affect how individuals value mortality risks from crime (or from dangerous working conditions or polluting industries) in comparison to risks viewed as more accidental or naturally-occurring; concerns about relative position may affect programs providing income transfers differently than programs promoting improved health. This research also foreshadows some of the issues in study design discussed in Section 5, emphasizing the need to be clear about whether values reflect only an individual’s private consumption or also reflect benefits to the household or community. To the extent that studies (intentionally or inadvertently) include other-regarding preferences related to disparities, the values used in the analysis of efficiency may incorporate some equity concerns, potentially blurring the distinction between the two types of analysis.

Behavioral research on social preferences is not as well developed as the research discussed in the prior sections, in part because of the complex interactions between social preferences and social pressures as well as the difficulty of extrapolating from experimental results to more complex real-world behaviors. Thus it is hard to offer much concrete advice to policy analysts that can be implemented immediately. However, this research suggests that these

³⁷ Solnick and Hemenway (2005) trace this concern with relative position back to work by Thorstein Veblen in 1899 and cite several theoretical treatments published since that time.

issues should be considered both when valuing particular nonmarket outcomes and when considering the distributional and other consequences of social programs.

This discussion leads to two recommendations for our prototypical policy analyst.

- 1) In addition to other psychological attributes, studies should be designed to test the effects of social preferences on benefit values, focusing on outcomes associated with social programs rather than private goods or services.
- 2) When reviewing studies for benefit transfer, analysts should consider whether the results are likely to reflect other-regarding as well as private preferences, and discuss the implications for both the analysis of economic efficiency and its relationship to equity concerns.

5. IMPROVING VALUATION STUDIES

Thoughtful and well-informed preferences are desirable when valuing outcomes in policy analysis for several reasons. First, resources are limited, and decisions to divert these resources to achieving particular social goals warrant careful consideration. Given that there is no market to test the match between choices, preferences, and values for many of the outcomes associated with social programs, it seems that decisions to expend resources on these types of public goods should be justified with particular care. Interactions in real markets can provide various types of feedback and learning that may be lacking in nonmarket valuation studies, hence the studies may not capture market dynamics that could lead to more well-informed choices. However, in real markets firms face incentives to cater to consumer biases. They may also conceal information or provide misleading information to protect their market share or increase profits. Thus analysts need to recognize that review of market decisions could also lead to erroneous understanding of underlying consumer preferences.

Second, these programs and policies are generally focused on achieving long-term improvements in welfare, suggesting that we want to value the outcomes based on individuals' thoughtful assessment of how they might benefit over the long-run. Third, and perhaps most relevant for the purposes of this paper, behavioral research suggests that individuals may at times make "mistakes" or "errors" in expressing preferences or values, following simple decision rules or making choices that do not fully coincide with their assessment of their own welfare (see DellaVigna 2009 for a recent review).³⁸ While researchers have developed a number of strategies to address these issues, studies must be carefully designed so that the results do not ultimately reflect the values held by the researcher rather than the individuals studied.

5.1 Stated-Preference Research

Stated-preference researchers have a long tradition of considering psychological factors in designing their studies, and behavioral concerns are well-integrated into related guidance. Early work by Mitchell and Carson (1989) on best practices discusses approaches for encouraging honest and meaningful responses and controlling biases that may result from strategic behavior.

³⁸ One underlying issue in this research is the effect of "constructed" preferences. Particularly when faced with unfamiliar goods, individuals are not likely to have pre-existing preferences. Rather, their preferences will be formed when related decisions are required. This raises the question of whether these preferences will then be stable over time and across different contexts, and of how education and experience relate to creating this stability.

In 1993, a National Oceanic and Atmospheric Administration (NOAA) expert panel (chaired by Kenneth Arrow and Robert Solow) issued guidelines for these studies in the context of natural resource damage assessment cases, particularly to assess nonuse values. They discuss several behavioral considerations, such as pretesting surveys to ensure that the responses are not biased by the desire to please an interviewer, and using a referendum format (rather than an open-ended valuation question or payment card) to alleviate problems related to framing. Since that time, substantial additional research has been conducted that provides insights into how to best conduct these studies.

For example, in 2005 and again in 2010, the journal *Environmental and Resource Economics* published special issues that address these concerns. The 2005 issue includes seven articles on anomalies and stated-preference research (Braga and Starmer 2005, Guria et al. 2005, Hanley and Shogren 2005, Kahneman and Sugden 2005, Knetsch 2005, List 2005, Sugden 2005a, and Sugden 2005b). For example, List finds that experienced market participants behave in ways that appear consistent with the traditional economic model, while those with less experience exhibit some of the anomalies found in laboratory studies. Hanley and Shogren describe many options for improving stated-preference research, such as using workshops to better inform respondents, allowing “maybe” responses to recognize that individuals may only know their preferences over a limited range, and including “cheap talk” that discourages free-riding. Researchers are also increasingly using choice experiments that ask respondents to value varying combinations of attributes rather than a single outcome, to mimic market choices and avoid overly narrow framing of decisions.

