

Lisa A. Robinson*

[Ir]rationality, Happiness, and Benefit-Cost Analysis: Introduction to the Special Issue

Abstract: Behavioral economics and happiness research have many important implications for the conduct of benefit-cost analysis as well as for policy design and implementation. By identifying ways in which we may act irrationally and providing new perspectives on the relationship between our circumstances and our sense of well-being, this research raises numerous questions regarding the evaluation of individual and societal welfare and the desirability of alternative policies. In this special issue, we present a series of articles that explore these concerns and provide significant new insights.

Keywords: behavioral economics; benefit-cost analysis; happiness; social welfare; subjective well-being.

JEL classifications: D03; D60; D61; H40; I18; I31.

1 Introduction

Behavioral economics and happiness research have captured the attention of scholars, policymakers, and the general public, by identifying ways in which we may act irrationally as well as providing insights into the relationship between our circumstances and our sense of well-being. This research has many implications for designing policies, for predicting how individuals will respond, and for evaluating the desirability of the outcomes. But what does it mean for benefit-cost analysis?

Benefit-cost analysis is built on the foundation of neoclassical welfare economics; behavioral economics and happiness research challenge that foundation. Does this mean benefit-cost analysis is useless? Or can it be adapted to incorporate the findings of this research? If so, what forms should this adaptation take?

The *Journal of Benefit-Cost Analysis* (JBCA) has been engaged in addressing these questions since its inception. This issue continues that tradition, providing a series of articles that build on past work to provide new insights. We expect these

Lisa A. Robinson, Harvard University, Centers for Risk Analysis and Health Decision Science, 718 Huntington Avenue, Boston, Massachusetts, 02115, USA, e-mail: robinson@hsph.harvard.edu

articles will stimulate continued discussion, both on the pages of this journal and elsewhere, as our understanding of these issues continues to evolve.

2 Background and definitions

What do we mean by “behavioral” economics and “happiness” research? Neither term is defined entirely consistently in the literature. All economics research is concerned to some extent with how people behave, and to many “welfare” sounds synonymous with “happiness.” The distinction between these frameworks and the standard economic model is not always clear, in part because some observed phenomena may have both behavioral and neoclassical explanations.

Behavioral economics has emerged largely from the increased integration of psychological research into the models used to explain or predict economic behavior. Those involved often distinguish their work by noting that, in its simplest form, the standard economic model assumes that people behave self-interestedly and rationally, while they consider how human behavior may deviate from these assumptions. These deviations may result from our limited ability to process information, incomplete self-control, or other-regarding preferences. At times, we may rely on simplifying heuristics and biases that lead us to act in ways that are contrary to our own welfare, even as self-defined. For example, numerous studies have found that preferences vary depending on the time frame, so that a decision made in the near term (such as eating dessert) is inconsistent with preferences over the longer run (such as losing weight).

Neoclassical economics often relies on expected utility theory as initially formulated by von Neumann and Morgenstern in the mid-1940s, while behavioral economics reflects challenges to that model, beginning most notably with work by Kahneman and Tversky and by Thaler in the late 1970s and early 1980s (see, for example, the collections of articles in Thaler, 1992 and Kahneman & Tversky, 2000). The standard model generally assumes that individuals assign utilities to consequences and prefer the choice that maximizes the expected value of this utility. In contrast, prospect theory (Kahneman & Tversky, 1979) and related models suggest that preferences depend on the reference point from which they are measured (with losses valued more than gains and diminishing sensitivity with increasing distance from the reference point) and that probabilities are evaluated nonlinearly (with changes in probabilities near zero and one more important than changes in intermediate probabilities).

This framework raises serious concerns regarding the appropriate treatment of preferences in benefit-cost analysis. Should the preferences revealed by individual

behavior always be used to evaluate welfare? Or should we override these revealed preferences in cases where it appears individuals are making mistakes? Much of the work on the implications of behavioral economics for benefit-cost analysis focuses on options for addressing these concerns.

While happiness research is sometimes described as related to behavioral economics, it is largely distinct. It considers how individuals rate their own happiness or life satisfaction; i.e., their subjective or self-reported well-being (SWB). It includes surveys that ask individuals to evaluate their overall life satisfaction, to indicate how they feel at the time they experience various activities or events, or to report on their sense of purpose or meaning (eudemonia). These evaluative, experiential, and eudemonic measures are fundamentally different and do not necessarily lead to the same conclusions; which measure is most appropriate depends in part on the question one wishes to answer.

