Journal of Benefit-Cost Analysis (2024), 15: S1, 84–109
doi:10.1017/bca.2024.28

ARTICLE

From Benefit–Cost Analysis to Social Welfare: A Pragmatic Approach

Maddalena Ferranna[1], James K. Hammitt[2,3] and Lisa A. Robinson[4]

[1]Department of Pharmaceutical and Health Economics, University of Southern California, Los Angeles, CA, USA
[2]Center for Health Decision Science, Harvard T.H. Chan School of Public Health, Boston, MA, USA
[3]Toulouse School of Economics, Université de Toulouse Capitole, Toulouse, France
[4]Center for Health Decision Science, Harvard T.H. Chan School of Public Health, Boston, MA, USA
Corresponding author: Maddalena Ferranna; Email: ferranna@usc.edu

Keywords: equity; social welfare functions; low- and middle-income countries; value-per-statistical-life; distributional weights

Abstract
Conventional benefit–cost analysis is well-established and widely used to assess interventions designed to improve public health and welfare. While it has many advantages, it has well-known limitations. Chief among these is its inattention to the distributional equity of the impacts. To measure individual well-being, the conventional approach relies on individuals' willingness to exchange their own income for the outcomes they experience. To measure societal welfare, it relies on simple aggregation of these values across individuals. This approach reflects a relatively narrow conception of welfare and ignores how impacts are distributed across advantaged and disadvantaged individuals. Social welfare analysis has been proposed as an alternative approach to address these limitations, but real-world applications are rare due largely to the complexity of the calculations. This article provides a pragmatic approach for conducting equity-sensitive benefit–cost analysis globally that addresses data limitations and other challenges, illustrated with example applications. It formally develops and implements equity weights that adjust for the decreasing marginal value of money and for additional moral considerations, prioritizing increases in welfare for those who are worse off.

1. Introduction

Conventional benefit–cost analysis (BCA) is the dominant form of economic evaluation used to compare the harms and improvements associated with interventions undertaken outside the health-care sector, including those that significantly affect public health. This conventional approach has many advantages. It focuses on the efficient use of resources, estimating the opportunity costs of investing in a defined policy rather than using the resources for other things. It is not paternalistic; it relies on individual preferences, assuming each individual is the best or most legitimate judge of their well-being. It measures harms and improvements using the same metric (monetary value), facilitating comparison and prioritization of interventions with different types of impacts. It is well-established and widely used. Substantial best practice guidance is available, and numerous analyses have been completed, easing application.

Conventional BCA has important limitations, however, many of which relate to its inattention to distributional concerns. To measure changes in individual well-being, it relies on estimates of individuals' willingness to exchange their own money for gains or losses they experience (their willingness to pay or willingness to accept compensation).[1] These values reflect individuals' budget constraints, leading to larger values for wealthier individuals.[2] To estimate changes in social welfare, it relies on simple aggregation of gains and losses across individuals, without considering preferences for distribution.

This conception of welfare is relatively narrow, failing to consider how the policy affects individuals with different levels of well-being. It also fails to address distributional equity, an important concern of decision-makers and other stakeholders as evidenced by the many long-standing guidance documents that require supplementing conventional BCA with information on distribution (e.g. Clinton, 1993, Robinson *et al*., 2019a, HM Treasury, 2022).

Social welfare analysis is an alternative to conventional BCA. Social welfare analysis first estimates the impact of the intervention on individuals using broader measures of well-being. These individual *well-being* gains and losses are then aggregated across individuals through the use of a social welfare function (SWF). The specific aggregation rule embodies ethical concerns for distributional equity. For example, the utilitarian SWF corresponds to the sum of individuals' well-being, which is assumed to increase at a diminishing rate as income rises. The prioritarian SWF corresponds to the sum of a nonlinear function of individuals' well-being that attaches more weight to well-being changes experienced by the less well-off. Compared to conventional BCA, social welfare analysis is sensitive to the distribution of policy impacts across the population and to the correlation between policy impacts and preexisting inequities. It has a long tradition (see Adler, 2019, Adler and Norheim, 2022, Adler, 2024, and references therein), and it is routinely used in fields such as optimal tax theory and climate change economics (Tuomala, 2016; Nordhaus, 2008).

In practice, social welfare analysis is often approximated by applying distributional weights to the estimates from conventional BCA (Adler, 2016, Fleurbaey and Abi-Rafeh, 2016). For example, weights to approximate a utilitarian SWF have been developed for use in U.K. policy analyses (HM Treasury, 2022) and in new U.S. guidance for assessing regulations and federal spending (OMB, 2023a, c). Weights to approximate a prioritarian SWF have been studied but have not yet been widely used in practice (Adler, 2016).

Despite the amount of scholarly work on distributional weights, equity-sensitive BCA is rarely adopted in practice. One reason is that it is more demanding than conventional BCA. First, analysts need descriptive information on the distribution of benefits and costs across the population; on who gains and who loses if the policy is implemented. While conventional BCAs sometimes include data on the distribution of benefits, the distribution of costs has

[1] As discussed in Robinson and Hammitt (2011) and elsewhere, care must be taken to ensure that these values are based on reasonably knowledgeable and thoughtful preferences. In some cases, conventional BCA might diverge from this approach, given concerns about cognitive capacity or other issues. For example, for young children, parental values are often used as proxies for the child's preferences (see Robinson et al., 2019b).

[2] In theory, BCA should rely on individuals' monetary values for the consequences each faces. In practice, analysts often use the same monetary value for consequences across the population due to gaps and limitations in the empirical research. Typically, these are population-average values. The result is that the value of a cost or a benefit is the same regardless of whether it is experienced by a wealthier or a poorer individual.

rarely been considered (Robinson *et al*., 2016). In addition, equity-sensitive BCA requires a measure of individual well-being and information on its distribution across the population. Such measures require data on the distribution of attributes relevant to well-being, which may include income, health, education, employment status, environmental quality, and so forth. Equity-sensitive BCA also requires defining an appropriate SWF. Most proposals for distributional weighting adopt a utilitarian SWF and account only for differences in income, although there is uncertainty about the sensitivity of marginal utility to changes in income (Acland and Greenberg, 2023, OMB, 2023b). In addition, the available distributional weights typically focus on developed countries, many of which have robust systems for collecting demographic and other data. For low- and middle-income countries, the data available on attributes of well-being are typically much more limited.

The goal of this article is to provide a simple approach to the computation of distributional weights that (i) goes beyond simply adjusting for the marginal utility of income, thereby recognizing that differences in other dimensions of well-being – health and longevity especially – also matter; and (ii) takes into account the limited data available for many countries, especially those that are low- or middle-income.[3] We focus on the application of weights to interventions implemented outside the health-care system that significantly affect health and longevity, such as those addressing environmental, transportation, occupational, nutritional, behavioral, and other risks.

We aim to aid practitioners in supplementing conventional BCAs conducted in countries with varying data availability, including high-, middle-, and low-income countries, to better inform decision-makers about the effects of interventions on broader measures of well-being and distributional equity. Our approach is intended to improve the available information without requiring more data, time, and resources than available to most analysts.

This article provides simple formulas BCA practitioners can use to compute weights tailored to the characteristics of the population affected by an intervention. It also includes tables of weights for different country income groups. The article is structured as follows. In Section 2, we briefly review the theory behind the construction of distributional weights and the link to social welfare analysis. In Section 3, we derive two types of weights to be used in equity-weighted BCA: utilitarian weights that adjust for differences in marginal utility of income and prioritarian weights that incorporate a greater concern for improving the well-being of those who are worse off. We discuss the data requirements and the calibration of the weights, especially for settings where data are limited. In Section 4, we illustrate the application of weights in a stylized example and contrast the results to those of a conventional BCA. The last section offers some concluding thoughts.

2. Social welfare analysis and distributional weights

Equity-sensitive BCA is a practical method to implement social welfare analysis. Social welfare analysis characterizes the value of an intervention as the associated change in social

[3] As of 2024, the World Bank categorizes low-income economies as those with a gross national income (GNI) per capita of $1,135 or less in 2022; lower middle-income economies as those with a GNI per capita between $1,136 and $4,465; upper middle-income economies as those with a GNI per capita between $4,466 and $13,845; high-income economies as those with a GNI per capita of $13,846 or more (https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups, as viewed 9 March 2024). These categories are defined using the Atlas method and market exchange rates to convert currencies to U.S. dollars.

welfare, where social welfare is a function of the distribution of well-being levels (or utilities) across the population. Many scholars have shown that the change in social welfare associated with an intervention can be approximated by the sum of individuals' monetary equivalents for the intervention (willingness to pay or willingness to accept compensation) weighted by the marginal social value of each individual's income (e.g. Arrow, 1963, Drèze and Stern, 1987, Blackorby and Donaldson, 1990, Adler, 2016).[4] Hence weighted BCA is a convenient, short-cut method to approximate the results of a social welfare analysis. The rationale is that income changes experienced by wealthy people have a smaller effect on social welfare than income changes experienced by less wealthy people. Moreover, social welfare may be more sensitive to improvements in well-being of individuals who are less well-off. The weight attached to individuals' monetary equivalent has two components that reflect these concerns: the marginal utility of income for the individual and the marginal social value of the individual's utility.

The computation of weights based on SWFs requires two ingredients: an interpersonally comparable measure of individual well-being; and a SWF to aggregate individuals' well-being levels. Interpersonal comparability of well-being is necessary to compare changes in well-being levels across individuals and to determine who is disadvantaged (in well-being terms).

