Symposium:
Mortality-risk Valuation and Age

Valuing Changes in Mortality Risk: Lives Saved Versus Life Years Saved

James K. Hammitt*

Introduction

Valuing changes in mortality risk is critical to economic evaluation of many environmental regulations. In recent decades, economists have made much progress in understanding the conceptual basis for valuing mortality risk and in developing and applying methods for empirical estimation. Nevertheless, the appropriate valuation remains controversial. Two approaches are commonly used, focusing on either the number of lives saved (known as the value per statistical life [VSL]), or the number of life-years saved (known as the value per statistical life-year [VSLY]). As conventionally applied, the VSL approach counts each life saved by a regulation as equally valuable, and the VSLY approach values each life saved in proportion to how long it is extended. The difference between the approaches can affect the apparent merits of regulatory programs that disproportionately affect people with differing life expectancies. For example, the major air-pollution regulations that dominate the regulatory agenda in the United States may primarily benefit people who have shorter-than-average life expectancies because they are older or in poor health. The calculated benefits of these rules may be much larger using a VSL than a VSLY approach.

The purpose of this symposium is to evaluate the VSL and VSLY approaches to valuing changes in mortality risk. There are three articles in the symposium. This article provides an overview of the theory behind the approaches and their applicability to estimating the benefits of environmental regulation. The articles by Aldy and Viscusi (2007) and Krupnick (2007) review the empirical evidence concerning valuation of mortality risk and its consistency with the two approaches. Aldy and Viscusi (2007) describe the evidence from "revealed-preference" studies of workers' tradeoffs between higher wages and safer working conditions. Revealed-preference studies are based on the assumption that individuals' real-world choices between alternatives that differ in mortality risk and monetary consequences reflect their actual preferences. Krupnick (2007) describes the evidence from "stated-preference" studies of fatal risks such as motor vehicle crashes and disease in which survey

*Harvard University, Cambridge, MA, USA.

Review of Environmental Economics and Policy, volume 1, issue 2, summer 2007, pp. 228–240
doi: 10.1093/reep/rem015

respondents are asked how they would choose between alternatives that differ in mortality risk and monetary consequences.

After providing background on US and other government approaches to the valuation of mortality risk for environmental regulation, the following section introduces and defines the relationships among VSL, VSLY, and survival curves, which describe the level of mortality risk and how it changes over time. The next section describes the standard single- and multiperiod theoretical models of the value of mortality risk. The following section considers some issues related to public-policy applications, including difficulties in empirically estimating how the value of mortality risk varies with life expectancy and alternative frameworks for evaluating changes in mortality risk within society. The final section is the conclusion.

Background and Key Concepts

Historically, the lives-saved and life-years-saved approaches have been applied in different contexts. The lives-saved or VSL approach has been the conventional method for evaluating changes in mortality risk due to environmental and transportation policy. In recent years, the US Environmental Protection Agency (EPA) has sometimes supplemented its VSL-based estimates of regulatory benefits with alternative analyses using the life-years-saved or VSLY approach. The VSLY approach, which focuses on the extent to which life is prolonged, is similar to the life-year-based methods that have been used to evaluate changes in mortality risk in other areas of public health, including medical treatment, disease screening, vaccination, and characterizing the burden of disease in a society (Gold et al. 1996; Murray and Lopez 1996). Under the current US administration, the Office of Management and Budget (OMB) has encouraged EPA and other agencies to analyze regulations using both the lives-saved and life-years-saved approaches (OMB 2003).

Although the VSL or VSLY is usually taken as constant over the population within an analysis, in some cases different values have been applied or recommended depending on the age of the affected population or the type of risk. For example, in sensitivity analyses, EPA has used a smaller VSL for evaluating risk to older adults from some air pollution regulations (Graham 2003); the European Commission has recommended using 30 percent smaller values for older adults and 50 percent larger values for risks of cancer than other causes of fatality (European Commission 2005); and OMB has recommended using a larger VSLY for older adults (Graham 2003).

The characterization of environmental regulations as saving lives (or saving life-years) is commonly used but can be misleading. Typically, a reduction in exposure to environmental pollution reduces some people's chances of dying from a specific cause in a defined time period by a small amount. When a large population is affected, it is virtually certain that fewer deaths will occur (from the specific cause in the defined time period). However, it is impossible to know either beforehand or afterward whose death will be or was averted. The beneficiaries of the risk reduction are unidentifiable.

