WORKING PAPERS

N° 1707

January 2026

"Judging Social Priority and the Marginal Utility of Income Among Individuals"

James K. Hammitt

Toulouse School of Economics

Judging Social Priority and the Marginal Utility of Income Among Individuals

James K Hammitt
Harvard University (Center for Risk Analysis & Center for Health Decision Science)
Toulouse School of Economics, University of Toulouse-Capitole
jkh@harvard.edu

January 2026

Abstract

Weighted benefit-cost analysis is receiving increased attention as a method to incorporate concerns about the distribution of policy effects across individuals. Weights are intended to reflect interpersonal differences in the effect of income on wellbeing (the marginal utility of income) and the social value of improving the wellbeing of different individuals. Lacking an objective method for comparing differences or levels of wellbeing between individuals, multiple approaches to estimating how the marginal utility of income depends on income or other factors have been developed, but each of these requires strong assumptions that are not always recognized. This suggests that weights must be chosen judgmentally. Holding income constant, weights are likely to be smaller for older people, due to shorter remaining life expectancy and other factors.

Keywords: weighted benefit-cost analysis; social welfare function; wellbeing; marginal utility

Acknowledgement: This work is based on presentations at the CEPR Economics of Longevity and Ageing conference (London, July 2025) and the TSE-SBCA European Webinar (virtual, September 2025) and was supported in part by the French National Research Agency (ANR) under the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010). I thank Matthew Adler, David Canning, Mark Fleurbaey, Marcus Pivato, and Stephane Zuber for valuable discussions and an anonymous reviewer for helpful comments.

1. Introduction

Social policies are often evaluated in terms of how much they enhance individuals' wellbeing, and how the resulting gains and losses in wellbeing are distributed among individuals in the population. Conventional benefit-cost analysis (BCA) measures allocative efficiency by comparing the total benefit against the total resources consumed. Distributional analysis describes the allocation of benefits and costs within the population. Judgments about which distributions are preferred and how much efficiency should be sacrificed for an improvement in distribution are typically made informally and without quantification.

An alternative approach to policy evaluation that has gained attention is the use of weighted BCA (wBCA), which weights benefits and costs to different individuals (or groups) to account for distribution (H.M. Treasury 2020, U.S. Office of Management and Budget 2023[1]). Weights can in principle depend on any characteristics of the affected individuals, including income, health, age, gender, race, ethnicity, geographic region, historical disadvantage, and others. They are often justified as a means to account for the common assumption of diminishing marginal utility of income (that a fixed increment to income provides a larger increase in wellbeing when added to a lower than to a higher income). The policy-induced change in total wellbeing in a population can be evaluated by summing the monetary value of the benefit or harm to each individual weighted by her marginal utility of income.

In economics, a utility function is often defined as a function that rationalizes an individual's preferences and choices (assuming she always chooses what she prefers, among the available alternatives). If an individual's preferences are coherent, in the sense of satisfying minimal axioms such as asymmetry and negative transitivity (alternatively, completeness and transitivity), then a preference relation over the consequences of these choices exists and can be represented by a utility function (e.g., Kreps 1988). Such a utility function is ordinal: it assigns higher values to more-preferred consequences. An ordinal utility function defines the indifference curves that represent an individual's tradeoffs between quantities of alternative goods and services, but the level of utility associated with each indifference curve, and the difference in utility between two indifference curves, convey no meaningful information.

[1] In 2025, the 2023 OMB guidance was revoked and the 2003 guidance, which does not endorse wBCA, was restored.

An ordinal utility function is unique up to a positive monotone transformation, which means that if $u(y)$ is a valid utility function of income, so is $v(y) = f[u(y)]$ where f is any strictly increasing function.[2] Hence the marginal utility of income, $v'(y) = f'[u(y)] \cdot u'(y)$, depends on the function f, which is unknown. This implies that the marginal utility of income cannot be inferred from behavior without making strong assumptions. From observing behavior, one can (in principle) estimate an individual's indifference curves. Choices provide information about the preference order over alternatives but do not provide any information about the intensity of preferences, i.e., of the difference in utility between a more- and a less-preferred outcome.