Metrics for assessing and comparing the well-being of individuals can be categorized into three broad groups (Adler and Decancq, 2022, Adler and Fleurbaey, 2016). The first category is represented by objective well-being, which includes the attainment of a list of objective goods (e.g. being healthy, having sufficient financial resources) (Sen, 1999). A popular example of an objective well-being index is the Human Development Index.[5] The second category includes subjective well-being, which is derived from individuals' reports of life satisfaction, experience of emotions, or other measures of mental states (e.g. Layard and De Neve, 2023). The third category of well-being measures is preference-based well-being, which depends on a bundle of attributes (e.g. income, longevity, health) and individuals' judgments of the trade-offs between attributes. The primary approaches to building a preference-based well-being measure include equivalent income, where an individual's income is adjusted by the value of nonmarket attributes (e.g. health) and the adjustment is based on the individual's preferences over those attributes (e.g. Decancq *et al.*, 2015; Fleurbaey and Ponthiere, 2023), and von Neumann-Morgenstern (vNM) utility, which assigns a score to bundles of attributes based on individuals' preferences between probability distributions over alternative possible lives (e.g. Adler, 2019, Cookson *et al.*, 2021).

Objective well-being measures are, by definition, interpersonally comparable as they depend on a common pre-specified combination of objective goods. However, they do not respect individual preferences. There is a lively debate on whether subjective well-being measures permit interpersonal comparisons since life satisfaction, happiness, and emotions may mean different things to different individuals, and individuals may interpret well-being

[4] In the discussion that follows, we often refer to willingness to pay for convenience, recognizing that willingness to accept compensation may also be used and that some monetary measures reflect a market equilibrium between supply and demand rather than willingness to pay or willingness to accept compensation for a change from the status quo. We also refer to income rather than wealth, recognizing that well-established estimates of income are more plentiful than estimates of wealth.

[5] See https://hdr.undp.org/data-center/human-development-index#/indicies/HDI.

scales differently (Adler, 2013, Graham, 2016, Sunstein, 2016). However, significant progress is being made in developing evaluation approaches that rely on measures of subjective well-being (e.g. HM Treasury, 2021a, 2021b, Frijters *et al.*, 2024). Preference-based well-being metrics respect individual preferences, and scholars have proposed approaches that allow interpersonal comparison of well-being (Adler and Decancq, 2022). However, individual preferences are often difficult to elicit, especially if they are heterogenous.

In this article, for pragmatic reasons, we focus on a single well-being measure, vNM utility. Like other preference-based well-being measures, vNM utility retains an important strength of conventional BCA: it is based on individuals' preferences for the gains and losses they would likely experience under each policy option. Within the realm of preference-based measures, we rely on vNM utility rather than equivalent income because it incorporates the judgment that the gain in well-being associated with an increase in income is a diminishing function of income. In other words, the marginal utility of income (or of most goods and services) is diminishing. As noted earlier, utilitarian weights, which account for the greater value of money to a poor person than a wealthy person, have received increased attention in recent years; these weights are consistent with the vNM framework.

The second ingredient to compute the weights is a SWF. Commonly used SWFs are the utilitarian and the prioritarian ones. Utilitarianism is the view that social welfare is the sum of individuals' well-being. Unlike conventional BCA, utilitarianism weights individuals' monetary equivalents to reflect the greater value of money to a poor person than a wealthy person, that is, to reflect the diminishing marginal utility of income. Prioritarianism is the view that social welfare is the sum of individuals' well-being transformed to give more weight to changes in well-being that accrue to those who are worse off and less weight to changes that accrue to those who are better off. Prioritarianism weighs individuals' monetary equivalents to reflect both the marginal utility of income and the greater marginal social value of improving the situation of disadvantaged people.

Which SWF should be used to compute the weights? The SWF reflects preferences for distributional equity. As such, it constitutes an ethical choice. Philosophers have long debated which theory of distributive justice is best. Public surveys can provide some clues about the moral preferences held by the population. Schokkaert and Tarroux (2022) conclude that there is a large degree of heterogeneity in individuals' preferences for redistributive justice and that the axioms underlying a prioritarian SWF are often not satisfied. In particular, preferences for redistribution depend on the good to be redistributed (e.g. income or health), on issues of desert and responsibility, and on individual personal circumstances, among other things (Alesina and Angeletos, 2005, Robson *et al.*, 2017, Almås *et al.*, 2020).

In this article, we take the view that it is not up to the analyst to choose a SWF. Rather, the analyst should provide alternative sets of weights, representing alternative preferences for distributional equity. Policymakers can then review the results and determine how to best incorporate them in their decision-making. We are agnostic about the "correct" SWF and develop both utilitarian and prioritarian weights.

It is worth noting that, in equity-sensitive BCA, the weights do not depend on the particular policy to be assessed. Thus, the weights can be computed in advance and then combined with the benefit and cost estimates for each group for each intervention being assessed. However, the weights do depend on the chosen well-being metric and SWF. We hope and expect that our analysis will be extended to other concepts of well-being and social welfare.

3. Constructing the weights

In constructing weights that can be applied to the results of a conventional BCA, we begin by assuming that the analyst reports the (unweighted) distributions of monetary values (either net benefits or benefits and costs) across groups of individuals differentiated by socio-demographic characteristics (e.g. income or wealth, age, health status). We assume further that decisions about how to conduct the conventional BCA (including the choice of discount rate, the values per statistical life or per case used to compute the value of mortality and morbidity risk reductions, and the estimates of the distributions of the effects) have already been made.

In other words, we assume that the starting point is a conventional BCA supplemented by a distributional analysis that takes into account how monetary values vary across individuals with different attributes. Our task in this article is to estimate weights to convert the results into measures of social welfare more broadly defined. If the conventional BCA instead follows the common practice of using standard (population-average) unit values without adjustment for all individuals, for example, using the same VSL regardless of income or age, analysts will need to first adjust the values to account for individual differences. We illustrate one approach to estimating income- and age-dependent VSL in the stylized example of Section 4.

3.1. Assumptions

As discussed earlier, computing weights requires selecting a well-being metric and the specification of a SWF. For the former, we adopt vNM utility as our measure of well-being and assume that individuals have identical preferences. Although this is a restrictive assumption, it is difficult to measure differences in preferences across individuals and unclear how much bias is introduced by assuming identical preferences.

To compute the vNM utility index to be used in the weights, we make the following additional assumptions.

First, we use *lifetime* utility as our individual well-being metric and assume that lifetime utility is the sum of utility in each time period that an individual is alive. We assume that lifetime utility depends on three attributes: longevity, the lifetime path of consumption (proxied by income), and the lifetime path of health. While additional attributes could be considered (e.g. nonwork time, environmental quality), given our global focus, we ignore these attributes due to concerns about data availability and quality in many countries.

Second, we consider inequalities across age and socioeconomic status, understanding that, in principle, well-being attributes (consumption/income, longevity, and health) are heterogeneously distributed across the population depending on the age and socioeconomic status of individuals. We ignore other sociodemographic characteristics (e.g. race, gender). Although individuals vary in consumption/income, longevity, and health, we assume they have homogenous preferences over these attributes (i.e. individuals have the same preferences over alternative bundles of consumption/income, longevity, and health). We use income quintiles as a proxy for socioeconomic status.

Third, we adopt the following functional form for the period utility function:

$$u = h\frac{y^{1-\epsilon} - \bar{y}^{1-\epsilon}}{1-\epsilon} \tag{1}$$

where y represents the period consumption/income, h is the health-related quality of life, $\bar{y}$ is subsistence consumption/income, and ϵ is the elasticity of the marginal utility of income.[6] This functional form is ubiquitous in the theoretical literature on the value of health (e.g. Murphy and Topel, 2006). An important feature of this formula is that the marginal utility of income is increasing in health-related quality of life. In other words, healthy people gain more utility from an increase in consumption than unhealthy people. Said otherwise, improving the health of a wealthy person increases well-being more than improving the health of a poor person. The complementarity between health and income is supported empirically (Finkelstein *et al.*, 2013). (For simplicity, we assume $h \geq 0$ where $h = 0$ corresponds to a health state no better nor worse than death.)[7]

For the specification of social welfare, we apply an Atkinson SWF of the form:

$$W = \sum_i \frac{v_i^{1-\gamma}}{1-\gamma} \tag{2}$$

where the sum is over all members of the affected society, the variable v_i represents the lifetime well-being of individual i, and the parameter $\gamma\,(\geq 0)$ represents the degree of inequality aversion ($W = \sum_i \log v_i$ if $\gamma = 1$).[8] The Atkinson weighting function is consistent with the view that differences in well-being are related to the ratio (rather than the difference) between utility levels. The Atkinson SWF is very flexible and captures different distributional equity preferences. When $\gamma = 0$, the SWF is utilitarian. When $\gamma > 0$, the SWF is prioritarian, that is, it attaches larger value to well-being changes experienced by the worse-off. The larger is γ, the greater is the concern about the well-being of the worse-off. When $\gamma \to \infty$, the approach becomes egalitarian and aims at improving the well-being of those at the bottom of the well-being scale.[9]

In the presence of risk, a prioritarian SWF depends on whether well-being is measured *ex ante* or *ex post.* An *ex ante* approach focuses on the distribution of *expected* lifetime well-being across individuals. Expected lifetime well-being considers the risk faced by individuals across their lifetime profile of income, health, longevity, and other attributes. In contrast, an *ex post* approach to well-being evaluation focuses on the *realized* distribution of lifetime well-being across individuals. Realized lifetime well-being considers the lifetime

[6] The functional form of the utility function assumes that the analyst has data on health-related quality of life. As explained later, we proxy h with a disability index derived from the Global Burden of Diseases Study (GBD 2019 Diseases and Injuries Collaborators 2020). However, if data on health-related quality of life are missing for the population of interest, we recommend to simply assume that $h = 1$.

[7] Other functional forms are possible. For example, Cookson et al. (2021) propose a period utility function that is additive in health-related quality of life and utility of income, $h + u(y)$; Lakdawalla and Phelps (2020) introduce risk aversion in health, $g(h)u(y)$, with $g' > 0$ and $g'' \leq 0$; Bommier and Villeneuve (2012) reject risk neutrality with respect to life duration and assume that lifetime well-being is not additively separable. Note that a few contributions suggest that health and income are substitutes instead of being complements (e.g. Tengstam, 2014). We reserve these extensions for future research.