Value per Statistical Life

The VSL is defined as the marginal rate of substitution between wealth and mortality risk in a defined time period. For example, if a typical individual would be willing to pay six

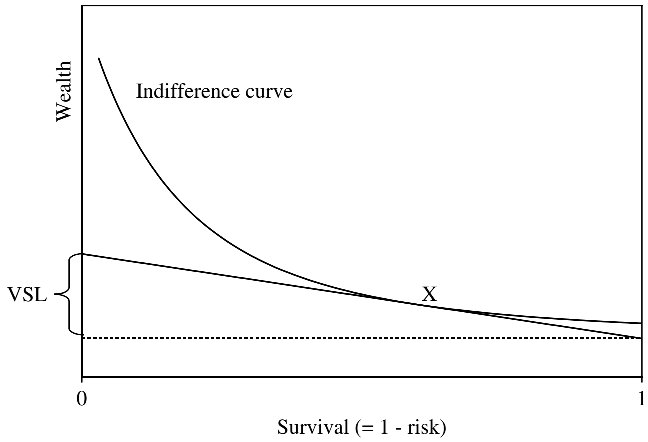

Figure 1. Value per statistical life is equal to the slope of the individual's indifference curve at the point (X) determined by her wealth and survival probability.

dollars to reduce her chance of dying this year by one in a million, her VSL is approximately six dollars/(one in a million), or six million dollars. The term *VSL* can be understood by recognizing that if each person in a population of one million paid six dollars to reduce his or her chance of dying this year by one in a million, a total of six million dollars would be paid, and one fewer death would be expected to occur this year.

An individual's VSL is her rate of substitution between wealth and survival probability (i.e., one—mortality risk) given her current risk and wealth (it may also depend on other factors such as anticipated future health and income). Just as speed can be measured in different units (e.g., meters per second and kilometers per hour), so too can the rate of substitution between risk and wealth. Howard (1984) proposed describing the rate of substitution between wealth and mortality risk in terms of a less misleading concept—a "micromort"—which is a risk change of one in a million. Using this terminology, our typical individual values a micromort at six dollars. Other authors use terms such as "value of a prevented fatality" (e.g., Jones-Lee 2004) to emphasize that the concept measures the valuation of risk reduction rather than the moral worth of an individual.

The conventional description of this rate of substitution as a VSL can be easily misinterpreted to suggest that the value of a person's life is equal to her VSL. As illustrated in figure 1, an individual's indifference curve for wealth and the probability of surviving a specified time period is convex to the origin. Her VSL is the *slope* of the indifference curve at her current wealth and survival probability. At the point marked X, for example, her VSL is equal to the vertical distance between the two ends of the line that is drawn tangent to her indifference curve at X. In contrast, the amount of money this individual would require to compensate her for a large reduction in survival probability is much larger than her VSL, and might even be infinite.

Value per Statistical Life-year

A human lifetime is finite. Death can be postponed but not prevented. Indeed, reducing the risk of dying at one time necessarily increases the risk of dying at some later time. (Death from a specific cause can be prevented, at least in principle, but doing so necessarily increases the chance of dying from some other cause.) It seems intuitive that the value to an individual of delaying her death depends on the duration of the delay, and that postponing death by several years would usually be preferred to postponing it by only days (even if those few days may be critically important). The VSLY approach values a reduction in mortality risk in proportion to the gain in life expectancy. Under this approach, reducing an individual's risk of dying in the current year produces a gain equal to the risk reduction multiplied by her life expectancy conditional on surviving the current year. In practice, future life-years are usually discounted. This means that the value of reducing the risk of dying this year is proportional to the expected present value of future life-years, which is often approximated by the present value of a series of years with duration equal to the individual's life expectancy.

Life expectancy varies with age, health, sex, and socioeconomic status. Most research on the effect of life expectancy on preferences for life saving focuses on the effects of age rather than the effects of life expectancy holding age constant.

There is a substantial literature that uses surveys to elicit preferences for public programs to reduce mortality risk to different populations. This question is distinct from asking about an individual's willingness to pay (WTP) to reduce his or her own mortality risk. These studies generally find a preference for reducing risk to younger populations with greater life expectancy. For example, Johannesson and Johansson (1997) estimate that their median respondent is indifferent between saving the life of one thirty-year-old, five fifty-year-olds, and thirty-five seventy-year-olds. They calculate that these choices imply that the value of an average future life-year declines with age, with the average future life-year of a thirty-year-old valued as equivalent to three future life-years of a fifty-year-old and ten future life-years of a seventy-year-old. Cropper et al. (1994) find that people prefer to save young adults over both older and younger adults. The mean respondent in their survey is indifferent between saving the life of eight sixty-year-olds and one twenty-year-old, eleven sixty-year-olds and one thirty-year-old, and seven sixty-year-olds and one forty-year-old. Similarly, in their estimates of the social burden of disease, Murray and Lopez (1996) assume that the average value of a life-year first rises then falls as children grow into adults and seniors.