Policy evaluation is often concerned with individuals' wellbeing. Wellbeing can be defined in any of several ways, including preference satisfaction, objective good, and experientialist (Adler 2019). The preference-satisfaction approach, ubiquitous in economics, defines wellbeing as equivalent to utility; consequences an individual prefers increase her wellbeing by definition. The objective-good approach defines some states as having higher wellbeing than others if the individual has certain goods or capabilities in the former but not the latter. Under this definition, wellbeing is independent of individuals' preferences; e.g., being more knowledgeable may be defined as better, even if the individual would prefer to be ignorant. The experientialist approach defines wellbeing as determined by the happiness, satisfaction, or other hedonic reactions to discrete experiences; it permits a distinction between decision utility (representing how one expects to respond to an experience) and experience utility (representing one's actual response) that can lead individuals to make choices that fail to maximize their experienced wellbeing.

Except for the objective-good definition, it is widely recognized that there is no objective method for making interpersonal utility comparisons, suggesting these comparisons must be judgmental.[3] It is less widely recognized that there is no objective method for making intrapersonal comparisons of differences in utility, e.g., of determining whether an individual's marginal utility of income decreases at higher levels, and by how much. The inability to make these determinations objectively implies that methods for taking account of distributional policy consequences, such as wBCA, require judgment about the appropriate weights.

[2] Strictly, $u(y)$ and $v(y)$ are indirect utility functions, representing the maximum utility the individual can achieve by optimal allocation of her income y across goods and services. For expositional simplicity, I assume that utility functions such as u and v, and transformation functions such as f, are differentiable.

[3] Even under the objective-good definition, judgment may be required to determine how multiple goods combine to produce wellbeing.

Wellbeing can be measured ordinally (unique up to a positive monotone transformation) or cardinally (unique up to a positive affine transformation[4]). Determining whether one individual or another gains or loses more wellbeing from a policy change requires a cardinal measurement. This is possible if and only if judgments about differences in wellbeing (for an individual, given different circumstances, or between individuals) satisfy certain axioms of measurement theory developed by Krantz, Luce, Suppes, and Tversky (1971), as described by Adler (2016, 2025).

The remainder of the paper is organized as follows. In Section 2, I describe how the weights used in wBCA can be chosen based on a social welfare function. Weights are often intended to capture differences in individuals' marginal utilities of income and possibly other factors. In Section 2.1, I discuss problems in measuring the marginal utility of income at different income levels for an individual, or equivalently, the marginal utility of income for different individuals having the same utility function. In Section 2.2, I discuss additional problems in comparing wellbeing and marginal utility between individuals whose utility functions may differ. Section 3 discusses weighting individuals based on their ages, and Section 4 concludes.

2. Choosing weights

Weights to reflect concerns about the distribution of project benefits and costs can be derived from a social welfare function (SWF). A SWF is a function of individuals' levels of wellbeing. It requires that differences in and/or levels of wellbeing are measured in an interpersonally comparable fashion so that it is meaningful to say that one person gains more wellbeing from an increment to income than another person, and/or that one person is better off than another.

Consider two prominent SWFs, utilitarian and prioritarian. Let wellbeing be defined by the preference-satisfaction approach and let $u_i(y_i, x_i)$ denote the wellbeing (utility) of individual i having income y_i and other relevant attributes (e.g., health) denoted by the vector x_i. The utilitarian SWF $W^U = \sum u_i$, where the summation is over all individuals who have standing (i.e., whose wellbeing is relevant to social welfare). The prioritarian SWF $W^P = \sum g(u_i)$, where g is a strictly increasing and concave function. The utilitarian SWF measures social welfare as the total of individuals' wellbeing. It is inequality-neutral with respect to wellbeing, meaning that it ranks any allocation of the same total wellbeing as equally good. In contrast, the prioritarian SWF is inequality-averse; any pure (non-leaky) transfer of wellbeing from a better-off to a worse-off individual (that does not change their order of

[4] Two functions are related by a positive affine transformation if one function can be obtained from the other by multiplying by a positive constant and adding a constant.

wellbeing) increases social welfare.[5] For the utilitarian SWF, evaluating whether a policy improves social welfare requires that differences (though not levels) in individuals' wellbeing are interpersonally comparable; for the prioritarian SWF, both differences and levels must be interpersonally comparable.

Whether a policy increases or decreases the value of a SWF can be evaluated using wBCA, where the weights are derived from the corresponding SWF. For small changes, the weight for individual *i* is proportional[6] to the marginal social value of an increase in her income. For the utilitarian SWF, the weight is proportional to her marginal utility of income, $u_i' = \frac{\partial u_i}{\partial y_i}$; for the prioritarian SWF, the weight is proportional to the product of the marginal utility of income and the slope of the transformation function *g* at the individual's level of wellbeing, $g'u_i' = \frac{dg(u_i)}{du_i}\frac{\partial u_i}{\partial y_i}$.