[8] To be accurate, the formula for the Atkinson social welfare function is $W = \sum_i \frac{(v_i - v_{zero})^{1-\gamma}}{1-\gamma}$, where v_{zero} is the lifetime well-being number of a life history specified as the zero point. In this article, we are implicitly setting the zero point equal to a lifetime with subsistence income in each period, implying that $v_{zero} = 0$. An often-used alternative to the Atkinson SWF is the Kolm-Pollack SWF (Adler and Norheim, 2022).

[9] As $\gamma \to 0$, the limit of the Atkinson social welfare function is leximin (Bosman, 2007).

profile of income, health, and other attributes experienced by an individual, from birth to death. There are axiomatic reasons to prefer either an *ex post* or an *ex ante* approach. In particular, the *ex ante* approach does not satisfy principles of stochastic dominance and can lead to temporally inconsistent preferences. On the other hand, the *ex post* approach violates the *ex ante* Pareto principle (see e.g. Adler, 2019, for a more thorough account of the properties of *ex ante* and *ex post* prioritarianism). We are agnostic to which approach is ethically preferable and construct weights using both *ex post* and *ex ante* approaches.

3.2. *Equity-sensitive BCA*

In this section, we briefly present mathematical expressions of the weights for utilitarian and prioritarian SWFs and discuss their main components. Appendix A provides a formal derivation of the equity-sensitive BCA rules and the corresponding weights. The next section describes the data required to estimate the weights. All weights are normalized with respect to the characteristics of the median individual in the population.

Let us consider a population of N individuals divided into "types" (i.e. individuals that share similar sociodemographic characteristics), and let N_i be the number of people of type i, with $\sum_i N_i = N$. Let us consider a policy that increases type i's survival probability by Δp_i and reduces their income by Δy_i (Δy_i could be negative, denoting an income gain). Note that $N_i \Delta p_i$ represents the expected number of deaths prevented among individuals of type i. Let VSL_i denote the value-per-statistical-life of individuals of type i.

According to conventional BCA, the total net benefits of such an intervention are equal to:

$$\sum_i N_i (\Delta p_i VSL_i - \Delta y_i) \tag{3}$$

In other words, conventional BCA ranks policies based on the unweighted sum of individuals' monetary equivalents for the policy, $\Delta p_i VSL_i - \Delta y_i$.

In the utilitarian case, the total net benefits of the intervention are given by:

$$\sum_i N_i w_i^U (\Delta p_i VSL_i - \Delta y_i) \tag{4}$$

where w_i^U represents the utilitarian weight associated with type i's monetary equivalent for the policy. Utilitarian weights are given by:

$$w_i^U = \left(\frac{y_i}{y_{med}} \right)^{-\epsilon} \tag{5}$$

where y_i is individual i's income, y_{med} is the population median income, and $\epsilon \geq 0$ is the elasticity of the marginal utility of income. The utilitarian weight adjusts for differences in income across the population. In particular, it attaches larger values to net monetary benefits experienced by lower income individuals. Note that the utilitarian weights do not depend directly on an individual's age but may depend indirectly on age if income depends on age.

In the *ex ante* prioritarian case, the total net benefits of the intervention are given by:

$$\sum_i N_i w_i^{EAP} (\Delta p_i VSL_i - \Delta y_i) \tag{6}$$

where w_i^{EAP} represents the *ex ante* prioritarian weight associated with type i's monetary equivalent for the policy. The *ex ante* prioritarian weights are equal to:

$$w_i^{EAP} = \left(\frac{y_i}{y_{med}}\right)^{-\epsilon}\left(\frac{QALE_i}{QALE_{med}}\frac{y_i^{1-\epsilon}-\bar{y}^{1-\epsilon}}{y_{med}^{1-\epsilon}-\bar{y}^{1-\epsilon}}\right)^{-\gamma} \tag{7}$$

where $QALE_i = \sum_{t=0}^{a_i} h_i(t) + \sum_{t=a_i+1}^{T} \pi_i(t;a_i)h_i(t)$ is the lifetime quality-adjusted life expectancy (from birth to death) of individual of type i who is currently at age a_i, $QALE_{med}$ is the population median quality-adjusted life expectancy, $\bar{y}$ is subsistence income, and γ is the inequality aversion parameter. QALE depends on the individual's expected longevity and on health-related quality of life h_i. It is age-dependent; for an individual currently at age a_i, QALE depends on past health-related quality of life (from age 0 to age a_i) and future expected quality of life (from age $a_i + 1$ to the maximum possible lifespan, set to age T), where $\pi_i(t;a_i)$ is the probability of surviving to age $t \in (a_i, T]$ conditional on being alive at age a_i. Note that QALE tends to increase with age because the risk of dying young is eliminated.

The *ex ante* prioritarian weight includes two terms: the marginal utility of income, as in the utilitarian weight; and the marginal social value of expected lifetime well-being, represented by the second term in Equation (7). Individual well-being is an increasing function of income[10] and QALE. As $\gamma > 0$, *ex ante* prioritarianism attaches larger weight to individuals with lower income and lower QALE. In addition, the larger the inequality aversion parameter γ, the larger the weight assigned to net benefits experienced by the poor and by individuals with low QALE.

In contrast, in Appendix A, we show that the total net benefits of the intervention under *ex post* prioritarianism are equal to:

$$\sum_i N_i \left\{ w_i^{EPP}(a_i)\Delta p_i VSL_i - \Delta y_i \sum_{t \geq a_i} \mu_i(t;a_i) w_i^{EPP}(t) \right\} \tag{8}$$

where $w_i^{EPP}(t)$, with $t \geq a_i$, represents the *ex post* prioritarian weight associated with the costs and benefits experienced by individuals of type i who die exactly at age t in the absence of the policy, and $\mu_i(t;a_i)$ is the proportion of people currently of age a_i who will die exactly at age $t \geq a_i$ without the policy.

The *ex post* prioritarian rule is considerably different from both the utilitarian and the *ex ante* prioritarian ones. In particular, while the utilitarian and the *ex ante* prioritarian weights are applied to individuals' monetary equivalents, $\Delta p_i VSL_i - \Delta y_i$, *ex post* prioritarian weights are not.[11] The latter depend on the *realized* longevity of individuals in the absence of the intervention since, from an *ex post* prioritarian perspective, what matters is the *realized* well-being of individuals. Specifically, a proportion Δp_i of type i individuals would have died at age a_i without the policy; their *ex post* prioritarian weight is thus $w_i^{EPP}(a_i)$, and it is applied to both the mortality risk change and the income change, $VSL_i - \Delta y_i$. A proportion $\mu_i(a_i;a_i) - \Delta p_i$ dies at age a_i both with and without the policy; their *ex post* prioritarian weight is thus $w_i^{EPP}(a_i)$, and it is applied only to the income change Δy_i (as they do not experience any survival benefit from the policy). Finally, a proportion $\mu_i(t;a_i)$ dies at age $t > a_i$ both with and

[10] For simplicity, we assume there is no economic growth and income is independent of age.

[11] More formally, conventional BCA ranks policies based on the sum of individuals' monetary equivalents for the policies. Both the utilitarian and the *ex ante* prioritarian policy ranking can be approximated by the sum of the weighted monetary equivalents. This is possible because (i) both the utilitarian and *ex ante* prioritarian objectives are a function of individuals' expected utilities; and (ii) an individual's monetary equivalent is a "sufficient statistic" for their expected utility (i.e. if two policies deliver the same expected utility for an individual, the monetary equivalent for the two policies is the same). In contrast, the *ex post* prioritarian objective cannot be expressed as a function of individuals' expected utilities or monetary equivalents.

without the policy; their *ex post* prioritarian weight is thus $w_i^{EPP}(t)$, and it is applied only to the income change Δy_i. As a consequence, even though individuals of type i are *ex ante* identical, they are different *ex post* since they may die at different ages.

The formula for *ex post* prioritarian weight for an individual who dies at age $t \geq a_i$ in the absence of the policy is:

$$w_i^{EPP}(t) = \left(\frac{y_i}{y_{med}}\right)^{-\epsilon} \left(\frac{H_i(t)}{H_{med}} \frac{y_i^{1-\epsilon} - \overline{y}^{1-\epsilon}}{y_{med}^{1-\epsilon} - \overline{y}^{1-\epsilon}}\right)^{-\gamma} \tag{9}$$

where $H_i(t) = \sum_{s=0}^{t} h_i(t)$ is the realized quality-adjusted longevity of the individual who dies at age t without the policy, and H_{med} is the population median quality-adjusted longevity. Like the *ex ante* prioritarian weights, the *ex post* ones adjust for the marginal utility of income and for the marginal social value of lifetime well-being.[12] Individuals with lower lifetime *realized* well-being are given higher weight. The latter now depends on realized quality-adjusted longevity (as opposed to expected quality-adjusted longevity). In other words, people who die young will have a larger weight than people who die old, even though, from an *ex ante* point of view, they have the same life expectancy.

Applying *ex post* prioritarian weights requires modeling changes in the number of people dying at each age since the weight depends on the age of death rather than at the current age of an individual.[13,14] This constitutes a major difference between the *ex ante* weights and *ex post* ones: The former depend on the current age of the individual experiencing the policy impacts; the latter depend on the baseline age of death of the individual who is experiencing the policy impacts. However, if the impacts of an intervention are independent of the age of death, one could derive *ex post* prioritarian weights as a function of current age by taking the expectation of $w_i^{EPP}(t)$ over the age of death:[15]

$$\overline{w}_i^{EPP}(a_i) = \sum_{t \geq a_i} \mu(t; a_i) w_i^{EPP}(t) \tag{10}$$

In this case, the *ex post* prioritarian rule simplifies as follows:

[12] Again, we assume income is independent of age.

