Journal of Benefit-Cost Analysis (2024), 1–17
doi:10.1017/bca.2024.21

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

The Value of Life in the Social Cost of Carbon: A Critique and a Proposal

John Broome[1,2]

[1]Faculty of Philosophy, University of Oxford, Oxford, UK
[2]Australian National University, Canberra, ACT, Australia
Email: john.broome@philosophy.ox.ac.uk

Keywords: EPA; IPCC; Kaldor–Hicks; life year; statistical life

Abstract
In its 2023 revision of the social cost of carbon, the U.S. Environmental Protection Agency values people's lives on the basis of their willingness to pay for them, without applying any distributional weights. It justifies this proposal on grounds of the Kaldor–Hicks criterion, which avoids interpersonal comparisons of wellbeing. But this criterion was discredited 70 years ago. Interpersonal comparisons of wellbeing cannot truly be avoided, and they should be used to determine distributional weights. One way of doing so is to identify as a numeraire a good that brings equal wellbeing to each person. A healthy life year is a reasonable, though only approximate, candidate for such a good. This article presents the point of view of a philosopher, regarding the practice of economists from outside the discipline.

1. Introduction

The U.S. Environmental Protection Agency (EPA) has recently reevaluated the social cost of carbon (EPA, 2023). The loss of lives due to climate change accounts for more than half of its new estimates.[1] To express this loss in monetary terms, the EPA values lives on the basis of what people are willing to pay to reduce their risk of dying. Since people in rich countries are willing to pay much more than people in poor countries to reduce their risk of dying, lives in rich countries are accorded much higher value than lives in poor countries. The EPA does not apply any adjustment or weighting to cancel out this disparity. It justifies its method of valuation by appealing to the Kaldor–Hicks criterion (EPA, 2023, p. 163), which is supposed to avoid making interpersonal comparisons of wellbeing. But I shall explain that this criterion was discredited 70 years ago.

Interpersonal comparisons of wellbeing cannot truly be avoided. They are needed to determine distributional weights that should be applied to people's willingnesses to pay.

[1] See EPA (2023, p. 80–1). This table gives figures drawn from different models, which are averaged to give the recommended figures (EPA, 2023, p. 100–1).

© The Author(s), 2024. Published by Cambridge University Press on behalf of Society for Benefit-Cost Analysis. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

More recently, the U.S. Office of Management and Budget (OMB) revised its guidance for major Federal regulations, including those issued by EPA, and authorized the use of distributional weights (OMB, 2023, p. 65–7).[2] One way of determining these weights is to identify a good that brings equal wellbeing to each person, and use it as a numeraire or standard of value. This article argues that a healthy life year is an acceptable, though only approximate, candidate for such a good.

There was a precedent for the EPA's valuation in a report presented to the Intergovernmental Panel on Climate Change (IPCC) in 1995 (IPCC, 1995). It was rejected by the world's governments. Section 2 tells this story. Section 3 presents the concept of a "statistical life" and explains how its use can be misleading. Section 4 describes an issue of fairness that is raised by this concept, and sets it aside. Section 5 goes on to the EPA's thinking, and describes how the Kaldor–Hicks criterion was discredited. Section 6 proposes the healthy life year as a numeraire. Section 7 sketches how it might be used in practice.

This article stems from a lecture that was given to a conference of philosophers and economists at the Brocher Foundation. It presents the point of view of a philosopher, regarding the practice of economists from outside the discipline.

2. Trouble at the IPCC

In 1995, the IPCC published its second assessment report. It includes a two-page box entitled "Attributing a monetary value to a statistical life" (IPCC, 1995, p. 196–7). This box suggests "as a reasonable first approximation" that the value of a statistical life in each country is proportional to income in that country. This implies that the value of a statistical life in poor countries is very much lower than in rich countries.

Because it is an intergovernmental panel, the IPCC submits its reports to its member governments for approval. The governments of a group of poor countries objected strongly to the implication that, for example, an American life is worth about 20 Indian lives. After a furore lasting several months, the report's Summary for Policymakers (IPCC, 1995, p. 5–15) was amended to accommodate the view of these countries, but the box in the main text remained unaltered.

The authors of this box were misunderstood. The responsible author, David Pearce (1996), said afterward: "The relevant chapter reports values of statistical life based on actual studies in different countries." The words "actual studies" makes it clear that the authors intended to report only empirical information about how much money people in different countries are willing to pay to reduce risk to their lives. People are willing to pay more in rich countries than in poor ones because they have more money to spend.

The authors claimed that willingness to pay is the "theoretically preferred method" of estimating what they call "the value of statistical life." They meant the *monetary* value of statistical life. But the word "value" is often used with a different meaning, as a synonym for "goodness." The "value of statistical life" could be understood to refer to the amount of goodness – or benefit – that is gained by saving a statistical life. So when the IPCC authors

[2] The OMB (2023, p. 67) notes that these weights should not be applied to the value of statistical life within the U.S. U.S. agencies already take the value of each U.S. life to be the same, even though a rich person is willing to pay more to avoid risk than a poor person is. This is implicitly to apply distributional weighting within the U.S. But the EPA does not take the value of each person's life to be the same throughout the world. It does not apply distributional weighting even implicitly between countries.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

assigned a much lower value to the life of an Indian than to the life of an American, they could be understood as saying that saving an Indian does much less good, or is much less beneficial, than saving an American. To put it another way: that an Indian is worth much less than an American.

That was not what the authors meant. They recognized that, before it could express the amount of goodness or benefit that is achieved by saving a life, the monetary value of statistical life needs to be adjusted. The adjustment can be made, they say, when the values of different people's lives are aggregated together to arrive at an estimate of the overall value of saving lives. The authors say:

> Aggregation in this context is a political and ethical process… The aggregation process makes it possible to correct for factors not reflected in individual WTP [willingness to pay] estimates, such as the injustice in the underlying income distribution or different responsibilities for the climate change problem. The process may thus result in changes in monetary values based on income levels as they are computed conventionally. For example, if VOSL [value of statistical life] is scaled in proportion to income, as suggested above, and aggregation weights are inversely related to income, weighted VOSLs will effectively be equal across countries.