The 2010 issue deals specifically with behavioral economics, and includes several articles that provide insights for stated-preference research (particularly Carlsson 2010, Hepburn et al. 2010, Johansson-Stenman and Konow 2010, Knetsch 2010, Smith and Moore 2010, and Shogren et al. 2010). For instance, Carlsson notes that while stated-preference researchers have much to learn from behavioral economics, behavioral economists can also learn from stated-preference research. He discusses four topics: revealed versus normative preferences (see Section 5.2), learning and constructed preferences, context dependence, and hypothetical bias. He indicates, for example, the importance of providing opportunities for learning, including practice and repetition. He also notes that while comparing results from studies with real and hypothetical payouts can be useful, there are several reasons for exercising caution in these comparisons.

The proliferation of these types of articles suggests that the intersection between behavioral research and the conduct of stated preference studies is fertile ground, with abundant lessons for researchers. However, it also means that there is a need to integrate the results of the numerous new studies to provide updated guidelines for best practices, that include guidance on how to best tailor stated-preference studies to different contexts and outcomes.

5.2 Revealed-Preference Research

The issues raised by behavioral research also affect revealed-preference studies. In particular, Beshears et al. (2008) distinguish between revealed (or positive) preferences (the choices people actually make) and normative preferences (the choices people think they should make). They note that actual choices result from combining normative preferences with the sorts of decisionmaking biases and anomalies found in behavioral research.

While Beshears et al. focus largely on how revealed and normative preferences diverge in savings and investment decisions, their findings have implications for many other contexts. They

identify five factors that increase the likelihood of a disparity: (1) passive choice (acceptance of defaults); (2) complexity (delayed choice, avoidance of complicated options, or misunderstanding of options due to cognitive difficulties); (3) limited personal experience (lack of learning through feedback); (4) third-party marketing (manipulating preferences); and (5) intertemporal choice (inconsistent time preferences, discussed in Section 3).

Beshears et al. then describe six approaches that can, in combination, help identify normative preferences:

...“*Structural estimation* specifies a positive model with a precise set of economic and psychological motives (perhaps including non-Bayesian thinking and other decision-making errors). This model is then estimated using data, and the resulting positive preferences are mapped into normative preferences using normative axioms.

Active decisions eliminate some biases generated by default regimes. Under an active decision regime, individuals are required to *explicitly state their preference* without being influenced by (or being able to rely on) a default option. In some circumstances, this preference elicitation will be more reliable and more socially efficient than allowing consumers to express their preferences by opting into or out of a pre-chosen default.

In most stationary economic environments, initial choices are likely to be further from normative optimality than choices made after many periods of experience. One should therefore give more weight to *asymptotic choices* when attempting to infer normative preferences.

When homogeneous individuals make noisy, error-prone decisions, their individual decisions do not reflect normative preferences, but their aggregate behavior can. Hence, normative preferences can sometimes be inferred from the central tendencies of *aggregated preferences*.

Self-reported preferences reveal something about an agent's goals and values. Normative economics should allow self-reports to have some standing. This is particularly true when self-reports can be used to distinguish confident consumer decisions from decisions that were made in a state of confusion.

Informed opinions come in two forms. External observers may offer expert advice, and decision-makers may themselves gain more expertise when they receive training or education. When trained/educated decision-makers make a choice, we call this an *informed preference*. Economists measuring normative preferences should give disproportionate weight to the actors who are most likely to know what they are doing.” (Beshears et al. 2008, p. 1793).

When revealed-preference studies are used to value nonmarket outcomes, researchers first identify a market good that includes the nonmarket outcome as one of its attributes. They then apply statistical methods to distinguish the value of the nonmarket outcome from the value of other attributes. For example, they may look at the trade-off between wages and job-related risks to value mortality risk reductions, or at residential decisions, recreational choices, or purchases of consumer goods to value various nonmarket attributes (see Freeman 2003 for more detailed examples). It seems sensible to consider whether both the market outcome and preferences toward the nonmarket attribute of concern may be significantly affected by the factors identified by Beshears et al.

Although the data and methods used in revealed-preference studies of nonmarket outcomes have been scrutinized (see, for example, Dockins et al. 2004), to the best of our knowledge these reviews have not systematically addressed the full range of behavioral issues listed above. More work is needed to determine the extent to which revealed-preference studies reflect normative preferences, rather than decisionmaking biases, when used to value the diverse outcomes associated with social programs. Thus while there is substantial synergy between the findings of behavioral economics and the evolving conduct of stated-preference research, more consideration of the implications of behavioral anomalies for revealed-preference research may also be desirable.

5.3 Conclusions and Implications

Under the standard economic model, the goal of nonmarket valuation is to mimic a market: to estimate the amount of money that individuals would be willing to exchange for these outcomes if they could be directly bought and sold. Under standard assumptions, these choices would be consistent with individual welfare. However, implementing this approach presents challenges given the findings of behavioral research. If individuals make mistakes, and their choices do not correspond with their welfare (as self-defined), relying on choices for valuation may not lead to the identification of the policy option that best enhances social welfare. Improving related studies so that they result in well-informed, thoughtful preferences is thus an important goal for researchers and policy analysts.

These issues have received substantial attention over the years in stated preference research, although the proliferation of new studies suggests that more review and integration is now needed. They have received relatively little attention when revealed preference methods are used, although similar issues may arise in market contexts.