[13] For example, suppose that an intervention affects individuals at age 80 and that they all gain \$100 from the intervention. However, some of them will die in 1 year (even with the intervention) and the others will die 10 years later (even with the intervention). From an *ex post* prioritarian point of view, \$100 to the former increase social welfare more than \$100 to the latter, because the former will die at a younger age and have smaller lifetime wellbeing. Thus, the *ex post* prioritarian weights differ across individuals even though they are all at age 80, and they all obtain the same benefit from the policy.

[14] In many cases, it is not possible to know which individuals' lives are prolonged by an intervention, that is, one can know the expected number of deaths at each age with and without the intervention, but not whether the intervention extends many people's lives by a relatively short amount or extends the lives of a smaller number of people by a larger amount (Hammitt et al., 2020). Consistent with common practice, we assume the fraction of people of each age whose lives are prolonged by an intervention are those whose deaths would be attributable to the absence of the intervention.

[15] Continuing the example of footnote 13, let π be the proportion of people who survive for 10 more years, $w^{EPP}(81)$ the *ex post* prioritarian weight if they die in 1 year, and $w^{EPP}(90)$ the weight if they die in 10 years. The *ex post* prioritarian weight conditional on their current age is: $\overline{w}^{EPP}(80) = \pi w^{EPP}(90) + (1-\pi) w^{EPP}(81)$. The *ex post* prioritarian value of the intervention is $\pi w^{EPP}(90) * \$100 + (1-\pi) w^{EPP}(81) * \$100$. Since everyone gains the same from the policy, the value can be rewritten as: $\overline{w}^{EPP}(80) * \$100$. Thus, in this case, knowing the *ex post* prioritarian weight conditional on current age is sufficient to estimate the value of the intervention. The results would not hold if the (net) benefits of the intervention depended on the age of death (e.g. the 80 years old who dies in a year gains \$100 from the intervention and the 80 years old who dies in 10 years gains \$200).

$$\sum_i N_i \left\{ w_i^{EPP}(a_i)\Delta p_i VSL_i - \Delta y_i \overline{w}_i^{EPP}(a_i) \right\} \tag{11}$$

We will illustrate this simplification in the example below.

Table 1. *Input data for estimating weights*

Variable	Description	Benchmark estimate	Rationale
$\bar{y}$	Subsistence income (level at which individuals are indifferent between survival and death)	\$1.00 per day	Below the “extreme poverty” level (less than \$2.15 per day) as defined by the World Bank.
ϵ	Elasticity of marginal utility of income	1.5	Consistent with recent reviews
γ	Inequality aversion	Range	Ethical choice, illustrate results of different values and defer to decision-makers
y_i	Distribution of income in the population of interest	Derived from World Bank data	Use of publicly available World Bank indicators if specific data for the population of interest are missing
h_i	Health-related quality of life in the population of interest	Derived from Global Burden of Diseases (GBD) Study	Use of age-specific disability indexes derived from GBD data as a proxy for health-related quality of life if specific data for the population of interest are missing (health-related quality of life = 1-disability weight)
π_i	Baseline survival rates in the population of interest	Derived from UN World Population Prospects and literature on the economic gradient of longevity	Reliance on studies on the distribution of life expectancy by income quintile to determine baseline survival rates by income

3.3. Data requirements and calibration

To estimate the proposed weights, we need the following: information about the distribution of income and of mortality and morbidity risks in the population of interest; estimates of the subsistence income $\overline{y}$ and the elasticity of marginal utility of income ϵ; and the inequality aversion parameter γ. We summarize our proposed estimates in Table 1, and describe them in more detail below.

Let us start with the calibration of $\overline{y}$, ϵ, and γ. The subsistence income $\overline{y}$ is a level of income so low that individuals are indifferent between survival and death. The World Bank estimates that almost 700 million people around the world live in "extreme poverty," defined as subsisting on less than \$2.15 per day. It is reasonable to assume that the subsistence income is smaller than the extreme poverty line. We propose to use \$1 a day as a benchmark.

There is an extensive literature on the elasticity of the marginal utility of income. For high-income countries, the proposed estimates often range between 1 and 2. For example, OMB (2023a, b) proposes a value of 1.4 for the United States, based on an update of a meta-analysis of U.S. and U.K. data by Acland and Greenberg (2023). Groom and Maddison (2019) combine results from different estimation methodologies and suggest a value of 1.5 for the U.K. Using subjective well-being data from a large set of countries, Layard *et al*. (2008) conclude that the average elasticity is 1.26. Studies in lower income countries are rare. Using subjective well-being data, Gandelman and Hernández-Murillo (2014) conclude that the elasticity is close to 1 in most countries, independent of their development level. A recent paper by Bergstrom and Dodds (2023) relies on schooling decisions in rural Mexico and estimates that the elasticity is 1.6. As a benchmark, we propose to use a value equal to 1.5.

The degree of inequality aversion γ is an ethical choice. The literature on the elicitation of ethical preferences can provide some indication of reasonable values for γ. However, being an ethical choice, we propose to apply a range of values. In the example below, we assume that γ is equal to 1 or 2.

Finally, the construction of the weights requires data on the distribution of income as well as mortality and morbidity risks across the population of interest. Data quality is a concern; often the analyst has estimates of the average but not the whole distribution. However, the analyst can leverage publicly available datasets such as the World Bank Indicators and the United Nations World Population Prospects to construct approximate distributions of income, mortality, and morbidity. We illustrate the use of available data in the next section.

3.4. Weights by individuals' socioeconomic status, age, and country's income level

As an illustration, we use the previous formulas to construct weights by individuals' socioeconomic status (proxied by income quintile) and age, for countries in each World Bank income category: low-income, lower-middle-income, upper-middle-income, and high-income. For each country type, we divide the population into five homogeneous groups corresponding to the quintiles of the income distribution. We further divide the population into broad age groups (0–19, 20–64, 65+).[16] Thus, we have 15 sociodemographic

[16] While these groups roughly correspond with the ranges defined as "children," "working age adults," and "elderly" in high income countries, we recognize that individuals enter the labor force and are likely to die at younger ages in lower income countries.

Table 2. Economic and demographic variables

Variable	Data			
	Low-income	Lower-middle-income	Upper-middle-income	High-income
Average income per-capita[a]	$1,900	$7,400	$17,400	$51,700
Income share held by each quintile[b]	6.6 % (1st); 10.7 % (2nd); 14.8 % (3rd); 21.1 % (4th); 46.8 % (5th)	6.7 % (1st); 10.9 % (2nd); 15.2 % (3rd); 21.5 % (4th); 45.7 % (5th)	6.5 % (1st); 10.9 % (2nd); 15.2 % (3rd); 21.5 % (4th); 45.9 % (5th)	7.7 % (1st); 12.8 %(2nd); 17.0 %(3rd); 22.5 %(4th); 40.0 % (5th)
Average life expectancy at birth[c]	65.0 years	70.4 years	77.9 years	82.4 years
Average quality-adjusted life expectancy at birth[d]	56.6 years	60.6 years	67.9 years	70.0 years
Life expectancy at birth by income quintile[e]	62.3 years (1st); 64.1 years (2nd); 65.0 years (3rd); 66.8 years (4th); 69.7 years (5th)	67.5 years (1st); 69.6 years (2nd); 70.4 years (3rd); 72.2 years (4th); 75.3 years (5th)	74.5 years (1st); 76.9 years (2nd); 77.9 years (3rd); 80.0 years (4th); 83.4 years (5th)	76.1 years (1st); 79.3 years (2nd); 82.4 years (3rd); 84.3 years (4th); 88.0 years (5th)
Quality-adjusted life expectancy at birth by income quintile[f]	54.5 years (1st); 55.9 years (2nd); 56.6 years (3rd); 57.9 years (4th); 60.0 years (5th)	58.6 years (1st); 60.0 years (2nd); 60.6 years (3rd); 61.9 years (4th); 64.1 years (5th)	65.5 years (1st); 67.2 years (2nd); 67.9 years (3rd); 69.3 years (4th); 71.7 years (5th)	65.5 years (1st); 67.8 years (2nd); 70.0 years (3rd); 71.2 years (4th); 73.7 years (5th)

Table 2. *Continued*

Variable	Data			
	Low-income	Lower-middle-income	Upper-middle-income	High-income
Median age of the population[g]	18 years	26 years	36 years	41 years

[a]Gross National Income per-capita, PPP, constant 2017 international $ (World Bank, 2024).
[b]Simple averages of income shares by country income group. Original data are drawn from the World Bank Poverty and Inequality Platform (World Bank, 2022).
[c]United Nations, World Population Prospects 2022, 2024 life table survivors (UN, 2022).
[d]Own computations based on survival rates from the 2022 World Population Prospects (UN, 2022) and disability indexes from the 2019 Global Burden of Disease Study.
[e]For low- and middle-income countries, the distribution of life expectancy at birth by income quintile is based on an Indian study (Asaria et al., 2019). We assume that the ratio between life expectancy in each quintile and median life expectancy in India applies to all low- and middle-income countries. For high-income countries, the distribution is based on a U.S. study (Chetty et al., 2016) and derived from Adler (2017), Appendix C.
[f]Own computations based on the disability indexes from the 2019 Global Burden of Disease Study and the income-dependent mortality rates (see note e).
[g]United Nations, World Population Prospects 2022 (UN, 2022).

subgroups. We then determine the utilitarian and prioritarian weights attached to each subgroup for each country income category. To simplify, we assume that there is no social mobility, that is, individuals do not move across the income distribution as they age, and that there is no income growth. As a consequence, socioeconomic status is age-independent.

