

Benefit–Cost Analysis and Distributional Weights: An Overview

Matthew D. Adler*

Introduction

Benefit–cost analysis (BCA)¹ evaluates governmental policies by summing individuals' monetary equivalents, and is insensitive to distributional concerns. BCA does not take into account whether those made better off by a policy have higher or lower incomes, or higher or lower levels of nonincome welfare-relevant attributes (e.g., health), than those made worse off.

Arguably, distributional considerations should be incorporated into BCA via “distributional weights.” A scholarly literature dating from the 1950s endorses distributional weights and analyzes how to specify them (Meade 1955, chap. 2; Weisbrod 1968; Dasgupta and Pearce 1972; Dasgupta, Sen, and Marglin 1972; Little and Mirrlees 1974; Squire and van der Tak 1975; Boadway and Bruce 1984, pp. 271–81; Brent 1984; Ray 1984; Drèze and Stern 1987; Drèze 1998; Cowell and Gardiner 1999; Yitzhaki 2003; Johansson-Stenman 2005; Creedy 2006; Liu 2006; Fleurbaey et al. 2013; Boadway 2016).

Distributional weights were adopted, for a time, at the World Bank (Little and Mirrlees 1994). They are currently recommended by the UK's official BCA guidance document (HM Treasury 2003, pp. 91–94). However, it appears that distributional weights have rarely if ever been used by BCA practitioners in the U.S. government, and the parallel U.S. guidance document does not mention them (Office of Management and Budget 2003).

This article, which is part of a symposium on Distributional Considerations in Benefit–Cost Analysis, provides an introduction to distributional weights.² The fulcrum for my discussion will be the concept of the social welfare function (SWF). The SWF is a fundamental construct in many areas of welfare economics, including optimal tax theory, growth theory, and analysis of climate change. BCA with distributional weights, in turn, is a practicable method for implementing an SWF. This is the view of BCA running through the literature on distributional weights, and is presented here.

This account of BCA is quite different from the familiar view that sees BCA as a tool for implementing the criterion of Kaldor–Hicks efficiency (potential Pareto superiority). The

*Richard A. Horvitz Professor of Law and Professor of Economics, Philosophy, and Public Policy, Duke University. Many thanks to Marc Fleurbaey, James Hammitt, Alex Pfaff, and Nicolas Treich for their comments.

¹Some prefer the term cost–benefit analysis (CBA). In conformity with this journal's style, I use BCA here.

²The other articles in the symposium are Fleurbaey and Rafah (2016), which uses insights from welfare economics to examine how distributional weights can be introduced into benefit–cost analysis, and Robinson, Hammitt, and Zeckhauser (2016), which focuses specifically on the role of distributional considerations in U.S. regulatory analyses.

Kaldor–Hicks criterion has the advantage of avoiding interpersonal comparability, but has various flaws, described in a literature beginning with Scitovsky (1941; see also Gorman 1955; Chipman and Moore 1978; Sen 1979; Boadway and Bruce 1984; Boadway 2016).³

The article is *not* a comprehensive survey of the literature on distributional weights. Rather, it aims to explain the key ideas. I first describe the SWF concept, with a particular focus on two specific SWFs: the utilitarian and isoelastic/Atkinson SWFs. The article then discusses the functional form of weights matching these two SWFs. Next, it provides a concrete example, involving risk-reduction policies and the “value of statistical life.”

The article concludes by considering two objections to weights. One concerns the possibility of interpersonal comparisons given heterogeneous preferences. The second is that distributional considerations are better handled via the combination of the tax system and unweighted BCA.

The discussion aims to be accessible. The fundamentals of distributional weighting are illustrated with a simple, one-period model. Space constraints preclude a treatment of certain additional issues that arise in the intertemporal context, in particular the relation between distributional weights and discounting. Key mathematical formulas are provided in a Technical Appendix. A more rigorous, mathematical analysis of many of the topics discussed in the main text is provided in the online supplementary materials. The reader should consult these materials, along with cited references, as backup for the discussion.

Social Welfare Functions: An Overview

The concept of the SWF originated with work by Bergson and Samuelson. It was reenergized by Sen, in response to Arrow’s impossibility theorem, and was the basis for Mirrlees’ groundbreaking scholarship on optimal tax theory. It now permeates many subdisciplines within economics (although less so governmental practice).⁴

Key Elements of the SWF Framework

The SWF framework has two key elements: an interpersonally comparable utility function, which transforms any given outcome (a possible consequence of policy choice) into a list or “vector” of utility numbers, one for each person in the population; and some rule for ranking these vectors.

To illustrate, imagine that there are three people in the population and two outcomes are being compared. Jim has a particular bundle of attributes in outcome x and a different bundle in outcome y . The same is true of Sue. Laura has the same attributes in both outcomes, and thus is unaffected by the choice between them (her well-being does not change). “Attributes,” here, means the characteristics that determine an individual’s well-being, such as income, health, leisure, the quality of the environment, and so forth.⁵

³The debate about the Kaldor–Hicks criterion is well known and will not be recapitulated here.

⁴See Adler (2012, pp. 79–88), for a summary of scholarly development of the SWF concept.

⁵It is, of course, infeasible for a decision maker and the policy analyst advising her to consider how policies affect the totality of individuals’ attributes. Rather, policy analysis will focus on a subset of attributes: for example, income and leisure; income, health, and leisure; etc.

Our utility function assigns Jim's bundles of attributes in x and y the utility values 10 and 11, respectively; it assigns Sue's bundles the values 30 and 25, respectively; and it assigns 40 to Laura's bundle. Thus outcome x is mapped onto the utility vector (10, 30, 40) and y is mapped onto the vector (11, 25, 40)—with the first entry the utility number for Jim, the second for Sue, and the third for Laura (see table 1).

There are various possible rules for comparing utility vectors (Bossert and Weymark 2004; Adler 2012, chap. 5; Weymark 2016). One such rule is the *leximin* rule: it compares the utility levels of the worst-off individuals; if those are equal, then the second-worst-off; and so forth. Leximin here prefers y , since Jim (the worst off) has utility 11 instead of 10.

A different rule is the utilitarian rule, which sums utilities. Utilitarianism prefers x , since the sum of utilities is 80 rather than 76. Although utilitarianism is sensitive to the distribution of income (given the declining marginal utility of income), utilitarianism does not take account of the distribution of utility itself. One outcome is ranked higher than a second as long as the sum total of utility gains for those better off in the first outcome exceeds the sum total of utility losses for those worse off—regardless of the comparative utility levels of the two groups.

Leximin is sensitive to the distribution of utility. In the case at hand, even though Sue's loss from y is greater than Jim's gain, leximin prefers y . However, leximin is absolutist, in the sense that it is willing to incur arbitrarily large utility losses for better-off individuals in order to realize a utility gain (however small) for an individual who is worse off and would remain so after the gain.

The isoelastic/Atkinson SWF lies in between leximin and utilitarianism. This SWF is parameterized by an inequality aversion parameter γ , which can take any positive value. With γ set to zero, the isoelastic formula becomes utilitarianism. As γ increases, the isoelastic SWF gives increasing priority to utility changes affecting worse-off individuals (those at lower utility levels). In the example in table 1, if γ is less than or equal to approximately 1.6, the isoelastic SWF prefers outcome x to y . Sue stands to lose more utility moving to y (5 units) than Jim stands to gain (1 unit); and with a low value of γ , the isoelastic SWF assigns a greater social value to her loss than to Jim's gain even though Jim is worse off than Sue in both outcomes. With γ greater than 1.6, the isoelastic SWF prefers y to x . The priority given to utility changes affecting those who are worse off is now large enough that moving Jim from 10 to 11 is seen as more socially valuable than avoiding Sue's move from 30 to 25.

The utilitarian and isoelastic SWFs are both popular in the SWF literature, and with good reason. Both satisfy the Pareto principle: if at least one person's utility increases, and no one else's decreases, the value of the SWF increases. These SWFs also satisfy an axiom of anonymity/impartiality, meaning indifference between any given utility vector and all rearrangements of its

Table 1 Outcomes as utility vectors

	Outcomes	
	x	y
Jim	10	11
Sue	30	25
Laura	40	40

component utility numbers. Anonymous/impartial SWFs focus only on the pattern of well-being, and not the identities of the people who end up at particular well-being levels.