The possible reasons they mention for adjusting monetary values are injustice in the income distribution and different responsibilities for climate change. Oddly, they do not mention that money is more valuable – in the sense of goodness – at the margin to poor people than to rich people: an extra dollar brings more benefit to the poor than to the rich. But that is the true reason why monetary values require adjustment if they are to measure benefit properly. Still, the authors did recognize that adjustments may be required, though they did not insist that they are required.

So the authors did not intend to imply that lives in poor countries are worth much less than lives in rich countries. But in using the word "value" incautiously they exposed themselves to this interpretation, and they did not take enough trouble to repudiate it. Obviously, the governments of poor countries could not accept this implication as they interpreted it. The misunderstanding could have been corrected by a little editing, but corrections were not made and the damage was done. The IPCC was no longer trusted on the issue of valuing lives.

The ambiguity of "value" is awkward. From here on in this article, when I particularly refer to monetary value, I shall write "monetary value," and I shall write "value" unqualified to refer to value in the sense of goodness. However, since monetary value can be taken to be a measure of value in the sense of goodness, in some contexts the unqualified word "value" will encompass monetary value too.

Subsequent IPCC assessment reports have made some attempt to repair the damage done in 1995, but they have largely avoided the issue. The third report (IPCC, 2001) treats the value of life only briefly, and recommends using the same monetary value in every country. The fourth (IPCC, 2007) touched on the issue even more briefly, simply referring to this recommendation from the third. The fifth (IPCC, 2014, p. 226) contains a box on the value of life, which suggests that a quality-adjusted life year might be given equal value in every country. Section 6 of this article takes up that suggestion. The Summary for Policymakers of the fifth report makes a general recommendation in favor of adjusting monetary values: "Ethical theories based on social welfare functions imply that distributional weights, which take account of the different value of money to different people, should be applied to monetary measures of benefits and harms" (IPCC, 2014, p. 5). I can find no mention of the value of life in the sixth report (IPCC, 2022).

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

In the 1995 report, cost–benefit analysis was taken seriously as a contribution toward deciding how much and how quickly climate change should be mitigated (IPCC, 1995, particularly chapter 5). But it has received less attention in subsequent reports. The economist's approach to decision making is to set a value on the costs and benefits of mitigation and weigh them against each other. A program for decarbonizing can be built on this basis, as costs and benefits develop over time. Nordhaus (2008) is an example. But by now, under the influence of the IPCC, this approach has been completely supplanted in determining what the rate of decarbonization should be.

Its place has been taken by methods from physics, which are centered on the idea of a carbon budget aimed at a particular target temperature. This target is set by a political process that is not transparent, but definitely does not involve weighing costs against benefits in any but the vaguest sense. Given the target, physicists estimate a carbon budget, defined as the amount of carbon that can be released into the atmosphere without exceeding the target temperature, or more accurately without creating more than a specified small probability of exceeding the target temperature. The very idea of a carbon budget implies that there must come a time of net zero carbon emissions, after which no more carbon is added to the atmosphere. These days, large-scale planning for climate change is focused on the date when net zero will be achieved. For instance, the US has committed to 2050 as its net-zero date (U.S. Department of State, 2021).

Central to the methods of economics, by contrast, is the social cost of carbon. This measures the benefit of reducing carbon emissions, which must be weighed against the cost of doing so. But we have no figure for the social cost of carbon that commands a consensus. A large part of it – a half or more – is made up of what are called "mortality costs," which measure the harm climate change does by killing people.[3] These mortality costs depend on the value of a statistical life, which has been mired in controversy since the 1995 IPCC report. For this and other reasons, the methods of economics have lost their authority in the international politics of climate change. For example, the social cost of carbon is not mentioned in the Synthesis Report of the IPCC's Sixth Assessment Report (IPCC, 2023).

The absence of economics from the top level of planning seriously diminishes the world's hope of decarbonizing successfully. Economics has important insights that cannot be ignored if climate change is to be overcome. It gives proper importance to efficiency and to the role of carbon pricing in achieving efficiency. Planning without regard to economics is so inefficient and so wasteful that it may prevent us from achieving net zero when it is needed.

An example appears in the UK National Health Service's (2020) plan to reach net zero carbon emissions by 2045. The NHS plans to use some of its resources – which were originally intended for improving the health of British people – to reduce emissions. Yet it appears to take no account of the cost of doing so. It does not even explain why it aims to achieve net zero emissions at that date. This needs to be explained. Even if Britain as a whole must achieve net zero emissions, it does not follow that every institution in Britain must do so.

Even if economics cannot successfully identify a social cost of carbon on grounds of value, it could nevertheless play an important role in achieving efficiency. At the highest level of planning, a target date for achieving net zero can be set on the basis of a target temperature and a carbon budget – as actually happens. Then the social cost of carbon can be set at whatever level it needs to be at in order to achieve this target (cf. Stern *et al.*, 2022,

[3] See note 1 above.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

p. 192). This figure can be used as a shadow price in more detailed decision making, for instance in the UK NHS and elsewhere. It could also be made a real carbon price in the market, by means of either a carbon tax or some other method.

Fixing a social cost of carbon on the basis of a target may be the best we can do in our present circumstances. We have no consensus on a social cost of carbon based on valuation, whereas there does seem to be an international political consensus on 2050 as the date when global net zero needs to be achieved (United Nations, undated). I do not say this consensus is well founded. Nevertheless it does supply a generally accepted starting point that can be used for determining a social cost of carbon. Unless we can establish a sound basis in valuation, we may have no alternative.

The focus of this article is on whether the social cost of carbon can indeed be given a sound basis in valuation. I shall therefore not pursue any further the idea of basing it on a target.

3. "Statistical" lives

What is the point of the word "statistical" in the "value of statistical life"? It is a reminder that, when a public policy extends lives as climate policy does, it is not usually known in advance whose lives will be extended. In evaluating a policy, what is being valued is the lives of random people among the population.

Why is this reminder needed? A policy to improve transport by building a railway line, say, benefits people by reducing their travel time. But the measure of benefit is not usually called "the value of a statistical hour," even though it is not known in advance who will receive it. The word is used particularly for the value of lives because those who use it are keen to insist that in determining the monetary value of a statistical life they are not setting a monetary value on an actual human lives.

