NBER WORKING PAPER SERIES

VALUING EXCESS DEATHS CAUSED BY CLIMATE CHANGE

R. Daniel Bressler
Geoffrey Heal

Working Paper 30648
http://www.nber.org/papers/w30648

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2022

Heal acknowledges support from Columbia Business School. Bressler acknowledges support from the Columbia School of International and Public Affairs and the Open Philanthropy Project. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research nor any other entity.

Valuing Excess Deaths Caused by Climate Change
R. Daniel Bressler and Geoffrey Heal
NBER Working Paper No. 30648
November 2022
JEL No. H00,Q00

ABSTRACT

Valuing deaths caused by climate change in Benefit Cost Analysis (BCA) is complex and controversial, having caused disagreement and acrimony in past high-profile settings. Furthermore, it is of first order consequence to the value of the social cost of carbon (SCC). Despite this, the underlying considerations remain under-analyzed. We address this by assessing the theory behind different approaches to BCA, and by evaluating how they fare when applied to global externalities like climate change. The pure Kaldor-Hicks approach to BCA – measuring costs in market dollars unadjusted for diminishing marginal utility and valuing premature deaths in rich areas more than poor areas – relies on assumptions that are debated in domestic contexts, but, as we show, clearly do not hold in the context of climate change. We show that this approach is equivalent to defining a Negishi weighted social welfare function. Furthermore, we show that if costs are measured in purchasing power parity adjusted money – as is typical for the SCC – then the Kaldor-Hicks potential compensation criterion no longer necessarily holds. We conclude that the first-best BCA approach in the climate context is welfare weighting. This approach accounts for diminishing marginal utility using empirical estimates for the curvature of the utility function, and it better captures what a social planner naturally cares about: real net benefits and the welfare people get from those net benefits. The current U.S. practice – identical to the pure Kaldor-Hicks approach except that it gives a uniform population average value to all premature deaths – is preferred over the pure Kaldor-Hicks approach because it implicitly welfare weights premature mortality costs. However, the fully welfare weighted approach is first-best because it accounts for diminishing marginal utility across all costs, not just premature mortality risk.

R. Daniel Bressler
Columbia University
420 W 118th St
#1410
New York, NY 10027
rdb2148@columbia.edu

Geoffrey Heal
Graduate School of Business
582 Kravis Hall
Columbia University
665 W 130th St
New York, NY 10027
and NBER
gmh1@gsb.columbia.edu

1. Introduction

A large and growing literature projects that climate change is likely to cause significant premature death around the world (Bressler et al. 2021; Carleton et al. 2022; K. R. Cromar et al. 2022, 20; Gasparrini et al. 2017; Hales et al. 2014). Recent estimates suggest that temperature-related mortality impacts are the largest damage source in the social cost of carbon (SCC) (Bressler 2021; Rennert et al. 2022). This represents a significant revision from previous estimates, which projected that mortality impacts represented only a small fraction of damages in the SCC (K. Cromar et al. 2021).

This new literature revives old debates around how global premature deaths caused by climate change should be valued in benefit-cost analysis (BCA). These debates came to a head in 1995 when a chapter from the draft IPCC AR2 report valued premature deaths in developed countries ten times more than premature deaths in developing countries (Pearce 1995a). This led to significant backlash in the IPCC plenary session, where many countries expressed outrage. The environmental minister of India called this approach "absurd and discriminatory" and called for it to be "purged from the process." In the end, the proposal did not pass the IPCC review process, and the negotiated Summary for Policymakers implicitly rejected the chapter's approach (Bruce, Lee, and Haites 1996; Grubb 2005; Pearce 1995b).

The stakes of this debate are high. The value that is placed on premature mortality is of first-order consequence to the value of the SCC (Carleton et al. 2022). Furthermore, the consequences of this debate go well beyond the SCC. BCA as practiced by national governments has historically focused on measuring domestic benefits and costs.[1] For most countries, the SCC is likely to be among the first BCA calculations that explicitly projects and monetizes premature deaths in other countries caused by national policy decisions.[2] Thus, the decision for how to monetize premature deaths in the context of climate change may be precedent-setting for how other future global externalities are assessed in BCA.[3] Although the IPCC AR2 controversy occurred 27 years ago, the considerations involved with the question of how to monetize

[1] In the U.S., for instance, the vast majority of regulatory impact assessments only quantify domestic benefits and costs (Gayer and Viscusi 2016).

[2] In the U.S., there does not appear to be precedent for a Regulatory Impact Statement (RIA) explicitly projecting and monetizing premature deaths caused in other countries by U.S. policy decisions. A few RIAs have considered impacts in other countries but have not monetized them. For instance, the Mercury and Air Toxic Standard (MATS) RIA discussed the foreign health benefits of U.S. mercury reductions qualitatively, but did not quantify them (Howard and Schwartz 2017; US EPA 2011). Likewise, the 1996 NASA *Final tier 2 environmental impact statement for International Space Station* report considered mortality impacts from falling debris in the U.S. and other countries, but did not monetize the impact (NASA 1996).

[3] For instance, future regulatory impact assessments may determine the costs and benefits of safety regulations for BSL-4 laboratories, which would affect the very small probability of a lab-leak of a pandemic pathogen. The value that the U.S. places on premature mortality in other countries significantly affect the costs and benefits of this sort of regulation as well.

premature mortality from climate change remains under-discussed and under-analyzed in the literature. This paper seeks to remedy this.

We proceed as follows. In section 2, we briefly discuss the current practice and theoretical foundations of BCA, where our goal is not to be comprehensive, but to sufficiently discuss the theories underlying current approaches to set up section 3. Section 3 discusses BCA issues that are raised specifically in the context of valuing GHG emissions, and how alternative approaches to BCA fare when applied to global externalities like climate change. Section 4 draws conclusions. Section 5 provides a technical appendix with results that we refer to throughout the paper.

To summarize our conclusions: first, as our analysis makes clear, the issue of monetizing deaths caused by climate change is both complex and controversial. Therefore, it is important to follow the National Academy's suggestions for transparency by showing impacts in their natural physical units.[4] For mortality impacts, this is the number of excess deaths caused by GHG emissions. After the number of excess deaths are shown, the assumptions about how these deaths are monetized should be clear and transparent.

Second, we argue that analysts should respect individual choices and preferences, and thus minimize their own value judgements in BCA. However, the approach that best does so varies depending on the underlying science and economics of the policy being considered. In section 2, we show that there are compelling rationales for monetizing excess deaths at a higher value for richer individuals than poorer individuals when either (1) the population receiving the benefit is the same population that bears the cost, or (2) the population that bears the cost is compensated through another means, such as an optimal tax redistribution system. As we discuss, there is a debate in the scholarship as to whether these rationales are applicable in most domestic contexts. Despite this, countries that use BCA, including the U.S. and the U.K., have historically monetized and continue to monetize all premature deaths at a single population average value.[5] [6] In section 3, we show that these rationales are clearly inapplicable when applied

[4] See Committee on Assessing Approaches to Updating the Social Cost of Carbon et al. (2017), e.g.: "In order to provide a satisfactory degree of transparency, it is desirable for the damages module to report impacts in physical units when possible, such as crop yield changes, mortality, or species effects. These natural-unit measures are more straightforward to compare to the impact literature and require fewer intermediary assumptions to estimate than their monetized counterparts."

[5] For the U.S., see, e.g., Sunstein (2004): "For over two decades, executive orders have required regulatory agencies to engage in cost-benefit analysis of major regulations, and Congress has imposed similar requirements in several statutes. To conduct cost-benefit analysis, agencies must assign monetary values to human lives that are potentially saved by a proposed regulation. How do they come up with the numbers that they use? Do some deaths count for more than others?... No agency values the lives of poor people less than the lives of rich people. No agency distinguishes between whites and African Americans or between men and women. For statistical lives, the governing idea is that each life is worth exactly the same. With respect to cost-benefit analysis, much is disputed. But on the idea of a uniform value per life saved, there is a solid consensus, at least in terms of regulatory practice." This practice remains the same today (Sunstein 2022).

[6] For the U.K., see Her Majesty's Treasury (2022): "On grounds of equity … the valuation of a statistically prevented fatality (VPF) are based on average values from representative samples of the population (who differ in their incomes,

to climate damages caused by greenhouse gas emissions. When those in, e.g., the U.S. burn fossil fuels, they get all the energy benefits, but pay only a small portion of the cost from the subsequent climate damages. Furthermore, there is no supra-national tax authority. This implies that, e.g., the people of Niger are not compensated for excess deaths incurred from U.S. emissions decades in the future (Kolstad et al. 2014). Thus, an approach that monetizes excess deaths at a higher value for richer individuals than poorer individuals seems to lack foundation in the climate context, as it amounts to using an implausible thought experiment to justify benefiting some people (usually those who are better-off) at the expense of others (usually those who are worse-off).

Given that these rationales are inapplicable in the climate context, a final way to justify valuing excess deaths at a higher value for the rich compared to the poor is simply to adopt the Kaldor-Hicks potential compensation criterion – which specifies that a policy is net beneficial if those who benefit from a policy could fully compensate those who are harmed and still remain gainers, regardless of if such compensation ends up occurring (also known as a "Potential Pareto Improvement") – as an analytic first principle. However, countries such as the U.S. and the U.K. have long decided not to do this; as we show in section 2, the Kaldor-Hicks potential compensation criterion does not necessarily hold under the current practice because all premature deaths are valued at a single uniform average value.[7] Thus, the comparatively less compelling context of climate change is a poor first place to adopt the Kaldor-Hicks potential compensation criterion as a first principle.[8] What remains of the Pure Kaldor-Hicks approach in such circumstances is an implicitly defined Negishi-Weighted social welfare function (as shown in section 5), which applies Negishi (i.e. inverse marginal utility) weights that undoes diminishing marginal utility so that, e.g., the social welfare function counts a marginal dollar to a billionaire the same as a marginal dollar to someone below the poverty line. Furthermore, we show in section 3.4 that the Kaldor-Hicks potential compensation criterion only holds if BCA is conducted using units of money that represent an actual unit of exchange. If costs and benefits are measured in purchasing power parity (PPP) adjusted money – as is typical when calculating the SCC – then the Kaldor-Hicks potential compensation criterion no longer necessarily holds either.

---

preferences, age, states of health and other circumstances). These values are used when analysing and planning the provision of assets, goods and services at a population or sub-population level."

[7] This inconsistency is not a mere technicality; it is of central importance because mortality net benefits play a pivotal role in the BCA of major regulations. See, e.g., Hemel (2022): "Lifesaving regulations are not an administrative-state sideshow—they are the main act. Really expensive regulations generally do one of three things. They (a) reduce the risk of death or serious illness from air pollution, (b) reduce the risk of death or serious injury from motor-vehicle crashes, or (c) reduce greenhouse gas emissions. Note that a primary—probably the primary—reason why we worry about greenhouse gas emissions is that global warming will lead to death and serious illness on a vast scale, so (c) is largely subsumed by (a)."

[8] Note also that Circular A-4, the U.S. guidance on regulatory benefit-cost analysis, rejects Kaldor-Hicks as a first principle in BCA, stating that "when numeric adjustments are made for life expectancy or quality of life, analysts should prefer use of population averages rather than information derived from subgroups dominated by a particular demographic or income group" (OMB 2003).