Moreover, the utilitarian and isoelastic SWFs are separable—meaning that the ranking of outcomes is not influenced by the utility levels of unaffected people. In the above example, Laura is unaffected. She happens to be at level 40 in both outcomes; but note that the utilitarian SWF would prefer x to y in *any* case where Jim’s and Sue’s utilities are as in table 1 and Laura has the same utility level in the two outcomes, *regardless of what that level is*. Similarly, the isoelastic SWF with γ less than or equal to 1.6 would prefer x to y regardless of Laura’s level, and the isoelastic SWF with γ greater than 1.6 would prefer y to x regardless of Laura’s level. Separability is a big practical advantage in policy analysis—enabling the analyst to focus her efforts on determining the utilities of those whose well-being would be changed by a policy, and not also to worry about how the policy would alter their position relative to the potentially vast number of unaffected.

The leximin SWF is also Paretian, anonymous, and separable, and is popular among some SWF theorists. But it cannot be represented by a mathematical formula, which creates difficulties in mimicking this SWF with distributional weights. Moreover, the leximin SWF can be approximated by the isoelastic SWF with a large value of γ .

Interpersonal Comparisons of Well-Being

In order to achieve a stable ranking of outcomes, all SWFs require some degree of interpersonal comparability of well-being. If we start with a particular rule for ranking utility vectors and then transform the utility vectors associated with outcomes so that intrapersonal comparisons of utility levels, differences, and ratios are preserved but interpersonal comparisons are not, the ranking of outcomes may change as well⁶ (Bossert and Weymark 2004; Weymark 2016).

More specifically, the utilitarian SWF requires interpersonal comparability of well-being differences. Consider table 2, which shows possible renumberings of Jim and Sue’s utility. In the first renumbering, we rescale Jim’s utility by a Jim-specific ratio transformation, and we rescale Sue’s by a Sue-specific ratio transformation. In the second renumbering, we rescale Jim’s and Sue’s utility by a common linear transformation. In the third renumbering, we rescale Jim’s and Sue’s utility by a common ratio transformation.

With the original scheme of utility assignments, Sue’s utility difference between the outcomes is -5 and Jim’s is 1 . The first renumbering changes the relative magnitude of these differences; the second and third renumberings do not. And now observe that the first renumbering alters the utilitarian ranking of the outcomes, while the second and third renumberings do not.

The isoelastic SWF is *more* demanding in terms of interpersonal comparability than the utilitarian SWF. While the utilitarian SWF requires interpersonal comparability of differences, the isoelastic SWF requires interpersonal comparability of levels, differences, and ratios. Note that the second renumbering changes the well-being ratios between Sue and Jim in the two outcomes. Given a particular value of inequality aversion γ , the isoelastic SWF might prefer x to y using the original numbering but not the second renumbering. (This occurs, for example, with $\gamma = 0.5$.) However, the third renumbering preserves well-being ratios. If the isoelastic SWF

⁶More precisely, all ethically *plausible* SWFs require some degree of interpersonal comparability. A “dictatorial” SWF, which focuses solely on the well-being of one particular individual, requires nothing more than intrapersonal well-being information regarding that person.

Table 2 Interpersonal comparisons and the renumbering of utility

	Original			Renumbering 1			Renumbering 2			Renumbering 3		
	x	y	Diff	x	y	Diff	x	y	Diff	x	y	Diff
Jim	10	11	1	200	220	20	2	12	10	30	33	3
Sue	30	25	-5	60	50	-10	202	152	-50	90	75	-15
Sum	40	36		260	270		204	164		120	108	

Notes: In the first renumbering, Jim's utilities are each multiplied by a Jim-specific factor a_{Jim} equaling 20, while Sue's are multiplied by a Sue-specific a_{Sue} equaling 2. In the second renumbering, Jim's and Sue's utilities are subject to a common linear transformation $au + b$, with the common scaling factor a equaling 10 and the constant b equaling -98 . Finally, in the third renumbering, Jim's and Sue's utilities are subject to a common ratio transformation, each multiplied by the common positive number 3. The bold columns labeled "Diff" show the difference between each person's utility in y and x . The last row displays the sum of utilities.

with a particular value of γ prefers one outcome over the other using the original numbering, than it does so with the third renumbering.

The Normative Basis for the SWF

We can now address two related questions. First, what is the normative basis for a determination that one SWF is "better" than another, or that the SWF framework for social decision is "better" than alternative frameworks? Second, what is the basis for assigning utility numbers that are interpersonally comparable?

The view adopted here is that the SWF is a template for ethical/moral preferences (Bergson 1948, 1954; Samuelson 1947, p. 221; Harsanyi 1977, chap. 4). The term "ethical" is more common among economists, "moral" among philosophers, but the two terms are essentially synonyms—denoting a standpoint of impartiality, whereby the decision maker gives equal weight to everyone's interests (or at least the interests of everyone within some population). An SWF constitutes a systematic framework for structuring ethical/moral (henceforth, "moral") preferences: a framework that a decision maker who has adopted the standpoint of impartiality might wish to use in specifying her moral tastes.

Is there some deeper criterion of moral truth that establishes whether someone's moral preferences are "correct"? That is a question of metaethics that is debated by philosophers but is beyond the scope of welfare economics and, in any event, not relevant to the discussion here. Whatever the nature of ultimate moral truth, a decision maker motivated by impartial concern will need to figure out what her moral preferences are. The SWF framework is a plausible format for regimenting those preferences: a format that conforms to various axioms that seem (to those in the SWF tradition) morally very attractive, and that the decision maker may find attractive as well.

While the SWF framework itself is a tool for specifying the moral preferences of some decision maker, the *inputs* for the framework are utility numbers measuring the well-being of everyone within some population (relative to which the decision maker has adopted an attitude of impartiality). Very plausibly, there is a close connection between someone's welfare and that individual's *personal*, that is, self-interested, preferences. Phillip is better off in outcome

x than y if, and only if, Phillip has a personal preference for x (Adler 2012, chap. 3; Fleurbaey et al. 2013).⁷

Moreover, if members of the population have *identical* personal preferences, interpersonal comparisons become straightforward, as we shall now see. My analysis of distributional weights will start with this simplifying assumption of identical personal preferences—postponing until later in the article the question of how the SWF framework should handle heterogeneity of personal preferences.

BCA and Utilitarian Distributional Weights

A simple one-period model will illustrate utilitarian distributional weights. In any given outcome, each individual has a bundle of attributes consisting of the following: her consumption, that is, the total amount of money she expends on marketed goods and services; the market prices she faces, assumed to be the same for all individuals; and her nonconsumption attributes, such as health, leisure, or environmental quality.⁸ Each policy choice, including the status quo choice of inaction, leads for sure to some outcome. There are a fixed number of individuals in the population.⁹

BCA and Monetary Equivalents

Unweighted BCA assigns each outcome a value equaling the sum of individual monetary equivalents relative to the status quo and ranks outcomes in the order of these values. Monetary equivalents can be defined as “equivalent variations” (equilibrating changes to status quo consumption) or “compensating variations” (equilibrating consumption changes in each outcome) (Freeman 2003, chap. 3). Equivalent variations are theoretically preferable because BCA with compensating variations can violate the Pareto principle, while BCA with equivalent variations cannot. My presentation will henceforth focus on the equivalent variation—now using the term “monetary equivalent” to mean specifically that. In practice, BCA analysts often employ compensating variations, which can be seen as a rough-and-ready proxy for equivalent variations.

The Utility Function

Consider now the utilitarian SWF. Recall that in order to arrive at interpersonally comparable utilities we are starting with the simplified assumption of common personal preferences. Moreover, utilitarianism requires a *cardinal* interpersonally comparable utility function, one that contains information about utility differences.

⁷To be clear, the choice of a conception of well-being (e.g., the view that well-being is equivalent to the satisfaction of personal preferences), to be measured by utility numbers, is a normative choice for the decision maker—no less so than the prior choice to regiment her moral preferences via the SWF format. However, the implementation of the particular conception chosen may depend on empirical facts. In particular, what people’s personal preferences are is an empirical question.

⁸In general, an individual’s income need not equal her consumption, but in a one-period model without bequests they are identical, and so the reader can substitute “income” for “consumption” if she likes.

⁹Space constraints preclude a discussion here of the SWF approach with variable population size.