For instance, the IPCC (1995, p. 196) authors says:

> It is important to understand that what is valued is a change in the risk of death, not human life itself.

The OMB (2023, p. 29) says:

> The latter phrase ['the value of life'] can be misleading because it suggests erroneously that the monetization exercise tries to place a "value" on individual lives.

The EPA (2023, p. 164, note 175) repeats the OMB's words.

Kip Viscusi (2008, p. 586) says:

> This designation [the 'value of life'] is in many respects a misnomer. What is at issue is usually not the value of life itself but rather the value of small risks to life.

Robinson *et al.* (2019, p. 16) say:

> It [the value of statistical life] is not the value that the individual, the society, or the government places on averting a death with certainty. Rather, it represents the rate at which an individual views a change in money available for spending as equivalent to a small change in his or her own mortality risk.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

Despite all these claims to the contrary, to value a risk to life is indeed to value a life. A risk is bad only because it is a risk of something bad, and its badness stems from the badness of what it is a risk of. You cannot value a risk without implying a value for what it is a risk of. You cannot value a person's risk of death without implying a value for her death itself. Expected utility theory makes the connection precise: the utility of saving a person from a chance p of dying is p times the utility of saving her from a certain death.

Let us do the expected utility calculation for a simple case. Imagine a population of n people who are exactly alike in all respects that affect the value of their lives. Imagine some threat will take the life of a random one of them if it is not averted. But imagine the threat can be averted at a cost c, saving the life of the one random person. That is to say, a statistical life can be saved at cost c.

The expected utility from the social point of view conditional on averting the threat is

$$u(a, a, \ldots a, c).$$

The rth place in the vector shows the condition of the rth person, signified by a if she is alive or d if dead. The last place shows the cost of the action taken: c if the threat is averted or 0 if it is not.

The expected utility conditional on not averting the threat is the average of utility if the first person dies, utility if the second person dies, and so on up to utility if the nth person dies. That is:

$$(1/n)\{u(d, a, \ldots a, 0) + u(a, d, \ldots a, 0) + \ldots + u(a, a, \ldots d, 0)\}.$$

Since all the people are alike in all respects that affect the value of their lives, each term within the brackets $\{\}$ is the same. So this expected utility is

$$(1/n)\{nu(d, a, \ldots a, 0)\} = u(d, a, \ldots a, 0).$$

The gain in expected utility from averting the threat is the difference:

$$u(a, a, \ldots a, c) - u(d, a, \ldots a, 0).$$

It is better to avert the threat if this gain is positive, and better not to avert it if the gain is negative. The social monetary value of a statistical life is the amount society should be willing to spend to save a statistical life, which is the cost C at which the gain in utility from doing so is zero. That is, C such that $u(a, a, \ldots a, C) = u(d, a, \ldots a, 0)$.

Now imagine a different, focused threat that will certainly take the life of the first person if it is not averted at cost c. The expected utility conditional on averting this threat is

$$u(a, a, \ldots a, c)$$

and the expected utility conditional on not averting the threat is

$$u(d, a, \ldots a, 0).$$

So the gain in expected utility from averting the threat is the same:

$$u(a, a, \ldots a, c) - u(d, a, \ldots a, 0).$$

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

The social monetary value of the first person's life is the cost C such that $u(a, a, \ldots a, C) = u(d, a, \ldots a, 0)$. This is the same as the social monetary value of a statistical life.

In practice not all lives have the same social monetary value. The social monetary value of a statistical life is then approximately the average social monetary value of all individual lives. It is exactly the average if utility is linear in cost. We may conclude that expected utility theory implies it is wrong to treat the value of saving a statistical life as different in kind from the value of saving the life of a particular, identified individual (cf. Hammitt & Treich, 2007, p. 50–1).

The value of a statistical life is truly just the value of a life. Without compunction, I shall write freely of the value of life, often omitting "statistical."

4. Fairness

It is natural to feel some discomfort with the conclusion I drew from expected utility theory in Section 3. The discomfort can be made more poignant by varying the example a little. Imagine now that both threats appear simultaneously. Imagine society has resources to avert one of the threats but not both, so it must choose between saving an identified person or saving a statistical life. According to expected utility theory, it should be indifferent between the two. But intuitively, fairness favors saving the identified person. It is fairer for the risk of death to be distributed across everyone rather than concentrated on one person. So expected utility theory is apparently mistaken.

This is exactly the argument Peter Diamond (1967) made against expected utility theory. He used the same example, except that he included only two people rather than many. But I think we should not give up expected utility theory so easily. The best way to handle Diamond's objection is to separate fairness from other sorts of goodness. From here on in this article, I shall use "goodness" to refer to good things apart from fairness (Broome, 2004, p. 37–9). Fairness and goodness understood this way are separate considerations that should both be taken into account in decision-making. Expected utility theory provides a good account of the structure of goodness apart from fairness (Broome, 1991, chapter 5). The arguments of Section 3 are concerned with goodness apart from fairness, and they therefore remain valid.

I do not mean to suggest that fairness is not itself a good thing. It is. My reason for separating it out is that fairness requires a very different response from us than does goodness of other sorts. To put it roughly, goodness requires maximizing but fairness requires equalizing. Goodness is promoted by increasing people's wellbeing, whereas fairness among equally deserving people is promoted by equalizing their wellbeing or, failing that, by equalizing their chances of wellbeing. Indeed, inequality among equally deserving people is bad only because it is unfair. When inequality is bad, its badness is the badness of unfairness (Broome, 1991, chapter 9).

To illustrate the point further, here is a different example of valuing lives. Suppose two people are threatened with death, and the available resources can save only one of them. Suppose that if the first is saved she will go on to live in good health, but if the second is saved her future health will be poor. To add detail, suppose both people have had equally good lives up to the time when a choice has to be made between them, but at that time the second person is struck down by a disease. The disease is independent of the threat to her life, but it leaves her in poor health.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

Which person's life should be saved? More good is done by saving the first. If we attended only to goodness, we would save that person. But saving the first just because her health is better is unfair to the second person. She has already suffered the misfortune of being struck by a debilitating disease, and it would be unfair to require her also to give up her life just because of this misfortune. It would be fairest to hold a lottery at equal odds to decide whom to save. Goodness requires saving the first person; fairness requires a lottery. Fairness may win this conflict between fairness and goodness, and if so a lottery should be held.

