

ARTICLE

The Elasticity of Marginal Utility of Income for Distributional Weighting and Social Discounting: A Meta-Analysis

Daniel Acland¹ and David H. Greenberg²

¹Goldman School of Public Policy, University of California, Berkeley, CA, USA

²Department of Economics, University of Maryland, Baltimore, MD, USA

Corresponding author: Daniel Acland; Email: acland@berkeley.edu

Keywords: distributional weights; elasticity of marginal utility; meta-analysis; social discount rate

JEL classifications: D30; D61; D63

Abstract

Estimates of the elasticity of the marginal utility of income are necessary for determining distributional weights to correct for diminishing marginal utility of income, which is particularly important in light of increasing concern about accounting for distributional impacts in regulatory review. The elasticity is also necessary for computing the social discount rate using the Ramsey formula. Despite many attempts to estimate the elasticity of the marginal utility of income, considerable uncertainty exists about the magnitude of this key parameter. In this paper, we use meta-analysis of estimates of the elasticity from the US and UK to shed light on the appropriate elasticity values to use for both distributional weighting and discounting. Relying on our findings, we tentatively conclude that it is reasonable to base the social discount rate and distributional weights on an elasticity of 1.6, with lower- and upper-bound sensitivity testing at 1.2 and 2.0. This estimate results in distributional weights which appear plausible, and which we believe can contribute to a consensus on how to conduct distributional weighting. Moreover, the resulting social discount rate is within the range typically recommended when the Ramsey formula is used.

1. Introduction

Benefit–cost analysis (BCA) measures individual welfare using willingness to pay (WTP) for whatever impacts an individual experiences as a result of a policy. It aggregates this measure across the affected population without regard to the income of the individuals experiencing the impacts. A dollar’s worth of welfare has the same weight in BCA regardless of whether it is the welfare of a poor person or the welfare of a wealthy person. Two concerns have been expressed with respect to this “unit weighting” of WTP. The first is that a dollar’s worth of consumption is typically considered to have a greater impact on the welfare of the poor than on the welfare of the wealthy—due to diminishing marginal utility of income—

which means that any given quantum of welfare is represented by a smaller number of dollars if it accrues to a poor person than if it accrues to a wealthy person. As a result, BCA is actually biased against the poor and thus may promote inequitable strategies for increasing aggregate welfare. The second concern is that, even if the bias in WTP as a measure of welfare were not an issue (which is to say, even if the marginal utility of income were constant with respect to income), BCA would still place equal weight on the welfare of the poor and the welfare of the wealthy, and thus take no account of the possibility that decision makers and members of society may simply place greater moral weight on the welfare of the poor, a phenomenon that has been referred to as “diminishing marginal moral value of welfare” (Adler, 2013).

The idea of addressing either or both of these two concerns using distributional weights to inflate the WTP of the poor and deflate the WTP of the wealthy has been around for at least six decades (see, e.g., Arrow, 1963; Feldstein, 1972; Harberger, 1978; Boadway & Bruce, 1984; Drèze & Stern, 1987; Blackorby & Donaldson, 1990; Layard & Glaister, 1994; and Adler, 2013). Recently, interest in distributional weighting has risen, in response to a memorandum from the Biden White House calling for ways to address distributional concerns in regulatory review. The memorandum is silent about what kinds of strategies could be used, but it seems clear that substantial changes to current practice are called for, and distributional weighting does not seem beyond the pale. In practice, however, the application of distributional weights has barely gotten started. The UK Treasury Green Book (2018) provides guidance on computing and applying weights, but we are not aware of any BCA conducted within the British government that adopts the guidance. Similarly, there are no instances of distributional weighting in regulatory review by agencies of the US federal government (Robinson *et al.*, 2016).

Fleurbay and Abi-Rafeh (2016) provide an important possible explanation for this absence of weighting in practice: “The real reason for the lack of success of weighted BCA is most likely the belief among BCA practitioners that there is no simple recipe for determining the... weights.” There are two complicating factors. First, there is a belief among some that selecting weights involves a value judgment about the relative importance of efficiency and equity that should not be made by analysts (Banzhaf, 2011; Adler, 2013). Second, even if there were general acceptance of the proposition that it is appropriate for analysts to select weights, there would be no consensus on precisely what weights to select.

We have addressed the first issue at length in Acland and Greenberg (2023), where we argue that value judgments only enter when addressing differential moral concern for the welfare of the poor versus the wealthy, so that the question of whether to apply weights to address that concern (what we call “equity-weights”) is quite different from the question of whether to apply weights to address the bias introduced by diminishing marginal utility of income (what we call “utility-weights”). The essence of our argument is that utility-weighting is simply a correction of a methodological problem with the measurement of individual welfare, and thus that it is entirely appropriate for analysts to select weights. We argue—though we are not the first—that weights should be selected based on empirical observations of individual preferences, which is the topic of this paper. By contrast, additional weighting by analysts to address disproportionate concern for the welfare of the poor is inappropriate because a) it does indeed insert moral or ethical values into the analysis, thus taking the assessment of tradeoffs between welfare and equity out of the hands of democratically accountable decision makers, to whom that responsibility should be assigned in a deliberative democracy; and b) it combines information about welfare and information about equity into a single metric and thus renders invisible the separate information about each value that decision makers need in order to make their own

assessments of the relative importance of the welfare and equity impacts of different policies. In this paper, we take this position as given and refer the unpersuaded reader to our earlier paper.¹

The purpose of this paper is to address Fleurbaey and Abi-Rafeh's second reason for the lack of progress on the practical application of distributional weights, that there can be no consensus on the appropriate weights to apply. In the case of utility-weighting, we believe there is a solid foundation for a consensus.

There is already a reasonable degree of consensus on the correct way to compute weights, using the ratio of the marginal utility of income of each income level to that of a reference income level, usually the median (Cowell & Gardiner, 1999; Adler, 2013; Nurmi & Ahtiainen, 2018). Using ratios overcomes the problem of the scale-invariance of utility of income. Furthermore, in the literature on distributional weighting, and indeed the literature on diminishing marginal utility itself, there is a considerable degree of agreement that the most practicable utility function to use for computing marginal utilities is an iso-elastic, or constant relative risk aversion utility function, which we explicate below (Kaplow, 2010; Adler, 2013; Nurmi & Ahtiainen, 2018). This utility function has a single parameter, the elasticity of marginal utility of income, with which the appropriate weights can easily be computed. What remains is to determine the appropriate value of that parameter. To accomplish this task, we have conducted a rigorous and comprehensive meta-analysis of 1,711 estimates of the elasticity, from 158 independent studies over four decades.

Contrary to a commonly expressed concern about this parameter, the estimate from our meta-analysis is tightly bounded, with a mean of 1.61 and a 95% confidence interval of 1.18 to 2.05. The 1.61 estimate implies that in distributionally weighted BCA the benefits and costs of households of average size, with incomes at the 25th percentile, should receive a weight that is a little over two and a half times that for households at the median, while the benefits and costs of those at the 75th percentile should receive a weight that is half that for households at the median. We discuss the plausibility of these weights in Section 6.2.