Finally, we conclude that the first-best BCA approach for assessing GHG emissions is to apply welfare weights that account for the diminishing marginal utility of income using empirical estimates for the curvature of the utility function.[9] This converts benefits and costs into empirically informed welfare-adjusted dollars instead of market dollars that are implicitly Negishi weighted.[10] If one compares social welfare functions, we argue that a social welfare function that accounts for diminishing marginal utility does a better job of representing observed preferences than the Negishi weighted social welfare function, which does not. Furthermore, the welfare weighting approach better captures what a typical social planner would naturally care about: the ability of money to buy real net benefits and the welfare people get from those benefits, as opposed to measuring net benefits in units designed to satisfy an implausible thought experiment. We conclude that the current U.S. practice is preferred over the pure Kaldor-Hicks approach because it implicitly welfare weights premature mortality costs, as we show in equations (8-10). However, the fully welfare weighted approach is first-best compared to the current U.S. practice because it treats the relative value of premature mortality risk reductions and other net benefits more consistently by accounting for diminishing marginal utility across all net benefits, not just premature mortality risk.

2. The Current Practice of Benefit-Cost Analysis

Since at least the 1980s, executive orders have required agencies to conduct BCA of major regulations, and Congress has imposed similar requirements in several statutes (McGartland 2021; Sunstein 2004). BCA involves determining and adding up all the benefits and costs involved with regulations to calculate an estimate of regulatory "net benefits," to ensure that the benefits of regulations exceed the costs. As discussed in Hemel (2022), there are two major dimensions across which the current practice of BCA varies, which are closely related to one another:[11]

1. Whether the costs and benefits are presented in dollars that are weighted to account for the diminishing marginal utility of income (i.e., that an extra dollar is worth less to a

[9] The U.K. (Her Majesty's Treasury 2022) suggests distributional weighting using a utility curvature parameter estimated from Layard, Mayraz, and Nickell (2008). The German government estimates the SCC (Matthey and Bünger 2019) using the equity weighting approach described in D. Anthoff (2007).

[10] I.e., weighted by the inverse of the elasticity of marginal utility, sometimes referred to as "unweighted" – see section 5 for details and further discussion.

[11] See Hemel (2022) for a fuller discussion. There are of course other dimensions over which BCA varies, such as whether excess mortality is represented in terms of excess deaths (as is done in the U.S.) or as reductions in years of life (as is more often done in the U.K.), although the concepts discussed in this paper also apply in the latter case as well.

billionaire than to someone below the poverty line), or whether benefits and costs are presented in market[12] (implicitly Negishi-weighted) dollars.
2. Whether premature deaths are valued at the same population-wide value or at a value that varies with individual differences in the estimated willingness to pay to avoid mortality risk.

There are three major approaches to BCA that vary in the way they address these questions:

1. Measure all benefits and costs in market dollars, thus not accounting for diminishing marginal utility. Premature deaths are valued based on the estimated individual willingness to pay to avoid premature mortality. This approach has been called "textbook BCA" (Hemel 2022). As we show in equation (3), BCA under this approach is equivalent to achieving positive net benefits in a Negishi-Weighted Social Welfare function. For conceptual clarity, we will refer to this approach as *"Pure Kaldor-Hicks BCA."*

2. As with pure Kaldor Hicks BCA, measure benefits and costs in market dollars, with one exception: value all premature deaths at the same population-wide value. This is the way that BCA is currently practiced in the U.S. (although as noted previously, this is also the case elsewhere, such as the U.K.) so we refer to this as *"Current U.S. Practice BCA."*

3. All benefits and costs are measured using individual willingness to pay, but then they are welfare weighted to account for diminishing marginal utility.[13] Under standard parameter values, the elasticity of marginal utility, η, is equal to the elasticity of the VSL, ϵ.[14] This implies that although individuals may value their statistical lives differently in market dollar terms, they value statistical lives equally in welfare terms (see equations 8-10). We call this approach *"Welfare Weighted BCA."*[15]

[12] Note that when BCA is applied in the international context, whether the numéraire is represented in money at current exchange rates or in PPP-adjusted money becomes an additional important consideration, as discussed in part 3. In the domestic context, money can also be adjusted for purchasing power in different regions and metropolitan areas, and estimates for this in the U.S. are made publicly available by the U.S. Bureau of Economic Analysis (Aten 2017; Aten, Figueroa, and Martin 2012). Despite this, BCA in the domestic context is typically done using money at market exchange rates without adjusting for purchasing power.

[13] The government of the United Kingdom has endorsed the use of welfare weighting for analysis of highly impactful regulations with substantial distributional effects (Her Majesty's Treasury 2022), and the German Government uses equity weighting for its SCC estimates (Matthey and Bünger 2019).

[14] Hemel (2022) refers to BCA that uses inverse income-weights and an income elasticity of VSL equal to 1 as "standard hard-weighted cost-benefit analysis." This weighting approach is consistent with the German government's approach to weighting for the SCC (Matthey and Bünger 2019), and a VSL income-elasticity of 1 has been presented as a standard estimate for policy analysis (Masterman and Viscusi 2018; Robinson, Hammitt, and O'Keeffe 2019; Viscusi and Masterman 2017). Eeckhoudt and Hammitt (2001) and Kaplow (2005) both developed models in which agents have iso-elastic utility functions, and the income elasticity parameter (coefficient of relative risk aversion) nearly approximates the elasticity of VSL. Although there remains uncertainty in the empirical literature around the values of both η and ϵ, estimates for the two values contain significant overlap -- see, e.g., Acland and Greenberg (2022**)**.

[15] Weights can also be constructed not to match empirically estimated diminishing marginal utility of income, but to reflect normative considerations such as societal distributional goals, which some call "equity weights" (Acland and

This section precedes as follows. We begin in 2.1 by clarifying an important often-conflated distinction between Value of a Statistical Life (VSL) versus the value applied to premature deaths in BCA. In 2.2, we discuss the Pure Kaldor-Hicks approach to BCA. In 2.3, we discuss the current U.S. practice, and we show that the Kaldor-Hicks potential compensation criterion does not necessarily hold under this practice. In 2.4, we show that there are strong arguments for valuing the deaths of the rich more than the deaths of the poor in some domestic contexts, but this depends on the underlying science and economics of the policy being considered. In 2.5, we discuss welfare weighted BCA.

2.1 VSL Versus the Value Applied to Premature Deaths in BCA

Before discussing the main content of this section, it is important to emphasize an important distinction that often gets conflated: VSL (Value of a Statistical Life) is related to, but not necessarily synonymous with, the value that is applied to premature deaths in BCA. VSL is an evidence-driven estimate of how much an average individual would pay, in monetary terms, to avoid a specified mortality risk. Methodologically, it is often estimated in labor markets where there are tradeoffs between small amounts of additional near-term safety and wage compensation. Thus, VSL is not intended to represent the government's own assessment of "the value of life," but rather it is intended to capture the tradeoffs individuals make.

All approaches to BCA value premature deaths in a way that is informed by VSL estimates, but only the pure Kaldor-Hicks approach actually uses individual estimates of VSL to value premature deaths.[16] The current U.S. practice of monetizing premature deaths is informed by VSL in that it considers evidence from the empirical literature to estimate that, more or less, an average income person in the U.S. is willing to pay $1,100 to reduce their mortality risk by 1/10,000.[17] But it does not attempt to monetize premature mortality based on individuals' true VSLs, which vary with a number of individual factors such as income. Similarly, weighted BCA monetizes premature mortality by transforming market dollar-denominated VSL estimates into units that value all premature deaths equally in welfare terms (e.g., in the U.S. approach, welfare-

Greenberg 2022). As we believe that it is best to respect individual choices and preferences, and to thus minimize value judgements in BCA, we focus here on empirically informed welfare weights as opposed to positively informed equity weights. We provide a more thorough discussion of this later in this section.

[16] At least approximately, as all BCA is subject to informational constraints.

[17] Sunstein (2004) also points this out: "Agencies are not really able to identify a 'value of a statistical life'; instead they take advantage of information about how much people are willing to pay for facing statistical risks." And "The major point is that regulators do not really use a VSL; instead they use a mean WTP to eliminate a statistical risk. For example, agencies might say that they are using a VSL of $6 million, but when they do so, they are relying on evidence more or less establishing that the average person is paid $600 to face a risk of 1/10,000." Note that Sunstein is referring to VSL estimates from 2004, which have since been updated to account for income per capita growth and new empirical evidence (Timothy 2021; US Department of Health and Human Services 2021).

equivalent dollars for an average income person). As discussed above, under standard parameter values, this also assigns the same value to all premature deaths (after welfare weighting).

Thus, it is important for conceptual clarity to take care to make a distinction between "VSL" and "the value applied to premature deaths in BCA." In this paper, we only use the word "VSL" to refer to the empirically estimated willingness to pay of individuals to avoid mortality risk. Depending on the BCA approach taken, VSL may be merely informative of (and not synonymous with) the value attached to premature deaths in BCA.

2.2 Pure Kaldor-Hicks BCA

Pure Kaldor-Hicks BCA involves estimating all costs and benefits in market (implicitly Negishi weighted) dollars. This is equivalent to aggregating individual willingness to pay for net benefits measured using a market dollar numéraire (see section 5). This is straightforward for costs and benefits provided in markets because market prices are readily available. For non-market costs and benefits – the largest of which, in practice, is usually premature mortality risk (Hemel 2022)– estimates of individual willingness to pay must be used instead. In the case of mortality risk, the willingness to pay for mortality risk reductions is estimated as the value of statistical life (VSL), discussed above.

VSL estimates in the literature vary across many dimensions, especially cause of death[18] and income. Since those with more ability to pay typically have a higher willingness to pay to avoid mortality risk, VSL estimates for richer individuals tend to be higher than VSL estimates for poorer individuals.[19] Pure Kaldor-Hicks BCA attempts to use the evidence from the VSL literature to make best-possible estimates for the VSL of the populations where the premature mortality from the policy are projected to occur, and to use those estimates to assign a monetary value to premature deaths. For instance, if premature deaths are projected to occur in an affluent suburb where average incomes are ten times higher than in a poor inner-city neighborhood, then the Pure Kaldor-Hicks approach would value one premature death in the rich suburb the same as ten premature deaths in the inner-city neighborhood.[20] Thus, premature deaths in Pure Kaldor-

[18] For instance, willingness to pay to avoid death from cancer has typically been found to be higher than deaths from other causes (McGartland 2021; Sunstein 2004).

[19] An income elasticity of 1, for instance, implies that VSL varies proportionally with income, such that a richer individual with an income 10 times higher than a poorer individual will have a VSL that is 10 times higher than the poorer individual.

[20] Assuming an income elasticity of 1. The pure Kaldor-Hicks approach – in theory – seeks to assign values to mortality risk at an even more granular level than the neighborhood, down to the levels of the individual. However, like all BCA, it is constrained by the information available, so in practice individual-level mortality risks are hard to estimate. Also see Sunstein (2004): "It would follow that within the United States, wealthy populations would show a higher VSL than poorer populations. If a program is designed to combat health risks in wealthy suburbs, the VSL should be above the population-wide median; if the protected population is mostly in poor areas, the VSL should be below this median. Currently agencies pay no attention to this possibility in undertaking cost-benefit analysis."

Hicks BCA are not valued equally in market dollar terms or in welfare-terms. They are instead valued using a best-possible estimate of individual willingness to pay.

