

Positive v. Normative Justifications for Benefit-Cost Analysis

James K. Hammitt

Harvard University (Center for Risk Analysis)
718 Huntington Ave., Boston, MA 02115 USA

Toulouse School of
Economics (LERNA-INRA) 21, allée de Brienne,
31000 Toulouse FRANCE

tel: +1 617 432 4343, +33 (0)5 61

12 86 22, email: jkh@harvard.edu

July 2012

Abstract What is the rationale for benefit-cost analysis (BCA)? The answer is critical for

determining how BCA results should be interpreted, their implications for policy, and how BCA should be conducted. There are at least two possible bases for justifying BCA, positive and normative. The positive basis is that BCA identifies policy changes that satisfy the Kaldor-Hicks compensation test, so that those who benefit could hypothetically compensate those who are harmed. The normative basis is that BCA identifies social improvements, e.g., by approximating a utilitarian calculus or promoting more consistent decisions by protecting against cognitive error. When human behavior differs from that which is assumed in standard economic models, the justifications may conflict. Individuals whose behavior differs from the models may disprefer a change in circumstances that normative models predict they should prefer. The positive justification is consistent with respect for individual autonomy and provides clarity about methodological choices in the analysis but can require endorsing cognitive and behavioral errors that individuals would wish to avoid. The normative justification implies rejecting policies that the population prefers and requires determining what preferences are normatively acceptable.

JEL classification: D61, D81, H40, Q50 Keywords: benefit-cost analysis, behavioral economics

Introduction

Behavioral-economic and related research has revealed that people may systematically behave in ways that are inconsistent with standard economic models. This raises questions about economic-welfare theory and the role, conduct, and interpretation of benefit-cost analysis (BCA). In short, what is the justification (or rationale) for BCA? What questions does it address, and why are the answers relevant to policy choice?

In general terms, BCA provides information about whether a policy change is likely to improve social well-being, conceived as an aggregate of the well-being of individuals comprising the society. Two classes of justification exist, positive and normative. The positive justification is that BCA seeks to determine whether the policy change in question satisfies the Kaldor-Hicks compensation test. It leaves unanswered the normative question of whether satisfying the compensation test is either necessary or sufficient for a policy to constitute a social improvement and also the extent to which other factors such as equity should offset or outweigh the results of BCA.

There are several normative justifications: that utilitarianism is an appropriate framework for policy choice and BCA approximates the utilitarian calculus; that BCA protects against cognitive mistakes in policy choice; and that BCA promotes consistent policy choice which enhances individuals' well-being on average.

The existence of systematic differences between human behavior and standard economic theory can drive a wedge between the positive and normative justifications. In particular, a policy that BCA suggests will satisfy the Kaldor-Hicks compensation test (i.e., a policy for which a set of transfer payments exists such that everyone in the population views himself at least as well off with the policy and transfer payments than without) may not deliver: people whose behavior is inconsistent with standard models may judge themselves to be worse off. Conversely, a policy that everyone prefers may not satisfy the compensation test. The tension between positive and normative perspectives is nicely posed by the parable "Trouble in Happyville" (Portney 1992): Imagine you are the Director of Environmental Protection for the town of Happyville. There is a naturally occurring contaminant in the town's drinking water that all of the residents believe is carcinogenic

and may account for the towns' above-average cancer rate. Each resident is willing to pay \$1,000 to cover the cost of treatment that will eliminate the contaminant. You have consulted with the world's top risk analysts and each has reported that, while one can never be certain a particular substance does not cause cancer, each would stake her professional reputation on the conclusion that this contaminant is benign. You have repeatedly and skillfully communicated these judgments to the citizenry, but each of them nevertheless prefers to spend the money to treat the water. What should you do? If you call for the water to be treated, you are knowingly denying every resident the other benefits he could achieve with \$1,000 but every resident will believe himself to be better off. If you reject the treatment option, you are knowingly imposing a policy that every resident believes is contrary to his well-being. This essay highlights the extent to which questions about how to conduct and

interpret BCA are dependent on the unresolved question of its fundamental justification or rationale. This dependence is illustrated in the context of environmental, health, and safety policy.

The paper is organized as follows. The following section describes some of the background, possible justifications for BCA, and proposals for reconciling the tension that arises if behavior deviates from standard economic models. The next section describes key concepts of BCA that are particularly relevant to this question: the relationship between choices, preferences, and well-being, and the distinction between predicting and evaluating the desirability of policy consequences. The subsequent section examines several areas in which individual behavior seems to conflict with normative models, including the role of information, non-proportionality between willingness to pay and risk reduction, ambiguity aversion, discounting future consequences, and the divergence between measures of willingness to pay for an improvement and willingness

¹ Salanie and Treich (2009) generalize Portney's example to permit the regulator to choose among multiple control levels and to allow consumers to adjust their consumption of town water. They suggest that a paternalistic regulator should clean the water more than would otherwise be optimal whenever residents misperceive its quality. When consumers overestimate the risk, they will incur costs in avoiding consumption of town water that can be reduced by over-control; when they underestimate the risk, they will consume more than the optimal level.

to accept compensation to forgo the improvement. The final section summarizes how the choice between positive and normative justifications has implications for the conduct and interpretation of BCA and concludes that the choice of justification is related to the larger question of the role of government: whether the government should attempt to serve the public interest as it is currently perceived by citizens or to provide leadership that is ultimately in the public's interest, even if that conflicts with public opinion.

Justifications for BCA

There are multiple justifications for BCA, i.e., rationales for why its results should influence policy choice. Some justifications are positive: they assert that BCA answers a positive question, and implicitly that the answer to this question is normatively significant for policy choice. Other justifications are normative: they assert that BCA helps identify policies that are socially beneficial but do not specify a well-defined positive question that BCA seeks to answer. Different rationales have different implications for how BCA should be conducted and how its results should be interpreted. The most prominent positive justification derives from the text-book description

of BCA as a method that identifies policy changes that satisfy the Kaldor-Hicks compensation test. The compensation test asks whether, if monetary transfers could be made without cost, those who benefit from a policy change could compensate those who are harmed so that everyone would judge himself better off with the policy change and transfer payments than without. This is a question that can, at least in principle, be answered using scientific methods. This positive justification leaves unanswered the normative question of whether passing the Kaldor-Hicks compensation test is either a necessary or sufficient condition for the policy change to constitute a social improvement

Nyborg (2012, p. 15) makes a similar distinction: “Most textbooks in cost-benefit analysis assume, implicitly or explicitly, that the purpose of the analysis is to rank projects according to their social preferability. My starting point is different. I assume, instead, that the primary role of an economic project analysis is to provide factual input to a democratic decision process.” She explains that positive and normative rationales require different sorts of analyses: a normative analysis is necessarily based on normative premises and cannot be objective; a positive analysis must be as objective as possible if it is to prove useful to decision makers holding different normative premises.

³when transfer payments are not made. The justification is also silent on the role that other factors, such as effects on equity, equality of opportunity, and fairness, should have in policy choice. There are several normative justifications for BCA. They have in common an

⁴assertion that BCA identifies policies that improve social well-being (where improvement must be defined in some way that is external to BCA) but the truth of this assertion cannot be empirically verified. One normative justification is derived from utilitarianism in which the objective is to maximize the sum of well-being in the society and it is assumed that well-being can be measured using standard economic concepts such as compensating and equivalent variation. This justification leads naturally to the suggestion that the monetary values of benefits and costs should be weighted depending on whether they fall on rich or poor, given the intuition that marginal utility declines with wealth. Whether total population utility is increased cannot be verified without a method to measure individual utility in an interpersonally comparable fashion, which does not yet exist. Other normative justifications are motivated by pragmatism. One such claim is

that BCA is a practical method that approximates the result of an ideal but impractical measure of social welfare (Adler and Posner 2006). A related justification is that BCA helps promote consistent decision making by avoiding random errors and protecting against cognitive mistakes that can arise when a decision maker tries to evaluate a policy change by holistic judgment (Sunstein 2000). Neither claim can be empirically verified without the ability to calculate the ideal measure of social welfare (for Adler and Posner) or to identify the better policy decision (for Sunstein).

³ The concept of a positive justification for BCA, as used here, is distinct from Abramowicz's (2004) proposal for "predictive cost-benefit analysis." Abramowicz suggests that information markets could provide information about whether a policy is likely to satisfy a Kaldor-Hicks compensation test. Specifically, such markets could predict whether a retrospective but otherwise conventional BCA of the policy in question will find that it yields net benefits. Abramowicz's proposal does not address the questions about how BCA should be conducted and its relevance for policy that are addressed here.

⁴ See Kahneman and Sugden (2005) for discussion of the possibility to use self reports of experienced utility as such a measure.

An alternative consistency claim asserts that, if policies that maximize net benefits are routinely chosen, everyone will be better off in the long run because those who benefit and those who are harmed will tend to vary across decisions (Boardman et al. 2006). The claim that winners and losers will vary across decisions is clearly true, but it seems doubtful that the expected gain from using BCA over a set of decisions is distributed uniformly over the population. Because the monetary values of beneficial and harmful consequences tend to increase with income, BCA will systematically favor the interests of rich over poor, compared with equal weighting of individuals' utilities. Moreover, the strength of these pragmatic claims obviously depends on the alternatives with which BCA is compared. Proponents of the positive justification tend to view BCA as providing

information that is useful to decision makers, but to recognize that factors in addition to those incorporated in BCA (such as equity) are also important, and so satisfying the compensation test is neither necessary nor sufficient for policy choice. In contrast, normative justifications seem to make a more explicit claim that a finding of net benefits should be determinative, because BCA, on this interpretation, is sufficient to distinguish between policies that do and do not improve social well-being.