In addition to distributional weights, as discussed below, the marginal utility of income plays a critical role in determining the social discount rate whenever the rate relies on the optimal growth rate method originally developed by Ramsey (1928). This is because the Ramsey formula accounts for the fact that as incomes grow over time, the marginal value of income decreases, so that costs and benefits that occur in the future will generate less welfare for those who experience them than those that occur in the present.² Our estimate of the elasticity of the marginal utility of income implies a social discount rate of approximately 4.0%, with a lower bound of 3.3% and an upper bound of 4.8%.

There have been two meta-analyses that cover some of the same ground as ours yet differ in important ways. Like us, Groom and Maddison (2019) conduct a meta-analysis of the

¹ It is important to note that we emphatically do not consider our recommendation for utility-weighting to be a fully adequate response to the Biden Memorandum. It does not actually account for the distribution of welfare at all, only the bias caused by diminishing marginal utility of income. Many other strategies have been proposed to account for the actual welfare-distributional impacts of policies, including some that we have proposed (Acland & Greenberg, 2023). Our utility-weighting recommendation in no way obviates the need for other tools for addressing distributional concerns. That said, we are unaware of any alternative proposals for overcoming the bias we are concerned with.

² The component of the Ramsey formula that accounts for diminishing marginal utility of income as incomes grow over time is implicitly a form of utility-weighting, as it converts the dollar valuations of those with higher income (future generations) into the comparable dollar valuations of a reference income level, that of the current generation.

elasticity of the marginal utility of income. However, their analysis is limited to six estimates for the United Kingdom (five of their own and one from Layard *et al.*, 2008). In contrast, we have attempted to find as many estimates as possible for the US and UK, obtaining over 1,700 in all. The meta-analysis conducted by Havranek *et al.* (2013) is based on a large number of estimates. However, it is focused on the elasticity of intertemporal substitution, which is the inverse of the elasticity of the marginal utility of income. The inverse of an average is not equal to the average of the inverse. Consequently, it is limited to only one of several methods for obtaining estimates of the elasticity of the marginal utility of income, while our meta-analysis includes several others. Moreover, as discussed in greater detail later, it has a different objective than ours (examining reasons why the elasticity of intertemporal substitution varies across countries rather than obtaining an up-to-date value that can be used for utility-weighting and discounting by BCA practitioners), and it is based on a very different methodology.

An important contribution of our study is that we conduct meta-regression to assess the significance of the various apparent differences among estimates generated using different methodologies and/or data, in different locations, and at different times. A key result is that regardless of methodology, time period, or location, there does not appear to be a wide disparity in estimates, and those that exist are statistically insignificant, suggesting that one can have some confidence in adopting our recommended elasticity values.

The following section discusses the role the elasticity of the marginal utility of income plays in determining distributional weights and discount rates. Section 3 reviews the variety of methods that have been used to estimate the elasticity. Section 4 describes the sample of elasticity estimates used in the meta-analysis, while Section 5 describes the methods that were used. Section 6 presents the findings from the meta-analysis. Section 7 discusses our conclusions from the meta-analysis.

2. Determining utility-weights and discount rates

2.1. Utility weights

As stated in the introduction, the strategy we adopt for addressing the diminishing marginal utility of income is to assign weights to the WTP of each individual that are based on the ratio of their marginal utility of income to that of some reference income—such as the median—and multiply estimates of costs and benefits to individuals at different income levels by these weights (Cowell & Gardiner, 1999; Adler, 2013; Nurmi & Ahtainen, 2018). These weights convert the dollar valuation of policy impacts for individuals at any given income level to the valuation they would have if their income was at the median. As stated, we refer to such weights as “utility-weights.” Once utility-weighting is completed, and diminishing marginal utility is thereby accounted for, the utility-weighted net benefit of a policy or program is the net benefit that would be estimated if everyone affected by the policy had median income. On this basis, the welfare impact of a policy or program can be assessed without bias.³

As just indicated, determining utility weights requires estimates of the marginal utility of income at different income levels, and in order to make this feasible, it is necessary to make an assumption about the functional form of utility. In the methodologies that have been used

³ Because the weights are based on a ratio of marginal utilities, the fact that the utility function is invariant to affine transformations does not affect the weights.

to estimate the marginal utility of income, which are described in Section 3, the standard assumption is that preferences can be represented by an isoelastic utility function:

$$U(y) = \frac{y^{1-\varepsilon}}{1-\varepsilon} \text{ if } \varepsilon \neq 1 \text{ and } U(y) = \ln(y) \text{ if } \varepsilon = 1, \quad (1)$$

where y is income and ε is the elasticity of marginal utility with respect to income.⁴

Given this utility function, the marginal utility of income is

$$\frac{\partial U}{\partial y} = \frac{1}{y^\varepsilon}. \quad (2)$$

Because ε is the elasticity of marginal utility of income with respect to income, it indicates the percentage decrease in the marginal utility of income caused by a 1% increase in income. The isoelastic function has the desirable property of constant elasticity, which allows for straightforward empirical estimation.⁵

Given this utility function and an estimate of ε , utility weights can be readily determined. For instance, setting the utility weight for those at median income (the m th group) equal to one (which is to say, choosing median income as the reference income), the utility weight (w_i) for a higher or lower income group (the i th group) would be the ratio of the income for the m th group to the income of the i th group, taken to the ε th power:

$$w_i = \frac{\frac{\partial U_i}{\partial y_i} / \frac{\partial U_m}{\partial y_m}}{\frac{\partial U_i}{\partial y_i} / \frac{\partial U_m}{\partial y_m}} = \left(\frac{y_m}{y_i} \right)^\varepsilon. \quad (3)$$

The information necessary to compute $\frac{y_m}{y_i}$ for US households is readily obtainable from federal government statistics such as the US Census Bureau's Annual ASEC Survey, but computing utility weights for various income groups also requires a value for ε . While a substantial number of estimates of ε are available from previous research, findings vary considerably across studies and the estimates from any one individual study are far from definitive. This is one of our motivations for conducting a meta-analysis of the available estimates of ε .

⁴ We are aware that theoretically we should be concerned with "marginal utility of consumption" rather than "marginal utility of income." However, because data allowing individuals who are affected by programs and policies evaluated through BCA to be weighted on the basis of their consumption are limited, income is usually used instead. For purposes of utility-weighting, therefore, it is necessary to assume that the marginal utility of total consumption is diminishing in approximately the same way as marginal utility of income.

⁵ In the literature that promotes additional weighting to address differential moral concern for the welfare of the poor, the utility generated by income becomes the argument of another iso-elastic function that maps individual utility into what is sometimes referred to as "social welfare." (See, for example, Boadway & Bruce, 1984; Adler, 2013; and Nurmi & Ahtiainen, 2018.) The resulting, combined function, which nests one iso-elastic function inside another, can be thought of as representing the "social welfare" of any given level of individual income. Its elasticity can be thought of as the elasticity of "marginal moral value of income." If the elasticity of the marginal moral value of welfare is γ , the elasticity of the iso-elastic social welfare function is $\gamma + \varepsilon + \gamma\varepsilon$. Our approach is mathematically equivalent to setting $\gamma = 0$, but that mis-represents our position, because we are agnostic as to the appropriate moral weight to place on the welfare of the poor and are simply arguing that any such differential moral concern should not be addressed via weighting at all but should be addressed by some other form of analysis. We have made recommendations regarding an alternative form of quantitative analysis in Acland and Greenberg (2023). There are others that we review in the same paper.

