WORKING PAPERS

N° 1177

January 2021

"Fatality Risk Regulation"

James K. Hammitt and Nicolas Treich

Toulouse School of Economics

Fatality Risk Regulation

James K. Hammitt[a] and Nicolas Treich[b]

January 2021

Abstract

Environmental, transportation, occupational, and other regulations that reduce fatality risk are frequently evaluated using benefit-cost analysis (BCA). We examine how risk reductions are valued under BCA, utilitarian and prioritarian SWFs. The social value of risk reduction (SVRR) to an individual is the rate of increase of social welfare for a small decrease to the individual's current-period fatality risk. Under BCA, the SVRR is the individual's value per statistical life (VSL), which is increasing in wealth and baseline risk. Under utilitarian and prioritarian SWFs, the SVRR is far less sensitive to income; it can decrease with income for prioritarian SWFs that exhibit sufficient inequality aversion. The SVRR increases with or is independent of baseline risk. Like VSL, it can increase or decrease with age, but prioritarian SWFs assign larger SVRR to younger relative to older individuals than does the utilitarian SWF. Extensions to catastrophe aversion and nonfatal health risks are discussed.

Keywords: mortality, regulation, benefit-cost analysis, value of a statistical life, fair innings, age, income, catastrophe aversion, prioritarian, utilitarian

Prepared for publication as chapter 7 in M. Adler and O.F. Norheim (eds.), *Prioritarianism in Practice*, Cambridge University Press.

a. Harvard University (Center for Risk Analysis, T.H. Chan School of Public Health) and Toulouse School of Economics, University Toulouse Capitole
b. INRAE, Toulouse School of Economics, University Toulouse Capitole

7.1. Introduction

Many forms of government regulation are designed to decrease risks of death from particular hazards, leading to increases in life expectancy. In some cases, reducing risks of death (and of nonfatal illness or injury) is the primary objective (as in food, transportation, and occupational safety); in other cases, reducing health risks is but one of several objectives (as in much of environmental regulation). Even when it is not the sole objective, reduction of fatality risks can dominate the quantified benefits. For example, estimates of the benefits of the U.S. 1990 Clean Air Act amendments find that 85 to 95 percent of the quantified annual benefit is due to the reduction in mortality risk (U.S. Environmental Protection Agency 2011, Table 7-1). Lee and Taylor (2019) estimate that reductions in mortality risk account for up to 70 percent of the quantified benefits of recent major U.S. federal regulations. In part this dominance reflects the much greater value estimated for reducing fatal rather than nonfatal health risks. It also reflects the fact that many nonfatal health effects and environmental consequences are not included in the evaluation because monetary values have not been estimated for many of these numerous and diverse effects.

Regulations that affect fatality risks are frequently evaluated using benefit-cost analysis (BCA). In the United States, such evaluation has been required for federal regulations under a series of presidential executive orders beginning in 1981; regulatory impact analysis, including quantification of benefits, costs, and distributional effects, is required by nearly all OECD member states (OECD 2018). We describe the way fatality risks are evaluated using BCA and contrast it with evaluation under several prominent social-welfare functions.

As described in the following section, the approach to valuing mortality risk used in BCA centers on the 'value per statistical life' or VSL (called the 'value of a prevented fatality' or VPF in the United Kingdom). Under standard theory, the VSL is expected to vary across individuals who differ in wealth, baseline mortality risk, and other dimensions such as age. Differences in VSL between individuals imply that the benefits of reducing fatality risk depend on whose risk is reduced. Government actors and others are often uncomfortable with explicitly valuing risk reductions differently across individuals; there has been sharp public criticism when values have been differentiated based on income or age (e.g., Cameron 2010, Robinson 2007, Viscusi 2009). Perhaps as a result, it is common practice to use the same VSL for all affected individuals in a country (although it is also common to assume that VSL will increase over time in response to assumed future increases in income). The reluctance to allowing VSL to vary across individuals may also reflect the tension between efficiency and distributional concerns; conventional BCA measures allocative efficiency, setting concerns about the distribution of effects aside on the assumption they will be addressed through supplementary distributional analysis.

Although most distributional analysis of U.S. regulations is qualitative (Robinson et al. 2016), a growing literature explores the use of inequality measures to quantify distributional effects, especially in the context of environmental justice. This literature uses inequality indices to evaluate the distribution of risk from, or exposure to, environmental hazards and whether racial/ethnic or low-income groups are particularly burdened. It has been applied to multiple case studies (see Levy 2021 and Sheriff and Maguire 2020 for citations). Key issues include adapting measures of the distribution of goods (e.g., income) to the distribution of bads (e.g., mortality risk), evaluating absolute or proportional differences, the importance of parameterizing inequality aversion, and whether inequality is evaluated across groups that are ordered (e.g., by income) or unordered (e.g., by race or ethnicity) (Levy 2021). Levy et al. (2006) reviewed multiple indices and recommended use of the Atkinson index, which measures relative differences; Sheriff and Maguire (2020) recommended the use of Fleurbaey's equally distributed equivalent measure with the Kolm-Pollak inequality index. The Kolm-Pollak index is sensitive to absolute differences and hence yields the same result when measuring differences in mortality risk as differences in survival probability; in contrast, the Atkinson index is sensitive to proportional differences and yields different results.[1]

In contrast to conventional BCA, SWFs can integrate concerns about equitable distribution with those about efficiency. We present and describe the results of using SWFs to characterize the value of reducing mortality risk through regulation. For consistency with the literature that applies BCA to environmental, health, and safety risks, we measure individual wellbeing using von Neumann-Morgenstern utility functions; for simplicity, we assume that utility functions are homogenous (i.e., that individuals have the same utility function but may differ in wealth, baseline risk, or other characteristics). Adler and Decancq (chapter 3, this volume) describe alternative measures of individual wellbeing.

We examine three benchmark SWFs: utilitarian, ex ante prioritarian, and ex post prioritarian. The utilitarian SWF is the dominant approach in the SWF literature; prioritarian SWFs add a concern for the distribution of wellbeing (Adler, chapter 2, this volume). In the context of fatality risk, the

[1] Atkinson and Kolm-Pollak measures of inequality are described by Decancq and Schokkaert (chapter 5, this volume).

difference between ex ante and ex post perspectives is important, and so we examine both ex ante and ex post prioritarian SWFs.[2]

This chapter is closely related to other chapters in this volume concerned with health risk, particularly chapter 6 on health care (Cookson, Norheim, and Skarda) and chapter 12 on COVID-19 (Bloom, Ferranna, and Sevilla).

In the following section, the standard one-period model for VSL is presented together with analysis of how VSL varies with individual characteristics, specifically wealth and baseline risk. Sections 7.3 and 7.4 apply the benchmark SWFs, both utilitarian and prioritarian, to evaluate the social value of mortality-risk reduction and how it depends on individuals' wealth and baseline risk. Section 7.3 analyzes the one-period model introduced in Section 7.2 and Section 7.4 describes a multiperiod model, which permits consideration of the effects of individuals' ages. Section 7.5 presents simulations calibrated to the U.S. population that illustrate how the value of mortality-risk reduction varies by age and income under BCA and the benchmark SWFs. Section 7.6 describes two extensions: alternative SWFs that can incorporate aversion to catastrophes that produce many deaths, and evaluating effects of health and nonfatal health risks using BCA and the benchmark SWFs. Section 7.7 concludes.

7.2. Valuing mortality risk changes in BCA: the value per statistical life

Benefit-cost analysis is based on measuring changes in individuals' wellbeing as monetary values. These values are positive for individuals who are better off with a policy and negative for those who are worse off. The sum of these values across the population is described as the net social benefit; if it is positive, then hypothetical compensation payments (from those who gain to those who are harmed) could be arranged such that everyone would be better off with the policy change and payment or receipt of compensation than without the policy change and compensation.

There are two alternative definitions of the monetary value of the effect of a change in mortality risk (or other policy effect). The compensating surplus is the change in wealth that compensates for the change in risk in the sense that the individual's wellbeing is unaffected by the combination of the change in risk and wealth. Compensating surplus is the individual's willingness to pay (WTP) for a risk reduction and willingness to accept compensation (WTA) for a risk increase. The equivalent surplus is

[2] Adler (chapter 2, this volume) uses the term SWF to mean a rule for ranking well-being vectors and the term uncertainty module to mean a method for incorporating uncertainty about the vector. We use the term SWF to mean the combination of an SWF and an uncertainty module and hence describe ex ante prioritarianism and ex post prioritarianism as different SWFs.

the change in wealth that is equivalent to the change in risk, in the sense that either change alone has the same effect on her wellbeing. Equivalent surplus is the individual's WTA to forgo a risk decrease and WTP to prevent a risk increase. For small changes in risk, the two values are similar in theory, although stated-preference studies find that the equivalent surplus (WTA to forgo a risk reduction) is often substantially larger than the compensating surplus (WTP for a risk reduction), by a factor on the order of five (Tunçel and Hammitt 2014).

The value per statistical life (VSL) is conceptualized as an individual-specific quantity equal to the individual's marginal rate of substitution of wealth for a reduction of her mortality risk in a specified period. It is the ratio of her compensating surplus for a change in mortality risk to the magnitude of the change, for infinitesimally small changes in risk. The time period is generally short (e.g., a year). Risk is defined as the probability of dying in a specified period (as a hazard rate per unit time). The value of surviving a period may depend on its duration; the value of surviving the current decade generally exceeds the value of surviving the current hour.

The conventional one-period model (Drèze 1962, Jones-Lee 1974, Weinstein et al. 1980) assumes the individual wishes to maximize her expected utility, which depends on her probability of surviving the current period and her utility conditional on surviving and on not surviving (each of which may depend on wealth). Specifically, expected utility

$$U = p\,u(c) + (1 - p)\,v(c) \tag{7.1}$$

where p is the probability of surviving the period, $u(c)$ and $v(c)$ are utility conditional on surviving and not surviving the period, and c is wealth. The utility of wealth conditional on surviving the period depends on the lottery over future conditions the individual will face if she survives, including her future longevity, health, and income. The utility of wealth conditional on dying during the period is often called a bequest function; it represents the utility gained by knowing one's wealth will be passed on to one's children or other heirs in the event of one's death. Any monetary effect of dying, such as medical bills or funeral expenses, can be incorporated in the definition of the bequest function.

It is conventional and usually reasonable to assume that survival is preferred to death ($u(c) > v(c)$), the marginal utility of wealth is strictly larger given survival than as a bequest ($u'(c) > v'(c)$) and the marginal utility of a bequest is greater than or equal to zero ($v'(c) \geq 0$). In words, it is never harmful to die with more wealth but it is better to spend it while living. It is also assumed the individual is risk averse (or risk neutral) with respect to wealth, i.e., $u''(c) \leq 0$ and $v''(c) \leq 0$.[3]

The VSL or marginal rate of substitution of wealth for survival probability is the change in wealth accompanying an infinitesimal change in survival probability that leaves the individual no better and no worse off than before (i.e., holding expected utility U constant). VSL can be obtained by differentiating equation (7.1) to obtain

$$VSL = -\frac{dw}{dp} = \frac{u(c)-v(c)}{pu'(c)+(1-p)v'(c)} = \frac{\Delta u(c)}{Eu'(c)}. \quad (7.2)$$

Equation (7.2) shows that VSL is the ratio of the gain in utility from surviving the period rather than dying ($\Delta u(c)$) to the expected opportunity cost of wealth ($EU'(c)$), which is the rate at which utility could be obtained by spending on other goods and services rather than on reducing mortality risk.

As illustrated in Figure 7.1, VSL is (the absolute value of) the slope of the individual's indifference curve between wealth and survival probability at her current wealth and mortality risk, and the monetary value of a small risk change is approximately equal to the product of VSL and the magnitude of the risk change. Under the assumptions described above, indifference curves are downward sloping and convex, as illustrated. Because increases in wealth and survival probability are both desirable, indifference curves are downward sloping (the individual is willing to give up some of one attribute to obtain more of the other). The indifference curves are convex because holding survival probability fixed, an increase in wealth increases VSL[4] and holding wealth fixed, an increase in survival probability decreases VSL.[5]

A slope is quantified as rise over run; by convention, VSL is measured for a run (change in survival probability) from 0 to 1. As is clear from Figure 7.1, the amount an individual would accept as compensation for a decrease in survival probability from 1 to 0 is much larger than her VSL (at survival probability one) and may be infinite (if the indifference curve does not intersect the vertical axis); the amount she would pay to increase survival probability from 0 to 1 is much less than her VSL

[3] A single prime denotes first derivative and a double prime denotes second derivative.
[4] An increase in wealth increases the numerator of equation (7.2) (because the marginal utility of wealth is larger conditional on survival than on death) and (weakly) decreases the denominator (because the utility conditional on survival and the bequest exhibit weakly decreasing marginal utility).
[5] An increase in survival probability has no effect on the numerator of equation (7.2) but increases the denominator as it puts more weight on the (larger) marginal utility of wealth conditional on survival and less weight on the (smaller) marginal utility of a bequest. Pratt and Zeckhauser (1996) call this the 'dead-anyway effect' because if survival probability is low, spending is likely to come out of the bequest and the opportunity cost of spending is small.