In a comprehensive review of empirical studies of public preferences to reduce mortality risk or improve health, Dolan et al. (2005) also conclude that there is a social preference for aiding younger rather than older people and that life-years lived by younger individuals are counted as more valuable than life-years lived by their elders. However, Dolan et al. (2005) find that there is no social preference for aiding men over women (despite men's shorter life expectancies) and that there is evidence of a preference for aiding members of disadvantaged social groups and those of lower socioeconomic status (who have shorter life expectancies).

Survival Curves

When evaluating policies that affect mortality risk, it is useful to think in terms of survival curves. A survival curve plots the probability that an individual is still alive as a function of her age (or calendar date). A survival curve can be constructed beginning at any age or date. The average survival curve for the US population beginning at age sixty is illustrated by the solid line in figure 2 (Arias 2002). The height of the curve at the initial age or date

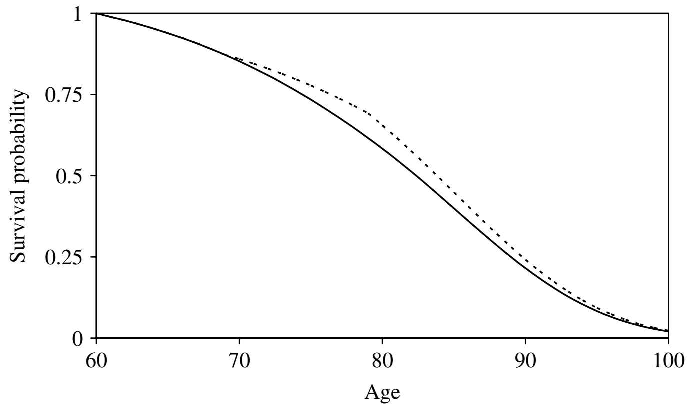

Figure 2. Survival curves from age sixty. Dashed line corresponds to decreasing mortality risk at each age from seventy-one through eighty by one-third.

is one and the curve falls as age or time increases. The slope of the curve depends on the mortality risk, with steeper decreases corresponding to periods of higher mortality risk. Life expectancy at any age is the area under the survival curve that begins at that age. For the solid curve illustrated in figure 2, life expectancy at age sixty is twenty-two years.

Any pattern of change in mortality risk over time can be characterized as a shift in the survival curve. For example, a one-year reduction in mortality risk (e.g., from reducing the motor vehicle crash rate) flattens the survival curve for that year and hence increases its height for later time periods. A risk reduction having only delayed effects (e.g., reducing exposure to a pollutant that causes cancer to develop after a latency period) has no effect on the curve for the time between the change in exposure and the end of the latency period, but flattens the curve and increases its height for subsequent periods. To illustrate, consider an intervention that decreases annual mortality probability by one-third, persists for ten years, and begins after a ten-year lag. The dashed line in figure 2 illustrates the effect of this intervention for sixty-year-olds. There is no change in survival probability during the lag period (ages sixty through sixty-nine). The survival curve is flatter over the period during which risk is reduced (ages seventy through seventy-nine) and remains higher than the baseline survival curve in later years (ages eighty and above). For every sixty-year old affected, this intervention saves one life-year and 0.08 lives (between ages seventy and seventy-nine).

Survival curves can be constructed for an individual or a population. Any change in the survival curve produces a unique expected number of life-years saved or lost (the change in the area under the curve) and a unique expected number of lives saved or lost at each point in time (the vertical shift in the curve as a function of time). The total number of lives saved during a time period (which may include saving the same life multiple times) depends on the period examined; for periods longer than a maximal human lifespan (e.g., 150 years), the number of lives saved (among a cohort) must be zero. Note that there is no unique change in the survival curve corresponding to a specified number of life-years saved or to a specified number of lives saved in a finite time period. The survival curve and how it shifts are the fundamental concepts; the numbers of life-years saved and lives saved in a specified time period are alternative and partial summary measures of the shift.

One can estimate an individual's WTP for a specified shift in her survival curve by dividing the shift into a series of instantaneous changes in risk and summing her WTP for each of these "blips" (Rosen 1988; Cropper and Sussman 1990; Johannesson et al. 1997). WTP depends on the time of payment and on the conditions under which the individual can save and borrow against future income. It may also depend on when (if ever) the individual learns of the change in her survival curve. WTP will typically be larger the earlier she learns of the shift, and can adjust by reallocating her planned future consumption and risk-reducing expenditures (Cave 1988).

