On Tools and Tool Boxes:
Valuing Mortality Risks in Benefit-Cost Analysis

November 2010 Draft

Prepared by:

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

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

Abstract: Policymakers access many different tool boxes when building the foundation for addressing environmental, health, and safety risks, drawing from science, law, philosophy, and other disciplines. We focus on the tool box labeled "benefit-cost analysis" and discuss how the tool commonly (but confusingly) labeled the "value per statistical life" fits within that framework. We first explore the role of benefit-cost analysis, arguing that it is best viewed as a pragmatic approach for providing information about individuals' preferences for specific government programs versus other ways to exchange their income for improvements in their own well-being. Within this framework, we explore the approach used for valuing mortality risks. Uncertainty about the extent to which these values change across individuals and across types of risks is a data limitation rather than a moral quandary. However, while benefit-cost analysis provides useful information, it is not, and should not be, the sole basis for decisions.

1.0 Introduction

Public policy decisions are rarely simple. Challenges range from disputes about the underlying facts and the degree of scientific uncertainty to disagreement about the appropriate role of government and competing policy goals. Policies focused on reducing environmental, health, and safety risks pose particularly difficult challenges. We have no market that provides direct evidence on how we value reductions in these risks, and consideration of these risks poses particular ethical and moral challenges.

As a result, policymakers must access many different tool boxes when building the foundation for these decisions, drawing from the fields of science, law, philosophy, and other disciplines. In this article, we focus on the tool box conventionally labeled "benefit-cost analysis" (BCA) and discuss how the tool commonly (but confusingly) labeled the "value per statistical life" (VSL) fits within that framework. We first briefly explore the role of benefit-cost analysis. We argue that it is most useful as a pragmatic approach for combining information on physical impacts with information on individual preferences for spending on the outcomes of concern. It is not, and should not be, the only source of information used to support policy decisions. We next explore how changes in mortality risks should be valued within this framework, discussing confusion over the definition of the VSL as well as difficulties in its measurement. We conclude with some recommendations for future work.

More generally, we argue that the appropriate approach for "valuing lives" should be determined based on the context. In BCA, we are largely interested in estimating individuals' willingness to trade money for small changes in their own risks. This framing mimics the effects of policy decisions to some extent, given that they involve choosing whether to expend resources on risk reductions rather than on alternative desired outcomes. However, policymakers are also interested in other impacts, such as the extent to which the policy benefits those who are worst off. As a result, BCA must be supplemented by other forms of assessment that weigh the importance of nonquantified impacts including these types of moral considerations.

2.0 The Benefit-Cost Analysis Tool Box

While BCA could be used to describe any process for evaluating the improvements and harms that might result from a policy change, it is typically used to describe analyses that rely on monetary values to measure and compare policy impacts. Typically, these analyses begin by describing current and potential future baseline conditions in the absence of the policy, then compare predicted policy outcomes to this baseline. Although there is no principled distinction between benefits and costs (costs can be portrayed as negative benefits and vice-versa), costs are typically defined as the opportunity costs of the real resources expended to develop, implement, and operate a program or comply with policy requirements, including market impacts. Benefits typically include the monetary value of the outcomes that are the goal of the policy, such as improved environmental quality, reduced illness or injury, or increased longevity. Ideally, any significant side effects (cost-savings or ancillary benefits) are included, and the implications of nonquantified effects and uncertainty are carefully assessed.

Such analysis is now widely used to inform policy decisions. In the U.S., it is required for economically-significant Federal regulations under government-wide requirements (Clinton 1993, OMB 2003), and it is increasingly used to support legislative and regulatory decisions in OECD-member countries (OECD 2009). While a recent evaluation found that the World Bank's use of BCA has declined, it provided several recommendations for increasing and improving its application (IEG 2010). Several U.S. state governments also require BCA as part of their regulatory processes (Schwartz 2010). In addition, current work by the Institute for Policy Integrity at the New York University School of Law suggests a growing role for BCA in developing countries.[1]

