Skills of the trade: valuing health risk reductions in benefit-cost analysis

Lisa A. Robinson, James K Hammitt

Lisa A. Robinson* and James K. Hammitt

Skills of the trade: valuing health risk reductions in benefit-cost analysis

Abstract: Many public policies and private actions affect the risk of injury, illness, or death, yet changes in these risks are not easily valued using market prices. We discuss how to value these risk reductions in the context of benefit-cost analysis. We begin with a pragmatic focus, describing the analytic framework and the approaches currently used for valuation, including estimates of willingness to pay, cost of illness, and monetized quality-adjusted life years. We then turn to some conceptual issues that illustrate areas in need of further exploration.

Keywords: benefit-cost analysis; cost of illness; health risk valuation; value per statistical life; willingness to pay; quality-adjusted life year.

─────────
*Corresponding author: Lisa A. Robinson, Senior Fellow, Mossavar-Rahmani Center for Business and Government, Harvard Kennedy School, USA, e-mail: lisa.a.robinson@comcast.net
James K. Hammitt: Harvard University (Center for Risk Analysis), USA and Toulouse School of Economics (LERNA-INRA), France

1 Introduction

While many public policies and private actions affect the risks of injury, illness, or death, changes in these risks are difficult to value. How to best derive these values depends on the analytic goals and is often constrained by the data available. We discuss valuation in the context of benefit-cost analysis. We begin with a pragmatic focus, introducing the approaches currently used and their implementation. We then turn to some conceptual issues that illustrate areas in need of further exploration.

2 General framework

While benefit-cost analysis provides a framework for exploring a variety of issues, its primary goal is to estimate the extent to which the aggregate benefits of a policy exceed its costs, often comparing alternatives to identify which option (if any) is likely to be most economically efficient. Typically, these analyses sum

benefits and costs over the affected population. How the positive and negative impacts are distributed across members of society, and the equity of this distribution, is usually considered separately.

Benefit-cost analyses are rarely the sole determinant of policy decisions given that some impacts cannot be quantified, policymakers and stakeholders care about equity as well as efficiency, and policy choices are often constrained by law and other factors. Moreover, because of uncertainty about the predicted effects of the policy as well as the valuation of those effects, benefit-cost analysis may be unable to distinguish among policy alternatives that are substantially similar (e.g., involve small differences in pollution control levels). If well-conducted, however, these analyses provide a well-established framework for carefully exploring the likely outcomes of policy options and yield important information on how these outcomes are valued by the affected population.

Conventionally, benefit valuation is based on neoclassical welfare economics. Individuals are assumed to derive utility (which may be interpreted as a sense of well-being or satisfaction) from the goods and services they consume. Each individual is assumed to be the best judge of his or her own welfare. Thus, one's willingness to exchange money for various goods and services can be used to measure the utility one receives from their consumption. Within this framework, the monetary value of a risk reduction is most appropriately defined as the change in wealth that has the same effect on one's utility as the risk change.

For a policy that reduces health risks, this trade-off is represented by individual willingness to pay (WTP): the maximum amount of money an individual would voluntarily exchange to obtain the improvement, given his or her budget constraints.[1] WTP is a different concept than cost or price. The resource cost of producing a good or service may be greater or less than WTP, and its price is a market outcome determined by the interaction of consumers and suppliers. If the cost of production exceeds WTP for all consumers, then no one would be willing to buy the good and it would not be produced. If the good is produced and sold in a competitive market, then the producers' supply cost is presumably no higher than the market price and the consumers' WTP is no lower than the price.

For outcomes that are not traded in markets, economists estimate WTP using stated or revealed preference methods. Stated preference methods typically employ

---

1 This trade-off also can be represented by individual willingness to accept compensation (WTA): the smallest amount an individual would accept to forego the improvement. WTP is the compensating variation and WTA is the equivalent variation for the improvement. In the sections that follow, we refer to WTP for convenience, because WTP studies dominate the health risk valuation literature and policy analyses typically address improvements from the status quo rather than compensation for forgoing the improvement. We discuss the distinction between WTP and WTA in the final section.

survey techniques to ask respondents about their WTP for the outcome of concern. They include contingent valuation surveys, which elicit WTP for the scenario(s) described in the survey. They also include choice experiments (or conjoint analyses) which present respondents with several scenarios involving outcomes with differing attributes and prices. Estimates of WTP are derived from the way in which respondents rank, rate, or construct equivalent sets of alternatives. Stated preference methods are attractive because researchers can tailor them to directly value the outcome(s) of concern; i.e., the survey can describe particular types of health risks that result from specific causes and also describe the characteristics of those affected. However, conducting a study that yields accurate and reliable results requires careful design and implementation. A key concern is that respondents may have little incentive to respond accurately, since the payment is hypothetical.

Revealed preference methods involve valuing nonmarket goods based on observed behaviors or prices and preferences for related market goods. For example, wage-risk (or hedonic-wage) studies examine the additional compensation associated with jobs that involve higher risks of fatal or nonfatal injury, using statistical methods to separate the effects of these risks from the effects of other job and personal characteristics.[2] While this indirect use of market data has the advantage of relying on actual transactions, it may be difficult to find a market good that can be used to value the outcome of concern for a particular policy analysis. Because both methods have advantages and limitations, the choice often depends on the characteristics of the outcome being valued; e.g., whether it can be modeled as an attribute of a market good in a revealed preference study. Comparing the results from both methods, or using a combined approach, can be informative where feasible.

Typically, policy analysts rely on existing valuation studies for benefit-cost analyses, rather than incurring the substantial time and expense associated with conducting new primary research. This approach, referred to as "benefit transfer," requires careful review of the literature to identify high-quality studies that are suitable for use in a particular context. "Quality" can be evaluated by considering the likely accuracy and reliability of the data and methods used, referencing guidance on best practices.[3] "Suitability" or "applicability" involves considering

---

2 Another revealed preference approach considers averting behaviors; i.e., defensive measures or consumer products used to protect against perceived health risks (see, for example, reviews by Blomquist, 2004; Viscusi and Aldy, 2003). These studies are applied infrequently in benefit-cost analysis due to concerns about their limitations, including the difficulty of estimating the size of the risk change associated with many behaviors and the need to separately estimate the value of key inputs such as the time spent in the activity.

3 For example, for US regulatory analysis, best practice guidance is provided in U.S. Office of Management and Budget (OMB) (2003) and US Environmental Protection Agency (EPA) (2010b). Boardman, Greenberg, Vining, and Weimer (2011) also discuss related issues.

the similarity of the risks and the populations affected. There are no firm guidelines for determining quality or suitability; benefit transfer relies heavily on the informed judgment of the analyst, and requires clear disclosure and discussion of related uncertainties and their implications for decision-making.