The value of a specified change in a survival curve can be described alternatively as the average value per life-year saved or the average value per life saved, by dividing WTP for the change in the survival curve by the corresponding expected number of life-years saved or the expected number of lives saved in some specified time period, respectively. The choice between these summary descriptions is arbitrary and largely a matter of convenience. The average VSL and average VSLY may depend on the initial survival curve, the shift in the curve, and other factors. Economic theory does not predict that either VSL or VSLY will be constant across interventions, individuals, or time. For policy evaluation, if either VSL or VSLY were reasonably stable across the interventions to be evaluated, that measure would seem to be the more convenient summary.

In practice, it may be difficult to estimate the effect of an intervention on a survival curve, and one may or may not be able to estimate either the numbers of lives saved or life-years saved. For mortality risks where the victims are identifiable *ex post*, such as motor vehicle crashes and deaths from a signature disease (e.g., mesothelioma from asbestos exposure), the number of lives saved as a function of time can be estimated, but the number of life-years saved cannot without additional information on the life expectancy of the affected population. Similarly, time-series studies of deaths from fine particulate matter or other air pollutants yield estimates of how the number of deaths in a short time period (e.g., a week) depends on ambient pollution and thus how many lives may be saved by reducing that pollution. However, these studies provide little information on the life expectancy of the affected population and hence, on the number of life-years saved. One cannot distinguish whether the additional deaths that occur in periods of higher pollution are suffered by people who would have otherwise died within a few months or survived many years. Time-series studies provide a lower bound on lives saved if some of the people who die from air pollution would otherwise have survived only a few more days (Rabl 2006).

In contrast, cohort studies that follow over time populations that are exposed to different levels of air pollution (e.g., those living in different cities) provide estimates of the survival curve and how it depends on air pollution. These studies can be used to estimate the number of life-years saved but not the number of lives saved. The identical shift in the population survival curve can be the result of extending either the lives of many people for a short time or the lives of fewer people for a longer period (Rabl 2003). To clarify, consider a stylized example in which survival curves are estimated for two cities that differ in air pollution. In the "polluted" city, half the population dies at age sixty and the other half dies at age seventy. In the "clean" city, half the population dies at age seventy and the other half at age eighty. Clearly, the lower pollution in the clean city is associated with an increase of ten statistical life-years per person. However, it is impossible to determine from the survival curves alone whether the difference arises because everyone lives ten years longer in the clean city than they would in the dirty city, or because the people who live to eighty in the clean city would have died at sixty in the dirty city. In the first case, the number of lives saved (at ages sixty and seventy) is equal to the population; in the second, it is equal to half the population.

For risks associated with exposure to chemicals that may cause cancer or other disease, epidemiological data are often limited or nonexistent and risk estimates are based on studies of laboratory animals. In these cases, estimation of either lives saved or life-years saved requires strong assumptions about how to extrapolate from effects observed in highly exposed laboratory animals to effects in less highly exposed humans, including the type of cancer or other disease, the probability of lethality, the latency of the disease, and the life expectancy of the affected population.

Economic Models of the Value of Changes in Mortality Risk

Economic models of the value of changes in mortality risk are based on the standard economic concept of individual sovereignty, i.e., the notion that the individual is usually the best judge of her own interests. Hence, the monetary value of a change in risk depends on the individual's preferences for risk reduction and other opportunities for spending. The standard model was developed by Drèze (1962), Schelling (1968), Mishan (1971), Jones-Lee (1974), and Weinstein et al. (1980). Both one-period and multiperiod models have been studied.

The One-period Model

In the simple one-period model, the individual faces a tradeoff between spending money to reduce her chance of dying during the period and spending the money on consumption in the current or future periods. If she survives the current period, she faces an uncertain future longevity that is determined by her future mortality risks. If she dies in the current period, she gains utility from living for the part of the period before she dies and from leaving part of her wealth as a bequest to her heirs or others. Her utilities of surviving and dying during the current period depend on her wealth. Both her total and marginal utilities of wealth are assumed to be higher if she survives the period (and can allocate her wealth toward future consumption) than if she dies and leaves her wealth as a bequest.

In the one-period model, VSL depends on the length of the period. An increase in survival probability of fixed magnitude (e.g., one in a million) is more valuable if it increases her probability of surviving the next decade rather than surviving only the next day. Johansson (2002) shows that the magnitude of VSL may be critically dependent on the duration of the period to which the risk reduction applies and that the only consistent definition of VSL may be as a rate of substitution between wealth and mortality risk at a defined instant (e.g., current risk).

Spending money to reduce current mortality risk has two effects on the individual's well-being. First, it increases her chance of surviving the current period and hence gaining utility from living in future periods. Second, it reduces the wealth she has available for spending on current and future consumption. To evaluate the effects of life expectancy and other factors on VSL one must consider both effects.