When all market and non-market (including mortality) costs and benefits are measured in market (Negishi weighted) dollars, the benefits and costs can then be aggregated and assessed using the Kaldor-Hicks potential compensation criterion. A regulation passes the Kaldor-Hicks test, and yields a Kaldor-Hicks improvement, if those who benefit could fully compensate those who are harmed and still remain gainers, regardless of whether such compensation actually occurs. As we show in equation (3), this is equivalent to achieving positive net benefits in a Negishi-Weighted Social Welfare function. Kaldor-Hicks improvements are often referred to as "potential Pareto improvements" because a Pareto improvement would occur if the winners of the policy were to fully compensate the losers of the policy. Pareto improvements – where at least some individuals are better off, and no individuals are worse off than before – are normatively attractive, and underly the theoretical support for efficient markets in the first and second fundamental welfare theorems of economics.

But the Kaldor-Hicks potential compensation criterion raises the question: why does the presence of a mere *potential* Pareto improvement matter if the compensation required to make an *actual* Pareto improvement does not in fact occur?

Proponents of the potential compensation criterion address these issues in a few different ways. One line of argument is that regulations should focus on maximizing total resources, which can then be redistributed through tax and transfer programs, which are better tools for redistributing income than alternative means.[21] However, this argument is disputed; others counter that, in fact, potential compensation is not synonymous with wealth maximization because it accounts neither for price changes caused by compensation payments[22] nor the costs of redistribution.[23] [24] [25] But critics and proponents of the argument alike have noted that the

[21] See, e.g., Kaplow and Shavell (2009).

[22] Boadway and Bruce (1984) show that because relative prices can change when post-project compensation is actually paid, a project can pass the formal Kaldor-Hicks benefit-cost test even when it does not in fact make Pareto improvements possible. This is because the prices people will face after compensation is paid may well differ from the prices they face before it is paid. If so, then if the project's "winners" are actually forced to pay the compensation, these winners may prefer the pre-project status quo to the post-compensation outcome.

[23] Hicks foresaw this problem in his original 1939 paper (Hicks 1939): "Since almost every conceivable kind of compensation (re-arrangement of taxation, for example) must itself be expected to have some influence on production, the task of the welfare economist is not completed until he has envisaged the total effects of both sides of the proposed reform; he should not give his blessing to the reform until he has considered these total effects and judged them to be good. If, as will often happen, the best methods of compensation feasible involve some loss in productive efficiency, this loss will have to be taken into account;"

[24] R. Boadway and Keen (1993) observe that when a policy alters the value of the public good/private good bundle associated with a given amount of income, individuals may alter their labor-leisure choices. When some individuals switch from labor to leisure in the presence of a distortionary tax, total wealth goes down. But Kaldor-Hicks BCA does not count that decline in total wealth as a cost.

[25] Hemel (2022) states: "Textbook CBA implicitly recognizes that we can almost always increase total wealth by redistributing less. Since textbook CBA doesn't account for changes in redistribution, it also doesn't count changes in

assumption that tax and transfer programs are always the best tools for redistribution will not always hold (Kaplow and Shavell 1994; Liscow 2013; Markovits 2004; Sanchirico 2000b, 2000a) and critics have questioned the relevance of such arguments if redistribution is politically and practically possible through other tools (such as regulation) but not tax and transfer programs (Liscow 2018, 2021; Revesz 2018).

Another line of argument in favor of pure Kaldor-Hicks is that—over time—net "winners" and "losers" from regulatory changes will alternate, such that no groups are made systematically worse off from a broad portfolio of regulations.[26] Critics of this argument argue that the potential compensation approach biases BCA in favor of the rich, and thus decisions made with this approach are systematically biased in favor of the rich as well.[27]

Because of these concerns, Arrow et al. (1996) argue that the "modern version" of the compensation principle should ask "whether compensation is *likely* to occur, rather than whether it could *possibly* occur".

2.3 The Current U.S. Practice: The Kaldor-Hicks Potential Compensation Criterion Does Not Necessarily Hold

The current U.S. practice of BCA is identical to the pure Kaldor-Hicks approach described above with one major exception: instead of monetizing excess deaths using individual or group-level estimates of VSL, it attaches a single population-wide value to all premature deaths regardless of income or other factors. Although a poorer individual will tend to have a lower willingness-to-pay to avoid mortality risks than a richer individual, U.S. regulatory analysis does not make income adjustments to reflect this. Federal agencies have long used a single population average VSL to monetize excess deaths for all Americans, regardless of their income, race, gender, or other demographic factors. For instance, the U.S. Department of Transportation (Timothy 2021) and Homeland Security (Houser and Sunstein 2021) use a 2021 nationwide average VSL of $11.8 million and the Department of Health and Human Services uses a 2021 central VSL estimate of $11.6 million (US Department of Health and Human Services 2021). According to the 2000 EPA Science Advisory Board report that addressed this topic, monetizing excess deaths based on an income varying VSL "raises difficult moral, ethical, and political issues about which the Committee is not and cannot be in full agreement.." and thus "…the Committee believes that EPA should not abandon its approach — described in its Guidelines for Preparing Economic Analyses — of using a uniform VSL across populations that vary in the above socioeconomic attributes" (US EPA 2000).

---

total wealth resulting purely from changes in redistribution. In this sense, textbook CBA is symmetrical: it applies the same treatment to the benefits of redistribution and the costs of redistribution (which is to say, it ignores both)."

[26] See Polinsky (1972) and Hicks (1941).

[27] See Liscow (2018): "Thus, rather than allocating resources to the poor, who are most in need," Kaldor-Hicks efficiency will "tend to do the opposite: allocating resources to the rich, who are willing to pay the most."

How does this affect the Kaldor-Hicks potential compensation criterion? It implies that a rule that passes the Kaldor-Hicks potential compensation criterion will not necessarily be net-beneficial under current U.S. practice, and vice versa. For instance, consider an agency deciding whether to allow a developer to build a factory whose primary cost is causing excess deaths from pollution in a downwind poor community. Although the U.S. average VSL is roughly $11 million, assume that this poor community's actual average VSL is only $5 million. Current U.S. practice would value the excess deaths in this community at the nationwide average $11 million value, whereas pure Kaldor-Hicks BCA would value those deaths at $5 million. If the BCA of allowing the factory to be built narrowly passes a Kaldor-Hicks potential compensation test valuing each life at $5 million, then it would not pass a benefit-cost test under current U.S. practice that values each life at $11 million. Although the winners of the policy (e.g., the factory's shareholders and consumers of the factory's products) could potentially compensate the losers (the poor community) for their incurred damages, this policy would have net costs under current U.S. practice.

Now consider the same example but assume instead that the factory primarily causes excess deaths in a downwind rich community, who's average VSL is in fact $30 million. If the BCA for allowing the factory to be built narrowly passes the benefit-cost test under current U.S. practice, in which each excess death is valued at $11 million, then it would fail the Kaldor-Hicks potential compensation test where each excess death is valued at $30 million. In this case, although the winners of the policy cannot potentially compensate the losers – because the losers have a high VSL and demand large compensation for mortality damages – allowing the factory to be built would still pass the benefit-cost test under current U.S. practice.[28]

Of course, there are many examples where policies would pass both the Kaldor-Hicks potential compensation test and the benefit-cost test under current U.S. practice. And "slam-dunk" policies – where benefits exceed costs by a wide margin – are likely to pass both tests. However, the key point is that while the current U.S. practice is informed by the Kaldor-Hicks potential compensation criterion and the principle of potential Pareto improvements, it is not synonymous with it because of the way it values premature deaths at a single national average value.

2.4 There are Strong Arguments for Valuing Deaths Higher for The Rich Than the Poor in Some Domestic Settings Depending on the Underlying Science and Economics

Although valuing all premature deaths at a single population average VSL has been the standard practice in the U.S. for decades, Cass Sunstein makes compelling arguments for why

[28] This example illustrates a point raised by Sunstein, which is that a uniform population average VSL of the sort that the government now uses threatens to "overprotect" the poor and "underprotect" the wealthy (Sunstein 2004).

this may not be a good idea (Sunstein 2004). He makes these arguments on two related grounds: welfare and autonomy.

Sunstein argues that: "In principle, government should not force people to buy protection against statistical risks at a price that seems excessive to them. At least as a general rule, people should not be required to pay $70 to reduce a risk of 1/100,000 if they are willing to pay no more than $50." As an example of Sunstein's point, take potential regulations involving automobile safety. The government might justify regulations that increase automobile safety by weighing the benefits and costs using a nationwide average VSL of $11 million. For the average American, such a regulation may be welfare-improving, but for poorer Americans who have lower willingness to pay to avoid mortality risks, such a regulation may force them to buy vehicles with more expensive safety features than they would be willing to purchase themselves. For these poorer Americans, the costs of additional safety features outweigh the value that they place on the reduction in mortality risk. Such a regulation amounts to a "forced exchange," where some individuals are forced by the regulation to make purchases that they would not make themselves, which lowers their welfare.[29]

Such "forced exchanges" also have important implications for autonomy. When regulators force their own views on individuals instead of making rules in a way that respects individuals' own free choices, Sunstein argues that this is "insulting their dignity,"[30] and Al McGartland argues that this is a violation of the "consumer sovereignty principle."[31] As Sunstein argues, the welfare and autonomy arguments tend to point in the same direction.[32]

---

[29](Sunstein 2004): "Why should government force people to pay for things that they do not want? Begin with welfare. By hypothesis, a forced exchange on terms that people dislike will make them worse off."

[30] (Sunstein 2004): "WTP might be defended instead on the ground of personal autonomy. On this view, people should be sovereign over their own lives, and government should respect personal choices about how to use limited resources (again so long as those choices are informed). When people decline to devote more than $60 to the elimination of a 1/100,000 risk, it is because they would prefer to spend the money in a way that seems to them more desirable. If regulators do not use people's actual judgments, then they are insulting their dignity."

[31] (McGartland 2021): "For there to be an increase in social welfare, Kaldor–Hicks requires a 'potential' Pareto improvement, which occurs when those who gain from the economic change would be willing to compensate the losers and still be better off. The rules of BCA, including the use of the consumer sovereignty principle, follow from this criterion. This means that BCA must value benefits and costs based on consumers' values, not what a policy maker in Washington thinks the benefits and costs are worth."

[32] See Sunstein (2004): "The remedy for unjust distributions, and for that form of coercion, is not to require people to buy regulatory benefits on terms that they find unacceptable. Suppose that people are willing to pay only $60 to eliminate a 1/100,000 risk because they are not rich, and that if they had double their current wealth, they would be willing to pay $120. Government does people no favors by forcing them to pay the amount that they would pay if they had more money. I have suggested that for those who do not believe that regulatory decisions should be based on welfare, considerations of autonomy point in the same direction. Those who refuse to pay a certain amount to eliminate a risk of 1/100,000 might want to use their resources for other things—medical care, children, food, recreation, entertainment, savings. If people are entitled to a kind of sovereignty over the conduct of their own lives, then they should be permitted to make such allocations as they choose. It is most standard to justify use of WTP on

However, as Sunstein argues, both the welfare and autonomy arguments break down when those that are benefitting from the regulation are not the ones paying the costs of the regulation.[33] Take the example of the polluting factory above whose primary cost is causing deaths in the downwind poor community. Assume that the downwind community does not get any of the benefit of the factory's production. Monetizing excess deaths in the poor community at the nationwide average VSL of $11 million causes the project to fail the BCA test. This increases the welfare of the community because they do not bear the costs of the pollution and they are not getting any benefits from the factory. Furthermore, assigning the nationwide average VSL to the poor community instead of their true VSL of $5 million is not a violation of their autonomy or their consumer sovereignty, as it saves them from paying a cost on terms that they would find acceptable.[34] The community left up to its own devices would certainly prefer for the plant not to be built, as they only bear its costs without getting benefits.