2.1 Efficiency vs. Equity

The widespread use of BCA suggests that decisionmakers find the results useful; however, even its advocates recognize its limitations. Early neoclassical theorists (beginning with Kaldor 1939) argued for separating analysis of economic efficiency (i.e., the BCA) from consideration of equity, a view maintained in much “best practice” guidance today (e.g., OMB 2003). Thus BCA focuses on identifying which policy (if any) would yield the largest net benefits, regardless of who bears the costs or who receives the benefits. A “hard” benefit-cost test would base policy decisions solely on this efficiency criterion, to ensure that limited resources are invested in those activities that most increase net social welfare.[2] This approach is often described as maximizing the social welfare “pie,” which then can be redistributed as desired. Those who support this perspective generally argue that equity goals can be achieved more comprehensively and at lower cost using direct money transfers (through the tax system and programs that provide supplementary income) rather than by altering each specific policy focused on improved health or other outcomes.

Graham (2008) and others note that this “hard” benefit-cost test is rarely, if ever, applied in practice. Several decades of experience suggest that BCA provides useful information but must be supplemented with careful consideration of nonquantified impacts. Such impacts are important for assessing efficiency: many policies lead to health or environmental impacts that cannot be fully quantified due to the need for more scientific research. In addition, decisionmakers and stakeholders clearly care about other concerns within the context of environmental, health, and safety policy, and are not willing to defer responsibility for addressing these concerns to those involved in tax and income support programs. Finally, the priorities and constraints embedded in the legal system have important implications for the application of BCA to public policy. Thus a “soft” benefit-cost test, that takes into account nonquantified factors and other concerns, is more frequently advocated.

In recent years, increased attention has been paid to the need for more rigorous assessment of equity. The approach for investigating these concerns is not as well-developed nor as formalized as the approach for assessing efficiency; often the distribution of the outcomes is described but not evaluated in detail.[3] Several scholars have proposed approaches for more careful analysis, based at least in part on existing research on social welfare functions, optimal

---

[1] Personal communication with M. Livermore, Executive Director, Institute for Policy Integrity, October 2010.

[2] More formally, this is the Kaldor-Hicks potential compensation test, which argues that policies should be pursued if the “winners” could compensate the “losers,” regardless of whether such compensation actually occurs.

[3] Proposals to weight costs and benefits to reflect equity concerns generally have not been accepted in the U.S., due to the lack of agreement on the appropriate weights as well as concerns about transparency. Other countries appear more willing to implement equity weighting; for example, the United Kingdom suggests weighting by income level in policy analyses (HM Treasury 2003).

taxation, the marginal utility of income, and related topics (e.g., Sunstein 2007, Adler 2008, Graham 2008, Farrow 2009, Loomis 2009, Zerbe 2009, Johansson-Stenman and Konow 2010). The feasibility and usefulness of these approaches have not yet been carefully tested.

The rapid growth of research in behavioral economics and the psychological aspects of decisionmaking has added another dimension to this discussion, suggesting that individuals do not always make choices that maximize their own welfare. For example, Sunstein (2007) argues that BCA should be used to counterbalance the tendency towards cognitive failures. However, Hammitt (2009) notes that, while we might want experts to fix obvious mistakes, in many cases the correct answer is uncertain. Experts also are not immune to the types of decisionmaking anomalies identified in behavioral research. The most appropriate role for BCA may be to provide a pragmatic and well-established approach for collecting, analyzing, and communicating information on the preferences of those potentially affected by the policy options. While recognizing that individuals may make errors, the task of the analyst then becomes identifying the best available evidence on the values that result when individuals are well-informed and have the opportunity for reflection.

2.2 Additional Other-Regarding Preferences

Equity is not the only type of other-regarding (or social) preference typically excluded from quantified BCA results. In general, BCA relies on values that reflect an individual's own preferences for exchanging income for the benefits he or she would receive. While this perspective is, at times, characterized as assuming that only selfish, self-regarding (or private) preferences matter in economic analysis, in reality it reflects difficulties inherent in determining whether and how to appropriately measure and include other types of preferences.

For example, altruism has received substantial attention within the traditional neoclassical economic framework (e.g., Jones-Lee 1991, Bergstrom 2006). When altruism is pure or non-paternalistic (i.e., we respect other's preferences), counting both altruistic and private values simply scales costs and benefits upwards without affecting the overall analytic conclusions. If altruism instead varies depending on the outcome, then it can affect the balance between benefits and costs.[4] In particular, some research suggests that health and longevity are viewed more paternalistically than most other goods (e.g., Jacobsson et al. 2007). However, determining whether and how the monetary values of various outcomes are differentially affected by paternalistic altruism is difficult, and there has been little empirical work. Hence altruism is typically not reflected in the quantified BCA results and instead must be included when considering the implications of nonquantified concerns.