Many philosophers and economists try to embed fairness within goodness rather than treat fairness and goodness as two separate considerations. Prioritarianism is one attempt to do so (e.g., Adler & Norheim, 2022). Prioritarians take the social value of a person's wellbeing to be a strictly concave function of her wellbeing. The result is diminishing marginal social value of wellbeing. An extreme version of prioritarianism has a maximin social value function. In the example, this gives an equal social value to saving the life of each of the two people. It implies the example is a choice between two options of equal value. Given that, a lottery could be justified as a tie-breaker. So this extreme version of prioritarianism seems able to explain what is intuitively fair in the example.

However, it is plainly wrong to give equal social value to saving the lives of the two people. Varying the example makes this plain. Suppose a single person's life is threatened, and there are two different ways of saving it. One is to apply a treatment that would leave the person in good health; the other to apply an equally costly treatment that would leave her in poor health. These options are parallel to the two options in the original version of the example, but in this case, there is no issue of fairness because only one person is involved. Extreme prioritarianism gives both options the same value. But that would make it permissible to hold a lottery between the two treatments, whereas actually it is plainly not permissible. In this case where fairness is not at issue, it is plain that the first treatment should be given. So the extreme version of prioritarianism is mistaken.

Trying to embed fairness within goodness fails to give a proper account of fairness. In making decisions, fairness and goodness should be taken as separate considerations. This article is concerned with just one of them: goodness, excluding fairness. Now I return to that subject.

5. The EPA and the Kaldor–Hicks criterion

The U.S. Environmental Protection Agency (2023) has recently settled on a new estimate of the social cost of carbon. It is an estimate based on valuation. It is not derived from or connected to the U.S.'s commitment to net zero emissions by 2050 (U.S. Department of State, 2021).

That overall commitment is based on physics and a temperature target. It remains in place. The EPA's social cost of carbon based on valuation is intended for lower-level decision making through cost–benefit analysis. It is clearly inconsistent and inefficient for a government to make lower-level decisions on a basis that is not aligned with its top-level, overall commitment. If the EPA's estimate were well founded it might in due course supersede the target-based approach, and allow the overall commitment to be revised on the basis of valuation. That would restore the authority of cost–benefit analysis in top-level planning. However, I shall explain that the EPA's estimate is not well founded, so it will not be able to achieve this result.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

More than half of the EPA's estimate consists of mortality costs: the value of lives lost to climate change.[4] The EPA values lives by unweighted willingness to pay (EPA, 2023, p. 163–7). It estimates the willingness to pay in each country by assuming that it is proportional to the country's per capital national income (EPA, 2023, p. 165). This implies, as the IPCC authors unwittingly implied in 1995, that the value of life in a poor country is much less than in a rich one.

In defense of its method, the EPA points out that people in rich countries are willing to pay more to protect their lives than are those in poor countries. It says:

> It is important to stress that this metric does not reflect the 'value' that this approach places on mortality risks in different parts of the world. Rather, it reflects an estimate of the willingness to pay for mortality risk reductions by the average resident of countries or regions conditional on their income. (EPA, 2023, p. 165)

So far so good. The EPA is making it clear that when it speaks of the value of life, it refers to monetary value as measured by willingness to pay and not to value in the sense of goodness. It puts scare quotes on "value" in this passage in order to stress that it has the second sense. In the next three paragraphs, I shall copy this use of scare quotes for the same purpose.

Next, in calculating the mortality cost of carbon emissions, the EPA adds up all the lives lost in countries around the world, valuing each according to the willingness to pay in its own country. So, for example, each Swiss life counts about the same as 20 Bangladeshi lives. These figures are to be used in U.S. decision making, which should be aimed at achieving the best result. So the EPA's figures show exactly how it "values" mortality risks in the sense that refers to goodness, despite what it says. It implies that saving Swiss lives is 20 times as good as saving Bangladeshi lives: a Swiss person is worth 20 Bangladeshis.

Whereas the IPCC authors were willing to apply weights to willingnesses to pay before using them for policy making, the EPA is not. It actually commits the sin that the IPCC authors were mistakenly criticized for. The implication that lives in poor countries are worth less than in rich countries cannot be accepted by most of the world. This decision of the EPA's will harm rather than improve the authority of economics in decision making about climate change.

To be sure, the EPA would deny that it "values" Swiss lives more than Bangladeshi ones. When it adds up monetary values of statistical lives across the world, this is not because it treats them as "values." It is because it is applying the Kaldor–Hicks criterion, which is supposed to determine right policy without making any "valuations." It is supposed to avoid interpersonal comparisons of "value." The EPA tells us that this criterion also underlies all other elements of the EPA's cost benefit analyses (EPA, 2023, p. 167). This is shocking. The Kaldor–Hicks criterion was thoroughly discredited 70 years ago.

What is the Kaldor–Hicks criterion and how was it discredited? It is supposed to be a test for whether a change in the economy is an improvement. The criterion is whether those who gain from the change would be able to make compensating money transfers to those who lose, in such a way that no one ends up worse off than before the change, and someone ends up better off. The criterion does not require the compensation actually to be paid. When the criterion is satisfied, the change is often said to be a "potential Pareto improvement." But it is not necessarily a genuine improvement.

[4] See note 1 above.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

The opposite claim that it *is* necessarily a genuine improvement was advanced in 1939 by Nicholas Kaldor (1939) and immediately endorsed by John Hicks (1939) with some equivocation (Hicks, 1939, p. 712). Neither author offered an argument in defense of their claim. They were more concerned to point out that their criterion for an improvement does not require the wellbeings of different people to be compared. Even economists who are sceptical about such comparisons – a growing number at that time – are able to apply their criterion. But the fact that a criterion can be applied does not mean it is correct, and Kaldor and Hicks did not argue for its correctness. Hicks (1939, p. 711–12) did discuss the question of whether, given that a change happens, compensation should be paid. But that is a different question.

Kaldor and Hicks claimed that the criterion is sufficient for a change to be an improvement, but neither claimed that it is necessary. Their criterion was therefore not even claimed to be fit to be the sole basis of cost–benefit analysis.