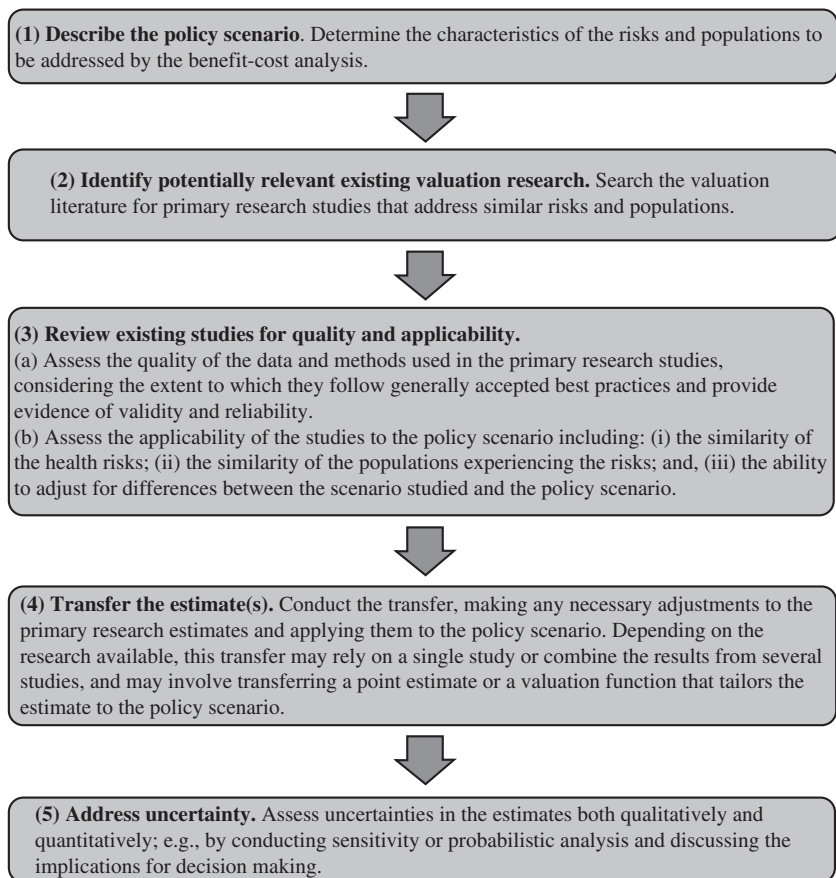Benefit transfer generally consists of five steps, as listed in Figure 1.

(1) Describe the policy scenario. Determine the characteristics of the risks and populations to be addressed by the benefit-cost analysis.

(2) Identify potentially relevant existing valuation research. Search the valuation literature for primary research studies that address similar risks and populations.

(3) Review existing studies for quality and applicability.
(a) Assess the quality of the data and methods used in the primary research studies, considering the extent to which they follow generally accepted best practices and provide evidence of validity and reliability.
(b) Assess the applicability of the studies to the policy scenario including: (i) the similarity of the health risks; (ii) the similarity of the populations experiencing the risks; and, (iii) the ability to adjust for differences between the scenario studied and the policy scenario.

(4) Transfer the estimate(s). Conduct the transfer, making any necessary adjustments to the primary research estimates and applying them to the policy scenario. Depending on the research available, this transfer may rely on a single study or combine the results from several studies, and may involve transferring a point estimate or a valuation function that tailors the estimate to the policy scenario.

(5) Address uncertainty. Assess uncertainties in the estimates both qualitatively and quantitatively; e.g., by conducting sensitivity or probabilistic analysis and discussing the implications for decision making.

Figure 1 Benefit transfer.

Most benefit-cost analyses value fatal and nonfatal risk reductions separately, because relatively few primary research studies yield values that account for both types of effects. This is generally referred to as the "damage function" approach, which involves identifying the benefits of the policy, valuing them separately, then aggregating them to determine the net effect – while taking care to avoid double-counting. In the valuation literature, mortality risk reductions are relatively well-studied while nonfatal risks have received less attention. As a

result, analysts at times use monetary measures other than WTP in the latter case. We discuss approaches for valuing these risks in more detail below.

3 Valuing mortality risk reductions

Conventionally, the value of mortality risk reductions is expressed as the value per statistical life (VSL). Many policies lead to relatively small changes in risks at the individual level, often expressed as "statistical cases." A statistical case, or a statistical life, involves aggregating small risk changes across individuals.

Analysts typically start with a risk assessment that estimates the impact of each policy option on the risk of dying in a particular time period. For example, for the US population, the annual likelihood of dying at each year of age increases from about 10 in 10,000 to about 100 in 10,000 between age 20 and age 65 years, conditional on surviving to that age (Arias, 2010). The risk assessment might find that a policy would, on average, decrease this annual risk by 1 in 10,000 for each member of the population, perhaps providing information on how the risk change varies by age or other factors. If this population contains 10,000 individuals, this means that one fewer person is expected to die each year after the policy is implemented.[4] 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 × 10,000 individuals annually=1.0 statistical case

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

For valuation, the starting point is individual WTP for the small risk change within the defined time period. VSL is measured in units of dollars per statistical life saved, which can be calculated by dividing this WTP by the risk change (see Hammitt, 2000).[5] 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:

---

4 This is true only in the short run because everyone eventually dies. Policies extend lives rather than saving them; the deaths prevented by the policy are delayed to a later date and possibly shifted to another cause.

5 VSL can also be viewed as aggregating individual WTP across a population; i.e., a $600 average individual WTP×10,000 affected individuals=$6.0 million. Because both the risk reduction and WTP are likely to vary by individual, it is more accurate to first determine each individual's WTP for the risk reduction he or she would receive then aggregate. However, the data needed to calculate these values on an individualized basis are often unavailable.

$$\$600\ WTP \div 1/10{,}000\ risk\ change = \$6.0\ million\ VSL$$

The key parameter is individual WTP for the 1 in 10,000 risk reduction (i.e., the $600); it is expressed as the VSL (i.e., the $6.0 million) largely for convenience. It is not the value of saving an individual's life with certainty. This distinction between the value of small risk changes and the value of a "life" has led to much confusion (Cameron, 2010; Robinson, 2007; Viscusi, 2009), in turn leading to proposals to change the VSL terminology to the "value of mortality risks" (VMR) or the "value of risk reduction" (VRR) (EPA, 2010a; Kling & Swackhamer, 2011). In the UK, the conventional term is "value per prevented fatality" (VPF) (HM Treasury, 2003).

An alternative approach to valuing mortality risk reduction is to use the value per statistical life year (VSLY). This approach assumes that an individual's WTP to reduce current mortality risk is dependent on the corresponding gain in life expectancy, i.e., the reduction in his or her chance of dying this year multiplied by his or her life expectancy conditional on surviving the year. (Future life years are typically discounted to account for time preferences.) VSLY has the advantage of acknowledging that lives "saved" by a policy are only extended; the amount of extension will depend on characteristics of the affected individuals as well as the risk reduction provided by the intervention.