A statement by eleven leading regulatory economists (Arrow et al. 1996) is largely consistent with the positive justification ("Most economists would argue that economic efficiency, measured as the difference between benefits and costs, ought to be one of the fundamental criteria for evaluating proposed environmental, health, and safety regulations"), although it also invokes a pragmatic justification (BCA "can provide an exceptionally useful framework for consistently organizing disparate information, and in this way, it can greatly improve the process and, hence, the outcome of policy analysis").

Both positive and normative justifications, as described here, assume that BCA is conducted to inform some impartial benevolent despot, often called a social planner. Sugden (2011) criticizes this framing because no such decision maker exists. He suggests that BCA should be addressed to actual members of a society who wish to coordinate their behavior to mutual advantage. Nyborg (2012, p. 15) adopts a similar perspective: "I will assume that several decision-makers, whose ethical and/or political views differ, take part in [the policy choice]. Consequently, the purpose of the analysis is not to determine one final ranking of projects, but rather to enable each decision-maker to arrive at a wellfounded ranking in accordance with her normative views."

Nyborg (2012, p. 19) suggests a position like this is common in BCA texts: “the most common conclusion still seems to be that CBA can still be used normatively – although with care. The main problem with this view is, as I see it, that it is usually not clear what it means to us e ca re.”

Two alternative justifications for BCA that accommodate behaviors that violate standard economic models have been recently offered. Sugden (2005) suggests that the objective of BCA should be to estimate the type of economic surplus that markets produce. In particular, if preferences are context-dependent, individuals’ choices and the resulting economic surplus will depend on the decision context. Sugden suggests that the surplus that should be incorporated in BCA should be the surplus that would result if perfectly competitive markets were available. Typically, this will be the surplus “at the last moment at which trade is possible” (p.148), e.g., immediately before consumption or at the moment of commitment to production of a custom good. This “closest market analog” approach seems to be an alternative positive justification: it seeks to predict the surplus that would be created if ideal competitive markets were available and people behaved as they do, regardless of whether that behavior promotes their long-run wellbeing. Whether predictions of behavior in the simulated ideal markets are empirically verifiable is not clear, however. Smith and Moore (2010) question how these predictions can be made without specifying what alternative to conventional coherent preferences to use in making them.

Bernheim and Rangel (2009) propose to base welfare analysis on actual choices and to accommodate choices that are inconsistent with a standard utility function by recognizing that an individual’s choices depend on both the objects between which he is choosing and “ancillary conditions” that are known to the individual but not to the analyst. They suggest that welfare analysis be conducted using an “unambiguous choice relation,” which is the most informative characterization of preferences that is consistent with a list of an individual’s observed choices. Because the ancillary conditions associated with a policy are unknown to the analyst, this framework accepts that there may be ambiguity in determining whether an individual is benefitted or harmed by a policy intervention. This approach is also a positive justification: it seeks to implement a

⁶revealed-preference criterion without imposing assumptions about whether observed behavior is actually in the individual's best interests.

Key Concepts

This section describes some key concepts of BCA that are critical to assessing the implications of alternative justifications. These include the relationship between an individual's choices, preferences, and well-being, and the distinction between predicting the consequences of a policy and evaluating their desirability.

Choices, Preferences, and Well-being

⁷BCA and similar approaches to evaluating policy confront two issues: how to evaluate individual well-being and how to evaluate policies that improve some people's well-being while reducing others' (in comparison with either continuing the status quo policy or with alternative policies that would have improved the others' well-being by more than the candidate policy). Conventional BCA assumes that the individual is the best judge of how changes in circumstances will affect his well-being (consumer sovereignty) and measures change in well-being by each individual's compensating variations, i.e., willingness to pay (WTP) for an improvement and willingness to accept compensation (WTA) for a harm. Other concepts of well-being include capabilities (Sen 1992), opportunities (Sugden 2004), quality-adjusted life years for health and longevity (Gold et al. 1996), and conceptions of a good or moral life proffered by many religions.

In identifying well-being with preferences, conventional BCA relies on the standard economic assumption that people act to maximize their own well-being (subject to the constraints they face). This assumption underlies the first welfare theorem: when markets are perfectly competitive, any market equilibrium is Pareto efficient (resources cannot be reallocated to improve someone's well-being without reducing someone

⁶ Smith and Moore (2010) suggest this proposal is unworkable because it requires many observations of people's choices in different contexts.

⁷ BCA can also be conducted using equivalent variations, i.e., WTA to forgo an improvement and WTP to prevent a harm.

else's). If people do not maximize their own well-being (given available income and market prices), the theorem fails.

Behavioral-economic research provides evidence that people often behave in ways that do not maximize well-being as it is represented by individual-specific utility functions like those incorporated in conventional economic models. Humans often seem to evaluate changes rather than positions, to evaluate changes as proportions of some reference rather than absolute magnitudes, and to be influenced by the way a choice is described or “framed” (e.g., Kahneman 2011, Kahneman and Tversky 1979, 2000, Kahneman et al. 1982, Tversky and Kahneman 1991). Beshears et al. (2008) assert that choices are jointly determined by a mixture of normative preferences that “represent the agent’s actual interests” and other factors including analytic errors, myopic impulses, inattention, passivity, and misinformation. They identify five factors that increase the likelihood that individual choices diverge from normative preferences: passive choice, complexity, limited personal experience, third-party marketing, and intertemporal choice.

It is well-recognized that people sometimes make consequential mistakes (as well as whimsical choices and flippant survey responses) that do not reflect their interests. Some people, in some contexts, may lack the requisite cognitive capability to make sensible choices (e.g., children, Alzheimer’s patients). Hence the standard for BCA is that choices should reflect individuals’ “well-informed preferences.” But how to determine which choices reflect well-informed preferences and which do not is a difficult question, especially if one accepts that some choices that violate normative models of choice may nevertheless reflect well-informed preferences.

One response is that a choice or behavior is mistaken if the individual, on being confronted with its implications, recognizes it as a mistake. Famously, when Maurice Allais asked L. James Savage to choose between the two paired lotteries of the Allais paradox, Savage evidently followed the common pattern, thus violating his axioms of expected utility theory. Once Allais pointed this out to him, Savage decided he had made a mistake and changed his answers. More recently, Gneezy et al. (2006) report experimental evidence in which samples of individuals value a lottery less than comparable samples value its worst outcome. They note that such violations do not occur when the same individuals value both the lottery and its worst outcome, when the

comparison is transparent. Whether an individual recognizes a choice or pattern of choices as undesirable may depend on his education or training. For example, Binmore (2009, pp. 23-24) suggests “ Just as it is easier to predict how educated kids will do arithmetic, so the spread of decision theory into our universities and business schools will eventually make it easier to predict how decisions get made in economic life.” It may be difficult, however, to distinguish between legitimate education and a form of training or “brainwashing” that conditions individuals to a specific behavior .

Another approach is to investigate the extent to which choices are constrained by factors that are not incorporated in standard economic models, such as constraints on cognitive capacity and decision-making energy (Smith and Moore 2010). If these constraints can be identified and modeled, the preferences that are inferred may be more consistent with a standard utility function. An alternative position is that choices that violate a normative model cannot reveal well-informed preferences.

8Do behavioral deviations from the predictions of standard economic models of normative behavior reflect cognitive errors or are standard economic models oversimplified, ignoring important and legitimate concerns? Similarly, do well-known differences in risk perception between experts and laypeople reflect naïve public evaluation or inadequacies of expert models that give insufficient attention to attributes other than probability and severity of consequence (Starr 1969, Slovic 1987, 2000)? Surely, both answers are correct: some deviations from expert models are due to error, and some reflect limitations of the models which are, by design, simplified representations. A more useful question is which of the deviations between behavior and economic models reflect errors that individuals would wish to correct, were they aware of them, and which reflect inadequacies of standard models in describing normatively appropriate behavior?

8 Gowdy (2007) argues that experimental economics and neuroscience provide strong evidence that many assumptions of standard economic models are false (e.g., that “money is a universal substitute for anything” and that “preferences are stable and [full y] self-regarding.” He advocates that BCA should be “based on scientific evidence about how humans actually behave, not on armchair theorizing about how people should act in the stylized world of Walrasian general equilibrium.”

9Differences between human behavior and standard economic models drive a wedge between positive and normative justifications for BCA. If people always behaved in accordance with standard economic theory, then any policy that satisfied the KaldorHicks compensation test would provide an opportunity for enhancing everyone's well-being (i.e., it would expand the "social pie") and the central question about the use of BCA would be how to balance efficiency (as measured by aggregate net benefits) against distribution of well-being within society and other concerns such as liberty and responsibility (Saint-Paul 2011). In effect, the questions would be about how to divide the pie (e.g., under what conditions should transfer payments be required) and, since transfers are not costless, under what conditions does an adverse distributional effect outweighs a beneficial efficiency effect (Okun 1975). If people do not always behave in accordance with standard theory, then policies combined with compensation payments that are predicted to yield Pareto improvements may not deliver; affected individuals may not perceive themselves to be better off.