This discussion has implications for both researchers and policy analysts:

- 1) Researchers should continue to incorporate the results of behavioral research into the design of valuation studies where useful and relevant, with the goal of identifying well-informed, thoughtful preferences, regardless of whether stated- or revealed-preference methods are used.
- 2) When evaluating study quality for benefit transfer, analysts should consider possible decisionmaking bias along with other uncertainties, including any evidence that the study population did not fully understand the context or made other cognitive errors.

6. DETERMINING THE ROLE OF BENEFIT-COST ANALYSIS

Many reviews of the implications of behavioral economics compare and contrast a simplified version of the standard economic model with the findings of behavioral research. However, the standard model both can and does incorporate more complex considerations than suggested by these reviews. Some of these complexities reflect elaboration of the model; e.g., to incorporate lifecycle consumption, household production, or social-welfare functions. Others integrate behavioral findings, such as the effects of psychological factors on the valuation of nonmarket outcomes. Overall, the conduct of benefit-cost analysis reflects a number of pragmatic concerns, using economic theory as a starting point, incorporating behavioral considerations, and focusing on information that is useful for decisionmaking. In practice, the gap between the traditional and

behavioral models is narrower than it might appear, and there is substantial synergy between nonmarket valuation and behavioral research.

The discussion in the prior sections of this paper suggests additional avenues to explore. Below, we first summarize the recommendations from the previous sections. We then discuss the more general implications for the conduct and interpretation of benefit-cost analysis, focusing on its practical implementation.

6.1 Summary of Conclusions and Recommendations

In this paper, we discuss the implications of behavioral economics for analyzing the benefits and costs of social policies and programs. While recognizing that behavioral research can aid in predicting program participation rates and effectiveness, as well as in predicting real resource costs conditional on these responses, we focus on the valuation of program benefits and costs. We take the role of the analyst who is asked to assess an exogenously-determined program or policy, and seeks advice on incorporating the results of behavioral research into the assessment given currently available data. We consider four areas: context-dependent valuation of nonmarket benefits (Section 2); present bias and exponential versus hyperbolic discounting (Section 3); social preferences such as altruism and equity (Section 4); and the conduct of nonmarket valuation studies (Section 5).

Because behavioral economics is a large and rapidly growing field, the significance or pervasiveness of many of its findings are not yet clear, and these findings have not yet been combined into a widely-accepted model that supplements or supplants the standard economic framework. As a result, we often raise questions for further research or note options for the analyst rather than trying to resolve the underlying issues. Throughout, we attempt to maintain the traditional perspective of benefit-cost analysis as describing the preferences of affected individuals (contingent on baseline conditions, including the current distribution of income and other resources). While we recognize that individual choices may include errors; i.e., diverge from what the individuals might choose given better information and more experience or reflection, we urge caution in substituting the judgment of the analyst for the information collected from representative members of the population.

In Table 1, we summarize our key recommendations.

Table 1. Summary of Major Recommendations

VALUING PSYCHOLOGICAL ATTRIBUTES
<ol style="list-style-type: none"> 1) Studies should be designed to test the effects of the psychologically-salient attributes found in behavioral research on benefit values, as well as the effects of other physical and psychological attributes. 2) Such studies should consider both the holistic effect of the full range of attributes relevant to a particular context, and the effects of varying the attributes one-by-one to develop adjustments for transferring values to other contexts. 3) Values should be not be rejected unless the study does not meet basic criteria for quality or adhere to generally-accepted principles for best practices. Rejected values, and the basis for rejecting them, should be clearly documented. 4) When values are uncertain, sensitivity, probabilistic, or breakeven analysis should be used to test the effects of this uncertainty on the results. 5) Where quantitative estimates are not available, the potential effects of both psychologically-salient and physical attributes should be discussed qualitatively.
ESTIMATING TIME PREFERENCES
<ol style="list-style-type: none"> 1) Analysts should provide a schedule of undiscounted costs and benefits, in the form of a bar chart (or other graphic) or a table, so that the decisionmaker can inspect this pattern. Simply summarizing costs and benefits as a net present value does not provide adequate information given related uncertainties. 2) In choosing a discount rate or set of rates, as well as a functional form (exponential or hyperbolic) analysts must be clear about what, conceptually, they intend to represent for review by decisionmakers. The quality of the underlying data and related assumptions should also be described. 3) Analysts should consider the impact of alternative rates on their results, in particular emphasizing the discount rates at which benefits do, and do not, exceed costs, and/or at which the relative ranking of different policy options change. Comparing these switch points to rates reported in the empirical literature will provide insights into the implications of related uncertainties for decisionmaking.
SEPARATING PRIVATE FROM SOCIAL PREFERENCES
<ol style="list-style-type: none"> 1) In addition to other psychological attributes, studies should be designed to test the effects of social preferences on benefit values, focusing on outcomes associated with social programs rather than private goods or services. 2) When reviewing studies for benefit transfer, analysts should consider whether the results are likely to reflect other-regarding as well as private preferences, and discuss the implications for both the analysis of economic efficiency and its relationship to equity concerns.
IMPROVING VALUATION STUDIES
<ol style="list-style-type: none"> 1) Researchers should continue to incorporate the results of behavioral research into the design of valuation studies where useful and relevant, with the goal of identifying well-informed, thoughtful preferences, regardless of whether stated- or revealed-preference methods are used. 2) When evaluating study quality for benefit transfer, analysts should consider possible decisionmaking bias along with other uncertainties, including any evidence that the study population did not fully understand the context or made other cognitive errors.