Deriving a set of weights requires some method for evaluating how the marginal utility of income differs by income level (and perhaps how it differs by attributes included in x_i). Acland and Greenberg (2023) describe this as a technical problem that can be solved using economic methods, in contrast to the normative problem of whether to count gains in wellbeing to different individuals as socially equally valuable (consistent with a utilitarian SWF) or to count gains in wellbeing to some as socially more valuable than gains to others (consistent with a prioritarian or some other SWF). As noted above, marginal utility cannot be inferred from behavior without making additional, strong assumptions. Hence determining how the marginal utility of income differs among circumstances for an individual, or between individuals, is not simply a technical problem. Several methods for measuring marginal utility and their limitations are discussed in Section 2.1. One assumption that is often adopted is that individuals have the same utility function (homogeneous utility) and hence differences in wellbeing depend only on differences in consumption and other circumstances. If the possibility of different (heterogenous) utility functions is recognized, the problem of comparing wellbeing interpersonally is compounded, as discussed in Section 2.2.

2.1 Measuring marginal utility

Several methods for estimating the marginal utility of income have been proposed or applied, but each has substantial limitations. The methods include: assuming an increment of a selected attribute

[5] The utilitarian SWF is the limiting case of the prioritarian SWF as the transformation function approaches linearity; the Rawlsian lexi-min SWF is the limiting case as the transformation function becomes increasingly concave.

[6] The magnitude of the weights is unimportant; what matters is the ratio of the weights for different individuals. For non-marginal changes in income, the weights depend on the average marginal utility and average slope of the transformation function over the relevant intervals.

has equal effects on individuals' wellbeing; evaluating effects of income on self-reported subjective wellbeing; deriving marginal utility from risk aversion or intertemporal choices; and evaluating the progressivity of income-tax schedules. I discuss each in turn.[7]

One approach to evaluating incremental utility of income is to assume that an increment to some attribute of wellbeing (i.e., an increment to x) provides the same increase in wellbeing to all individuals. Using each individual's indifference curves, one can infer the change in income Δy providing the same increment to wellbeing (the compensating or equivalent variation).[8] An individual who is willing to pay or accept a larger amount of money for the increment Δx reveals that her marginal utility of income is smaller than that of an individual whose willingness to pay or accept for that increment is smaller. For example, one can assume that a small increase in the probability of surviving the current period, or a small increase in life expectancy, has the same effect on wellbeing for all individuals and then use estimates of the individuals' value per statistical life (VSL) or value per statistical life year (VSLY) to infer the marginal utility of income (e.g., Courard-Hauri and Lauer 2012, Canning 2023). Note that this approach is equivalent to conducting BCA using the attribute that is assumed to have the same effect on wellbeing as the numéraire (e.g., using net health benefits; Hammitt 2021).

A second approach is to use individuals' self-reported subjective wellbeing (SWB) or "happiness" as a measure of utility, and to evaluate changes or differences in subjective wellbeing as a function of changes or differences in income (Layard et al. 2008, Layard and De Neve 2023). The critical assumption underlying this approach is that individuals use the SWB scale (typically 0 to 10 or 0 to 100) in a way that is linear (each unit corresponds to the same difference in wellbeing for the individual) and is interpersonally comparable (one unit corresponds to the same difference in wellbeing for every individual). An additional concern is that the time period over which individuals are asked to report their SWB is not always clearly defined, e.g. whether it is a measure of anticipated total future or lifetime wellbeing or wellbeing in the current period (Benjamin et al. 2024).

A third approach is to infer the rate at which marginal utility decreases with income from behavior involving risk or intertemporal allocation. Under the conventional model of decision making under uncertainty (expected utility), individuals are assumed to prefer the action that yields the probability distribution over potential consequences with the largest expected utility. Risk aversion implies the

[7] Groom and Maddison (2019) describe these methods and review estimates using each approach.
[8] For an infinitesimal change, $\Delta y_i/\Delta x_i$ equals the marginal rate of substitution between income and the attribute x.

utility function is concave, i.e., exhibits diminishing marginal utility, and estimates of the degree of concavity are often used as estimates of the elasticity of the marginal utility of income (e.g., Acland and Greenberg 2023). This approach is logically flawed.[9] An individual who maximizes the expected utility of a lottery simultaneously maximizes all positive monotone functions of the expected utility. Equivalently, each individual has a certainty equivalent for every lottery (a sure outcome such that she is indifferent between the lottery and its certainty equivalent) and a preference order over sure outcomes, which is represented by an ordinal utility function.