Given common personal preferences, *expected utility theory* provides a straightforward path to cardinal and interpersonally comparable utilities. This theory shows that if someone has a well-behaved ranking of lotteries over attribute bundles, there will exist a so-called von Neumann–Morgenstern (vNM) utility function, the mathematical expectation of which corresponds to the ranking. Moreover, this utility function, albeit not unique, is unique “up to a positive linear transformation” (Kreps 1988).

In the model now being discussed, assume that individuals indeed have common personal preferences with respect to lotteries over consumption/market price/nonconsumption attribute bundles and that these preferences are represented by some vNM utility function, denoted as $u(\cdot)$. Imagine that the utilitarian SWF, with $u(\cdot)$ as the utility function, achieves a particular ranking of outcomes. If any other vNM utility function also represents the common preferences, it is straightforward to show that the utilitarian SWF with this second function will achieve the very same ranking of outcomes, because the second function must be a positive linear transformation of $u(\cdot)$.

In short, with common personal preferences and the utilitarian SWF, we can simply use any vNM function representing the common preferences as the basis for interpersonal comparisons and the input to the SWF.¹⁰

Marginal Utility and Distributional Weights

Take any vNM utility function representing the common preferences. Let the weighting factor for each individual i be her *marginal utility* of consumption in the status quo: the change in utility that occurs when a small amount is added to the individual’s status quo consumption, divided by the consumption change. Denote this as MU_i . Assign a given outcome the sum of monetary equivalents multiplied by these weighting factors. Then if all the outcomes are sufficiently small variations around the status quo—prices, individuals’ consumption amounts, and their nonconsumption attributes do not change very much—the ranking of outcomes by the utilitarian SWF is well approximated by this formula.

This formula is very intuitive. BCA measures well-being changes in *money*: a given change in welfare-relevant attributes is measured as the equivalent change in the individual’s consumption (total monetary expenditure). But equal money changes do not necessarily correspond to equal changes in interpersonally comparable utility. A small money change for a given individual can be translated into a utility change by using the individual’s marginal utility of money as an adjustment factor.

How do we calculate MU_i ? Recall that we are trying to estimate individuals’ marginal utilities in the status quo—where individuals have various consumption and nonconsumption attributes but face a common price vector. We can therefore ignore the full content of the utility function $u(\cdot)$ and focus on estimating the utility values of consumption–nonconsumption bundles at status quo prices.

There are many different empirical methods for estimating these utility values. Let me briefly describe one approach. Consider the ranking (by the common preferences) of lotteries over consumption amounts, given that nonconsumption attributes are fixed at one or another specific level. Each such ranking is captured by a “conditional” utility function, which takes

¹⁰For a defense of vNM theory as the basis for interpersonal comparisons, see Adler (2016).

consumption alone as its argument. The conditional utility function will be concave, convex, or linear in consumption if individuals are, respectively, risk averse, risk prone, or risk neutral over consumption lotteries, holding fixed nonconsumption attributes at the designated level.

Assume that we have estimated this conditional utility function for the various levels of nonconsumption attributes that individuals experience in the status quo. Assume, moreover, that we have ordinary (nonlottery) willingness-to-pay/willingness-to-accept data, indicating the change in consumption that suffices to compensate individuals for moving from one level of the nonconsumption attributes to another. Putting both sorts of data together, we are in a position to estimate utility as a function of both consumption and nonconsumption attributes, and thus the weighting factor MU_i for each individual i .

The conditional utility function depends upon the level of nonconsumption attributes. This allows for the possibility that the degree of individuals' risk aversion/proneness with respect to consumption gambles depends upon the specific level of nonconsumption attributes. But two useful simplifications can now be introduced. First, we might assume that preferences over consumption gambles are invariant to the level of nonconsumption attributes. If so, there is a single conditional utility function, and it can be shown that MU_i is just equal to the slope (derivative) of this function at i 's status quo consumption level, multiplied by a scaling factor that inflates or deflates this value to take account of her nonconsumption attributes.

A second simplification is to assume that preferences over consumption gambles, holding fixed nonconsumption attributes, take the constant relative risk aversion (CRRA) form (Gollier 2001, chap. 2). The CRRA form is extremely popular in the literature on preferences for consumption (or income or wealth) gambles. It allows us to capture the degree of risk aversion or proneness with respect to consumption in a single parameter, λ . Indeed, much of the existing work on distributional weights incorporates the CRRA form.¹¹

The risk-aversion parameter λ of the CRRA utility function and the inequality-aversion parameter γ of the isoelastic SWF are conceptually quite different. λ is a number that captures individuals' *personal* preferences over consumption gambles. It is useful in estimating distributional weights both for the utilitarian SWF (which has no γ parameter) and for the isoelastic SWF (which does). In contrast, the inequality-aversion parameter γ captures the *moral* preferences of a certain kind of social planner (namely, one who morally prefers to give some priority to utility changes affecting worse-off individuals rather than simply aggregating utilities in utilitarian fashion). Appendix table 1 gives formulas for BCA with distributional weights to mimic a utilitarian SWF, both in the general case and with the two simplifications just mentioned.

Weights That Depend Only on Consumption?

As can be seen from Appendix table 1, the weighting factor MU_i for individual i depends on both her status quo consumption and her status quo nonconsumption attributes. Much work on distributional weights is yet more simplified—making someone's weighting factor just a function of his consumption (or income or wealth) (see, e.g., HM Treasury 2003).

Distributional weights of this consumption-only form can be theoretically supported in only two special cases: (1) those affected by the policy are heterogeneous with respect to status quo consumption but relatively homogeneous with respect to status quo nonconsumption

¹¹For estimates of λ , see Kaplow (2005).

attributes or (2) the utility function not only satisfies the invariance requirement, but takes a special additively separable form.

Generalizations: Uncertainty and Inframarginal Changes

The model presented here has assumed that each policy choice, for certain, leads to a particular outcome. More realistically, each policy choice is a probability distribution over outcomes (with these probabilities capturing the decision maker's uncertainty about which outcome will result), and the status quo itself is a probability distribution over outcomes. In this case, BCA defines an individual's monetary equivalent for a policy as the change to her consumption in all status quo outcomes that makes her indifferent between the status quo and the policy. The utilitarian SWF, in turn, can now be approximated by a formula that smoothly generalizes the one given above: the sum of each individual's monetary equivalent multiplied by a weighting factor (denote this as EMU_i) equaling her expected marginal utility of consumption, given the probability distribution over consumption and other attributes that she faces in the status quo.

The use of weighted BCA to mimic a utilitarian SWF also generalizes to the case where policies are no longer "small" variations from the status quo. An individual's weight, now, is policy specific. It depends not only on her status quo attributes, but also the magnitude of the monetary equivalent corresponding to a particular policy.

BCA and Isoelastic Distributional Weights

Let us continue using the simple one-period model of choice under certainty from the previous section. However, we now assume that the decision maker has "isoelastic" moral preferences, taking the form of an isoelastic SWF.

The isoelastic SWF's ranking of policies that are "small" deviations from the status quo—like the utilitarian ranking—can be approximated by the sum of individual monetary equivalents, each multiplied by a weighting factor that is just a function of the individual's status quo attributes. Recall that isoelastic moral preferences give priority to well-being changes affecting worse-off individuals. Recall, too, that the degree of such priority is embodied in an inequality-aversion parameter γ , which can take any positive value. These features of the isoelastic SWF are reflected in the corresponding distributional weights. Individual i 's isoelastic distributional weight is her utilitarian distributional weight, MU_i —her marginal utility of consumption, given her status quo attributes—multiplied by an additional term, $MMVU_i$.

$MMVU_i$ is the marginal moral value of utility for individual i . If we increase i 's utility by some small delta (an increment to utility), there is a "moral benefit": the outcome in which that increase occurs becomes morally better, *ceteris paribus*. This moral benefit, divided by delta—the moral benefit per unit of utility—is $MMVU_i$.

If the utility level of individual i is larger than that of individual j , $MMVU_i$ is less than $MMVU_j$. This reflects the priority that the isoelastic SWF gives to worse-off individuals. A small increment to individual j 's utility yields a larger moral benefit than the very same increase in individual i 's utility.¹²

¹²In contrast, in the case of the utilitarian SWF, the moral benefit per unit of utility is constant. A given well-being increment yields the same moral improvement regardless of the utility level of the person receiving the increment, and so there is no need to multiply MU_i by a $MMVU_i$ term.