2.2. Social discount rates

Some of the approaches to determining the social discount rate involve using market rates of interest as proxies for the social opportunity cost of capital, using the after-tax return on savings as a proxy for the social opportunity cost of foregone consumption, or using a mix of the two (see Boardman *et al.*, 2018). An alternative approach, one that is recommended in the UK Green Book (2018) and the European Union's guidelines for BCA (European Commission, 2014), involves an optimal growth rate method originally developed by Ramsey (1928). Ramsey's work produced the following equation, which is known as the Ramsey formula or Ramsey rule, for determining the social discount rate, d :

$$d = \rho + g\varepsilon, \quad (4)$$

where ρ is the rate of pure time preference, g is the rate of growth of income, and ε is the absolute value of the elasticity of the marginal utility of income with respect to changes in income. It is argued that the Ramsey rule is appropriate for the appraisal of public projects because by consistently following, it society will maximize the present value of utility from its current and future per capita consumption (Feldstein, 1964; Lind, 1982 and Moore *et al.*, 2004). Equation (4) presumes that the consumption that results from future income should be discounted (1) because of impatience in delaying consumption, as reflected by ρ ; and (2) in order to address the fact that the welfare value of the marginal dollar of income will in fact be lower in the future, assuming that future income will exceed current income because of economic growth (that is, assuming that g is positive). Thus, $g\varepsilon$ is the predicted percentage reduction in the marginal utility of income. The greater the extent to which individuals have a declining marginal utility of consumption—that is, the more positive the value of ε —the larger the discount rate will be. As a consequence, future increments to income and consumption will be given less weight in BCAs than current increments to income and consumption. But to use the Ramsey formula requires an estimate of ε . The meta-analysis will hopefully contribute to this goal.

3. Methods for estimating the elasticity of the marginal utility of income

The studies from which we take the estimates of ε that we use in our meta-analysis use a variety of methods. We next briefly describe and assess each of these methods.⁶

3.1. Indirect behavioral evidence: lifetime consumption models

The most frequently used approach by far to estimate ε relies on either micro survey panel data or (much more frequently) macro time-series data. The theory that underpins this approach is that households allocate consumption over time in order to maximize a multi-period utility function. Exactly how they allocate their consumption depends on their rate of time preference and the curvature of their utility of income function. Under the assumption that individuals allocate such that their rate of time preference is equal to the interest rate,

⁶ While most attempts to estimate ε assume constant elasticity with respect to income, there are important exceptions (e.g., Blundell *et al.*, 1994; Crossley & Low, 2011 and Demuyck, 2018). This does not affect findings from the meta-analysis, however.

observing how consumption is affected by changes in the interest rate allows inferences to be made about the shape of the utility function.

In utilizing time series data, researchers typically use a log-linearized consumption Euler equation: $\ln(g_t) = c + br_t + e_t$, where g_t is the growth rate in consumption at time t ; c is the constant; r_t is the real rate of return on assets (e.g., the stock market return or the treasury bill return); coefficient b provides an estimate of the elasticity of intertemporal substitution (EIS),⁷ which is the inverse of ε ; and e_t is the error term. Because the causality goes in both directions (i.e., consumption growth affects the return on assets), instruments are usually used for r_t . As can be seen, if r is viewed as a proxy for the social discount rate, d in Equation (4), the regression that is used bears a close kinship with Equation (4), the Ramsey equation, once Equation (4) is divided by ε , and the ρ parameter, the rate of pure time preference, is set to zero. Indeed, sometimes, but much less often, there is a direct attempt to estimate ε by regressing g_t on r_t , although this approach is more subject to weak instruments (see Yogo, 2004).

The estimation is sensitive to model specification and to the choice of the interest rate used to proxy for d (Evans, 2008). Estimation also requires several assumptions, arguably the most important of which is that capital markets are perfect (e.g., market participants are rational and there are no constraints on borrowing or on information) and that the utility function is additively separable. The estimates may depend on the period over which they are computed, especially if market turbulence or changes in financial regulations occur. They may also be sensitive to which goods are used in the demand estimation or consumer preferences for the goods may change over time. Additionally, when the estimates rely on macro-data, rather than micro-data, it is not possible to take account of changes in demographic composition or household behavioral changes over the lifespan.

3.2. Indirect behavioral evidence: consumer demand analysis

This approach focuses on how utility is affected by the consumption of a particular good or broad category of goods (e.g., all foods). Essentially, the estimate of ε is the ratio of the good's income elasticity to its own-price elasticity, with an adjustment for the share of the total budget accounted for by the good. That is, the following formula, which was developed by Frisch (1959), is used: $\varepsilon = y(1 - wy)/p$, where y is the income elasticity of demand, p is the compensated own-price elasticity of demand, and w is the budget share accounted for by the good. The empirical approaches vary, sometimes using single-equation models and sometimes based on complete demand systems. Time series data are typically used.

Like the lifetime consumption approach, findings may be sensitive to the specific specification of the regression model used to estimate ε . The composition of the category of goods is assumed to be additively separable—that is, additional utility from consuming the good does not depend on the quantity of any other good that is consumed. Thus, the good selected is often “all foods.” In the case of food, for example, the assumption of additive separability would hold if utility from food consumption for an individual is unaffected by

⁷ Because the EIS indicates households' willingness to substitute consumption between time periods in response to the interest rate, it is of critical interest in determining monetary policy, which is an important reason there have been so many attempts to estimate it.

the value of the house in which the individual lives.⁸ Obviously, there is no guarantee that the assumption is met, although some researchers have tested it and failed to reject it (Selvanathan, 1988).

3.3. Indirect behavioral evidence: relative risk aversion models

A few researchers have attempted to use insurance data to estimate the so-called coefficient of relative risk aversion, a measure that determines WTP to avoid risks to income, which corresponds to the elasticity of the marginal utility of income if relative risk aversion is independent of wealth.⁹ The model presumes a rational insurance market in which consumers have knowledge of the probability that they will make a claim. The model used to estimate the coefficient of relative risk aversion was developed by Szpiro (1986) and uses time series insurance data.

Szpiro (1986) begins with the following relationship, which is derived in a technical appendix:

$$I = W - \lambda/r(W),$$

where I is the amount of insurance, W is wealth, $r(W)$ is absolute risk aversion, and λ is insurance loading (i.e., the ratio of premiums to claims minus one). Szpiro (1986) then establishes the relationship between absolute risk aversion and relative risk aversion ($\varepsilon(W)$) as

$$r(W) = \frac{\varepsilon(W)}{W}.$$

Under the assumption of constant relative risk aversion (which is the same as isoelastic utility)—an assumption which Szpiro (1986) empirically examines and finds cannot be statistically rejected—substituting this expression into the first equation yields

$$I = W - (\lambda W/\varepsilon).$$

The negative sign in this model reflects the assumption that while the amount of insurance increases with wealth, it declines with insurance loading.

Because the amount of insurance, I , cannot be directly observed, but total claims, Q , and the claim rate, q , can be, Szpiro (1986) substitutes Q/q for I to get $Q = qW - (q/\varepsilon)(\lambda W)$, resulting in the regression equation

$$Q_i = \beta_0 W_i + \beta_1 \lambda_i W_i + e_i,$$

The elasticity can then be computed as $\varepsilon = -\frac{\beta_0}{\beta_1}$.