Consider a lottery where the possible consequences are monetary values given by the random variable $\tilde{x}$. Under expected utility, an individual is indifferent between the lottery $\tilde{x}$ and its certainty equivalent $\hat{x}$; i.e., $Eu(\tilde{x}) = u(\hat{x})$, where *u* is a cardinal utility function. Let $v(x) = f[u(x)]$ be a cardinal measure of wellbeing when *x* is a certain consequence and *f* is a strictly positive monotone function. Then $v(\tilde{x}) = f[Eu(\tilde{x})] = f[u(\hat{x})] = v(\hat{x})$. Choice among lotteries provides no information about which possible *v* provides a cardinal measure of the individual's wellbeing. The marginal utility of income, $v'(x) = f'[u(x)] \cdot u'(x)$, is unknown because *f* and *v* are unknown. Expected utility allows one to rank lotteries in preference order but this provides no information about the magnitude of the difference in wellbeing between lotteries. Note that this argument can be extended to alternative theories of decision making under uncertainty, such as prospect theory and rank dependent utility (Kahneman and Tversky 1979, Tversky and Kahneman 1992, Quiggin 1982, 1993). These theories provide alternative models that yield a certainty equivalent for each lottery; once these certainty equivalents are obtained, preferences can be represented by an ordinal utility function that provides no information about strength of preference.[10]

Adler (2016, 2019, 2025) and Broome (2004) argue that any utility function *u* that represents preferences over lotteries should be interpreted as measuring wellbeing. They do so by adding the assumption that preferences are consistent with the Bernoulli axiom, i.e., individuals are risk-neutral in wellbeing, and hence the only admissible transformations *f* are positive affine.[11]

[9] Wakker (1994) provides a useful history of the distinction between utility or wellbeing functions that do and do not incorporate preferences over risky outcomes.

[10] The separation between risk preferences (characterized by *u*) and wellbeing or strength of preference (characterized by *v*) is similar to the separation between risk preferences and intertemporal preferences by Kreps and Porteus (1978) and Selden (1978).

[11] Broome (2004) defines measured wellbeing as consistent with the Bernoulli axiom, while Adler (2016, 2019, 2025) uses the axiom to conclude that wellbeing as measured by judgments about differences and as measured by preferences over lotteries are equivalent cardinal measures.

Wakker (1994) provides an axiomatization of rank-dependent utility[12] in which the utility function (which can be interpreted as a measure of wellbeing) can be estimated using a series of derived tradeoffs, which are judgments about whether substituting x_2 for x_1 in one lottery is preferred to substituting y_2 for y_1 in another lottery.[13] For the degenerate lotteries that yield each outcome for sure, this derived tradeoff is a judgment about preference for one certain change or the other. Note that such a preference cannot be revealed through choice, because the starting points (x_1 and y_1) may differ; the derived-tradeoff approach requires that judgments about differences in wellbeing are consistent with the axioms of measurement theory (Krantz et al. 1971).

Inferring marginal utility from evaluation of intertemporal choices suffers from a similar logical error. The notion behind using intertemporal choice to evaluate marginal utility is to assume an individual wishes to maximize the discounted present value of the utility of consumption in multiple periods. To illustrate, consider the two-period case in which the individual seeks to maximize $v(c_1, c_2) = u(c_1) + \delta u(c_2)$ subject to a budget constraint $c_1 + \rho c_2 = y$, where $u(c_t)$ is the utility of consuming c_t in period t = 1, 2, δ is a utility discount factor (the reciprocal of one plus the utility discount rate), and ρ is an interest factor (the reciprocal of one plus the interest rate). The first-order condition for optimality requires $\frac{u\prime(c_1)}{u\prime(c_2)} = \frac{\delta}{\rho}$. Knowing or estimating the ratio on the right-hand side, one can infer the proportional change in marginal utility between the two consumption levels. But as in the case of risk aversion, this procedure does not identify how the marginal utility of consumption differs. Let $w(c_1, c_2)$ be a strictly positive monotone transformation of $v(c_1, c_2)$, i.e., $w(c_1, c_2) = f[v(c_1, c_2)]$ where $f' > 0$. In this model, w represents wellbeing while the function u could represent preferences for having a smoother consumption path. The marginal wellbeing of an increase in c_1, $\frac{\partial w(c_1,c_2)}{\partial c_1} = f'[v(c_1, c_2)]\frac{\partial v(c_1,c_2)}{\partial c_1}$, is unknown because f and w are unknown.