VSLY is typically estimated by dividing a VSL estimate for a population by the average (discounted) remaining life expectancy for that population.[6] For example, if a study yields a mean VSL of $6.0 million, the mean individual in that study is age 40, and mean (population) life expectancy for an individual who reaches age 40 is an additional 35 years, then estimated VSLY would be $279,000 using a 3% discount rate (or $171,000=$6.0 million divided by 35 years using a 0% discount rate). Using this approach, the benefit of reducing mortality risk is calculated as the gain in life years in the population multiplied by the VSLY.

The assumption that VSLY is independent of the gain in life expectancy implies that VSL is proportional to remaining life expectancy. This is not supported by empirical research, however. Much of this research focuses on the extent to which an individual's WTP for risk reduction varies with age, which is closely correlated with life expectancy. As discussed in Aldy and Viscusi (2007), Hammitt (2007) and Krupnick (2007), both theory and empirical evidence suggest that VSL may change with age, with some models suggesting it increases over younger ages and decreases at older ages. The extent to which VSL increases, decreases, or remains constant with age is uncertain, with different studies yielding substantially different results. As a result, two US expert panels recommended against the use of

---

6 As noted below, an increasing number of studies directly estimate the value of life extension, but more work is needed to establish robust estimates suitable for use in policy analysis.

a constant VSLY in policy analysis (Cropper & Morgan, 2007; National Academies, 2008), suggesting that more research is needed.

Researchers have conducted well over a hundred VSL studies. Recent reviews include Viscusi and Aldy (2003) for revealed preference research and Lindhjem, Navrud, Braathen, and Biausque (2011) for stated preference research. Historically, this literature was dominated by revealed preference studies which usually address job-related risks. However, this is no longer true, as the number of stated preference studies has increased rapidly, addressing environmental, traffic safety, and other risks.[7] Despite conventional wisdom that stated preference studies yield higher benefit estimates than revealed preference studies (e.g., because the payment is hypothetical), the opposite tends to be true for VSL. While VSL estimates from both revealed and stated preference studies vary widely, stated-preference estimates tend to be smaller.

Because the value of mortality risk reductions is often important in regulatory decision-making, many government agencies have issued related guidance. In the US, the Office of Management and Budget (OMB) (2003) allows agency discretion, while indicating that research suggests that the VSL is generally between roughly $1 million and $10 million (no dollar year reported). Two regulatory agencies have developed formal guidance. The US Environmental Protection Agency (EPA) recommends a central estimate of $7.4 million (in 2006 dollars) (EPA, 2010b) and the US Department of Transportation (DOT) recommends a central estimate of $6.2 million (in 2011 dollars) (DOT, 2011). Both suggest these estimates should be updated to reflect inflation and real income growth over time. EPA also adjusts for any lag between changes in exposure and changes in risk, and some DOT components add averted costs (such as those associated with medical treatment) to the base VSL. These agencies are now considering altering their approaches to reflect new research as well as evolving standards for best practices (Cropper, Hammitt, and Robinson, 2011; EPA, 2010a; Kling & Swackhamer, 2011).

Other countries have also developed guidance. Canada recommends a value of $6.11 million (2004 Canadian dollars, Treasury Board, 2007), similar to the US agencies.[8] Both Canada and the US rely primarily on revealed preference studies that consider the trade-off between wages and mortality risks.

Somewhat smaller values are used in other countries, in part because they rely on stated preference studies which tend to yield lower VSLs than revealed prefer-

---

7 Viscusi and Aldy (2003) review over 60 wage-risk studies conducted globally. The Lindhjem et al. database [described in more detail in Organisation for Economic Co-operation and Development (OECD), 2012], includes 74 stated preference surveys from around the world, of which 37 address traffic safety, 30 address public health or health risks from unidentified causes, and 17 address environmental risks. (Some surveys address more than one risk type.)

8 One Canadian dollar was worth roughly 1 US dollar at the time this manuscript was prepared.

ence studies, as summarized in Table 1. The UK's guidance references a value of $1.145 million used by its Department of Transport (2000 British pounds, HM Treasury, 2003), about $1.8 million.[9] It notes that its Health and Safety Executive has doubled this value when addressing cancer risks, but that evidence is lacking to support this adjustment. While the UK's guidance does not indicate the basis for its estimate, a recent study that explores updating this approach notes that it is based on stated preference research, supplemented by estimates of averted costs (NERA Economic Consulting, 2011). The European Commission (2009) recommends values of 1 million to 2 million Euros (or $1.3 million to $2.6 million), but notes that context-specific values may be used in lieu of this default. This European guidance does not reference the source of its estimates. Finally, the Organisation for Economic Cooperation and Development (OECD) (2012) recommends a value of $1.5 million to $4.5 million (2005 dollars) for analyses that address all OECD countries, based on stated preference research, and provides guidance on developing estimates for individual member countries. While few VSL studies have been conducted in lower income countries, analysts often extrapolate values from wealthier countries, adjusting for income differences (Hammitt and Robinson, 2011).

Although we expect values to differ across countries due to income as well as social and cultural factors (such as attitudes towards risks and health system characteristics), most of the variation in the values above reflects the choice between relying on revealed preference or stated preference studies. More work is needed to understand the reasons why these methods yield such different values.

In the near term, analysts may find it easiest to review the most current version of the government guidance referenced above and apply the recommended values. However, this approach ignores the opportunity to use studies that may be better tailored to the risks and populations associated with the policy; e.g., to rely on studies of air pollution-related risks rather than job-related risks when assessing related policies (Robinson and Hammitt, 2011a).[10] However, such

---

9 All values are converted to US dollars based on current exchange rates using the Google currency converter (which relies on Citibank data), as viewed May 2, 2012.

10 In practice, VSL is generally not adjusted for income differences within a particular population, but is adjusted for population-wide income changes over time or income differences across countries. In the U.S., there has been substantial debate about adjusting VSL to reflect age differences, particularly over the use of lower values for older individuals (i.e., the "senior discount") (Robinson, 2007, 2009; Viscusi, 2009). While this lack of adjustment for population characteristics is often framed as an equity issue, applying the same value regardless of the characteristics of those affected ignores their preferences for spending on health risk reductions rather than on other goods and services.

Table 1 Examples of recommended VSL estimates.

Country/Agency	VSL estimates	Basis	Adjustments
US Environmental Protection Agency (2010b)	**$7.4 million** (Standard deviation: $4.7 million, 2006 dollars)	Primarily revealed preference studies of job-related risks	Inflation and real income growth, latency or cessation lag
US Department of Transportation (2011)	**$6.2 million** (Standard deviation: $2.8 million, 2011 dollars)	Primarily revealed preference studies of job-related risks	Inflation and real income growth, averted costs
Canada (Treasury Board, 2007)	**$6.11 million**[a] (2004 Canadian dollars)	Primarily revealed preference studies of job-related risks	Inflation
UK (HM Treasury, 2003)	**£1.145 million**[b] (2000 British pounds)	Stated preference studies plus averted costs	Inflation and real income growth
European Commission (2009)	**1 million to 2 million Euros**[c] (Year not reported)	Not reported	May use more context-specific estimates
Organisation of Economic Cooperation and Development (2012)	**$1.5 million to $4.5 million** (2005 dollars)	Stated preference studies	Inflation and real income growth

Notes: [a]$6.17 million US dollars based on current exchange rate. [b]$1.8 million US dollars based on current exchange rate. [c]$1.3 million to $2.6 million US dollars based on current exchange rate.

tailoring may require more primary research, and will often involve relying on stated preference methods, emphasizing the importance of improving our understanding why such research often results in lower values.