Now consider the same case except that both the rich community (with an average VSL of $30 million) and the poor community (average VSL of $5 million) are downwind from the factory. Both communities bear costs but receive no benefits. Under the Pure Kaldor-Hicks approach, excess deaths caused by the same pollution would be valued 6 times more in the rich community than in the poor community, whereas under current U.S. practice, deaths in both communities would be valued the same. Considering that the welfare and autonomy arguments no longer hold in this case, this unequal result may be difficult to justify.

However, even in this case, one can still justify using individual estimates of VSLs to monetize premature deaths if one believes that the tax system optimally redistributes income, or if one assumes that the government's various regulations end up balancing out in the end so that no group is systematically left worse off (Sunstein 2004). In any case, this requires one to believe that both of these arguments are empirically true, which as we mentioned above, is fiercely contested. Weighing all of these factors, Sunstein concludes that monetizing premature mortality risks in the same way in *all* cases, without regarding the scientific and economic specifics of the regulation, does not make sense.[35]

---

welfare grounds, but the same approach is at least equally defensible as a means of respecting the autonomy of persons."

[33] See (Sunstein 2004): "The central claim… is that the argument for using WTP is strongest when the beneficiaries of regulation must pay all of its cost… The argument for using WTP is weaker when the beneficiaries of regulation pay only a fraction of that cost."

[34] See (Sunstein 2004): "The difficulty is that a high VSL, one that exceeds what WTP studies show for poor people, may produce outcomes that are in the best interest of poor people, in the sense that the result is a welfare improvement for them. And if poor people do not bear all of the costs of programs that benefit them, the autonomy argument for WTP is greatly reduced; they are enjoying a benefit (partly) for free, and it does not insult anyone's autonomy to give them a good on terms that they find acceptable."

[35] See Sunstein (2004): "It is therefore reasonable to reject the confident view of economically inclined analysts who believe that accurate VSLs, based on actual WTP (and hence individuated), should always be the basis of regulatory

Daniel Hemel provides an additional argument for caution in using VSLs that vary with demographic characteristics: "expressive harms," i.e. "the harmful—though not easily quantifiable—effects of agency procedures and policies that predictably send a message to some that their lives are worth less" (Hemel 2022). Hemel argues that the mere fact that Regulatory Impact Assessments attach a higher value to some premature deaths than others may itself lead to an additional cost that should be considered, even if there are strong reasons to support such an approach, such as those raised by Sunstein.[36] Hemel also argues that monetizing the deaths of the rich more than the poor is likely to sow mistrust not just in the government writ large, but more specifically it may undermine the enterprise of BCA itself.[37]

In addition to expressive harms, it is important to note that in the cases where poor communities are bearing the costs but not benefits from a policy – such as the factory pollution example described above – a Pure Kaldor-Hicks approach to BCA *does* actually value premature deaths in their communities less than richer communities, and *does* cause them to be worse off than if the U.S. Current Practice of valuing all lives at the nationwide average were used. Thus, this case is not just a question of expressive harms to members of poorer communities, but of substantive harms incurred directly by the policy.[38] When poor communities bear both the costs and benefits from a policy, this can cut both ways, but this merely further supports Sunstein's arguments that a "one-size-fits-all" approach to monetizing premature mortality is not good policy, at least when expressive harms are not considered. And it provides a rationale for a BCA approach that is better suited to such circumstances.

---

policy. But it is similarly reasonable to reject the confident view of skeptics who believe that a uniform VSL, refusing to make distinctions among persons, is best on distributive grounds."

[36] (Hemel 2022): "First, the very fact that agency officials assign lower-dollar VSLs to lower-income individuals might itself give rise to expressive harms. Agencies publicly release regulatory impact analyses (the documents that detail their CBAs), and they typically summarize key elements of the CBA in the preambles to proposed and final rules published in the Federal Register. Although it is unlikely that many people will read regulatory impact analyses and Federal Register notices in their original form, news reporters likely will read these documents. Before deciding to use lower dollar VSLs for lower-income individuals, practitioners of CBA need to think about the consequences of news headlines declaring, for example, that the EPA and the DOT are discounting poorer people's lives. The concern is not purely about public relations. Virtually all will agree that it is a bad thing if millions of Americans think that the federal government values their interests less than the interests of other, richer Americans—and bad for reasons beyond the fact that agency officials may endure a few difficult news cycles. We derive utility from believing that federal officials are looking out for our interests and disutility from believing that they are not. Public confidence in government is a difficult-to-quantify value, but it is not a trivial value. One consideration in the decision to use (or not to use) income-elastic VSLs should be whether the practice will be interpreted as communicating a lack of concern for lower-income individuals' interests." Relatedly, McGartland (2021) explains that using a single uniform VSL "allows agencies to avoid difficult communications concerning this controversial issue."

[37] (Hemel 2022): "Assigning different-dollar VSLs to individuals of different income levels may affect perceptions not only of government writ large but of CBA specifically. One outcome of the senior-death-discount episode was to provide CBA critics with a predictably effective avenue of attack against the practice of assigning dollar figures to benefits and costs… the EPA's age adjusted VSL experience—as well as the anticipated backlash from using income-elastic VSLs in the future—should at least prompt second thoughts about whether the potential gains from using different-dollar VSLs justify the risk to the entire CBA enterprise."

[38] To the extent that the policy choice is guided by BCA.

2.5 Welfare Weighted BCA

A final approach to BCA involves welfare weighting, which explicitly accounts for the distribution of damages by using weights that reflect that individuals have different marginal utilities depending on their current circumstances (Mirrlees 1978). Practically, welfare-weighted BCA begins the same way as pure Kaldor-Hicks BCA: by estimating all costs and benefits using market dollar-denominated willingness to pay, including the estimated individual willingness to pay to avoid premature mortality. After this, weights are applied to account for the fact that a marginal dollar of income[39] is valued more by lower income people than higher income people: money in the hands of poorer individuals is renormalized to a common welfare unit (usually money in the hand of a person with median income, although any reference individual x can be used – see section 5 for details); the same renormalization is done to money in the hands of richer individuals. Welfare weights, which are typically chosen to reflect empirical estimates of the diminishing marginal utility of income across individuals, can be estimated from multiple lines of evidence.[40]

Weights can also be constructed to upweight marginal benefits to lower income individuals and to downweight marginal benefits to higher income individuals not to match empirically estimated diminishing marginal utility of income, but to reflect normative considerations such as societal distributional goals, which some call "equity weights" (Acland and Greenberg 2022). As we believe that it is best respect individual choices and preferences, and to minimize value judgements in BCA, we focus here on welfare weights.[41] The U.K. has endorsed the use of welfare weighting for analysis of highly impactful regulations with substantial distributional effects (Her Majesty's Treasury 2022), and Germany uses equity weighting for its SCC estimates (Matthey and Bünger 2019).

Under standard parameter values, the weights that account for diminishing marginal utility offset the differences in willingness to pay to avoid mortality so that premature deaths are valued at the same population-wide value after weighting.[42] In the standard case, this implies that although individuals may have quite divergent money-valued VSLs due to income differences, diminishing marginal utility perfectly explains this, which results in individuals across the income spectrum valuing their life equally in welfare terms.

[39] Or consumption.
[40] In the standard isoelastic utility setting, the η parameter, which represents the elasticity of the marginal utility of consumption, controls the curvature of the utility function. This η parameter can be estimated using multiple lines of evidence, including studies that consider risk aversion, "happiness" and subjective wellbeing studies, and studies of intertemporal substitution. See, e.g., Acland and Greenberg (2022) table 1 and Groom and Maddison Pr. (2019).
[41] The basis for welfare weights (the diminishing marginal utility of income) is empirically verifiable, whereas there is no purely empirical basis for the normative judgments that are reflected in equity weights.
[42] See footnote 14.

As with the other BCA approaches, benefits and costs are then aggregated to determine if the benefits of some policy exceed the costs. The process of weighting for diminishing marginal utility transforms the numéraire from units of market money into welfare-weighted units of money, and thus the aggregated net benefits can be said to represent a social welfare function. If the diminishing marginal utility parameter is estimated empirically (welfare weighting), then this social welfare function can be justified on descriptive grounds of producing results that are equivalent to the aggregation of changes in individual welfare (utility).

As discussed above, the pure Kaldor-Hicks approach is justified by respect for autonomy and consumer sovereignty (when the same individuals bear the benefits and costs of a policy), and on the grounds of the potential compensation criterion (more generally). However, if those benefiting from a policy are different than those paying the costs of the policy, and if the potential for Pareto efficient compensation does not result in such compensation actually occurring, the pure Kaldor-Hicks approach seems to lack foundation. In such a circumstance, it amounts to using an implausible thought experiment to justify benefiting some people (usually those who are better-off) at the expense of others (usually those who are worse-off). What remains of the pure Kaldor-Hicks approach in such circumstances is implicitly defining a Negishi-Weighted social welfare function (equation 3) that assumes that there is no diminishing marginal utility of income, which empirical evidence suggests is not the case. If one compares social welfare functions, a social welfare function that accounts for diminishing marginal utility does a better job of representing observed preferences than one that does not.

An advantage of welfare weighted BCA relative to the current U.S. practice is that it accounts for diminishing marginal utility across all costs and benefits, not just premature mortality. The current U.S. practice implicitly weights mortality risk, but all other benefits and costs remain Negishi weighted and measured in market dollars. This has the effect that for the poor, safety from premature mortality risk will be overvalued relative to other benefits whereas for the rich, safety will be undervalued relative to other benefits, as discussed previously. Welfare weighted BCA avoids this issue by applying the same weighting factor to all benefits and costs.

Furthermore**,** weighted BCA can also be justified by the Kaldor-Hicks potential compensation criterion, provided that we make the assumption that compensation is possible with welfare-weighted dollars, as we show in equation (7). This is because weighted BCA transforms costs and benefits from units of market money into welfare-weighted money by applying welfare weights. After this transformation, the benefits and costs can then be aggregated and assessed using the same Kaldor-Hicks potential compensation criterion described above. If the weighted benefits exceed the weighted costs, then the winners of the policy will be able to compensate the losers with their surplus welfare-weighted dollars to reach a potential Pareto improvement. Compared to the pure Kaldor-Hicks approach where potential

compensation is paid in units of market dollars, this has less real-world applicability, because market dollars are a real unit of exchange whereas welfare-adjusted dollars are not. However, a hypothetical about welfare-weighted dollar compensation that does not occur is arguably a more plausible basis for policy analysis and policymaking than a hypothetical about Negishi-weighted dollar compensation that also does not occur. Furthermore, as we discuss in more detail in section 3.4, the pure Kaldor-Hicks approach to BCA only works when it uses money that represents an actual unit of exchange. If money is transformed to account for purchasing power or other factors, then the weighted BCA approach has the same claim to following the Kaldor-Hicks potential compensation principle as welfare-weighting: that is, it only necessarily satisfies the Kaldor-Hicks potential compensation principle if we grant ourselves a hypothetical unit of exchange.