We first discuss the data used in the estimation, then summarize them in Table 2. Data on average income per-capita are taken from the World Development Indicators (World Bank, 2024). Income is measured in terms of purchasing power parity Gross National Income (GNI) (constant 2017 international dollars).[17] We rely on the World Bank Poverty and Inequality Platform for information on the distribution of income in many countries around the world (World Bank, 2022). We use data on the income share held by each quintile (i.e. how much of the country's aggregate income is earned by each income quintile). We average across countries belonging to the same income category to estimate the distribution of income in the representative countries.

To determine QALE, we combine the life tables from the 2022 World Population Prospects (UN, 2022) with disability metrics from the 2019 Global Burden of Disease Study (GBD) (GBD 2019 Diseases and Injuries Collaborators 2020).[18] Both datasets provide data for single countries and for World Bank income groups; our computations are based on the latter. We first use data from the 2019 GBD and compute age-specific disability indexes by dividing the "years of healthy life lost due to disability" (YLD) at each age by the prevalence of illnesses for each country income group at that age.[19] We assume that one minus the disability index approximates health-related quality of life at different ages and in different country income groups (disability is measured on a zero-to-one scale, with one representing death and zero full health). The 2022 World Population Prospects provide life table survivors, that is, estimates of survivors by age for a hypothetical cohort of 100,000 newborns who would be subject for all their lives to the mortality rate of a given year (UN, 2022). We use 2024 as the reference year and compute the probability at birth of surviving to different ages. We apply the 2024 survival curve to all age groups and neglect that previous birth cohorts may have different survival chances (i.e. the probability of surviving to age 40 of the 2024 birth cohort may be larger than the probability of surviving to age 40 of the 2004 birth cohort).[20] We then combine those survival probabilities and the computed health index at different ages to determine the QALE measures.[21] Table 2 reports both the average life expectancy at birth derived from

[17] The World Bank defines income groups based on GNI per capita computed using the Atlas method and exchange rates; the 2024 categories are based on 2022 income levels in U.S. dollars. The GNI per capita estimates used in the illustrative example are instead averages for each category based on purchasing power parity in 2017 international dollars and, in some cases, are above the U.S. dollar thresholds used to define the categories.

[18] https://vizhub.healthdata.org/gbd-results/

[19] YLD at each age is calculated as the sum over diseases of the product of age-specific prevalence and disease-specific disability weight, with adjustment for comorbidities. Hence dividing YLD by prevalence (at each age) yields an average disability weight.

[20] The 2022 World Population Prospects provide estimates up to the year 2021 and population projections from 2022 to 2050. In the study, we use the medium-variant population projection associated to the year 2024.

[21] Let $\pi_i(t)$ be the probability (at birth) of surviving to age t and $h_i(t)$ be the health-related quality of life at age t. We compute quality-adjusted life expectancy at birth using the following formula: $\sum_{t=0}^{100}\pi_i(t)h_i(t)$, where 100 is the maximum length of life. The probabilities $\pi_i(t)$ are derived from the UN World Population Prospects. Quality adjusted life expectancy at age a_i is $\sum_{t=0}^{a_i}h_i(t)+\sum_{t=a_i+1}^{100}\frac{\pi_i(t)}{\pi_i(a_i)}h_i(t)$, where $\pi_i(t;a_i)\equiv\frac{\pi_i(t)}{\pi_i(a_i)}$ is the probability of surviving to age t conditional on being alive at age a_i.

the World Population Prospects and the QALE that we compute. Figure B1 (in Appendix B) depicts the age-specific quality-adjusted remaining life expectancies in the various country groups. To apply *ex post* prioritarian weights, we need information on past quality of life. Figure B2 depicts the number of years spent without a disability for an individual that dies at different ages, again assuming that one minus the disability index approximates the average health status of individuals.[22] The QALE used in the *ex ante* prioritarian weights is the sum of age-specific remaining QALE and age-specific past quality of life (Figure B3).

Furthermore, we assume that life expectancy differs by socioeconomic status. We assume that the survival curves from the 2022 World Population Prospects apply to individuals with median income and adjust the mortality rates for the other income quintiles using scaling factors. To determine the scaling factors, we use two studies. For the high-income category, the scaling factors are loosely based upon Chetty *et al.* (2016, Figure 2), who estimate life expectancy at age 40 in the United States for different income groups. For low- and middle-income countries, we use results from Asaria *et al.* (2019, Table 1), who estimate the distribution of life expectancy at birth by income quintile in India.[23] The scaling factors are such that the ratio between life expectancy for a given quintile and life expectancy for the median individual are approximately equal to the ratios estimated in the two studies.[24] Table 2 reports the resulting life expectancy at birth by income quintile and country income group. Lacking data on the distribution of disabilities by socioeconomic status, we assume that survival rates are income-dependent, while the health indexes are not. As a consequence, remaining quality-adjusted life expectancy at a given age is income-dependent, but past quality of life for someone who dies at a given age is not income-dependent. Table 2 reports the resulting QALEs at birth by income quintile and country income group.

To compute the weights, we also need to define the standardization method (i.e. the denominator). In the utilitarian case, we simply use median income for each country type. In the *ex ante* prioritarian case, we use median income and define median QALE as the value for an individual with median income at median age. Since lower-income countries have a younger population (Table 2), the reference level differs across country income groups. The *ex post* prioritarian weights are also computed with respect to the realized quality of life of an individual at median age.

Table 3 summarizes the weights by country type, age, and income quintile. It is worthwhile to recall that these are the weights to be used in the assessment of national policies. Were one interested in global policies, the weights should be rescaled using a

---

[22] Based on the disability data gathered from GBD (GBD 2019 Diseases and Injuries Collaborators 2020), individuals' realized health profile does not differ much across country income groups, that is, a high-income-country individual who dies at age 80 experiences approximately the same lifetime health profile as a low-income-country individual who dies at age 80. Of course, a larger proportion reaches older ages in higher-income countries than in lower-income countries.

[23] Life expectancies at birth are, respectively, equal to 65.1 years (1st income quintile), 67.0 years (2nd income quintile), 67.9 years (3rd income quintile), 69.6 years (4th income quintile), and 72.7 years (5th income quintile).

[24] For low-income countries, the scaling factors are, respectively, 1.10 (1st income quintile), 1.03 (2nd income quintile), 0.95 (4th income quintile), and 0.87 (5th income quintile). For lower-middle-income countries, the scaling factors are 1.12, 1.03, 0.94, and 0.84. For upper-middle-income countries, the scaling factors are 1.2, 1.05, 0.91, and 0.76. For high-income countries, the scaling factors are those suggested by Adler (2017) based on Chetty et al. (2016): 1.5, 1.2, 0.9, and 0.7. Consistent with Chetty et al. (2016), this implies a 10-year age difference in life expectancy between the richest and the poorest quintiles.

Table 3. *Utilitarian weights,* ex ante *prioritarian weights by current age, and* ex post *prioritarian weights by age of death*

Subgroups	Utilitarian weight	*Ex ante* prioritarian weight ($\gamma = 1$)			*Ex ante* prioritarian weight ($\gamma = 2$)			*Ex post* prioritarian weight ($\gamma = 1$)			*Ex post* prioritarian weight ($\gamma = 2$)		
		Young	Middle aged	Old	Young	Middle aged	Old	Young	Middle aged	Old	Young	Middle aged	Old
Low-income country													
1st income quintile	3.4	7.2	6.9	6.2	15.6	14.2	11.3	21.9	2.7	1.8	142.6	2.2	1.0
2nd income quintile	1.6	2.0	1.9	1.7	2.5	2.3	1.9	6.3	0.8	0.5	24.2	0.4	0.2
3rd income quintile	1.0	1.0	1.0	0.9	1.0	0.9	0.8	3.2	0.4	0.3	9.9	0.2	0.1
4th income quintile	0.6	0.5	0.5	0.4	0.4	0.4	0.3	1.6	0.2	0.1	4.3	0.1	0.03
5th income quintile	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.4	0.1	0.03	0.8	0.01	0.01
Lower-middle income country													
1st income quintile	3.4	4.3	4.2	3.8	5.5	5.1	4.3	18.3	2.2	1.5	98.3	1.4	0.7
2nd income quintile	1.6	1.8	1.7	1.6	2.0	1.9	1.6	7.8	0.9	0.7	36.8	0.6	0.3
3rd income quintile	1.0	1.0	1.0	0.9	1.0	1.0	0.8	4.4	0.5	0.4	19.7	0.3	0.1
4th income quintile	0.6	0.6	0.6	0.5	0.5	0.5	0.4	2.5	0.3	0.2	10.5	0.2	0.1
5th income quintile	0.2	0.2	0.2	0.1	0.1	0.1	0.1	0.7	0.1	0.1	2.9	0.04	0.02
Upper-middle income country													
1st income quintile	3.6	4.2	4.1	3.9	4.9	4.7	4.2	23.9	2.8	1.9	159.4	2.3	1.0
2nd income quintile	1.6	1.8	1.7	1.6	1.9	1.8	1.6	10.2	1.2	0.8	63.2	0.9	0.4
3rd income quintile	1.0	1.0	1.0	0.9	1.0	1.0	0.9	6.0	0.7	0.5	35.7	0.5	0.2
4th income quintile	0.6	0.6	0.6	0.5	0.6	0.5	0.5	3.4	0.4	0.3	19.9	0.3	0.1
5th income quintile	0.2	0.2	0.2	0.2	0.2	0.2	0.1	1.1	0.1	0.1	5.8	0.1	0.04
High-income country													
1st income quintile	3.3	3.7	3.7	3.4	4.2	4.1	3.6	19.6	2.7	1.8	117.3	2.3	1.0
2nd income quintile	1.5	1.6	1.6	1.5	1.7	1.7	1.5	8.8	1.2	0.8	51.1	1.0	0.5
3rd income quintile	1.0	1.0	1.0	1.0	1.0	1.0	0.9	5.7	0.8	0.5	32.4	0.6	0.3

