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

January 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

What is the goal of benefit-cost analysis (BCA)? What question does it address? Answering these questions is critical to understanding how BCA results should be interpreted, their implications for public policy, and how BCA should be conducted.

There are at least two possible bases for justifying BCA, positive and normative. The positive basis 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 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 compensation payments than without. The positive justification for BCA is that it answers the question: is there a set of monetary transfers such that the policy change plus transfers would constitute a Pareto improvement over the status quo? This 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.

The normative basis asserts that BCA is a method to identify policy changes that constitute social improvements, where improvement must be defined in some way that is external to BCA. Several normative justifications can be offered, depending on how social improvement is defined. One justification is derived from a form of 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.

Another set of normative justifications is 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). An alternative consistency claim is 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 individual's utilities. Obviously, the strength of these pragmatic claims depends on the alternatives with which BCA is compared.

BCA 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' (either in comparison with continuing the status quo policy or by forgoing alternative policies that would have improved the others' well-being by more than the selected policy). 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 willingness to pay (WTP) for an improvement and willingness to accept compensation (WTA) for a decrement. Other concepts of well-being include capabilities (Sen 1992), quality-adjusted life years (QALYs) for health and longevity (Gold et al. 1996), and conceptions of a good or moral life proffered by many religions.

BCA compares changes in well-being among people using WTP and WTA. By contrast, cost-effectiveness analysis (CEA), as widely practiced in public health and medicine, evaluates changes in health using QALYs and changes in other consequences using monetary units (Gold et al. 1996). The assumption is that interpersonal comparisons of health should be evaluated differently than interpersonal comparisons of other consequences; i.e., health should be evaluated by treating healthy years as equally valuable and other consequences should be evaluated by treating dollars (or other monetary units) as equally valuable.

Background

Conventional BCA relies on standard economic assumptions about human behavior. The most important of these is 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 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.

Do these 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?

Differences 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 Kaldor-Hicks compensation test 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. 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 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.) 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.

How should policy makers and analysts respond when confronted with public preferences that depart significantly from the normative preferences embodied in economic models? Paul Portney provocatively posed this question in his 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 ten 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 still 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 each resident the other benefits he could achieve with $1,000 but each resident will believe himself to be better off. If you reject the treatment option, you are knowingly imposing a policy that each resident believes is contrary to his well-being.

Distinguishing behavioral and cognitive errors from oversimplified models is critical to understanding how BCA results should be interpreted, what significance they should have for policy making, and how BCA should be conducted. In the following sections, I describe how BCA relies on two types of inputs (scientific predictions of consequences and individuals' preferences over consequences) and examine how the divergence between scientific models and human behavior has implications for conduct and interpretation of BCA. Conclusions are offered in the final section.

Inputs to Benefit-Cost Analysis

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 risk-assessment 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 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. 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 somewhat 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.

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).

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

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 (Hammitt 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 (Pratt and Zeckhauser 1996).

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 or 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).

Stated-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 x $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.

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 arises in other contexts as well. It is but one example of the problem of 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). 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 can be reconciled with economic theory by adopting a more refined theoretical model. As noted, if the problem is one of imperfect risk communication, certain visual aids or other devices may help (Corso et al. 2001, Hammitt and Graham 1999). 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 combining their prior beliefs about the hazard with information provided in the survey, which 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 we have 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 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 of that value of *p*, 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 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 UN 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.”

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 be in the 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 fatalities each year from burning coal to produce electricity in the U.S. alone are estimated to be of the same magnitude.

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 one death but with probability 0.01 the technology will be catastrophic and cause 1,000 deaths. The other alternative is known to cause 101 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 899 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, the risks are 11/N and 101/N, respectively).

If one generalizes from this single decision to a policy 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 1,010 deaths for sure. If one always chooses the uncertain technology, the number of deaths will be between 10 (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 systematic or 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 either expand or suppress the technology if 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.

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 taking alternative decisions. In the worst case, information can be ignored and so it causes no harm. In practice, individuals may be 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 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.

Hyperbolic Discounting

Another domain in which behavior appears to systematically differ from standard models is the choice among actions with near-term and more distant future consequences. Standard models assume 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., d^t, where d is the discount factor and t the number of periods). A key property of this model is that a one-period change in the time until arrival has the same proportional effect on the evaluation whether the consequence will occur sooner or later. In contrast, empirical studies suggest that individuals' evaluations of future consequences are more sensitive to one period changes for near-term than for long-term 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 consumers' 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 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 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") and that they 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. (2001) 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. 2001). In contrast, the evaluation of future consequences by their discounted monetary values may be justified as an intertemporal budget constraint: to the extent that individuals can shift financial resources forward or backward in time, the effect 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 increases or decreases his current savings (to offset a future cost or benefit, respectively), or decreases or increases his current borrowing. As an example, an individual's willingness to pay to reduce mortality risk in a future period is 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).

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 block 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 have not been 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 large differences have been attributed to the hypothetical nature of stated-preference choices or to other study limitations, though Horowitz and McConnell (2002) found in their meta-analysis that limitations such as using response formats that are not incentive-compatible or student subjects did not account for the disparities.

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, or WTA to forgo, 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 (Tversky and 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 population. When these perceptions conflict with normative models, the normative models are irrelevant. 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 logical, and not susceptible to empirical testing, which places in 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.

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 council him to continue to communicate the risk as clearly as possible, and also to educate the citizens about the opportunity costs of the water-treatment 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

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 Americaine," *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.

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

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

Bleichrodt, H., and L. Eeckhoudt, "Willingness to Pay for Reductions in Health Risks when Probabilities are Distorted," *Health Economics* 15: 211-214, 2006.

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.

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.

Frederick, S., G. Loweinstein, 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, “Maxmin Expected Utility with Non-unique Prior,” *Journal of Mathematical Economics* 18: 141-153, 1989.

Gold, L.S., B.N. Ames, and T.H. Slone, “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.

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.1539-6924.2011.01617.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, "Willingess to Accept, Willingness to Pay, and the Income Effect," *Journal of Economic Behavior and Organization* 51: 537-545, 2003.

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

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.

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. "Stationary 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.

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.

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,” *Marking Science* 11: 287-209, 1992.

Quiggin, J., *Generalized Expected Utility Theory: The Rank-Dependent Model*, Springer, 1993.

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.

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., “Perception of Risk,” *Science* 236(4799), 280-285, 1987.

Slovic, P, and A. Tversky, “Who Accepts Savage’s Axiom?” *Behavioral Science* 19: 308-373, 1974.

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

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

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

Sunstein, C., and R.J. Zeckhauser, “Overreaction to Fearsome 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, "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.

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

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.