(at survival probability near zero). Indeed, conventional estimates of VSL greatly exceed individual wealth, and so an individual would be unable to pay as much as her VSL to avoid certain death.

Government agents and the public can be uncomfortable applying different values of risk reduction to different people. This concern is exacerbated by the term value per statistical life, which is easily misunderstood as measuring the total value of an individual's life (Cameron 2010, Simon et al. 2019). Benefit-cost analysis of policies that affect a domestic population almost invariably value risk reductions to different individuals using a common VSL, independent of income, age, health, and other factors. For policies that reduce current mortality risk by the same amount for each person, or for which individual risk reductions are uncorrelated with factors such as age and income that affect VSL, the use of a common value is unproblematic.[6] When comparing policies that disproportionately affect different subpopulations (e.g., young v. old, rich v. poor), use of a common VSL may be inconsistent with the principles on which BCA is based and can produce biased results. Use of a common VSL may produce results that better accord with preferences for equity that might be obtained using a weighted BCA. Baker et al. (2008) examined the question of when a common VSL might exist and suggest using weighted BCA.

Social-welfare functions provide a means for explicitly incorporating concerns for both efficiency and equity in a policy evaluation. In Section 7.3, we describe the application of SWFs to valuing fatality risk in the context of the standard one-period model described above. In Section 7.4, we extend the analysis to life-cycle models that allow comparison of the value of reducing risk to individuals of different ages. We present simulation results illustrating the magnitudes of the effects in Section 7.5.

7.3. SWFs applied to a single-period model

Using the single-period model described above (equation (7.1)), Adler et al. (2014) evaluated how changes to individuals' survival probabilities affect social welfare as measured by alternative social-welfare functions. They defined the 'social value of risk reduction' (SVRR) as the rate of increase in social welfare resulting from a small increase to an individual's survival probability, i.e., as the partial derivative of social welfare with respect to an individual's survival probability. SVRR for an SWF is the analog of VSL for BCA, as VSL measures the rate of increase of net social benefits associated with a small increase to an individual's survival probability. Exploiting this analogy, we describe VSL as the SVRR under BCA. Adler et al. (2014) examined how SVRR depends on wealth and baseline risk under

[6] Under BCA, the social value of reducing mortality risk by a small amount for many people equals the sum of the affected individuals' monetary values, which are approximately equal to the product of each individual's risk reduction and VSL. If the individual risk reductions are equal to each other or are uncorrelated with individuals' VSLs, this sum is equal to the sum of the average risk reduction multiplied by the average VSL.

different SWFs; as noted above, VSL is increasing in these attributes. In this section, we describe the results for three benchmark SWFs: the utilitarian, ex ante prioritarian and ex post prioritarian SWFs.

To account for possible dependence of the realization of the mortality risk among individuals (examined in Section 7.6.1), the formalization is slightly more complex than introduced in Section 7.2. Specifically, assume there are *N* individuals indexed by *i*. Let *s* denote the state of nature and π_s denote the probability of state *s*. An individual's outcome in each state includes her wealth c_i (which is assumed to be constant across states) and realization of the mortality risk, i.e., whether she survives or dies during the period. An individual's expected utility $U_i = \sum_s \pi_s \left[l_i^s u(c_i) + (1 - l_i^s) v(c_i) \right]$ where l_i^s is an indicator variable equal to 1 if she survives in state *s* and 0 otherwise. The utility functions conditional on survival and death, *u* and *v*, are assumed to be interpersonally comparable, to satisfy the standard conditions described in Section 7.2, and to be identical across individuals (to allow us to focus on the effects of other interpersonal differences). An individual's probability of surviving the period equals the sum of the probabilities of the states in which she survives, $p_i = \sum_s \pi_s l_i^s$.

An SWF ranks vectors that describe the wellbeing of all individuals. In the context of mortality risk, individuals' wellbeing is uncertain (they may live or die) and so the SWF requires a procedure for ranking risks over wellbeing vectors, an uncertainty module (Adler, chapter 2, this volume). Two natural uncertainty modules for a given SWF are (a) to calculate the expected value of the SWF and (b) to apply the SWF to the vector of expected utilities associated with a policy. In the case of utilitarianism, these are mathematically equivalent. In the case of prioritarianism, they are not.

The utilitarian SWF is the sum of individual utilities. Let S^U denote social welfare under the utilitarian SWF. Because it is the sum of individual utilities, its ex ante and ex post forms are equal:

$$S^U = \sum_s \pi_s \left[\sum_i l_i^s u(c_i) + (1 - l_i^s) v(c_i) \right] = \sum_i \left[\sum_s \pi_s [l_i^s u(c_i) + (1 - l_i^s) v(c_i)] \right]$$

$$= \textstyle\sum_i p_i u(c_i) + (1 - p_i) v(c_i). \qquad (7.3)$$

Note that the first expression (the ex post form) equals the expected value (across states) of the sum of individuals' realized (ex post) utilities. The second expression (the ex ante form) equals the sum across individuals of their expected utilities, before the mortality risks are resolved. The last expression is a simplification of the second expression, using the fact that an individual's survival probability is the sum of the probabilities of the states in which she survives.

In contrast to the utilitarian SWF, the ex post and ex ante prioritarian SWFs differ from each other. Let $g(\cdot)$ be the prioritarian transformation function, which is assumed to be continuously differentiable, strictly increasing, and strictly concave.

The ex post prioritarian SWF is

$$S^{EPP} = \sum_s \pi_s [\sum_i g[l_i^s u(c_i) + (1 - l_i^s) v(c_i)]] = \sum_i p_i g[u(c_i)] + (1 - p_i) g[v(c_i)], \qquad (7.4)$$

which is the expected value (across states) of the sum of individuals' realized utilities as transformed by the function *g*.

The ex ante prioritarian SWF is

$$S^{EAP} = \sum_i g[p_i u(c_i) + (1 - p_i) v(c_i)], \qquad (7.5)$$

which is the sum of the individuals' expected utilities as transformed by the function *g*.

All three of the benchmark SWFs are separable across individuals. Social welfare is the population sum of some individual-specific quantity: expected utility in the case of the utilitarian SWF, expected transformed utility in the case of the ex post prioritarian SWF, and transformed expected utility in the case of the ex ante prioritarian SWF. This implies the effect on social welfare of a change to an individual's risk is independent of the risk and other characteristics of everyone else.

As noted above, the social value of risk reduction (SVRR) for an individual is defined as the rate of change in social welfare in response to a small change to the individual's survival probability, $\frac{\partial S}{\partial p_i}$, for each SWF. VSL is the analogous rate of change of net social benefits under BCA.

For the utilitarian SWF (equation (7.3)), the SVRR for individual *i* is $u(c_i) - v(c_i)$. Note this is the numerator of the VSL (equation (7.2)), which equals the difference in utility between survival and death.

For the ex post prioritarian SWF (equation (7.4)), the SVRR is $g[u(c_i)] - g[v(c_i)]$. This is similar to the SVRR for the utilitarian SWF except the realized utilities are transformed by the function *g*.

For the ex ante prioritarian SWF (equation (7.5)), the SVRR is $g'\left[p_i u(c_i) + (1 - p_i) v(c_i)\right]\left[u(c_i) - v(c_i)\right]$. In words, it equals the SVRR for the utilitarian SWF (the second term, in brackets) multiplied by a factor that depends on the individual's expected utility; this priority factor is larger for individuals with smaller expected utility (because *g* is concave).

Recall that VSL is increasing in the individual's wealth and mortality risk. Consider the effects of mortality risk and wealth on the SVRR under the three benchmark SWFs, summarized in Table 7.1 (the effects of health, shown in the last column, are discussed in Section 7.6.2). Unlike VSL, the SVRR under the utilitarian and ex post prioritarian SWFs are independent of mortality risk, because these SWFs are expected values (of individual wellbeing or transformed wellbeing, respectively) and hence are linear functions of an individual's survival probability. In contrast, the SVRR under the ex ante prioritarian SWF is increasing in mortality risk (because that decreases the individual's expected utility and hence increases her social priority). The sensitivity to baseline risk of VSL and SVRR under the ex ante prioritarian SWF implies that BCA and the ex ante prioritarian SWF exhibit a preference for risk equity, defined as a more equal distribution of risk holding the expected number of fatalities constant. In contrast, the utilitarian and ex post prioritarian SWFs are neutral with respect to risk equity.[7]

Turning to wealth, SVRR under the utilitarian SWF is increasing in wealth, as is VSL. However, the utilitarian SVRR is less sensitive to wealth than is VSL. Recall that SVRR under the utilitarian SWF equals the numerator of VSL (equation (7.2)), which increases with wealth under the assumption that the marginal utility of wealth conditional on survival is greater than as a bequest. The denominator of VSL decreases with wealth, increasing the sensitivity of VSL to wealth compared with the sensitivity of the SVRR under the utilitarian SWF.[8] The effect of wealth on the SVRR under the two prioritarian SWFs depends on the curvature of the transformation function g. Although greater wealth increases an individual's utility, it decreases her social priority and the combined effect of wealth on SVRR can be positive, negative, or zero. If g is close to linear, the prioritarian SWFs are close to the utilitarian SWF and the effect of wealth on SVRR is positive; if g is sufficiently concave, the effect is negative.

To summarize, the effects of mortality risk and wealth on the social value of risk reduction differ between BCA and the three benchmark SWFs. Baseline risk increases the SVRR under BCA and the ex ante prioritarian SWF but it has no effect on the SVRR under the utilitarian and ex post prioritarian SWFs. In contrast, wealth increases the SVRR under BCA and the utilitarian SWF, but its effect under the two prioritarian SWFs can be positive, negative, or zero. The only SWF under which the SVRR can be equal across individuals differing in wealth and baseline risk is the ex post prioritarian SWF, and then only if the curvature of the transformation function exactly offsets the positive effect of wealth on individual utility.

---

[7] Risk equity is discussed in Section 7.6.1.

[8] If the utility of wealth conditional on survival and on death is risk neutral ($u'' = v'' = 0$), then the denominator is unaffected by wealth and SVRR under the utilitarian SWF and VSL are equally sensitive to wealth.

7.4. SWFs applied to a lifetime model

The effect of an individual's age on the social value of reducing her current-period mortality risk is of significant policy interest. Over most of the lifespan, mortality risk increases with age at an accelerating rate. Many important sources of mortality risk increase sharply with age, including cardiovascular disease and many cancers. One of the largest environmental-health risks, exposure to fine-particulate air pollution, is believed to increase mortality risk more for older than younger individuals; the effect of exposure on mortality risk is consistent with the assumption that exposure increases baseline risk by a constant factor, a so-called constant proportional hazard. In contrast, motor-vehicle safety programs often produce larger risk reductions for young adults, who face larger fatality risk from this hazard than do older individuals.

Because life expectancy conditional on surviving the current period tends to be negatively associated with age,[9] it seems intuitive that the social value of risk reduction is larger for younger than for older individuals. Consistent with this perspective, cost-effectiveness analysis of health interventions often measures effectiveness by the expected increase in quality-adjusted life years. Under this metric, decreasing mortality risk to a younger individual (with greater life expectancy) generally produces greater benefits. Some commentators go further, endorsing a 'fair innings' argument that "everyone be given an equal chance to have a fair innings, to reach the appropriate threshold but, having reached it, they have received their entitlement. The rest of their life is the sort of bonus which may be canceled when this is necessary to help others reach the threshold" (Harris 1985). The fair innings perspective recommends giving greater priority to the young than would be justified by their larger expected increase in quality-adjusted life years (Williams 1997).

As described below, neither VSL nor the SVRR under any of the benchmark SWFs we consider necessarily decrease with age; in each case, there are conditions under which the social value of risk reduction is larger for an older than an otherwise identical younger individual. However, both the ex ante and ex post prioritarian SWFs give more priority to a younger individual than does the utilitarian SWF, a property that Adler et al. (2021) denote 'priority to the young'.