Table 3. _Continued_

Subgroups	Utilitarian weight	*Ex ante* prioritarian weight ($\gamma = 1$)			*Ex ante* prioritarian weight ($\gamma = 2$)			*Ex post* prioritarian weight ($\gamma = 1$)			*Ex post* prioritarian weight ($\gamma = 2$)		
		Young	Middle aged	Old	Young	Middle aged	Old	Young	Middle aged	Old	Young	Middle aged	Old
4th income quintile	0.7	0.6	0.6	0.6	0.6	0.6	0.6	3.7	0.5	0.4	20.7	0.4	0.2
5th income quintile	0.3	0.3	0.3	0.2	0.2	0.2	0.2	1.5	0.2	0.1	8.4	0.2	0.1

Notes: The age brackets considered in the table are: 0–19 (young), 20–64 (middle-aged), and 65+ (old). In the *ex ante* prioritarian case, we report the weights assigned to an individual at average age within a given age bracket. For low-income countries, the average ages are 9 (0–19), 36 (20–64), and 72 (65+); for lower-middle-income countries, the average ages are 9 (0–19), 38 (20–64), and 73 (65+); for upper-middle-income countries, the average ages are 10 (0–19), 41 (20–64), and 73 (65+); for high-income countries, the average ages are 10 (0–19), 42 (20–64), and 75 (65+). Similarly, in the *ex post* prioritarian case, we use the average age of death for people who die in a given age group. These correspond to 5 (0–19), 49 (20–64), and 78 (65+) for low-income countries; 5 (0–19), 51 (20–64), and 78 (65+) for lower-middle-income countries; 5 (0–19), 52 (20–64), and 81 (65+) for upper-middle-income countries; 6 (0–19), 53 (20–64), and 83 (65+) for high-income countries.

common denominator to reflect cross-country differences. The utilitarian weights are similar across countries. This is because the relative distribution of income does not differ much across country income categories. The distribution of income in the high-income countries is slightly more equal, which explains the lower weight to individuals in the first income quintile. For lack of suitable data, we assumed that income is constant across age groups. Thus, the utilitarian weights are age-independent.

Ex ante prioritarian weights assign considerably more value to costs and benefits experienced by low-income individuals. Within an income quintile, they assign slightly more weight to younger people than to older people. Individuals in low-income quintiles are among the worst-off from a lifetime perspective because they have low income and low life expectancy (due to the income gradient of mortality). Within an income bracket, younger people receive a slightly larger weight because they face a larger lifetime mortality risk than older people (i.e., they have lower chances to reach old age than people who are closer to old age). This is particularly true in lower-income countries, with high infant and childhood mortality rates. Increasing the inequality aversion parameter increases the weight attached to impacts experienced by young and low-income individuals.

As in the *ex ante* case, *ex post* prioritarian weights assign more value to impacts experienced by low-income individuals than comparable impacts experienced by high-income individuals. In addition, *ex post* prioritarianism attaches very large weights to those who die young and poor. For example, let us consider a low-income setting with inequality aversion equal to 1. For each income quintile, we find that the weight attached to an individual who dies young is more than 10 times larger than the weight attached to an individual who dies at old age. The weight attached to a poor (1st income quintile) individual who dies young is about 700 times larger than the weight attached to a rich (5th income quintile) individual who dies at old age. This result is driven by the fact that *ex post* prioritarian weights depend on *realized* income, longevity, and quality of life. Thus, from an *ex post* prioritarian point of view, it is a tragedy if the individual dies very young when the rest of the population enjoys a longer lifespan. The larger is the inequality aversion parameter γ, the larger is the weight assigned to costs and benefits experienced by individuals who die young and poor in the absence of the intervention. It is worthwhile to stress that this result hinges on the well-being concept we use, lifetime well-being, which increases in longevity.

Tables C1–C4 in Appendix C report *ex ante* and *ex post* prioritarian weights for each age integer. Tables C5 and C6 report the expected *ex post* prioritarian weights conditional on current age. Table B1 in Appendix B reports the expected *ex post* prioritarian weights by age group, income quintile, and country income group.

4. Stylized example

To clarify the application of equity-sensitive BCA, we present a stylized example. The data concerning the costs and benefits of the interventions are fictional.

Suppose we want to determine the value of alternative policies in a low-income country context. We assume the interventions last 1 year and have only short-term mortality and income impacts.[25] We furthermore assume that the conventional BCA provides data on the

[25] For interventions with longer-term impacts than in this example, the weights would be adjusted to account for economic growth and for changes in population health.

 https://doi.org/10.1017/bca.2024.28

distribution of benefits and costs by socioeconomic status. We also assume that the population affected by the policy amounts to 30 million individuals and that the distribution of age by income is constant.[26] Based on the survival curves derived from the World Population Prospects for low income-countries, 52.4 % of the population is aged 0–19, 44.4 % is aged 20–64, and 3.2 % is aged 65+.

We consider four different interventions, summarized in Table 4. All four interventions reduce mortality risk by 1-in-100,000 (i.e. they save 300 people) and cost $100 million. The risk reductions are assumed to be concentrated on the young. However, the interventions differ in how the health benefits and the policy costs are distributed across income quintiles. Intervention A uniformly reduces the mortality risk for all individuals, that is, it will save 60 young people in each income quintile. The costs of intervention A are also uniformly distributed across the population: Everyone pays an equal share of $100 million or $3.3 per

Table 4. *Distribution of deaths averted and policy costs*

	Policy A		Policy B		Policy C		Policy D	
Subgroup	Deaths averted	Per capita costs	Deaths averted	Per capita costs	Deaths averted	Per capita costs	Deaths averted	Per capita costs
Age 0–19								
1st income quintile	60	$3.3	60	$1.1	100	$3.3	100	$1.1
2nd income quintile	60	$3.3	60	$1.8	80	$3.3	80	$1.8
3rd income quintile	60	$3.3	60	$2.5	60	$3.3	60	$2.5
4th income quintile	60	$3.3	60	$3.5	40	$3.3	40	$3.5
5th income quintile	60	$3.3	60	$7.8	20	$3.3	20	$7.8
Age 20–64								
1st income quintile	0	$3.3	0	$1.1	0	$3.3	0	$1.1
2nd income quintile	0	$3.3	0	$1.8	0	$3.3	0	$1.8
3rd income quintile	0	$3.3	0	$2.5	0	$3.3	0	$2.5
4th income quintile	0	$3.3	0	$3.5	0	$3.3	0	$3.5
5th income quintile	0	$3.3	0	$7.8	0	$3.3	0	$7.8
Age 65+								
1st income quintile	0	$3.3	0	$1.1	0	$3.3	0	$1.1
2nd income quintile	0	$3.3	0	$1.8	0	$3.3	0	$1.8
3rd income quintile	0	$3.3	0	$2.5	0	$3.3	0	$2.5
4th income quintile	0	$3.3	0	$3.5	0	$3.3	0	$3.5
5th income quintile	0	$3.3	0	$7.8	0	$3.3	0	$7.8

[26] Since higher-income people are expected to live longer, the distribution of age by income group may be skewed, with more high-income individuals among the old than among the young. Since we are considering broad age groups, we disregard this effect.

capita. Intervention B uniformly reduces the mortality risk for all young individuals as for intervention A. However, the distribution of costs is proportional to income. Interventions C and D benefit primarily poor young individuals, that is, they are expected to avert more deaths among low-income individuals than among high-income individuals. However, they differ in terms of distribution of costs. The costs of intervention C are uniformly distributed across income quintiles, while the costs of intervention D are distributed proportional to income.

Using the previous notation, $N_i \Delta p_i$ is the number of deaths averted among individuals of type i (characterized by a specific age and a specific income quintile), and Δy_i is the per-capita cost of the intervention. Those data are included in Table 4. To apply the conventional BCA rule (1), the only missing ingredient is the individual-specific value-per-statistical-life VSL_i. We assume that the ratio of VSL to income is constant and equal to 160 times GNI per capita for each income quintile (i.e., we set the income elasticity of VSL to 1). To adjust for age differences, we derive income-specific value-per-statistical-life-year (VSLY) by dividing the VSL by undiscounted remaining quality-adjusted life expectancy at the average age of the population in the country group (the latter proxied by half the quality adjusted life expectancy at birth). The resulting VSLY is then multiplied by age- and income-specific remaining quality-adjusted life expectancy to determine the age-and income-specific VSLs.