Because welfare weighted BCA involves first measuring all costs and benefits in units of market money (as in pure Kaldor-Hicks BCA) and then weighting those costs and benefits to account for differences in marginal utility, it has similar drawbacks to pure Kaldor-Hicks BCA in terms of its demandingness on granular and accurate information and projections, and its demandingness on analysts' time. As Hemel (2022) argues, an advantage of the current U.S practice over both Pure Kaldor-Hicks BCA and Welfare Weighted BCA is that it measures non-mortality net benefits in the most straightforward and least demanding way (measuring net benefits in market money without welfare weighting), and it also measures mortality net benefits in the most straightforward way and least demanding way (i.e. applying a single population average VSL to all projected premature deaths). For mortality impacts, both the pure Kaldor-Hicks approach and the welfare weighted approach require analysts to know the incomes (and ideally other factors over which VSL varies) of those who face altered premature mortality risk from a policy in addition to the change in on premature deaths. Whereas the current U.S. practice only requires analysts to project the change in premature deaths, which are then all monetized at the same uniform value.

3. How Climate Change Complicates BCA

In this section, we show that the rationales for taking a pure Kaldor-Hicks, and thus applying higher values to the deaths of the rich and the deaths of the poor, are much more challenged in the context of valuing climate damages than in the domestic contexts discussed in the previous section. In 3.1, we show that using an income-elastic value for premature deaths creates especially large discrepancies in the international context compared to the domestic context. Whereas the discrepancy between the value given to premature deaths using an income-elastic value between rich and poor U.S. states is on the order of 2:1, the discrepancy between rich and poor countries is on the order of 100:1. This places an especially large burden on the underlying rationale for taking such an approach to valuing premature deaths caused by GHG

emissions. In sections 3.2-3.4, we show that such rationales are in fact much weaker in the climate context than in domestic contexts. In section 3.2, we show that those that pay the climate cost of emissions are largely not those that get the benefit of burning fossil fuels. In section 3.3, we show that arguments around compensation to losers are much less credible in the global GHG emissions context, in particular because there is no supra-national tax authority. And finally in section 3.4, we show that if income projections account for purchasing power parity, as is typical in the underlying socioeconomic projections used to calculate the SCC, then, as with welfare-weighting, BCA no longer supports the Kaldor-Hicks potential compensation criterion unless if we grant ourselves a hypothetical unit of exchange. For these reasons, we conclude that the current U.S. approach and the welfare-weighting approach are superior to the pure Kaldor-Hicks approach when calculating the SCC, and we argue that the welfare-weighting approach is first-best.

3.1 A Globally Varying Income-Elastic Premature Death Value Creates Especially Large Discrepancies and Raises the Specter of Expressive Harms

Within the U.S., the largest income discrepancy between states is between Mississippi and New York, where New York has roughly twice the per capita income of Mississippi. Thus, under a pure Kaldor-Hicks approach to BCA that uses an income-varying value to monetize premature mortality, just over two deaths in Mississippi would be valued the same as one premature death in New York[43] (although as discussed above, the U.S. does not do this, but instead values all lives at the same nationwide average value).

Globally, however, the differences are much greater. For instance, Belgium has 89 times the GDP per capita of the Congo.[44] Thus, 89 premature Congolese deaths are valued the same as a single premature death in Belgium.

Furthermore, this approach leads to results that seem inconsistent with the public opinion and government policies in countries with official SCC estimates including the U.S. and Germany. To cite another jarring statistic, under this approach, a single Russian death is valued the same as 2.5 Ukrainian deaths.[45]

[43] Assuming a VSL income elasticity of 1, and 2021 GDP per capita values calculated from U.S. Bureau of Economic Analysis and U.S. Census Bureau data https://www.bea.gov/data/gdp/gdp-state https://www.census.gov/data/tables/time-series/demo/popest/2020s-state-total.html#par_textimage

[44] https://data.worldbank.org/indicator/NY.GDP.PCAP.CD; we compare the Democratic Republic of the Congo to Belgium.

[45] https://data.worldbank.org/indicator/NY.GDP.PCAP.CD; The point here is not to argue that premature deaths should be valued higher or lower in any given country (in fact it is to argue against this), but to illustrate that approaches that choose to monetize deaths across countries based on regionally varying income will inevitably result in some uncomfortable results that clearly clash with public opinion and current policy.

These unpalatable results and large discrepancies place an especially large burden on the underlying rationale for using income varying values to monetize premature deaths in this context, which we cover in more detail below. The 1995 IPCC case mentioned in the introduction of this paper suggests that such a practice is likely to cause backlash because of both its substantive (insofar as BCA guides decision making) and expressive harms. Especially because the poorest countries whose deaths are valued the least – such as most countries in sub-Saharan Africa and South Asia – represent only a small fraction of the historical responsibility for climate change (Ritchie 2019).

3.2 Those that Benefit from Emissions Reductions Are Not the Same as those Paying for Emissions Reductions

As discussed above, there is a strong rationale for using an income varying value to monetize premature mortality when comparing policies in which a single population bears both the benefits and the costs of a project because this approach respects the empirically estimated tradeoffs that this population makes between its own consumption benefits and mortality risks. This rationale can be applied to the international realm as well. For instance, consider an aid organization whose mission is to build roads in the Congo, but insists on using the German average VSL (87x the Congolese average VSL) to assess the benefits and costs of its road-building projects. This approach far overstates the costs associated with traffic fatalities relative to the consumption benefits that the Congo would receive from the new road. New roads would likely never be built under this approach that would have provided significant welfare benefits to the people in the Congo.

Now consider a different international aid organization with a broader mission: to improve health by making investments anywhere in the world. Would it make sense for this aid organization to use an income varying value? If it did, then it would be heavily privileging rich countries relative to poor countries, as Sunstein (2004) argues: "The fact that a poor person in a poor nation would be willing to pay $1 to eliminate a risk of 1/10,000, whereas a wealthy person in a wealthy nation would be willing to pay $100, cannot plausibly be used to defend the view that an international agency should devote its resources to the latter rather than the former." [46]

[46] See also Sunstein (2004): "Any judgment about the appropriate VSL, and about individuation, must be heavily pragmatic; it must rest on the consequences of one or another choice. Whether government should use a higher or lower VSL across demographic lines cannot be answered simply. An important implication involves the assessment of VSL across nations. A poor nation would do well to adopt a lower VSL than a wealthy nation; for China or India, it would be disastrous to use a VSL equivalent to that of the United States or Canada. But this point should not be taken to support the ludicrous proposition that donor institutions, both public and private, should value risk reduction in a wealthy nation above equivalent risk reduction in a poor nation."

The case of climate change is theoretically analogous to the second case as opposed to the first. In the first case, the aid organization only considers projects (roads) where all of the benefits and the costs accrue to the Congolese. But GHG emissions produced in the U.S. are not like this. When individuals in the U.S. produce GHG emissions, they get the benefits of the energy produced, but only pay a small portion of the climate cost of those emissions. The climate cost, including the mortality cost, is spread throughout the whole world, and it is projected to be most severe in poorer countries (Bressler et al. 2021; Carleton et al. 2022). In this situation, countries outside the U.S. are not facing a tradeoff between consumption benefits and climate costs; they are only paying costs. Like the international aid organization in the second case, the question facing the U.S. is how much to value the deaths its emissions are causing in different parts of the world. Such a decision cannot be justified on the grounds of respecting autonomy, consumer sovereignty, or the tradeoffs that individuals in poor countries make between consumption benefits and mortality risks, because they are facing no such tradeoff.

3.3 Arguments Around Compensation to Losers Fail in the Context of GHGs

As discussed above, the pure Kaldor-Hicks approach aggregates individuals' willingness-to-pay using market (Negishi weighted) dollars and applying an income-elastic value to monetize excess mortality. An important line of argument in favor of this approach is that although the winners may not immediately compensate the losers, potential compensation is likely to occur in other forms in the future through tax-system transfers and other future rules where the winners and losers may be swapped. As the IPCC argued, however, full compensation for the contribution of domestic greenhouse gas emissions towards increasing premature mortality risk in other countries up to centuries in the future is conceptually difficult to estimate, and unlikely to occur (Kolstad et al. 2014).[47] Thus, arguments around compensation, which are plausible arguments in favor of the pure Kaldor-Hicks approach in the domestic context, fall apart in the climate change context.

To further illustrate the salience of this point: in 1991, then World Bank Chief Economist Larry Summers caused controversy due to a leaked World Bank memo he signed that stated:

[47] See Kolstad et al. (2014): "It is sometimes assumed that CBA is conducted against the background of efficient markets and an optimal redistributive taxation system, so that the distribution of income can be taken as ideal from society's point of view. If that were true, it might reduce the need for distributional weights. But this is not an acceptable assumption for most projects aimed at climate change. Credit and risk-sharing markets are imperfect at the world level, global coordination is limited by agency problems, information is asymmetric, and no supra-national tax authority can reduce worldwide inequalities. Furthermore, intergenerational transfers are difficult. In any case, the power of taxation to redistribute income is limited because redistributive taxes create inefficiency (Mirrlees 1971). Even optimal taxation would therefore not remove the need for distributional weights. Thus, the assumption that incomes are (second-best) optimally redistributed does not neutralize the argument for welfare weights in aggregating costs and benefits."

"The measurements of the costs of health impairing pollution depends on the foregone earnings from increased morbidity and mortality. From this point of view, a given amount of health impairing pollution should be done in the country with the lowest cost, which will be the country with the lowest wages. I think the economic logic behind dumping a load of toxic waste in the lowest wage country is impeccable and we should face up to that" (Hausman and McPherson 2006). This memo caused significant controversy, and Summers later claimed that it was written sarcastically and that it was a mistake (Rosenberg 2001; United States Senate 1993).

However, when compensation is considered, the rationale for valuing excess deaths from climate change less in poor countries and more in rich countries is demonstrably weaker than the rationale for dumping toxic waste in poor countries. In the dumping example, when a rich country dumps its toxic waste in a poor country, that poor country will demand compensation for the cost of taking on the rich country's waste. The poor country would estimate the number of excess deaths that such toxic waste is expected to cause, use their own internal VSL to monetize those deaths along with other costs, and then demand compensation that at least exceeds those costs. If the rich country agrees to pay at least this amount, a dumping deal can be agreed to that leaves both countries better off, and a Pareto improvement is achieved. If the rich country is not willing to pay the compensation, the poor country can turn down the offer and ask the rich country to look elsewhere for toxic waste dumping grounds.

For GHG emissions, the possibility that countries will pay compensation for damages is not credible, as the IPCC argues. As a result, GHG emissions are the theoretical equivalent of dumping toxic waste in other countries without giving them any payment for it. If compensation were credible, then using an SCC that monetizes premature deaths in poor countries at a much lower value and premature deaths in rich countries at a much higher value would theoretically match the World Bank dumping example, and Pareto improvements would occur. Even though rich countries would cause damage in poor countries from "dumping" GHGs in the atmosphere, those poor countries would be end up getting compensated for the damages, and all parties could be better off. However, such compensation is not credible in the GHG case, and no Pareto improvement follows. Thus, the strength of the case for using different values for rich and poor (already hotly disputed in the domestic context) is at its nadir in the GHG context.

3.4 Purchasing-Power Parity Is Incompatible with a Pure Kaldor-Hicks Approach

Finally, there is a practical issue facing the implementation of a pure Kaldor-Hicks approach to estimating the SCC. The income projections most commonly used in the SCC literature – the SSPs (Riahi et al. 2017) and the recently produced RFF-SPs (Rennert et al. 2021)

– calculate country-level income projections using purchasing power parity (PPP) instead of market exchange rates. PPP transforms current prices into adjusted prices using weights that account for the ability of money in different places to purchase fixed bundles of goods and services. Like welfare-weighting, it tends to up-weight money in low-income areas and down-weight money in high-income areas.