Predicting and Evaluating Consequences

Conventional BCA requires two types of evaluation: predicting the consequences of alternative policies and assessing the desirability of the consequences. Prediction of policy consequences is a positive exercise; the task is to make the most accurate prediction possible. Predictions need not be limited to point estimates; uncertainty can be represented by a probability distribution over possible consequences. Many of the tools used to predict consequences are scientific models. In the context of environmental, health, and safety policies, these include economic and natural-science models used to predict how regulations or other policies influence firm and individual behavior and ultimately environmental quality and human exposures to hazards and health risks.

Note that the consequences of a policy often depend on behavior as well as on natural laws. For example, energy use in lighting, transport, and other services is

9 If transfers can be made more efficiently through a general taxation and welfare system rather than by accompanying every policy intervention with specific transfers, then the tax and welfare system might be adjusted to compensate for the effects of a portfolio of policy interventions.

influenced not only by technical improvements in efficiency but also by the rebound effect through which consumption of the service increases as its marginal cost falls (Tsao et al. 2010). Similarly, the health effects of ambient air pollution depend on people's behavior such as limiting outdoor activity on days with high pollution; as air quality is improved, people may reduce their self-protective behavior leading to less reduction in exposure and smaller health gains than if their behavior did not respond. In his wellknown paper, Peltzman (1975) provided evidence that mandated technical improvements in automobiles provided smaller reductions in injuries than anticipated, because drivers took some of the benefits of safer cars in the form of reducing their efforts to drive carefully. Even the seemingly technological costs of producing pollution-control equipment depend on workers' behavior and preferences over income, working conditions, and leisure time.

Prediction is a positive exercise that is, in principle, subject to empirical verification. If people's behavior differs from conventional economic models, these deviations should be taken into account when predicting consequences – there is no rationale for knowingly mispredicting consequences. For this purpose, descriptive models that accurately forecast the consequences of policy are more useful than normative models that explain how people ought to respond to the policy.

In contrast to prediction, evaluating the desirability of alternative consequences (or probability distributions over consequences) is a normative problem. In conventional BCA, the preferences of the affected population are deemed to govern and the standards that preferences must satisfy to be deemed acceptable are modest, typically consisting of basic properties of coherence (e.g., transitivity) so that choices can be represented as maximizing a utility function (*de gustibus non est disputandum*, i.e., analysts should accept preferences as they are, since they cannot be judged objectively right or wrong). In evaluating risks, conventional BCA assumes preferences are consistent with maximizing the expected value of a utility function, which is a more demanding characterization of coherence that is widely but not universally accepted (for critiques, see Allais 1953, Slovic and Tversky 1974, Machina 1987, Cohen and Jaffray 1988, Manski 2009).

Note that while the use of individual preferences to evaluate policies is normative, the exercise of measuring individual preferences is positive: the analyst seeks to describe

individuals' preferences as accurately as possible. BCA is populist: the preferences that determine desirability of alternative consequences are those of the affected population, not those of an analyst, bureaucrat, or other expert (at least under the positive justification).

Monetary values of consequences, such as reductions in health risks, are ascertained using revealed- or stated-preference approaches. Revealed-preference methods infer people's preferences on the assumption that they prefer the choices they make to the available alternatives. (Revealed-preference methods are also applied in laboratory settings, in which case there is the additional question of how accurately behavior in the laboratory corresponds to behavior in the field.) Stated-preference methods rely on people's statements about which options they prefer, usually in surveys. In many cases, the preferences inferred from either approach appear to be inconsistent with standard economic models. For example, purchasers of lower cost, less efficient air conditioners and refrigerators seem to reveal personal discount rates of 20 percent or more, substantially higher than apparent borrowing costs (Hausman 1979, Gately 1980). As another example, stated-preference estimates suggest that people require two to ten times as much compensation to forgo a gain as they would be willing to pay for the gain, despite theoretical models that suggest these values should be (in most cases) nearly equal (Horowitz and McConnell 2002).

Differences between human behavior and standard economic models not only affect the rationale for BCA but can also affect estimates of parameters that are required inputs. For example, if people evaluate risks using some form of probability weighting (as in prospect and rank-dependent-expected-utility theories; Kahneman and Tversky 1979, Tversky and Kahneman 1992, Quiggin 1993), then estimates of rates of substitution between money and changes in risk may be biased (Bleichrodt and Eeckhoudt 2006). As another example, if behavior is affected by loss aversion (Tversky and Kahneman 1991), estimates of the price elasticity of a good may depend on whether the price increases or falls compared with some reference value (Putler 1992 found the price elasticity of eggs was more than twice as large for a price increase than for a price decrease). To predict the consequences of a policy change, models that incorporate probability weighting and loss aversion should be used (when these phenomena are

present), but to evaluate well-being it may be desirable to adjust the estimates of rates of substitution and price elasticities that come from these behaviors to better estimate normative preferences that “represent the agent’s actual interests” (Beshears et al. 2008).

Differences between Positive and Normative Perspectives

Differences between standard economic models and human behavior are numerous. In some cases, these differences appear to reflect cognitive or behavioral errors that people would wish to avoid if they were aware of them. These include sensitivity of decisions to framing, many differences between willingness-to-pay and willingness-to-accept values, hyperbolic rather than exponential discounting of future consequences, nonlinear response to changes in probability, and perhaps distinguishing consequences according to whether they result from acts of omission or commission, loss aversion, and ambiguity aversion. In other cases, differences between behavior and economic models reflect oversimplified models and can be reconciled by adopting more realistic models. As an example, the assumption of perfect competition implies that consumers have full and complete information about available market goods and that they can collect and process new information at zero cost. If this were true, health and safety regulations on consumer products, including food, medicines, and motor vehicles, would be unnecessary at best and harmful if they were ever binding on someone’s choice. More realistic models that recognize that collecting and processing information are costly imply that consumers are rationally ignorant about many aspects of consumer products and that it is efficient for them to delegate some decisions about health and safety standards to a government or other authority (Downs 1957).

¹⁰Consider some examples of how these differences between behavior and economic models pose questions about how to conduct and interpret BCA. Specifically, consider the role of information, non-proportional response of WTP to risk reduction, ambiguity aversion, hyperbolic discounting, and the divergence between estimates of WTP and WTA.

¹⁰ For further discussion of the evidence concerning differences between individual behavior and normative models with a focus on valuing health risks, discounting, and other-regarding preferences, see Robinson and Hammitt (2011a, b).

Information

In the standard economic model, provision of (accurate) information is never harmful and is possibly beneficial. This follows because people can adapt their behavior to better meet their goals when they have more accurate information about factors that influence the consequences of alternative decisions. In the worst case, information can be ignored and so it causes no harm.

In practice, individuals may not process information optimally and may be misled. For example, they may over-emphasize attributes that appear particularly salient and under-emphasize others, notably the probability that a bad outcome will occur (Sunstein 2002, Sunstein and Zeckhauser 2011). In some cases, individuals may perceive a difference in risk between alternatives even when scientific studies support the hypothesis of no difference. Possible examples include the common belief that synthetic chemicals are more likely to be carcinogenic than naturally occurring chemicals, even though equal fractions (about half) of all synthetic and natural chemicals that have been tested in conventional high-dose rodent bioassays show evidence of carcinogenicity (Gold et al. 2002) or that genetically modified foods are less healthful than those produced through selective breeding (Cantley and Lex 2010). Firms can exploit these tendencies through advertising, making some attributes more salient and others less, hence biasing consumer behavior away from their normative interests. If information provision can be harmful, there may be cases in which it should be suppressed, or at least not publicized. This issue is recognized in the law of evidence in U.S. and other courts, where a factor in determining whether certain evidence is admissible is the balance between its probative and prejudicial value. Similarly, regulations exist to limit advertising, e.g., of tobacco on television and to children. In another case, the US Department of Agriculture (USDA) prohibited a private beef producer from testing its cattle for BSE (mad-cow disease). The producer hoped to reassure potential importers in Japan and South Korea that its beef was BSE-free, but USDA argued the test was of no value and, presumably, if this firm were permitted to test

¹¹ other firms would be disadvantaged and/or there would be pressure to vastly increase the federal testing program.

Non-proportionality of WTP to Risk Reduction

Normative economic models of decisions about risky outcomes (i.e., expected utility) imply that WTP for a small reduction in the probability of suffering an adverse health or other consequence should be nearly proportional to the magnitude of the probability change. Positive models, such as prospect and rank-dependent-expected utility theories (Kahneman and Tversky 1979, Tversky and Kahneman 1992, Quiggin 1993), yield the same result except under rather implausible conditions where an individual's response may be highly nonlinear for probabilities in the relevant range (Hammit 2000, Corso et al 2001). In contrast, stated-preference studies usually find that WTP is substantially less than proportional to the probability change, implying sharply decreasing marginal WTP as the risk reduction increases.

In the normative model, the primary factor leading to non-proportionality is the income effect; as WTP becomes large relative to the individual's assets, the rate at which he is willing to pay for risk reduction should decrease because the marginal utility of forgone consumption increases. In addition, for risks of death, the rate at which he is willing to pay should also decrease because of the dead-anyway effect (Pratt and Zeckhauser 1996): The expected opportunity cost of spending on risk reduction is the probability-weighted average of the opportunity costs conditional on avoiding and suffering the adverse effect. If the incremental opportunity cost of resources is smaller for a bequest than for consumption while living, then as the probability of death increases, the expected opportunity cost decreases and WTP increases.