$MMVU_i$ is not merely decreasing with utility level; its specific magnitude depends upon the value of γ . Adler (2012, pp. 392–399) describes various thought experiments that the decision maker might undertake in choosing a value of γ . Appendix table 1 provides formulas for BCA with isoelastic distributional weights.

Further Differences Between Isoelastic and Utilitarian Distributional Weights

Apart from the extra $MMVU_i$ term, there are three additional respects in which isoelastic weighting differs from utilitarian weighting. First, the isoelastic SWF requires interpersonal comparability of utility ratios. Even given the assumption of common personal preferences, a vNM utility function representing such preferences is only unique up to a positive linear transformation, and thus is adequate for utilitarian weights but not isoelastic weights. Given common personal preferences, a utility function unique up to a positive ratio transformation (and thus sufficient to determine both isoelastic and utilitarian weights) can be produced by taking the vNM function and then assigning the number zero to a “threshold bundle”: one which individuals regard as being just at the threshold of a life worth living (e.g., a bundle with very low consumption and bad health) (Adler 2012, chap. 5).

A second additional difference from utilitarian weights concerns policy choice under uncertainty (Fleurbaey 2010; Adler 2012, chap. 7). It turns out that there are two ways to apply an isoelastic SWF to a set of policies, each of which is a probability distribution over outcomes: the “ex ante” approach and the “ex post” approach. The ex ante approach is readily operationalized by taking monetary equivalents under uncertainty, and by setting the weighting factor equal to the EMU_i term multiplied by an $MMVEU_i$ term (marginal moral value of expected utility) analogous to the $MMVU_i$ term in the certainty case. The ex post approach is more complicated to mimic with distributional weights.

Finally, the utilitarian SWF yields consumption-only weights, even with population heterogeneity in nonconsumption attributes, if the utility function takes a special additively separable form. That is not true for the isoelastic SWF.

An Illustrative Example: VSL and Weights

This section uses the value of statistical life (VSL) to illustrate distributional weighing. VSL is the marginal rate of substitution between survival probability and income. This is the concept used to monetize the benefits that accrue from reducing fatality risks. VSL is central to the practice of BCA by the Environmental Protection Agency and other U.S. governmental agencies (Cropper, Hammitt, and Robinson 2011).

I use the workhouse, one-period model of VSL that is standard in the literature (Eeckhoudt and Hammitt 2004). Each individual in the status quo earns some income and has some probability of surviving through the end of the current period (e.g., the current year) and consuming her income; if she doesn’t survive, the income is bequeathed. Policies change individuals’ survival probabilities or incomes. In other words, each individual in any given outcome has an attribute bundle consisting of an income amount plus a single binary nonincome attribute: surviving the current period or dying. The status quo and policies are lotteries over such outcomes. Each individual has personal preferences over (income, die/survive) bundles,

represented by a vNM utility function. Her VSL in the status quo reflects her status quo income and survival probability, plus these preferences.

In order to enable the calculation of distributional weights, I assume that individuals have common personal preferences over (income, die/survive) bundles, represented by a common vNM function. I also assume that marginal utility in the death state is zero: the utility of income conditional on the attribute “die” is a flat line. Since utility here is supposed to represent personal benefit, this assumption seems compelling.¹³ It also means that we can calibrate the common vNM function by knowing the “subsistence” level of income: the level which, if combined with the attribute “survive,” is so low that individuals are indifferent between that bundle and dying. I add the simplifying assumption that individuals’ preferences over consumption lotteries in the “survive” state are CRRA. Note finally that by determining the subsistence level of income, we have at the very same time identified a natural zero point for purposes of the isoelastic SWF.

Given an individual’s status quo income and survival probability, we can now assign her (1) a VSL value; (2) a utilitarian weight, equaling her expected marginal utility of income (EMU_i); and (3) an isoelastic weight, equaling EMU_i multiplied by the marginal moral value of expected utility ($MMVEU_i$), with $MMVEU_i$ in turn a function of the coefficient of inequality aversion (γ) that the policymaker chooses. For small changes in income and survival probability, an individual’s monetary equivalent is approximately the income change plus the product of her VSL and the change in survival probability. The utilitarian ranking of policies is, in turn, approximated by the sum of monetary equivalents multiplied by the utilitarian weights and the ex ante isoelastic ranking by the sum of monetary equivalents multiplied by the isoelastic weights [see Adler, Hammitt, and Treich (2014) for a related analysis].

VSL and Distributional Weights for Two Populations

Assume that we are considering policies that will affect two populations: a better-off group with a higher income and lower status quo fatality risk, and a worse-off group with a lower income and higher status quo risk. Specifically, let each member of the first group have an annual income of \$100,000 and face an annual all-cause fatality risk of 0.005, while each member of the second group has an annual income of \$20,000 and faces an annual all-cause fatality risk of 0.01.¹⁴

Table 3 calculates VSL values for the better- and worse-off individuals (for short, “Rich” and “Poor”), as well as the ratios of these values, the ratios of the Poor and Rich utilitarian distributional weights, and the ratios of their isoelastic weights—given different assumptions about the coefficient of risk aversion λ , the subsistence level, and the degree of inequality aversion γ .

Note that, especially at higher values of λ , the Rich VSL is many multiples of the Poor VSL. Such high multiples are inconsistent with observed values of the income elasticity of VSL—perhaps reflecting the limitations of the simple analytic model of VSL used here. Alternatively, low observed income elasticities of VSL may reflect real-world violations of the axioms of expected utility.¹⁵

In any event, what bears special note about table 3 is how weights counteract the higher VSL values of Rich individuals. For example, at $\lambda = 2$ and a subsistence level of \$1,000, the Rich VSL

¹³I may have moral or altruistic reasons to care about the amount of income left to my heirs, but their consumption after my death does not change my well-being.

¹⁴The annual all-cause fatality risk for the entire U.S. population is approximately 0.008 (Xu et al. 2016).

¹⁵On income elasticity of VSL, see Kaplow (2005), Evans and Smith (2010), and Viscusi (2010).

Table 3 Distributional weights and VSL

	$\lambda = 0.5$ Sub = \$1000	$\lambda = 0.5$ Sub = \$5000	$\lambda = 1$ Sub = \$1000	$\lambda = 1$ Sub = \$5000	$\lambda = 2$ Sub = \$1000	$\lambda = 2$ Sub = \$5000	$\lambda = 3$ Sub = \$1000	$\lambda = 3$ Sub = \$5000
VSL _{Rich}	\$180,905	\$156,059	\$462,831	\$301,079	\$9,949,749	\$1,909,548	\$502,462,312	\$20,050,251
VSL _{Poor}	\$31,369	\$20,202	\$60,520	\$28,006	\$383,838	\$60,606	\$4,030,303	\$151,515
VSL _{Rich} /VSL _{Poor}	5.8	7.7	7.6	10.8	25.9	31.5	124.7	132.3
U-Weight _{Poor} /U-Weight _{Rich}	2.2	2.2	5	5	24.9	24.9	124.4	124.4
Iso-Weight _{Poor} /Iso-Weight _{Rich} $\gamma = 0.5$	3.6	4.2	6.2	7.3	25.5	28.1	124.8	128.6
Iso-Weight _{Poor} /Iso-Weight _{Rich} $\gamma = 1$	5.8	7.8	7.7	10.8	26.1	31.7	125.3	133
Iso-Weight _{Poor} /Iso-Weight _{Rich} $\gamma = 2$	15.1	27.1	11.9	23.5	27.3	40.3	126.2	142.2

Note: “U-Weight” indicates the utilitarian distributional weight for Rich or Poor (depending on the subscript) and “Iso-Weight” indicates the isoelastic weight. Each column corresponds to different assumptions about the coefficient of risk aversion λ for the CRRA utility function and the subsistence level of income.

is 25.9 times that of the Poor, but the Poor utilitarian weight is 24.9 times that of the Rich.¹⁶ Thus, while traditional BCA would assign Rich a monetary equivalent for a given small risk reduction that is 25.9 times larger than Poor’s monetary equivalent for the same risk reduction, utilitarian weighted BCA would assign Rich an adjusted monetary equivalent that is $25.9/24.9 = 1.04$ times larger than Poor’s.