4 Valuing morbidity and injury risk reductions

The approaches typically used to value nonfatal risk reductions are more diverse than the approaches used for mortality risks, due to significant gaps in the research literature. Given the large variety of health risks that may be of interest in benefit-cost analysis, in this section we focus on sources that analysts can review to identify appropriate WTP studies. We then discuss alternative measures often used as proxies when WTP estimates are not available.

4.1 Willingness to pay

We are not aware of a recent, comprehensive review of the valuation literature for nonfatal risks, although some researchers have reviewed portions of this literature. For nonfatal illnesses, respiratory and cardiovascular conditions associated with air pollution appear to have been subject to more study than other conditions; a recent review is provided in Hunt and Ferguson (2010). For acute effects, a meta-analysis by Van Houtven, Powers, Jessup, and Yang (2006) lists 17 contingent valuation studies from nine countries published from the late 1970s through the early 2000s.

For injuries, a few WTP studies address a single type (generally relatively severe, such as a spinal cord or brain injury) while others bundle a range of injuries into a few categories. For example, Viscusi and Aldy (2003) identify 40 wage-risk studies that define the nonfatal risk variable as either the overall injury rate, the rate for injuries severe enough to result in a lost workday, or the rate of lost workdays. Whether these studies are applicable to a particular benefit-cost analysis will depend on the nature of the risks affected by the policy.

Analysts will generally want to search resources, such as the EconLit and EVRI bibliographic databases and the comprehensive bibliography of contingent valuation studies compiled in Carson (2012), to identify WTP studies for the illnesses or injuries of concern.[11] Sources of best practice guidance that can be adapted for use as criteria in reviewing these studies are numerous: examples include Alberini and Kahn (2006), Bateman et al. (2002), Champ, Boyle, and Brown (2003), EPA (2010b) and Freeman (2003).

Because it is often difficult to identify WTP studies of reasonable quality that are applicable to the risks addressed by a particular benefit-cost analysis, analysts frequently use other measures as proxies. Below, we discuss the two most commonly used (alone or in combination): cost of illness (COI) and monetized quality-adjusted life years (QALYs).

4.2 Direct and indirect costs of illness

The costs of illness are the real resource costs of incurred cases of illness, injuries, and deaths. Such estimates are often used to compare the costs of differ-

---

11 EconLit: http://www.aeaweb.org/econlit/index.php; EVRI: https://www.evri.ca/Global/HomeAnonymous.aspx.

ent health conditions or alternative treatments for a particular condition. They also may be used to estimate the costs averted by policy interventions. We first discuss how these estimates are constructed, then describe their relationship to WTP.

COI estimates typically include two components: direct medical costs and indirect productivity losses. Medical costs may include those paid by patients, their families, and/or third parties such as insurance companies and employers. They typically include costs associated with physician services, medication, hospital stays, and other treatment-related activities. Additional costs, such as those related to processing insurance claims, are at times included.

Many studies also consider the indirect costs associated with lost productivity. These costs may stem from absence from work or from decreased productivity while at work, and may include other costs such as those associated with idling assets or training replacement workers. Compensation is generally used to value productive time, assuming that workers are paid the value of their marginal product; this is referred to as the "human capital" approach. While the measure of compensation varies across studies, it often includes both pre-tax wages and fringe benefits (see for example, Grosse, Krueger, and Mvundura, 2009). Some studies consider unpaid work (such as volunteer and household services), in which case the value of lost productivity is often based on the compensation of employed workers with similar responsibilities.

Developing and applying COI estimates in benefit-cost analysis requires first identifying the types of costs that may be averted by the policy options and then locating sources that estimate these costs. Many of the available data sources have significant limitations in this context. Policies often result in small changes in overall incidence, yet COI studies usually focus on average costs per case rather than on marginal costs per case averted. Policies also typically prevent a illness or injury from occurring, averting related costs over many years for severe injuries and chronic illnesses. However, many COI studies are prevalence- rather than incidence-based, focusing on costs within a particular year rather than over time. One important resource is a series of papers commissioned by the US Agency for Healthcare Research and Quality and the National Cancer Institute, which were published in a special issue of *Medical Care* (Yabroff, Brown, Lawrence, Barnett, & Lipscomb, 2009). These papers are a rich source of information on best practices and available data sources.[12]

---

12 In particular, Lund et al. (2009) provides a comprehensive inventory of sources of cost estimates. For injuries, Finkelstein, Corso, and Miller (2006) present incidence-based estimates

As noted earlier, averted costs per case are not equivalent to individual WTP for a risk reduction. The logic behind using costs to value benefits is that, if a policy allows society to avoid these costs, then the benefits are at least equal to the averted expenditures. However, this need not be true in all cases. For example, the relationship between medical costs and WTP is distorted by the presence of insurance. For insured conditions, an individual's out-of-pocket costs are likely to understate the costs averted. Adding costs paid by third parties may overestimate WTP for that treatment (although not necessarily for averting the condition itself), as individuals may receive treatment that they would not have willingly funded themselves.[13]

In addition, the valuation of time losses in these analyses focuses on lost productivity, relying on simplistic assumptions regarding the operation of the labor market and individual preferences for work and leisure. From a market perspective, production may not be noticeably affected if an ill worker is quickly replaced by someone who would be otherwise unemployed. From the individual's perspective, lost income will reflect post-tax rather than pre-tax earnings and will depend on factors such as the availability of sick pay or disability insurance. In addition, this focus on monetary compensation ignores other aspects of the utility or disutility associated with lost work time, which may be substantial due to imperfections in the labor market. For example, requirements to work 40 h per week limit individuals' ability to achieve their preferred balance between work and leisure time.

Furthermore, these costs reflect outcomes that differ in some respects from the policy outcomes of concern. For example, costs do not fully incorporate the value of pain and suffering nor other quality of life impacts associated with a health impairment and its affects on work and other activities. Risk aversion may also affect the values: costs are measured *ex post*, whereas risk reductions are typically valued from an *ex ante* perspective in benefit-cost analysis.

While COI estimates are believed to often understate WTP, it can be difficult to demonstrate the extent to which this is the case. Such comparisons require controlling for other factors that influence the values, including both how the health outcomes are defined in each case and the methodological choices made by the researchers. Thus the direction and magnitude of the bias that results is difficult to estimate.

---

of lifetime costs per case, and the Centers for Disease Control and Prevention (CDC) provide an injury cost calculator as part of the Web-based Injury Statistics Query and Reporting System (WISQARS) website (http://www.cdc.gov/injury/wisqars/index.html).

13 The cost of this treatment is nevertheless a social cost, and savings in this cost can be added to estimates of WTP when calculating the total benefits of the policy as long as care is taken to avoid double-counting.