Some have argued that these SWB measures should replace the estimates of willingness to pay (WTP) and willingness to accept compensation (WTA) conventionally used for valuation in benefit-cost analysis (e.g., Kahneman & Sugden, 2005; Dolan & Kahneman, 2008; Layard, 2010). Such SWB measures are at times assigned monetary values for use in policy analysis, by estimating the change in income that has the same effect on SWB as the outcome of concern (e.g., Fujiwara & Campbell, 2011).

The articles in the current volume, as well as in previous *JBCA* issues, continue these debates and provide numerous important and useful insights.

3 Previous work

Behavioral economics and happiness research were attracting growing attention when *JBCA* began publication. At the same time that the John D. and Catherine T. MacArthur Foundation helped found *JBCA* and the Society for Benefit-Cost Analysis, it also funded a project to develop principles and standards for conducting benefit-cost analysis. One of the resulting *JBCA* articles explicitly addressed the implications of behavioral economics for benefit-cost analysis (Robinson & Hammitt, 2011). The authors summarize the status of related research, discuss several challenges this research poses, and consider how to accommodate these challenges within the benefit-cost analysis framework. In particular, they address differences between WTP and WTA estimates, psychological reactions to risk, inconsistent time preferences, social (other-regarding) preferences, and nonmarket valuation. Differing perspectives have been provided by other *JBCA* authors, such as Brennan (2014), who argues that behavioral economics suggests that benefit-cost

analysis should be replaced with other decision-making frameworks. However, he indicates that strong evidence is needed before benefit-cost analysis is abandoned.

Several *JBCA* articles explore the issues raised by behavioral research in specific contexts. For example, Jin et al. (2015) and Cutler et al. (2015) address difficult questions related to estimating the welfare effects of smoking policies. Although the details of the approaches differ, each research team estimates the benefits of these policies by identifying a group whose decisions appear more likely to be fully informed and rational, then uses data from that group to estimate benefits for all smokers. Another example is a set of articles by Viscusi (2015), Hammitt (2015), and Knetsch (2015), who provide varying views on the implications of differences between estimates of WTP and WTA that may arise as a result of reference dependence and loss aversion.

The articles in the current issue were inspired in part by Raj Chetty's (2015) Ely Lecture, "Behavioral Economics and Public Policy: A Pragmatic Perspective." Chetty argues that rather than debating the challenges that behavioral economics poses to the foundational assumptions of neoclassical economics, we should focus on whether incorporating the results of such research helps us predict outcomes and make policy decisions. Many benefit-cost analysis practitioners are likely to agree on the importance of using whatever data are available to improve predictions: one oft-stated goal of such analysis is to predict conditions without and with the policy as accurately as possible. If behavioral research helps, then its implications should be considered.

The more difficult issue is how to evaluate welfare and to use the results to select among alternative policies. Typically, benefit-cost analysts assume that each individual is the best judge of his or her own welfare, which means that values should be based on the preferences of the affected individuals. But if people behave in ways that are inconsistent with their own preferences, or their preferences are unstable, how are we to evaluate whether a policy provides net benefits?

Chetty suggests three approaches for addressing this issue that he views as non-paternalistic. His starting assumption is that the utility associated with the experience itself, rather than utility at the point when the decision is made, should be used to evaluate welfare. His first approach involves directly measuring experienced utility, using data on self-reported happiness (i.e., SWB). While Chetty notes the need for further work to address potential biases in these measures, he believes that they provide useful qualitative information. His second approach is based on the concept of sufficient statistics. In this case, the analyst relies on preferences revealed by behavior in contexts where individuals are known to make choices that maximize their experienced utilities. The third approach involves building structural models that characterize how demand varies depending on behavioral biases, then using

the models to infer experienced utility by extrapolating to the case with no bias. Chetty concludes that rather than defaulting to the neoclassical model, researchers should explicitly account for uncertainty related to whether the data are generated by a neoclassical or behavioral model.

While useful, whether these approaches are in fact nonpaternalistic seems open to debate. Applying them requires normative judgments about which measures and which assumptions and models are appropriate. As Robinson and Hammitt (2011) note, the danger is that such approaches “may be abused if revealed 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 how we conduct benefit-cost analysis.” Chetty’s call for explicitly accounting for the associated uncertainty is well in line with established best practices, and will hopefully encourage analysts to be clear about the judgments they make and the implications of these judgments for their analytic conclusions.

4 Articles in the special series

Our special series on “[Ir]rationality, Happiness, and Benefit-Cost Analysis” consists of the seven articles summarized below. These articles suggest that we are making substantial progress in better understanding these issues, while also highlighting many areas where more work is needed.