Table 5 summarizes the VSL estimates and the net benefits of the interventions by income quintile and age group. The net benefits are the difference between the monetary value of deaths averted among individuals from a given quintile and the total costs paid by them, that

Table 5. *Value-per-statistical-life and net benefits by demographic group*

Income quintile	VSL	Net benefits			
		Policy A	Policy B	Policy C	Policy D
Age 0–19					
1st	\$185,100	\$0.6 million	\$7.6 million	\$8.0 million	\$15.0 million
2nd	\$299,500	\$7.5 million	\$12.4 million	\$13.5 million	\$18.4 million
3rd	\$414,100	\$14.4 million	\$17.1 million	\$14.4 million	\$17.1 million
4th	\$590,300	\$24.9 million	\$24.4 million	\$13.1 million	\$12.6 million
5th	\$1,308,200	\$68.0 million	\$54.0 million	\$15.7 million	\$1.6 million
Age 20–64					
1st	\$105,000	−\$8.9 million	−\$2.9 million	−\$8.9million	−\$2.9 million
2nd	\$172,200	−\$8.9 million	−\$4.8 million	−\$8.9 million	−\$4.8 million
3rd	\$239,800	−\$8.9 million	−\$6.6 million	−\$8.9 million	−\$6.6 million
4th	\$346,600	−\$8.9 million	−\$9.4 million	−\$8.9 million	−\$9.4 million
5th	\$783,600	−\$8.9 million	−\$20.8 million	−\$8.9 million	−\$20.8 million
Age 65+					
1st	\$21,500	−\$0.6 million	−\$0.2 million	−\$0.6 million	−\$0.2 million
2nd	\$39,500	−\$0.6 million	−\$0.3 million	−\$0.6 million	−\$0.3 million
3rd	\$58,400	−\$0.6 million	−\$0.5 million	−\$0.6 million	−\$0.5 million
4th	\$94,700	−\$0.6 million	−\$0.7 million	−\$0.6 million	−\$0.7 million
5th	\$243,100	−\$0.6 million	−\$1.5 million	−\$0.6 million	−\$1.5 million

Table 6. *Results of conventional BCA compared to equity-sensitive BCA*

Approach	Value (million $)				Recommended intervention
	Policy A	Policy B	Policy C	Policy D	
Conventional BCA	67.8	67.8	17.1	17.1	A and B equally
Utilitarianism	−8.9	51.0	9.5	69.4	Policy D, not policy A
Ex ante prioritarianism ($\gamma = 1$)	−43.6	70.8	10.1	124.5	Policy D, not policy A
Ex ante prioritarianism ($\gamma = 2$)	−108.8	115.0	12.8	236.6	Policy D, not policy A
Ex post prioritarianism ($\gamma = 1$)	453.9	489.5	614.6	650.2	Policy D
Ex post prioritarianism ($\gamma = 2$)	2,436.4	2,462.2	3,542.4	3,568.2	Policy D

is, $N_i(\Delta p_i VSL_i - \Delta y_i)$. Not everyone gains from the interventions. The distribution of net benefits is policy dependent. All things considered, policy D is the most progressive since within each age group the lowest income individuals are expected to gain the most (or lose the least).

To apply the equity-sensitive BCA rules (4), (6), and (11), we supplement the data on benefits and costs by age and income quintile with the weights to be applied in a low-income context. Those weights are provided in Tables 3 and B1 (as stressed earlier, in the *ex post* prioritarian case we use the expected weights $\overline{w}_i^{EPP}(a_i)$ since the intervention costs are independent of the age of death).

Table 6 summarizes the results of conventional BCA and equity-sensitive BCA. Conventional BCA sums the net benefits across the population. According to conventional BCA, policies A and B are equally desirable and preferable to policies C and D. Thus, the distribution of policy costs has no impact on the BCA outcome, and policies that primarily benefit the disadvantaged are valued less than similar policies with a more regressive distribution of health benefits.[27]

In this specific example, all equity-sensitive BCAs agree on the ranking of policies and recommend policy D over the alternatives. Moreover, both utilitarianism and *ex ante* prioritarianism find that policy A is not beneficial, in stark contrast with the BCA outcome. *Ex post* prioritarianism recommends policy D, but attaches positive value also to policy A. The difference between *ex ante* and *ex post* prioritarianism is driven by the weight attached to the young. From an *ex post* point of view, avoiding deaths at a young age deserves priority. Since the benefits of the policy are concentrated among the young, in this specific example, *ex post* prioritarianism concludes that all proposed policies are worthwhile.

[27] The benefits of policies A and B exceed those of policies C and D because we use a VSL that increases with individuals' income. If a common VSL were used for all young people, the net benefits of all four policies would be identical.

5. Summary and conclusions

While conventional BCA has many strengths, it also has important limitations. Chief among these is its inattention to the distributional equity of the impacts. To measure individual welfare, it relies on individuals' willingness to exchange their own income for the outcomes they experience. To measure societal welfare, it relies on summing these values across individuals. Many alternative approaches have been proposed to address these limitations, but none is yet widely used. A major barrier to their application is the difficulty of implementation.

To encourage more attention to the distributional implications of interventions, we derive multiple sets of weights to be used in conjunction with the results of conventional BCAs to improve the information available on distributional equity globally. We discuss how the weights can be estimated using publicly available datasets and results from the literature. We derive both utilitarian and prioritarian weights. The utilitarian weights account for the greater value of money to a poor person than a rich person. The prioritarian weights account also for preferences for prioritizing increases in welfare for those who are worse off. The worse off include individuals with low socioeconomic status, poor health, or low life expectancy.

The weights that we derive are independent of the specific intervention under analysis. However, they depend on (i) the characteristics of the population affected by the intervention, and (ii) the chosen well-being metric and SWF. We derive weights for representative country profiles based on the income groups defined by the World Bank using publicly available data. As long as the country profile reasonably approximates the population affected by the intervention, the estimated weights can be directly applied to the results of a conventional BCA.

For the theoretical underpinnings of the estimated weights, we measure well-being with a vNM utility function and we adopt an Atkinson SWF whose concavity spans from utilitarianism to prioritarianism (in our example, but could also address egalitarianism). These assumptions are supported by previous studies. Future research should investigate alternative measures of well-being and alternative SWFs and the sensitivity of the estimated weights to these normative choices. In addition, because of data constraints, we made some further simplifying assumptions, including well-being depends only on three attributes, income, health, and longevity; lifetime well-being is time separable; health and income are complements; and the distribution of income is age-independent. These assumptions could be relaxed in future research.

In our setup, utilitarian weights account for differences in income across the population. Prioritarian weights account for differences in income and differences in expected or realized quality-adjusted longevity. Since young individuals have a lower expected or realized longevity, prioritarianism tends to attach larger weights to young people than older people.[28] *Ex post* prioritarianism confers extremely large weight to individuals who die very young on the grounds that dying young is a tragedy from a lifetime perspective. This result runs against the argument that, especially in resource-constrained settings, saving an adult life is more highly valued due to the economic support provided by working-age individuals. In our framework, the larger contribution of adult individuals could be captured by including the

[28] Note that our analysis assumes no change in age-specific mortality rates over time. If anticipated decreases in future mortality rates were incorporated, the priority given to younger individuals would decrease because their life expectancy would increase.

positive externalities associated with saving an adult life (e.g. in the form of increased production or of improved well-being of family members). Similarly, in our stylized example, we disregard the non-market productive contribution of older adults, for example, through childcare, volunteering, or informal care for sick family members. There is also a large debate on the value of preserving the life of newborns and infants. Our framework is not able to capture all the nuances of such a debate.

The choice between utilitarianism and prioritarianism (and the degree of priority to the worse off) is ultimately an ethical choice. In our framework, this was captured by a single parameter γ representing the degree of inequality aversion. We computed weights under three values: $\gamma = 0$ (utilitarianism), $\gamma = 1$, and $\gamma = 2$. The larger is γ, the larger is the priority to the worse off. We avoided making any judgement about the most appropriate value for γ, and we proposed weights for multiple values of γ. The literature on social preferences can provide guidance on the inequality aversion attitudes most frequently held by individuals (Venmans & Groom, 2021; Schokkaert & Tarroux, 2022).

Supplementary material. The supplementary material for this article can be found at http://doi.org/10.1017/bca.2024.28.

Acknowledgments. The authors thank Matthew Adler, Richard Cookson, Marc Fleurbaey, Ole F. Norheim, J.P. Sevilla, and Brad Wong for their thoughtful comments on previous drafts of this article and many helpful discussions. The authors are also grateful to Wafa Aftab, David Canning, and Christian Krekel for valuable inputs to this project. This work was funded in part by the Bergen Centre for Ethics and Priority Setting, through grants from the Norwegian Agency for Development Cooperation and the Norwegian Research Council. It builds on a 2022 conference held by the Brocher Foundation, "Healthy, Wealthy, and Wise - The Ethics of Health Valuation," organized by Nir Eyal (Rutgers University), Samia Hurst (University of Geneva), Lisa A. Robinson (Harvard University), and Daniel Wikler (Harvard University). The authors thank the Brocher Foundation and conference participants for their support and helpful comments. This special issue was supported by the Brocher Foundation, with supplemental funding from the Rutgers University Center for Population-Level Bioethics, the University of Geneva Institute for Ethics, History, and the Humanities, and the University of Bergen Centre for Ethics and Priority Setting. More information on the Brocher Foundation is available here: https://fondation-brocher.ch/.

Competing interest. The authors declare none.

References

Acland, Daniel and David H. Greenberg. (2023). The Elasticity of Marginal Utility of Income for Distributional Weighting and Social Discounting: A Meta-Analysis. *Journal of Benefit-Cost Analysis 14*(2): 386–405.

Adler, Matthew D. (2013) Happiness Surveys and Public Policy: What's the Use? *Duke Law Journal 62*: 1509–1601.

Adler, Matthew D. (2016). Benefit-Cost Analysis and Distributional Weights: An Overview. *Review of Environmental Economics and Policy 10*(2): 264–285.

Adler, Matthew D. (2017). A Better Calculus for Regulators: From Cost-Benefit Analysis to the Social Welfare Function. Duke Law School Public Law and Legal Theory Working Paper No. 2017–2019. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2923829

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

Adler, Matthew D. (2024). Social Welfare Functions and Health Policy: A New Approach. *Journal of Benefit-Cost Analysis, First View*. 1–26.

Adler, Matthew D. and Koen Decancq. (2022). Well-Being Measurement. In Adler, Matthew D. and Norheim, Ole F. (Eds.) *Prioritarianism in Practice*, pp. 128–171. New York: Cambridge University Press.

Adler, Matthew D. and Ole F. Norheim (Eds.) (2022). *Prioritarianism in Practice*. New York: Cambridge University Press.

Adler, Matthew D. and Marc Fleurbaey (Eds.) (2016). *Oxford Handbook of Well-Being and Public Policy*. New York: Oxford University Press.

Alesina, Alberto and George-Marios Angeletos. (2005). Fairness and Redistribution. *American Economic Review 95*(4): 960–980.