Relying on COI estimates where suitable WTP estimates are lacking may be a reasonable approach, since it provides information on the costs that may be averted by the risk reduction. However, given the limitations of this approach, it is important that the analyst clearly discuss related uncertainties and the implications for decision-making. In addition, when applying these estimates, analysts often calculate the number of statistical cases of illness or injury averted then multiply by the average cost per case. Given that both risks and costs may vary depending on the characteristics of those affected by the policy, a more disaggregate approach may be desirable where possible. Such an approach would involve calculating costs averted for each subgroup then multiplying the results by the risk change likely to be experienced by that group, and exploring the effects of uncertainty.

4.3 Quality-adjusted life years

The QALY is a nonmonetary measure that integrates the duration and severity of injury or illness. QALYs were originally developed for use in ranking and prioritizing public health problems and in analyzing the cost-effectiveness of health policy and medical treatment decisions (Zeckhauser and Shepard, 1976). They are also widely used to compare health status across individuals or population groups.[14] In these contexts, QALYs are generally not assigned a monetary value, but monetization is needed to apply these estimates in benefit-cost analysis.[15] Below, we first discuss how QALY estimates are constructed, then discuss their monetary value and relationship to WTP.

4.3.1 Measuring QALYs

QALYs are derived by multiplying the amount of time an individual spends in a health state by a measure of the health-related quality of life (HRQL) associated with that state. HRQL is measured on a scale anchored at 0 and 1, where 1 corresponds to full health and 0 to a state that is as bad as dead (values > 1 are not

---

[14] A closely related measure, the disability-adjusted life year (DALY), was developed as an egalitarian indicator of the gap between current and "ideal" health to support assessment of the global burden of disease (Murray & Lopez, 1996). The DALY approach is currently undergoing revision (see: http://globalburden.org/).

[15] In cost-effectiveness analysis, valuation is implicit, because monetary thresholds are needed to compare with the cost-effectiveness ratio to determine whether the intervention is worth implementing.

possible but values < 0 are used for states that are judged to be worse than dead). QALYs can be summed across health states to determine the total QALYs associated with a particular condition. The results can then be added across the health conditions and individuals affected by a policy to determine the total QALYs potentially gained or lost as a result of its implementation. For example, QALY gains can be summed across averted cases of chronic bronchitis, heart disease, asthma and other conditions, as well as premature mortality, to determine the total impact of a policy that would reduce air pollution. These steps are illustrated in Figure 2.

Assume that, in the absence of the policy, the average individual is likely to survive for 10 more years with a health-related quality of life of 0.7. With the policy, assume that the average individual affected is likely to survive for 15 more years with a health status of 0.9.

Then the QALY gain attributable to the policy is the difference between 15 years with a health status of 0.9 (13.5 QALYs) and 10 years with a health status of 0.7 (7 QALYs), which equals 6.5 QALYs.

This gain can be described as having two components.

- The decreased morbidity during the 10-year survival period would lead to an average gain of 2.0 QALYs (moving from 0.7 to 0.9 provides an HRQL increase of 0.2, multiplied by 10 years).

- The individual would survive for an additional five years, with an HRQL of 0.9 over this time period. The increase in life expectancy leads to an additional gain of 4.5 QALYs (moving from 0 to 0.9 provides an HRQL increase of 0.9, multiplied by 5 years).

The total gain is thus 6.5 QALYs (2.0 QALYs plus 4.5 QALYs) for the average individual affected.

If the affected population includes 500 such individuals, then the total gain attributable to the regulation would be 3,250 QALYs (6.5 QALYs multiplied by 500 individuals).

Figure 2 Example of QALY calculations.

The process described above can be implemented using new primary research; e.g., by surveying the affected population to determine their current health status and their preferences for changes in this status. However, analysts usually follow other approaches that require less time and funding to implement. These include relying on estimates from previously completed studies, many of which can be found in the Tufts Cost Effectiveness Analysis (CEA) Registry (Thorat et al., 2012), using the benefits transfer process discussed earlier.[16]

---

16 The Tufts CEA Registry can be accessed at:https://research.tufts-nemc.org/cear4/.

A frequently used option is to apply one of several generic HRQL indices, examples of which include the EQ-5D, the Health Utilities Index (HUI), and the Quality of Well-Being (QWB) scale.[17] Each employs a classification system with several dimensions to describe health; e.g., in the case of the EQ-5D: mobility, self-care, usual activities, pain, and anxiety and depression. A particular health state is rated within each dimension; for example, as causing no, some, or extreme mobility problems. Each attribute of the health state (such as having "some" mobility problems) is then weighted based on a population survey developed especially for that index. These indices have the advantage of standardizing the approach for describing each health state and including pre-established preference weights for each attribute. The results will vary, however, depending on which index is applied, given differences in the attributes they include and in how the attributes are weighted.

In recent years, researchers have used these indices to develop catalogues of weights based on population-wide surveys. Some large national surveys (such as the US Medical Expenditure Panel Survey or MEPS) now include one or more of the generic indices. Relying on such surveys can be particularly useful for policy analysis, because they provide consistently-derived estimates across a wide range of outcomes and enable researchers to control statistically for the effects of other factors (such as age) on HRQL. For example, Sullivan and Ghushchyan (2006) estimate EQ-5D scores for a large number of health conditions based on MEPS.[18]

Developing approaches for measuring QALYs and testing their implementation is an active area of research. An expert panel report (Institute of Medicine, 2006) provides more detailed discussion of these measures and their application. Lipscomb, Drummond, Fryback, Gold, and Revicki (2009) discuss the status of these measures, and note that there continue to be diverse opinions on many technical issues such as the HRQL dimensions that should be considered, the types of survey questions that should be used to explore these dimensions, the elicitation of preferences, and the statistical analysis of the results. Thus the approaches described above continue to evolve, and new options are under development.

---

17 More information on these and other indices, including cross-index comparisons, is available on the National Health Measurement Study website: http://www.healthmeasurement.org/NHMS.html.

18 This article, and a calculator that allows users to retrieve EQ-5D scores by International Classification of Disease code and demographic characteristics, are available online: http://www.ohsu.edu/epc/mdm/webResources.cfm. The underlying scoring function used to calculate HRQL from the EQ-5D for the US population was developed by Shaw, Johnson, and Coons (2005).

4.3.2 Valuing QALYs

To apply these QALY estimates to value morbidity risks in benefit-cost analysis, another step is needed: they must be assigned a monetary value. Because reliance on WTP estimates is preferable when possible, the VSL should be used to value mortality risk reductions (e.g., the 4.5 QALY gain in the example in Figure 2). When WTP estimates are not available for morbidity (e.g., the 2.0 QALY gain in Figure 2), analysts at times simply to multiply the gain (discounted to reflect time preferences) by the VSLY (e.g., $279,000 if we follow the earlier calculations, starting with a $6.0 million VSL and using a three percent discount rate).