There is strong intuitive evidence against Kaldor's claim. It is easy to construct counterexamples to it. Here is one. Dives in his castle finds his view slightly marred by the hovel of Lazarus, who lives at his gate. Dives, being rich, would be willing to pay a substantial amount to have the hovel moved to a less salubrious neighborhood outside his field of view. Lazarus, being poor, would be willing to accept this amount; it would more than compensate him for moving his home. Would it be an improvement to move Lazarus without the compensation's being paid? It would not. The benefit to Dives would be slight: a small improvement in his view. The harm to Lazarus would be large: living in a less salubrious neighborhood. On balance the change is bad.

This intuitive evidence was reinforced within 2 years by Tibor Scitovsky's (1941) demonstration that Kaldor's claim contradicts itself. Scitovsky demonstrated that there can be two states of the economy, *A* and *B*, such that *A* is an improvement on *B* according the criterion, and *B* is an improvement on *A* according to the criterion. This implies a contradiction.

Scitovsky suggested augmenting the criterion in a way that removes this objection. According to Scitovsky's criterion, a change from *A* to *B* is an improvement if it satisfies the Kaldor–Hicks criterion and a change from *B* to *A* does not. In effect, this is to accept Kaldor's claim except in cases where it contradicts itself in the way Scitovsky described.

Scitovsky's criterion remains subject to the intuitive objection. To the story of Dives and Lazarus, we can add a further assumption that Lazarus, were he living in the less salubrious neighborhood, would not be able to pay Dives enough to persuade him to allow Lazarus to move his hovel to Dives's gate. The conclusion according to Scitovsky's criterion would still be the same: moving Lazarus to the less salubrious neighborhood would be an improvement. Intuitively, this is still incorrect.

Fourteen years later, Terence Gorman (1955) put the last nail in the coffin. He demonstrated that Scitovsky's two-way criterion, like the original Kaldor–Hicks criterion, also implies a contradiction. There can be three states of the economy, *A*, *B* and *C*, such that *A* is an improvement on *B* according to the criterion, *B* is an improvement on *C* according to the criterion, and *C* is an improvement on *A* according to the criterion. This implies a contradiction.

In sum, there was no argument for the Kaldor–Hicks criterion, whereas there are strong counterexamples to it, and it is self-contradictory. It could not have been more thoroughly discredited. This happened 70 years ago.

But it is not the end of the story. Despite its assertion, the EPA does not actually apply the Kaldor–Hicks criterion. Practical cost–benefit analyses never do. The criterion they

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

conventionally apply is whether the total of what the gainers from a change would be willing to pay to have the change is greater than the total of what the losers would be willing to accept as compensation for the change. To put it another way: whether the total of compensating variations for the change is positive. If it is, conventional cost–benefit analysis takes the change to be an improvement. Surprisingly perhaps, this is not the Kaldor–Hicks criterion. If the gainers were to transfer money to the losers, it would alter the distribution of wealth and consequently alter market prices. People's compensating variations are calculated at the prices that prevail before any transfer. If the gainers tried to compensate the losers, prices might alter in a way that makes it impossible for the gainers to compensate the losers fully. This was demonstrated by Robin Boadway (1974). So the criterion used by the EPA and typically used in cost–benefit analysis is not the Kaldor–Hicks criterion.

This makes things worse. The only justification offered for the criterion of a positive total of compensating variations is that it is supposed to be equivalent to the Kaldor–Hicks criterion. Blackorby and Donaldson (1990) have explored in detail the relation between the two criteria. They show that (if lump-sum transfers are possible) a necessary condition for the Kaldor–Hicks criterion to be satisfied is that the total of compensating variations is positive. This is not a sufficient condition. Yet the Kaldor–Hicks criterion was proposed as a sufficient condition for a change to be a improvement, and not as a necessary condition. So, even if the Kaldor–Hicks criterion were indeed a sufficient condition for a change to be a improvement, the condition that the total of compensating variations is positive would be neither a necessary nor a sufficient condition for a change to be a improvement.

The basic problem is that these criteria are trying to do the impossible. When a change is good for some people and bad for others, whether or not it is good overall depends on comparing the benefit to some against the loss to others. Some people's wellbeing must be compared with others'. A cost–benefit analysis implicitly makes that comparison. Yet these criteria refuse to make it explicitly. They therefore cannot succeed. The EPA should not use them.

6. Interpersonal comparisons

Some cost–benefit analysts are sceptical about comparing the wellbeings of different people, and that is why they find these criteria attractive. But interestingly, the value of life provides a means of overcoming their scepticism.

One approach to comparing wellbeings is to find some single good that can fairly be taken to bring the same wellbeing, or have the same value in the sense of goodness, to each person. This good can be used to calibrate everyone's scale of wellbeing, to bring each person on to the same scale. The good of life might play this role.

It seems to be widely supposed that each person's life has the same value in this sense. For instance, the authors of the criticized box in the 1995 IPCC report say:

> At how much ought a statistical life be valued according to ethical or other criteria? An obvious implication of this [the 'prescriptive or normative'] approach is that all lives will be treated equally. Each statistical life saved should have the same value. (IPCC, 1995, p. 96)

The IPCC (2001) recommended using the same value of life for every person in the world.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

The EPA ascribes the same value to every American life, despite the fact that not every American is willing to pay the same amount of money for a decrease in risk to her life. As Bressler and Heal (2022) point out, this valuation by the EPA does not accord with the Kaldor–Hicks criterion, which the EPA officially endorses. Indeed, the OMB (2023, p. 67, note 129) recognizes that the EPA's valuation implicitly applies distributional weights to the monetary value of life. Other countries' cost–benefit analyses also ascribe equal value to each member of their population (Bressler & Heal, 2022, p. 3).

Suppose we accept this idea and adopt life as a numeraire that has the same value for everyone. Then we give each person's life the same monetary value. We might base this monetary value on (say) the population average of willingness to pay to decrease risk. This is what the EPA does for the American population. But merely giving each person's life the same monetary value is not enough. If everyone's monetary value of life is equalized, but other monetary values are not correspondingly adjusted, people's relative values will be distorted. For some, the monetary value of their lives will be raised above other things that they actually value more, and for others it will be depressed below other things they actually value less. If cost–benefit analysis is done on the basis of this valuation, the result will be inefficient.[5] It will turn out that some people could be benefited without any cost in resources, by shifting resources away from reducing risk to their lives and toward other things that they value more. Others could be benefited by shifting resources in the opposite direction.