The effects of age and life expectancy are implicit in the single-period model introduced in Section 7.2. Utility conditional on survival $u(c)$ implicitly depends on the individual's uncertain prospects if she survives the period, including her future longevity, health, income, and other factors. Hence $u(c)$ can be interpreted as her expected utility conditional on surviving the period. Normally, greater life expectancy will increase the utility of surviving the current period, the utility gain from survival (the

[9] Life expectancy is decreasing in age if the hazard rate (the current mortality risk) does not decrease with age.

numerator of equation (7.2), holding the bequest $v(c)$ constant), and hence the SVRR under the utilitarian and ex post prioritarian SWFs. However, it need not increase SVRR under the ex ante prioritarian SWF or VSL. Under the ex ante prioritarian SWF, the increase in the utility of survival can be more than offset by the decrease in the social priority of the individual. For VSL, the increase in the numerator can be offset by an increase in the expected opportunity cost of spending (the denominator of equation (7.2)). For example, if the individual will have no future income and must support herself from current wealth, the opportunity cost of spending conditional on survival $u'(c)$ will be larger when life expectancy is greater. Depending on the relative magnitudes of its effects on $u(c)$ and on $u'(c)$, the effect of life expectancy on VSL may be positive, zero, or negative.

Further progress in evaluating how the SVRR depends on age requires a more-explicit model of how wellbeing depends on age and life expectancy. The conventional multi-period model assumes wellbeing depends on the present value of future utility, where utility in each period depends on consumption (and implicitly on health and other factors). Although the utility of a bequest might depend on the age at death (or on the age of one's children, which is correlated with age at death), applications of this model set the utility of a bequest to zero (primarily for simplicity).

Under the conventional multi-period model, the individual seeks to maximize her expected lifetime wellbeing, subject to a budget constraint. The age-dependence of VSL depends on the conditions under which an individual can reallocate income[10] to consumption in different periods, e.g., whether she can borrow against future income or save past income to adjust her consumption over time. It also depends on how the rate at which she discounts future utility compares with the interest rate, which affects returns to saving and costs of borrowing. When income increases with age (as it usually does over younger ages), if the individual cannot borrow against future income her VSL first rises then falls with age, exhibiting an inverted-U pattern (Shepard and Zeckhauser 1984). When the individual can both save past income and borrow against future income and the interest rate and utility-discount rate are equal, optimal consumption is constant over time and VSL decreases with age (Shepard and Zeckhauser 1984). In contrast, when the utility discount rate is smaller than the interest rate, it is optimal to defer consumption, at least at younger ages, and the opportunity cost of spending decreases with age. In this case, VSL again follows an inverted-U pattern (Ng 1992).

Adler et al. (2021) analyzed the social value of risk reduction for individuals of different ages using the standard multi-period model under the same benchmark SWFs described in Section 7.3:

[10] It is conventional in these models to assume that income in each period is exogenous. A richer model could allow the individual to alter her income trajectory by investing in education or choosing how many hours to work each period.

utilitarian, ex ante and ex post prioritarian. In the base case, individuals face identical time paths of mortality risk and income, but their current ages differ because of differences in birth date. The SWFs are functions of individuals' lifetime utility. In the model, an individual experiences the utility of consuming in a period only if she survives that period; if not, her lifetime utility includes only the periods before the current period. This assumption is not limiting (periods can be arbitrarily short) but it influences the interpretation of some of the results.

The SVRR depends on whether the individual can save or borrow and, if she can, on whether the policy that reduces mortality risk is anticipated (which affects her planned consumption path) or comes as a surprise. In the base case, an individual's consumption is assumed to be equal to her (exogenously determined) income in each year that she survives; she can neither save nor borrow. In an extension (Adler et al. 2019), the individual is assumed to be able to save or borrow at actuarially fair rates, so she can allocate her expected lifetime wealth to consumption across all periods, but the policy comes as a surprise and so does not affect her consumption plan.

7.4.1. Effects of age and priority to the young

As noted above, the effect of age on VSL depends on whether the individual can finance current consumption by borrowing against future income and on how she discounts future utility compared with the interest rate. Under some conditions, VSL decreases with age; under others, it increases over younger ages then decreases over older ages. Similarly, SVRR under the benchmark SWFs can but does not necessarily decrease with age. However, compared with the utilitarian SWF, both the ex ante and ex post prioritarian SWFs give more priority to reducing risk to younger than to older individuals.

Consider a simple two-period example. At birth each individual has probability p_1 of surviving her first period ("youth") and, conditional on surviving the first period, probability p_2 of surviving her second period ("adulthood"). Her utility of surviving the first period is u_1 and utility of surviving the second period is u_2 (utility for a period she does not survive is zero). A young person, in her first period, has expected lifetime utility $p_1 u_1 + p_1 p_2 u_2$. An old person, in her second period, has already survived her first period and has expected lifetime utility $u_1 + p_2 u_2$. A reduction in current-period risk increases p_1 for the young person and increases p_2 for the old person.

Under the utilitarian SWF, the SVRR for the young person is $u_1 + p_2 u_2$ and the SVRR for the old person is u_2. The difference, SVRR(young) – SVRR(old) = $u_1 + (p_2 - 1) u_2$, can be interpreted as the sum of two effects, a 'duration effect' u_1, which reflects the longer future life the young person may live (two periods rather than one) and a 'risk effect' $(p_2 - 1) u_2$, which reflects the fact that the young person has a smaller chance of living during her second period, because she might not survive her first period, while the old person already has. The duration effect is positive and the risk effect is negative, hence the SVRR may be larger or smaller for the young than for the old person. Typically the difference in SVRR will be positive but it can be negative if u_1 is small compared with u_2 and p_2 is small compared with one.

Similarly, under the ex ante and ex post prioritarian SWFs, the SVRR for the young person is not necessarily larger than for the old person. However, the ratio of the SVRR for the young person to the SVRR for the old person is always larger under either of the prioritarian SWFs than it is under the utilitarian SWF. In this sense, the two prioritarian SWFs exhibit a property that Adler et al. (2021) call 'ratio priority to the young'.

Continuing with the two-period example, the ratio SVRR(young)/SVRR(old) for the utilitarian SWF equals $\frac{u_1+p_2u_2}{u_2} = \frac{u_1}{u_2} + p_2$. For the ex ante prioritarian SWF, the ratio SVRR(young)/SVRR(old) equals $\frac{g'(s_1u_1+p_1p_2u_2)(u_1+p_2u_2)}{g'(u_1+p_2u_2)u_2} = \frac{g'(p_1u_1+p_1p_2u_2)}{g'(u_1+p_2u_2)}\left(\frac{u_1}{u_2} + p_2\right)$. The last term (in parentheses) is the ratio for the utilitarian SWF. It is multiplied by a term that depends on the slope of the transformation function *g* evaluated at the expected lifetime wellbeing of the young person (in the numerator) and at the expected lifetime wellbeing of the old person (in the denominator). Because the expected lifetime wellbeing of the young person is less than that of the old person, the numerator is larger than the denominator and so the ratio SVRR(young)/SVRR(old) is larger than under the utilitarian SWF. For the ex post prioritarian SWF, the ratio SVRR(young)/SVRR(old) equals $\frac{g(u_1)-g(0)+p_2[g(u_1+u_2)-g(u_1)]}{g(u_1+u_2)-g(u_1)} = \frac{g(u_1)-g(0)}{g(u_1+u_2)-g(u_1)} + p_2$. This ratio is larger than the ratio for the utilitarian SWF because the concavity of *g* implies that the first term on the right-hand side is larger than $\frac{u_1}{u_2}$.

These results can be extended to a model with more than two periods. For example, let each period be one year. Consider two individuals whose age-specific mortality risks and consumption levels are identical; the only difference between them is their year of birth. Let A_0 and A_1 denote the ages of the younger and older individuals, respectively. Consider the utilitarian SWF. Reducing current-period mortality risk increases each individual's chances of being alive in all future years and hence gain utility from consumption in those years. The duration effect is that, although both individuals increase their chances of living at ages greater than A_1, the younger individual also stands to gain the utility of living at ages between A_0 and A_1; the older individual has already lived these years. The risk effect is that the chance of experiencing the utility of living at each age greater than A_1 is smaller for the younger individual, who may die before reaching age A_1. If the utility of living each year is constant or decreasing with age (e.g., because income is constant, consumption is perfectly smoothed, or utility decreases as health declines with age) and the probability of surviving a period does not increase with age, then the duration effect exceeds the risk effect and the SVRR is larger for the younger individual. But if the utility of living per year or the probability of surviving a year can increase with age, then the risk effect can outweigh the duration effect and the utilitarian SWF can give priority to reducing risk to the older individual.

Under the ex ante prioritarian SWF, the difference in social welfare associated with reducing current risk to the younger rather than the older individual includes the same duration and risk terms as the utilitarian SWF. In addition, there is a third term which has the same sign as $g'(V_0) - g'(V_1)$, where V_0 and V_1 are the expected lifetime utilities of the younger and older individuals, respectively (Adler et al. 2021). The older individual's expected lifetime utility is larger than that of the younger individual's ($V_0 < V_1$) because the older individual has already survived the risk of dying between ages A_0 and A_1. Hence the difference $g'(V_0) - g'(V_1)$ is positive. This term is denoted the 'priority to the young' term and it ensures that the ex ante prioritarian SWF gives greater priority to the young than does the utilitarian SWF. Indeed, Adler et al. (2021) prove a stronger result: that the ratio of the SVRR for the younger individual to the SVRR for the older individual is strictly larger under the ex ante prioritarian SWF than under the utilitarian SWF.

The ex post prioritarian SWF also assigns greater priority to reducing current risk to the younger than to the older individual. Although the formula differs from that for the ex ante prioritarian SWF, the same three effects are present: duration, risk, and priority to the young. The stronger result also holds, i.e., that the ratio of the SVRR for the younger individual to the SVRR for the older individual is strictly larger under the ex post prioritarian SWF than under the utilitarian SWF.

While both ex ante and ex post prioritarian SWFs give priority to the young compared with the utilitarian SWF, it is not true that one of these always gives more priority to the young than the other. To demonstrate, return to the two-period example and let $u_1 = u_2 = 1$, $p_2 \approx 1$, and $g(u) = \sqrt{u}$. The ratio SVRR(young)/SVRR(old) equals approximately $2/\sqrt{p_1}$ under the ex ante prioritarian SWF and approximately $\sqrt{2}/(\sqrt{2}-1)$ under the ex post prioritarian SWF. Both ratios are of course larger than the corresponding ratio under the utilitarian SWF, which equals approximately 2. The ratio for the ex ante prioritarian SWF depends on the value of p_1 but the ratio for the ex post prioritarian SWF does not. For values of p_1 larger than $(2-\sqrt{2})^2 \approx 0.34$ the ratio for the ex ante prioritarian SWF is smaller than the ratio for the ex post prioritarian SWF; for values of p_1 smaller than this threshold the situation is reversed.

7.4.2. Effects of income and baseline risk

When analyzing the effects on the SVRR of differences in income or in risk, one must distinguish between differences that last for only a single period, for multiple periods, or throughout the lifetime. Moreover, the effects of a change for a single period can depend on whether that period is in the past, the present, or the future.

The effects of single-period and all-period differences in income and survival probability on the SVRR are summarized in Table 7.2, for the case in which the individual can neither borrow nor save, so consumption equals income in each period. Note that the SVRR under BCA and the utilitarian SWF are unaffected by differences in past-period income but SVRR under the prioritarian SWFs is affected. Prioritarian SWFs are history-dependent; they give priority to individuals with lower lifetime utility (expected utility in the case of the ex ante and possible realized utility in the case of the ex post prioritarian SWF). Greater income in a previous period increases lifetime utility and hence decreases the priority associated with improving an individual's wellbeing under both prioritarian SWFs. Because it has no effect on the gain from surviving the current period, greater income in a previous period decreases the SVRR.

Greater income in the current or a future period increases SVRR under BCA, the utilitarian, and the ex post prioritarian SWF. In all three cases, greater income increases wellbeing conditional on survival.[11] In contrast, the effect is ambiguous under the ex ante prioritarian SWF; higher income increases the utility gain from survival but decreases the priority given to the individual. If the transformation function *g* is not too concave, the ex ante prioritarian is similar to the utilitarian SWF and higher current or future income increases the SVRR; if the transformation function is sufficiently concave, the decrease in priority more than offsets the increase in wellbeing. Combining the effects of changes in past, present, or future income implies that the effect of an increase in income in all periods increases the SVRR under BCA and the utilitarian SWF but has an ambiguous effect under the two prioritarian SWFs. For the ex ante case, higher expected lifetime utility reduces social priority, and this effect can be large enough to offset the gain in expected lifetime utility if the transformation function is sufficiently concave. For the ex post case, the ambiguity arises because an increase in current or future income increases SVRR but an increase in past income decreases SVRR; the net effect depends on the relative magnitudes of these effects.