3.4. Relating subjective well-being to income

A number of happiness surveys have been conducted in recent years in which respondents rate their life satisfaction or “happiness” on a scale of 1–5 or 1–10. To use these findings to estimate ε , analysts must first make an assumption about the functional relationship between

⁸Note that using “all foods” in this analysis is not the same as using food as a proxy for all consumption. Under additive separability, the estimate of ε does not depend on what category of consumption is used.

⁹Put another way, the constant relative risk aversion (CRRA) utility function is identical to the iso-elastic utility function used to estimate ε in other methodologies.

the reported life satisfaction of individuals and their utility. In particular, the assumption typically made is that happiness is linearly linked to utility through a transformation: $h_{it} = u_{it} + v_{it}$, where h_{it} is the measure of happiness for individual i at time t , u_{it} is individual i 's utility at time t , and v_{it} is an error term that is independent of the circumstances affecting utility (see Layard *et al.*, 2008). Given this transformation, and a measure of the utility of the surveyed respondents, ε can be estimated through a modification of Equation (1). This approach appears to have been used only once (Layard *et al.*, 2008) but this study is frequently cited and is the single estimate upon which the UK government bases its recommended utility-weights (Treasury, 2018).

The extent to which individuals accurately measure their happiness is obviously problematic, as is the extent to which what they say reflects their cardinal utility. Moreover, the relationship between measured happiness and utility may not be linear. Nonetheless, Layard *et al.* (2008) used different data sets to conduct their analysis for different countries (including both the US and the UK) and obtained surprisingly consistent results.

3.5. Asking experts

One way to determine the value of ε is to survey persons who are familiar with the concept, an approach that to the best of our knowledge has been used only once (Drupp *et al.*, 2018). It is not clear how even experts could knowledgeably guess at the value unless they were aware of at least a few previous studies or had values in mind for d , ρ , and g and used Equation (4) to obtain a value for ε .

3.6. Progressivity of income taxes

An estimate of ε can be obtained by using progressivity in income tax schedules under an assumption of equal sacrifice. In other words, that society's (or at least policymakers') aversion to inequality is such that income tax rates are set so that the loss in welfare is equal among taxpayers at different income levels. As shown by Groom and Maddison (2019) and Evans (2008), operationally, the equal sacrifice assumption implies that the following equation must hold for all income levels:

$$U(y) - U[y - T(y)] = k,$$

where k is a constant, y is income before taxes, and $T(y)$ is the total tax liability according to the income tax schedule. By assuming an isoelastic utility function and then substituting it into the above equation, the following relationship is obtained:

$$\frac{y^{1-\varepsilon}}{1-\varepsilon} - \frac{([y - T(y)]^{1-\varepsilon})}{1-\varepsilon} = k.$$

By totally differentiating this expression with respect to y and rearranging the terms, the following equation is derived:

$$1 - MTR = (1 - ATR)^\varepsilon,$$

where MTR is the marginal tax rate and ATR is the average tax rate. Taking logs of this expression and solving for ε results in an equation that can be used to estimate ε :

$$\varepsilon = \frac{\ln(1 - MTR)}{\ln(1 - ATR)}.$$

This approach has been used fairly frequently, but the equal sacrifice assumption itself is subject to obvious criticisms. For example, in addition to possible concerns about inequalities in the income distribution, in setting tax rates policymakers likely also consider potential disincentive effects and the goals of various interest groups. Application of the model is also not entirely straightforward. For example, decisions must be made about which taxes to include and how to compute the marginal and average tax rates. Moreover, tax rates change over time and differ for different demographic groups.

3.7. Convex time budget experiments

This approach, which is fairly new and so far not widely used, involves a laboratory experiment in which subjects allocate a pot of money between two time periods, with an interest rate applied to any money allocated to the later time period. Log-linearizing the first-order condition gives the following relationship between allocation in the two periods, the interest rate, and the temporal distance between the two periods:

$$\ln\left(\frac{c_t}{c_{t+k}}\right) = \left(\frac{\ln\left(\frac{1}{1+\rho}\right)}{\varepsilon}\right) \cdot k + \left(\frac{1}{\varepsilon}\right) \cdot \ln(1+r),$$

where ρ is the individual rate of pure time preference, k is the distance between the two periods, r is the interest rate, and ε is the elasticity of marginal utility of income. Estimating this using censored regression allows for computing ε as the inverse of the coefficient on $\ln(1+r)$.¹⁰ The methodology presents a trade-off for researchers. In order to achieve incentive compatibility, researchers must use real stakes, which requires that the stakes be relatively low (for budget reasons) and that time horizons be short (for practical reasons). In order to elicit preferences using larger stakes and longer time horizons, which would be more ideal for BCA purposes, it is necessary to use hypothetical stakes, which are considered less reliable.

3.8. Discussion

Some might question the validity of some or all of these methodologies, relying as they do on formal models that are not universally accepted, and, in most cases, on strong assumptions. We acknowledge that each of the methodologies is subject to critique, and that in some cases there may be disagreement as to whether the formal models used reflect actual human decision making. Our approach to these critiques is twofold. First, we see ourselves as embarked on an enterprise that is fundamentally and necessarily rooted in existing formal models. Ours is primarily a contribution to economic analysis of public policy, which is founded on formal

¹⁰The rate of pure time preference can be computed as a function of the coefficient on the interest rate and the coefficient on the interval. In practice, most applications of this method have assumed quasi-hyperbolic discounting, which requires the addition, on the right-hand side, of a dummy variable for the early period being the immediate present. The short-term discount factor in the quasi-hyperbolic model, β , is computed as a function of the coefficient on the interest rate and the coefficient on this dummy variable. Most studies have failed to reject the exponential model in which $\beta = 1$.

economic theory as a matter of principle. Concave utility functions are embedded in almost all areas of formal economic theory. Given that we are operating within these constraints, it is natural to take the formal models that underlie the various estimation strategies at face value.

It is true that, having adopted certain structural models, it is necessary to make certain assumptions in order to justify the empirical methods (which we have attempted to enumerate), and it is true that the estimates of ε generated by the various methodologies are sensitive to violation of these assumptions. Accordingly, the second component of our twofold approach to addressing critiques of the methodologies is to estimate a meta-regression to explore whether estimates of ε are sensitive to the methodology used. We report the results of this meta-regression in Section 6, where we find that there is no statistically significant effect of the methodology used.

Finally, we take note of a theoretical finding that the elasticity of the value of a statistical life (VSL) with respect to income is an upper bound on the elasticity of marginal utility of income (Hammitt, 2017). Meta-analysis of the elasticity of VSL finds that the mean is approximately .6 (Viscusi & Masters, 2017), far below the large majority of estimates of ε , which might call the theoretical finding into question. Furthermore, the model from which this finding is derived rests upon assumptions about how individuals value reductions in fatality risk which we consider to be more problematic than the assumptions made in the above methodologies. In particular, the model reduces the very complex psychology of mortality to a simplistic expected utility model over wealth, ignoring, among other things, the many profound emotions evoked by the prospect of fatal injury. In addition, as the model places only a bound on the elasticity and does not allow for empirical estimation of ε itself, we have concluded that this approach is not relevant to our analysis.