Almås, Ingvild, Alexander W. Cappelen and Bertil Tungodden. (2020). Cutthroat Capitalism versus Cuddly Socialism: Are Americans More Meritocratic and Efficiency-Seeking than Scandinavians? *Journal of Political Economy 128*(5): 1753–1788.

Arrow, K. J. (1963). *Social Choice and Individual Values*. New York: John Wiley.

Asaria, Miqdad, Sumit Mazumdar, Samik Chowdhury, Papiya, Mazumdar, Abhiroop, Mukhopadhyay, Indrani, Gupta, (2019). Socioeconomic Inequality in Life Expectancy in India. *BMJ Global Health 4*: e001445.

Bergstrom, Katy and William Dodds. (2023). Using Schooling Decisions to Estimate the Elasticity of Marginal Utility of Consumption. *Journal of Public Economics 224*: 104948.

Blackorby, Charles and David Donaldson. (1990). A Review Article: The Case Against the Use of the Sum of Compensating Variations in Cost-Benefit Analysis. *Canadian Journal of Economics 23*(3): 471–494.

Bommier, Antoine and Bertrands Villeneuve. (2012). Risk Aversion and the Value of Risk to Life. *The Journal of Risk and Insurance 79*(1): 77–103.

Bosmans, Kristof. (2007). Extreme Inequality Aversion Without Separability. *Economic Theory 32*: 589–594.

Chetty Ray, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron and David Cutler. (2016). The Association between Income and Life Expectancy in the United States, 2001-2014. *Journal of the American Medical Association 315*(16): 1750–1766.

Clinton, William J. (1993). Executive order 12866: Regulatory Planning and Review. *Federal Register 58*: 51735–51744.

Cookson, Richard, Ieva Skarda, Owen Cotton-Barratt, Matthew D. Adler, Miqdad Asaria and Toby Ord. (2021). Quality Adjusted Life Years Based on Health and Consumption: A Summary Wellbeing Measure for Cross-Sectoral Economic Evaluation. *Health Economics*, *30*: 70–85.

Decancq, Koen, Marc Fleurbaey and Erik Schokkaert. (2015). Inequality, Income, and Well-Being. In Atkinson, Anthony B. and Bourguignon, François (Eds.) *Handbook of Income Distribution*, Vol. *2*, pp. 67–140. Amsterdam: Elsevier.

Drèze, Jean and Nicholas Stern. (1987). The Theory of Cost Benefit Analysis. In Alan J. Auerbach, and Martin, Feldstein, (Eds.) *Handbook of Public Economics*, Vol. *2*, pp. 909–989. Elsevier.

Finkelstein, Amy, Erzo F.P. Luttmer and Matthew J. Notowidigdo. (2013). What Good is Health Without Wealth? The Effect of Health on the Marginal Utility of Consumption. *Journal of the European Economic Association 11*(1): 221–258.

Fleurbaey, Marc and Gregory Ponthiere. (2023). Measuring Well-Being and Lives Worth Living. *Economic Theory 75*: 1247–1266.

Fleurbaey, Marc and Rossi Abi-Rafeh. (2016). The Use of Distributional Weights in Benefit–Cost Analysis: Insights from Welfare Economics. *Review of Environmental Economics and Policy 10*(2): 286–307.

Frijters, Paul, Christian Krekel, Raúl Sanchis and Ziggi Ivan Santini. (2024). The WELLBY: A New Measure of Social Value and Progress. *Humanities and Social Sciences Communication 11*: 736.

Gandelman, Néstor and Rubén Hernández-Murillo. (2014). Risk Aversion at the Country Level. Working Paper 2014–005B, Federal Reserve Bank of St. Louis. https://files.stlouisfed.org/files/htdocs/wp/2014/2014-005.pdf

Graham, Carol. (2016). Unequal Life Chances and Choices: How Subjective Well-Being Metrics Can Inform Benefit-Cost Analysis. *Journal of Benefit-Cost Analysis 7*(1): 121–146.

Groom, Ben and David Maddison Pr. (2019). New Estimates of the Elasticity of Marginal Utility for the UK. *Environmental and Resource Economics 72*: 1155–1182.

Hammitt, James K., Peter Morfeld, Jouni T. Tuomisto and Thomas C. Erren. (2020) Premature Deaths and Statistical Lives: Identification, Quantification, and Valuation of Mortality Risks. *Risk Analysis 40*(4): 674–695.

HM Treasury. (2021a). Wellbeing Guidance for Appraisal: Supplementary Green Book Guidance. https://www.gov.uk/government/publications/green-book-supplementary-guidance-wellbeing

HM Treasury. (2021b). Wellbeing Discussion Paper: Monetisation of Life Satisfaction Effect Sizes: A Review of Approaches and Proposed Approach. https://www.gov.uk/government/publications/green-book-supplementary-guidance-wellbeing

HM Treasury. (2022). The Green Book: Central Government Guidance on Appraisal and Evaluation. https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government

Lakdawalla, Darius N. and Charles E. Phelps. (2020). Health Technology Assessment with Risk Aversion in Health. *Journal of Health Economics*, *72*: 102346.

Layard, Richard, Guy Mayraz and Stephen Nickell. (2008). The Marginal Utility of Income. *Journal of Public Economics 92*: 1846–1857.

Layard, Richard and Jan-Emmanuel De Neve. (2023). *Wellbeing: Science and Policy*. Cambridge: Cambridge University Press.

Murphy, Kevin M. and Robert H. Topel. (2006). The Value of Health and Longevity. *Journal of Political Economy 114*(5): 871–904.

Nordhaus, William. (2008). *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven: Yale University Press.

Robinson, Lisa A. and James K. Hammitt. (2011). Behavioral Economics and the Conduct of Benefit-Cost Analysis: Towards Principles and Standards. *Journal of Benefit-Cost Analysis 2*(2): 1–51.

Robinson, Lisa A., James K. Hammitt and Richard J. Zeckhauser. (2016). Attention to Distribution in U.S. Regulatory Analyses. *Review of Environmental Economics and Policy 10*(2): 308–328.

Robinson, Lisa A., James K. Hammitt, Michele Cecchini, Kalipso Chalkidou, Karl Claxton, Maureen Cropper, Patrick Hoang-Vu Eozenou, David de Ferranti, Anil B. Deolalikar, Frederico Guanais, Dean T. Jamison, Soonman Kwon, Jeremy A. Lauer, Lucy O'Keeffe, Damian Walker, Dale Whittington, Thomas Wilkinson, David Wilson, and Brad Wong. (2019a). *Reference Case Guidelines for Benefit-Cost Analysis in Global Health and Development*. Funded by the Bill & Melinda Gates Foundation. https://www.hsph.harvard.edu/bcaguidelines/wp-content/uploads/sites/2447/2019/05/BCA-Guidelines-May-2019.pdf.

Robinson, Lisa A., William J. Raich, James K. Hammitt and Lucy O'Keeffe. (2019b). Valuing Children's Fatality Risk Reductions. *Journal of Benefit-Cost Analysis 10*(2): 156–177.

Robson, Matthew, Miqdad Asaria, Richard Cookson, Aki Tsuchiya and Shehzad Ali. (2017). Eliciting the Level of Health Inequality Aversion in England. *Health Economics 26*(10): 1328–1334.

Sen, Amartya K. (1999). *Commodities and Capabilities*. Delhi: Oxford University Press.

Schokkaert, Erik and Benoît Tarroux. (2022). Empirical Research on Ethical Preferences: How Popular is Prioritarianism? In Adler, Matthew D., and Ole F. Norheim (Eds.), *Prioritarianism in Practice*, pp. 459–517. New York: Cambridge University Press.

Sunstein, Cass R. (2016). Cost-Benefit Analysis, Who's Your Daddy? *Journal of Benefit-Cost Analysis 7*(1): 107–120.

Tengstam, Sven. (2014). Disability and Marginal Utility of Income: Evidence from Hypothetical Choices. *Health Economics*, *23*(3): 268–282

Tuomala, Matti. (2016). *Optimal Redistributive Taxation*. New York: Oxford University Press.

United Nations Department of Economic and Social Affairs, Population Division. (2022). World Population Prospects: The 2022 Revision. https://population.un.org/wpp/

U.S. Office of Management and Budget (OMB). (2023a). Circular A4: Regulatory Analysis https://www.whitehouse.gov/wp-content/uploads/2023/11/CircularA-4.pdf

U.S. Office of Management and Budget (OMB). (2023b). OMB Circular No. A-4: Explanation and Response to Public Input https://www.whitehouse.gov/wp-content/uploads/2023/11/CircularA-4Explanation.pdf

U.S. Office of Management and Budget (OMB). (2023c). OMB Circular No. A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs https://www.whitehouse.gov/wp-content/uploads/2023/11/CircularA-94.pdf

Venmans, Frank and Ben Groom. (2021). Social Discounting, Inequality Aversion, and the Environment. *Journal of Environmental Economics and Management 109*: 102479.

GBD 2019 Diseases and Injuries Collaborators (2020). Global Burden of 369 Diseases and Injuries in 204 Countries and Territories, 1990-2019: A Systematic Analysis for the Global Burden of Disease Study 2019. *The Lancet Global Health Metrics 396*(10258): 1204–1222.

World Bank. (2022). Poverty and Inequality Platform, Key Inequality Indicators, Decile shares, Selected Economies and Years. https://pip.worldbank.org/key-inequality-indicators

World Bank. (2024). Gross National Income Per Capita, PPP (constant 2017 international dollars) https://data.worldbank.org/indicator/NY.GNP.PCAP.PP.KD

Cite this article: Ferranna, Maddalena, James K. Hammitt, and Lisa A. Robinson. 2024. "From Benefit–Cost Analysis to Social Welfare: A Pragmatic Approach." *Journal of Benefit-Cost Analysis* 15: 84–109, doi:10.1017/bca.2024.28