Other-regarding preferences can take a number of forms, such as a desire to help those who are less-well off, to reduce differences between oneself and others, to reward or penalize others depending on perceived fairness of their actions, or to maintain one's relative as well as absolute position. The fact that these values are generally not quantified in BCA is not a statement about their importance or merit, rather it reflects the complex nature of these concerns and the difficulties associated with measuring their impact on monetary values. These types of challenges are one of the many reasons why BCA advocates often suggest that it should be treated as only one of many sources of information for policy decisions.

---

[4] This discussion assumes that the values to be counted in the benefit-cost analyses are those held by the individuals who bear its costs and/or receive its benefits. Decisionmakers may, of course, have altruistic motives or other interests that differ from those of the affected individuals.

In the remainder of this article, we put aside these issues of the appropriate weighing of nonquantified impacts as well as equity and other-regarding preferences, discussing how we might measure individual willingness to pay for reductions in one's own mortality risk within the traditional BCA framework. While the most appropriate monetary values for nonmarket outcomes are uncertain, assigning values has the advantage of providing information on the intensity as well as the direction of preferences. In particular, this information can be used to compare the extent to which affected individuals prefer the policy outcomes to other ends they could achieve using the same resources. Our underlying assumption is that BCA represents a pragmatic approach – a tool box – for providing certain types of information on the impacts of policy options. Tools from other boxes provide important complementary and supplementary information for decisionmaking.

3.0 The Value per Statistical Life Tool

In BCA as typically conducted, we are interested in the willingness of those affected to trade income for the outcomes of concern. For mortality risks, we have no market that fully discloses their willingness to make this exchange. Perhaps the closest direct market data involves consideration of human capital. Under this approach, the value of averting mortality is considered equal to the value of lost production, commonly estimated by assuming that workers' pay equals the value of their marginal product. However, this approach does not fully capture the monetary value of risk reductions, leading to estimates that are substantially less than what individuals indicate they would be willing to pay.[5]

Given this lack of direct market evidence, analysts instead rely on data on related marketed goods or observed behavior ("revealed preferences") or data from survey research ("stated preferences").[6] Below, we first discuss confusion over the definition of the VSL, and then discuss gaps in the available research. Both pose problems for analysts, and lead to confusion over positive and normative issues.

3.1 Definitional Problems

Many policies lead to relatively small changes in health risks at the individual level, often expressed as "statistical cases" for ease of presentation. A statistical case, or statistical life, involves aggregating small risk changes across several individuals. Generally, analysts start with a risk assessment that estimates the impact of each policy option. For example, for the U.S. population, the annual likelihood of dying increases from about 10 in 10,000 to about 100 in 10,000 for each year of age between age 20 and age 65, conditional on surviving to that age

[5] For example, the present value of future earnings and household production is about $1.1 million in the U.S. at age 40; excluding nonwage (household) production decreases this value to about $0.8 million (using a 3 percent discount rate, Grosse 2003). In contrast, for workers (with an average age of about 40), Viscusi (2004) estimates the U.S. VSL as $5 million based on the trade-off between wages and job-related risks. (Both estimates are for the year 2000.)

[6] In the interest of brevity, we do not discuss approaches based on a constant value per statistical life year (VSLY) or on monetized quality-adjusted life years (QALYs). Problems related to the use of a constant VSLY are discussed in Hammitt (2007), Aldy and Viscusi (2007), Krupnick (2007), Cropper et al. (2007), and NAS (2008). The construction and use of QALY measures, and issues related to their inconsistency with the BCA framework, are discussed in Hammitt (2002) and IOM (2006).

(Arias 2010). The risk assessment may find that a policy will, on average, decrease this annual risk by 1 in 10,000 for each member of the potentially affected population (perhaps providing some information on how this risk reduction varies by year of age or other factors). If this population contains 10,000 individuals, this means that one less person is expected to die each year after the policy is implemented. However, risk assessors cannot predict in advance (nor necessarily determine afterwards) which individual's life will be (or has been) extended by the policy; the risk reduction is a "statistical" case – a sum of probabilities.