As the above summary indicates, behavioral research is not yet sufficient to support specific recommendations for adapting the practice of benefit-cost analysis in many areas. Our recommendations are thus similar to current best practice recommendations, which take into account uncertainty in the underlying data and analyses. In some cases (such as time preferences and the design of stated-preference studies), the behavioral research base is relatively large and additional criteria-driven review may be helpful. In other areas (such as social preferences), the behavioral research is limited and does not necessarily address outcomes relevant to the evaluation of social programs. In the near term, the implications of behavioral research can be explored by testing the sensitivity of the results to different values, using probabilistic or

breakeven analysis, or discussing impacts qualitatively. In the longer term, more significant changes may be warranted as new research becomes available.

6.2 Benefit-Cost Analysis as Applied Pragmatism

Conducting benefit-cost analysis is challenging, regardless of whether framed by the traditional or behavioral economic model. However, denoting values in monetary terms has the advantage of providing information on the intensity as well as the direction of preferences. It also mimics the effects of policy decisions, which often require trading money for the outcomes of concern.³⁹ Thus we believe that valuation research provides useful information, despite its limitations – as long as the analyst is clear about the implications of these limitations for the conclusions of the analysis.

Those conducting such analysis may be most comfortable describing it as a positive, descriptive exercise, that provides decisionmakers with factual information on the preferences of those potentially affected by the benefits and costs of the policy. Behavioral economics questions aspects of this model, suggesting that individuals do not always act rationally, are motivated in part by social preferences, and make choices that do not necessarily coincide with their own long-term welfare even as self-defined.

One possible response to these findings is to substitute a paternalistic approach for the positive or descriptive model, placing the analyst and/or policymaker in the position of deciding which preferences are rational and welfare-enhancing and which are not. This is a task worthy of Solomon, because the dividing line will often be vague and preferences that appear nonsensical on the surface may become more sensible as one digs deeper. The alternative is to continue to work on developing information on the preferences that emerge when individuals are well-informed and have the opportunity for reflection. This task may require the endurance of Hercules, however. While we can clearly improve our understanding of these preferences, the complex attributes of social programs as well as the complexities of human judgment mean that our understanding of these preferences will always be somewhat imperfect.

Fortunately, the approaches currently used in benefit-cost analysis provide several practical tools for dealing with these uncertainties. These tools include quantitative sensitivity, probabilistic, and breakeven analysis, qualitative discussion of the implications of concerns that cannot be easily quantified, and explicit recognition that benefit-cost analysis should be only one of several inputs into policy decisions.

These concerns suggest a more pragmatic role for benefit-cost analysis than indicated by either the simplified version of the standard economic model or by behavioral economics.⁴⁰ Imagine again our prototypical policy analyst, asked to provide information that supports a recommendation about whether or not to proceed with an exogenously-determined policy option.

³⁹ In addition, given the Kaldor-Hicks potential compensation criteria, Hammitt (2009) notes, “[t]he justification for using a monetary metric is that money can be transferred among affected individuals at low or modest cost (e.g., through changes in the tax code and income-support programs), but other possible metrics are more difficult or costly to transfer” (p. 194).

⁴⁰ There is a large literature on the rationale for conducting benefit-cost analysis that is too extensive to summarize here, including more philosophical discussions (e.g., Adler and Posner 2006) as well as more detailed practical recommendations (e.g., Harrington et al. 2000). In particular, Sunstein (2000) argues that benefit-cost analysis “is most plausibly justified on cognitive grounds – as a way of counteracting predictable problems in individual and social cognition” (p. 1059), describing how such analysis can be used to counterbalance related errors.

If well-conducted, following the benefit-cost analysis framework provides the analyst with several advantages. While some of these advantages are not exclusive to benefit-cost analysis, in combination they help develop useful information for decisions.

- 1) Benefit-cost analysis offers a well-established and tested approach for identifying and assessing the physical impacts of different policy options in addition to estimating their economic value, supported by substantial research and guidance documents that address best practices as well as by numerous examples of previously completed studies.
- 2) Benefit-cost analysis' focus on monetization provides information on the intensity of preferences as well as on individuals' willingness to make the types of trade-offs implicit in many social policy decisions (which involve exchanging money for social outcomes rather than for other goods and services) while not precluding the consideration of nonquantifiable effects and uncertainty.
- 3) Benefit-cost analysis promotes a broad perspective, by incorporating the preferences of individuals affected by the costs and/or benefits of the policy options rather than being limited to the preferences of those most directly involved in the decision.
- 4) Benefit-cost analysis encompasses several tools for incorporating uncertain outcomes and values, including sensitivity, probabilistic, and breakeven analysis as well as qualitative discussion.
- 5) Benefit-cost analysis aids in focusing data collection and research and provides rough "stopping rules" for these efforts. These include the principle of proportionate analysis (matching the level of analytic effort to the stakes of the decision) and value of information considerations (focusing on research that is likely to affect the decision), essentially applying a benefit-cost test to the analysis itself.
- 6) Benefit-cost analysis provides a transparent record of the data, assumptions, and analyses considered in the decision, if well-documented.
- 7) Benefit-cost analysis helps decisionmakers and stakeholders to clarify areas of agreement and disagreement, separating data from assumptions and allowing those who disagree to test the effects of alternative analytic approaches.