The issue with PPP as it relates to the Kaldor-Hicks potential compensation principle is that it transforms net benefits from actual units of exchange (i.e. money at current market exchange rates) to another numéraire that – like welfare-weighted dollars – is hypothetical and is not an actual unit of exchange. PPP, like welfare-weighted dollars, adjusts the numéraire (money at current market exchange rates) to capture something that is economically important. In the case of PPP, this is the purchasing power of money. In the case of welfare weights, this is diminishing marginal utility. Both the PPP and welfare weighting conversions arguably convert the money numéraire into units that are more relevant to what a typical social planner would care about: the ability of money to buy goods and services and the welfare people get from those goods and services. But in both cases, this conversion undermines the appeal to the Kaldor-Hicks potential compensation criterion. Thus, an approach that estimates all mortality and non-mortality net benefits in PPP adjusted dollars has the same claim to following the Kaldor-Hicks potential compensation principle as welfare-weighting: that is, it necessarily satisfies the Kaldor-Hicks potential compensation principle only if we grant ourselves a hypothetical unit of exchange. If, as in the pure Kaldor-Hicks approach to BCA, we restrict compensation to units that actually exist as a unit of exchange, then a policy that yields positive net benefits measured in PPP-adjusted money will not necessarily pass the Kaldor-Hicks potential compensation test.

We can show this with an example: the current market exchange rate of a Congolese Franc to a Euro is 2,000:1.[48] That is, if one were to exchange Francs for Euros in financial markets, one could exchange 2,000 Francs for 1 Euro. However, when purchasing power parity is accounted for between the Congo and Germany, this rate becomes 1,110:1.[49] That is, 1,110 Francs could buy the same portion of a fixed basket of goods and services in the Congo as 1 Euro could buy in Germany.

Now assume that one conducts a BCA of some policy in units of PPP adjusted Euros, and this policy yields 1M Euros in benefits to Germany, but -1.5M Euros in costs in the Congo. Thus, this policy appears to yield -0.5M in net benefits. One may conclude that this policy does not pass the Kaldor-Hicks Potential Compensation test because there are net costs, and no

[48] As of October 25th, 2022; see https://g.co/finance/EUR-CDF

[49] The PPP rate between Francs and German Euros is calculated as: $\left(\frac{Congolese\ PPP\ Conversion\ Factor}{German\ PPP\ Conversion\ Factor} * \frac{Market\ Franc}{Market\ Euro}\right) = \left(\frac{0.5}{0.9}\right) * 2000 = 1{,}110$. See https://data.worldbank.org/indicator/PA.NUS.PPPC.RF for 2021 PPP conversion factors.

potential Pareto improvement can be reached. We provide a summary of these calculations below in table 1.

However, this is incorrect. This policy actually *does* pass the Kaldor-Hicks potential compensation test. The problem with the above analysis is that it does not consider that the exchange rate between PPP adjusted money – like the exchange rate between welfare adjusted dollars – is hypothetical, and not the exchange rate that is available in the market. Consider that the Germans can convert 0.9M Euros out of their 1M Euros in benefits into 1,800M Congolese Francs at the 2,000:1 Franc:Euro market exchange rate, and transfer those 1,800M Francs to the Congo. After this transaction, the Germans end up with 0.1M Euros in net benefits. Recall that the policy caused -1.5M PPP adjusted Euros in costs to the Congolese. After the transfer, the Congolese have an extra 1,800M Francs. Converting this into PPP adjusted Euros yields: $1{,}800\ Francs\left(\frac{1\ PPP\ adjusted\ Euro}{1{,}110\ Francs}\right) = 1.62M\ PPP\ adjusted\ Euros$. Thus, after the transfer, the Congolese have 0.12M PPP adjusted Euros in net benefits. Thus, as long as the transfer occurs, both the Congolese and the Germans enjoy net benefits (0.22M PPP adjusted Euros in total), and thus a Pareto improvement results.

	Germany	**Congo**	**Net Benefit (Germany + Congo)**
Benefits in PPP-Adjusted Euros	*1.00M Euros*	*-1.50M Euros*	*-0.50M Euros*
Transfer in Local Currency at Market Exchange Rates	-0.90M Euros	1,800M Francs	
Convert Transfers at Market Exchange Rates to PPP-Adjusted Euros	-0.90M Euros	1.62M Euros	
Benefits in PPP-Adjusted Euros Post-Transfer	*0.10M Euros*	*0.12M Euros*	*0.22M Euros*

Table 1: BCA using PPP-Adjusted Euros

A simpler and more straightforward way to see this is just to convert the numéraire for the BCA into units of money at market exchange rates from the beginning, which we show below in table 2. Since in this analysis we are conducting the PPP adjustment based on German Euros so that 1 PPP adjusted German Euro is equivalent to 1 market Euro, we only need to convert the Congolese costs from PPP adjusted Euros into market-exchange rate Euros, which we do with the following equation:
$-1.5M\ PPP\ Adjusted\ Euros\left(\frac{1{,}110\ Francs}{1\ PPP\ Adjusted\ Euro}\right)\left(\frac{1\ Euro}{2000\ Francs}\right) = -0.83M\ Euros.$

Thus, once we do this conversion, we see that the benefits to Germany (1 M Euros) exceed the costs to the Congo (0.83M Euros), yielding net benefits (0.17M Euros), thus implying that there is indeed a potential Pareto Improvement, as we saw above.

	Germany	Congo	Net Benefit (Germany + Congo)
Benefits in PPP-Adjusted Euros	*1.00M Euros*	*-1.50M Euros*	*-0.50M Euros*
Benefits in Euros at Current Market Exchange Rates	*1.00M Euros*	*-0.83M Euros*	*0.17M Euros*

Table 2: BCA using Market Exchange Rate Euros

As this example shows, if benefits and costs are measured in a numéraire that doesn't represent an actual available unit of exchange, such as PPP and welfare adjusted monetary units, then BCA conducted with such a numéraire cannot claim to be supported by the Kaldor-Hicks potential compensation criterion, unless if we grant ourselves a hypothetical unit of exchange. The only way to ensure that the Kaldor-Hicks potential compensation criterion is satisfied is to measure all costs and benefits in money at market exchange rates. Thus, in the context of the SCC, one would need to estimate the SCC using socioeconomic projections and damage estimates measured in money at market exchange rates as opposed to PPP adjusted money to claim to be following the Pure Kaldor-Hicks approach to BCA. The downside of doing so, as noted previously, is that it does not track what an analyst or policymaker would naturally care about: the change in real consumption that is experienced by those affected. And, as discussed before, this is not common, as the income projections most commonly used in calculating the SCC (SSPs and RFF-SPs) make projections using purchasing power parity.

4. Conclusions

As has been illustrated throughout this piece, the issue of monetizing deaths caused by climate change is both complex and controversial. Therefore, we conclude that it is important to follow the National Academy's suggestions for transparency by showing impacts in their natural physical units.[50] For mortality impacts, this is the number of excess deaths caused by greenhouse gas emissions. The incremental mortality impact of GHG emissions in physical units is the mortality cost of carbon, which represents the number of excess deaths caused by a marginal ton of carbon dioxide emissions (Bressler 2021). After the number of projected excess deaths are shown in physical units, the assumptions for how these deaths are monetized should be clear and transparent.

Second, we conclude that there are compelling rationales for monetizing excess deaths at a higher value for richer individuals than poorer individuals in some contexts: when either (1) the population receiving the benefit is the same population that pays the cost, or (2) compensation is

[50] See (Committee on Assessing Approaches to Updating the Social Cost of Carbon et al. 2017), e.g.: "In order to provide a satisfactory degree of transparency, it is desirable for the damages module to report impacts in physical units when possible, such as crop yield changes, mortality, or species effects. These natural-unit measures are more straightforward to compare to the impact literature and require fewer intermediary assumptions to estimate than their monetized counterparts."

likely to occur through some other form, such as an optimal tax redistribution system. As discussed above, many scholars have argued that these rationales do not apply even in most domestic contexts. Perhaps because of such criticism, the U.S. and the U.K. have historically monetized (and continue to monetize) all excess deaths at a single population average VSL when conducting BCA.

In the context of monetizing global deaths from climate change, these rationales clearly fail to apply. When individuals in the U.S. burn fossil fuels, they get all of the energy benefits, but pay only a small portion of the cost from climate damages. Furthermore, there is no supra-national tax authority to optimally redistribute income at the global level, so that, e.g., the people of Niger are compensated for excess deaths incurred from U.S. emissions decades in the future. As these rationales fail to apply, a final way to justify valuing excess deaths at a higher value for the rich compared to the poor is simply to take the Kaldor-Hicks criterion as an analytic first principle. However, given that the U.S. has long decided not to do this when conducting BCA in contexts when the rationale to do so is much more compelling, the far less compelling context of climate change is a poor first place to take the Kaldor-Hicks criterion as a first principle.

Finally, we conclude that approaches that value excess deaths equally either in dollar terms (matching the current BCA practice in the U.S.) or in welfare-adjusted money terms (matching the standard welfare weighting approach) are more appropriate when monetizing premature mortality from global externalities like climate change. Because the welfare weighting approach treats the relative value of mortality risk reductions and other benefits more consistently than the current U.S. approach, the current U.S. practice is second-best compared to the first-best fully welfare weighted approach.

5. Technical Appendix

The utility of person i as a function of their consumption c in time t is given by $u_i(c_{i,t})$.[51] Assuming isoelastic utility, this term becomes:

$$u_i(c_{i,t}) = \frac{{c_{i,t}}^{1-\eta}}{1-\eta}$$

Following, e.g. Anthoff, Hepburn, and Tol (2009) and Nordhaus (2017), define the Social Welfare Function (SWF) W as the aggregate of individual utility across the population of n people in time t across all the time periods considered in the analysis from time $t = 1$ to $t = T$, discounted by the discount factor r_t:

[51] Utility can also be represented as a function of income. We choose to represent utility as a function of consumption here, but the same analysis could be applied in the former case as well.

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} u_i(c_{i,t})\, r_t$$

The Social Cost of Carbon (SCC) is equivalent to the marginal damage caused by marginal carbon dioxide emissions in some period, for instance in the period $t = 1$:

$$SCC_1 = \frac{\partial W}{\partial E_1} = \frac{\partial \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} u_i(c_{i,t})\, r_t}{\partial E_1}$$

Assuming an exogenous discount factor:

$$SCC_1 = \frac{\partial W}{\partial E_1} = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\partial u_i(c_{i,t})}{\partial E_1} r_t$$

Applying the Chain rule:

$$SCC_1 = \frac{\partial W}{\partial E_1} = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\partial u_i(c_{i,t})}{\partial c_{i,t}} \frac{\partial c_{i,t}}{\partial E_1} r_t \qquad (1)$$

Equation (1) is in units of welfare. As discussed in Anthoff, Hepburn, and Tol 2009, this SCC value can be used directly in BCA, but this tends to be inconvenient because other benefits and costs are usually measured in consumption-equivalent dollars. For practical application in BCA, this SCC should be converted into dollars. There are two ways to do this:

(1) Negishi Weighting
(2) Welfare Weighting
(3) Equity Weighting[52]

Negishi weights (Negishi 1960) weight each person i's utility by the inverse of their marginal utility. I.e., the Negishi weight $\alpha_{i,t} = 1/\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}$. Applying this to the welfare-denominated SCC in equation (1):

$$SCC_1 = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \alpha_{i,t} \frac{\partial u_i(c_{i,t})}{\partial c_{i,t}} \frac{\partial c_{i,t}}{\partial E_1} r_t = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{1}{\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}} \frac{\partial u_i(c_{i,t})}{\partial c_{i,t}} \frac{\partial c_{i,t}}{\partial E_1} r_t = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\partial c_{i,t}}{\partial E_1} r_t \qquad (2)$$

Thus, Negishi weighting does two things here: (1) It converts the SCC into units of consumption-equivalent dollars, which makes the SCC more convenient to compare with other benefits and costs. (2) It defines a new SWF in consumption-equivalent units that undoes the

[52] As we discuss in the main text, we focus on welfare weights in this paper because we believe that it is best respect individual choices and preferences, and to minimize value judgements in BCA.

transformation of consumption units into welfare units. We can see this by applying the antiderivative to equation (2) to extract the Negishi-Weighted SWF W:[53]

$$\int \frac{\partial W}{\partial E_1} dE_1 = \int \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\partial c_{i,t}}{\partial E_1} r_t dE_1 = W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} c_{i,t}\, r_t \tag{3}$$

Thus, when the utility function is concave to account for diminishing marginal utility from consumption, the Negishi weight undoes this transformation so that utility is now linear instead of concave in consumption so that, e.g., a marginal dollar given to a rich individual with millions of dollars of consumption each year counts the same as a marginal dollar given to a poor individual with thousands of dollars of consumption each year.