The calculation is straightforward:

1/10,000 risk reduction x 10,000 individuals annually
= 1.0 statistical case

(1)

Thus "saving" a statistical life is not the same as saving an identifiable individual from certain death.

In benefit-cost analysis, we are interested in individual willingness to pay (WTP) for these small risk changes within a defined time period.[7] We make decisions almost every day that reflect this trade-off between money and risk; e.g., choosing whether to accept a riskier job or buy a house in a safer neighborhood, or how much to spend on car safety features or on padding and helmets for sports.

Once we estimate individual WTP, the VSL is then calculated by dividing this WTP by the risk change (see Hammitt 2000). For example, if an individual is willing to pay $600 for a 1 in 10,000 reduction in his or her risk of dying in the current year, his or her VSL is calculated as:

$600 average individual WTP ÷ 1/10,000 annual risk change
= $6.0 million VSL

(2)

Alternatively, individual WTP for small risk reductions can be aggregated across a population:

$600 average individual WTP x 10,000 affected individuals annually
= $6.0 million VSL[8]

(3)

In either case, what analysts are really seeking is individual WTP for the 1 in 10,000 risk reduction; i.e., the $600. We express it as the VSL (i.e., the $6.0 million) largely for convenience.

The use of the VSL terminology had led to problems, however. Those not immersed in the technical details of these analyses often interpret it as the value of an individual's life; i.e., of

---

[7] Willingness to accept compensation may also be used for valuation; however, we refer to WTP throughout this article for simplicity because it is generally easier to measure, should be about equal to WTP under standard economic theory, and is more frequently used in BCA.

[8] As discussed in more detail in Hammitt and Treich (2007) and Cameron (2010), this approach is a shortcut that involves premature aggregation. Both risks and WTP are likely to vary by individual. Conceptually, we are interested in the sum of the individual's WTP for the risk reduction each individual is expected to receive. While this sum may differ from what results when we multiply the population average VSL by the total risk reduction, we generally lack the data needed to calculate this total on an individualized basis.

averting the death of an identifiable individual with certainty.[9] These problems have been particularly evident when U.S. government agencies have reduced the VSL used in their analyses (see Robinson 2007, Viscusi 2009). Although the smaller values reflected research on individual WTP for one's own risk reductions, they were interpreted as the "government" placing lower values on "life." Cameron (2010) investigates the problems caused by this terminology and suggests substituting the term " 'willingness to swap' (WTS) alternative goods and services for a microrisk reduction" in these risks (pp. 162-163).

3.2 Measurement Problems

The VSL is relatively well-studied: Viscusi and Aldy (2003) summarize over 50 studies that investigate the trade-offs between wages and job-related risks; Braathen et al. (2009) summarize over 70 surveys that investigate these values in the context of environmental, traffic, and public health risks. Despite this large number of studies, significant uncertainty remains, in particular related to how the VSL varies depending on individual characteristics (e.g., age, income), physical characteristics of the risk (e.g., latency, morbidity prior to death), and psychological responses to risk characteristics (e.g., controllability, voluntariness, fear or dread). In the discussion that follows, we focus primarily on how individual characteristics (age and income) affect the VSL, because it is within this realm where the most controversies have erupted over the appropriateness of "valuing life" in BCA.

Policy analysts typically apply similar VSL estimates across different types of risks and population subgroups. For example, U.S. Federal agencies, including the Environmental Protection Agency (EPA), the Department of Transportation, the Food and Drug Administration, and the Department of Homeland Security, generally rely on population-average estimates from studies of job-related risks, adjusting only for real income growth over time (not across population subgroups).[10] Such standardization means that the BCA can fall short of the goal of accurately reflecting the preferences of those affected by policies.

The reluctance to adjust the VSL for the age of those affected stems in part from a controversy that erupted over the so-called "senior discount:" the U.S. EPA's use of lower estimates for older individuals affected by air pollution in sensitivity analyses (see Robinson 2007, Viscusi 2009). This controversy led EPA to discontinue the use of VSL age adjustments, a stance that was ultimately incorporated into government-wide guidance (OMB 2003).