The series begins with **The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics** by B. Douglas Bernheim (2016).¹ This paper directly responds to the need for a unified framework that addresses the foundational issues raised by behavioral research. It builds on, but in some respects diverges from, Bernheim’s substantial previous work with Antonio Rangel. Bernheim notes that “the dominant approach to those questions has been rooted in the paradigm of revealed preference, which teaches us to infer objectives and welfare (“the good and the bad”) from choices. But behavioral economics teaches us that choices are not always consistent (“the ugly”).” He sets out to develop a practical and unified approach for addressing the implications of behavioral research.

Bernheim begins by discussing the premises that underlie the standard definition of economic welfare and the ways in which behavioral research challenges these premises. After exploring related issues in detail, he arrives at revised

¹ The article by B. Douglas Bernheim is available open access, thanks to generous funding from the John D. and Catherine T. MacArthur Foundation.

premises that selectively defer to choice. Along the way, he explores the use of SWB measures, noting the shortcomings of these measures while indicating that they can provide useful information on choices. He does not, however, advocate using these metrics to measure welfare. More generally, Bernheim's work refines and redefines the terms of the debate, providing new perspectives on the issues introduced above as well as those discussed in the subsequent articles.

Rational Benefit Assessment for an Irrational World, by W. Kip Viscusi and Ted Gayer (2016) addresses the need to develop principles for incorporating behavioral research into benefit-cost analysis. The authors argue that rational choice models, such as expected utility theory, should continue to be the default, because in most contexts consumers are better able to make decisions that affect their own well-being than are analysts or policymakers. They suggest a cautious approach to applying assumptions that deviate from revealed preferences, given that consumers may be harmed by policies that override their preferences and restrict their choices. Viscusi and Gayer also raise several concerns about the use of SWB measures that relate to the construction of the scales and the lack of a theoretical foundation as well as the difficulties inherent in using these measures in analyses that are designed to rely on comparable money metrics.

Viscusi and Gayer indicate that a behavioral transfer test should be used before applying findings from narrow contexts to broader populations. When such a test indicates that behavioral biases are well documented and systematic, policy decisions should focus on achieving the outcome that would result from fully informed and rational decision-making.

In **Do We Need a New Behavioral Benchmark for BCA?**, Jason Shogren and Linda Thunström (2016) provide a somewhat different perspective. They argue that an interval method should be used, which relies on a range of reasonable estimates derived from both the rational choice model and behavioral models given the alternative exchange institutions that underpin choice and welfare (e.g., highly competitive, nonmarket allocation). This approach is necessary because behavioral economics encompasses diverse models that address a wide range of biases and other issues. If the range of benefit estimates is greater than costs, decision-makers can be relatively confident about the welfare effects of accepting or rejecting the policy. If the costs are within the benefits interval, then more discussion will be needed to determine which policy best promotes welfare.

Shogren and Thunström also consider the role that stated preference studies can play in supporting this method. They note that researchers have made substantial progress in accounting for psychological factors, which can help bridge the gap between decision and experienced utility. For example, oaths can be used as a commitment device to aid in retrieving thoughtful responses. To the extent that the gap

cannot be bridged by such research, they conclude that new structural models will be needed to help explain the remaining differences.

In **Cost-Benefit Analysis, Who's Your Daddy?**, Cass Sunstein (2016) reminds us that what we really want to measure is welfare itself; if we could do so, there would be no need for benefit-cost analysis. He notes that WTP estimates have several shortcomings and may not fully reflect the welfare effects of many outcomes such as unemployment. WTP estimates may reflect errors individuals make in forecasting welfare, and may not fully take into account the welfare effects associated with the distribution of the consequences.

While SWB research is promising, Sunstein finds it is not yet sufficiently developed to replace the use of WTP measures. The available studies are not adequate to allow analysts to map policy consequences onto well-being scales. In addition, there is substantial debate regarding the quality, reliability, and applicability of different scales, as well as the extent to which they fully account for all aspects of well-being. He suggests that SWB measures provide useful information, however, and are deserving of more attention in policy decisions.

The following article, **Unequal Life Chances and Choices: How Subjective Well-Being Metrics Can Inform Benefit-Cost Analysis**, by Carol Graham (2016), discusses SWB measurement in more detail. She explores the findings of related research in several contexts, focusing on variation across socioeconomic cohorts. After briefly surveying the literature, she investigates differences between the rich and the poor in the United States and in Latin America. She finds that income groups differ noticeably in how they think about the future, which will affect the value they place on future consequences regardless of whether SWB or WTP measures are used. Because the poor tend to focus more on the present, they may respond less to policies focused on the future than those who are wealthier.

