Harvard Center for
Risk Analysis

March 2008
Volume 16, Issue 1

Risk in Perspective

Valuing "Lives Saved" vs. "Life-Years Saved"

James K. Hammitt

"Theory and empirical evidence suggest that neither VSL nor VSLY is constant over an individual's life. Accurate valuation requires using values that depend on characteristics of the affected individuals."

Introduction

There is long-standing debate whether to count "lives saved" or "life-years saved" when evaluating policies to reduce mortality risk. Historically, the two approaches have been applied in different domains. Environmental and transportation policies have often been evaluated using lives saved, while life-years saved has been the preferred metric in other areas of public health including medicine, vaccination, and disease screening. For benefit-cost analysis, the monetary value of risk reductions can be calculated either by multiplying expected lives saved by the "value per statistical life" (VSL) or by multiplying expected life-years saved by the "value per statistical life-year" (VSLY).

The choice between metrics can affect the apparent merits of regulatory programs that affect people with different life expectancies. For example, the air-pollution measures that dominate the U.S. regulatory agenda may disproportionately benefit people who are older or in poor health and have shorter-than-average life expectancies. The benefits of these rules may appear larger when counting lives rather than life-years saved. Conversely, policies to protect children such as car-seat requirements may appear to be larger when counting life-years saved. The U.S. Office of Management and Budget has encouraged federal agencies to analyze regulations using both the lives-saved and life-years-saved approaches.

In this issue of *Risk in Perspective*, I describe how the value of any intervention that alters mortality risk can be expressed using either lives saved or life-years saved and the appropriate VSL or VSLY. However, because theory and empirical evidence suggest that *neither* VSL nor VSLY is constant over an individual's life, accurate valuation requires using values that depend on characteristics of the affected individuals. The extent to which VSL and VSLY vary with life expectancy, age, health, and other factors is not clear; better empirical evidence is required.

Key Concepts

Describing environmental, health, and safety interventions as "saving lives" or "saving life-years" can be misleading. Reduction in exposure to a hazard typically reduces some people's chances of dying from that hazard by a small amount, thereby "saving lives" (at least from that hazard). Reducing the risk of dying now increases the risk of dying later, so these lives are not saved forever but life-years are gained. It is usually impossible to know either beforehand or afterward whose death will be or was averted, so the lives or life-years that are saved are anonymous.

Value per Statistical Life (VSL)

An individual's value per statistical life (VSL) is defined as her rate of tradeoff between wealth and small changes in mortality risk in a defined time period. For example, if a typical individual is willing to pay at most $5 to reduce her chance of dying this year by 1 in a million, her VSL is $5 ÷ (1 in a million) or $5 million. The term "value per statistical life" can be understood by recognizing that if each of 1 million people were willing to pay $5 to reduce his or her chance of dying this year by 1 in a million, a total of $5 million dollars would be pledged and one fewer death would be expected.

VSL is defined for very small changes in risk. Theoretically, an individual's VSL is the slope of an indifference curve representing her preferences for wealth and survival probability. As illustrated in Figure 1, an individual's rate of tradeoff between wealth and risk will depend on the size of the risk change. A typical person who would pay $5 to reduce her chance of dying this year by 1 in a million could not afford to pay $1 million to reduce her chance of dying this year by one in five (e.g., from 21 percent to 1 percent). Similarly, she would not accept certain death this year in exchange for $5 million. (As illustrated in Figure 1, she might not accept certain death in exchange for any amount of wealth.)

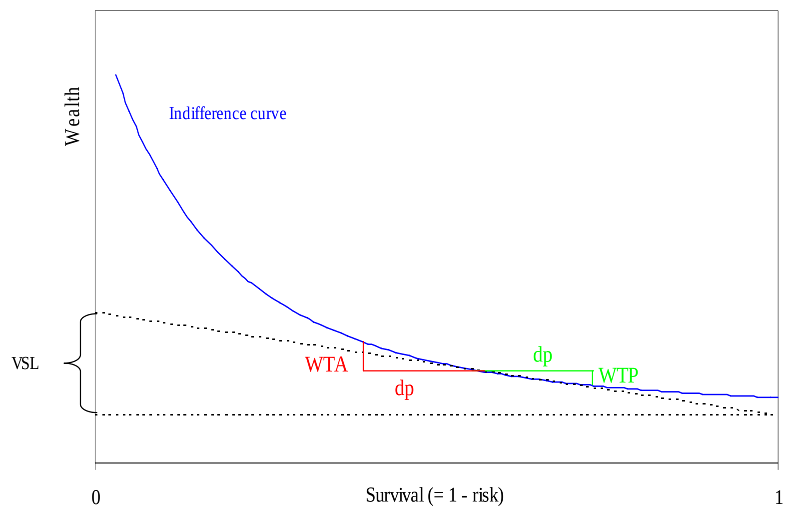

FIGURE 1. Value per statistical life (VSL) is equal to the slope of the individual's indifference curve at her wealth and survival probability, i.e., WTP/dp for a small increase in survival probability and WTA/dp for a small decrease in survival probability. As the risk change dp increases, willingness to pay (WTP) increases less than proportionately to dp and willingness to accept (WTA) increases more than proportionately to dp.