Under this formulation, benefit-cost analysis becomes one of many sources of information for decisionmaking rather than providing "the" answer to a policy question. It is instead a method for collecting, organizing, and evaluating information relevant to the decision.

In summary, when taken to the extreme, traditional economic theory can be interpreted as suggesting that decisions are always rational and welfare-enhancing, while behavioral economics could be interpreted as suggesting that individual choices are so unstable and prone to error as to render the results of benefit-cost analysis meaningless. Both extremes substantially overstate the implications of related research. Analysts recognize that benefit-cost analysis as traditionally conducted is one of many sources of information, and must be supplemented by consideration of other concerns including nonquantified effects and equity. Behavioral researchers acknowledge that behavioral anomalies may be limited in occurrence and/or in significance, and that these anomalies can be counterbalanced to varying degrees by education and experience.

The introduction to this paper suggests that the core issue raised by behavioral economics is whether choices in fact reflect welfare-enhancing preferences. The answer to this question appears to be "it depends." In some cases, behavioral research has identified attributes that individuals may truly value, that are not explicitly incorporated into the standard economic

model. In other cases, decisions may reflect simplifying heuristics or emotional responses that drive a wedge between choices and underlying, normative preferences. We are far from fully understanding these distinctions, but the questions raised by behavioral economics provide useful insights for how we might further explore these issues. While benefit-cost analysis is useful for many reasons, it will inevitably have limitations. Clear discussion of its implications and uncertainties will always be required.

REFERENCES

- Adler M.A. 2008. "Risk Equity; A New Proposal." *Harvard Environmental Law Review*. 32, 1-47.
- Adler, M.A. and E.A. Posner. 2006. *New Foundations for Benefit-Cost Analysis*. Cambridge: Harvard University Press.
- Alberini, A., M. Cropper, A. Krupnick, and N. Simon. 2004. "Does the Value of a Statistical Life Vary with Age and Health Status? Evidence from the U.S. and Canada." *Journal of Environmental Economics and Management*. 48(1), 769-792.
- Allcott, H. and S. Mullainathan. 2010. "Behavior and Energy Policy." *Science*. 327, 1204-1205.
- Andreoni, J., W.T. Harbaugh, and L. Vesterlund. 2008. "Altruism in Experiments." *The New Palgrave Dictionary of Economics*. (Second Edition). (S.N. Durlauf and L.E. Blume, eds.) New York: Palgrave Macmillan.
- Bergstrom, T.C. 2006. "Benefit-Cost in a Benevolent Society." *American Economic Review*. 96(1), 339-351.
- Bernheim, B.D. and A. Rangel. 2007. "Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision-Makers." In *Economic Institutions and Behavioral Economics* (P. Diamond and H. Vartiainen, eds.). Princeton: Princeton University Press.
- Bernheim, B.D. and A. Rangel. 2009. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics*. 124(1), 51-104.
- Berns, G.S., D. Laibson, and G. Loewenstein. 2007. "Intertemporal Choice – Toward an Integrative Framework." *Trends in Cognitive Sciences*. 11(11), 482-488.
- Beshears, J., J.J. Choi, D. Laibson, and B.C. Madrian. 2008. "How Are Preferences Revealed?" *Journal of Public Economics*. 92, 1787-1794.
- Braga, J. and C. Starmer. 2005. "Preference Anomalies, Preference Elicitation, and the Discovered Preference Hypothesis." *Environmental and Resource Economics*. 32(1), 55-89.
- Camerer, C.F., G. Loewenstein, and M. Rabin (eds.). 2004. *Advances in Behavioral Economics*. Princeton: Princeton University Press.
- Carlsson, F. 2010. "Design of Stated Preference Surveys: Is There More to Learn from Behavioral Economics?" *Environment and Resource Economics*. 46(2), 167-177.