¹¹ BSE testing in beef cattle may be of little or no value because BSE has a substantial incubation period and it may not be possible to detect evidence of infection in beef cattle at the age of slaughter. The producer sued USDA and ultimately the court ruled USDA could not prevent it from conducting private tests. See "Federal judge rules Creekstone can test for BSE," *High Plains Journal*, April 5, 2007 (www.hpj.com/archives/2007/apr07/apr9/FederaljudgerulesCreekstone.cfm, accessed July 25, 2012) and "Creekstone Farms Premium Beef v. USDA," Civil Action No. 060544 (JR), US District Court for the District of Columbia, filed March 29, 2007.

Among stated-preference studies that estimate WTP for health-risk reductions of different magnitudes, most find that WTP is substantially less than proportional to the probability change. Hammitt and Graham (1999) tried to identify all stated-preference studies that estimated WTP for a numerically specified reduction in risk of some adverse health effect published between 1980 and 1998. Of the 14 studies they identified that provided enough information to test for the sensitivity of WTP to the magnitude of risk reduction, WTP increased with risk reduction in a statistically significant manner in 11 studies, but never in proportion to risk reduction. More recently, some studies have found that estimated WTP is proportional to probability change when appropriate visual aids are used to communicate risk (Corso et al. 2001, Hammitt and Haninger 2010) but other studies using similar methods have continued to find that WTP varies significantly less than proportionately (e.g., Alberini et al. 2004, Haninger and Hammitt 2011).

¹²Stated-preference studies that find WTP increases less than proportionately to risk reduction imply that marginal WTP for risk reduction decreases sharply as the size of the reduction increases. For example, Alberini et al. (2004) estimated mean annual WTP for a 1 in 1000 reduction in the risk of dying in the next decade as \$483 and mean WTP for a 5 in 1000 reduction as \$770. Taken at face value, these estimates imply respondents value an initial 1 in 1000 increase in the probability of surviving the next decade at \$483 per year but value four additional increases at an average of only \$72 per year ($= [770 - 483] / 4$), 15 percent as much as they value the initial risk reduction. If used in BCA, these values suggest the population would strongly prefer two policies that each reduced their mortality risk over the next decade by 1 in 1000 (valued at $\$966 = 2 \times \483) to a single policy that reduced risk by 5 in 1000. Such a preference might be consistent with having a strong preference for taking action to reduce a risk with relatively little concern for the efficacy of the action, but it is inconsistent with the standard economic model that evaluates actions by their consequences.

¹² This inference assumes that having received the first risk reduction would not decrease WTP for the second; an alternative interpretation is that WTP for the second risk reduction is sharply less than for the first. The initial interpretation is supported by results showing WTP for a risk reduction is independent of small changes in the level of the initial risk (Hammitt and Haninger 2010, Haninger and Hammitt 2011).

Alternatively, if near-proportionality of WTP to reduce risk is taken to be normatively required and the non-proportional results of many stated-preference studies are attributed to cognitive errors on the part of survey respondents or to problems with study design that lead respondents to reject the described scenario, then the question arises how to estimate the rate at which people are willing to pay for risk reduction. Using the Alberini et al. study, the marginal rate of substitution between income and mortality risk is \$4.8 million per life saved using the smaller risk reduction and \$1.5 million using the larger risk reduction. Which estimate (or what alternative estimate) should one apply in BCA?

The problem of non-proportionality of WTP to risk reduction is but one example of the problem of apparently inadequate sensitivity to scope (i.e., to the magnitude of the good) that has been observed in stated-preference studies for many years, perhaps most famously when these methods were used to help estimate damages caused by the 1989 Exxon-Valdez oil spill in Alaska (Diamond and Hausman 1994, Hanemann 1994). Unlike the case of WTP to reduce the probability of an adverse consequence, theory provides much less guidance about how much WTP should increase with improvements in environmental quality, or more generally with increases in the quantity or quality of a good. In another health context, Hammitt and Haninger (2007) estimated that WTP to reduce the risk of acute illness from food-borne pathogens was implausibly insensitive to the severity and duration of the illness: The marginal rate of substitution between money and risk was estimated as \$8,300 per expected case for a one-day episode of mild illness and \$16,100 per expected case for a week-long episode requiring hospitalization.

The non-proportionality of estimated WTP to risk reduction may be reconciled with economic theory by adopting a more refined theoretical model. As noted, if the problem is one of cognitive constraints as suggested by Smith and Moore (2010), it may be possible to design elicitation methods to relax these constraints (e.g., certain visual aids may help ameliorate problems in risk communication; Corso et al. 2001). If deviations can be well-explained by a descriptive model, such as non-linear probability evaluation, it may be possible to adjust elicited values to correct for this distortion (Bleichrodt et al. 2001). Alternatively, respondents may be valuing a change in risk that is not the risk change specified in the survey, but rather a risk change based on rationally

updating their prior beliefs about the hazard with information provided in the survey. These updated beliefs do not vary as much between respondents as do the risk reductions presented in the survey (Viscusi 1985, 1989). In this case, it may be possible to elicit the risk reduction that each respondent is valuing to determine his rate of substitution between money and risk.

Ambiguity Aversion

It has long been recognized that people tend to prefer situations in which the probabilities of the possible outcomes are known to situations in which the probabilities are unknown. Knight (1921) described the first situation as one of risk and the second as one of uncertainty. Ellsberg (1961) provided the seminal paper.

From a Bayesian perspective, the distinction between Knightian risk and uncertainty appears meaningless. From this perspective, probability is always personal; different people can attach different probabilities to the same event, depending on their knowledge and prior beliefs, without entailing any inconsistency. As individuals collect more information about the risk, they should update their probability assessments using Bayes' rule. As common information accumulates, the updated (posterior) probabilities will tend to converge: in the limit, individuals' probabilities for the risk will be equal.

Cases characterized by so-called "objective probabilities" (Knightian risk) such as tossing dice, spinning roulette wheels, and the like seem to be cases in which logic and experience lead individuals to common probabilities. (These cases may be more accurately described as chaotic processes, i.e., non-linear deterministic processes for which the outcome is sensitively dependent on initial conditions. Uncertainty about the outcome results not from randomness but from insufficient knowledge of the initial conditions.) Cases relevant to policy (e.g., the sensitivity of Earth's climate to greenhouse gases, the effects of acid deposition on forest growth, the human-health effects of low-dose exposure to a chemical that causes cancer in laboratory animals at high doses) are manifestly not cases in which logic and experience lead to common probabilities, nor are they cases in which there is sufficient evidence so that posterior probabilities converge.

Under the normative expected utility model (Savage 1954), ambiguity plays no role. Individuals are assumed to behave as if they have both probabilities and utilities for

all possible consequences and to choose the action that has the largest expected utility. Uncertainty about the probability that an event (e.g., developing cancer) will occur can be represented as a probability distribution on the probability of the event, but in calculating expected utility, the “second-order” uncertainty about the probability of cancer integrates out and plays no role.

As an example, consider two risk analysts attempting to characterize uncertainty about whether exposure to methyl mercury through fish consumption increases the risk of heart attack. For simplicity, assume that if methyl mercury increases heart-attack risk, the exposure-response function is accurately estimated (from epidemiological studies); the key uncertainty is whether the relationship is causal (Rice et al. 2010). Let p denote the probability that the relationship is causal. One risk analyst finds the evidence for causality to be ambiguous and assesses a probability distribution for p that is asymmetric triangular between 0 and 1 with mode 0. A second risk analyst interprets the evidence as compelling and assesses a probability distribution that is symmetric triangular between $2/9$ and $4/9$. To calculate the probability that an individual with specified methyl mercury exposure will suffer a heart attack because of his exposure, both risk analysts follow the same procedure: for each possible value of p , multiply the probability of heart attack (given by the exposure-response function), the value of p , and the probability that value of p is correct, then sum. Because the two analysts' distributions for p have identical means ($1/3$), they will calculate identical probabilities of heart attack and their differing interpretations of the ambiguity of the evidence will have no effect on their common estimate of the number of heart attacks under alternative policies.

In contrast to this normative model, much of the risk-perception literature suggests that individuals are less tolerant of environmental and health risks that are perceived to be uncertain or ambiguous. Indeed, the literature on risk perception suggests that the qualitative aspects of risk that influence perception and tolerance can be summarized by two attributes: uncertainty and dread. Risks that are perceived as more uncertain tend to be unobservable, newly recognized, not understood scientifically, and to have delayed consequences (Slovic 1987, 2000). Further evidence of ambiguity or uncertainty aversion is the popularity of the precautionary principle, which is incorporated in a number of international environmental

agreements and in the French constitution. Although there are many statements of the principle, a common theme is that greater uncertainty about a risk should lead to more stringent regulation and smaller exposure. For example, the 1982 United Nations World Charter for Nature (A/RES/37/7, 28 October 1982) states: “Activities which are likely to pose a significant risk to nature shall be preceded by an exhaustive examination; their proponents shall demonstrate that expected benefits outweigh potential damage to nature, and *where potential adverse effects are not fully understood, the activities should not proceed*” (emphasis added). This statement seems consistent with some models of ambiguity aversion, such as the maxmin expected utility model in which decision makers are assumed to hold multiple probability distributions and to evaluate each potential act using the most pessimistic distribution (i.e., the distribution that minimizes the expected utility of that act; Gilboa and Schmeidler 1989). It also seems unworkable, since potential adverse effects of any activity are unlikely to ever be “fully understood.”