Value per Statistical Life-Year (VSLY)

It seems intuitive that the value to an individual of delaying her death depends on the duration of the delay. Postponing death by years is usually preferred to postponing it by only days (even if those days are precious). 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 increase in the chance of surviving the year multiplied by her life expectancy conditional on survival. The value of this gain is equal to the expected number of life-years saved multiplied by the VSLY. (Future life-years are usually discounted and the value of reducing the risk of dying this year is proportional to the expected present value of future life-years.)

Survival Curves

Changes in mortality risk are most accurately described using survival curves. Survival curves can be constructed for an individual or a population. An individual survival curve plots the probability that an individual will remain alive as a function of age (or calendar date). A population survival curve plots the fraction of a population that remains living as a function of age or date. A survival curve can be constructed beginning at any age or date. The height of the curve begins at one and declines as age and time increase. The slope of the curve depends on the mortality risk, with steeper decreases in periods of higher mortality risk.

The survival curve for the U.S. population beginning at age 60 is illustrated by the solid line in Figure 2. Life expectancy at any age is the area under the survival curve that begins at that age. For the solid curve, life expectancy at age 60 is 22 years.

Any pattern of change in mortality risk over time can be described by the corresponding shift in the survival curve. Moreover, any change in a survival curve implies a unique expected number of life-years saved (the change in the area under the curve) and a unique expected number of lives saved at each point in time (the vertical shift in the curve at that 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 much longer than a century or so, 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 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.

To illustrate, consider an intervention that decreases annual mortality probability by one third, persists for 10 years, and begins after a 10 year lag. The dashed line in Figure 2 illustrates the effect of this intervention for the average 60 year old. There is no change in survival probability during the lag period (ages 60 through 69). The survival curve is flatter over the period during which risk is reduced (ages 70 through 79) and remains higher than the baseline survival curve for later years (ages 80 and above). For each 60 year old affected, this intervention saves one life-year and 0.08 lives (between ages 70 and 79).

Valuing Reductions in Mortality Risk

Each person's willingness to pay (WTP) for a specified shift in her survival curve can be estimated by dividing the shift into a series of instantaneous changes in risk and summing her WTP for each of these "blips" (Johannesson et al, 1997). WTP for each blip depends on the size of the risk reduction, time of payment, conditions under which the individual can save and borrow against future income, and other factors.

WTP for a shift in the survival curve can be described using either VSLY or VSL. The individual's average VSLY for a shift is her WTP for the shift divided by the expected number of life-years saved. Her average VSL for a shift is her WTP divided by the expected number of lives saved in a specified time period. If an individual is willing to pay at most $80,000 for the (very large) shift in the survival curve illustrated in Figure 2, her average VSLY for this shift is $80,000 (= $80,000 ÷ 1 life-year saved) and her average VSL is $1 million (= $80,000 ÷ 0.08 lives saved).

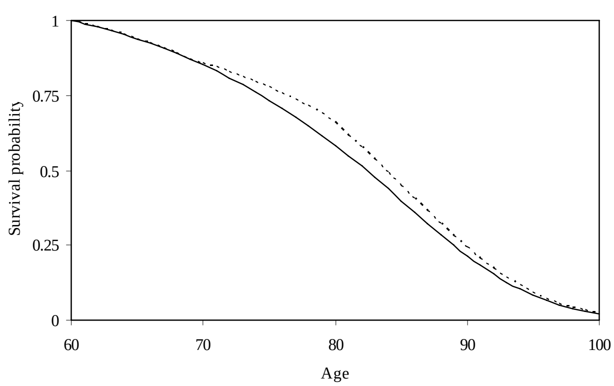

Figure 2. Survival curves from age 60. Solid line is for the U.S. population. Dashed line corresponds to decreasing mortality risk at each age from 71 through 80 by one third.

Estimating WTP to Reduce Mortality Risk

Average VSL and average VSLY for a shift in a survival curve may depend on the initial survival curve, the details of the shift, and other factors. Economic theory does not predict that either VSL or VSLY will be constant across interventions or individuals. In theory, both VSL and VSLY may change with age, life expectancy, anticipated future health, income, and other factors.

Most studies of WTP to reduce mortality risk estimate the value of a reduction in current risk, i.e., VSL. These studies do not provide information on how VSL varies with life expectancy *per se*, but some provide information on how it varies with age. Although it seems intuitive that WTP to reduce current risk decreases with age, as the number of future life-years at stake declines, VSL may remain constant or even increase with age. One reason is that income and wealth usually rise with age, at least over part of the lifecycle. Higher income and wealth increase ability to pay and hence increase VSL. In addition, as life expectancy declines an individual may have less reason to conserve her wealth for future needs. Some economic models suggest that VSL follows an inverted U, rising through middle age and falling at older ages.