The differing values of money to different people mean that different people's willingnesses to pay for a particular good are not a good indication of the relative values of that good to those different people. However, a single person's willingnesses to pay for different goods are a good indication of the relative values of those goods to that single person. For that reason they are an excellent basis for cost–benefit analysis. Cost–benefit analysis will lead to inefficiency unless the relative values of goods to each person are respected. For the sake of efficiency, each person's relative willingnesses to pay need to be maintained. But all of her willingnesses to pay need to be multiplied by a distributional weight that reflects the value of money to the person. This is exactly what the IPCC (2014, p. 5) recommends and the OMB (2023, p. 65–7) now judges acceptable.

The way to use life properly as the numeraire is to let it determine distributional weights for each person. These weights must be such that, when people's willingness to pay are weighted by them, each person's weighted willingness to pay for reducing risk to life comes out the same. These same weights must be applied to willingnesses to pay for all other goods too. Then the relative values of different goods to each person will be preserved. Inefficiency will be averted.[6]

However, I now need to say it cannot be true that each person's life has the same value, despite the widespread agreement that it is. It is not even true for a single person that each of her potential lives has the same value. A longer life is better for her than a shorter one. What is described as "the value of life" is truly the value of extending life. Extending life is valuable only because a longer life is better than a shorter one. So recognizing that life is valuable is recognizing that a longer life is better than a shorter one. It is therefore inconsistent with valuing everyone's life equally. Saving some people's lives adds more time to their lives than saving other people's lives does. It is therefore more valuable.

[5] This point is well known. See Bressler and Heal (2022, p. 11–13), who attribute it to Cass Sunstein (2004).
[6] This idea was proposed by Somanathan (2006) and explored in more detail by Baker *et al.* (2008).

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

This suggests that a better numeraire, which we should take to have the same value for each person, is a year of life rather than life as a whole.[7] But this too cannot be correct because even for a single person a year of life does not always have the same value. A year of life in good health is better than one in poor health. This is why people strive for good health.

So an even better numeraire is a year of life in good health: a healthy life year. This is in effect the numeraire commonly used in health economics and public health analysis. These disciplines commonly treat some version of the quality-adjusted life year, or QALY, as having the same value everywhere. The IPCC (2014, p. 226) makes the same suggestion. The World Health Organization similarly uses the disability-adjusted life year, or DALY, in its measures of the burden of disease (Roser *et al.*, 2021).

Even the healthy life year is not exactly a correct numeraire. Not all healthy life years are equally good for a person. Some are better than others because in some years a person's life goes better than in others, in ways that are unconnected with her health. It is even plausible that a healthy life year for a rich person is typically better than a healthy life year for a poor person, because life generally goes better for rich people. The World Happiness Report (Helliwell *et al.*, 2023, Table 2.1) finds a significant correlation between happiness in a country and the log of the country's per capita GDP. This might add a shadow of justification to the greater value the EPA gives to the lives of people in rich countries than to those of people in poor countries. But it cannot come close to justifying the EPA's claim that the value of life in a country is proportional to income in the country. The World Happiness Report shows happiness increasing far less than proportionally to GDP.

Fleurbaey and Hammitt (this issue) raise an objection to taking a healthy life year as numeraire. They say it "will tend to be unfair to individuals with a low life expectancy." People who have a short life expectancy – feeling short of life – will tend to give more value to extending their lives, relative to other goods, than will those who have a longer life expectancy. So if all healthy life years are given the same value, other goods consumed by people with shorter life expectancy will be valued less than those consumed by people with a longer life expectancy. Fleurbaey and Hammitt think this is unfair to those with lower life expectancy.

They must mean to cast doubt on the assumption that a healthy life year has the same value whoever it comes to. If it does have the same value, we should accord it the same value. Then, if we accept that a person's preferences indicate the relative values of different goods to her, we should accept the consequences that follow for the valuation of these goods. So the question is whether Fleurbaey and Hammitt have identified a reason to doubt that a healthy life year has the same value for each person.

People with lower life expectancy fall into two classes. First, there are those who have already lived a long life and do not have much more to go. These people have presumably had a chance to plan their lives, and they should be receiving in their old age the goods they had planned to receive. If they now find they would prefer to sacrifice more of those goods for extending their lives, I do not see why this should make their life years more valuable compared with other people's. It would be odd to think that later years of life are more valuable than earlier ones.

[7] The life year was the unit of account used in a recent large and comprehensive study of the mortality costs of climate change (Carleton *et al.*, 2022). In part of the study, the authors took the value of a life year to be the same for each person. But in the finally published version, this part was unfortunately relegated to an appendix (Carleton *et al.*, 2022, online appendix Table H3.)

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

Second, there are the people who have lower life expectancy at birth. Typically these will include the inhabitants of poorer countries. These people are short of many goods including years of their lives. Valuing life years equally means that their whole lives are valued less than the whole lives of people from richer countries. And indeed their lives are less good than the lives of people from richer countries, because a shorter life is less good than a longer one. So the relative values accorded to whole lives are as they should be, and along with that, so are the relative value of other goods.

To be sure, the situation is unfair. As I said in Section 4, when people are equally deserving, fairness requires them to have equal wellbeing. It is consequently unfair for people in some countries to have shorter lives that people have in other countries. The unfairness is in the situation, not the valuation. It should be corrected by redistribution toward poorer countries. Because fairness is separate from goodness, the redistribution required by fairness is over and above whatever redistribution would be implied by maximizing overall goodness, as measured in the numeraire of healthy life years. This is no reason to revise the valuations given by this numeraire. That is the lesson of Section 4.

I am therefore unconvinced by Fleurbaey and Hammitt's argument. I continue to think that, as a first approximation, it is acceptable to assume that a healthy life year has equal value to each person. By comparison, conventional cost–benefit analysis without distributional weights implies that a dollar has equal value to each person. That is a much worse approximation. So let us continue with the healthy life year.