An increase in life expectancy (conditional on surviving the current period) typically increases the utility gained by surviving the current period. However, an increase in life expectancy may also increase the opportunity cost of spending on risk reduction. If, for example, the individual must support herself from her accumulated savings and has no prospect of future income, she must husband her resources much more carefully if she faces a long life expectancy rather than a short one. In this case, the marginal utility forgone by spending on risk reduction increases with life expectancy and the net effect of an increase in life expectancy—increasing both the benefit from surviving the current period and the opportunity cost of spending to increase the chance of survival—is ambiguous; an increase in life expectancy conditional on surviving the current period may increase, decrease, or have no effect on VSL (Hammitt 2002). In contrast, if actuarially fair annuities are available, the individual can reduce her potential bequest in the event of death and increase consumption in the event of survival, which means that the marginal utility of spending is independent of whether she survives or dies. Similarly, if she anticipates future labor, pension, or other income conditional on survival, then the opportunity cost of spending on risk reduction may not increase with life expectancy. In these cases, VSL will increase with life expectancy.

Note that the effect of anticipated future health on VSL is also ambiguous. Better future health increases the utility from surviving the current period (survival in good health is preferred to survival in poor health), but it may also increase the marginal utility of wealth. Poor health may restrict one's options for increasing utility through spending (e.g., a bed-ridden individual may have fewer opportunities to improve her utility by spending money than she would have if she were healthier). In this case, worse anticipated health will decrease the opportunity cost of spending to reduce current mortality risk, which may increase VSL (Hammitt 2000, 2002).

Multiperiod Models

In more complicated models that include multiple periods, the effect of greater life expectancy on VSL is also ambiguous; VSL can remain constant, increase, or decrease with age (Rosen 1988; Johansson 2002). Multiperiod models have generally been used to analyze how VSL varies with age (and thus indirectly with life expectancy) rather than to evaluate how VSL changes with life expectancy holding age constant.

The results of multiperiod models are sensitive to assumptions about the relationship between the interest rate, the rate at which future utility is discounted, and the survival curve; the terms under which the individual can borrow against future income; and the dependence of the marginal utility of consumption on age. If the utility of consumption does not vary with age and the individual can borrow against future income and discounts future utility (including the discount for the probability of dying) at the market interest rate, then it is optimal to allocate consumption uniformly across her lifespan. In this case, VSL is proportional to the expected present value of future life years (Rosen 1988). If she cannot borrow against future income, VSL follows an inverted U shape, first rising (as earnings and savings rise) then falling with age. The age at which VSL peaks and the extent to which it shifts with age are sensitive to her discount rate. If she discounts future utility at the market interest rate, VSL may peak around age forty (Shepard and Zeckhauser 1984); if she discounts utility at a lower rate, VSL may peak at older ages because it is optimal to invest more when young and consume more when old (Ng 1992).

If, as these models suggest, an individual's VSL first rises then falls with age, then her VSLY cannot be constant over her lifespan. Life expectancy typically decreases with age. During the period in which VSL increases with age, the average VSLY (i.e., VSL divided by life expectancy) must increase with age. In contrast, during periods when VSL falls with age, VSLY may remain constant if VSL decreases in proportion to life expectancy (or to the expected present value of future life-years, if life-years are discounted).

Applications to Public Policy

Benefit-cost analysis is conventionally understood as a means of determining whether a policy change or intervention yields a potential Pareto improvement, i.e., whether the monetary value of the benefits to those who gain exceeds the monetary value of the costs to those who are harmed. If this net-benefit criterion is satisfied, then the winners could, in principle, compensate the losers, and the policy change would yield a Pareto improvement (Stokey and Zeckhauser 1978). Following this approach, policies that affect mortality risk should be evaluated using the affected individuals' own valuations of the changes they face, and empirical research on how valuation of mortality risk varies with life expectancy, age, health, income, and other factors is critical for improved policy analysis.

There is a substantial empirical literature on how the value of mortality risk depends on age, sex, income, and other factors correlated with life expectancy (the effects of age are reviewed in the accompanying papers by Aldy and Viscusi (2007) and Krupnick (2007)). In contrast, there is little information on how it varies with life expectancy holding these factors constant (Evans and Smith 2006). Moreover, empirical estimation of how the value of mortality risk depends on life expectancy or correlated factors such as age is complicated by several factors. First, information on perceived life expectancy is usually not available and, when it is, may not be reliable. Some large population surveys have asked respondents about their perceived probability of surviving to certain ages (e.g., seventy-five and eighty-five in the Health and Retirement Study (http://hrsonline.isr.umich.edu)). These responses provide useful information about perceived longevity, but whether they can be interpreted as measures of subjective probability is debatable (e.g., Hurd and McGarry 1995, 2002; Smith et al. 2001; Viscusi and Hakes 2003). Stated-preference surveys could ask respondents about life expectancy and subjective probabilities of surviving to older ages, but usually have not.