4. Available estimates of the elasticity of the marginal utility of income

Because each of the approaches discussed above has its own weaknesses and requires fairly strong assumptions, the best approach is not apparent to us. Nevertheless, as discussed in the following section, it is possible that available estimates tend to converge on what hopefully is the correct value. We explore this possibility through a meta-analysis. To conduct such a meta-analysis, it is necessary to first do a literature search and collect as many of the available estimates of ε as possible. We included all the estimates we found that used the methods described in Section 3.

In conducting our literature search, we focused on the US and the UK, two countries with similar cultures and economies, where one could anticipate that the elasticity values might be similar. As will be seen, pooling the available estimates of ε for these two countries provides us with a considerable number of sample points. Although estimates of ε do exist for other countries, especially for those in continental Europe and for Japan, the vast majority of work is for the US and UK. Moreover, as shown by Havranek *et al.* (2013), inter-country variation in the estimates is considerable, for a variety of reasons. As a result, our findings may not be valid for any countries other than the US and the UK. Given the preponderance of results from these two countries, we think it is important to avoid the possibility that inclusion of estimates from other countries might give the impression of a false level of generalizability.

Many of the studies we examined provide multiple estimates of ε , and the database we constructed for the meta-analysis contains these multiple estimates. For reasons discussed later, however, we cluster all the estimates for a single study in our meta-analysis. A few

Table 1. Studies and estimates using each method of estimating ε .

	Number of studies	Number of estimates
Indirect behavioral evidence: lifetime consumption models	139	1,691
Indirect behavioral evidence: consumer demand analysis	13	13
Indirect behavioral evidence: relative risk aversion models	3	4
Relating subjective well-being to income	2	2
Progressivity of income taxes	6	6
Asking experts	1	1
Convex time budget experiments	4	9

studies, however, used more than one of the approaches described in Section 3, and a few provided estimates for both the US and UK. In these instances, we treat each approach used and each covered country as a separate study.

In all, we found 168 studies that provided one or more estimates of ε . The number of studies and the number of estimates based on each approach is reported in Table 1.

As is evident, studies based on the lifetime consumption model provide the vast majority of estimates of ε , with studies based on consumer demand analysis occupying a distant second place. All but one of the studies based on the lifetime consumption model in our database were taken from a multiple nation meta-analysis of EIS conducted by Havranek *et al.* (2013), who did an extremely thorough literature search of EIS estimates. In fact, the one estimate we found that was not included in Havranek *et al.* (2013) was from a study conducted after their search was completed. Their meta-analysis, which covered all nations for which EIS estimates existed, differs from ours by focusing on why EIS estimates vary across countries and by excluding estimates that use approaches other than the lifetime consumption model, as that was beyond the scope of their study. As will be seen in the following section, their statistical approach also differs greatly from ours, and has come to be considered inferior (Hedges *et al.*, 2010).

5. Methods used in the meta-analysis of available estimates of ε

Using several meta-analytic tools (see Hedges, 1994 and Shadish & Haddock, 1994), we pooled the available estimates of ε described in Section 4 and estimated their mean value. We also estimated means for several subsets of estimates. In calculating mean values, best practice in meta-analysis involves taking account of the fact that some estimates are statistically more precise than others—as implied by their smaller variances—and calculating a weighted mean, the weight being the inverse of the variance of each of the individual estimate of ε . By using this weighting scheme, estimates of ε that are of greater statistical significance contribute more to the pooled mean than estimates that are less so.

In our meta-analysis, a complication in computing the means of the available estimates of ε occurs because many of the studies we use provide multiple estimates. Because they are estimated by the same researchers or team of researchers and typically with the same data sets, within study estimates are likely to be correlated, and thus, not independent. In addition, if each estimate is given equal weight in a meta-analysis, studies with a greater number of

estimates will be more influential in determining mean values than those with fewer estimates. Using a procedure developed by Hedges *et al.* (2010), we were able to determine that controlling for correlation among estimates within studies affected our estimates of the mean elasticity at only four significant digits. Under this circumstance, Hedges *et al.* (2010) provide a simple methodology for computing the mean and standard error of the elasticity. Thus, unlike Havranek *et al.* (2013), who treated each available estimate in their database equally, our results adjust for the number of estimates in each study and the correlations among those estimates, as well as for the statistical precision of each estimate.

The Hedges *et al.* (2010) procedure determines the weight given to each estimate within a given study, w_j , as follows:

$$w_j = \frac{1}{k_j \sum_{i=1}^{k_j} v_{ij}}, \tag{5}$$

where $j = 1 \dots m$ indexes studies, k_j is the number of estimates in study j , $i = 1 \dots k_j$ indexes the estimates within study j , and v_{ij} is the variance of estimate i in study j . Thus, there is only a single weight per study, which is applied to each estimate within the study; and estimates in studies with more estimates receive a smaller weight than those in studies with fewer estimates, given similar variances.

Given the value of w_j for estimates in each study, the weighted mean of the estimates across studies, b , is computed as follows:

$$b = \frac{\sum_{j=1}^m \sum_{i=1}^{k_j} w_j \varepsilon_{ij}}{\sum_{j=1}^m k_j w_j}. \tag{6}$$

Thus, although the sum of the elasticities tends to be greater in studies with a greater number of estimates, their weight tends to be smaller. The mean, b , can alternatively be estimated via a meta-regression in which b is the intercept term and values of w_j are used as weights. The advantage of a meta-regression is that each factor that may influence estimates of ε (e.g., whether they are based on US or UK data and the estimation method that was used) can be investigated, while holding other factors constant. The results from such a meta-regression are presented later.

The variance of the weighted mean, v^R is

$$v^R = \frac{\sum_{j=1}^m w_j^2 (\bar{\varepsilon}_j - b)^2}{\left(\sum_{j=1}^m w_j\right)^2}, \tag{7}$$

where $\bar{\varepsilon}_j$ is the unweighted average of the estimates of epsilon within study j . The square root of v^R is the robust standard error.

In computing the mean values, we used 158 of the 168 studies that provided estimates of ε . Four of the six progressivity of income taxes and two of the 13 consumer demand studies did not provide standard errors. Hence, w_j could not be computed for them, and they had to

be excluded. In addition, in convex time budget experiments, each respondent is presented with multiple tradeoffs, each of which is treated as a separate observation. For example, in the most extreme such study (Coble & Lusk, 2010), there were 47 study participants but 4,418 total choices, with each choice treated as an independent sample point. This procedure greatly inflates the values of w_j for the time budget experiments. In fact, our two largest values of w_j came from convex time budget studies, the largest of which was over four times the largest weight among studies using any of the other methodologies. This rendered the convex time budget studies unsuitable for inclusion in the meta-analysis. The implications of leaving out these 10 studies is investigated below.

Consideration was also given to omitting estimates of ε with implausible values. For example, negative values are clearly implausible, and perhaps very small positive values are as well. And large positive values, say above 4 or 5, may also be implausible. For example, as suggested by Equation (4), if g is near 2, as it appears to be, estimates this large imply an exceedingly high social discount rate. While very small and very large estimates do exist in our database, as shown below, excluding values such as those below, say, 0.5 or above, say, 3 would inevitably be arbitrary and seemed likely to introduce bias. Instead, we calculate means with the studies that reported the lowest and highest 10% (and, alternatively, the lowest and highest 20%) of all the estimates excluded, as a robustness check. However, we put the greatest emphasis on the mean value that includes all the estimates we obtained, with the exception of the 10 studies discussed in the previous paragraph. In our judgement, this is the most reliable estimate of ε to use in computing utility weights.