Costs of Accommodating Ambiguity Aversion

In conducting BCA and in choosing policy, should analysts and decision makers incorporate aversion to ambiguity and uncertainty? In Portney’s “Trouble in Happyville” example, should the choice depend on whether the residents believe the contaminant will cause cancer or are merely uncertain about whether it does and wish to follow a precautionary principle?

Note that acting in accordance with uncertainty aversion can paradoxically increase risk if decision makers regulate uncertain risks more stringently (relative to expected values) than they regulate more certain risks (Nichols and Zeckhauser 1988). In the United States, the abandonment of new nuclear-power generation and greater reliance on coal-fired power since the 1970s is consistent with choosing a relatively certain risk for a more uncertain one. Although one cannot be sure, subsequent experience with nuclear power world-wide suggests that the damage to human health and ecosystems is much greater than it would have been had the use of nuclear power expanded. Cumulative world-wide fatalities from nuclear power used for producing electricity since the 1950s are estimated to not exceed a few tens of thousands, with the vast majority of cases due to cancer associated with radiation released through the Chernobyl accident

(including cancers that have not yet developed). Estimated *annual* fatalities from burning coal to produce electricity are estimated to be of the same magnitude in the U.S. alone, with equal or greater numbers in Europe and Asia.

To illustrate how regulation in accordance with ambiguity or uncertainty aversion can increase risk, consider a hypothetical example. One must choose which of two technologies to adopt for a particular purpose, e.g., fuels for generating electricity or pesticides for growing cotton. The risk associated with one alternative is uncertain; with probability 0.99 use of this technology will cause no deaths but with probability 0.01 the technology will be catastrophic and cause 1,000 deaths. The other alternative is known to cause 100 deaths. If the choice between the technologies is to be made only once, there is not much to be said except one must weigh the large chance that the uncertain technology will be better than the known technology (saving 100 lives) against the small chance that it will be very much worse (killing 900 more). If the population at risk is the same and everyone is exposed to equal risk, then individuals who behave in accordance with standard economic models would prefer the uncertain technology to the known technology, as it yields smaller risk (if N is the number of people exposed, the risks are $10/N$ and $100/N$, respectively).

If one generalizes from this single decision to a strategy about which choice to make in situations of this kind, then it becomes clear that choosing in accordance with uncertainty aversion is likely to produce greater harm. Assume a choice identical to the example must be made ten times, i.e., one must choose between alternatives with risk profiles identical to the uncertain and known technologies for ten different applications. If one always chooses the certain technology, there will be 1,000 deaths for sure. If one always chooses the uncertain technology, the number of deaths will be between zero (if none of the uncertain technologies prove catastrophic) and 10,000 (if all prove catastrophic). If the event that the uncertain technology is catastrophic is probabilistically independent across the ten situations, then the probability that there are fewer deaths if one always chooses the uncertain technology is 0.996, i.e., it is nearly certain that a policy of choosing the uncertain technology (that causes fewer expected deaths) will cause fewer actual deaths.

Note that the assumption that the risks are independent (or at least not strongly positively dependent) is critical to the argument. In practice, the risks and uncertainty about them will be estimated using some type of risk assessment, which brings together available theory, data, and other evidence about the characteristics of these technologies (Bedford and Cooke 2001, Cox 2002, Hammitt 2008). To the extent that errors in risk assessments are dependent across technologies, the risks may be positively correlated and so the policy of always choosing the uncertain technology will not provide as much benefit as when the risks are independent. Errors in risk assessment may often be positively dependent. For example, risk assessment often uses conservative assumptions about exposure to hazards and the shape of the exposure-response function at low exposures (where risks are too small to be measured but must be extrapolated from measurements at higher exposures). This conservative bias may yield systematically greater over-estimates of risk for technologies with greater uncertainty (Nichols and Zeckhauser 1988). Another source of positive dependence is when the same or similar models or parameters are used to simulate environmental fate and transport, exposure-response functions, or the monetary value of reducing mortality risk.

Note that there may be opportunities to learn about the risks of an uncertain technology over time and to use this information to improve decisions. If feedback about the risks of the technology occurs relatively rapidly, it may be possible to introduce the more uncertain technology on a limited scale, gain experience, then expand or suppress the technology depending on whether it proves less or more hazardous than initially anticipated. If the effects of the technology manifest slowly (e.g., storage technologies for radioactive waste) or cannot be implemented on a limited scale (e.g., changes in greenhouse-gas emissions), it may not be possible to learn enough from limited-scale experimentation to provide useful information for the decision.

Discounting future consequences

In BCA of policies for which the benefits and costs are concentrated at substantially different times, the rate at which future benefits and costs are discounted can have an enormous effect on the conclusions. Standard BCA practice assumes that the importance of a future consequence can be evaluated by its present value, calculated by

discounting the monetary value of the consequence at the time it will occur by a factor that declines geometrically with the number of periods before occurrence (i.e., $[1/(1 + r)]^t$, where r is the discount rate and t the number of periods). This practice is inconsistent with some behavior and some normative arguments in at least two dimensions: the choice of the discount rate and whether the rate is constant or decreasing over time. These issues are discussed in the following subsections.

Hyperbolic Discounting

A key property of the standard discounting model is that a one-period change in the time until arrival has the same proportional effect on the evaluation, regardless of how long in the future the event will occur (i.e., delaying a consequence one period decreases its present value by the factor $1/(1 + r)$). In contrast, empirical studies suggest that individual's evaluations of future consequences are more sensitive to one period changes for proximate than for distant consequences. In particular, individuals may sharply distinguish between present and future consequences. These patterns can be described by hyperbolic discounting (Angeletos et al. 2001).

Overweighting immediate relative to future consequences can bias choices away from alternatives characterized by up-front costs and future benefits, such as investment in more costly but more energy-efficient equipment (e.g., motor vehicles, light bulbs, electrical appliances) as well as precautionary measures to reduce the chance or magnitude of environmental harms, such as reducing greenhouse-gas emissions or more secure containment of hazardous waste. Systematic underweighting of future benefits is an important justification for energy-efficiency standards for household appliances, lighting, and Corporate Average Fuel Economy (CAFE) standards for motor vehicles. Whether consumer's choices between more and less energy-efficient products are systematically biased through underweighting future cost savings is unclear. Early studies (Hausman 1979, Gately 1980) estimated that consumers discounted future energy savings at annual rates of 20 percent or more, but later studies have estimated lower rates and suggested alternative explanations (e.g., Hausman and Joskow 1982, Jaffe and Stavins 1994a, 1994b). Other factors that could explain apparently high discount rates and are consistent with standard economic models include misperception of or uncertainty about

future energy prices, credit constraints, resale market imperfections in which energy efficiency is inadequately capitalized, and heterogeneity or uncertainty about usage patterns (low-use consumers may be better served by inexpensive but less efficient models). Moreover, in choosing when to replace existing equipment with a more efficient model, one must consider the possibility of future improvements in efficiency; even if one could save money by replacing an existing appliance now, the net benefits might be increased by waiting to purchase an even more efficient model later. In short, determining the optimal purchase plan for durable goods is a complicated problem in which many factors might lead to a divergence between consumer behavior and an analytic model.

In a general-population survey, Allcott (2010) finds conflicting evidence about whether American consumers are likely to underweight future fuel costs when buying a motor vehicle. Consumers report having given little attention to fuel efficiency when purchasing their most recent motor vehicle (40 percent “did not think about fuel costs at all”), a result that is compatible with hyperbolic discounting. Other findings suggest that other factors may be important. Allcott reports that consumers misperceive the incremental cost savings from fuel economy (underestimating the marginal effect for low-efficiency and overestimating it for high-efficiency vehicles, consistent with “MPG illusion;” Larrick and Sole 2008). Expectations about future fuel costs are diverse, though consumers tend to anticipate that gasoline prices will rise faster than the rate implied by oil price futures.

If consumer behavior differs from the standard discounting model, one must ask if the model is normatively appropriate. Frederick et al. (2002) note that neither Samuelson (1937) nor Koopmans (1962), who introduced and provided an axiomatic basis for the discounted utility model, respectively, claimed it had either normative or positive validity. Here it is necessary to distinguish between discounting utility and discounting monetary values. Clearly, the model in which future utility is discounted imposes strong assumptions about separability of preferences for consequences in different periods (Frederick et al. 2002). In contrast, the evaluation of future consequences by their discounted monetary values may be justified by an intertemporal budget constraint: to the extent that individuals can shift financial resources forward or backward in time, the

¹³effect on current well-being of a future benefit or cost is equal to the future monetary value discounted at the consumer's interest rate. The relevant interest rate may depend on whether the consumer must increase or decrease his savings (to offset a future cost or benefit, respectively), or must decrease or increase his current borrowing. As an example, an individual's willingness to pay to reduce mortality risk in a future period is theoretically equal to his willingness to pay for that risk reduction at the time it would manifest discounted to the present at the relevant interest rate. This implies that future health or environmental benefits need not be discounted at the same rate as future costs, because the future rate of substitution between health or environment and costs may differ from the present rate (e.g., Cropper and Sussman 1990, Gravelle and Smith 2001, Hammitt and Liu 2004, Sterner and Perrson 2008).