We have to go further than health economists and public health analysts do. They concentrate on health only, and aim at attaining the population's highest level of health. But cost–benefit analysis takes account of all good things and all sorts of costs. The values of all of these goods and bads are initially measured in terms of willingnesses to pay. These should then be multiplied by distributional weights, and we now have a method of determining those weights. They should be such that the value of a healthy life year comes out the same for each person.

7. Practical matters

A theoretical basis of cost–benefit analysis using healthy life years as numeraire is worked out by David Canning (2013). But how should we determine the distributional weights in practice? We must start by determining, for each country, what people are willing to pay in their own currency for a healthy life year. There is little direct data about this, so this figure will have to be found by extrapolating from the data we have (Robinson *et al.*, 2019, p. 35). This consists mainly in empirical estimates of willingness to pay for reducing risk to whole lives, obtained from studies in developed countries. They give us estimates of the monetary value of a whole life in those countries. From these, estimates of the monetary value of a healthy life year in those countries can be derived. Then these estimates can be projected on to other countries.

The first step of this extrapolation is deriving an estimate of the monetary value of a healthy life year on the basis of data on the monetary value of a whole life. This can be done by assuming that people evaluate their lives as an aggregate of the goodness of the life years that compose them. Specifically, assume that they value a life at the discounted sum of its quality-adjusted life years. "Quality" refers to the quality of health only; by a "quality-adjusted life year" I mean a QALY as it is standardly understood. The discount rate is the rate at which a

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

person discounts her own future, not the rate at which the value of life years should be discounted in policy making. It is reported to be around 6 or 8% (Jones-Lee *et al.*, 2015, p. 11).

This assumption implies that the monetary value of a healthy life year for a population is approximately the monetary value of a life for the population, divided by the population's average discounted quality-adjusted life expectancy (compare Jones-Lee *et al.*, 2015). The divisor in this fraction is not a readily available quantity. Approximations to quality-adjusted life expectancies are available for all countries from the World Health Organization, in the form of what it calls "healthy life expectancies" (WHO, undated). However, the WHO's figures are undiscounted, and they are published for only two ages – at birth and at 60 – whereas we need an average for the population. So some computation will be needed to estimate the divisor. The task is made easier by the fact that it is not needed for every country. An estimate for one country will suffice, provided it is one for which a reliable monetary value of a whole life is available. We need only one estimate for the monetary value of a healthy life year, which can then be projected to other countries.

The second step of the extrapolation is achieved by assuming a relationship between people's income and their willingness to pay to extend life. Specifically, we assume a particular elasticity of the monetary value of a healthy life year with respect to per capita income. Some data is available for the income elasticity of the monetary value of a whole life. It is reviewed by Robinson *et al.* (2019, p. 30–4). These authors arrive at 1.5 as an average figure for this elasticity when extrapolating from the US to low-income or middle-income countries, and at a lower figure when extrapolating from the US to other high-income countries. The EPA takes it to be 1 (EPA, 2023, p. 165). The income elasticity of the monetary value of a healthy life year will be somewhat lower, because poorer countries generally have shorter quality-adjusted life expectancies. But this effect is small. For instance, Switzerland, with about 20 times the per capita income of Bangladesh, has a life expectancy only about 15% longer. To extrapolate monetary values to countries where data is unavailable, this elasticity should be applied to their national income per capita measured in real terms – that is to say, at purchasing-power parity.

Once these two steps have been completed, we shall have determined monetary values of healthy life years for each country in its own currency. Then we can use those to determine exchange rates among currencies that make these values equal everywhere. Let us call these "value parity" exchange rates.

Take an example. The rupee to dollar value parity rate will be much higher than the rupee to dollar purchasing-power parity rate. Purchasing-power parity makes the rupee price of goods the same as their dollar price. But since people in India are poorer than people in the US, goods are more valuable to them. So purchasing-power parity undervalues Indian people's goods, including years of their lives, relative to U.S. goods and life years. At the value parity rate, healthy years of life will be accorded equal value, and other goods in India will be accorded much more value that they are given at present.

Calculating value parity exchange rates is like running in reverse the extrapolation of monetary values described by Robinson *et al.* (2019), from countries that have direct data to those that do not. Whereas the extrapolation of monetary values ascribes less monetary value to healthy life years in poorer countries, this reverse calculation ascribes more value to all goods in poorer countries. The ratio between a country's value parity rate and its purchasing-power parity rate is in effect a distributional weight applied to all the goods in the country.

The crucial parameter in determining this weight is the elasticity of the monetary value of a healthy life year with respect to per capita income. Since we are assuming that a healthy life

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

year is equally good for everybody, this is nothing other than the income elasticity of the marginal utility of money. As an illustration: if this elasticity is 1, the value parity rates of exchange will make income per capita equal in every country.

The considerations I mentioned above suggest that 1.4 would not be an unreasonable figure for the elasticity. This is for international comparisons. Interestingly, 1.4 is just the elasticity that the OMB (2023, p. 67) says is a "reasonable estimate" to use as a basis for distributional weighting within the US.

A final note. I suggest value parity exchange rates should be used to bring the average monetary value of a healthy life year into equality in all countries, whilst maintaining the relative prices of goods within countries in the interests of efficiency. But it will not achieve complete efficiency because it does not apply distributional weighting within countries. Within countries there will still be the inefficiency mentioned in Section 6.

Acknowledgements. This article developed from a lecture I gave at a 2022 conference held by the Brocher Foundation: "Healthy, Wealthy, and Wise – The Ethics of Health Valuation," organized by Nir Eyal (Rutgers University), Samia Hurst (University of Geneva), Lisa A. Robinson (Harvard University), and Daniel Wikler (Harvard University). I thank the Brocher Foundation and conference participants for their support and helpful comments. I have also received valuable and extensive comments from Michael Geruso, James Hammitt, Dean Spears, and three anonymous referees. I particularly thank Daniel Bressler for introducing me to the EPA's report discussed in this article. Bressler and Heal (2022) is another comment like this one on the EPA's treatment of the value of life. Its argument is similar and in many respects developed in more detail.