Another approach to estimating marginal utility is to assume that income taxes in a jurisdiction are selected to maximize a SWF subject to a revenue constraint (Kaplow 2024). Note that this approach identifies the relative marginal social values of income to different individuals, which could differ from the relative marginal utility of income if the SWF is not utilitarian. For example, if taxes are designed to maximize a prioritarian SWF, the rates would be related to the product of the marginal utility of income and the slope of the transformation function g at the individual's level of wellbeing.

[12] Rank-dependent utility (Quiggin 1982, 1993) is a generalization of expected utility in which probabilities are transformed in a way that depends on the rank of the associated outcome among all outcomes in the lottery.
[13] Moreover, all four outcomes must hold the same rank among all outcomes in their respective lotteries.

Optimal tax rates must also account for incentives, and these could distort the tax schedule compared with what would be optimal if individuals could not adjust their taxable income.

2.2 Heterogeneous utility functions

If individuals' utility functions exhibit different degrees of curvature (i.e., different rates at which marginal utility decreases), they cannot have the same utility, or the same marginal utility, at every income level. In the simple case where one utility function is globally more concave than the other, the utility functions can have the same utility level at no more than two income levels, and the same marginal utility at no more than one income level. Hence differences in curvature will affect comparisons of wellbeing at specified consumption levels and the increase in wellbeing from a specified increment to consumption. If differences in curvature reflect differences in risk aversion, then evaluations of non-risky policies can be affected by risk aversion, which seems perverse; one can argue that wellbeing levels at non-risky consumption levels should not depend on risk aversion. If the Bernoulli axiom is adopted, however, risk aversion with respect to income is equivalent to diminishing marginal wellbeing with respect to income, and hence individuals who differ in risk aversion will necessarily differ in marginal utility of income and level of wellbeing.

A common approach to making utility functions interpersonally comparable is to assume that wellbeing is the same for all individuals at each of two reference bundles (i.e., levels of income and other attributes of utility) and assigning wellbeing levels of 0 and 1 to the less- and more-preferred bundle, respectively. Then wellbeing for the individual whose utility is more concave (e.g., because she is more risk averse) will be greater than for the other individual at intermediate consumption levels (between the reference bundles). An alternative approach (Fleurbaey and Zuber 2021) is to assume that wellbeing and the marginal utility of income are equal at some point where the utility functions are tangent (e.g., the poverty line). This has the property that pure income transfers from individuals above the point of tangency to individuals below it always increase utilitarian social welfare, and that individuals with more concave utility are everywhere worse off (except at the tangency point). Choosing a point of tangency requires judging that different individuals are equally well-off given the defined circumstances; for discussion of some of the difficulties in defining poverty lines and assessing their dependence on cost of living and other factors, see Moatsos and Lazopoulos (2021) and Ravallion (2020).

3. Age weighting

Weights in wBCA can be based on attributes other than, or in addition to, income. Societies are often concerned with the wellbeing of older people and large shares of government spending are

allocated to income support and health care for older people. In the United States, for example, Social Security and Medicare (that predominantly benefit people older than 65) account for about one-third of federal spending (U.S. Office of Management and Budget 2024).

The marginal utility of income may vary with age for several reasons. First, older people tend to have a shorter (remaining) life expectancy than younger people, and hence less opportunity for future consumption. This decreases the marginal utility of income. If individuals have little future income and must support themselves from past savings (i.e., in the absence of annuities or pensions), the marginal utility of income increases with life expectancy. Someone with a long life expectancy must husband her resources carefully while someone with a short life expectancy can spend profligately with little risk of outliving her assets.

The change in marginal utility of income over time can depend on saving behavior. Consider the optimal lifetime consumption path in the standard additive, multiperiod model. The individual seeks to maximize the expected present value of the utility of consumption in each period, $\sum_t s_t \delta^t u(c_t)$, where t indexes future periods, s_t is the probability of surviving to period t, δ is a utility discount factor, and $u(c_t)$ is the utility of consuming c_t in period t. If the individual can borrow and save at a common interest rate, her budget constraint is $\sum_t s_t \rho^t c_t \leq b$, where ρ is the interest factor and b is the expected present value of lifetime income. If she discounts future utility at the interest rate ($\delta = \rho$), optimal consumption and the marginal utility of income are constant across periods. It is plausible to assume that individuals should discount future utility at a rate smaller than the interest rate. If so, $\delta > \rho$, c_t increases and the marginal utility of income decreases over time.