While there is some evidence that the VSL declines at older ages, newer work suggests that this relationship is uncertain (Hammitt 2007, Aldy and Viscusi 2007, Krupnick 2007). As a result, two U.S. expert panels more recently advised against making VSL age adjustments at this time (Cropper et al. 2007, NAS 2008), indicating that further research is needed. In addition, informal conversations with U.S. government staff indicate that they are reluctant to engage in battles over the use of age-differentiated values; experience with the senior discount debate

---

[9] For more discussion of identifiable vs. statistical lives, see Hammitt and Treich (2007).

[10] In addition, agencies also often adjust for latency and add cost-savings not included in the VSL estimate.

suggests that it will divert significant time and attention away from the primary goals of their policy initiatives.[11]

This lack of adjustment for age could lead to misinformation on the preferences of those affected, however, and could also lead to misallocation of resources if decisions were based solely on the BCA results. For example, in the U.S., mortality risk reductions account for most of the quantified benefits of air pollution regulations, and the majority of these risk reductions accrue to older individuals (see, for example, EPA 2010). If older individuals are in fact willing to pay less for risk reductions than the population average, this approach could lead to overinvestment in policies benefiting the elderly. However, in actuality, using lower values may at times have negligible effects on policy decisions. The benefits of many (but not all) air regulations exceed their costs by such a large amount that the VSL would need to be reduced substantially for costs to exceed benefits. More importantly, the Clean Air Act requires that the U.S. EPA establish standards based largely on consideration of risk and/or available technology, limiting the use of BCA.

Emerging research suggests, in contrast, that values for children are likely to exceed adult values (e.g., OECD 2010). In this case, the use of (adult) population averages could lead to underinvestment in programs that benefit children. These larger values have not yet been widely used in policy analysis, and it is unclear whether increasing the VSL for children would lead to the same sort of controversy as decreasing the VSL for older individuals.

For income, the research evidence is more consistent, at least among relatively wealthy populations. Several studies suggest that, in the U.S., a one percent change in income is likely to lead to about a 0.4 to 0.6 percent change in the VSL (e.g., EPA 1999, Viscusi and Aldy 2003).[12] While several U.S. agencies use these elasticity estimates to adjust the VSL for changes in population-wide real income over time, none make adjustments for cross-sectional income differences. Instead, the VSL is based on the average income of the individuals included in the underlying valuation studies, regardless of the income levels of those affected by the policies.

The effect of income is more uncertain when extrapolating across individuals or countries with significantly different income levels. There have been few, if any, VSL studies conducted in low-income countries. Research on the income elasticity of the VSL suggests that it may be well in excess of 1.0 when extrapolating from high- to low-income populations (Hammitt and Robinson 2010). Given uncertainty in this elasticity, such extrapolation can lead to estimates that vary by orders of magnitude. For example, extrapolating from a U.S. VSL of $6.3 million (where per capita gross national income is about $46,000) to Tanzania (where per capita gross national income is about $1,200) leads to VSL estimates that range from about $165,000 to about $4,300, depending on whether an elasticity of 1.0 or 2.0 is used (2007 dollars, based on purchasing power parity). As discussed in Hammitt and Robinson (2010), the lower value falls below rough estimates of the present value of future income (human capital) and consumption, contrary to theory. This comparison suggests that these latter estimates should be used instead of the extrapolated VSL as a lower bound. Regardless, at present, analysts will need to test to effects of

---

[11] Interestingly, VSL age adjustments appear less controversial in other countries. For example, while the current Canadian guidance for impact assessment does not discuss age adjustments (Treasury Board 2007), Canadian agencies have included these adjustments in some regulatory analyses (e.g., Chestnut et al. 1999) without the sort of public outcry that resulted in the U.S. Age adjustments have also been included in some World Bank analyses, such as Lvovsky et al. (2000) and World Bank (2002).

[12] The formula for calculating the income-adjusted VSL is: $VSL_b=VSL_a*(Income_b/Income_s)^{elasticity}$.

a range of elasticities on their results, which will lead to substantial uncertainty about the extent to which the policies are cost-beneficial.

In the U.S., the use of estimates based on averages is often viewed as providing more equitable treatment, or equal protection, for different groups in policy decisions. However, whether this approach is in fact equitable depends on how one views the incorporation of individual preferences into these analyses. If the individuals affected have preferences for spending on their own risk reductions that differ from the population average, an analysis based on the average VSL will not reflect their preferences. In addition, these population averages are anchored in the distribution of health, income and other characteristics that existed at the time of the underlying studies, and this distribution will change over time.