6. Findings from the meta-analysis

6.1. Elasticity of marginal utility of income

Table 2 shows the overall weighted mean estimate for all 158 included studies and for several subgroups of studies. The overall weighted mean of ε is 1.61, which is statistically significant at the 1% level and has a confidence interval of 1.18 to 2.05.

Table 2. Weighted means of estimates of ε .

	Mean	s.e.	No. of estimates	No. of studies
All studies	1.612	0.221	1,711	158
Excludes lowest and highest 10%	1.612	0.221	1,367	131
Excludes lowest and highest 20%	1.676	0.269	1,026	108
US studies	1.718	0.269	1,432	110
UK studies	1.340	0.103	267	48
Lifetime consumption studies	1.660	0.275	1,691	141
Consumer demand studies	1.209	0.301	13	11
All other studies	1.457	0.043	7	7
Micro-data studies	1.779	0.127	482	44
Aggregate data studies	1.547	0.256	1,229	115
20th-century publication date	1.892	0.194	1,449	124
21st-century publication date	1.112	0.192	262	35

Had it been possible to include the 10 omitted studies in computing the overall mean it is highly unlikely that the 1.61 estimate would have changed by very much. These studies account for only 6% of all the available studies. Moreover, although one of the elasticity estimates is 5.5 and two are just above zero, the remaining seven elasticity estimates are between 1.2 and 2.0, fairly close to the 1.61 estimate. To obtain a rough idea of whether the omitted studies would be likely to affect this mean if we knew the standard errors needed for their inclusion, we arbitrarily assume that the standard error of each estimate of ε is equal to 1. Using these standard errors and Equation (6), the estimated mean that incorporates the omitted ten studies is virtually identical to the estimate for the 158 included studies. If the standard error of each estimate is assumed to be 0.5 (or 1.5) instead of 1, the overall mean is again hardly changed from the value for the sample of 158 studies.

Table 2 permits several comparisons of mean estimates. (We test the statistical significance of the results using meta-regression below.) One important comparison is between the mean for all the available studies and the mean estimate that excludes the highest 10% of all the estimates and the lowest 10% of all the estimates. There is no difference in the mean elasticity at three significant digits. Trimming the highest and lowest 20% of estimates increases the mean by a small amount, from 1.612 to 1.676.

Perhaps the most important comparison of weighted means in Table 1 is between studies based on the lifetime consumption approach, the relatively few studies that use the consumer demand approach, and the even smaller number of studies that use other approaches. The mean for consumer demand studies is quite a bit smaller than the mean for the lifetime consumption studies, 1.2 versus 1.7. Because there were only seven studies that use other methods, they were combined. Their mean is closer to that for the lifetime consumption studies. These differences suggest that there could be value in additional studies that use methodologies other than lifetime consumption.

The differences in means between US and UK studies (1.7 versus 1.3) and between older and newer studies (1.9 versus 1.1) appear noteworthy. The difference between studies that rely on micro-data (usually surveys) and those based on aggregate data is relatively small, 1.8 versus 1.5. Arguably, the larger mean that is based on micro-data may be the more reliable one. For example, estimates that rely on aggregate data can be biased because they cannot control for demographic factors (Attanasio & Weber, 1993).

In addition to these pairwise comparisons, we conducted a meta-regression of elasticity estimates on each of the variables in Table 2, clustered at the study level to account for correlations among estimates from the same study. Despite the fact that some of the pairwise comparisons appear noteworthy, when all variables were controlled for simultaneously, none of the differences proved statistically significant. We conclude that while the results in Table 2 might suggest further study, the meta-regression results provide some confidence in using a mean value of 1.61 based on the full sample of estimates.¹¹

¹¹ We note that the elasticity of marginal utility of income for the U.S. implied by inverting the estimate of the elasticity of intertemporal substitution found by Havranek *et al.* (2013) is 1.68, which is quite close to our estimate of the elasticity of marginal utility of income for the U.S. and U.K. combined. We believe this similarity is largely coincidental. The 1.68 number is the inverse of the mean of EIS estimates. In our meta-analysis we have computed the mean of the inverses of EIS estimates (along with the other estimates of epsilon that we include). In general, we would not expect the mean of inverses of a variable to be equal to the inverse of the mean of that variable. Essentially, inverting their estimate of the EIS does not provide an estimate of the elasticity of marginal utility of income. Thus, it seems likely that the close similarity between the 1.68 number from Havranek *et al.* and our estimate of epsilon is largely coincidental.

A major concern in conducting meta-analysis is publication bias (see Begg, 1994), which results if the estimates from the studies included in the meta-analysis are unrepresentative of all those that exist. One potential cause of publication bias is that published studies, those that are most likely to be in the meta-database, are more likely to have statistically significant findings, suggesting a negative observed relationship between the estimates and their standard errors for studies in the database because studies with wide confidence intervals would be missing. This does not appear to occur in our case. Of the estimates of ε in our database that rely on the lifetime consumption model, over half (54.7%) are statistically insignificant at the 5% level. Of the estimates that rely on other methods, almost a quarter (22.7%) are statistically insignificant.

Another potential source of publication bias is that studies with implausible findings are more likely to be rejected for publication. This seems unlikely for studies using the lifetime consumption approach. For example, 14.2% of the estimates in our database that use this method are negative, while 26.0% exceed 5 in value. In contrast, none of the estimates of ε that used other approaches exceed these thresholds, suggesting that publication bias could exist for them. One way to examine the possibility that estimates with implausible values were unpublished is to determine the number of missing studies that would be required to appreciably change our mean estimate of ε . To use Equation (6) in these calculations, we assume that the missing studies would have the same average values for w_j and $k_j w_j$ as the studies in our sample. We computed the number of studies with estimates equal to 4.0 (−1.0) necessary to increase(decrease) our estimate of the mean elasticity to the upper(lower) bound of our confidence interval and in both cases found that there would need to have been approximately 100 missing studies, which we consider to be implausible, particularly because the missing studies would be more likely to contain a mix of exceptionally high and low estimates of ε rather than all high or all low. We conclude that a strong bias due to missing studies with implausible estimates is improbable.

6.2. Implied utility weights and discount rate

Table 3 shows the utility-weights implied by the mean and upper and lower bound values that we suggest for ε . In addition to the values for ε , the weights are based on Equation (3) and data on income for US households obtained from the US Census Bureau's Annual ASEC Survey. As defined by the Census Bureau, income is estimated prior to taxes and excludes non-monetary receipts (e.g., SNAP payments and goods that are bartered). In addition, the income

Table 3. Utility weights at selected income percentiles.