Second, revealed-preference studies that estimate the compensation workers demand for occupational hazards provide little information about the groups that may be most susceptible to air pollution and other environmental health risks, since the elderly and individuals with poor health are less likely to be employed or to work in jobs with significant safety risks. Within these groups, those who are employed may not be representative (e.g., employed elderly workers may be healthier than average for their age).

Third, it may be difficult to accurately separate effects of age from effects of financial resources and obligations when estimating VSL. Data on personal or household wealth are usually not included in risk-valuation studies, and controlling for income may not adequately capture differences in financial resources available to younger and older individuals. For many households—especially those headed by older individuals—wealth is large compared with income. For families headed by a thirty-five- to forty-four-year-old, mean net worth is 4.1 times as large as mean income. For families headed by a sixty-five- to seventy-four-year-old, mean net worth is 11.6 times as large as mean income. The ratio of median net worth to median income also increases substantially with age, increasing from 1.4 to 5.7 for these age groups (author's calculations using data from Bucks et al. 2006). Failure to control for these large differences in wealth may lead analysts to underestimate the extent to which VSL decreases with age.

A number of scholars have proposed alternatives to valuing risk reduction using the affected individuals' current values. Pratt and Zeckhauser (1996) argue that empirical estimates of VSL should be adjusted to reflect the values that would be agreed to by a society of expected-utility maximizers choosing behind a "veil of ignorance," as proposed by Harsanyi (1953, 1955) and Rawls (1971). Such individuals would know the characteristics of the society of which they would become members (e.g., the distribution of health and wealth), but would not know their own particular characteristics. This proposed approach requires adjusting VSL to offset the "dead-anyway effect" that induces individuals facing higher current mortality risk to be willing to spend more for current risk reduction. Because they are unlikely to survive and have the opportunity to spend on future consumption, the *private* opportunity cost of spending is low for such individuals. However, the *social* opportunity cost of spending is not reduced by an individual's high risk. The proposed adjustment of VSL would reduce the valuation of risks to people with high current risk, who tend to have shorter life expectancies conditional on surviving the current period.

In many areas of public health, interventions are evaluated in terms of the expected change in "quality adjusted life-years" (QALYs) in the population. This approach, which is similar to using a common VSLY for all interventions that alter mortality risk, has been justified as an approximation to benefit-cost analysis (e.g., Garber et al. 1996; Garber and Phelps 1997). Similarly, Adler (2006) suggests that life-years, or QALYs, may provide a better approximation to (unobservable) individual welfare than WTP measures, especially when welfare is understood as reflecting the decisions individuals would make behind a veil of ignorance.

Evaluating interventions using the expected change in life-years, or QALYs, has been justified using the argument that everyone in society has an equal claim to health, and so, whether those who gain could compensate those who are harmed by an intervention is irrelevant (see, e.g., Culyer 1989; Garber et al. 1996; Hammitt 2002; Williams 1993). A related rationale for assigning higher priority to interventions that reduce mortality risk to younger rather than to older people is the "fair innings" argument, based on the notion that everyone is entitled to a "normal" lifespan (e.g., the biblical three-score years and ten) and that those who die younger are "cheated" while those who die older are "living on borrowed time" (Williams 1997).

Conclusions

Ideally, the effects of environmental policy on mortality risk can be characterized as shifts in individual survival curves. The monetary value to an individual of a specified change in her survival curve caused by an intervention will depend on the magnitude and timing of changes in her mortality risk. Any shift in a survival curve can be described by either the unique change in life expectancy or the unique expected number of lives saved in a specified time period. Hence the value of the shift can be described either by the average saved VSL or the average saved VSLY. Note that an individual may assign different monetary values to alternative changes in the survival curve that produce the same number of life-years saved or the same number of lives saved in a specified time period, resulting in different average VSL and VSLY.

In theory, both individual VSL and individual VSLY may change with age, life expectancy, anticipated future health, income, and other factors. It is not possible for both VSL and VSLY to remain constant over the life cycle since life expectancy changes with age, health status, and other factors. If empirical evidence were to suggest that either VSL or VSLY were reasonably constant across individuals and the types of shifts in survival curves produced by policy changes, then it would be convenient to use that concept as the primary measure for policy evaluation.