Graham concludes that SWB measures provide important insights for policy-making, but are a complement to rather than a substitute for standard WTP measures. Each approach has limitations, and well-being measures do not resolve issues related to irrational decision-making, inattention to future consequences, or gaps between revealed and stated preferences. They do, however, provide very different insights into the equity of policy impacts across socioeconomic groups.

Bad Air Days: The Effects of Air Quality on Different Measures of Subjective Well-being, by Paul Dolan and Kate Laffan (2016), provides a detailed example of the application of SWB measures, in this case to the effects of air pollution. They test the use of different measures, including evaluative, experiential, and eudemonic approaches, and find a strong (negative) association between air pollutant concentrations and three positive measures of well-being. However, they find that once they control for health status, there is no association with a negative

experiential measure of well-being. They conclude that their findings suggest that focusing solely on health effects may underestimate the impact of air pollution on well-being.

Dolan and Laffan also monetize these measures for the subsample of the population for which income data are available. They find that these valuations vary greatly depending on the well-being measure used. They conclude that further research is needed to develop more robust estimates of these values. In addition, those interested in applying monetized SWB estimates will need to consider the implications of different measures for their results.

The final article, **Behavioral Economics, Happiness Surveys, and Public Policy**, by Matthew Adler (2016), explores the relationships between standard economic theory, behavioral economics, and SWB measures from a more philosophical, normative perspective. Given that behavioral research finds that individuals often fail to comply with the norms of rationality, he notes that we should not expect that individuals' answers to questions about their SWB tell us much about their rational preferences. As a result, relying on such surveys is not a cure for the difficulties that arise from the problems identified by behavioral research.

Adler explores the normative concepts that underlie these conclusions. He defines rationality as decisions that maximize expected utility, including Bayesian updating over time. He then explores conceptions of well-being given this definition, which vary in the extent to which they include both mental states (e.g., cognition, emotion, memories) and other states (e.g., health, relationships, achievements), as well as in whether they are self-defined or externally-defined based on varying views of human capacity or welfare. He concludes that SWB surveys provide information only on selected aspects of well-being and hence do not provide a full accounting of welfare.

4.1 Concluding thoughts

Where does this leave us? The articles in this series make substantial progress in clarifying the contributions of behavioral economics and happiness research to the conduct of benefit-cost analysis, and sharpen our understanding of points of debate and areas where more work is needed. They address both theory and practice, developing a more robust conceptual framework for integrating the findings of behavioral research as well as pragmatic methods for applying these findings in benefit-cost analysis.

When behavioral issues lead individuals to act in ways that appear contrary to their preferences, several authors suggest that well-informed rational preferences should be used in valuing policy outcomes. However, these articles indicate that

analysts should exercise caution in extrapolating findings across contexts and Bernheim raises several issues related to this framing that deserve more attention. In combination, the articles suggest that more detailed and comprehensive review of behavioral research is needed to identify the characteristics of contexts and individuals that make different types of decision-making errors more likely to occur. The articles also emphasize the importance of clearly identifying the assumptions used in the analysis and carefully assessing the implications of related uncertainties.

The series suggests that SWB measures are not yet well enough developed to substitute for WTP measures, with differing views on whether such substitution may eventually be desirable. Some authors raise issues about the underlying framework that may be insurmountable, while others appear more optimistic. In addition, many note that the reporting of SWB is likely to be affected by the same sorts of problems identified by behavioral researchers in other contexts. Most authors, however, appear to agree that SWB measures provide useful information that should be considered along with the results of the benefit-cost analysis.

To some extent, these debates reflect differences in how the goals of benefit-cost analysis are defined. While the main goal is often described as seeking to determine which policy option is the most economically efficient (using the Kaldor–Hicks potential compensation criterion as a guide), in reality its goals are perhaps more pedestrian. The implications of behavioral economics and happiness research are only some of the many challenges analysts face. Typically, gaps and limitations in the available research mean that not all outcomes of interest can be quantified and monetized and that net benefits are uncertain. In addition, decision-makers and stakeholders care about factors other than economic efficiency, including ethical, legal, political, and practical concerns. Thus it seems more consistent with current practice to describe benefit-cost analysis as a well-established and useful framework for providing information for decision-making, rather than as a normative guide for choosing which policy should be implemented. It helps those engaged in the policymaking process anticipate outcomes that might otherwise be unexpected, understand the preferences of those affected, and develop evidence to support the decision. To the extent that behavioral economics and happiness research aid in achieving these goals, incorporating the findings appears useful.

We hope you enjoy reading the articles that follow, and look forward to your contributions to further discussions of these issues in the pages of *JBCA*.