Adding isoelastic weights further counteracts the Rich/Poor VSL divergence. Continuing with the scenario of $\lambda = 2$ and a subsistence level of \$1,000, note that at a low value (0.5) of inequality aversion γ , the Poor/Rich ratio of isoelastic weights is 25.5. Rich’s adjusted monetary equivalent is now even closer to Poor’s (the ratio is $25.9/25.5 = 1.02$). At larger values of γ , isoelastic weights “overcompensate” for the difference between Rich and Poor VSL: the Poor/Rich ratio of isoelastic weights exceeds the Rich/Poor VSL ratio.

The utilitarian and isoelastic weights would also, of course, affect the relative weighting of income reductions incurred by Rich or Poor. If someone’s income is reduced by amount Δc , traditional BCA assigns the same value (Δc) to that reduction, regardless of the individual’s attributes. Weighted BCA assigns the reduction a value equaling Δc multiplied by the distributional weight. Thus the ratio between the weighted value of a reduction in Poor’s income and the weighted value of the very same reduction in Rich’s income is simply the ratio of distributional weights—as displayed in the fourth row of the table for the utilitarian case and in subsequent rows for isoelastic weights.

Using Traditional BCA and Weighted BCA to Evaluate Risk-Reduction Policies

Let us now consider four types of policies, assuming that the two groups have the same number of members: (1) uniform risk reduction and cost incidence (the policy produces the same

¹⁶A subsistence level of \$1,000 is in the range of the “extreme poverty” level of \$1.90/day now used by the World Bank (Ferreira et al. 2015). Empirical evidence on λ is mixed; a value of 2 or even substantially higher is not empirically implausible (Kaplow 2005).

reduction in fatality risk for both Rich and Poor individuals, specifically a 1-in-100,000 reduction for each individual of the risk of dying during the current year, and the dollar costs of the policy are borne equally by the Rich and Poor groups); (2) uniform risk reduction and redistributive incidence (both Rich and Poor individuals receive the 1-in-100,000 risk reduction, but all the costs are borne by the Rich group); (3) concentrated risk reduction and cost incidence (only Poor individuals receive the 1-in-100,000 risk reduction, and they pay the costs); and (4) regressive risk transfer (Poor individuals suffer an increase in risk of 1-in-100,000, with Rich individuals receiving a risk reduction of the same amount—as might occur with a decision to relocate a hazardous facility closer to where Poor rather than Rich individuals reside).

Table 4 shows how each of these policy choices would be evaluated by traditional BCA (summing monetary equivalents), BCA with population-average rather than differentiated VSL values, BCA with utilitarian weights, and BCA with isoelastic weights. The use of BCA with population-average values lacks any firm theoretical justification but is now standard practice in the U.S. government (Robinson 2007), and so it is included in the analysis.

In the case of uniform risk reduction and cost incidence, utilitarian BCA is only willing to impose a relatively low per capita cost on all individuals (as compared with traditional or population-average BCA), and isoelastic BCA a yet lower cost: paying more for a 1-in-100,000 fatality risk reduction would still be net beneficial for the Rich (given their larger VSL values), but would be too large a net welfare loss for the Poor. Conversely, in the case of uniform risk reduction and redistributive incidence, utilitarian BCA is willing to impose a larger per capita cost (now borne exclusively by Rich individuals) than traditional or population-average BCA, and isoelastic BCA a yet larger cost.

In the case of concentrated risk reduction and cost incidence, population-average BCA violates the Pareto principle: government may end up imposing a per-capita cost on Poor individuals much larger than they are willing to pay for a 1-in-100,000 risk reduction.¹⁷ Note that traditional BCA avoids this unpleasant result, but so do the weighted versions.

Although the use of population-average BCA conflicts with the Pareto principle, it does have a key intuitive advantage: in a case such as regressive risk transfer, traditional BCA approves the transfer while population-average BCA is neutral. Interestingly, utilitarian BCA also approves the regressive transfer, and isoelastic BCA will do so for a sufficiently low value of the inequality-aversion parameter, γ —but as γ increases, the isoelastic SWF “switches” and favors a progressive risk transfer from Poor to Rich.

Objections to Distributional Weighting

Next I examine two potential concerns about distributional weighting: first, that interpersonal comparisons of well-being, and thus weights, are undermined by heterogeneous preferences, and second, that distributional considerations are best addressed through the tax system.

¹⁷This break-even cost is approximately VSL_{Poor} times the risk reduction, which in the scenario covered by the table is \$3.84.

Table 4 The effect of distributional weights on different kinds of risk policies

	Uniform risk reduction and cost incidence	Uniform risk reduction and redistributive incidence	Concentrated risk reduction and cost incidence	Regressive risk transfer
	Maximum per capita cost imposed uniformly on Rich and Poor	Maximum per capita cost imposed on Rich	Maximum per capita cost imposed on Poor	“Yes” if the transfer is assigned a positive sum of monetary equivalents or weighted equivalents; “No” if it is assigned a negative sum; “Neutral” if assigned a zero sum
BCA without weights	\$51.67	\$103.33	\$3.84	Yes $VSL^{Rich}/VSL^{Poor} = 25.9 > 1$
BCA with population average VSL	\$42.33	\$84.67	\$42.33	Neutral Cost to Poor = Benefit to Rich = $VSL^{Avg} \times 1/100,000 = \42.33
BCA with utilitarian weights	\$7.54	\$194.97	\$3.84	Yes $(VSL^{Rich}/VSL^{Poor}) \div (U-Weight^{Poor}/U-Weight^{Rich}) = 1.04 > 1$
BCA with isoelastic weights, $\gamma = 0.5$	\$7.45	\$197.21	\$3.84	Yes $(VSL^{Rich}/VSL^{Poor}) \div (Iso-Weight^{Poor}/Iso-Weight^{Rich}) = 1.02 > 1$
BCA with isoelastic weights, $\gamma = 1$	\$7.37	\$199.50	\$3.84	No $(VSL^{Rich}/VSL^{Poor}) \div (Iso-Weight^{Poor}/Iso-Weight^{Rich}) = 0.99 < 1$
BCA with isoelastic weights, $\gamma = 2$	\$7.22	\$204.23	\$3.84	No $(VSL^{Rich}/VSL^{Poor}) \div (Iso-Weight^{Poor}/Iso-Weight^{Rich}) = 0.95 < 1$

Notes: These calculations assume the scenario of $\lambda = 2$ and a subsistence level of \$1,000. The population-average VSL of \$4.233 million is calculated by assuming that individuals are uniformly distributed from income levels of Poor = \$20,000 in \$10,000 increments to Rich = \$100,000, and that background fatality risk decreases linearly from the Poor level of 0.01 to the Rich level of 0.005. To understand the regressive risk transfer column, note that BCA without weights approves the transfer if $VSL^{Rich} \times (1/100,000) > VSL^{Poor} \times (1/100,000)$, which is equivalent to $VSL^{Rich}/VSL^{Poor} > 1$. Similarly, BCA with utilitarian weights approves the transfer if $VSL^{Rich} \times U-Weight^{Rich} \times (1/100,000) > VSL^{Poor} \times U-Weight^{Poor} \times (1/100,000)$, which is equivalent to $(VSL^{Rich}/VSL^{Poor}) \div (U-Weight^{Poor}/U-Weight^{Rich}) > 1$. This explains the formulas in the first and third rows, and a similar explanation can be provided for the other formulas.

Heterogeneous Preferences

It is important to distinguish between the heterogeneity of *moral* preferences and the heterogeneity of *personal* preferences.

Heterogeneity of moral preferences is no threat to the view of distributional weighting presented here. Some citizens or officials will oppose BCA as a criterion for assessing governmental policies, while others will endorse BCA. Within the latter group, some will be persuaded by the Kaldor–Hicks defense of BCA, while others will find the SWF framework more attractive.

Weighting is a procedure that operationalizes the moral preferences of this last group in a systematic form.

Thus critiques of distributional weights as “value laden” are inapposite. Of course, the decision to use weights, and the choice of SWF to be mirrored by weights, are “value laden”—but so is the decision to use BCA at all, or to use the unweighted variant. These decisions, too, involve a whole series of contestable moral judgments. Governmental officials inevitably make such judgments, or work for higher-ups who do.

The use of distributional weights does raise questions of institutional role. An unelected bureaucrat might feel that it would be legally problematic, or democratically illegitimate, for her to specify weights. Who in government gets to act on contestable moral preferences is a complicated (and itself contestable) question of law and democratic theory. Suffice it to say that the advice welfare economists and moral theorists provide about the specification of weights is addressed to officials with the legal and democratic authority to act on such advice—whoever exactly those officials may be.