4.0 The Way Forward

The discussion above raises issues both inside and outside of the BCA tool box. First, it is clear that BCA needs to be supplemented by information on moral values, including equity and other types of other-regarding preferences. While new work may eventually allow the BCA framework to expand beyond its current focus on private preferences, we do not yet have a good understanding of how social preferences differentially affect individual WTP for mortality risks or other policy outcomes. Second, better communication of what the VSL does, and does not, measure is needed. Changing its name may be helpful; analysts also need to state more clearly that it measures individual willingness to exchange income for small changes in their own health risks. Finally, more research is needed on how this willingness to pay varies depending on individual characteristics as well as the risks themselves.

REFERENCES

Adler M.A. 2008. “Risk Equity; A New Proposal.” *Harvard Environmental Law Review*. 32, 1-47.

Aldy, J.E. and W.K. Viscusi. 2007. “Age Differences in the Value of Statistical Life: Revealed Preference Evidence.” *Review of Environmental Economics and Policy*. 1(2), 241-260.

Arias, E. 2010. “United States Life Tables, 2006.” *National Vital Statistics Reports*, 58(21). Centers for Disease Control and Prevention, U.S. Department of Health and Human Services.

Bergstrom, T.C. 2006. “Benefit-Cost in a Benevolent Society.” *American Economic Review*. 96(1), 339-351.

Braathen, N.A., H. Lindhjem, and S. Navrud. 2009. *Valuing Lives Saved from Environmental, Transport and Health Policies: A Meta-Analysis of Stated Preference Studies*. ENV/EPOC/WPNEP(2008)10/FINAL. Paris: Organization for Economic Co-operation and Development.

Cameron, T.A. 2010. “Euthanizing the Value of a Statistical Life.” *Review of Environmental Economics and Policy*. 4(2), 161-178.

Chestnut, L.G., D. Mills, and R.D. Rowe. 1999. *Air Quality Valuation Model Version 3.0 (AQVM 3.0) Report 2: Methodology*. Prepared for Environment Canada and Health Canada by Stratus Consulting Incorporated.

Clinton, W.J. 1993. "Executive Order 12286: *Regulatory Planning and Review*."

Cropper, M.L. et al. 2007. "SAB Advisory on EPA's Issues in Valuing Mortality Risk Reduction." Memorandum from the Chair, Science Advisory Board, and the Chair, Environmental Economics Advisory Committee, to EPA Administrator Stephen L. Johnson. EPA-SAB-08-001.

Farrow, S. 2009. "Incorporating Equity in Regulatory and Benefit-Cost Analysis." Prepared for "New Ideas for Risk Regulation" conference. (Available at: http://www.rff.org/Events/Pages/New-Ideas-for-Risk-Regulation.aspx.)

Graham, J.D. 2008. "Saving Lives through Administrative Law and Economics." *University of Pennsylvania Law Review*. 157, 395-540.

Grosse, S.D. 2003. "Appendix I: Productivity Loss Tables." In *Prevention Effectiveness: A Guide to Decision Analysis and Economic Evaluation.* (A.C. Haddix, S.M. Teutsch, and P.S. Corso, eds.) New York: Oxford University Press.

Hammitt, J.K. 2000. "Valuing Mortality Risk: Theory and Practice." *Environmental Science and Technology*. 34, 1396-1400.

Hammitt, J.K. 2002. "QALYs versus WTP." *Risk Analysis*.22(5), 985-1001.

Hammitt, J.K. 2007. "Valuing Changes in Mortality Risk: Lives Saved versus Life Years Saved." *Review Of Environmental Economics and Policy*. 1(2), 228-240.

Hammitt, J.K. 2009. "Response: *Saving Lives: Benefit-Cost Analysis and Distribution*." *University of Pennsylvania Law Review PENNumbra*. 157, 189-199.

Hammitt, J.K. and L.A. Robinson. 2010. "The Income Elasticity of the Value per Statistical Life: Transferring Estimates Between High and Low Income Populations." *Journal of Benefit-Cost Analysis.*

Hammitt, J.K. and N. Treich. 2007. "Statistical vs. Identified Lives in Benefit-Cost Analysis," *Journal of Risk and Uncertainty*. 35(1), 45-66.