- Chabris, C.F., D.I. Laibson, and J.P. Schuldt. 2008. "Intertemporal Choice." *The New Palgrave Dictionary of Economics* (Second Edition). (S.N. Durlauf and L.E. Blume, eds.) New York: Palgrave Macmillan.
- Charness, G. and M. Rabin. 2002. "Understanding Preferences with Simple Tests." *Quarterly Journal of Economics*. 117(3), 817-869.
- Congdon, W.J., J.R. Kling, and S. Mullainathan. 2011. *Policy and Choice: Public Finance through the Lens of Behavioral Economics*. Washington, D.C.: Brookings Institution Press.
- Corso, P.S., J.K. Hammitt, and J.D. Graham. 2001. "Valuing Mortality-Risk Reduction: Using Visual Aids to Improve the Validity of Contingent Valuation." *Journal of Risk and Uncertainty*. 23(2), 165-184.
- DellaVigna, S. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*. 47(2), 315-372.
- Diamond, P. and H. Vartiainen (eds). 2007. *Behavioral Economics and Its Applications*. Princeton and Oxford: Princeton University Press.
- Dockins, C., R.R. Jenkins, N. Owens, N.B. Simon, and L.B. Wiggins. 2002. "Valuation of Childhood Risk Reduction: The Importance of Age, Risk Preferences and Perspective." *Risk Analysis*. 22(2), 335-346.
- Dockins, C., K. Maguire, N. Simon, and M. Sullivan. 2004. *Value of Statistical Life Analysis and Environmental Policy: A White Paper*. Prepared by the National Center for Environmental Economics, U.S. Environmental Protection Agency, for presentation to the Environmental Economics Advisory Committee of the Science Advisory Board.
- Ellsberg, D. 1961. "Risk, Ambiguity, and the Savage Axioms." *The Quarterly Journal of Economics*. 75(4), 643-669.
- Farrow, S. 2009. "Incorporating Equity in Regulatory and Benefit-Cost Analysis." Unpublished manuscript prepared for New Ideas for Risk Regulation conference.
- Frank, R.H. 2005. "Positional Externalities Cause Large and Preventable Welfare Losses." *American Economic Review*. 95(2), 137-141.
- Frank, R.H. and C.R. Sunstein. 2001. "Cost-Benefit Analysis and Relative Position." *University of Chicago Law Review*. 68(2), 323-374.
- Frederick, S., G. Loewenstein, and T. O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*. 40, 351-401.
- Freeman, A.M. III. 2003. *The Measurement of Environmental and Resource Values: Theory and Methods* (Second Edition). Washington D.C.: Resources for the Future.

- Fudenberg, D. 2006. "Advancing beyond *Advances In Behavioral Economics*." *Journal of Economic Literature*. 44, 694-711.
- Fudenberg, D. and D.K. Levine. 2006. "A Dual-Self Model of Impulse Control." *American Economic Review*. 96(5), 1449-1476.
- Gold, M.R., J.E. Siegel, L.B. Russell, and M.C. Weinstein (eds.). 1996. *Cost-Effectiveness in Health and Medicine*. New York and Oxford: Oxford University Press.
- Gollier, C. and R. Zeckhauser. 2005. "Aggregation of Heterogeneous Time Preferences." *Journal of Political Economy*. 113(4), 878-896.
- Graham, J.D. 2008. "Saving Lives through Administrative Law and Economics." *University of Pennsylvania Law Review*. 157, 395-540.
- Green, J.R. and D.A. Hojman. 2009. *Choice, Rationality, and Welfare Measurement*. Unpublished Manuscript.
- Guria, J., J. Leung, M. Jones-Lee, and G. Loomes. 2005. "The Willingness to Accept Value of Statistical Life Relative to the Willingness to Pay Value: Evidence and Policy Implications." *Environmental & Resource Economics*. 32, 113-127.
- Hammit, J.K. 2000. "Valuing Mortality Risk: Theory and Practice." *Environmental Science and Technology*. 34, 1396-1400.
- Hammit, J.K. 2007. "Valuing Changes in Mortality Risk: Lives Saved versus Life Years Saved." *Review of Environmental Economics and Policy*. 1(2), 228-240.
- Hammit, J.K. 2009. "Response: *Saving Lives: Benefit-Cost Analysis and Distribution*." *University of Pennsylvania Law Review PENNumbra*. 157, 189-199.
- Hammit, J.K. and J.D. Graham. 1999. "Willingness to Pay for Health Protection: Inadequate Sensitivity to Probability?" *Journal of Risk and Uncertainty*. 18(1), 33-62.
- Hammit, J.K. and L.A. Robinson. 2010. "The Income Elasticity of the Value per Statistical Life: Transferring Estimates Between High and Low Income Populations." *Journal of Benefit-Cost Analysis*.
- Hammit, J.K. and N. Treich. 2007. "Statistical vs. Identified Lives in Benefit-Cost Analysis," *Journal of Risk and Uncertainty*. 35(1), 45-66.
- Hanemann, W.M. 1991. "Willingness to Pay and Willingness to Accept: How Much Can They Differ?" *American Economic Review*. 81, 635-647.

Hanley, N. and J.F. Shogren. 2005. "Is Benefit-Cost Analysis Anomaly-Proof?" *Environmental and Resource Economics*. 32(1), 13-34.

Harrington, W., L. Heinzerling, and R.D. Morgenstern. 2009. *Reforming Regulatory Impact Analysis*. Washington, DC: Resources for the Future.

Harvey, C.M. 1994. "The Reasonableness of Non-Constant Discounting." *Journal of Public Economics* 53, 31-51.

Hepburn, C., S. Duncan, and A. Papachristodoulou. 2010. "Behavioural Economics, Hyperbolic Discounting, and Environmental Policy." *Environment and Resource Economics*. 46(2), 189-206.

HM Treasury. 2003. *The Green Book: Appraisal and Evaluation in Central Government*. London: TSO.

Horowitz, J.K. and K.E. McConnell. 2002. "A Review of WTA/WTP Studies." *Journal of Environmental Economics and Management*. 44(3), 426-447.