Empirical estimation of the effect of life expectancy on valuation of mortality risk is impeded by difficulty in estimating individual life expectancy. While life expectancy varies with age, health status, sex, race, and income, these factors may influence valuation of mortality risk for reasons unrelated to life expectancy, making it difficult to identify the pure effect of life expectancy on valuation. As detailed in the accompanying papers in this symposium, most of the empirical evidence concerns how valuation of mortality risk varies with age. The revealed-preference studies reviewed by Aldy and Viscusi (2007) suggest that VSL follows an inverted U shape, consistent with some theoretical multiperiod models. However, the extent to which VSL varies with age is rather uncertain. Over the younger ages during which VSL increases, life expectancy decreases, and so VSLY must increase more rapidly than VSL. But over the older ages during which VSL decreases, the rate of decrease in VSL appears to be somewhat more rapid than the rate at which life expectancy decreases, suggesting that VSLY decreases with age. An important limitation of the wage-differential studies is that they provide little information about the preferences of elderly and less healthy individuals who are unlikely to be in the labor force, but may be those who are most affected by improvements in air quality or other environmental conditions. The stated-preference studies, reviewed by Krupnick (2007), do not provide as clear a picture. The studies that are arguably most relevant to estimating how VSL varies with age suggest that VSL is reasonably constant over adult ages, with the exception that it falls modestly at ages of perhaps seventy and above. This pattern suggests that VSLY increases with age.

If one chooses to rely on the conventional economic justification for benefit-cost analysis, that is, the identification of potential Pareto improvements, then policies that affect mortality risk should be evaluated using the weighted average of the individual-specific values of the affected population, where the weights are proportional to the individual-specific risk changes (in practice one would use relatively homogenous population subgroups rather than individuals). The individual-specific values of the changes in individual survival curves induced by the policy may be summarized using either VSL or VSLY approaches; the choice is arbitrary. The critical research need is to develop improved empirical estimates of the dependence of valuation on life expectancy and other factors. For current decisions, however, the papers in this symposium provide an excellent review of current empirical evidence about how VSL varies with age.

Acknowledgment

I thank Anna Alberini, Maureen Cropper, John Graham, Michael Greenstone, Per-Olov Johansson, Michael Jones-Lee, Emmett Keeler, Ari Rabl, Kerry Smith, and Milton Weinstein for thoughtful comments, and Suzanne Leonard for improving the presentation.

References

Adler, M. D. 2006. QALYs and policy evaluation: a new perspective. *Yale Journal of Health Policy, Law and Ethics* 6, no. .

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, no. 2: 241–60.

Arias, E. 2002. United States life tables, 2000. *National Vital Statistics Reports* 51, no. (3), Centers for Disease Control and Prevention, U.S. Department of Health and Human Services, December 19.

Bucks, B. K., A. B. Kennickell, and K. B. Moore. 2006. Recent changes in U.S. Family finances: evidence from the 2001 and 2004 survey of consumer finances. *Federal Reserve Bulletin* 92: A1–A38.

Cave, J. A. K. 1988. Age, time, and the measurement of mortality benefits, RAND R-3557, Santa Monica, CA.

Cropper, M. L., and F. G. Sussman. 1990. Valuing future risks to life. *Journal of Environmental Economics and Management* 19: 160–74.

Cropper, M. L., S. K. Aydede, and P. R. Portney. 1994. Preferences for life-saving programs: how the public discounts time and age. *Journal of Risk and Uncertainty* 8: 243–65.

Culyer, A. J. 1989. The normative economics of health care finance and provision. *Oxford Review of Economic Policy* 5: 34–58.

Dolan, P., R. Shaw, A. Tsuchiya, and A. Williams. 2005. QALY maximisation and people's preferences: a methodological review of the literature. *Health Economics* 14: 197–208.

Drèze, J. 1962. L'utilité Sociale d'une Vie Humaine. *Revue Française de Recherche Opérationelle* 6: 93–118.

European Commission. 2005. Recommended Interim Values for the Value of Preventing a Fatality in DG Environment Cost Benefit Analysis, http://ec.europa.eu/environment/enveco/others/value_of_life.htm, last updated January 26.

Evans, M. R., and V. K. Smith. 2006. Do we really understand the age-VSL relationship? *Resource and Energy Economics* 28: 242–61.

Garber, A. M., and C. E. Phelps. 1997. Economic foundations of cost-effectiveness analysis. *Journal of Health Economics* 16: 1–31.

Garber, A. M., M. C. Weinstein, G. W. Torrance, and M. S. Kamlet. 1996. Theoretical foundations of cost-effectiveness analysis. In *Cost-effectiveness in health and medicine*, ed. M. R. Gold, J. E. Siegel, L. B. Russell, and M. C. Weinstein, 25–53. Oxford: Oxford University Press.