Heterogeneity of personal preferences creates additional complexity in the construction of the utility function, which generates the vectors of interpersonally comparable utility numbers that are the inputs into the SWF. We are presuming throughout this article that the decision maker whose moral preferences are represented by the SWF has adopted the following (very plausible) conception of well-being: each person’s well-being is determined by her personal preferences. Until this point, we have further assumed that individuals in the population of concern have identical personal preferences over attribute bundles. What happens if that assumption is relaxed?

If so, the utility function must assume a more complicated form than was used earlier. It must now take into account both the individual’s attributes and her preference structure with respect to attribute bundles. If two individuals in some outcome have the same attributes, but different personal preferences, this more complicated utility function may assign them different utility numbers.

How to construct an interpersonally comparable utility function that makes an individual’s utility a function of both her attributes and her preferences is a topic of ongoing research. One possibility deploys the concept of *equivalent income* (Fleurbaey 2016; Fleurbaey and Rafeh 2016). Each person’s income is adjusted to take account of her nonincome attributes (health, leisure, etc.) and her preferences. This adjusted (“equivalent”) income is then used as the measure of her well-being.

A different possibility is suggested by John Harsanyi’s concept of “extended preferences,” which I have further developed in recent work (Harsanyi 1977; Adler 2016). In effect, the utility function is constructed by taking vNM functions representing the different possible preferences over bundles in the population and choosing scaling factors for each one.

However we construct a utility function that allows for heterogeneous personal preferences—whether by using equivalent incomes, extended preferences, or in some other manner—we can generate a corresponding scheme of distributional weights.

The Tax System

Are distributional considerations better handled through a combination of the tax system and unweighted BCA, as opposed to BCA with distributional weights? I examine two arguments for this position.

If taxes and transfers are optimal, distributional weights are uniform

Assume that the decision maker with legal authority over some nontax body has moral preferences reflected in a utilitarian or isoelastic SWF and that the legislature (which controls the tax system) has the very same moral preferences. Moreover, the tax system is a “lump sum” system, meaning taxation has zero administrative cost and no adverse incentive effects (e.g., with respect to individual labor supply).

Under these assumptions, utilitarian or isoelastic distributional weights will be identical for all individuals. In the case where the legislature and decision maker share utilitarian moral preferences, the legislature will have enacted taxes and transfers that equalize individuals' MU_i values. And in the case where they have isoelastic moral preferences (with the same value of γ), the legislature will have enacted taxes and transfers that equalize individuals' $MU_i \times MMVU_i$ values (Drèze and Stern 1987, pp. 935–937; Salanié 2003, pp. 79–83). If so, the decision costs of calculating weights are unnecessary, since BCA with identical weights is equivalent to BCA without weights.

However, the assumptions can readily fail. First, the decision maker may have different moral preferences from the legislature. For example, a President elected by one political party may have more strongly egalitarian preferences than the competing party that controls the legislature. If so, and even if the tax system is a “lump sum” system, the status quo distribution of income could well be such that distributional weights according to the SWF that the President endorses are unequal. Thus the President (given his moral preferences) would have good moral reason to direct nontax bodies within his control to use BCA with distributional weights, and he might well also believe that doing so is democratically and legally legitimate.

Indeed, preference heterogeneity on the part of government officials is reflected in the institutional history of BCA in the United States. Since 1981, a presidential order has directed Executive Branch agencies to use BCA as a general tool for evaluating regulation, but proponents of BCA have never succeeded in persuading Congress to embody this directive in a statute (Wiener 2006, pp. 461–466). Thus it is hardly science fiction to suppose that a President might prefer to move BCA, now in place, in a more egalitarian direction, but be unable to secure tax code changes reflecting this preference.

Second, even if the legislature and decision maker share a moral commitment to a utilitarian SWF or to an isoelastic SWF with a particular value of γ , the absence of a lump-sum tax and transfer system may mean that distributional weights in the status quo are not equal by the lights of the shared SWF. The most systematic scholarly thinking about non-lump-sum taxation has occurred in the field of optimal tax theory (see, e.g., Salanié 2003; Kaplow 2008). It is not generally true in this context that the optimal taxes (in light of a given SWF) will yield a pattern of income and nonincome attributes such that distributional weights (in light of that same SWF) are equal (Layard 1980; Fleurbaey et al. 2013). The optimal-tax setup typically ignores administrative costs; adding these to the tax system may also help to produce unequal distributional weights at the optimum (Quiggin 1995.)

Whether or not distributional weights are uniform, the decision maker can produce a Pareto-superior result by a combination of unweighted BCA and tax changes

A second argument against the use of distributionally weighted BCA has been developed by Kaplow (see, e.g., 1996, 2008). Consider a nontax decision maker who endorses some Paretian

SWF (be it isoelastic, utilitarian, or any other SWF that respects the Pareto principle). Assume that some policy P^* is chosen by traditional BCA without weights, as opposed to an alternative policy P . Thus those who gain from P^* have enough to compensate those who lose. Kaplow shows how these gains can be redistributed from the winners to the losers to yield a Pareto-superior result even with a non-lump-sum tax system.¹⁸ Note, now, that if P^* together with a change to the tax code is Pareto superior to P , any Pareto-respecting SWF will favor that package as opposed to the enactment of P .

Kaplow's analysis relies on a key technical assumption of labor separability. This assumption has no clear empirical basis, and a different assumption is adopted by much scholarship in environmental economics (Johansson-Stenman 2005; Fullerton 2009, p. xxi). Moreover, the analysis does not consider the administrative costs of changes to the tax system or—once again—the possibility that a nontax decision maker who favors some SWF may not control the tax system. Imagine that a decision maker has the power to choose between P and P^* . If the tax code does not change, her morally preferred SWF favors P over P^* . Although P^* bundled with a change to the tax code is Pareto superior to P , the decision maker does not believe the legislature will make this tax change. She is therefore morally justified in picking P .

Given these caveats, Kaplow's analysis does not demonstrate that a decision maker favoring some SWF should refrain from putting in place a matching schedule of distributional weights. Rather, the analysis has a less dramatic—but still important—implication. It suggests that the “policies” being considered should include not merely nontax policies, but also such policies combined with changes to the tax code where the decision maker believes these changes to be politically feasible. In some cases, P will be favored by distributionally weighted BCA (mirroring the SWF) over P^* ; but P^* together with a politically feasible change to the tax system will be favored by distributionally weighted BCA over P , even given the actual administrative costs of this change and the predicted incentive effects of the change on labor supply (zero or not). If the combined option of P^* plus the feasible tax change is Pareto superior to P , distributionally weighted BCA will prefer the first option to the second, and it might do so even absent Pareto superiority.

In short, decision makers favoring an SWF mirrored by distributional weights should consider an expanded menu of policy options—to include tax system changes, if feasible, as well as nontax policies—but the existence of the tax system does not show that they should abandon weighting entirely.

Conclusion

This article has shown how to “put structure” on the problem of distributional weights and has addressed some recurrent objections. The decision to use weights to mimic a particular SWF does require an ethical/moral judgment, but so does the decision to use unweighted BCA—a procedure that is intensely controversial outside the community of economists (Ackerman and Heinzerling 2004). The specification of utilitarian weights is quite straightforward; these simply reflect the marginal utility of consumption/income. The specification of “isoelastic” weights is somewhat more complicated, requiring a further moral judgment regarding the appropriate degree of inequality aversion γ and the specification of a “zero point” (subsistence level) so that

¹⁸Specifically, he works within the setup from optimal tax theory, which takes account of the labor-disincentive effects of an income tax.

utility will be measurable on a ratio scale. Still, some will feel that the utilitarian SWF is insufficiently sensitive to distribution and that this additional complexity is justified.

Traditional BCA assigns a higher monetary value to risk reductions incurred by richer individuals, since VSL increases with income. The use of weights tends to counteract this feature of traditional BCA without producing a conflict with the Pareto principle (a key flaw in BCA with a population-average VSL).

The absence of a lump-sum tax system, and the economic and political costs of tax reform, undercut the objection that distributional concerns should be wholly relegated to the tax code. The SWF approach, and thus weights, do require interpersonal well-being comparisons. These are straightforward with common personal preferences. Heterogeneous personal preferences pose additional complications but are not an insuperable obstacle to interpersonal comparisons.