[11] Recall an individual gains utility from the current period only if she survives it.

The effects of a change in survival probability in a past period have no effect on the SVRR under BCA or any of the benchmark SWFs. This is because the risk of dying in these periods has already been resolved, so the values of the probabilities are irrelevant to lifetime wellbeing.

An increase in current survival probability has no effect on SVRR under the utilitarian and ex post prioritarian SWFs because it has no effect on lifetime wellbeing conditional on surviving the current period. In contrast, it decreases the SVRR under BCA and the ex ante prioritarian SWF. For BCA, an increase in current survival probability has no effect on the utility gain from survival (the numerator of equation (7.2)) but increases the expected opportunity cost of reduced consumption (the denominator of equation (7.2), the dead-anyway effect). Under the ex ante prioritarian SWF, the change in survival probability has no effect on the expected gain in utility from surviving the current period but increases expected lifetime utility and hence decreases social priority.

The effect of an increase in a future-period survival probability is to increase SVRR under BCA[12], the utilitarian and the ex post prioritarian SWF; in all three cases, the change increases expected lifetime wellbeing conditional on surviving the current period for BCA and the utilitarian SWF, and expected transformed lifetime wellbeing for the ex post prioritarian SWF. In contrast, the effect is ambiguous under the ex ante prioritarian SWF because the increase in expected lifetime wellbeing can be more than offset by the associated decrease in priority if the transformation function *g* is sufficiently concave.

Because the effect of an increase in current or future survival probabilities is to increase SVRR under the utilitarian and ex post prioritarian SWFs, so is the effect of a permanent increase in survival probabilities. In contrast, because the effects of an increase in current or future survival probabilities differ from each other under BCA and can differ under the ex ante prioritarian SWF, the effects of a permanent change are ambiguous and depend on which effect is larger.

In an extension to this model, Adler et al. (2019) analyze the case in which the individual can save and borrow against future income, which allows her to reallocate consumption over time. Specifically they assume the individual can sign a contract at birth pledging her income in each year she survives to an insurer that will pay her a defined annuity (that she will spend on consumption) each year she survives. The expected present value of her planned consumption path is equal to the expected present value of her (exogenous) income path; these depend on her time path of mortality risk and

[12] As noted above, in the single-period model an increase in life expectancy conditional on survival can decrease VSL if it increases the opportunity cost of spending. In the model of Adler et al. (2021), the individual receives income each period she survives and sets consumption equal to income, so an increase in life expectancy has no effect on the opportunity cost of spending.

the interest rate. The individual's optimal consumption path increases, decreases, or does not change with age if she discounts future utility at a rate that is smaller than, larger than, or equal to the interest rate, respectively.

The effect of being able to reallocate consumption over time on the SVRR depends on whether the policy that decreases mortality risk is anticipated (perhaps with some probability) or not. An anticipated change in mortality risk could affect the terms on which the individual can exchange her income stream for an annuity stream. Adler et al. (2019) analyze the case in which the policy that reduces risk is unanticipated; it comes as a surprise and does not affect the terms on which the individual can purchase her annuity or her planned consumption path. Hence the findings that the ex ante and ex post prioritarian SWFs exhibit priority to the young (compared with the utilitarian SWF), remain true in this extension to the model. The effect of a change in income on the SVRR does not depend on whether the income change occurs for a period in the past, present, or future; in each case, a higher income in one period increases the expected present value of lifetime income. This implies the effect of a change in income in any period is qualitatively the same as a permanent increase in income under the base-case model; it increases SVRR under BCA and the utilitarian SWF and has an ambiguous effect on SVRR under the two prioritarian SWFs (see Table 7.2). In contrast, the effect on the SVRR of a change in risk for a single period or for all periods differs from the base-case; it is ambiguous under all of the evaluation frameworks. The intuition is that two individuals of the same age having identical lifetime income paths but different survival-probability paths will have different paths of planned consumption.

This section reveals a diversity of qualitative effects of differences in income and baseline risk on SVRR across the different SWFs but gives little indication of their magnitudes. In the following section we report simulation exercises to help evaluate how large the effects of different evaluation methods may be on the SVRR of reducing mortality risk to individuals of different age and income.

7.5. Simulations

Adler et al. (2021) calculated the SVRR under BCA and each of the benchmark SWFs for a simulated population calibrated to the U.S. population. At each age, the distribution of income is based on U.S. Census data and the survival probability is based on the CDC lifetable, adjusted upward for higher and downward for lower incomes.[13] It is assumed there is no income mobility and an individual's income path over time corresponds to the median income of a fixed quintile of the income

[13] Annual mortality risk in each income quintile (lowest to highest) equals population risk multiplied by the following factors, derived from Chetty et al. (2016): 1.5, 1.2, 1, 0.9, 0.7.

distribution at each age (i.e., the 10th, 30th, 50th, 70th, or 90th percentile).[14] The simulation is conducted for ages 20 through 100. Consumption is assumed to be equal to income in each year. Results for ages 20-24 and 75 and older are in part artifacts of setting income at these ages equal to their values at ages 25 and 74, respectively, due to limitations in how the income data are reported.

The within-period individual utility function is ln(*c*) – ln($1,000), where *c* is annual income or consumption in dollars and the threshold level of $1,000 is based on the World Bank's estimate of subsistence income; the utility discount rate is zero. For the prioritarian SWFs, the transformation function $g(\cdot)$ is the Atkinson (isoelastic) function $g(u) = \frac{u^{1-\gamma}}{1-\gamma}$ with two values of the inequality-aversion parameter, $\gamma = 1, 2$.[15]

Figure 7.2 presents the SVRR by age and income quintile for BCA and the benchmark SWFs. For BCA, the SVRR is VSL. For this simulation, it tends to decrease with age while also exhibiting an inverted-U pattern, which is more pronounced for the higher-income quintiles. Because consumption is set equal to income, the inverted-U reflects the pattern of income with age rather than any saving or borrowing. The variation with income (and associated mortality risk) is large; as shown by Figure 7.3, at age 40 the VSL for an individual in the highest-income quintile is about 30 times as large as for an individual in the lowest-income quintile.

The SVRR is monotonically decreasing as a function of age for all the SWFs. Under the utilitarian SWF, the SVRR is increasing in income but the effect is much smaller than for VSL; the ratio of the SVRR for the highest- to the lowest-income quintile is between 2.5 and 3 for ages between about 20 and 80 (see Figure 7.3). The increase with income reflects two factors: expected future utility is larger because it is increasing in consumption, and life expectancy conditional on surviving the period is larger because it is increasing in income. For the prioritarian SWFs with inequality-aversion coefficient $\gamma = 1$ (both ex ante and ex post), the effect of income is very small, though the SVRR is increasing with income. In contrast, with the larger coefficient of inequality aversion $\gamma = 2$, the ordering by income is reversed with SVRR larger for lower- than for higher-income individuals. Individuals in lower-income quintiles receive greater priority because their consumption is lower and their expected lifetime is shorter.

---

[14] Income data are from the U.S. Census Bureau, Current Population Survey, data on individual money income in 2016 (https://www.census.gov/data/tables/time-series/demo/income-poverty/cps-pinc/pinc-01.html). This table reports income distribution by 5-year age increments from ages 25 to 74 but pools ages 15 to 24 years and 75 years and older. Survival data are from the United States Life Tables for 2014 (https://www.cdc.gov/nchs/products/life_tables.htm).

[15] When $\gamma = 1$, $g(u) = \ln(u)$.

The effects of income and the associated difference in mortality risk on the SVRR are illustrated more clearly in Figure 7.3, which shows the ratio of SVRR for the highest-income quintile to SVRR for the lowest-income quintile as a function of age. The left panel includes the ratios for VSL and for income in addition to those for the SVRR under the benchmark SWFs; the right panel excludes VSL and income so that the differences among the SWFs are more clearly seen.

For VSL, the ratio of SVRR for the highest to SVRR for the lowest-income quintile is greater than 20 at most ages. Its variation over time mirrors and exaggerates that of the ratio of income between the highest- and lowest-income quintiles, which increases from a value of about 8 to a peak of about 13 near age 65 then decreases. For the SWFs, the proportional effect of income rises very slightly with age (except for the sharp drop at ages greater than about 90 years); it is largest for the utilitarian SWF (a factor of about 3). The results under the prioritarian SWFs are more sensitive to the coefficient of inequality aversion than to the difference between the ex ante and ex post forms. With inequality aversion equal to 1, the ratio of SVRR for the highest to the lowest income quintile rises with age from about 1 to 1.5. With inequality aversion equal to 2, the SVRR is smaller for the highest than for the lowest-income quintile; the ratio rises from about 0.5 to 0.7 (which implies the effect of income decreases with age, contrary to the patterns for the other SWFs). Recall that income is assumed to be constant for ages younger than 25 and older than 74, so patterns at these younger and older ages are influenced by this perhaps unrealistic assumption.

The effects of age for median-income individuals are illustrated by Figure 7.4. Compared with the SVRR at age 60, the ex post prioritarian SWF produces the largest SVRR at young ages; the highest curve is with the inequality-aversion coefficient γ equal to 2 and the second highest is with the coefficient equal to 1. The ex ante prioritarian SWF with inequality-aversion coefficients of 1 and 2 and the utilitarian SWF produce very similar patterns of SVRR as a function of age that are nearly indistinguishable on this plot. As noted earlier, the evaluation under BCA is the outlier, with the SVRR increasing with age for younger ages then decreasing at older ages.

Adler (2019, chapter 5) presents a similar simulation in which he evaluates policies that reduce current mortality risk but at a cost that decreases individuals' current consumption. The population is divided into quintiles by income, with survival curves similar to those in the simulation described above. Wellbeing is measured using the same logarithmic utility function. In contrast, income is assumed to be constant over an individual's lifetime. Using this framework, Adler evaluates policies that reduce current mortality risk (for one year) by 1 per 100,000 on average. The risk reduction may be distributed uniformly across individuals aged 20, 30, 40, 50, and 60 years (present in equal numbers) or targeted to either the youngest quintile (20 year olds) or the poorest quintile. A targeted policy reduces mortality risk by 5 per 100,000 for individuals in the targeted quintile and zero for all others. The one-year cost is distributed either uniformly or in proportion to income. Adler calculates the maximum cost of each policy (average reduction in individual consumption) such that the policy is preferred to the status quo, using BCA, the utilitarian, and ex post prioritarian SWFs (using the Atkinson transformation function with inequality-aversion coefficient $\gamma = 2$). A policy increases social welfare if its average cost is less than the breakeven cost, i.e., more desirable policies have larger breakeven costs.

Results are summarized in Table 7.3. Under BCA, the breakeven cost is unaffected by its allocation. It is smaller when the risk reduction is targeted to the poorest quintile (whose VSL is smaller). The breakeven cost is larger when the risk reduction is targeted to the youngest quintile, who have a larger VSL than do older cohorts. (In the simulation, VSL decreases with age because consumption is assumed to be constant over time.)

Under the utilitarian and ex post prioritarian SWFs, the breakeven cost is always larger when costs are allocated in proportion to income rather than uniformly, because the utility loss of a reduction in consumption is smaller at higher incomes. The proportional difference is larger for the ex post prioritarian than for the utilitarian SWF, because the former exhibits aversion to inequality in wellbeing. Under both SWFs, the policy that targets risk reduction to the youngest quintile has a higher breakeven cost than does the uniform risk reduction. In contrast, the policy that targets risk reduction to the poorest quintile has a smaller breakeven cost than the uniform risk reduction under the utilitarian SWF, but a higher breakeven cost under the ex post prioritarian SWF. The poorest quintile experiences a smaller than average gain in wellbeing from risk reduction, but the ex post prioritarian SWF assigns sufficient priority to this quintile to offset the difference in wellbeing compared with a uniform risk reduction.