Institute of Medicine. 2006. *Valuing Health for Regulatory Cost-Effectiveness Analysis*. (W. Miller, L.A. Robinson, and R.S. Lawrence, eds.). Washington, D.C.: The National Academies Press.

Jacobsson, F., M. Johannesson, and L. Borgquist. 2007. "Is Altruism Paternalistic?" *The Economic Journal*. 117(520), 761-781.

Johansson-Stenman, O. 2008. "Mad Cows, Terrorism and Junk Food: Should Public Policy Reflect Perceived or Objective Risks?" *Journal of Health Economics*. 27(2), 234-248.

Johansson-Stenman, O. and J. Konow. 2010. "Fair Air: Distributive Justice and Environmental Economics." *Environmental and Resource Economics*. 46(2), 147-166.

Jones-Lee, M.W. 1991. "Altruism and the Value of Other People's Safety." *Journal of Risk and Uncertainty*. 4(2), 213-219.

Just, R.E., D.L. Hueth, and A. Schmitz. 2004. *The Welfare Economics of Public Policy*. Cheltenham and Northampton: Edward Elgar.

Kahneman, D. 2000a. "Experienced Utility and Objective Happiness: A Moment-Based Approach." In *Choices, Values, and Frames* (D. Kahneman and A. Tversky, eds.). Cambridge: Cambridge University Press.

Kahneman, D. 2000b. "Evaluation by Moments: Past and Future." In *Choices, Values, and Frames* (D. Kahneman and A. Tversky, eds.). Cambridge: Cambridge University Press.

Kahneman D., J.L. Knetsch, and R.H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives*. 5(1), 192-206.

Kahneman D. and J.L. Knetsch. 1992. "Valuing Public Goods: The Purchase of Moral Satisfaction." *Journal of Environmental Economics and Management*. 22(1), 57-70.

Kahneman, D. and R. Sugden. 2005. "Experienced Utility as a Standard of Policy Evaluation." *Environmental and Resource Economics*. 32(1), 161-181.

Kahneman, D. and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*. 47(2), 263-91.

Kahneman, D. and A. Tversky (eds.). 2000. *Choices, Values, and Frames*. Cambridge: Cambridge University Press.

Knetsch, J.L. 2005. "Gains, Losses, and the US-EPA *Economic Analysis Guidelines*: A Hazardous Product?" *Environmental and Resource Economics*. 32(1), 91-112.

Knetsch, J.L. 2010. "Values of Gains and Losses: Reference States and Choices of Measure." *Environment and Resource Economics*. 46(2), 179-188.

Kniesner, T.J. and W.K. Viscusi. 2005. "Value of a Statistical Life: Relative Position vs. Relative Age." *American Economic Review*. 95(2), 142-146.

Laibson, D.I. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*. 112(2), 443-478.

Laibson, D.I. 2010. "Behavioral Economics." Presented at American Economic Association Annual Meeting, Continuing Education Program. [Slides available at: <http://www.economics.harvard.edu/faculty/laibson/AEA2010%20Mini%20Course>; webcast available to AEA members at: <http://www.aeaweb.org/webcasts/assa2010.php>.]

Laibson, D.I., A. Repetto, and J. Tobacman. 2007. *Estimating Discount Functions with Consumption Choices over the Lifecycle*. Unpublished manuscript.

Layard, R. 2010. "Measuring Subjective Well-Being." *Science*. 327, 534-535.

Lindhjem, H., S. Navrud, and N.A. Braathen. 2010. *Valuing Lives Saved from Environmental, Transport, and Health Policies: A Meta-Analysis of Stated Preference Studies*. Prepared for the Organization for Economic Co-operation and Development. ENV/EPOC/WPNEP(2008)10/FINAL.

List, J.A. 2005. "Scientific Numerology, Preference Anomalies, and Environmental Policymaking." *Environmental and Resource Economics*. 32, 35-53.

- Loomis, J.B. 2009. *Importance of Incorporating Distributional Issues in Benefit Cost Analysis*. Unpublished manuscript.
- Luttmer. 2005. "Neighbors as Negatives: Relative Earnings and Well-Being." *Quarterly Journal of Economics*. 120(3), 963-1002.
- Mitchell, R.C. and R.T. Carson. 1989. *Using Surveys to Value Public Goods: The Contingent Valuation Method*. Washington, DC: Resources for the Future.
- Munro, A. 2009. "Introduction to the Special Issue: Things We Do and Don't Understand about the Household and the Environment." *Environmental and Resource Economics*. 43(1), 1-10.
- National Oceanic and Atmospheric Administration. 1993. "Natural Resource Damage Assessments under the Oil Pollution Act of 1990." *Federal Register*. 58(10), 4601-4614.
- Rabin, M. 1998. "Psychology and Economics." *Journal of Economic Literature*. 36, 11-46.
- Read, D. 2001. "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty*. 23(1), 5-32.
- Riddel, M. and W.D. Shaw. 2006. "A Theoretically-Consistent Empirical Model of Non-Expected Utility: An Application to Nuclear-Waste Transport." *Journal of Risk and Uncertainty*. 32, 131-150.
- Robinson, L.A. 2004. *Current Federal Agency Practices for Valuing the Impact of Regulations on Human Health and Safety*. Prepared for the Committee to Evaluate Measures of Health Benefits for Environmental, Health, and Safety Regulation, Institute of Medicine.
- Robinson, L.A. 2008a. "Assessing Regulatory Costs and Benefits." Prepared for: Nriagu, J. (ed.). *Encyclopedia of Environmental Health* (forthcoming). Oxford UK: Elsevier Ltd.
- Robinson, L.A. 2008b. *Valuing Mortality Risk Reductions in Homeland Security Regulatory Analyses*. Prepared for U.S. Customs and Border Protection, Department of Homeland Security, under subcontract to Industrial Economics, Incorporated.
- Robinson, L.A. and J.K. Hammitt. 2009. *The Value of Reducing Air Pollution Risks in Sub-Saharan Africa*. Prepared for the World Bank under subcontract to ICF International.
- Robinson, L.A., J.K. Hammitt, J.E. Aldy, A. Krupnick, and J. Baxter. 2010. "Valuing the Risk of Death from Terrorist Attacks." *Journal of Homeland Security and Emergency Management*. 7(1), Art. 14.
- Rubinstein, A. 2003. "'Economics and Psychology'? The Case of Hyperbolic Discounting." *International Economic Review*. 44(4), 1027-1216.