Discount rates

There is a large and active literature on the appropriate discount rate to use for evaluating policies with long time horizons, such as climate change. Although recent literature has advanced to give more attention to uncertainty about future economic growth (Gollier and Weitzman 2010), one of the key issues has been and remains whether to use a discount rate based on positive analysis of behavior (as in Nordhaus 1994) or one based on normative considerations (as in Cline 1992). Discount rates based on analysis of observed interest rates tend to be larger than those based on normative models, and hence BCA using the normative approach tends to support more stringent near-term reductions in greenhouse-gas emissions.

This difference in perspective is highlighted by the Stern Review (Stern 2007) and critiques. Discount rates are often analyzed in terms of the Ramsey equation, which equates the discount rate (equal to the return on capital) to the time discount rate plus the consumption elasticity multiplied by the consumption-growth rate. The time discount rate (or utility discount rate) characterizes the weight placed on the utility of future generations relative to the present generation. Normative arguments including those of Ramsey (1928) suggest a value of zero or slightly positive to account for the risk of

¹³ Hence consumers should allocate their consumption over time so that their consumption discount rate equals the relevant market interest rate.

human extinction. The consumption elasticity is a measure of aversion to inequality of well-being between generations (or the rate at which the marginal utility of consumption decreases), with larger values corresponding to greater inequality-aversion.

¹⁴In its main results, the Stern Review used a time discount of 0.1 percent per year and a consumption elasticity of 1. Combining these with its assumption about future economic growth (approximately 1.3 percent per year) yields a discount rate of approximately 1.4 percent per year and the conclusion that sharp reductions in greenhouse-gas emissions are warranted. In contrast, Nordhaus' well-known DICE model assumes a time discount of 1.5 percent per year and a consumption elasticity of 2. These values are chosen in part so that the resulting discount rate, about 5.5 percent per year, is consistent with positive estimates of the real return to capital over many decades. With these parameter values, DICE suggests more moderate reductions in near-term greenhouse-gas emissions. The stringency of emission reductions can be conveniently characterized by the optimal carbon tax or shadow price of carbon emissions. For both models, the optimal carbon tax rises with time. For DICE, the optimal value in 2015 is approximately \$35/ton carbon while for the Stern Review it is approximately \$350/ton carbon.

Nordhaus (2007) shows that the difference in optimal carbon tax between DICE and the Stern Review can be almost entirely explained by the difference in the discount rate. Changing the time discount and consumption elasticity in DICE to 0.1 percent per year and 1 (as in the Stern Review) yields a discount rate of 2.0 percent and an optimal 2015 carbon tax of \$360/ton carbon. Maintaining the time discount at 0.1 percent per year but increasing the consumption elasticity to 3 yields a discount rate of 5.6 percent per year, comparable to the 5.5 percent rate using the standard assumptions in DICE. In this case, the optimal 2015 carbon tax is \$36/ton carbon, increasing over time in a pattern close to that produced using the standard discounting assumptions in DICE. In short, the optimal restriction on greenhouse-gas emissions is sensitive to the discount rate, but not to the time-discount rate and consumption elasticity (for combinations of these parameters that yield similar discount rates). Moreover, using a normatively-based

¹⁴ Numerical values in the remainder of this subsection are from Nordhaus (2007).

discount rate that is inconsistent with current returns to capital suggests that other policies should also be changed to enhance social welfare. For example, using the discount rate from the Stern Review implies the optimal global savings rate is on the order of twice its current value (Nordhaus 2007).

Divergence of Willingness to Pay and Willingness to Accept

In principle, benefit-cost analysis uses estimates of both willingness to pay (WTP) and willingness to accept compensation (WTA). A policy change can be evaluated using compensating variation by aggregating WTP for the change by those who benefit with WTA for the change by those who are harmed. Alternatively, one could use equivalent variation, aggregating WTA to forgo the policy change by those who would benefit with WTP to prevent the change by those who would be harmed. In practice, however, most benefit-cost analyses use only measures of WTP, usually WTP for the change by those who gain and WTP to prevent the change by those who are harmed.

The conventional reliance on estimated WTP when WTA is conceptually appropriate arises because the two concepts are generally anticipated to be nearly equal under standard economic theory yet empirical estimates of WTA are often much larger than comparable estimates of WTP. The reasons for this divergence are not well understood and analysts have relied on WTP judging it to be more reliable, a position endorsed by U.S. government guidance (OMB 2003).

Revealed-preference estimates typically do not distinguish WTP from WTA, so most of the evidence for differences between these values comes from stated-preference studies and experiments. These studies have found that WTA often exceeds WTP by factors of two to ten or more (Knetsch and Sinden 1984, Horowitz and McConnell 2002, 2003). The disparity is often described as an endowment effect, because the value individuals attach to a good seems to depend on whether they own it or not (Kahneman et al. 1990). The extent to which these results reflect actual preferences or artifacts of study design are controversial. Plott and Zeiler (2005) suggest the disparity reflects experimental subjects' misconceptions and that it can be eliminated when experiments are conducted using an incentive-compatible elicitation mechanism and other features,

though Isoni et al (2011) report experimental evidence that suggests these experimental features are not determinative. In a meta-analysis, Horowitz and McConnell (2002) found that study limitations such as using response formats that are not incentive-compatible or student subjects did not account for the disparities. List (2003, 2004) provides evidence that market experience reduces or eliminates the disparity. This result holds even for individuals who are not professional traders (who may purchase goods for resale rather than private consumption; Kahneman et al. 1990).

Under standard economic models, indifference curves are smooth, not kinked, and hence WTP and WTA for a good are equal at the margin. For non-marginal changes, WTA is expected to be larger than WTP, yet the difference should be small unless the change is large enough to create a substantial income effect or the good is one for which the individual cannot adjust the quantity (e.g., because it is a public good) and there are no adequate market substitutes whose quantity he can adjust to compensate for unwanted changes in the quantity of the primary good (Hanemann 1991). These qualifications do not explain most of the empirical examples, which involve goods of low value for which many substitutes are available.

The disparity between WTP and WTA seems to be better explained by loss aversion and factors that determine the reference point (Knetsch 2010). For example, in evaluating an oil spill or other environmental damage, people are likely to view the environment in the absence of the spill as the reference point, and to evaluate both the spill and efforts to clean it up as in the domain of losses. Hence WTA for the spill (compensating variation) and WTA to forgo cleanup (equivalent variation) will tend to be large compared with a situation in which the status quo is viewed as the appropriate reference, perhaps WTP for restoration of salmon runs in rivers from which the fish have long been absent.

If loss aversion explains much of the WTP-WTA disparity, then a key question is whether loss aversion is normatively acceptable. In some cases, the reference point can be easily manipulated and any resulting change in decision seems likely to be judged as a cognitive error. For example, it seems difficult to imagine that an individual offered a choice between a coffee mug and \$5 would wish to make his choice dependent on whether or not he was initially given the mug by the experimenter, though this is the

pattern observed in experiments (Kahneman 1991). In a case such as the “Trouble in Happyville” example where both the benefits and the costs of an action accrue to the same people, it seems that each individual should prefer either the consequences of spending the money and treating the water, or saving the money and not treating the water; consideration of whether the consequences are compared with the status quo or with treated water should not alter his decision. When the benefits and costs accrue to different people, there are strategic reasons for those who are harmed to exaggerate the harm, and those who benefit to underplay the value of the gain (if they will be made to pay). Yet loss aversion is observed even with incentive-compatible decisions (such as the coffee mug experiments). It seems to reflect an availability effect (Tversky and Kahneman 1974), in which the loss of the specific good (e.g., the coffee mug) is highly salient and the opportunity loss of the goods one could purchase with the money (including a substitute mug) are less salient.

Implications for Benefit-Cost Analysis

How should BCA be interpreted, what are its implications for policy choice, and how should it be conducted? The answers depend on whether BCA is justified as a positive or normative exercise.

If BCA is conceived as a positive exercise, with the goal of determining whether policy consequences satisfy the condition that those who benefit could theoretically compensate those who are harmed, then the objective is to measure benefits and harms exactly as they are perceived by the affected individuals. When these perceptions conflict with normative models, the normative models are irrelevant for the analysis. Under this interpretation, analysts should measure individual preferences as accurately as possible. This includes fully incorporating any inter-individual differences in valuation, such as those related to income, age, or other factors, which are ignored in conventional BCA and may incite strong political reaction (recall the controversies over valuing mortality risks differently by age; Viscusi 2009). Predicting policy consequences and measuring individual perceptions are scientific questions that are, in principle, susceptible to empirical testing. In addition, this approach respects individual autonomy (consumer sovereignty). However, if BCA is conceived as a positive exercise, the question of its

significance for policy remains. While a practice of choosing policies that satisfy the Kaldor-Hicks compensation test allows for the possibility that everyone in the population will gain in the aggregate, there is no guarantee that such an objective will be achieved and the possibility that other social objectives, such as fair distribution of outcomes or equality of opportunity, may be compromised.