Other factors can also affect how the marginal utility of income changes with age. Health tends to decline with age. If medical-care and other health-related expenses are covered by insurance or other sources, it is reasonable to assume that the marginal utility of income (consumption) is smaller when health is worse, as poor health limits the ways one can enhance utility through consumption (except for health care). Several empirical papers support this assumption (Finkelstein et al. 2013, Sloan et al. 1998, Viscusi 2019, Viscusi and Evans 1990). In contrast, adverse health shocks that increase health-care costs will increase the marginal utility of income. Household or family structure is also relevant: older people are less likely to be financially responsible for children, decreasing the marginal utility of income.

Weights can incorporate not only differences in the marginal utility of income, but also differences in the social value of improving wellbeing of different individuals (e.g., under a prioritarian SWF). For comparing older with younger people, it is useful to consider judgments about allocation of

lifesaving resources to individuals of different ages. The fair innings principle (Harris 1985, Daniels 1988, Williams 1997, Bognar 2008, 2015, Ottersen 2013) asserts that older people have a weaker claim than younger people to (public) lifesaving resources. As first articulated by Harris (1985, p. 91), "The fair innings argument requires that everyone be given an equal chance to have a fair innings, to reach the appropriate threshold but, having reached it, they have received their entitlement. The rest of their life is the sort of bonus which may be canceled when this is necessary to help others reach the threshold." Bognar (2015) challenges the notion of a threshold and defines the fair innings principle as an expression of inequality aversion with respect to longevity, consistent with a prioritarian SWF applied to life years.

Although the fair innings principle applies to lifesaving, since longevity is a component of wellbeing it implies that the component of weights that reflects the social value of improving wellbeing should be smaller for older than for younger people. Assuming that individual wellbeing is an increasing function of completed longevity and average annual income, Adler et al. (2021) show that weights associated with improving wellbeing are a decreasing function of age under both utilitarian and prioritarian SWFs. Weights increase with income under the utilitarian SWF but can increase or decrease with income under a prioritarian SWF.

Studies of public attitudes toward the fair innings principle and inequality aversion with respect to completed longevity yield mixed results. Adler et al. (2025) asked a panel of 2000 French residents a series of questions in which they were required to allocate a single treatment to either of two patients, who differ in age. When choosing to allocate a treatment that would provide the same life extension to an older or a younger patient, about one-half the respondents would allocate the treatment to the younger patient while one-third would treat the patients equally and one-fifth would allocate treatment to the older patient. Of the 65 percent of respondents whose responses to multiple questions that varied the benefits of treating the patients are consistent with a SWF, 40 percent were consistent with fair innings, 30 percent would treat the patients equally, 13 percent would treat whichever patient would gain more life years (or would treat the patients equally when the gains were similar), and 17 percent would treat the older patient. In a systematic review of studies in the United Kingdom, McNamara et al. (2020) found mixed support for inequality aversion, which tends to be stronger for differences in longevity than in health and stronger when longevity inequalities are positively associated with inequalities in socio-economic status. Using a survey with nearly 2000 participants in Ontario, Canada, Hurley et al. (2000) found a bi-modal distribution of inequality aversion for health-adjusted life expectancy: nearly half their respondents display little or no inequality aversion while the other have display substantial inequality aversion.

In summary, holding income constant, weights seem likely to be smaller for older than younger individuals. The marginal utility of income seems likely to be smaller for older individuals, for a variety of reasons. In addition, the fair innings principle suggests that the social value of improving wellbeing is smaller for older than for younger individuals. This principle, however, draws significant but far from universal support by the public.

4. Conclusion

Given that there is no solid approach to estimate marginal utility of income and how it differs between individuals, one must recognize that the choice of weights for use in wBCA must be judgmental. In making these judgments, it may be useful to consider the issue from multiple perspectives. One perspective would be to consider the social value of increasing income for individuals at different income levels (and different levels of other characteristics) directly. An alternative perspective would be to divide the problem into two components: how income affects wellbeing, and the social value of enhancing wellbeing for people at different wellbeing levels, recognizing that judgments about the first component cannot be confirmed or refuted by empirical evidence. Judgments may be informed by empirical evidence about how SWB correlates with income and other characteristics, such as health status, time spent in different activities, and quality of social relations, that one would anticipate to correlate with wellbeing. Drupp et al. (2018) surveyed 200 economic experts to elicit their judgments about the elasticity of marginal utility (and other factors relevant to long-term discounting).