Gold, M. R., J. E. Siegel, L. B. Russell, and M. C., Weinstein, eds. 1996. *Cost-effectiveness in health and medicine*. Oxford: Oxford University Press.

Graham, J. D. 2003. Benefit-Cost Methods and Lifesaving Rules, Memorandum to the President's Management Council, Office of Management and Budget, May 30.

Hammitt, J. K. 2000. Valuing mortality risk: theory and practice. *Environmental Science and Technology* 34: 1396–1400.

———. 2002. QALYs versus WTP. *Risk Analysis* 22: 985–1001.

Harsanyi, J. C. 1953. Cardinal utility in welfare economics and in the theory of risk taking. *Journal of Political Economy* 61: 434–35.

———. 1955. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy* 63: 309–21.

Howard, R. A. 1984. On fates comparable to death. *Management Science* 30: 407–22.

Hurd, M. D., and K. McGarry. 1995. Evaluation of the subjective probabilities of survival in the health and retirement study. *Journal of Human Resources* 30: S268–S292.

———. 2002. The predicative validity of subjective probabilities of survival. *Economic Journal* 112: 966–85.

Johansson, P.-O. 2002. On the definition and age dependency of the value of a statistical life. *Journal of Risk and Uncertainty* 25: 251–63.

Johannesson, M., and P.-O. Johansson. 1997. Is the valuation of a QALY gained independent of age? Some empirical evidence. *Journal of Health Economics* 16: 589–99.

Johannesson, M., P.-O. Johansson, and K.-G. Lofgren. 1997. On the value of changes in life expectancy: blips versus parametric changes. *Journal of Risk and Uncertainty* 15: 221–39.

Jones-Lee, M. W. 1974. The value of changes in the probability of death or injury. *Journal of Political Economy* 82: 835–49.

———. 2004. Valuing international safety externalities: does the "Golden Rule" Apply? *Journal of Risk and Uncertainty* 29: 277–87.

Krupnick, A. 2007. Mortality-risk valuation and age: stated preference evidence. *Review of Environmental Economics and Policy* 1, no. 2: 261–82.

Mishan, E. J. 1971. Evaluation of life and limb: a theoretical approach. *Journal of Political Economy* 79: 687–705.

Murray, C. J. L., and A. D. Lopez, 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.

Ng, Y.-K. 1992. The older the more valuable: divergence between utility and dollar values of life as one ages. *Journal of Economics* 55: 1–16.

Office of Management and Budget. Executive Office of the President. 2003. Circular A-4, September 17, (http://www.whitehouse.gov/omb/circulars/a004/a-4.pdf).

Pratt, J. W., and R. J. Zeckhauser. 1996. Willingness to pay and the distribution of risk and wealth. *Journal of Political Economy* 104: 747–63.

Rabl, A. 2003. Interpretation of air pollution mortality: number of deaths or years of life lost? *Journal of the Air & Waste Management Association* 53: 41–50.

———. 2006. Analysis of air pollution mortality in terms of life expectancy changes: relation between time series, intervention, and cohort studies. *Environmental Health: A Global Access Science Source* 5: 1–11.

Rawls, J. 1971. *A theory of justice*. Cambridge: Belknap Press of Harvard University Press.

Rosen, S. 1988. The value of changes in life expectancy. *Journal of Risk and Uncertainty* 1: 285–304.

Schelling, T. C. 1968. The life you save may be your own. In *Problems in public expenditure analysis*, ed. S. B. Chase. Washington, DC: Brookings.

Shepard, D. S., and R. J. Zeckhauser. 1984. Survival versus consumption. *Management Science* 30: 423–39.

Smith, V. K., D. H. Taylor Jr, and F. A. Sloan. 2001. Longevity expectations and death: can people predict their own demise? *American Economic Review* 91: 1126–134.

Stokey, E., and R. Zeckhauser. 1978. *A Primer for Policy Analysis*. New York, NY: Norton.

Viscusi, W. K., and J. K. Hakes. 2003. Risk ratings that do not measure probabilities. *Journal of Risk Research* 6: 23–43.

Weinstein, M. C., D. S. Shepard, and J. S. Pliskin. 1980. The economic value of changing mortality probabilities: a decision-theoretic approach. *Quarterly Journal of Economics* 94: 373–96.

Williams, A. 1993. Cost-benefit analysis: applied welfare economics or general decision aid. In *Efficiency in the public sector*, eds. A. Williams, and E. Giardina. London: Edward Elgar.

———. 1997. Intergenerational equity: an exploration of the "Fair Innings" argument. *Health Economics* 6: 117–32.