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

Independent Evaluation Group (IEG). 2010. *Cost-Benefit Analysis in World Bank Projects.* International Bank for Reconstruction and Development and the World Bank.

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

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

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

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

Kaldor, N. 1939. "Welfare Propositions of Economics and Interpersonal Comparisons of Utility." *Economic Journal.* 49(195), 549-552.

Krupnick, A. 2007. "Mortality-risk Valuation and Age: Stated Preference Evidence." *Review of Environmental Economics and Policy.* 1(2), 261-282.

Loomis, J.B. 2009. *Importance of Incorporating Distributional Issues in Benefit Cost Analysis.* Prepared for the Benefit-Cost Analysis Center, University of Washington, with funding from the John D. and Catherine T. MacArthur Foundation. (Available at: http://evans.washington.edu/research/centers/benefit-cost-analysis/principles-and-standards-draft-papers.)

Lvovsky, K., G. Hughes, D. Maddison, B. Ostro, and D. Pearce. 2000. *Environmental Costs of Fossil Fuels: A Rapid Assessment Method with Application to Six Cities.* The World Bank Environment Department, Pollution Management Series, Paper Number 78.

National Academy of Sciences (NAS). 2008. *Estimating Mortality Risk Reduction and Economic Benefits from Controlling Ozone Air Pollution*. Committee on Estimating Mortality Risk Reduction Benefits from Decreasing Tropospheric Ozone Exposure. Washington D.C.: National Academies Press.

Organization for Economic Co-operation and Development (OECD). 2009. *Indicators of Regulatory Management Systems.*

Organization for Economic Co-operation and Development (OECD). 2010. *Valuation of Environmental Health Risks For Children (Draft Report).* ENV/EPOC/WPNEP(2008)9/FINAL.

Robinson, L.A. 2007. "How U.S. Government Agencies Value Mortality Risk Reductions." *Review of Environmental Economics and Policy.* 1(2), 283-299.

Robinson, L.A. and J.K. Hammitt. 2010. *Behavioral Economics and the Conduct of Benefit-Cost Analysis: Towards Principles and Standards.* Prepared for the Benefit-Cost Analysis Center, University of Washington, with funding from the John D. and Catherine T. MacArthur

Foundation. (Available at: http://evans.washington.edu/research/centers/benefit-cost-analysis/principles-and-standards-draft-papers.)

Schwartz, J.A. 2010. *52 Experiments with Regulatory Review: The Political and Economic Inputs into State Rulemakings.* New York University School of Law, Institute for Policy Integrity Report Number 6.

Sunstein, C.R. 2007. "Willingness to Pay vs. Welfare." *Harvard Law and Policy Review.* 1, 303-330.

Treasury Board of Canada, Secretariat. 2007. *Canadian Cost-Benefit Analysis Guide: Regulatory Proposals (Interim).*

U.S. Environmental Protection Agency (EPA). 1999. *The Benefits and Costs of the Clean Air Act, 1990 to 2010*. EPA 410-R-99-001.

U.S. Environmental Protection Agency (EPA). 2010. *The Benefits and Costs of the Clean Air Act: 1990 to 2020*

U.S. Office of Management and Budget (OMB). 2003. *Circular A-4, Regulatory Analysis.*

Viscusi, W.K. 2004. "The Value of Life: Estimates with Risks by Occupation and Industry." *Economic Inquiry*. 42(1), 29-48.

Viscusi, W.K. 2009. "The Devaluation of Life." *Regulation & Governance.* 3, 103-127.

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

World Bank. 2002. *Improving Air Quality in Metropolitan Mexico City*. Washington, D.C: The World Bank Policy Research Working Paper Number 2785.

Zerbe, R. 2009. *Ethical Benefit-Cost Analysis as Art and Science: Ten Rules for Benefit-Cost Analysis.* Prepared for the Benefit-Cost Analysis Center, University of Washington, with funding from the John D. and Catherine T. MacArthur Foundation. (Available at: http://evans.washington.edu/research/centers/benefit-cost-analysis/principles-and-standards-draft-papers.)