If BCA is conceived as a normative exercise, then the normative basis must be specified. If chosen by the analyst, she must be explicit about the choice and why it is appropriate. As the choice of normative basis is a political rather than a scientific question, it seems appropriate for the choice to be made by the relevant political decision maker, e.g., by legislation or executive-branch guidance, though the prospects that political decision makers will provide a sufficiently precise statement for analysts to follow seem limited. When using a normative basis, the analyst must determine which parameter values are consistent with the corresponding normative model. This may require adjusting empirical estimates to correct for biases that result from behavior that is inconsistent with the normative model; in many cases, it is not clear how such adjustments are best made. This approach assumes the analyst is a better judge of individuals' well-being than the individuals themselves and opens her to charges of elitism or paternalism. Many of the questions involved in conducting BCA under a normative justification are not scientific but philosophical and not susceptible to empirical testing, which places the analyst in more of an advocacy than a scientific role. Nevertheless, benefit-cost analysis that rests on an accepted normative basis is by definition more useful for policy guidance than one that simply predicts if the policy passes the Kaldor-Hicks compensation test.

It is possible that future research may limit the differences in interpretation between positive and normative perspectives. For example, some important cases of individual behavior that conflict with economic models might be shown to be errors, or an increasingly sophisticated population may come to behave more like the models assume. Alternatively, some deviations might be explained by refining economic models to incorporate factors that are viewed as normatively legitimate but currently omitted. Such changes seem unlikely to substantially improve the correspondence between normative models and behavior, which implies there will remain an important question

about the basis for BCA, with corresponding implications for its conduct and interpretation.

The choice of justification is part of a larger question about the role of representative government: should government provide the citizenry what the citizenry believes it wants at the moment, much as a direct democracy (or a politician who slavishly follows public-opinion polls) might do, or should it provide leadership, directing the citizenry in a direction it does not yet know (and might never agree) is in its real interests? The tension is encapsulated in the debate between Thomas Jefferson and Edmund Burke (Wiener 1997). Jefferson (1820): “I know of no safe repository of the ultimate powers of society but the people themselves; and if we think them not enlightened enough to exercise their control with a wholesome discretion, the remedy is not to take it from them, but to inform their discretion by education.” Burke (1774): “Your representative owes you, not only his industry, but his judgment; and he betrays, instead of serving you, if he sacrifices it to your opinion.”

How would these sages advise the Director of Environmental Protection in Happyville? Jefferson might counsel him to continue to communicate the risk as clearly as possible, and also to educate the citizens about the opportunity costs of the watertreatment option. If the citizens persist in their preference, would he ultimately advise the Director to treat the water? And would Burke advise the Director to refuse?

References

Abramowicz, “Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis,” *University of Chicago Law Review* 71: 933-1020, 2004.

Adler, M.A., and E.A. Posner, *New Foundations for Benefit-Cost Analysis*, Harvard University Press, Cambridge, 2006.

Alberini, A., M. Cropper, A. Krupnick, and N. Simon, “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: 769-792, 2004.

Allais, M., “Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’Ecole Américaine,” *Econometrica* 21: 503-546, 1953.

Allcott, H., “Consumers’ Perceptions and Misperceptions of Energy Costs,” *American Economic Review* 101: 98-104, 2010.

Angeletos, G.-M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg, “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives* 15: 47-68, 2001.

Arrow, K.J., M.L. Cropper, G.C. Eads, R.W. Hahn, L.B. Lave, R.G. Noll, P.R. Portney, M. Russell, R. Schmalensee, V.K. Smith, and R. N. Stavins, “Is There a Role for Benefit-Cost Analysis in Environmental, Health, and Safety Regulation?” *Science* 272: 221-222, 1996.

Bedford, T., and R. Cooke, *Probabilistic Risk Analysis: Foundations and Methods*, Cambridge University Press, Cambridge, 2001.

Bernheim, B.D., and A. Rangel, “Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics,” *Quarterly Journal of Economics* 124: 51-104, 2009.

Beshears, J., J.J. Choi, D. Laibson, and B.C. Madrian, “How Are Preferences Revealed?” *Journal of Public Economics* 92: 1787-1794, 2008.

Binmore, K., *Rational Decisions*, Princeton University Press, Princeton, 2009. Bleichrodt, H., and L. Eeckhoudt, “Willingness to Pay for Reductions in Health Risks when Probabilities are Distorted,” *Health Economics* 15: 211-214, 2006. Bleichrodt, H., J.L. Pinto, and P.P. Wakker, “Making Descriptive Use of Prospect Theory to Improve the Prescriptive Use of Expected Utility,” *Management Science* 47: 1498-1514, 2001.

Boardman, A., D. Greenberg, A. Vining, and D. Weimer, *Cost Benefit Analysis: Concepts and Practice*, Prentice Hall, New York, 2006.

Burke, E. “Speech at the Conclusion of the Poll, November 3, 1774,” *The Writings and*

Speeches of Edmund Burke (W.M. Elofson and J.A. Woods, eds.), 1996.

Cantley, M., and M. Lex, “Genetically Modified Foods and Crops,” 39-64 in *The Reality of Precaution: Comparing Risk Regulation in the United States and Europe* (J.B.

- Wiener, M.D. Rogers, J.K. Hammitt, and P.H. Sand, eds.), *Resources for the Future Press*, Washington, D.C., 2010.
- Cline, W.R., *The Economics of Global Warming*, Institute for International Economics, Washington, DC, 1992.
- Cohen, M., and J.-Y. Jaffray, "Is Savage's Independence Axiom a Universal Rationality Principle?" *Behavioral Science* 33: 38-47, 1988
- Corso, P.S., J.K. Hammitt, and J.D. Graham, "Valuing Mortality-Risk Reduction: Using Visual Aids to Improve the Validity of Contingent Valuation," *Journal of Risk and Uncertainty* 23: 165-184, 2001.
- Cox, L.A. Jr., *Risk Analysis: Foundations, Models and Methods*, Kluwer Academic Publishers, Dordrecht, 2002.
- Cropper, M. L. and F. G. Sussman, "Valuing Future Risks to Life," *Journal of Environmental Economics and Management* 19: 160– 174, 1990.
- Diamond, P. A., and J.A. Hausman, "Contingent Valuation: When is Some Number Better than No Number?" *Journal of Economic Perspectives* 8: 45-64, 1994.
- Downs, A., "An Economic Theory of Political Action in a Democracy," *Journal of Political Economy* 65: 135-150, 1957.
- Ellsberg, D., "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics* 75: 643-669, 1961.
- Fredrick, S., G. Lowenstein, and T. O'Donohue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature* 40: 351-401, 2002.
- Gately, D., "Individual Discount Rates and the Purchase and Utilization of Energy -Using Durables: Comment," *Bell Journal of Economics* 11: 373-374, 1980.
- Gilboa, I., and D. Schmeidler, "Maximin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics* 18: 141-153, 1989.
- Gneezy, U., J.A. List, and G. Wu, "The Uncertainty Effect: When a Risky Prospect is Valued Less than its Worst Possible Outcome," *Quarterly Journal of Economics* 121: 1283-1309, 2006.
- Gold, L.S., B.N. Ames, and T.H. Slovic, "Misconceptions about the Causes of Cancer," pp. 1415-1460 in *Human and Environmental Risk Assessment: Theory and Practice* (D. Paustenbach, ed.), John Wiley and Sons, New York, 2002.
- Gold, M.R., J.E. Siegel, L.B. Russell, and M.C. Weinstein (eds.), *Cost-Effectiveness in Health and Medicine*, Oxford University Press, Oxford, 1996.
- Gollier, C., and M.L. Weitzman, "How Should the Discount Rate be Discounted When

Discount Rates are Uncertain? ” *Economics Letters* 107: 350-353, 2010.

Gowdy, J.M., “Toward an Experimental Foundation for Benefit -Cost Analysis,” *Ecological Economics* 63: 649-655, 2007.

- Gravelle, H., and D.S. Smith, "Discounting for Health Effects in Cost-Benefit and Cost-Effectiveness Analysis," *Health Economics* 10: 587–599, 2001.
- Hammitt, J.K., "Environmental Risk Regulation," Chapter 112 (613–617) in E. Melnick and B. Everitt (eds.), *Encyclopedia of Quantitative Risk Analysis and Assessment*, John Wiley and Sons, Chichester, 2008.
- Hammitt, J.K., "Evaluating Contingent Valuation of Environmental Health Risks: The Proportionality Test," *Association of Environmental and Resource Economists Newsletter* 20(1): 14-19, May 2000.
- Hammitt, J.K., and J.D. Graham, "Willingness to Pay for Health Protection: Inadequate Sensitivity to Probability?" *Journal of Risk and Uncertainty* 18: 33-62, 1999.
- Hammitt, J.K., and J.-T. Liu, "Effects of Disease Type and Latency on the Value of Mortality Risk," *Journal of Risk and Uncertainty* 28: 73-95, 2004.
- Hammitt, J.K., and K. Haninger, "Valuing Fatal Risks to Children and Adults: Effects of Disease, Latency, and Risk Aversion," *Journal of Risk and Uncertainty* 40: 57-83, 2010.
- Hammitt, J.K., and K. Haninger, "Willingness to Pay for Food Safety: Sensitivity to Duration and Severity of Illness," *American Journal of Agricultural Economics* 89: 1170-1175, 2007.
- Hanemann, W.M., "Valuing the Environment through Contingent Valuation," *Journal of Economic Perspectives* 8: 19-43, 1994.
- Hanemann, W.M., "Willingness to Pay and Willingness to Accept: How Much Can They Differ?" *American Economic Review* 81: 635-647, 1991.
- Haninger, K., and J.K. Hammitt, "Diminishing Willingness to Pay per Quality-Adjusted Life Year: Valuing Acute Foodborne Illness," *Risk Analysis* DOI: 10.1111/j.15396924.2011.01617.x, published online April 2011.
- Hausman, J.A., "Individual Discount Rates and the Purchase and Utilization of Energy Using Durables," *Bell Journal of Economics* 10: 33-54, 1979.
- Hausman, J.A., and P.L. Joskow, "Evaluating the Costs and Benefits of Appliance Efficiency Standards," *American Economic Review* 72: 220-225, 1982.
- Horowitz, J.K., and K.E. McConnell, "A Review of WTA/ WTP Studies," *Journal of Environmental Economics and Management* 44: 426-447, 2002.
- Horowitz, J.K., and K.E. McConnell, "Willingness to Accept, Willingness to Pay, and the Income Effect," *Journal of Economic Behavior and Organization* 51: 537-545, 2003.
- Isioni, A., G. Loomes, and R. Sugden, "The Willingness to Pay — Willingness to Accept Gap, the "Endowment Effect," Subject Misconceptions, and Experimental Procedures for Eliciting Valuations: Comment," *American Economic Review* 101: 991-1011, 2011.
- Jaffe, A.B., and R.N. Stavins, "The Energy Efficiency Gap: What Does it Mean?" *Energy*