Funding statement. This special issue of the journal was supported by the Brocher Foundation, with supplemental funding from the Rutgers University Center for Population-Level Bioethics, the University of Geneva Institute for Ethics, History, and the Humanities, and the University of Bergen Centre for Ethics and Priority Setting. More information on the Brocher Foundation is available at https://fondation-brocher.ch/.

References

Adler, Matthew D., and Ole Norheim. 2022. *Prioritarianism in Practice*. Cambridge: Cambridge University Press.

Baker, Rachel, Susan Chilton, Michael Jones-Lee, and Hugh Metcalf. 2008. "Valuing Lives Equally: Defensible Premise or Unwarranted Compromise?" *Journal of Risk and Uncertainty*, *36*: 125–138.

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

Boadway, Robin. 1974. "The Welfare Foundations of Cost-Benefit Analysis." *Economic Journal*, *84*: 926–939.

Bressler, Daniel, and Geoffrey Heal. 2022. "Valuing Excess Deaths Caused by Climate Change." NBER Working Paper 30648.

Broome, John. 1991. *Weighing Goods: Equality, Uncertainty and Time*. Oxford: Blackwell.

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

Canning, David. 2013. "Axiomatic Foundations for Cost-Effectiveness Analysis." *Health Economics*, *22*(2013): 1405–1416.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang. 2022. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." *Quarterly Journal of Economics*, *137*: 2037–2105.

Diamond, P. A. 1967. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility: Comment." *Journal of Political Economy*, *75*(1967): 765–766.

Environmental Protection Agency (EPA). 2023. "Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances." Washington, D. C.

Fleurbaey, Marc, and James Hammitt. This issue. "The Right Numeraire or the Just Weights? How to Make BCA Rational and Fair." *Journal of Benefit-Cost Analysis*.

Gorman, Terence, (1955), "The intransitivity of certain criteria used in welfare economics." *Oxford Economic Papers*, *7*: 25–35.

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press

Hammitt, James K., and Nicolas Treich. 2007. “Statistical vs. Identified Lives in Benefit-Cost Analysis.” *Journal of Risk and Uncertainty*, *35*(2007): 45–66.

Helliwell, J. F., H. Huang, M. Norton, L. Goff, and S. Wang. 2023. “World Happiness, Trust and Social Connections in Times of Crisis.” *World Happiness Report 2023*, chapter 2.

Hicks, John R. 1939. “The Foundations of Welfare Economics.” *Economic Journal*, *49*: 697–712.

Intergovernmental Panel on Climate Change. 2023. *AR6 Synthesis Report: Climate Change 2023*. Cambridge and New York: Cambridge University Press.

Intergovernmental Panel on Climate Change Working Group 3. 1995. *Climate Change 1995: Economic and Social Dimensions of Climate Change*. Cambridge and New York: Cambridge University Press.

Intergovernmental Panel on Climate Change Working Group 3. 2001. *Climate Change 2001: Mitigation*. Cambridge and New York: Cambridge University Press.

Intergovernmental Panel on Climate Change Working Group 3. 2007. *Climate Change 2007: Mitigation of Climate Change*. Cambridge and New York: Cambridge University Press.

Intergovernmental Panel on Climate Change Working Group 3. 2014. *Climate Change 2014: Mitigation of Climate Change*. Cambridge and New York: Cambridge University Press.

Intergovernmental Panel on Climate Change, Working Group 3. 2022. *Climate Change 2022: Mitigation of Climate Change*. Cambridge and New York: Cambridge University Press.

Jones-Lee, Michael, Susan Chilton, Hugh Metcalf, and Jytte Seested Nielsen. 2015. “Valuing Gains in Life Expectancy: Clarifying Some Ambiguities.” *Journal of Risk and Uncertainty*, *51*: 1–21.

Kaldor, Nicholas. 1939. “Welfare Propositions of Economics and Inter-Personal Comparisons of Utility.” *Economic Journal*, *49*: 549–552.

National Health Service. 2020. “Delivering a Anet Zero@ National Health Service.”

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

Office of Management and Budget (OMB). 2023. “Circular A–4: Regulatory Analysis.” 17 September 2003. Washngton, D. C.

Pearce, David. 1996. “Correction on Global Warming Cost Benefit Conflict.” *Environmental Damage Valuation and Cost Benefit News*, 1 January 1996.

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.

Roser, Max, Hannah Ritchie, and Fiona Spooner. 2021. *Burden of Disease.* Published online at OurWorldInData.org. https://ourworldindata.org/burden-of-disease.

Scitovsky, Tibor de. 1941. “A Note on Welfare Propositions in Economics.” *Review of Economic Studies*, *9*: 77–88.

Somanathan, E. 2006. “Valuing Lives Equally: Distributional Weights for Welfare Analysis.” *Economics Letters*, *20*: 122–125.

Stern, Nicholas, Joseph Stiglitz, and Charlotte Taylor. 2022. “The Economics of Immense Risk, Urgent Action and Radical Change: Towards New Approaches to the Economics of Climate Change.” *Journal of Economic Methodology*, *29*: 181–216.

Sunstein, Cass R. 2004. “Valuing Life: A Plea for Disaggregation.” *Duke Law Journal*, *54*: 61.

United Nations. Undated. https://www.un.org/en/climatechange/net-zero-coalition.

United States Department of State and United States Executive Office of the President. 2021. *The Long-Term Strategy of the United States: Pathways to Net-Zero Greenhouse Emissions by 2050*. Washington, D. C.

Viscusi, Kip. 2008. “The Value of Life.” In Durlauf, S. N. and L. E. Blume (Eds.) *The New Palgrave Dictionary of Economics*, pp. 586–589. London: Palgrave Macmillan.

World Health Organization. Undated. *The Global Health Observatory: Healthy Life Expectancy (HALE) at Birth.* https://www.who.int/data/gho/data/indicators/indicator-details/GHO/gho-ghe-hale-healthy-life-expectancy-at-birth.

Cite this article: Broome, John. 2024. “The Value of Life in the Social Cost of Carbon: A Critique and a Proposal.” *Journal of Benefit-Cost Analysis*, doi:10.1017/bca.2024.21

https://doi.org/10.1017/bca.2024.21 Published online by Cambridge University Press