These simulations show that the use of alternative evaluation methods, whether BCA or a specific SWF, can have a dramatic effect on the relative social value of reducing risk to individuals with different characteristics. The SVRR is increasing with income (and associated life expectancy) under BCA and the utilitarian SWF; the effect is much larger under BCA. In contrast, SVRR can increase or decrease with income under the ex ante and ex post prioritarian SWFs, depending on the degree of inequality aversion. All of the benchmark SWFs produce a larger SVRR for younger than older individuals in the simulations; this contrasts with BCA, which yields an SVRR that rises over younger ages then falls over older ages for some income quintiles. The ex post prioritarian SWF gives more priority to the young than does the ex ante prioritarian SWF in the simulations, though as described in Section 7.4 this is not always true theoretically.

7.6. Extensions

In this section we consider two extensions. The first is the use of SWFs that are not additively separable across individuals, which can exhibit catastrophe aversion. The second is a more general model that includes health and nonfatal health risks.

7.6.1. Catastrophe aversion and nonseparable SWFs

Hazards that lead to mass casualties, such as hurricanes, earthquakes, floods, volcanic eruptions, and other natural disasters, airliner crashes, and terrorist attacks, are often perceived to be of greater social concern than hazards that produce an equal or even greater number of deaths that occur in separate incidents, such as traffic crashes and ambient air pollution. Arguably, resistance to nuclear power plants and acceptance of coal-fired alternatives is rationalized by the perceived chance of a catastrophic loss of life in a nuclear accident as opposed to the routine mortality toll from fine-particulate air pollution.

Catastrophe aversion is defined as the preference for a higher probability of a smaller number of deaths over a lower probability of a larger number of deaths, holding the expected number of deaths constant. Keeney (1980) showed that, when risks are independent, catastrophe aversion conflicts with a preference for risk equity, defined as greater equality among individuals' risks. To illustrate, consider a population of N people, each facing independent mortality probability p_i. If the individual risks are all equal to $1/N$, the expected number of deaths is one. Because the individuals' risks are equal the risk distribution is perfectly equitable. Alternatively, if individuals 1 and 2 each face a risk of 1/2 and all the others face zero risk, the expected number of deaths is also 1. This distribution is less equitable, because two people face higher risk than everyone else. The actual number of deaths can be very different: in the first case, it can be any number between 0 and N; in the second case, any number between 0 and 2. Catastrophe aversion implies a preference for the second risk over the first, which conflicts with a preference for risk equity.

Adler et al. (2014) and Rheinberger and Treich (2017) generalized Keeney's definition of catastrophe aversion; they define it as a preference for any mean-preserving contraction[16] in the probability distribution of the number of fatalities. When risks are independent, a pure transfer of risk from one individual to another that decreases the difference between them increases risk equity and yields a mean-preserving spread in the distribution of fatalities, i.e., it makes the risk more catastrophic. Hence the conflict between catastrophe aversion and a preference for risk equity also exists under

[16] A mean-preserving contraction shifts probability away from the tails of a distribution without changing the mean. A mean-preserving spread shifts probability toward the tails.

this more general definition of catastrophic risk. When risks are dependent (before or after the risk transfer) and there are more than two people in the population, the relationship between risk equity and catastrophic risks is more complicated because a risk transfer can affect the correlation of risks within the population (Bernard et al. 2017).

Rheinberger and Treich (2017) reviewed the empirical literature and found that catastrophe aversion is supported less often than not, perhaps because of a greater concern for risk equity. Another rationale for preferring a more-catastrophic risk is common-fate preference. As Schelling (1968) observed, a family may prefer to travel together on a single airliner rather than traveling on separate flights; traveling separately reduces the chance of a catastrophe in which the entire family dies but increases the chance of losing at least one family member, leaving the survivors to grieve.

Of the evaluation frameworks evaluated to this point, BCA and the ex ante prioritarian SWF prefer policies in which risk is distributed more equally, and the utilitarian and ex post prioritarian SWFs are insensitive to how individual risks are distributed. This implies that, when risks are independent, BCA and the ex ante prioritarian SWF exhibit catastrophe acceptance and the utilitarian and ex post prioritarian SWFs exhibit catastrophe neutrality. In contrast, when risks are dependent the link between a preference for risk equity and catastrophe acceptance need not hold. For example, all four evaluation frameworks (BCA, the utilitarian, ex ante and ex post prioritarian SWFs) are indifferent between the case in which there is a $1/N$ chance that all N (identical) members of a population will die (with a complementary chance that all survive) and the case in which one of the N members to be selected at random will die for sure. In both cases, every individual faces a $1/N$ chance of death.

The utilitarian, ex ante and ex post prioritarian SWFs are each defined as sums over individuals; this implies the evaluation is unaffected by any correlation or dependence between individuals' outcomes. That is, the effect on social welfare of whether an individual survives or dies is the same regardless of whether other individuals survive or die. Other SWFs can account for such dependence. For example, the ex post transformed utilitarian and ex post transformed prioritarian SWFs are sensitive to dependence among individuals' outcomes.

These SWFs are obtained by introducing a strictly increasing function $h(\cdot)$ that transforms each possible social outcome and calculating the expected value of the transformed outcomes. For the single-period model defined in Section 7.3, the ex post transformed utilitarian SWF is given by

$$S^{EPTU} = \sum_s \pi_s h[\sum_i \{l_i^s u(c_i) + (1 - l_i^s) v(c_i)\}] \quad (7.6)$$

and the ex post transformed prioritarian SWF is given by

$$S^{EPTP} = \sum_s \pi_s h[\sum_i g\{l_i^s u(c_i) + (1 - l_i^s) v(c_i)\}]. \quad (7.7)$$

As shown by Adler et al. (2014), these SWFs exhibit catastrophe aversion if and only if the transformation function *h* is concave. The intuition is that introducing a concave *h* is like introducing risk aversion over the evaluation of the possible outcomes. This implies the ex post transformed utilitarian and prioritarian SWFs prefer policies that are associated with a mean-contracting spread of the evaluation of the outcomes, e.g., a less-catastrophic risk. The utilitarian SWF (equation (7.3)) calculates the expected value of the sum of individuals' utilities in each state. In contrast, the ex post transformed utilitarian SWF (equation (7.6)) calculates the expected value of this sum transformed by the function *h*. Analogously, the ex post prioritarian SWF (equation (7.4)) calculates the expected value of the sum of individuals' utilities transformed by the function *g*. The ex post transformed prioritarian SWF (equation (7.7)) calculates the expected value of this sum transformed by the function *h*.

An important special case of the transformation function *h* is Fleurbaey's (2010) equally distributed equivalent (EDE).[17] For any distribution of wellbeing in a population, the equally distributed equivalent is the level of individual wellbeing that, if held by everyone in a population, would yield the same social welfare. For the utilitarian SWF, the equally distributed equivalent is the average utility in the population and the transformation function h^{EDE} is the identity function. Hence the ex post transformed utilitarian SWF using h^{EDE} is catastrophe neutral. For the prioritarian SWF, the equally distributed equivalent is $g^{-1}\left[\frac{1}{N}\sum_1^N g(u_i)\right]$ where u_i is the wellbeing of individual *i*. Because *g* is concave, its inverse g^{-1} and hence h^{EDE} are convex. This implies that ex post transformed prioritarianism using h^{EDE} is catastrophe accepting. The intuition is that a more catastrophic situation is also more equitable ex post in the sense that more people live or die together, which is socially valued under this SWF.

7.6.2. Health

In the single-period model defined in Section 7.2, the effect of health is implicit; the utility of surviving the current period $u(c)$ depends on the individual's uncertain future health, longevity, income, and other factors. This model can be adapted to incorporate the effects of health and the possibility of suffering nonfatal illness or injury. In this section, we first examine the effect of future health on the social value of reducing current-period mortality risk.[18] Second, we generalize the

---

[17] EDE is discussed by Adler (chapter 2, this volume). Ferranna and Fleurbaey (chapter 8) compare ex ante, ex post, and expected EDE prioritarianism in the context of climate policy.

[18] Here we depart from the assumption that the utility function $u(c)$ is identical across individuals.

single-period model to include the possibility an individual may suffer an illness or injury that diminishes her health and may prove fatal. In this richer model, the set of possible actions is broader than for fatal risks. In addition to reducing the probability of suffering the condition, one can potentially reduce its severity or its lethality. We examine the social value of each of these actions under BCA and the benchmark SWFs, and how they depend on the baseline risk of illness, its severity, and the probability it is fatal.

In the single-period model introduced in Section 7.2, utility conditional on surviving the current period $u(c)$ can be interpreted as the certainty-equivalent utility conditional on survival, which depends on the individual's lottery over future health prospects. Define 'better future health' as any change in this lottery that increases its certainty equivalent $u(c)$ and hence increases the utility of survival. Better health increases SVRR under the utilitarian and ex post prioritarian SWFs but it can increase or decrease SVRR under the ex ante prioritarian SWF and BCA, as summarized in the last column of Table 7.1. Moreover, comparing two individuals who are identical except that one has better future health than the other, the ex post and ex ante prioritarian SWFs give relatively greater priority to reducing mortality risk for the individual with worse future health than does the utilitarian SWF. This result parallels the finding that the two prioritarian SWFs give priority to younger individuals, discussed in Section 7.4.1.

Recall the SVRR under the utilitarian SWF is $u(c_i) - v(c_i)$ and the SVRR under the ex post prioritarian SWF is $g[u(c_i)] - g[v(c_i)]$. In both cases, increasing utility conditional on survival $u(c_i)$ while holding constant the bequest $v(c_i)$ increases the SVRR. For the ex ante prioritarian SWF, the SVRR is $g'\left[p_i u(c_i) + (1-p_i)v(c_i)\right]\left[u(c_i) - v(c_i)\right]$. The second term (in brackets) is the SVRR under the utilitarian SWF (which is increasing in $u(c_i)$) but the first term is decreasing in $u(c_i)$. If the transformation function g is sufficiently concave, the decrease in the first term can more than offset the increase in the second term so that the SVRR decreases with future health. For BCA, if (as is plausible) better health increases the marginal utility of wealth conditional on survival, the corresponding increase in the denominator of equation (7.2) can more than offset the increase in utility (the numerator), decreasing VSL.

Under both the ex post and ex ante prioritarian SWFs, the ratio of the SVRR for someone with worse future health to the SVRR for an otherwise identical individual with better future health is larger than the corresponding ratio under the utilitarian SWF. That is, although the utilitarian and ex post prioritarian SWFs give priority to reducing risk to an individual with better future health (all else equal), the ex post prioritarian SWF gives less priority to such an individual. The ex ante prioritarian

SWF also gives less weight to such an individual, and can even favor reducing mortality risk to the individual with worse future health.

As shown in Table 7.1, the effects of health on SVRR differ from the effects of both baseline risk and income. Neither VSL nor the SVRR under any of the benchmark SWFs can be equal across individuals differing in baseline risk, wealth, and health.

To capture the possibility of illness or injury that may prove fatal, consider the single-period model developed by Rheinberger et al. (2016). An individual faces a probability p of illness. Conditional on illness, mortality risk is q.[19] If ill, health deteriorates by a fraction h defined such that $h = 0$ corresponds to full health and $h = 1$ to a health state as bad as dead (h is limited to the interval between 0 and 1). Utility depends on wealth c; the utility of a bequest is $v(c)$ and the utility of survival equals $(1 – h)\, u(c) + v(c)$.[20] As in the standard model for mortality risk (equation (7.1)), we assume survival is preferred to death ($u(c) > 0$), the marginal utility of wealth conditional on survival is greater than the marginal utility of a bequest ($u'(c) > 0$), the marginal utility of a bequest is non-negative ($v'(c) \geq 0$), and the individual is averse to financial risk conditional on survival and on death ($u''(c) \leq 0$ and $v''(c) \leq 0$).

Individual expected utility

$$U = (1-p)[u(c)+v(c)] + p(1-q)[(1-h)u(c)+v(c)] + pqv(c)$$
$$= [1-p(q+h-qh)]u(c)+v(c). \quad (7.8)$$

One implication of this specification is that the marginal utility of wealth decreases as the health deterioration increases; this is often but not always assumed. This model is most naturally interpreted as a model of chronic illness because the individual cannot recover. It provides a model for acute illness if the utility if ill is interpreted as expected future lifetime utility conditional on becoming ill in the current period and recovering. Nonfatal health conditions can be represented by setting lethality $q = 0$. Note that the roles of lethality q and health deterioration h are symmetric; they appear in the last expression of equation (7.8) only in the term $(q + h – q\, h)$. This has implications that are addressed below.