As noted above, weights can also be used to reflect judgments about the social value of improving wellbeing to individuals as a function of characteristics other than income, including health, age, gender, race, ethnicity, geographic region, historical disadvantage, and others. Given that one must rely on judgment, it seems reasonable to restrict the arguments of the weighting function to attributes that are reasonably observable, such as an objective list of attributes that affect the wellbeing an individual can achieve, such as primary goods advocated by Rawls (1971), capabilities advocated by Sen (1979, 1989) and Nussbaum (1988, 2011), or a measure of opportunity advocated by Sugden (2004). In applied work, income, health, and life expectancy have been used (Ferranna et al. 2024). Attempting to incorporate differences in the ability of individuals to convert resources into wellbeing seems impractical, as illustrated by the difficulty of identifying or rejecting the possibility of utility monsters (who can convert tiny gains in resources to enormous gains in wellbeing) and their opposite, unhappy rich princes (Nozick 1974). Ideally, the arguments of the weighting function should be exogenous or one should consider how to adjust for differences in attributes for which the

individual is at least partly responsible (e.g., income, education, and health depend in part on effort and investment).

In assigning weights to individuals judgmentally, Sher (2024) identifies an important constraint. In response to Saez's and Stantcheva's (2016) proposal to choose weights based on any attributes that are deemed justified, Sher (2024) shows that this approach yields intransitive policy rankings unless the weights are "structurally utilitarian," i.e., consistent with a "generalized utilitarian" SWF. This implies that the weight attached to an individual is determined by her marginal utility of consumption (in a model that distinguishes consumption from income, with the difference being income tax). Sher's result imposes constraints on the weights that must be respected in order to avoid incoherent policy rankings.

References

Acland, D.J., and D.H. Greenberg, "Distributional Weighting and Welfare/Equity Tradeoffs: A New Approach," *Journal of Benefit-Cost Analysis* 14(1): 68–92, 2023.

Adler, M.D., *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*, Oxford: Oxford University Press, 2012.

Adler, M.D.. "Extended Preferences," M.D. Adler and M. Fleurbaey, eds., *The Oxford Handbook of Well-Being and Public Policy*, pp. 476– 517. Oxford: Oxford University Press, 2016.

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

Adler, M.D., *Risk, Death, and Well-Being: The Ethical Foundations of Fatality Risk Regulation*, Oxford: Oxford University Press, 2025.

Benjamin, D.J., K. Cooper, O. Heffetz, and M. Kimball, "From Happiness Data to Economic Conclusions," *Annual Review of Economics* 16: 359‑91, 2024.

Bognar, G., "Age-Weighting," *Economics and Philosophy* 24: 167–189, 2008.

Bognar, G., "Fair Innings," *Bioethics* 29: 251–261, 2015.

Broome, J., *Weighing Lives*, Oxford: Oxford University Press, 2004.

Canning, D., "Conducting Cost Benefit Analysis in Expected Utility Units Using Revealed Social Preferences." AWI Discussion Paper No. 722, 2023.

Courard-Hauri, D., and S.A. Lauer, "Taking 'All Men Are Created Equal' Seriously: Toward a Metric for the Intergroup Comparison of Utility Functions Through Life Values" *Journal of Benefit-Cost Analysis* 3(3): article 3, 2012.

Daniels, N., *Am I My Parents' Keeper? An Essay on Justice Between the Young and the Old*, Oxford: Oxford University Press, 1988.

Drupp, M.A., M.C. Freeman, B. Groom, and F. Nesje, "Discounting Disentangled," *American Economic Journal: Economic Policy* 10 (4): 109–134, 2018.

Ferranna, M., J.K. Hammitt, and L.A. Robinson, "From Benefit-Cost Analysis to Social Welfare: A Pragmatic Approach," *Journal of Benefit-Cost Analysis* 15(S1): 84–109, 2024.

Finkelstein, A., E.F. Luttmer, and M.J. Notowidigdo, "What Good is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption," *Journal of the European Economic Association* 11(s1): 221–258, 2013.

Fleurbaey, M., and J.K. Hammitt, "The Right Numeraire or the Just Weights? How to make BCA Rational and Fair," *Journal of Benefit-Cost Analysis* 15(S1): 6–25, 2024.

Fleurbaey, M., and S. Zuber, "Fair Utilitarianism," *American Economic Journal: Microeconomics* 13(2): 370–401, 2021.