Policy 22: 804-810, 1994a.

Jaffe, A.B., and R.N. Stavins, "The Energy Paradox and the Diffusion of Conservation Technology," *Resource and Energy Economics* 16: 91-122, 1994b.

Jefferson, T., "Letter to Mr. Jarvis, September 28, 1820," *The Writings of Thomas Jefferson* (H.A. Washington, ed.), 1854.

Kahneman D., *Thinking, Fast and Slow*, Farrar Straus & Giroux, New York, 2011.

Kahneman, D., and A. Tversky (eds.), *Choices, Values, and Frames*, Cambridge

University Press, Cambridge, 2000. Kahneman, D., and A. Tversky, "Prospect Theory: An Analysis of Decision Under Risk,"

Econometrica 47: 263-291, 1979. Kahneman, D., and R. Sugden, "Experimented Utility as a Standard of Policy Evaluation,"

Environmental and Resource Economics 32: 161-181, 2005. Kahneman, D., J.L. Knetsch, and R.H. Thaler, "Experimental Tests of the Endowment

Effect and the Coase Theorem," *Journal of Political Economy* 98: 1325-1348, 1990.

Kahneman, D., P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, 1982.

Knetsch, J.L., "Values of Gains and Losses: Reference States and Choice of Measure," *Environmental and Resource Economics* 46: 179-188, 2010.

Knetsch, J.L., and J.A. Sinden, "Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value," *Quarterly Journal of Economics* 99: 507-521, 1984.

Knight, F.H., *Risk, Uncertainty, and Profit*, Houghton Mifflin Company, Boston, 1921.

Koopmans, T.C. "Statistical Ordinal Utility and Impatience," *Econometrica* 28: 287-309,

1960. Larrick, R.P., and J.B. Soll, "The MPG Illusion," *Science* 320(5883): 1593-1594, 2008. Machina, M.J., "Choice under Uncertainty: Problems Solved and Unsolved," *Journal of*

Economic Perspectives 1: 121-154, 1987. Manski, C.F., "Actualist Rationality," *Theory and Decision*. forthcoming. Nichols, A.L., and R.J. Zeckhauser, "The Perils of Prudence: How Conservative Risk

Assessments Distort Regulation," *Regulatory Toxicology and Pharmacology* 8: 61-75, 1988.

Nordhaus, W.D., "A Review of the Stern Review on the Economics of Global Warming," *Journal of Economic Literature* 45: 686-702, 2007

Nordhaus, W.D., *Managing the Global Commons: The Economics of Climate Change*, MIT Press, Cambridge, 1994.

Nyborg, K., *The Ethics and Politics of Environmental Cost-Benefit Analysis*, Routledge, Abingdon, Oxon, UK, 2012.

- Okun, A.M., *Equality and Efficiency: The Big Tradeoff*, The Brookings Institution, Washington, DC, 1975.
- OMB (U.S. Office of Management and Budget), "Circular A-4: Regulatory Analysis," Washington, D.C., September 17, 2003.
- Peltzman, S., "The Effects of Automobile Safety Regulation," *Journal of Political Economy* 83: 677-726, 1975.
- Portney, P.R., "Trouble in Happyville," *Journal of Policy Analysis and Management* 11: 131-132, 1992.
- Pratt, J.W., and R.J. Zeckhauser, "Willingness to Pay and the Distribution of Risk and Wealth," *Journal of Political Economy* 104: 747-763, 1996.
- Putler, D.S., "Incorporating Reference Price Effects into a Theory of Consumer Choice," *Marketing Science* 11: 287-209, 1992.
- Quiggin, J., *Generalized Expected Utility Theory: The Rank-Dependent Model*, Springer, 1993.
- Ramsey, F.P., "A Mathematical Theory of Saving," *Economic Journal* 38: 543-559, 1928.
- Rice, G., J.K. Hammitt, and J.S. Evans, "A Probabilistic Characterization of the Health Benefits of Reducing Methyl Mercury Intake in the United States," *Environmental Science and Technology* 44: 5216-5224, 2010.
- Robinson, L.A., and J.K. Hammitt, "Behavioral Economics and the Conduct of Benefit Cost Analysis: Towards Principles and Standards," *Journal of Benefit-Cost Analysis* 2(2): Article 5, DOI: 10.2202/2152-2812.1059, 2011a.
- Robinson, L.A., and J.K. Hammitt, "Behavioral Economics and Regulatory Analysis," *Risk Analysis* 31(9): 1408-1422, 2011b.
- Saint-Paul, G., *The Tyranny of Utility: Behavioral Social Science and the Rise of Paternalism*, Princeton University Press, Princeton 2011.
- Sala-i-Martin, X., and N. Treich, "Regulation in Happyville," *Economic Journal* 119: 665-679, 2009.
- Samuelson, P., "A Note on Measurement of Utility," *Review of Economic Studies* 4: 155-61, 1937.
- Savage, L.J. *The Foundations of Statistics*, John Wiley and Sons, New York, 1954. Sen, A., *Inequality Reexamined*, Harvard University Press, Cambridge, 1992. Slovic P. *The Perception of Risk*, Earthscan Publications, London, 2000. Slovic, P., and A. Tversky, "Who Accepts Savage's Axiom?" *Behavioral Science* 19: 308-373, 1974. Slovic, P., "Perception of Risk," *Science* 236(4799), 280-285, 1987. Smith, V.K., and E.M. Moore, "Behavioral Economics and Benefit Cost Analysis,"

Starr, C., "Social Benefit versus Technological Risk: What is Our Society Willing to Pay for Safety?" *Science* 165: 1232-1238, 1969.

Stern, N., *The Economics of Climate Change: The Stern Review*, Cambridge University Press, 2007.

Stern, T., and U.M. Persson, "An Even Stern Review: Introducing Relative Prices into the Discounting Debate," *Review of Environmental Economics and Policy* 2: 61-76, 2008.

Sugden, R., "Coping with Preference Anomalies in Cost-Benefit Analysis: A Market Simulation Approach," *Environmental and Resource Economics* 32: 129-160, 2005.

Sugden, R., "The Behavioural Economist and The Social Planner: To Whom Should Behavioural Welfare Economics be Addressed?" *Papers on Economics and Evolution* #1121, Max Planck Institute of Economics, Jena, Germany, 20 December 2011.

Sugden, R., "The Opportunity Criterion: Consumer Sovereignty without the Assumption of Coherent Preferences," *American Economic Review* 94: 1014-1033, 2004.

Sunstein, C., "Probability Neglect: Emotions, Worst Cases, and Law," *Yale Law Journal* 112: 61-107, 2002.

Sunstein, C., and R.J. Zeckhauser, "Overreaction to Fearful Risks," *Environmental and Resource Economics* 48: 435-449, 2011.

Sunstein, C.R., "Cognition and Benefit-Cost Analysis," *Journal of Legal Studies*, 29(2), 1059-1103, 2000.

Tsao, J.D., H.D. Saunders, J.R. Creighton, M.E. Coltrin, and J.A. Simmons, "Solid-State Lighting: An Energy-Economics Perspective," *Journal of Physics D: Applied Physics* 43: 354001 (17 pp), 2010.

Tversky, A., and D. Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5: 297-323, 1992.

Tversky, A., and D. Kahneman, "Judgment Under Uncertainty: Heuristics and Biases," *Science* 185: 1124-1131, 1974.

Tversky, A., and D. Kahneman, "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *Quarterly Journal of Economics* 106: 1039-1061, 1991.

Viscusi, W. K., "A Bayesian Perspective on Biases in Risk Perception," *Economic Letters* 17: 59-62, 1985.

Viscusi, W. K., "Prospective Reference Theory: Toward an Explanation of the Paradoxes," *Journal of Risk and Uncertainty* 2: 235-263, 1989.

Viscusi, W. K., "The Devaluation of Life," *Regulation and Governance* 3: 103-127, 2009. Wiener, J.B., "Risk in the Republic," *Duke Environmental Law and Policy Forum* 8: 1-

21, 1997.