---

[19] Rheinerger et al. (2016) denote the probability of illness by q and the conditional mortality risk by p.
[20] Note that $u(c)$ is the difference in utility between survival in full health and death. This differs from the model for mortality risk (equation (7.1)), in which $u(c)$ is the utility of wealth conditional on survival. This specification is convenient for representing ill health as a fractional deterioration toward a health state as bad as dead.

Decreases in any of the three parameters (risk of illness *p*, severity *h*, and lethality *q*) increase expected utility. We describe the marginal value of decreases in each of these parameters (and of an increase in wealth *c*) under BCA and under the three benchmark SWFs. We then describe how these marginal values depend on the values of the other parameters (e.g., how the value of reducing the risk of an illness depends on its severity and lethality). Finally, we compare the relative social value of decreases in each parameter. Derivations are reported by Hammitt (2020).

Under BCA, the social value of a marginal improvement in any dimension of the health risk equals the individual's WTP for that improvement. By analogy to VSL, this value can be normalized to a one-unit change and reported as a value per statistical case. Rheinberger et al. (2016) derive the marginal WTP for prevention (reducing *p*), treatment (reducing *q*), and palliative care (reducing *h*). They show that marginal WTP for each of these improvements is positive and evaluate its dependence on the values of the parameter itself and of the other parameters. These results are summarized in Table 7.4. WTP for an improvement in each dimension is increasing in wealth. Higher baseline risk (larger *p*) increases WTP for prevention (that reduces *p*), treatment (that reduces *q*), and care (that reduces *h*). Greater lethality (larger *q*) increases WTP for prevention (that reduces *p*) and treatment (that reduces *q*) but has an ambiguous effect on WTP for care (that reduces *h*). Greater severity (larger *h*) increases WTP for prevention (that reduces *p*) and care (that reduces *h*) but has an ambiguous effect on WTP for treatment (that reduces *q*).

The social value of a decrease in baseline risk, lethality, or severity for an individual is defined as the rate of increase of social welfare resulting from a decrease in the corresponding parameter, i.e., the negative of the partial derivative of social welfare with respect to that parameter. Similarly, the social value of an increase in an individual's wealth is the partial derivative of social welfare with respect to the individual's wealth. These definitions are analogous to the definition of the social value of risk reduction in the case of mortality.

Under the utilitarian SWF, social welfare is the sum of each individual's expected utility (equation (7.8)). The contribution to social welfare of a single individual[21]

$$V^U = U = [1 - p(q + h - qh)]u(c) + v(c). \tag{7.9}$$

The social value of a reduction in parameter θ is $-\frac{\partial V^U}{d\theta} = -\frac{\partial U}{d\theta}$, where θ = *p*, *q*, or *h*. A decrease in baseline risk, lethality, or severity increases individual expected utility, and so the social value of a

[21] We omit individual subscripts for simplicity.

decrease in each of these dimensions is positive. The sensitivity of the social values to the parameter values are summarized in Table 7.4.

The social value of an improvement (reduction) in baseline risk *p*, lethality *q*, and severity *h* is independent of the value of that parameter and is increasing in wealth. An increase in baseline risk *p* increases the social value of both treatment (reducing lethality *q*) and care (reducing severity *h*). The intuition is that, the more likely the individual is to become sick, the more likely it is that she will benefit from a reduction in lethality or severity. An increase in either lethality or severity decreases the social value of reducing the other, because it makes it less likely the individual will benefit from that reduction.

Under the ex ante prioritarian SWF, social welfare is the sum of each individual's expected utility *U* transformed by an increasing, concave function *g*. The contribution to social welfare of a single individual

$$V^{EAP} = g[U] = g\{[1 - p(q + h - qh)]u(c) + v(c)\}. \qquad (7.10)$$

The social value of a reduction in parameter θ is given by $-\frac{\partial V^{EAP}}{d\theta} = -g'(U)\frac{\partial U}{d\theta}$, i.e., the social value of reducing that parameter under the utilitarian SWF multiplied by a priority factor that is equal to the slope of the transformation function *g* evaluated at the individual's expected utility. Under this SWF, many of the effects are ambiguous because a change in a parameter can increase the social value of improving the lottery under the utilitarian SWF and at the same time increase the individual's expected utility, decreasing priority to the individual. The effects of an increase in wealth on the social values of prevention, treatment, and care are ambiguous for this reason. The sign of the effect depends on the concavity of the function *g*: if the function is nearly linear, the effects go in the same direction as those under the utilitarian SWF; if it is sharply concave, the directions of the effects can be reversed.

One effect that is unambiguous under the ex ante prioritarian SWF is that an increase in baseline risk *p* increases the social value of all three actions: prevention, treatment, and care. Although the social value of prevention under the utilitarian SWF is independent of baseline risk, an increase in risk decreases expected utility and hence increases the priority given to the individual under the ex ante prioritarian SWF. Because of the symmetry of lethality and health in the model, their effects are symmetric: an increase in either parameter increases the social value of reducing that parameter but has an ambiguous effect on the social value of reducing the other parameter. The difference from the results under the utilitarian SWF is that an increase in lethality or severity decreases expected utility and increases priority to the individual. If the function g is sufficiently concave, the increase in social priority can more than offset the decreased gain in expected utility.

Under the ex post prioritarian SWF, social welfare is the sum of each individual's expected transformed utility. The contribution to social welfare of a single individual

$$V^{EPP} = (1-p)g[u(c)+v(c)] + p(1-q)g[(1-h)u(c)+v(c)] + pqg[v(c)] \qquad (7.11)$$

and the social value of a reduction in parameter θ is given by $-\frac{\partial V^{EPP}}{d\theta}$. Results are almost identical to those under the utilitarian SWF, with the major exception that the effect of an increase in wealth on the social value of prevention, treatment, and care is ambiguous rather than positive; the social values of these changes can decrease if the function g is sufficiently concave. The only other difference from the results under the utilitarian SWF is that an increase in severity h increases the social value of care (that reduces h) while it has no effect under the utilitarian SWF; in contrast, an increase in lethality has no effect on the social value of treatment (that reduces q) under both the ex post prioritarian and utilitarian SWFs. This result breaks the symmetry of the effects of severity h and lethality q seen under the utilitarian and ex ante prioritarian SWFs. The reason is that the two parameters do not enter the ex post prioritarian SWF (equation (7.11)) in a symmetric fashion. Severity h is a parameter of utility conditional on illness, so it enters as part of the argument of the transformation function g, specifically as $g[(1-h)\, u(c) + v(c)]$. In contrast, lethality q is a probability and so it never appears in the argument of g.

Comparison of the relative social values of reducing each of the three dimensions of health risk reveals that results are identical for BCA, the utilitarian and the ex ante prioritarian SWF. This follows because the social value of reducing each parameter under BCA equals the social value under the utilitarian SWF divided by the expected marginal utility of wealth, and the social value under the ex ante prioritarian SWF equals the social value under the utilitarian SWF multiplied by the slope of the transformation function g evaluated at the individual's expected utility. Hence if the social value of reducing one parameter is larger than the social value of reducing another under the utilitarian SWF, the same is true under BCA and the ex ante prioritarian SWF.

Because of the symmetry of the roles of lethality q and health deterioration h in the model (except under the ex post prioritarian SWF), the relative social values of treatment (to reduce q) and care (to reduce h) is determined by their relative magnitudes. Under BCA, the utilitarian SWF and the ex ante prioritarian SWF, the marginal social value of decreasing one of these parameters is always greater for the larger of the two parameters. That is, treatment (to reduce q) has a higher social value than care (to reduce h) if and only if $q > h$.

The social value of prevention (to reduce *p*) is larger than the social value of treatment (to reduce *q*) if and only if $\frac{h}{1-h} > p - q$. This condition is satisfied for diseases that are uncommon and lethal (for which $p < q$). For diseases that are rarely fatal ($q \approx 0$), the condition is satisfied only if health deterioration is large compared with baseline risk; for common, mild, nonfatal illnesses, the social value of treatment can exceed that of prevention. The social value of prevention is larger than that of care (to reduce *h*) if and only if $\frac{q}{1-q} > p - h$. This condition is satisfied for diseases that are uncommon and severe (for which $p < h$), and for highly lethal diseases (with $q \approx 1$). For common, mild, nonfatal illnesses, the social value of care can exceed that of prevention.

Under the ex post prioritarian SWF, the social value of treatment (decreasing lethality *q*) exceeds the social value of palliative care (decreasing health deterioration *h*) if $q \geq h$. The social value of prevention (decreasing risk *p*) is greater than that of treatment (decreasing lethality *q*) if baseline risk *p* is near zero or health deterioration *h* is near one, i.e., for very rare or very severe disease. If lethality *q* and health deterioration *h* are both small, and risk *p* is not too small, then treatment has a larger social value than prevention. Prevention (to reduce risk) has a larger social value than care (to reduce severity) if *h* is small and *q* is larger than *p* or if *h* and *q* are large and *p* is small. Care has a larger social value than prevention if the disease is rarely fatal ($q \approx 0$) and if health deterioration is smaller than baseline risk.

In summary, in the case of fatal risk, better health conditional on survival increases SVRR under the utilitarian and ex post prioritarian SWFs but can increase or decrease SVRR under the ex ante prioritarian SWF and BCA. In the case of more general health risks, the social value of an improvement in any of the three dimensions of health risk (baseline risk, lethality and severity of condition) is positive under BCA and the three benchmark SWFs. The sensitivities of the social values of risk reduction to changes in parameters differ substantially across the evaluation methods. One general result is that the social values are all increasing in wealth under BCA and the utilitarian SWF, but can increase, decrease, or be unaffected by wealth under the ex ante and ex post prioritarian SWFs. Another result is that the social value of prevention is larger when the health condition is more lethal or more severe; it is also larger when the baseline risk is higher under BCA and the ex ante prioritarian SWF, but is unaffected by baseline risk under the utilitarian and ex post prioritarian SWFs.

7.7. Conclusion

Concern about how to value reductions in mortality risk to people of different ages, incomes, and other characteristics has been a perennial issue in evaluation of environmental, health, and safety policies. Typically, these policies have been evaluated using benefit-cost analysis and the value of reducing mortality risk is quantified as a value per statistical life (VSL). Analyses within a national population generally use a common VSL, independent of individual characteristics. In contrast, evaluation of medical and public-health interventions such as vaccination often use a constant value per unit increase in life expectancy, i.e., a constant value per life year or per life year adjusted for health (such as quality-adjusted or disability-adjusted life years).[22] Obviously the two approaches are mutually inconsistent as the gain in life expectancy associated with a reduction in current-period mortality risk is proportional to life expectancy conditional on surviving the current period, which typically decreases with age.

The use of social-welfare functions may help to defuse this tension. An important reason for the use of a constant value, either per expected life saved or per expected life year gained, is a desire to treat individuals equitably, perhaps equally. Conventional BCA attempts to separate evaluation of efficiency and equity, where efficiency is defined by the Kaldor-Hicks compensation test: a policy is efficient if the individuals who benefit from the policy could hypothetically compensate the individuals harmed by the policy with money, so that everyone would prefer the policy change plus compensation to the status quo. Recognizing that this compensation is hypothetical and is typically not made, it is conventional to supplement BCA (or at least to recommend it be supplemented) with an analysis of the distributional effects. The notion is that decision makers can choose policies recognizing both efficiency and distributional effects and make appropriate tradeoffs between these dimensions.

Unlike conventional BCA, social-welfare functions can explicitly incorporate concerns for distribution together with concerns for efficiency. In principle, there is no need, and perhaps no advantage, to supplement an SWF analysis with a distributional analysis. Agreement on an appropriate SWF (including the functional form and parameter values), or at least agreement on a limited set of SWFs and parameter values, might reduce some of the tension about balancing effects on efficiency and equity. Use of an SWF requires choice of an interpersonally comparable measure of wellbeing, which is not required for BCA.

The simulations reported in Section 7.5 suggest that differences in the social value of reducing risk to individuals with different incomes (and associated baseline mortality risks) are much smaller under any of the benchmark social-welfare functions than they are under BCA. Plausible SWFs treat individuals with different incomes more similarly than does BCA and for this reason may be perceived

[22] These analyses are often framed as cost-effectiveness analyses, where the health effect is valued in terms of the expected increase in quality-adjusted life years (or some other measure of health and longevity).

as more equitable. In contrast, the effect of differences in age can be smaller or larger under the benchmark SWFs than under BCA. If differences in the social value of risk reduction are justified by a particular model of equity that is embedded in the chosen SWF, they may be more acceptable to decision makers and the public than the differences that arise under BCA.