Groom, B., and D. Maddison, "New Estimates of the Elasticity of Marginal Utility for the UK," *Environmental and Resource Economics* 72: 1155–1182, 2019.

H.M. Treasury, *The Green Book: Central Government Guidance on Appraisal and Evaluation*, London, 2020.

Hammitt, J. K., "Accounting for the Distribution of Benefits and Costs in Benefit–Cost Analysis." *Journal of Benefit-Cost Analysis*, 12: 64–84, 2021.

Kahneman, D., and A. Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47: 263-291, 1979.

Kaplow, L., "Optimal Income Taxation," *Journal of Economic Literature* 62(2): 637–738, 2024.

Krantz, D.H., R. D. Luce, P. Suppes, and A. Tversky, *Foundations of Measurement*, vol. 1 (*Additive and Polynomial Representations*), New York: Academic Press, 1971.

Kreps, D.M., *Notes on the Theory of Choice*, New York: Routledge, 1988.

Kreps, D.M., and E.L. Porteus, "Temporal Resolution Of Uncertainty and Dynamic Choice Theory," *Econometrica* 46: 185-200, 1978.

Layard R., and J-E. De Neve, *Wellbeing: Science and Policy*, Cambridge: Cambridge University Press, 2023.

Layard R., G. Mayraz, and S. Nickell, "The Marginal Utility of Income," *Journal of Public Economics* 92: 1846–1857, 2008.

Moatsos, M., and A. Lazopoulos, "Global Poverty: A First Estimation of its uncertainty," *World Development Perspectives* 22: 100315, 2021.

Nozick, R., *Anarchy, State, and Utopia*, New York: Basic Books, 1974.

Nussbaum, M., "Nature, Function, and Capability: Aristotle on Political Distribution," In *Oxford Studies in Ancient Philosophy*, Oxford: Oxford University Press, 1988.

Nussbaum, M., *Creating Capabilities: The Human Development Approach*, Cambridge: Harvard University Press, 2011.

Ottersen, T., "Lifetime QALY Prioritarianism in Priority Setting," *Journal of Medical Ethics* 39: 175–180, 2013.

Quiggin, J., "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization* 3: 323-343, 1982.

Quiggin, J., *Generalized Expected Utility Theory*, Dordrecht: Kluwer, 1993.

Ravallion, M., "On Measuring Global Poverty," *Annual Review of Economics* 12(1): 167-188, 2020.

Rawls, J., *A Theory of Justice*, Cambridge: Harvard University Press, 1971.

Saez, E., and S. Stantcheva, "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review* 106(1): 24–45, 2016.

Selden, L., "A New Representation of Preferences over 'Certain x Uncertain' Consumption Pairs: The 'Ordinal Certainty Equivalent' Hypothesis," *Econometrica* 46: 1045-1060, 1978.

Sen, A., *Equality of What?* Stanford University: Tanner Lectures on Human Values, 1979.

Sen, A., "Development as Capability Expansion," *Journal of Development Planning* 19: 41–58, 1989.

Sloan, F.A., W.K. Viscusi, H.W. Chesson, C.J. Conover, and K. Whetten-Goldstein, "Alternative Approaches to Valuing Intangible Health Losses: The Evidence for Multiple Sclerosis," *Journal of Health Economics* 17(4): 475–497, 1998.

Sugden, R., "The Opportunity Criterion: Consumer Sovereignty Without the Assumption of Coherent Preferences," *American Economic Review* 94(4): 1014-1033, 2004.

Tversky, A., and D. Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5: 297-323, 1992.

U.S. Office of Management and Budget, *Budget of the U.S. Government: Fiscal Year 2025*, Washington, D.C.: U.S. Government Publishing Office, 2024.

U.S. Office of Management and Budget, *Circular A-4: Regulatory Analysis*, Washington, D.C., 2003.

U.S. Office of Management and Budget, *Circular A-4: Re/gulatory Analysis*, Washington, D.C., 2023.

Viscusi, W.K., "Utility Functions for Mild and Severe Health Risks," *Journal of Risk and Uncertainty* 58 (2–3): 143–166, 2019.

Viscusi, W.K., and W.N. Evans, Utility Functions that Depend on Health Status: Estimates and Economic Implications," *American Economic Review* 80(3): 353–374, 1990.

Williams, A., "Intergenerational Equity: An Exploration of the 'Fair Innings' Argument," *Health Economics* 6: 117–132, 1997.