The Negishi-weighted SWF is identical to the SWF that is implicitly used in the Kaldor-Hicks approach to BCA. We can see this because, as discussed above, a policy passes the Kaldor-Hicks potential compensation test if the winners gain enough from the policy to be able to potentially compensate the losers. This is equivalent to a policy change causing positive net benefits, i.e, leading to a positive change in the Negishi-Weighted SWF. That is, a policy passes the Kaldor-Hicks potential compensation test iff it leads to a positive change in the Negishi-Weighted SWF:

$$\Delta W > 0 \equiv \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \Delta c_{i,t}\, r_t > 0$$

We now consider (2) the welfare weighting approach. We start with equation (1), but instead of applying a Negishi weight, we apply a weight that retains the structure of the utility function while converting the SWF from units of welfare into units of dollars: $\mu_{x,1} = 1/\frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}$. This represents the inverse of the marginal utility that some reference person x gets from consumption in time $t = 1$, as in, e.g., Anthoff, Hepburn, and Tol (2009) and Errickson et al. (2021). Thus, this factor converts the SCC from units of welfare into units of money as it is valued on the margin by the reference person x in $t = 1$:

$$SCC_1 = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \mu_{x,0} \frac{\partial u_x(c_{x,1})}{\partial c_{x,1}} \frac{\partial c_{i,t}}{\partial E_1} r_t = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}}{\frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}} \frac{\partial c_{i,t}}{\partial E_1} r_t \tag{4}$$

[53] The antiderivative operator technically yields equation (3) plus some constant K: $\int \frac{\partial W}{\partial E_1} dE_1 = \int \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\partial c_{i,t}}{\partial E_1} r_t dE_1 = W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} c_{i,t}\, r_t + K$. However, in practice, the SWF in equation (3) will be used to compare the impact of multiple policies, and the K will just be subtracted out from both sides when policies are being evaluated, so for simplicity, we leave it out.

Now, the consumption-equivalent impact from a marginal emission to person i in time t , $\frac{\partial c_{i,t}}{\partial E_1}$, is multiplied by the ratio of the marginal utility of that person and the marginal utility of reference person x in time $t = 1$. Thus, equation (4) represents the welfare-weighted SCC in the general case.

As with Negishi weighting, welfare weighting converts the SCC into units of consumption-equivalent dollars, which makes the SCC more convenient to compare with other benefits and costs. In addition, the Negishi-weighted SCC from equation (2) is actually a special case of the welfare-weighted SCC from equation (4). The Negishi-weighted version assumes that every individual has the same marginal utility of consumption in every time period as the marginal utility of consumption as the reference person x, i.e. $\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}} = \frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}\ \forall i, t$, and thus $\frac{\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}}{\frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}} = 1\ \forall i, t$. In other words, the Negishi-weighted SCC treats a marginal dollar to every person equally regardless of whether that person is rich with high consumption or poor with low consumption. This may be because person x is the same as person $i\ \forall i$ so that each individual's marginal utility of consumption is offset, so that net benefits are measured according to each person's own current dollar valuation of marginal consumption (Abbott and Fenichel 2014). Or it may be because it is assumed that every person i's marginal utility of consumption is equivalent to some individual agent x's marginal utility of consumption.

If we assume that all people have isoelastic utility with a common η curvature parameter, we can simplify equation (4) further:

$$SCC_1 = \sum_{t=1}^{t=T}\sum_{i=1}^{i=n} \frac{\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}}{\frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}} \frac{\partial c_{i,t}}{\partial E_1} r_t = \sum_{t=1}^{t=T}\sum_{i=1}^{i=n} \frac{{c_{i,t}}^{-\eta}}{{c_{x,1}}^{-\eta}} \frac{\partial c_{i,t}}{\partial E_1} r_t = \sum_{t=1}^{t=T}\sum_{i=1}^{i=n} (\frac{c_{x,1}}{c_{i,t}})^{\eta}\ \frac{\partial c_{i,t}}{\partial E_1} r_t \qquad (5)$$

In this setting, we can see that the consumption-equivalent damage for each person in time t from a marginal emission in time $t = 1$, $\frac{\partial c_{i,t}}{\partial E_1}$, is multiplied by the welfare weight $(\frac{c_{x,1}}{c_{i,t}})^{\eta}$. Thus, in the isoelastic utility setting, Negishi weights implicitly assume that $\eta = 0$ so that $(\frac{c_{x,1}}{c_{i,t}})^{\eta} = 1\ \forall c$. This implies that $u_i(c_{i,t}) = \frac{{c_{i,t}}^{1-\eta}}{1-\eta} = c_{i,t}$, i.e., utility is linear in consumption and $\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}} = 1$.

As with the special case of Negishi weighting, the more general weighting approach defines a new SWF in consumption-equivalent units. We can see this by applying the antiderivative to equation (4) to extract the welfare-weighted SWF W:[54]

$$\int \frac{\partial W}{\partial E_1} dE_1 = \int \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}}{\frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}} \frac{\partial c_{i,t}}{\partial E_1} r_t dE_1 = W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \frac{\frac{\partial u_i(c_{i,t})}{\partial c_{i,t}}}{\frac{\partial u_x(c_{x,1})}{\partial c_{x,1}}} c_{i,t}\, r_t$$

When utility is isoelastic, the SWF W becomes:

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} (\frac{c_{x,1}}{c_{i,t}})^{\eta}\, c_{i,t}\, r_t \tag{6}$$

As with the Negishi-weighted approach, we can also apply the Kaldor-Hicks potential compensation criterion to this SWF, provided that we make the assumption compensation is possible in dollars that are weighted to the reference person x. Under this assumption, a policy passes the Kaldor-Hicks potential compensation test if it leads to a positive change in the welfare-Weighted SWF:

$$\Delta W > 0 \equiv \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} (\frac{c_{x,1}}{c_{i,t}})^{\eta}\, \Delta c_{i,t}\, r_t > 0 \tag{7}$$

Finally, we turn our attention to the current U.S. practice. We start with equation (6), which is the general SWF in consumption-equivalent units. However, we now define the new term $k_{i,t}$, which includes all net benefits in consumption-equivalents except for willingness to pay to avoid some mortality risk $p_{i,t}VSL_{i,t}$, where $p_{i,t}$ is the probability of mortality from mortality risk p for person i in time t and $VSL_{i,t}$ is person i's value of statistical life for mortality risk p in time t. As discussed above, the current U.S. practice is Negishi-weighted for all net benefits with exception of premature mortality risk, which is measured at a population average value. We define the term $VSL_{x,1}$ to be reference person x's VSL, which in the case of the U.S. is an average American. Thus, the current U.S. practice SWF is written formally as:

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} (k_{i,t} - p_{i,t} VSL_{x,1}) r_t \tag{8}$$

Now, we compare this to a fully welfare-weighted SWF:

[54] The antiderivative operator technically yields equation (3) plus some constant K. However, in practice, the SWF in equation (3) will be used to compare the impact of multiple policies, and the K will just be subtracted out from both sides when policies are being evaluated, so for simplicity, we leave it out.

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta} (k_{i,t} - p_{i,t} VSL_{i,t}) r_t \quad (9)$$

For the pure Kaldor-Hicks approach and for the welfare weighting approach, ideally $VSL_{i,t}$ estimates would be available for all individuals in society and for all risks. However, such estimates are typically not available. Analysts must make estimates for individual VSLs based on broader group-based estimates. Although, as discussed above, individual VSL estimates vary across many dimensions -- including the cause of death, income levels, individual risk preferences, and more -- it is common in the BCA literature to only vary estimates for income (or consumption) based on an estimated VSL elasticity parameter, ϵ. Because many populations do not have VSL estimates available, it is common to employ a benefits transfer methodology to estimate a particular population's VSL based on another population's VSL estimates, the ratio of the populations' per capita income or consumption, and the VSL elasticity parameter ϵ (Robinson, Hammitt, and O'Keeffe 2019). Below, we calculate population i's VSL in period t based on reference population x's VSL in period 1, population i's per capita consumption in period t $c_{i,t}$, population x's per capita consumption in period t, and the consumption elasticity of VSL ϵ:[55]

$$VSL_{i,t} = VSL_{x,1} \left(\frac{c_{i,t}}{c_{x,1}}\right)^{\epsilon}$$

This is equivalent to:

$$VSL_{i,t} = VSL_{x,1} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{-\epsilon}$$

Plugging this into the welfare-weighted SWF:

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta} [k_{i,t} - p_{i,t} VSL_{x,1} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{-\epsilon}] r_t$$

And then expanding:

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} [\left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta} k_{i,t} - p_{i,t} VSL_{x,1} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{-\epsilon}] r_t$$

Yields:

[55] Here, we represent the benefits transfer based on consumption instead of income to maintain consistency with the other equations, but the same logic holds if we use income everywhere instead of consumption, as mentioned in the first footnote in this section.

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} [\left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta} k_{i,t} - p_{i,t} VSL_{x,1} \left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta-\epsilon}] r_t$$

As discussed in the main text, under standard $\eta = \epsilon$, which implies that although individuals may value their statistical lives differently in current dollar terms, they value statistical lives the same in welfare terms. When this is the case:

$$W = \sum_{t=1}^{t=T} \sum_{i=1}^{i=n} [\left(\frac{c_{x,1}}{c_{i,t}}\right)^{\eta} k_{i,t} - p_{i,t} VSL_{x,1}] r_t \quad (10)$$

And this is equivalent to the current U.S. practice SWF from equation (8), except that non-monetary net benefits $k_{i,t}$ are welfare weighted. Thus, as discussed in the main text, the current U.S. practice welfare weights mortality net benefits, but Negishi-weights all other net benefits.

References

Abbott, Joshua K., and Eli P. Fenichel. 2014. "Following the Golden Rule: Negishi Welfare Weights without Apology." *Unpublished manuscript*.

Acland, Daniel, and David H. Greenberg. 2022. "Principles and Practices For Distributional Weighting: A New Approach." *SSRN Electronic Journal*. https://www.ssrn.com/abstract=4067472 (October 19, 2022).