This approach has several limitations. First, QALYs are not entirely consistent with the framework for benefit-cost analysis. As discussed in Hammitt (2002), the construction of QALYs assumes that how individuals value health states is independent of the duration of the state, the age at which they are experienced, the individual's remaining life expectancy, and his or her wealth and income. Moreover, QALYs do not explicitly account for changes in wealth or income that result from an injury or illness. These assumptions conflict with economic theory and WTP research.

In addition, the VSLY estimates the average value of future life years without accounting for future health. Because health tends to deteriorate with age, future QALYs are less than future life years, and dividing VSL by future QALYs would yield an average value per QALY larger than VSLY. (Some studies have attempted to estimate the average WTP per QALY by dividing VSL by discounted future QALYs; e.g., Hirth, Chernew, Miller, Fendrick, and Weissert, 2000.) Parallel to the discussion of VSLY above, the assumption of a constant value per QALY implies that VSL is proportional to future QALYs, which is not consistent with empirical estimates of how VSL varies with age.

Research that explicitly considers WTP per QALY indicates that this value is not a constant for reasons other than the age of those affected. Haninger and Hammitt (2011) find that the value depends on the magnitude of the expected QALY gain and the duration of the health effect. More studies are needed to explore this relationship; this type of research eventually may be useful in creating a (non-constant) valuation function for QALYs that could be applied in benefit-cost analysis. Such an approach is illustrated by meta-analyses of the value of acute effects (Van Houtven et al., 2006, updating Johnson, Fries, and Banzhaf, 1997), chronic illnesses. More recent research provides an opportunity to enhance this approach and to address.

Thus while monetized QALYs are sometimes used as a proxy when WTP estimates are not available, issues related to both the construction of the QALY measure and to the monetary valuation of incremental QALYs introduce

significant uncertainty. It is unclear whether this approach is likely to over- or understate individual WTP for reducing health risks, and by how much. Continued work on developing a valuation function for QALYs is needed. Given these concerns, an expert panel recommended against assigning monetary values to QALYs in regulatory analysis (Institute of Medicine, 2006).

As in the case of COI estimates, when using monetized QALYs it is important to discuss these uncertainties and the implications for decision-making.[19] The effect of these uncertainties can also be tested quantitatively; for example, by determining the value per QALY at which benefits would equal costs, comparing this value to the range suggested by available research, and considering whether it appears reasonable. The application of monetized QALYs also raises the same aggregation issues as the application of COI estimates. To the extent that it is possible to separately estimate risk changes and expected QALY gains by subgroup, such an approach will lead to more accurate estimates.

5 Conceptual issues

The sections above focus on describing the approaches commonly used to value health risk reductions in benefit-cost analysis. In this section, we turn to some conceptual issues that represent areas where additional research may be useful. These include differences between WTP and WTA and between risks to oneself and to the community, as well as the effects of decision-making heuristics and biases.

As discussed in the introduction, both individual WTP and WTA are consistent with the benefit-cost analysis framework. Under conventional assumptions, the choice of measure would not matter much in many contexts. In particular, Willig (1976) demonstrates that for private goods, where the individual can choose the quantity consumed, WTP and WTA should be similar as long as income effects are negligible. Hanemann (1991) finds that for public goods, where the individual cannot choose the quantity, WTP and WTA can diverge when there are no private goods that are close substitutes.

However, several empirical studies have found substantial divergence. For example, Horowitz and McConnell (2002) examine 45 studies of a wide range of goods (ranging from pens to nuclear waste repositories) and find that, on average,

---

19 While some analysts add COI estimates to QALY estimates, arguing that the later incorporates the value of averting pain and suffering, such an approach can lead to double-counting. For example, a QALY estimate that includes effects on usual activities may overlap with estimates of productivity losses.

WTA is about seven times higher than WTP. This ratio varies with characteristics of the good and tends to be higher for items not ordinarily purchased in the marketplace. Several economists, beginning with Kahneman, Knetsch, and Thaler (1991) and including subsequent work by Knetsch (2010), attribute such differences to loss aversion and the endowment effect. Individuals often value changes more strongly if they are viewed as a loss rather than as a gain; the endowment effect determines whether they view their present state or their state after the change as the reference point. If the reference state is one's current status, WTP for a gain will be smaller than WTA for a loss of the same magnitude.

This divergence creates challenges for benefit-cost analysts (Robinson and Hammitt, 2011b). It can be difficult to determine the appropriate reference state for a particular analysis and the choice of reference may be easily manipulated by describing or framing the decision in alternative but logically equivalent ways (Tversky and Kahneman, 1974). While some (e.g., Freeman, 2003) argue that the reference state should be based on property rights, such rights are not clearly defined in many policy contexts and may not be consistent with how an individual views his or her own reference state. It seems reasonable, however, to assume that individuals will often view their current status as the reference state when considering health-improving policies. In this case, WTP for the improvement, rather than WTA as compensation to forgo the improvement, would be the appropriate measure. The appropriate measure may be ambiguous in some settings, leading to difficulties in determining whether benefits exceed costs as well as in estimating the extent to which a policy is cost-beneficial. Whereas analysts could test the sensitivity of their results to the choice of WTP or WTA estimates, in practice this may not be possible without new primary research. Most studies that address health risk reductions focus on estimating WTP, and WTA estimates are often not available for comparison.

Another conceptual challenge relates to differences in values for one's own risk reductions vs. values for risk reductions that accrue to the community of which one is a part. Benefit-cost analysis typically addresses policies that have broad impacts, while many valuation studies focus on individual WTP for a decrease in one's own risk consistent with the notion of consumer sovereignty. This distinction raises questions about the appropriate treatment of other-regarding preferences, particularly altruism, as well as differences in WTP for public vs. private programs.[20]

---

20 This issue is distinct from the question of whether the outcome is a public or private good. In economics, public goods are nonexclusive (individuals cannot be prevented from consuming them) and nonrival (consumption by one does not reduce the amount available for others). However, public programs can provide private goods (e.g., health care consumed by the individual) as well as public goods (e.g., air pollution abatement that affects the community as a whole).

The treatment of altruism in benefit-cost analysis has received some atten-
tion, although few empirical studies explore its impact on the value of health
risk reductions. Economic theory suggests that the implications of altruism for
benefit-cost analysis depend on whether it is pure (non-paternalistic) or pater-
nalistic (Jones-Lee, 1991). A pure altruist respects the preferences of others: she
values both the benefits they receive and the costs they incur with the same rela-
tive weight as the affected individuals. In a society of pure altruists, a project to
provide a public good that satisfies a benefit-cost test using private values will
also satisfy a benefit-cost test using altruistic values, although the converse is
not necessarily true (Bergstrom, 2006). With pure altruism, it is not appropriate
to include altruistic WTP to reduce risk to others unless one also counts altruis-
tic losses associated with the costs that others bear (Bergstrom, 1982; Jones-Lee,
1991). In contrast, a paternalistic altruist cares only about some aspects of others'
well-being. For example, he may care about others' mortality risk but not about
the costs they must pay for risk reduction. Incorporating paternalistic altruism in
benefit-cost analysis can affect the sign of net benefits. However, altruism is typi-
cally not included when valuing health risk reductions in benefit-cost analysis, in
part because of theoretical complexity and in part because of the lack of empiri-
cal estimates.