Acknowledgments: I thank JBCA Editors-in-Chief Glenn Blomquist and William Hoyt for their encouragement and wise advice throughout the development of this series, and JBCA Managing Editor Mary Kokoski for guiding me through the submission and review system. I also thank the editors and authors for their

helpful comments on an earlier draft of this introduction. We are deeply grateful to the authors for the substantial time and thought they devoted to developing their articles for this series, and to the peer reviewers for their careful and detailed reviews.

References

- Adler, Matthew D. (2016). Behavioral Economics, Happiness Surveys, and Public Policy. *Journal of Benefit-Cost Analysis*, 7(2), in this issue.
- Bernheim, B. Douglas (2016). The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics. *Journal of Benefit-Cost Analysis*, 7(2), in this issue.
- Brennan, Timothy J. (2014). Behavioral Economics and Policy Evaluation. *Journal of Benefit-Cost Analysis*, 5(1), 89–109.
- Chetty, Raj (2015). Behavioral Economics and Public Policy: A Pragmatic Perspective. *American Economic Review: Papers & Proceedings*, 105(4), 1–33.
- Cutler, David M., Jessup, Amber, Kenkel, Don & Starr, Martha A. (2015). Valuing Regulations Affecting Addictive or Habitual Goods. *Journal of Benefit-Cost Analysis*, 6(2), 247–280.
- Dolan, Paul & Kahneman, Daniel (2008). Interpretations of Utility and their Implications for the Valuation of Health. *Economic Journal*, 118, 215–234.
- Dolan, Paul & Laffan, Kate (2016). Bad Air Days: The Effects of Air Quality on Different Measures of Subjective Well-being. *Journal of Benefit-Cost Analysis*, 7(2), in this issue.
- Fujiwara, Daniel & Campbell, Ross (2011). *Valuation Techniques for Social Cost-Benefit Analysis: Stated Preference, Revealed Preference and Subjective Well-Being Approaches*. United Kingdom: Prepared for HM Treasury.
- Graham, Carol (2016). Unequal Life Chances and Choices: How Subjective Well-Being Metrics Can Inform Benefit-Cost Analysis. *Journal of Benefit-Cost Analysis*, 7(2), in this issue.
- Hammitt, James K. (2015). Implications of the WTP–WTA Disparity for Benefit–Cost Analysis. *Journal of Benefit-Cost Analysis*, 6(1), 207–216.
- Jin, Lawrence, Kenkel, Don, Liu, Feng & Wang, Hua (2015). Retrospective and Prospective Benefit-Cost Analyses of US Anti-Smoking Policies. *Journal of Benefit-Cost Analysis*, 6(1), 154–186.
- Kahneman, Daniel & Sugden, Robert (2005). Experienced Utility as a Standard of Policy Evaluation. *Environmental and Resource Economics*, 32(1), 161–181.
- Kahneman, Daniel & Tversky, Amos (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263–291.
- Kahneman, Daniel & Tversky, Amos (Eds.) (2000). *Choices, Values, and Frames*. New York: Cambridge University Press.
- Knetsch, Jack L. (2015). The Curiously Continuing Saga of Choosing the Measure of Welfare Changes. *Journal of Benefit-Cost Analysis*, 6(1), 217–225.
- Layard, Richard (2010). Measuring Subjective Well-Being. *Science*, 327(5965), 534–535.
- Robinson, Lisa A. & Hammitt, James K. (2011). Behavioral Economics and the Conduct of Benefit-Cost Analysis: Towards Principles and Standards. *Journal of Benefit-Cost Analysis*, 2(2), 1–51. Reprinted in Scott Farrow and Richard O. Zerbe, Jr. (Eds.), *Principles and Standards for Benefit-Cost Analysis* (pp. 317–363). Cheltenham: Elgar Publishing.

- Shogren, Jason F. & Thunström, Linda (2016). Do We Need a New Behavioral Benchmark for BCA? *Journal of Benefit-Cost Analysis*, 7(2), in this issue.
- Sunstein, Cass R. (2016). Cost-Benefit Analysis, Who's Your Daddy? *Journal of Benefit-Cost Analysis*, 7(2), in this issue.
- Thaler, Richard H. (1992). *The Winner's Curse: Paradoxes and Anomalies of Economic Life*. Princeton: Princeton University Press.
- Viscusi, W. Kip (2015). Reference-Dependent Effects in Benefit Assessment: Beyond the WTA-WTP Dichotomy and WTA-WTP Ratios. *Journal of Benefit-Cost Analysis*, 6(1), 187–206.
- Viscusi, W. Kip & Gayer, Ted (2016). Rational Benefit Assessment for an Irrational World: Toward a Behavioral Transfer Test. *Journal of Benefit-Cost Analysis*, 7(2), in this issue.