Technical Appendix

The main text focuses on the case in which each policy choice leads with certainty to some outcome. This Appendix provides some key formulas for that case. For their derivation and further discussion, see the online supplementary materials and Adler (2012, chaps. 2, 3, and 5).

Utilitarian and Isoelastic SWF

Let \mathbf{B}_i^x be individual i 's bundle of attributes in a given outcome x and $u(\cdot)$ be the utility function mapping bundles onto utility numbers. Then outcome x corresponds to the utility vector $(u(\mathbf{B}_1^x), \dots, u(\mathbf{B}_N^x))$, with N individuals in the population.

The utilitarian and isoelastic SWFs—expressed as rules for ranking outcomes in light of their corresponding utility vectors—are as follows. The utilitarian SWF: x at least as good as y iff $\sum_{i=1}^N u(\mathbf{B}_i^x) \geq \sum_{i=1}^N u(\mathbf{B}_i^y)$. The isoelastic SWF, parameterized by $\gamma > 0$: x at least as good as y iff $(1 - \gamma)^{-1} \sum_{i=1}^N u(\mathbf{B}_i^x)^{1-\gamma} \geq (1 - \gamma)^{-1} \sum_{i=1}^N u(\mathbf{B}_i^y)^{1-\gamma}$. In the special case of $\gamma = 1$: x at least as good as y iff $\sum_{i=1}^N \ln u(\mathbf{B}_i^x) \geq \sum_{i=1}^N \ln u(\mathbf{B}_i^y)$.

The Utility Function

Assume individuals have common personal preferences over bundles and bundle lotteries. If these preferences satisfy the axioms of vNM theory, there exists a vNM function $u(\cdot)$ that represents these preferences, and another function $u^*(\cdot)$ is a vNM function iff there exists $a > 0$, b , s.t. $u^*(\mathbf{B}) = au(\mathbf{B}) + b$ for all \mathbf{B} .

Note that the utilitarian SWF is invariant to the substitution of $u^*(\cdot)$ for $u(\cdot)$: $\sum_{i=1}^N u(\mathbf{B}_i^x) \geq \sum_{i=1}^N u(\mathbf{B}_i^y)$ iff $\sum_{i=1}^N u^*(\mathbf{B}_i^x) \geq \sum_{i=1}^N u^*(\mathbf{B}_i^y)$. This is not true for the isoelastic SWF. However, we can define a subfamily of the vNM functions representing the common preferences, unique up to a ratio transformation, as follows. Let \mathbf{B}^{Zero} be a threshold bundle and let $u(\cdot)$ be a vNM function s.t. $u(\mathbf{B}^{\text{Zero}}) = 0$. Then if $u^*(\cdot) = au(\cdot)$, $a > 0$, the isoelastic SWF is invariant to the substitution of $u^*(\cdot)$ for $u(\cdot)$.

BCA and Equivalent Variations

Let $(c_i^x, \mathbf{p}^x, \mathbf{a}_i^x)$ denote individual i 's bundle in outcome x —with c_i^x individual i 's consumption and \mathbf{a}_i^x her nonconsumption attributes, respectively, in x , and \mathbf{p}^x the price vector in x . Let $(c_i^s, \mathbf{p}^s, \mathbf{a}_i^s)$ denote i 's status quo bundle. Individual i 's equivalent variation for outcome x is the amount Δc_i^x such that she is indifferent between the bundle $(c_i^s + \Delta c_i^x, \mathbf{p}^s, \mathbf{a}_i^s)$ and the bundle $(c_i^x, \mathbf{p}^x, \mathbf{a}_i^x)$. Traditional (unweighted) BCA using equivalent variations assigns each outcome x the numerical value $\sum_{i=1}^N \Delta c_i^x$, and ranks outcomes in the order of these values.

BCA with Distributional Weights

We simplify notation by suppressing reference to status quo prices and abbreviating $u(c, \mathbf{p}^s, \mathbf{a})$ as $u(c, \mathbf{a})$. MU_i is individual i 's status quo marginal utility of consumption. $MU_i = \frac{\partial u(c_i^s, \mathbf{a}_i^s)}{\partial c}$. BCA with utilitarian distributional weights assigns each outcome x the value $\sum_{i=1}^N MU_i \times \Delta c_i^x$.

Let $h^{\mathbf{a}}(\cdot) = h^{\mathbf{a}}(c)$ be a conditional utility function representing the ranking (by the common preferences) of lotteries over consumption amounts, with nonconsumption attributes fixed at level \mathbf{a} and prices at \mathbf{p}^s . “Invariance” means that this ranking is independent of \mathbf{a} , and thus there is a single conditional utility function $h(\cdot)$ s.t. $h^{\mathbf{a}}(\cdot) = h(\cdot)$ for all \mathbf{a} . Given invariance, there exists a positive constant $m(\mathbf{a})$ and a constant $k(\mathbf{a})$ such that $u(c, \mathbf{a}) = m(\mathbf{a}) \times h(c) + k(\mathbf{a})$ and $MU_i = h'(c_i^s) \times m(\mathbf{a}_i^s)$, with $h'(\cdot)$ the first derivative. In the case of invariance plus CRRA, $h(\cdot)$ is a CRRA function parameterized by risk aversion parameter λ . Then $h(c) = (1 - \lambda)^{-1} c^{1-\lambda}$ (or $\ln c$ in the special case of $\lambda = 1$) and $MU_i = (c_i^s)^{-\lambda} \times m(\mathbf{a}_i^s)$.

BCA with Isoelastic Distributional Weights

$MMVU_i$, the marginal moral value of utility for individual i , equals $u(c_i^s, \mathbf{a}_i^s)^{-\gamma}$. BCA with isoelastic distributional weights assigns each outcome x the value $\sum_{i=1}^N MMVU_i \times MU_i \times \Delta c_i^x$.

A Consolidated Table

Appendix table 1 provides formulas for utilitarian and isoelastic weights in the general case, with invariance and with invariance plus CRRA.

Appendix Table 1 Formulas for distributional weights

	General case	Invariance	Invariance plus CRRA
Utilitarian weights			
Outcome x is assigned the value $\sum_{i=1}^N MU_i \times \Delta c_i^x$	$MU_i = \frac{\partial u(c_i^s, \mathbf{a}_i^s)}{\partial c}$	$MU_i = h'(c_i^s) \times m(\mathbf{a}_i^s)$	$MU_i = (c_i^s)^{-\lambda} \times m(\mathbf{a}_i^s)$
Isoelastic weights			
Outcome x is assigned the value $\sum_{i=1}^N MMVU_i \times MU_i \times \Delta c_i^x$	$MU_i = \frac{\partial u(c_i^s, \mathbf{a}_i^s)}{\partial c}$ $MMVU_i = u(c_i^s, \mathbf{a}_i^s)^{-\gamma}$	$MU_i = h'(c_i^s) \times m(\mathbf{a}_i^s)$ $MMVU_i = u(c_i^s, \mathbf{a}_i^s)^{-\gamma}$ with $u(c_i^s, \mathbf{a}_i^s) = h(c_i^s) \times m(\mathbf{a}_i^s) + k(\mathbf{a}_i^s)$	$MU_i = (c_i^s)^{-\lambda} \times m(\mathbf{a}_i^s)$ $MMVU_i = u(c_i^s, \mathbf{a}_i^s)^{-\gamma}$ with $u(c_i^s, \mathbf{a}_i^s) = h(c_i^s) \times m(\mathbf{a}_i^s) + k(\mathbf{a}_i^s)$ and $h(c_i^s) = (1 - \lambda)^{-1} (c_i^s)^{1-\lambda}$