Anthoff, D. 2007. "Report on Marginal External Damage Costs Inventory of Greenhouse Gas Emissions." *NEEDS–New Energy Externalities Developments for Sustainability, Delivery* (5.4-RS).

Anthoff, David, Cameron Hepburn, and Richard S J Tol. 2009. "Equity Weighting and the Marginal Damage Costs of Climate Change." *Ecological Economics*: 46.

Arrow, Kenneth J. et al. 1996. "Intertemporal Equity, Discounting, and Economic Efficiency." *Climate Change 1995-Economic and Social Dimensions of Climate Change*.

Aten, Bettina H. 2017. "Regional Price Parities and Real Regional Income for the United States." *Social Indicators Research* 131(1): 123–43.

Aten, Bettina H, Eric B Figueroa, and Troy M Martin. 2012. "Regional Price Parities for States and Metropolitan Areas, 2006–2010." *U.S. Bureau of Economic Analysis*: 14.

Boadway, Robin, and Michael Keen. 1993. "Public Goods, Self-Selection and Optimal Income Taxation." *International Economic Review*: 463–78.

Boadway, Robin W., and Neil Bruce. 1984. *Welfare Economics*. B. Blackwell New York.

Bressler, R. Daniel. 2021. "The Mortality Cost of Carbon." *Nature Communications* 12(1): 4467.

Bressler, R. Daniel, Frances C. Moore, Kevin Rennert, and David Anthoff. 2021. "Estimates of Country Level Temperature-Related Mortality Damage Functions." *Scientific Reports* 11(1): 20282.

Bruce, James P., Hoesung Lee, and Erik F. Haites. 1996. *Climate Change 1995. Economic and Social Dimensions of Climate Change*. Intergovernmental Panel on Climate Change.

Carleton, Tamma et al. 2022. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." *The Quarterly Journal of Economics*: 69.

Committee on Assessing Approaches to Updating the Social Cost of Carbon, Board on Environmental Change and Society, Division of Behavioral and Social Sciences and Education, and National Academies of Sciences, Engineering, and Medicine. 2017. *Valuing Climate Changes: Updating Estimation of the Social Cost of Carbon Dioxide*. Washington, D.C.: National Academies Press. https://www.nap.edu/catalog/24651 (December 17, 2018).

Cromar, Kevin, Peter Howard, Váleri N. Vásquez, and David Anthoff. 2021. "Health Impacts of Climate Change as Contained in Economic Models Estimating the Social Cost of Carbon Dioxide." *GeoHealth* 5(8). https://onlinelibrary.wiley.com/doi/10.1029/2021GH000405 (January 31, 2022).

Cromar, Kevin R. et al. 2022. "Global Health Impacts for Economic Models of Climate Change: A Systematic Review and Meta-Analysis." *Annals of the American Thoracic Society* 19(7): 1203–12.

Eeckhoudt, Louis R., and James K. Hammitt. 2001. "Background Risks and the Value of a Statistical Life." *Journal of Risk and Uncertainty* 23(3): 261–79.

Errickson, Frank C. et al. 2021. "Equity Is More Important for the Social Cost of Methane than Climate Uncertainty." *Nature* 592(7855): 564–70.

Gasparrini, Antonio et al. 2017. "Projections of Temperature-Related Excess Mortality under Climate Change Scenarios." *The Lancet Planetary Health* 1(9): e360–67.

Gayer, Ted, and W. Kip Viscusi. 2016. "Determining the Proper Scope of Climate Change Policy Benefits in U.S. Regulatory Analyses: Domestic versus Global Approaches." *Review of Environmental Economics and Policy* 10(2): 245–63.

Groom, Ben, and David Maddison Pr. 2019. "New Estimates of the Elasticity of Marginal Utility for the UK." *Environmental and Resource Economics* 72(4): 1155–82.

Grubb, Michael. 2005. "Stick to the Target The Most Recent Attempt to Undermine Kyoto Does Not Make Sense." *PROSPECT-LONDON-PROSPECT PUBLISHING LIMITED*- 114: 14.

Hales, Simon et al. 2014. *Quantitative Risk Assessment of the Effects of Climate Change on Selected Causes of Death, 2030s and 2050s*. http://apps.who.int/iris/bitstream/10665/134014/1/9789241507691_eng.pdf (June 6, 2019).

Hausman, Daniel M., and Michael S. McPherson. 2006. *Economic Analysis, Moral Philosophy and Public Policy*. Cambridge University Press.

Hemel, Daniel. 2022. "Regulation and Redistribution with Lives in the Balance." *The University of Chicago Law Review*: 86.

Her Majesty's Treasury. 2022. "The Green Book: Central Government Guidance on Appraisal and Evaluation." *London: HM Treasury*.

Hicks, John R. 1939. "The Foundations of Welfare Economics." *The economic journal* 49(196): 696–712.

———. 1941. “The Rehabilitation of Consumers’ Surplus.” *The Review of Economic Studies* 8(2): 108–16.

Houser, David, and Cass R. Sunstein. 2021. “Best Practices for the Treatment of a Statistical Life in U.S. Department of Homeland Security Regulatory Analyses.” https://www.dco.uscg.mil/Portals/9/VSL%20Final%20for%20DIST%205April21.pdf.

Howard, Peter, and Jason Schwartz. 2017. “Think Global: International Reciprocity as Justification for a Global Social Cost of Carbon.” *Columbia Journal of Environmental Law* 42(S). https://journals.library.columbia.edu/index.php/cjel/article/view/3734 (October 27, 2022).

Kaplow, Louis. 2005. “The Value of a Statistical Life and the Coefficient of Relative Risk Aversion.” *Journal of Risk and Uncertainty* 31(1): 23–34.

Kaplow, Louis, and Steven Shavell. 1994. “Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income.” *The Journal of Legal Studies* 23(2): 667–81.

———. 2009. *Fairness versus Welfare*. Harvard university press.

Kolstad, Charles et al. 2014. “Social, Economic and Ethical Concepts and Methods.” *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovern- mental Panel on Climate Change*.

Layard, R., G. Mayraz, and S. Nickell. 2008. “The Marginal Utility of Income.” *Journal of Public Economics* 92(8–9): 1846–57.

Liscow, Zachary. 2013. “Reducing Inequality on the Cheap: When Legal Rule Design Should Incorporate Equity as Well as Efficiency.” *Yale LJ* 123: 2478.

———. 2018. “Is Efficiency Biased.” *U. Chi. L. Rev.* 85: 1649.

———. 2021. “Redistribution for Realists.” *Iowa L. Rev.* 107: 495.

Markovits, Richard S. 2004. “Why Kaplow and Shavell’s Double-Distortion Argument Articles Are Wrong.” *Geo. Mason L. Rev.* 13: 511.

Masterman, Clayton J., and W. Kip Viscusi. 2018. “The Income Elasticity of Global Values of a Statistical Life: Stated Preference Evidence.” *Journal of Benefit-Cost Analysis* 9(3): 407–34.

Matthey, Astrid, and Björn Christian Bünger. 2019. *Methodological Convention 3.0 for the Assessment of Environmental Costs: Cost Rates: Version 02/2019*. German Environment Agency (UBA).

McGartland, Al. 2021. “Quality Science for Quality Decisions: Protecting the Scientific Integrity of Benefit–Cost Analysis.” *Review of Environmental Economics and Policy* 15(2): 340–51.

Mirrlees, James A. 1971. “An Exploration in the Theory of Optimum Income Taxation.” *The review of economic studies* 38(2): 175–208.

———. 1978. “Social Benefit-Cost Analysis and the Distribution of Income.” *World Development* 6(2): 131–38.

NASA. 1996. “Final Tier 2 Environmental Impact Statement for International Space Station.”

Negishi, Takashi. 1960. “WELFARE ECONOMICS AND EXISTENCE OF AN EQUILIBRIUM FOR A COMPETITIVE ECONOMY.” *Metroeconomica* 12(2–3): 92–97.

Nordhaus, William D. 2017. “Revisiting the Social Cost of Carbon.” *Proceedings of the National Academy of Sciences* 114(7): 1518–23.

OMB. 2003. “Circular A-4: Regulatory Analysis.”

Pearce, Fred. 1995a. “Global Row over Value of Human Life.” *New Scientist* 147(1991): 7–7.

———. 1995b. “Price of Life Sends Temperatures Soaring.” *New Scientist* 146(1971): 5–5.

Polinsky, A. Mitchell. 1972. “Probabilistic Compensation Criteria.” *The Quarterly Journal of Economics*: 407–25.

Rennert, Kevin et al. 2021. “The Social Cost of Carbon: Advances in Long-Term Probabilistic Projections of Population, GDP, Emissions, and Discount Rates.” *Brookings Papers on Economic Activity*.

———. 2022. “Comprehensive Evidence Implies a Higher Social Cost of CO2.” *Nature*. https://www.nature.com/articles/s41586-022-05224-9 (September 9, 2022).

Revesz, Richard L. 2018. “Regulation and Distribution.” *NYUL Rev.* 93: 1489.

Riahi, Keywan et al. 2017. “The Shared Socioeconomic Pathways and Their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview.” *Global environmental change* 42: 153–68.

Ritchie, Hannah. 2019. “Who Has Contributed Most to Global CO2 Emissions?” *Our World in Data. Available at: https://ourworldindata. org/contributed-most-global-co2*.

Robinson, Lisa A., James K. Hammitt, and Lucy O’Keeffe. 2019. “Valuing Mortality Risk Reductions in Global Benefit-Cost Analysis.” *Journal of Benefit-Cost Analysis* 10(S1): 15–50.

Rosenberg, John. 2001. “Toxic Memo.” *Harvard Magazine*.

Sanchirico, Chris William. 2000a. “Deconstructing the New Efficiency Rationale.” *Cornell L. Rev.* 86: 1003.

———. 2000b. “Taxes versus Legal Rules as Instruments for Equity: A More Equitable View.” *The Journal of Legal Studies* 29(2): 797–820.

Sunstein, Cass R. 2004. “Valuing Life: A Plea for Disaggregation.” *DUKE LAW JOURNAL* 54: 61.

Sunstein, Cass R. 2022. “Inequality and the Value of a Statistical Life.” *SSRN Electronic Journal*. https://www.ssrn.com/abstract=4236366 (October 19, 2022).

Timothy, Darren. 2021. “Departmental Guidance on Valuation of a Statistical Life in Economic Analysis.” https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis.

United States Senate, ed. 1993. *Nomination of Lawrence H. Summers: Hearing before the Committee on Finance, United States Senate, One Hundred Third Congress, First Session, on the Nomination of Lawrence H. Summers, to Be Under Secretary of the Treasury for the International Affairs, March 18, 1993*. Washington: U.S. G.P.O. : For sale by the U.S. G.P.O., Supt. of Docs., Congressional Sales Office.

US Department of Health and Human Services. 2021. “Guidelines for Regulatory Impact Analysis, Appendix D: Updating Value per Statistical Life (VSL) Estimates for Inflation and Changes in Real Income.” In ASPE.

US EPA. 2000. *An SAB Report on EPA’s White Paper Valuing the Benefits of Fatal Cancer Risk Reductions*. EPA-SAB-EEAC-00013. US Environmental Protection Agency, Science Advisory ….

———. 2011. “Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards.”

Viscusi, W. Kip, and Clayton J. Masterman. 2017. “Income Elasticities and Global Values of a Statistical Life.” *Journal of Benefit-Cost Analysis* 8(2): 226–50.