Percentile	Household income, 2021	Elasticity estimate (ε)		
		1.6	1.2	2.0
25th	\$44,250	2.50	2.00	3.14
50th	\$78,400	1	1	1
75th	\$128,000	0.46	0.56	0.38
95th	\$272,000	0.14	0.22	0.08

Note: The following formula is used to compute the weights: $(y^m/y^j)^\varepsilon$ where y^j is the income of the j^{th} percentile group, y^m is the median income of the group at the 50th percentile, and ε is the elasticity of the marginal utility of income. Source for 2021 household income at each percentile: US Census Bureau's Annual ASEC Survey in September 2021 (Ruggles *et al.*, 2022).

of people in group quarters (prisons and elder-care facilities, for example) is not included. The welfare weights in Table 3 pertain to households at the 25th, 50th, 75th, and 95th household income levels in 2021. Because income is higher in larger households than in smaller households, and the average size of US households is 2.5, the income amounts used to compute the utility weights are a weighted average of the incomes of two-person and three-person households. The utility weights must equal 1 at the 50th percentile because at that percentile $y_m = y_i$ and hence $\frac{y_m}{y_i} = 1$.

Note that with an elasticity of 1.6, the weight on the 25th percentile household is approximately five times larger than that on the 75th percentile household. To assess whether these weights are plausible it is necessary to consider the lived experience of households at these levels of income. Household income is \$44,250 at the 25th percentile and \$128,000 at the 75th percentile, an approximately threefold difference. Based on estimates from the Center for Women's Welfare at the University of Washington of the income necessary for "self-sufficiency" in urban California, these incomes are approximately 50% and 150% of the necessary income, respectively. If their income were spread evenly across the six categories of expenditures used by the Center, the money available for food each month in the average household at the 25th percentile would be \$387. The cost to the average household of the Thrifty Food Plan, which is the US Department of Agriculture's estimate of the cheapest possible strategy for achieving minimum nutritional needs, is \$583 per month. The 25th percentile household has 66% of the money they need to achieve adequate nutrition. An additional \$100 per month, if spent on food, would close that gap to 83%. The welfare impact of that increase in food consumption is probably outside the lived experience of most readers but is perhaps imaginable. It seems reasonable to think that it would be five times as great as the welfare impact of whatever a household with a combined income of \$128,000 would spend the same amount of money on.

That said, at the bottom and top of the income distribution the weights become very large and very small, respectively. For a household with income of \$500,000 the weight is .05 and for a family with income of \$5,000 it is 81.8, a difference of 1,535%. We could again try to imagine and compare the lived experience of these two households, but the exercise might become quite abstract. In ongoing work on the practical application of utility-weighting, we have concluded that it is appropriate, in order to avoid controversy, to establish upper and lower thresholds on the weights. The choice of the lower threshold has a considerable effect on utility-weighted net benefits because the weights increase hyperbolically as income goes down, so that sensitivity analysis on the lower threshold is important. In our applied work we have chosen a benchmark upper threshold weight of 5 and a lower threshold of .5.

With respect to the social discount rate, Equation (4) indicates that computing this parameter requires estimates of the rate of time preference (ρ) and the rate of growth of consumption (g), in addition to an estimate of ε . Boardman *et al.* (2018) suggest that ρ be set to 1.0 and g to 1.9. A recent survey of approximately 180 so-called experts by Drupp *et al.* (2018) obtained very similar values (at the mean, $\rho = 1.1$ and $g = 1.7$). Using the values suggested by Boardman *et al.* and our estimated 1.6 value for ε implies a social discount rate, d , of 4.0%. If instead, our lower bound for ε of 1.2 is used, $d = 3.28\%$; at the upper bound of 2.0 for ε , $d = 4.8\%$. These values are in the range that government bodies and various scholars have found for the social discount rate when it is based on the Ramsey formula. For example,

the UK Green Book (2018) suggests a rate of 3.5% and the European Union recommends 4% (European Commission, 2017).

7. Conclusion

The meta-analysis reported in this paper contributes to the literature in three ways. First, in order to most effectively advance a consensus on the correct utility weights and the Ramsey discount rate, it is important to have the most comprehensive and up-to-date estimates possible. Second, unlike previous studies, our meta-analysis accounts for the statistical precision of each of the included estimates and addresses the correlation of estimates within studies. Third, we have estimated a meta-regression and found that there are no systematic differences in estimates of ε that would complicate the application of our recommendations.

Using meta-analysis, we establish a value of the elasticity of marginal utility of income, ε , intended to contribute to a consensus on how to use weights to correct for bias in BCA caused by the diminishing marginal utility of income. Based on estimates from all the studies that could be included in the meta-analysis, the overall weighted mean of ε was found to be 1.61. There is some evidence that this estimate could be a bit low, although certainly still less than 2.0. For example, when only studies that were based on micro-data (arguably more reliable than those based on aggregate data) were included in the meta-analysis, the mean of ε increased to 1.779. However, the most recent studies—those conducted since the turn of the century—had a mean value of only 1.1. Relying on these findings, we tentatively conclude that it is reasonable to base the social discount rate and utility weights on an elasticity of 1.61, with lower- and upper-bound sensitivity testing at 1.2 and 2.0.

This paper is the second in a series of three papers that we hope will move the practice of BCA in the direction of utility-weighting to correct for the bias introduced into BCA by diminishing marginal utility of income. In the first paper we make the case that utility-weighting *should* be done. In brief our argument is that BCA should be an unbiased measure of aggregate individual welfare, both because we consider that to be the right information to provide to decision makers and because we believe that, in a time of increasing concern about distributional impacts, if BCA continues to be biased against the poor it will become increasingly irrelevant. The theoretical finding that efficiency and desirable distributional outcomes in a frictionless and undistorted system of markets can be achieved by making all laws and regulations on the basis of the potential Pareto criterion, with appropriate transfers being made through the tax system, is simply not relevant to the actual process of policy making.

In this paper and the third, which is ongoing, we attempt to establish that utility-weighting *can* be done. The two main objections are that there is no basis for a consensus on the correct weights to use and that the information necessary to apply utility-weighting is not available or is prohibitively difficult to apply. In this paper we believe we have successfully countered the first objection. There is a widely agreed upon model for computing weights and a narrowly bounded estimate of the relevant parameter. In the third paper we attempt to confront the second objection. We apply utility weighting to a series of existing, real-world BCAs and develop some tools for applying weights in an unbiased way and conducting sensitivity analysis. We conclude that, indeed, utility-weighting can be accomplished with available data and the kinds of informed assumptions that are routinely made in economic analysis of all kinds.

Supplementary material. The supplementary material for this article can be found at <http://doi.org/10.1017/bca.2023.29>.