To determine whether the use of SWFs in evaluation of fatality risk regulations will help to better inform policy choice, much work remains to be done. Incorporating distributional effects in policy analysis requires more information on how policy effects are distributed than is necessary for BCA. In particular, information about how costs (reductions in disposable income) are distributed in a population is not needed for BCA but can be critical when applying an SWF (or conducting a distributional analysis). Because costs often appear in the form of increases in taxes or product prices, or decreases in wages, their incidence depends on how people respond to these market changes and can be difficult to estimate. Moreover, under prioritarian SWFs, the evaluation depends on history; the priority given to different individuals depends on their lifetime wellbeing, which depends on their wellbeing in previous periods. Collecting and organizing the information needed to adequately represent the distribution of lifetime wellbeing and its relationship to the distribution of possible consequences is a challenging topic for future work.

Acknowledgments

The authors thank Matthew Adler, Ole Norheim, and Maddalena Ferranna for useful comments. This work was supported by the Chair AMEP-IDEX at the University of Toulouse Capitole. Nicolas Treich acknowledges financial support from the ANR under grant ANR-17-EURE-0010 (the Investissements d’Avenir program) and the TSE-P SCOR chair.

References

Adler, M.D. (2019). *Measuring Social Welfare: An Introduction*, Oxford: Oxford University Press.

Adler, M.D., Ferranna, M., Hammitt, J.K. & Treich, N. (2021). "Fair innings? The utilitarian and prioritarian value of risk reduction over a whole lifetime," *Journal of Health Economics* 75: 102412.

Adler, M.D., Ferranna, M., Hammitt, J.K. & Treich, N. (2019). "Fair innings? The utilitarian and prioritarian value of risk reduction over a whole lifetime," Duke Law School Public Law and Legal Theory Working Paper No. 2019-79. Available at https://ssrn.com/abstract=3493976.

Adler, M.D., Hammitt, J.K, & Treich, N. (2014). "The social value of mortality risk reduction: VSL versus the social welfare function approach," *Journal of Health Economics* 35: 82-93.

Baker, R., Chilton, S., Jones-Lee, M. & Metcalf, H. (2008). "Valuing lives equally: defensible premise or unwarranted promise?" *Journal of Risk and Uncertainty* 36: 125–138.

Bernard, C., Rheinberger, C.M. & Treich, N. (2018). "Catastrophe aversion and risk equity in an interdependent world," *Management Science* 64: 4490-4504.

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

Chetty, R., Stepner , M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A. & Cutler, D. (2016). "The association between income and life expectancy in the United States, 2001-2014," *JAMA* 315: 1750-1766.

Drèze J. (1962). "L'Utilité sociale d'une vie humaine," *Revue Française de Recherche Opérationelle*, 6: 93–118.

Fleurbaey, M. (2010). "Assessing risky social situations," *Journal of Political Economy* 118: 649-680.

Hammitt, J.K. (2020). "The marginal social value of reducing health risk," mimeo, November.

Harris, J. (1985). *The Value of Life*, London: Routledge and Kegan Paul.

Jones-Lee, M.W. (1974). "The value of changes in the probability of death or injury," *Journal of Political Economy* 82: 835–849.

Keeney, R.L. (1980). "Equity and public risk," *Operations Research* 28: 527-34.

Lee, J.M. & Taylor, L.O. (2019). "Randomized safety inspections and risk exposure on the job: quasi-experimental estimates of the value of a statistical life," *American Economic Journal: Economic Policy* 11: 350–374.

Levy, J.I. (2021). "Incorporating Inequality into Regulatory Impact Analysis: Barriers and Opportunities," *Risk Analysis*, in press.

Levy, J.I., Chemerynski, S.M. & Tuchmann, J.L. (2006). "Incorporating Concepts of Inequality and Inequity into Health Benefits Analysis," *International Journal for Equity in Health* 5: Article 2.

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.

OECD. (2018). *OECD Regulatory Policy Outlook 2018*, Paris: OECD Publishing.

Pratt, J.W. & Zeckhauser, R.J. (1996). "Willingness to pay and the distribution of risk and wealth," *Journal of Political Economy* 104: 747–763.

Rheinberger, C.M., Herrera-Araujo, D. & Hammitt, J.K. (2016). "The value of disease prevention vs. treatment," *Journal of Health Economics* 50: 247-255,.

Rheinberger, C.M. & Treich, N. (2017). “Attitudes toward catastrophe,” *Environmental and Resource Economics* 67: 609–636.

Robinson, L.A. (2007). “How US government agencies value mortality risk reductions,” *Review of Environmental Economics and Policy* 1: 283–299.

Robinson, L.A., Hammitt, J.K. & Zeckhauser, R.J. (2016). “Attention to distribution in U.S. regulatory analyses,” *Review of Environmental Economics and Policy* 10: 308-328.

Schelling T.C. (1968). The life you save may be your own. In S.B. Chase, Jr., ed., *Problems in Public Expenditure Analysis*. Washington, DC: Brookings Institution, pp. 127–162.

Shepard, D.S. & Zeckhauser, R.J. (1984). “Survival versus consumption,” *Management Science* 30: 423-439.

Sheriff, G. & Maguire, K.B. (2020). “Health Risk, Inequality Indexes, and Environmental Justice,” *Risk Analysis* 40: 2661-2674.

Simon, N. B., Dockins, C., Maguire, K.B., Newbold, S.C., Krupnick, A.J. & Taylor, L.O. (2019). “What’s in a name? a search for alternatives to ‘VSL’,” *Review of Environmental Economics and Policy* 13: 155–161.

Tunçel, T. & Hammitt, J.K. (2014). “A New Meta-Analysis on the WTP/WTA Disparity,” *Journal of Environmental Economics and Management* 68: 175-187.

U.S. Environmental Protection Agency. (2011). *The Benefits and Costs of the Clean Air Act from 1990 to 2020*, Final Report, Rev. A, Washington, D.C.

Viscusi, W.K. (2009). “The devaluation of life,” *Regulation and Governance* 3: 103–127.

Weinstein, M.C., Shepard, D.S. & Pliskin, J.S. (1980). “The economic value of changing mortality probabilities: a decision-theoretic approach,” *Quarterly Journal of Economics* 94: 373-396.

Williams, A. (1997). “Intergenerational equity: an exploration of the ‘fair innings’ argument,” *Health Economics* 6: 117–132.

Table 7.1. Effects of baseline risk, wealth, and health on SVRR in the single-period model

	Mortality risk	Wealth	Health
BCA	+	+	+/-
Utilitarian SWF	0	+	+
Ex ante prioritarian SWF	+	+/-	+/-
Ex post prioritarian SWF	0	+/-	+

Note: Symbols describe the sign of the effect of an increase in the column parameter on the social value of a marginal decrease in the row parameter. Symbols denote increase (+), decrease (-), no effect (0), ambiguous effect (+/-). An effect is ambiguous if it depends on parameter values; the effect of wealth on SVRR for the two prioritarian SWFs is negative if the transformation function is strongly concave and positive if it is close to linear.

Table 7.2. Effects of income and baseline risk on SVRR in the multi-period model

SVRR	**Income: Single-period difference**	**Income: permanent difference**	**Survival probability: single-period difference**	**Survival probability: permanent difference**
BCA	Past period: *Independent* Current period: *Increasing* Future period: *Increasing*	*Increasing*	Current period: *Decreasing* Future period: *Increasing*	*Ambiguous*
Utilitarian	Past period: *Independent* Current period: *Increasing* Future period: *Increasing*	*Increasing*	Current period: *Independent* Future period: *Increasing*	*Increasing*
Ex ante prioritarian	Past period: *Decreasing* Current period: *Ambiguous* Future period: *Ambiguous*	*Ambiguous*	Current period: *Decreasing* Future period: *Ambiguous*	*Ambiguous*
Ex post prioritarian	Past period: *Decreasing* Current period: *Increasing* Future period: *Increasing*	*Ambiguous*	Current period: *Independent* Future period: *Increasing*	*Increasing*

Source: Adler et al. (2021) Table 2.

Table 7.3. Breakeven average cost ($) for policy that reduces average annual mortality risk by 1/100,000

Policy	BCA	Utilitarian	Ex post prioritarian
Uniform risk reduction			
Uniform cost	91	48	78
Cost proportional to income	91	77	159
Risk reduction for youngest quintile			
Uniform cost	132	71	178
Cost proportional to income	132	108	360
Risk reduction for poorest quintile			
Uniform cost	15	32	98
Cost proportional to income	15	51	201

Source: Adler (2019) Table 5.12.

Table 7.4. Social value of reducing alternative dimensions of health risk

	Parameter decreased	Incidence (p)	Lethality (q)	Severity (h)	Wealth (c)
BCA	Prevention (p)	+	+	+	+
	Lethality (q)	+	+	+/-	+
	Severity (h)	+	+/-	+	+
Utilitarian SWF	Prevention (p)	0	+	+	+
	Lethality (q)	+	0	-	+
	Severity (h)	+	-	0	+
Ex ante prioritarian SWF	Prevention (p)	+	+/-	+/-	+/-
	Lethality (q)	+	+	+/-	+/-
	Severity (h)	+	+/-	+	+/-
Ex post prioritarian SWF	Prevention (p)	0	+	+	+/-
	Lethality (q)	+	0	-	+/-
	Severity (h)	+	-	+	+/-

Note: Symbols describe the sign of the effect of an increase in the column parameter on the social value of a marginal decrease of the row parameter. Symbols denote increase (+), decrease (-), no effect (0), ambiguous effect (+/-).

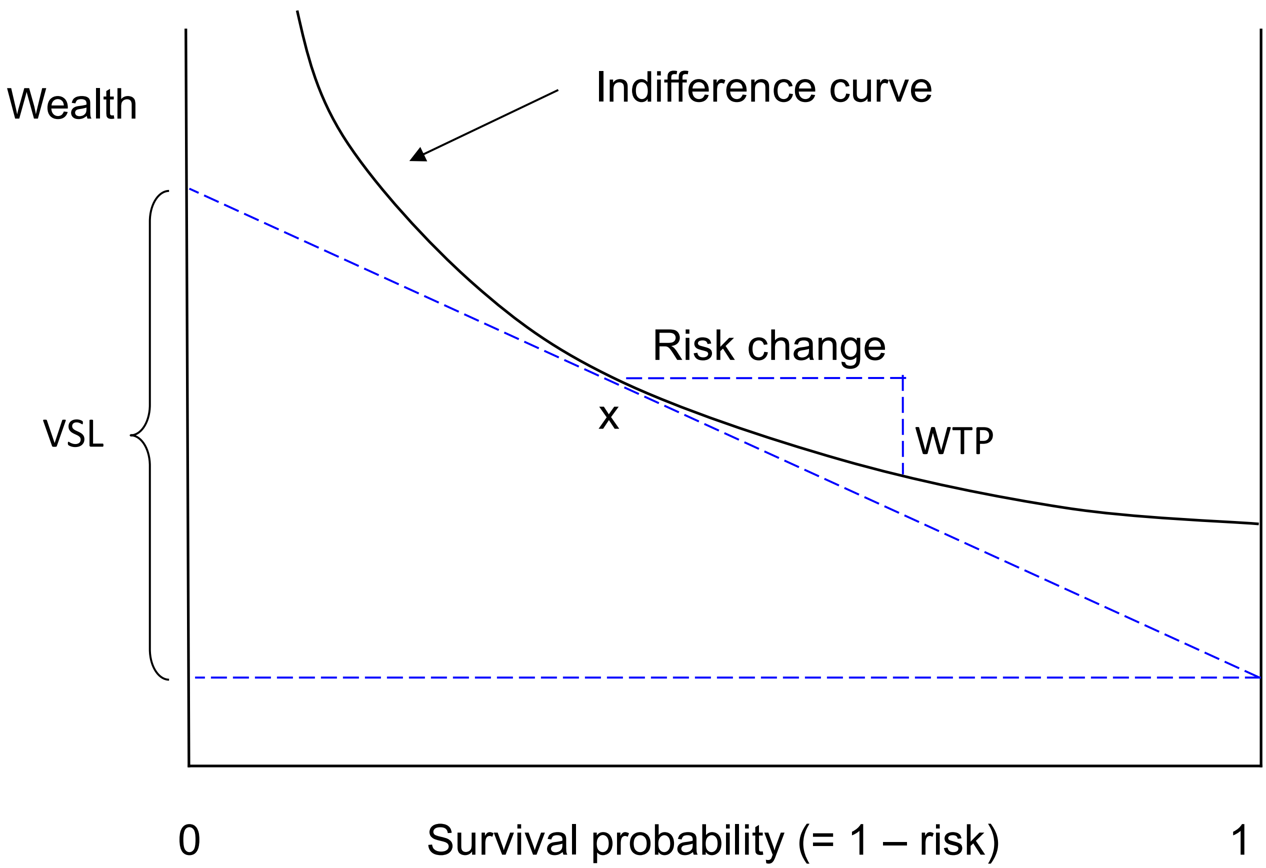

Figure 7.1. VSL is the slope of the indifference curve between the individual's wealth and survival probability at her current position