Behavioral economists are exploring additional ways in which preferences
may be "other-regarding" ("social" preferences) rather than "self-regarding"
("private" preferences). These include concerns about social welfare (increasing
total welfare by helping others, particularly those less-well off); difference aver-
sion (reducing differences between oneself and others); reciprocity (rewarding
or penalizing others depending on the perceived fairness of their actions); and
relative rather than absolute position ("keeping up with the Joneses"). To date,
these concerns have been investigated primarily in laboratory experiments; their
combined effect on individual WTP for health risk reductions has not been com-
prehensively assessed (Robinson and Hammitt, 2011b).

Some valuation research explores preferences for public programs that
benefit the community of which the individual is a part, rather than solely the
individual.[21] According to Svensson and Johansson (2010), research on road
safety suggests that WTP for public risk reductions is noticeably smaller than for
private reductions. A meta-analysis of mortality risk reduction studies (Lindhjem
et al., 2011) also finds that risk changes are valued less if they affect individu-

---

[21] This issue is not the same as considering internalities and externalities. An externality is an
outcome not reflected in market interactions, and is the traditional rationale for government
intervention. However, government may intervene for other reasons; for example, to address the
need for accurate information or to promote equity.

als in addition to the respondent's own household. The finding that WTP for a program to reduce risk to others as well as to oneself suggests a negative WTP to reduce risk to others, but it may be explained if survey respondents believe the public program is less likely to deliver the stated risk reduction than the private program. Bosworth, Cameron, and DeShazo's (2010) study of treatment vs. prevention programs finds that preferences depend on both the nature of the policy and the characteristics of those affected. As for altruism and other social preferences, more work is needed to determine what factors influence these differences and to develop appropriate adjustments for application in benefit-cost analysis. At the present time, analysts may wish to recognize these issues in discussing their results, but may not be able to quantify their impact.

A final issue is the rapidly growing research in behavioral economics that identifies other ways in which individuals' preferences diverge from the predictions of standard economic models (Robinson and Hammitt, 2011b). This research encompasses many decision-making biases and heuristics in addition to the loss aversion and reference dependence discussed earlier. For example, it suggests that preferences may depend on timing: individuals often act in ways that suggest they have higher discount rates in the near-term than in the long-term. Such hyperbolic discounting may at least in part reflect self-control problems; e.g., eating dessert now while wanting to ultimately lose weight. While it is tempting to label these anomalies as "irrational," in some cases they may reflect inadequate information or valued attributes that are not obvious to the researcher. Whether and how to incorporate these preferences into a benefit-cost analysis poses significant normative challenges (Hammitt, 2009). If our goal is simply to describe the preferences of those affected by the policy, then we should presumably use reported values as long as the underlying studies are well-conducted. If our goal is more normative; e.g., to promote those policies that save the most money over the long run or achieve the largest health risk reductions regardless of individual preferences, then we may want to act more paternalistically and be selective about what values we include.[22] Regardless, it is important to be clear about the extent to which we are taking the results of studies as given, or are selecting the results to apply in a particular analysis based on normative grounds.

6 Summary and conclusions

In the context of benefit-cost analysis, the value of health risk reductions is most appropriately measured by individual WTP. For mortality risks, WTP estimates

---

22 Adler and Posner (2006) refer to this as "laundering" preferences.

are plentiful; the main challenge is to identify the estimates most suitable for a particular context, given both the quality of the available research and its applicability to the policy scenario. For nonfatal risks, fewer WTP studies are available. If the analyst cannot locate studies of reasonable quality for the effects of concern, he or she may consider the use of COI estimates or monetized QALYs as rough proxies. These approaches are widely used in policy analysis, but both involve important limitations that should be discussed when presenting the analytic results.

In addition to the need for more WTP studies, there are a number of conceptual issues worthy of further attention. In particular, better understanding the divergence between estimates of WTP and WTA and between values for one's own risk reductions vs. risks to the community would be useful, as would more exploration of the role of decision-making heuristics and biases.

References

Adler, M. A., & Posner, E. A. (2006). *New foundations for benefit-cost analysis.* Cambridge, MA: Harvard University Press.

Alberini, A., & Kahn, J. R. (Eds.) (2006). *Handbook on contingent valuation*. Cheltenham UK and Northampton MA: Edward Elgar.

Aldy, J. E., & Viscusi, W. K. (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.

Bateman, I. J., Carson, R.T., Day, B., Hanemann, M., Hanley, N., Hett, T., et al. (2002). *Economic valuation with stated preference techniques.* Cheltenham UK and Northampton MA: Edward Elgar.

Bergstrom, T. C. (1982). When is a man's life worth more than his human capital? In: M. Jones-Lee (Ed.), *The value of life and safety*. Amsterdam: North Holland.

Bergstrom, T. C. (2006). Benefit-cost in a benevolent society. *American Economic Review 96*(1), 339–351.

Blomquist, G. (2004). Self-protection and averting behavior, values of statistical lives, and benefit cost analysis of environmental policy. *Review of the Economics of the Household 2*, 89–110.

Boardman, A. E., Greenberg, D. H., Vining, A. R., & Weimer, D. L.. (2011). *Cost-benefit analysis: concepts and practices*. Upper Saddle River, NJ: Pearson.

Bosworth, R., Cameron, T. A., & DeShazo, J. R. (2010). Is an ounce of prevention worth a pound of cure? comparing demand for public prevention and treatment policies. *Medical Decision Making 30*(4), E40–E56.

Cameron, T. A. (2010). Euthanizing the value of a statistical life. *Review of Environmental Economics and Policy 4*(2), 161–178.

Carson, R. T. (2012). *Contingent valuation: a comprehensive bibliography and history.* Cheltenham UK and Northampton MA: Edward Elgar.

Champ, P. A., Boyle, K. J., & Brown, T. C. (Eds.) (2003). *A primer on non-market valuation.* Boston MA: Kluwer Academic Press.

Cropper, M., Hammitt, J. K., & Robinson, L. A. (2011). Valuing mortality risk reductions: progress and challenges. *Annual Review of Resource Economics 3*, 313–336.

Cropper, M., & Morgan, M. G. (2007). SAB Advisory on EPA's issues in valuing mortality risk reduction. Memorandum to EPA Administrator Stephen L. Johnson from the EPA Science Advisory Board and Environmental Economics Advisory Committee. (EPA-SAB-08-001.)

European Commission. (2009). *Impact Assessment Guidelines.* SEC(2009) 92.

Finkelstein, E. A., Corso, P. S., & Miller, T. R. (2006). *The incidence and economic burden of injuries in the United States.* Oxford and New York: Oxford University Press.

Freeman, A. M. III. (2003). *The measurement of environmental and resource values: theory and methods (2nd ed.),* Washington D. C.: Resources for the Future.