References

- Ackerman, Frank, and Lisa Heinzerling. 2004. *Priceless: On knowing the price of everything and the value of nothing*. New York: New Press.
- Adler, Matthew D. 2012. *Well-being and fair distribution: Beyond cost-benefit analysis*. Oxford: Oxford University Press.
- . 2016. Extended preferences. In *The Oxford handbook of well-being and public policy*, eds. Matthew D. Adler and Marc Fleurbaey, 476–517. New York: Oxford University Press.
- Adler, Matthew D., James Hammitt, and Nicolas Treich. 2014. The social value of mortality risk reduction. *Journal of Health Economics* 35:82–93.
- Bergson, Abram. 1948. Socialist economics. In *A survey of contemporary economics*, vol. 1, ed. Howard S. Ellis, 412–448. Homewood: Richard D. Irwin.
- . 1954. On the concept of social welfare. *Quarterly Journal of Economics* 68:233–252.
- Boadway, Robin. 2016. Cost-benefit analysis. In *The Oxford handbook of well-being and public policy*, eds. Matthew D. Adler and Marc Fleurbaey, 47–81. New York: Oxford University Press.
- Boadway, Robin, and Neil Bruce. 1984. *Welfare economics*. Oxford: Basil Blackwell.
- Bossert, Walter, and John A. Weymark. 2004. Utility in social choice. In *Handbook of utility theory*, vol. 2, eds. Salvador Barbera, Peter J. Hammond, and Christian Seidl, 1099–1177. Boston: Kluwer Academic.
- Brent, Robert J. 1984. Use of distributional weights in cost-benefit analysis: A survey of schools. *Public Finance Quarterly* 12:213–230.
- Chipman, John S., and James C. Moore. 1978. The new welfare economics, 1939–1974. *International Economic Review* 19:547–584.
- Cowell, Frank A., and Karen Gardiner. 1999. Welfare weights. Suntory and Toyota International Centres for Economics and Related Disciplines, London School of Economics. [http://darp.lse.ac.uk/papersDB/Cowell-Gardiner_\(OFT\).pdf](http://darp.lse.ac.uk/papersDB/Cowell-Gardiner_(OFT).pdf).
- Creedy, John. 2006. Evaluating policy: Welfare weights and value judgments. Working Paper 971, Department of Economics, University of Melbourne, Melbourne, Australia. <http://ideas.repec.org/p/mlb/wpaper/971.html>.
- Cropper, Maureen, James K. Hammitt, and Lisa A. Robinson. 2011. Valuing mortality risk reductions: Progress and challenges. *Annual Review of Resource and Environmental Economics* 3:313–336.
- Dasgupta, Ajit K., and D. W. Pearce. 1972. *Cost-benefit analysis: Theory and practice*. New York: Harper and Row.
- Dasgupta, Partha, Amartya Kumar Sen, and Stephen A. Marglin. 1972. *Guidelines for project evaluation*. Vienna: United Nations Industrial Development Organization.
- Drèze, Jean. 1998. Distribution matters in cost-benefit analysis: Comment on K. A. Brekke. *Journal of Public Economics* 70:485–488.
- Drèze, Jean, and Nicholas Stern. 1987. The theory of cost-benefit analysis. In *Handbook of public economics*, vol. 2, eds. Alan J. Auerbach and Martin Feldstein, 909–989. Amsterdam: North-Holland.
- Eckhoudt, Louis R., and James K. Hammitt. 2004. Does risk aversion increase the value of mortality risk? *Journal of Environmental Economics and Management* 47:13–29.
- Evans, Mary F., and V. Kerry Smith. 2010. Measuring how risk tradeoffs adjust with income. *Journal of Risk and Uncertainty* 40:33–55.
- Ferreira, F. H. G., S. Chen, A. Dabalen, Y. Dikhanov, N. Hamadeh, D. Jolliffe, A. Narayan, E. B. Prydz, A. Revenga, P. Sangraula, U. Serajuddin, and N. Yoshida. 2015. A global count of the extreme poor in 2012. World Bank Policy Research Working Paper 7432. <http://econ.worldbank.org>.
- Fleurbaey, Marc. 2010. Assessing risky social situations. *Journal of Political Economy* 118:649–680.
- . 2016. Equivalent income. In *The Oxford handbook of well-being and public policy*, eds. Matthew D. Adler and Marc Fleurbaey, 453–475. New York: Oxford University Press.
- Fleurbaey, Marc, Stéphane Luchini, Christophe Muller, and Erik Schokkaert. 2013. Equivalent income and fair evaluation of health care. *Health Economics* 22:711–729.
- 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.
- Freeman, A. Myrick. 2003. *The measurement of environmental and resource values*, 2nd ed. Washington, DC: Resources for the Future.
- Fullerton, Don, ed. 2009. *Distributional effects of environmental and energy policy*. Farnham, UK: Ashgate.
- Gollier, Christian. 2001. *The economics of risk and time*. Cambridge, MA: MIT Press.
- Gorman, W. M. 1955. The intransitivity of certain criteria used in welfare economics. *Oxford Economic Papers* 7:25–35.
- Harsanyi, John. 1977. *Rational behavior and bargaining equilibrium in games and social situations*. Cambridge: Cambridge University Press.
- HM Treasury. 2003. *The green book: Appraisal and evaluation in central government*. Norwich, UK: The Stationery Office.
- Johansson-Stenman, Olof. 2005. Distributional weights in cost-benefit analysis: Should we forget about them? *Land Economics* 81:337–52.
- Kaplow, Louis. 1996. The optimal supply of public goods and the distortionary cost of taxation. *National Tax Journal* 49:513–533.
- Kaplow, Louis. 2005. The value of a statistical life and the coefficient of relative risk aversion. *Journal of Risk and Uncertainty* 31:23–34.
- . 2008. *The theory of taxation and public economics*. Princeton, NJ: Princeton University Press.
- Layard, Richard, 1980. On the use of distributional weights in social cost-benefit analysis. *Journal of Political Economy* 88:1041–1047.
- Liu, Liqun. 2006. Combining distributional weights and the marginal cost of funds. *Public Finance Review* 34:60–78.
- Kreps, David M. 1988. *Notes on the theory of choice*. Boulder, CO: Westview Press.
- Little, I. M. D., and J. A. Mirrlees. 1974. *Project appraisal and planning for developing countries*. New York: Basic Books.
- . 1994. The costs and benefits of analysis: Project appraisal and planning twenty years on. In *Cost-benefit analysis*, eds. Richard Layard and Stephen Glaister, 199–231. Cambridge: Cambridge University Press.
- Meade, J. E. 1955. *Trade and welfare: Mathematical supplement*. London: Oxford University Press.
- Office of Management and Budget. 2003. *Circular A-4: Regulatory analysis*. http://www.whitehouse.gov/omb/inforeg_default/.
- Quiggin, John. 1995. The suboptimality of efficiency. *Economics Letters* 47:389–392.
- Ray, Anandarup. 1984. *Cost-benefit analysis: Issues and methodologies*. Baltimore: Johns Hopkins University Press.
- Robinson, Lisa. 2007. How U.S. government agencies value mortality risk reductions. *Review of Environmental Economics and Policy* 1:283–299.
- Robinson, Lisa A., James K. Hammitt, and Richard Zeckhauser. 2016. Attention to distribution in U.S. regulatory analyses. *Review of Environmental Economics and Policy* 10(2):308–328.
- Salanié, Bernard. 2003. *The economics of taxation*. Cambridge, MA: MIT Press.
- Samuelson, Paul A. 1947. *Foundations of economic analysis*. Cambridge, MA: Harvard University Press.
- Sen, Amartya. 1979. The welfare basis of real income comparisons: A survey. *Journal of Economic Literature* 17:1–45.
- Scitovsky, T. 1941. A note on welfare propositions in economics. *Review of Economic Studies* 9:77–88.
- Squire, Lyn, and Herman G. van der Tak. 1975. *Economic analysis of projects*. Baltimore: Johns Hopkins University Press.
- Viscusi, W. Kip. 2010. The heterogeneity of the value of statistical life: Introduction and overview. *Journal of Risk and Uncertainty* 40:1–13.
- Weisbrod, Burton A. 1968. Income redistribution effects and benefit-cost analysis. In *Problems in public expenditure analysis*, ed. Samuel B. Chase, 177–209. Washington, DC: Brookings Institution.
- Wiener, Jonathan B. 2006. Better regulation in Europe. *Current Legal Problems* 59:447–518.

Weymark, John A. 2016. Social welfare functions. In *The Oxford handbook of well-being and public policy*, eds. Matthew D. Adler and Marc Fleurbaey, 126–159. New York: Oxford University Press.

Xu, Jiaquan, Sherry L. Murphy, Kenneth D. Kochanek, and Brigham A. Bastian. 2016. Deaths:

Final data for 2013. *National Vital Statistics Reports* (vol. 64, no. 2, February 16, 2016). http://www.cdc.gov/nchs/data/nvsr/nvsr64/nvsr64_02.pdf.

Yitzhaki, Shlomo. 2003. Cost-benefit analysis and the distributional consequences of government projects. *National Tax Journal* 56:319–336.