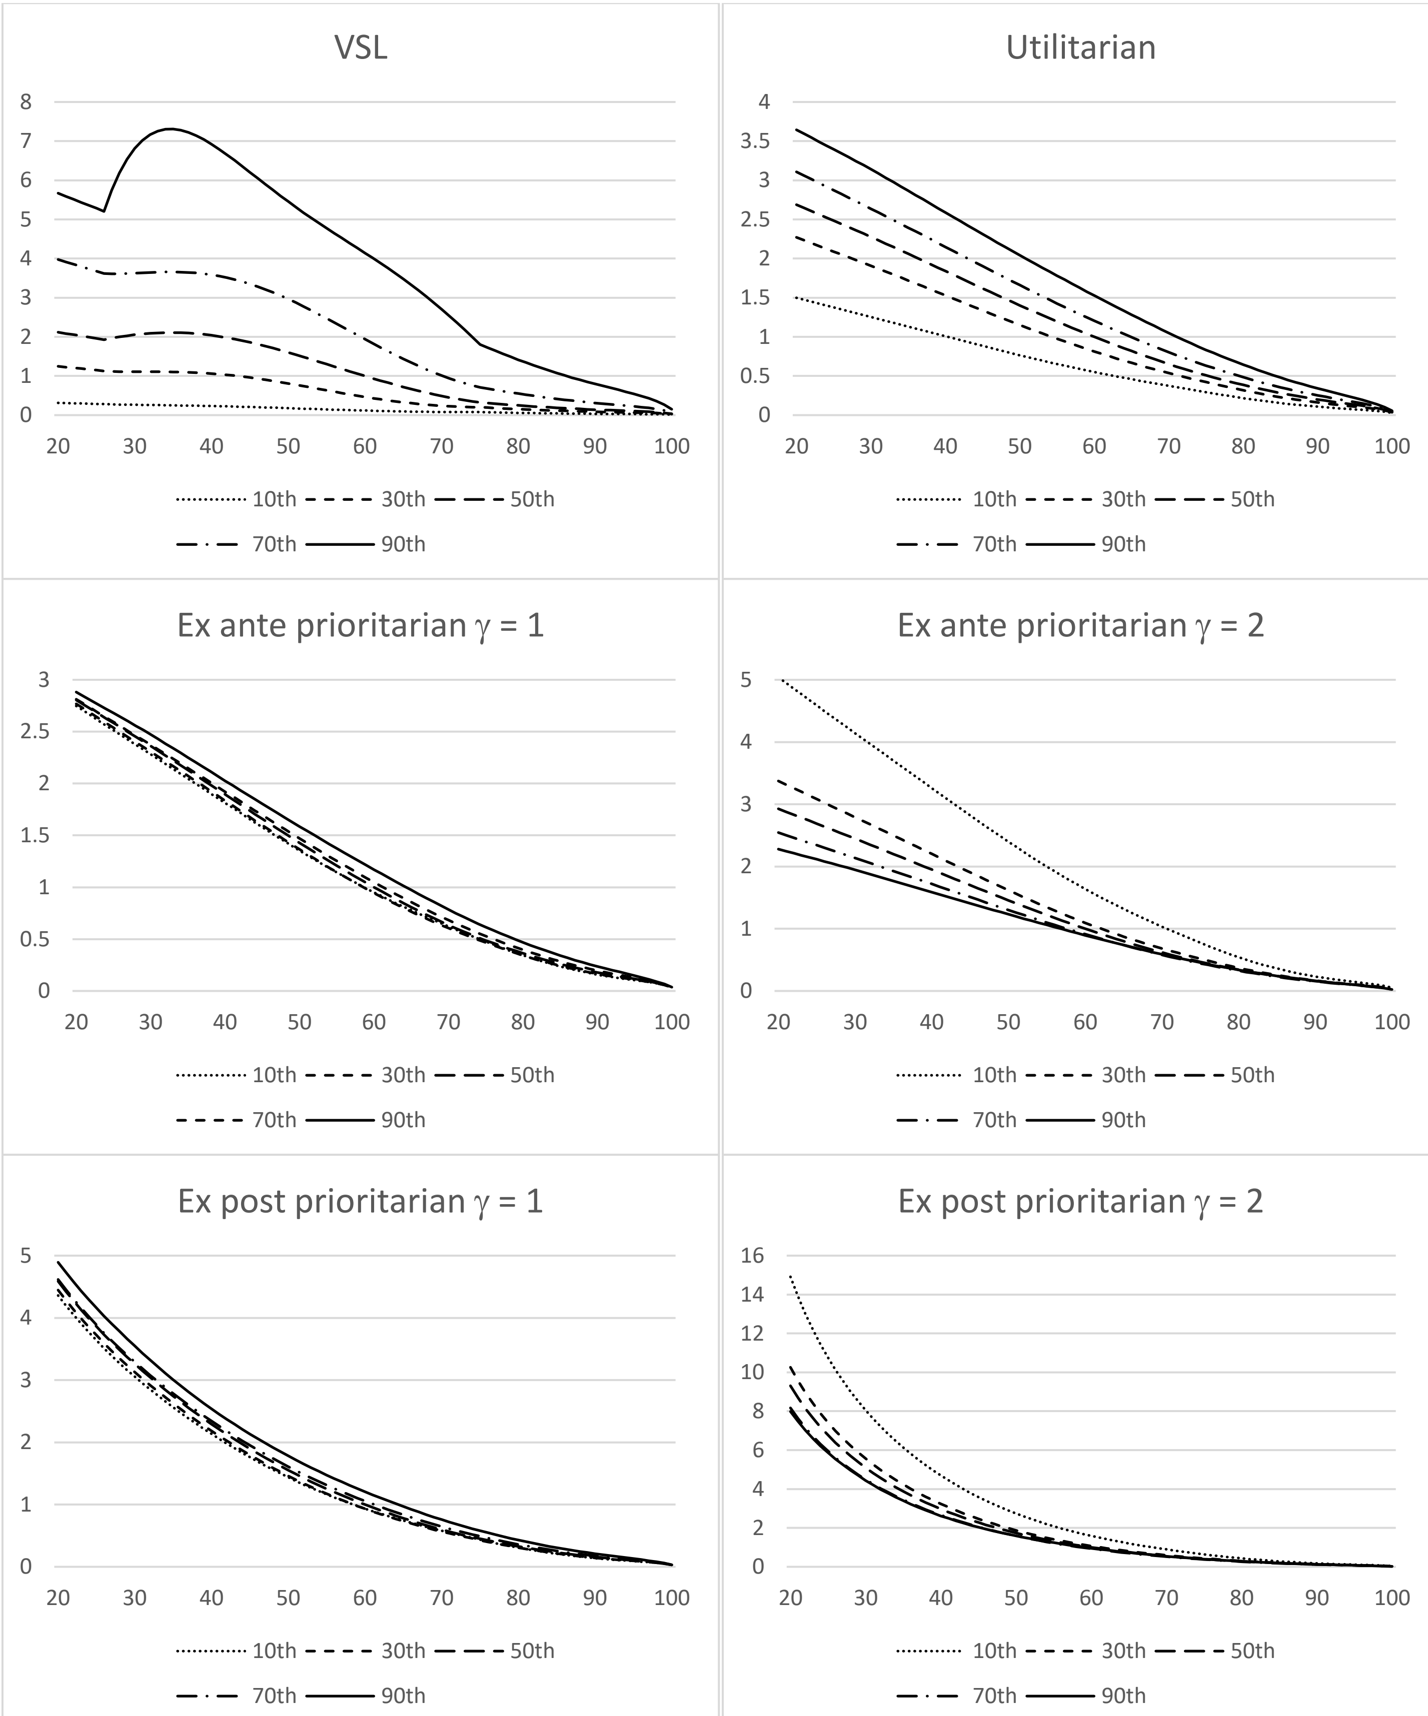

Figure 7.2. SVRR normalized for age and income (= 1 for age 60, median income) for BCA and alternative SWFs.
Source: Adler (2021)

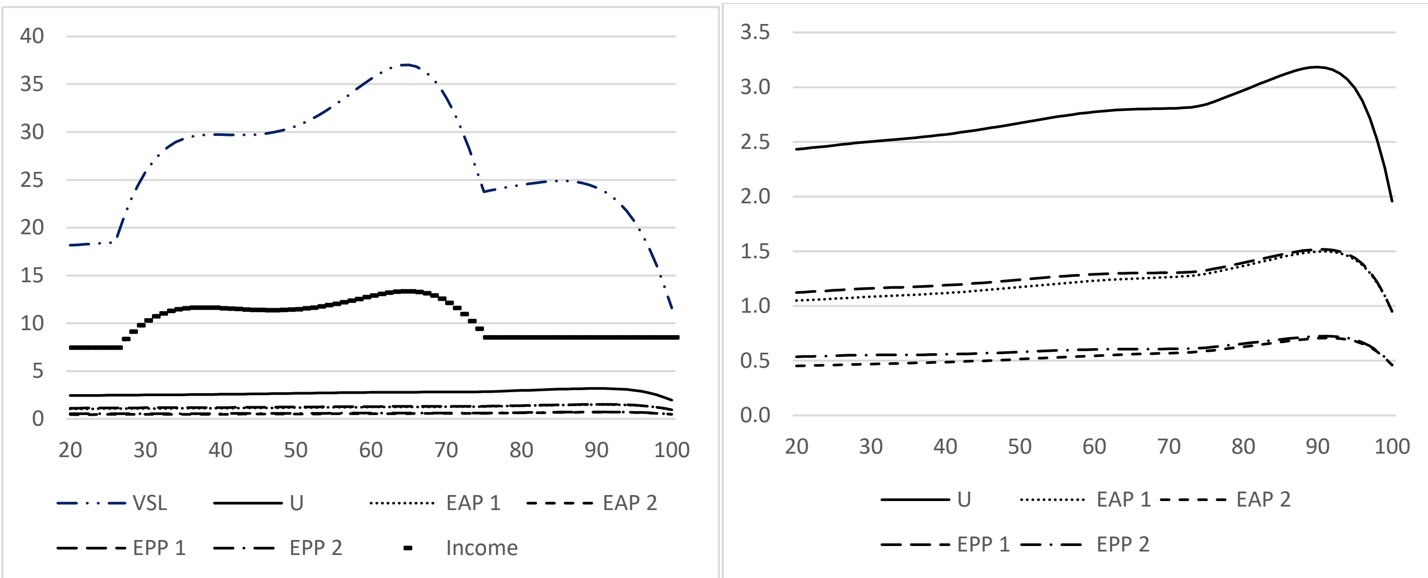

Figure 7.3. Ratio of SVRR at 90th percentile to 10th percentile of income under BCA and alternative SWFs. Left panel includes ratio of income at 90th to 10th percentile.
Source: After Adler et al. (2021)

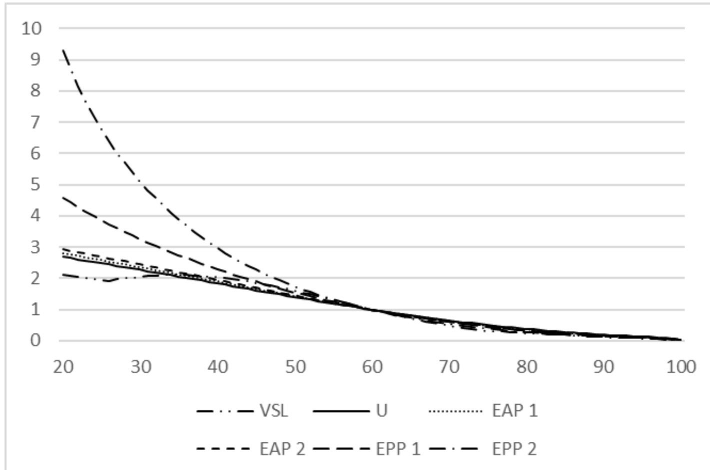

Figure 7.4. SVRR for median income normalized for age (= 1 for age 60) for BCA and alternative SWFs.
Source: After Adler et al. (2021).