- Samuelson, P.A. 1937. "A Note on the Measurement of Utility." *Review of Economic Studies*. 4(2), 155-161.
- Shogren, J.F. 2005. "Economics of Diet and Health: Research Challenges." *Acta Agriculturae Scand Section C*. 2, 117-127.
- Shogren, J.F., G.M. Parkhurst, and P. Banerjee. 2010. "Two Cheers and a Qualm for Behavioral Environmental Economics." *Environment and Resource Economics*. 46(2), 235-247.
- Slovic, P. 1987. "Perception of Risk." *Science*. 236(4799), 280-285.
- Smith, V.K. 2007. "Reflections on the Literature." *Review of Environmental Economics and Policy*. 1(1), 152-165.
- Smith, V.K. 2008. "Reflections on the Literature." *Review of Environmental Economics and Policy*. 2(2), 292-308.
- Smith, V.K. and E.M. Moore. 2010. "Behavioral Economics and Benefit Cost Analysis." *Environment and Resource Economics*. 46(2), 217-234.
- Smith, V.K., S.K. Pattanayak, and G. L. Van Houtven. 2006. "Structural Benefit Transfer: An Application Using VSL Estimates." *Ecological Economics*. 60, 361-371.
- Solnick, S.J. and D. Hemenway, 2005. "Are Positional Concerns Stronger in Some Domains than in Others?" *American Economic Review*. 95(2), 147-151
- Sugden, R. 2004. "The Opportunity Criterion: Consumer Sovereignty without the Assumption of Coherent Preferences." *American Economic Review*. 94(4), 1014-1033.
- Sugden, R. 2005a. "Coping with Preference Anomalies in Benefit-Cost Analysis: A Market Simulation Approach." *Environmental and Resource Economics*. 32(1), 129-160.
- Sugden, R. 2005b. "Anomalies and Stated Preference Techniques: A Framework for Discussion of Coping Strategies." *Environmental and Resource Economics*. 32(1), 1-12.
- Sugden, R. 2009. "Market Simulation and the Provision of Public Goods: A Non-Paternalistic Response to Anomalies in Environmental Evaluation." *Journal of Environmental Economics and Management*. 57(1), 87-103.
- Sunstein, C.R. 2000. "Cognition and Benefit-Cost Analysis." *Journal of Legal Studies*, 29, 1059-1103.
- Sunstein, C.R. 2007. "Willingness to Pay vs. Welfare." *Harvard Law and Policy Review*, 1, 303-330.
- Thaler, R.H. 1992. *The Winner's Curse*. Princeton: Princeton University Press.

Thaler, R.H. and S. Mullainathan. 2008. "How Behavioral Economics Differs from Traditional Economics." *The Concise Encyclopedia of Economics*.

Thaler, R.H. and C.R. Sunstein. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven: Yale University Press.

U.S. Environmental Protection Agency. 2000. *Guidelines for Preparing Economic Analysis*. EPA 240-R-00-003.

U.S. Office of Management and Budget. 1992. *Circular A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs*.

U.S. Office of Management and Budget. 2003. *Circular A-4: Regulatory Analysis*.

Viscusi, W.K. and J.E. Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World." *Journal of Risk and Uncertainty*. 27(1), 5-76.

Viscusi, W.K., W.A. Magat, and A. Forrest. 1988. "Altruistic and Private Valuations of Risk Reduction." *Journal of Policy Analysis and Management*. 7(2), 227-245.

Viscusi, W.K., W.A. Magat, and J. Huber. 1991. "Communication of Ambiguous Risk Information." *Theory and Decision*. 31, 159-173.

Weitzman, M.L. 2001. "Gamma Discounting." *The American Economic Review*. 91(1), 260-271.

Willig, R.D. 1976. "Consumer's Surplus without Apology." *The American Economic Review*. 66(4), 589-597.

Zerbe, R. 2009. *Ethical Benefit-Cost Analysis as Art and Science: Ten Rules for Benefit-Cost Analysis*. Unpublished manuscript.