Grosse, S. D., Krueger, K. V., & Mvundura, M. (2009). Economic productivity by age and sex: 2007 estimates for the United States. *Medical Care 47*(7), S94–S103.

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., & Robinson, L. A. (2011). The income elasticity of the value per statistical life: transferring estimates between high and low income populations. *Journal of Benefit-Cost Analysis 2*, Art. 1.

Hanemann, W. M. (1991). Willingness to pay and willingness to accept: how much can they differ? *American Economic Review 81*, 635–647.

Haninger, K., & Hammitt, J. K. (2011). Diminishing willingness to pay per quality-adjusted life year: valuing acute foodborne illness. *Risk Analysis 31*(9), 1363–1380.

Hirth, R. A., Chernew, M. E., Miller, E., Fendrick, A. M., & Weissert, W. G. (2000). Willingness to pay for a quality-adjusted life year: in search of a standard. *Medical Decision Making 20*, 332–342.

HM Treasury. (2003). *The green book: appraisal and evaluation in central government.* London: TSO. (Includes 2011 updates.)

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

Hunt, A., & Ferguson, J. (2010). *A review of recent policy-relevant findings from the environmental health literature.* Prepared for the Organisation for Economic Co-Operation and Development. ENV/EPOC/WPNEP(2009)9/FINAL.

Institute of Medicine. (2006). *Valuing health for regulatory cost-effectiveness analysis.* In: W. Miller, L. A. Robinson & R. S. Lawrence. (Eds.), Washington, D.C.: The National Academies Press.

Johnson, F. R., Fries, E. E., & Banzhaf, H. S. (1997). Valuing morbidity: an integration of the willingness-to-pay and health-status index literatures. *Journal of Health Economics 16*, 641–665.

Jones-Lee, M. W. (1991). Altruism and the value of other people's safety. *Journal of Risk and Uncertainty 4*(2), 213–219.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: the endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives 5*(1), 192–206.

Kling, C. L., & Swackhamer, D. L. (2011). *Review of 'valuing mortality risk reductions for environmental policy: a white paper'* (December 10, 2010). Memorandum to Lisa P. Jackson, EPA Administrator from the EPA Science Advisory Board and Environmental Economics Advisory Committee. EPA-SAB-11-011.

Knetsch, J. L. (2010). Values of gains and losses: reference states and choices of measure. *Environment and Resource Economics 46*(2), 179–188.

Krupnick, A. (2007). Mortality-risk valuation and age: stated preference evidence. *Review of Environmental Economics and Policy 1*(2), 261–282.

Lipscomb, J., Drummond, M., Fryback, D., Gold, M., & Revicki, D. (2009). Retaining, and enhancing, the QALY. *Value in Health 12*(Supp. 1), S18–S26.

Lindhjem, H., Navrud, S., Braathen, N. A., & Biausque, V. (2011). Valuing mortality risk reductions from environmental, transport, and health policies: a global meta-analysis of stated preference studies. *Risk Analysis 31*(9), 1381–1407.

Lund, J. L., Yabroff, K. R., Ibuka, Y., Russell, L. B., Barnett, P. G., Lipscomb, J., ... Brown, M. L. (2009). Inventory of data sources for estimating health care costs in the United States. *Medical Care 47*(7), S120–S126.

Murray, C. J. L., & Lopez, A. D. (Eds.) (1996). *The global burden of disease: a comprehensive assessment of mortality and disability from diseases, injuries, and risk factors in 1990 and projected to 2020*. Cambridge: Harvard University Press.

National Academies. (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.

NERA Economic Consulting. (2011). *Updating the VPF and VPIs: phase 1: final report*. Prepared for the United Kingdom Department for Transport.

Organisation for Economic Co-operation and Development. (2012). *Mortality risk valuation in environment, health and transport policies*. Paris: OECD Publishing.

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. (2009). Valuing lives, valuing risks, and respecting preferences in regulatory analysis. *Regulation and Governance 3*(3), 298–305.

Robinson, L. A., & Hammitt, J. K. (2011a). *Valuing health and longevity in regulatory analysis: current issues and challenges*. In: D. Levi-Faur (Ed.), *The handbook of the politics of regulation*. Cheltenham and Northampton: Edward Elgar.

Robinson, L. A., & Hammitt, J. K. (2011b). Behavioral economics and the conduct of benefit-cost analysis: towards principles and standards. *Journal of Benefit-Cost Analysis 2*(2), Art. 5.

Shaw, J. W., Johnson, J. A., & Coons, S. J. (2005). US valuation of the EQ-5D health states: development and testing of the D1 valuation model. *Medical Care 43*, 203–220.

Sullivan, P. W., & Ghushchyan, V. (2006). Preference-based EQ-5D index scores for chronic conditions in the United States. *Medical Decision Making 26*(4), 410–420.

Svensson, M., & Johansson, M. V. (2010). Willingness to pay for private and public road safety in stated preference studies: why the difference? *Accident Analysis and Prevention 42*, 1205–1212.

Thorat, T., Cangelosi, M., & Neumann, P. J. (2012). Skills of the trade: the Tufts cost-effectiveness analysis registry. *Journal of Benefit-Cost Analysis 3*(1), Art. 1.

Treasury Board of Canada, Secretariat. (2007). *Canadian cost-benefit analysis guide: regulatory proposals (Interim).*

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science 185*(4157), 1124–1131.

US Department of Transportation. (2011). *Treatment of the economic value of a statistical life in departmental analyses – 2011 interim adjustment.* Memorandum to Secretarial Officers and Modal Administrators from P. Trottenberg, Assistant Secretary For Transportation Policy, and R. Rivkin, General Counsel.

US Environmental Protection Agency. (2010a). *Valuing mortality risk reductions for environmental policy: a white paper (Review Draft).* Prepared by the National Center for Environmental Economics for consultation with the Science Advisory Board – Environmental Economics Advisory Committee.

US Environmental Protection Agency. (2010b). *Guidelines for preparing economic analysis.* EPA 240-R-10-001.

US Office of Management and Budget. (2003). *Circular A-4: regulatory analysis.*

Van Houtven, G., Powers, J., Jessup, A., & Yang, J.-C. (2006). Valuing avoided morbidity using meta-regression analysis: what can health status measures and QALYs tell us about WTP? *Health Economics 15*, 775–795.

Viscusi, W. K., & Aldy, J. E. (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of Risk and Uncertainty 27*(1), 5–76.

Viscusi, W. K. (2009). The devaluation of life. *Regulation and governance 3*, 103–127.

Willig, R. D. (1976). Consumer's surplus without apology. *The American Economic Review 66*(4), 589–597.

Yabroff, K. R., Brown, M. L., Lawrence, W. F., Barnett, P. G., & Lipscomb, J. (Eds.). (2009). Health care costing: data, methods, future directions. *Medical Care 47*(7), Supplement 1.

Zeckhauser, R., & Shepard, D. (1976). Where now for saving lives? *Law and Contemporary Problems 40*(4), 5–45.