Empirical estimates are derived using either “revealed-preference” or “stated-preference” methods. Revealed-preference studies are based on observation of behaviors that affect mortality risk and wealth. The most common are wage-differential studies that estimate the extra pay workers receive for more hazardous jobs. In contrast, stated-preference studies ask survey respondents about hypothetical choices. Many of these ask about choices between safer but more expensive food or transportation or about hypothetical medicines or disease-screening programs.

Recent wage-differential studies that examine how VSL varies with age support the inverted-U hypothesis (Aldy and Viscusi, 2007). These studies are limited in that they necessarily include only employed workers and thereby exclude the elderly and those in poor health. Stated-preference studies, which can include a broader population, yield mixed results. Some suggest little or no effect of age on VSL and others suggest a modest decrease at older ages (Krupnick, 2007).

It is not possible for both VSL and VSLY to remain constant over the lifecycle since life expectancy changes with age. The empirical evidence suggests that *neither* VSLY nor VSL is constant over the lifespan. If VSL follows an inverted U, then VSLY must increase at young ages when VSL is increasing and life expectancy is decreasing. VSLY could be roughly constant at older ages if the rate at which VSL decreases coincides with the rate at which life expectancy decreases. Some wage-differential studies suggest VSL declines more rapidly than life expectancy, which implies that VSLY decreases with age (Aldy and Viscusi, 2007). If VSL is roughly constant with age or decreases only modestly at older ages, as the stated-preference studies seem to suggest (Krupnick, 2007), then VSLY increases over the lifespan.

Estimating Lives Saved and Life-Years Saved

In practice, it may be more difficult to estimate the effect of an intervention on either lives or life-years saved. For mortality risks where 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 can be estimated but the number of life-years saved cannot without information about the life expectancy of the affected population. Similarly, time-series studies that analyze how daily air-pollution levels affect the number of deaths each day provide estimates of the number of lives that may be saved by reducing pollution, but not the number of life-years. The problem is the difficulty in knowing whether the people who succumb to these hazards have the same life expectancy as others of their age and sex or are more susceptible to these (or other) risks.

In contrast, cohort studies that monitor populations exposed to different levels of pollution (e.g., those living in different cities) provide estimates of the survival curve and how it depends on pollution. These studies can be used to estimate the number of life-years saved but not the number of lives saved. The same shift in the population survival curve can be the result of extending the lives of many people for a short time or the lives of fewer people for a longer time (Rabl, 2003). To illustrate, consider a stylized example: In a “polluted” city, half the population dies at age 60 and the other half at age 70. In a “clean” city, half the population dies at 70 and the other half at 80. Living in the clean city is associated with an increase of 10 statistical life-years per capita. However, it is impossible to determine from the survival curves alone whether the difference arises because everyone in the clean city lives 10 years longer than they would in the dirty city, or because the people who die at 60 in the dirty city would have lived to 80 in the clean city. In the first case, the number of lives saved (at ages 60 and 70) 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 may not exist and risk estimates are often 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.

Conclusion

The effects of environmental, health, and safety interventions on mortality risk are most accurately characterized as shifts of individuals' survival curves. The monetary value to an individual of a change in her survival curve depends on the magnitude and timing of changes in her mortality risk. Any shift in a survival curve implies expected numbers of both life-years gained and lives saved in a specified time period. Hence the value of the shift can be described by the corresponding average value per life-year saved (VSLY) or value per life saved (VSL); the choice between these measures is arbitrary.

Note that an individual may assign different monetary values to alternative changes in her survival curve, even if they 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.

Accurately valuing changes in mortality risk requires using values that are appropriate to the risk change, which may depend on the age, health, life expectancy, and other characteristics of the affected population. The existing empirical literature offers conflicting evidence about how VSL and VSLY change with age; better empirical evidence is needed.

Note: This essay is adapted from Hammitt (2007).

For further reading, see:

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

Hammitt, J.K., "Valuing Changes in Mortality Risk: Lives Saved Versus Life Years Saved," *Review of Environmental Economics and Policy* 1: 228-240, 2007.

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

Johannesson, M., P.-O. Johansson, and K.-G. Lofgren., "On the Value of Changes in Life Expectancy: Blips Versus Parametric Changes," *Journal of Risk and Uncertainty* 15: 221-239, 1997.

Krupnick, A., "The Senior Death Discount and Stated Preference Evidence," *Review of Environmental Economics and Policy* 1: 261-282, 2007.

Rabl, A., "Interpretation of Air Pollution Mortality: Number of Deaths or Years of Life Lost?" *Journal of the Air and Waste Management Association* 53: 41-50, 2003.

Robinson, L.A., "How US Government Agencies Value Mortality Risk Reductions," *Review of Environmental Economics and Policy* 1: 283-299, 2007.

Peer reviewer:
Henrik Andersson, Ph.D.

Harvard Center for Risk Analysis
Harvard School of Public Health
Email: hcra@hsph.harvard.edu
Website: http://www.hcra.harvard.edu