References

- Acland, D. and D. H. Greenberg. 2003. "Principles and Practices for Distributional Weighting: A New Approach," *Journal of Benefit-Cost Analysis*.
- Adler, M. D. 2013. "Cost-Benefit Analysis and Distributional Weights: An Overview." Duke Environmental and Energy Economics Working Paper EE, 13-04.
- Arrow, K. 1963. *Social Choice and Individual Values*, 2nd ed. New Haven, CT: Yale University Press.
- Attanasio, O. P. and G. Weber. 1993. "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies*, 60: 631–649.
- Banzhaf, H. S. 2011. "Regulatory Impact Analysis of Environmental Justice Effects." *Journal Land Use & Environmental Law*, 27: 1.
- Begg, C. B. 1994. "Publication Bias." In Cooper, H. N. and L. V. Hedges (Eds.) *The Handbook of Research Synthesis*. New York: Russell Sage Foundation.
- Blackorby, C. and D. Donaldson. 1990. "A Review Article: The Case Against the Use of the Sum of Compensating Variation in Cost-Benefit Analysis," *Canadian Journal of Economics*, 23(3): 471–494.
- Blundell, R., M. Browning, and M. Meghir. 1994. "Consumer Demand and The Life-Cycle Allocation of Household Expenditures," *Review of Economic Studies*, 61: 57–80.
- Boadway, R. W. and N. Bruce. 1984. *Welfare Economics*. Oxford, UK: Basil Blackwell Ltd.
- Boardman, A. E., D. H. Greenberg, A. R. Vining, and D. L. Weimer. 2018. *Cost-Benefit Analysis: Concepts and Practice*, 5th ed. New York, NY: Cambridge University Press.
- Coble, K. H. and J. L. Lusk. 2010. "At the Nexus of Risk and Time Preferences: An Experimental Investigation," *Journal of Risk and Uncertainty*, 41(1): 67–79.
- Cowell, F. A. and K. Gardiner. 1999. *Welfare Weights*. London: STICERD London School of Economics. [http://darp.lse.ac.uk/papersDB/Cowell-Gardiner_\(OFT\).pdf](http://darp.lse.ac.uk/papersDB/Cowell-Gardiner_(OFT).pdf)
- Crossley, T. F. and H. W. Low. 2011. "Is The Elasticity of Intertemporal Substitution Constant?" *Journal of the European Economic Association*, 9(1): 87–105.
- Demuyck, T. 2018. "Testing the Homogeneous Marginal Utility of Income Assumption." *Econometric Reviews*, 37 (10): 1120–1136.
- Drèze, J. and N. Stern. 1987. "The Theory of Cost-Benefit Analysis." In Auerbach, A. J. and M. Feldstein (Eds.) *Handbook of Public Economics* (Vol. 2, pp. 909–989). Amsterdam: Elsevier.
- Drupp, M., M. C. Freeman, B. Groom, and F. Nesje. 2018. "Discounting Disentangled." *American Economic Journal: Economic Policy*, 10(4): 109–134.
- European Commission. 2014. *Guide to Benefit-Cost Analysis of Investment Projects*.
- European Commission. 2017. *Better Regulation Toolbox*.
- Evans, D. J. 2008. "The Marginal Social Valuation of Income for the UK." *Journal of Economic Studies*, 35(1): 26–43.
- Feldstein, M. 1964. "Opportunity Cost Calculations in Cost-Benefit Analysis." *Public Finance*, XIX: 117–139.
- Feldstein, M. 1972. "Distributional Equity and the Optimal Structure of Public Prices," *American Economic Review*, 62(1): 32–36.
- Fleurbay, M. and R. Abi-Rafeh. 2016. "The Use of Distributional Weights in Benefit–Cost Analysis: Insights from Welfare Economics". *Review of Environmental Economics and Policy*, 10: 264–285.
- Frisch, R. 1959. "A Complete System for Computing All Direct and Cross Elasticities in a Model with Many Sectors," *Econometrica*, 27: 177–196.
- Groom, B. and D. Maddison. 2019. "New Estimates of the Elasticity of Marginal Utility for the UK." *Environmental and Resource Economics*, 72: 1155–1182. <http://doi.org/10.1007/s10640-018-0242-z>
- Hammit, J. K. 2017. "Extrapolating the Value Per Statistical Life Between Populations: Theoretical Implications." *Journal of Benefit-Cost Analysis*, 8(2): 215–225.
- Harberger, A. C. 1978. "On the Use of Distributional Weights in Social Cost-Benefit Analysis," *Journal of Political Economy*, 86(2): S87–S120.
- Havranek, T., R. Horvath, Z. Irsova, and M. Rusnak. 2013. "Cross-Country Heterogeneity in Intertemporal Substitution." William Davidson Institute Working Paper Number 1056. Ann Arbor Michigan: University of Michigan.

- Hedges, L. V. 1994. "Fixed Effect Models." In Cooper, H. N. and L. V. Hedges (Eds.) *The Handbook of Research Synthesis*. New York: Russell Sage Foundation.
- Hedges, L. V., E. Tipton, and M. C. Johnson. 2010. "Robust Variance Estimation in Meta-Regression with Dependent Effect Size Estimates." *Research Synthesis Methods*, 1: 39–65.
- Kaplow, L. 2010. "Concavity of Utility, Concavity of Welfare, and Redistribution of Income." *International Tax and Public Finance*, 17, 25–42.
- Layard, R. and S. Glaister. 1994. *Cost-Benefit Analysis*. Cambridge: Cambridge University Press.
- Layard, R., G. Mayraz, and S. Nickell. 2008. "The Marginal Utility of Income." *Journal of Public Economics*, 92: 1846–1857.
- Lind, R. 1982. "A Primer on the Major Issues Relating to the Discount Rate for Evaluating National Energy Options." In Lind, R (Ed.) *Discounting for Time and Risk in Energy Policy*. Washington, DC: Routledge.
- Moore, M., A. E. Boardman, A. R. Vining, D. L. Weimer, and D. H. Greenberg. 2004. "Just Give Me a Number." *Journal of Policy Analysis and Management*, 23(4): 789–812.
- Nurmi, V. and H. Ahtiainen. 2018. "Distributional Weights in Environmental Valuation and Cost–Benefit Analysis: Theory and Practice," *Ecological Economics*, 150: 217–228.
- Ramsey, F. P. 1928. "A Mathematical Theory of Savings." *Economic Journal*, 38(152): 543–559.
- Robinson, L. A., J. K. Hammitt, and R. J. Zeckhauser. 2016. "Attention to Distribution in US Regulatory Analyses," *Review of Environmental Economics and Policy*, 10: 308–325.
- Ruggles, S., S. Flood, R. Goeken, M. Schouweiler, and M. Sobek. 2022. IPUMS USA: Version 12.0 [dataset]. Minneapolis, MN: IPUMS, 2022. <https://doi.org/10.18128/D010.V12.0>
- Selvanathan, S. 1988. "A System-Wide Analysis of International Consumption Patterns: Advanced Econometric Series." Ph.D. thesis, University of Western Australia.
- Shadish, W. R. and C. K. Haddock. 1994. "Combining Estimates of Effect Size." In Cooper, H. N. and L. V. Hedges (Eds.), *The Handbook of Research Synthesis*. New York: Russell Sage Foundation.
- Szpiro, G. 1986. "Measuring Relative Risk Aversion: An Alternative Approach." *Review of Economics and Statistics*, 68: 156–159.
- Treasury, H. M. 2018. *Green Book: Appraisal and Evaluation in Central Government*, London: HMSO.
- Viscusi, W. K. and Masters, C. L. 2017. "Income Elasticities and Global Values of a Statistical Life," *Journal of Benefit-Cost Analysis*, 8(2): 226–250.
- Yogo, M. 2004. "Estimating the Elasticity of Intertemporal Substitution When Instruments Are Weak," *Review of Economics and Statistics*, 86(3): 797–810.

Cite this article: Acland, D., and D. H. Greenberg. 2023. "The Elasticity of Marginal Utility of Income for Distributional Weighting and Social Discounting: A Meta-Analysis." *Journal of Benefit-Cost Analysis* 14: 386–405, doi:10.1017/bca.2023.29