Environmental Economy and Policy Research

Discussion Paper Series

Ethics and Climate Change Cost-Benefit Analysis: Stern and after

by

Jonathan Aldred

2009

Number: 44.2009

Ethics and Climate Change Cost-Benefit Analysis: Stern and after

Jonathan Aldred

Abstract: The Stern Review on the economics of climate change (hereafter 'Stern') has received much attention regarding its potential political impact.[1] It has also been extensively discussed among academic economists because its conclusions are more radical, in terms of action to mitigate climate change, than previous orthodox economic analyses. This paper aims to contribute to the debate by critically assessing three key features of climate change CBA, with particular reference to Stern: monetary valuation of human life, discounting, and the treatment of risk and uncertainty. This assessment reveals some common themes, which are examined in the final section. They concern the tensions which emerge in climate change CBA between ethical arguments, market-based preferences, and the claims of CBA to scientific objectivity.

Keywords: Climate change, Stern report, Cost benefit analysis, discounting, uncertainty

Address for correspondence
Dr. Jonathan Aldred
Department of Land Economy
University of Cambridge
Tel: +44 (0)1223 362329
e-mail**:** **jsa1001@cam.ac.uk**

[1] Nicholas Stern, *The Economics of Climate Change: the Stern Review* (Cambridge University Press, 2007).

Ethics and Climate Change Cost-Benefit Analysis: Stern and after

Jonathan Aldred

1. Introduction

The Stern Review on the economics of climate change (hereafter 'Stern') has received much attention regarding its potential political impact.[2] It has also been extensively discussed among academic economists because its conclusions are more radical, in terms of action to mitigate climate change, than previous orthodox economic analyses.

But a perhaps more lasting intellectual impact will follow from the explicit recognition Stern gives to the role of ethical assumptions in cost-benefit analysis (CBA). On its own, the recognition is not large, and ethical assumptions permeate Stern much more widely than it acknowledges, but nevertheless it helps to break a powerful taboo – economic policy analysis should not mention ethics. Before Stern, it was almost unknown for orthodox economists to debate different ethical assumptions underpinning economic analysis in general and CBA in particular. There are signs that this may be beginning to change.[3]

This paper aims to contribute to the debate by critically assessing three key features of climate change CBA, with particular reference to Stern: monetary

2 Nicholas Stern, *The Economics of Climate Change: the Stern Review* (Cambridge University Press, 2007).

3 See for instance Wilfred Beckerman and Cameron Hepburn 'Ethics of the discount rate in the Stern Review on the Economics of Climate Change', *World Economics,* Vol.8, No. 1 (2007)**,** pp. 187-210.

valuation of human life, discounting, and the treatment of risk and uncertainty. This assessment reveals some common themes, which are examined in the final section. They concern the tensions which emerge in climate change CBA between ethical arguments, market-based preferences, and the claims of CBA to scientific objectivity.

2. Valuing life

CBA does not value life directly, but in terms of risk of death. The standard view of market economics is, unsurprisingly, that the money values of life, or decreased risk of death, should reflect values in markets. On this view, how much people care about an increased or decreased risk of death can be deduced by observing their behaviour in markets. In the *wage differentials* method, economists observe how much extra people must be paid to persuade them to take a dangerous job, involving some well-defined risk. If, on average, each person must be paid an extra £2000 to compensate them for undertaking a job or task involving a risk of death of 1 in 1000, then economists say that the *value of statistical life* is £2 million – £2000 multiplied by 1000. (In the UK, CBAs typically assume a valuation of life of around £1.5 million, based on wage differentials).[4] An alternative method is simply to ask people, in surveys, how much they would be willing to pay to reduce some specified risk, or how much they would require in compensation to tolerate some new risk. For example, people might be asked how much they would be willing to pay in increased tax to pay for reducing some hazardous pollutant in city air. Or how much

4 More precisely, £1.145million at 2000 prices. The figure apparently includes gross lost output, medical and ambulance costs. See *The Green Book: Appraisal and Evaluation in Central Government* (HM Treasury, 2003), Annex 2, p. 63.

money would be the minimum acceptable compensation for tolerating the siting of a hazardous waste dump nearby.

There has been much argument at the IPCC about the monetary valuation of human life. However, the argument was not about the notion of a money value for human life *per se*, but whether lives in rich countries should be assigned a higher value than those in poor countries. If the raw values from most valuation methods (such as the survey and wage differentials methods just described) are fed into the CBA as they stand, then the lives of poor people will be given a lower value than those of the rich. More specifically, the decreased risk of death, arising from policies which reduce climate change, attracts a lower monetary value in poor countries. Compared to rich countries, the extra wage earnt in dangerous jobs will be lower, and so too will the amount people are willing to pay to avoid an increased risk of death. On any unadjusted measure of their value drawn directly from market prices or surveys, poor people will be assigned a lower value. In consultation prior to the second IPCC report, most governments completely rejected this view, but a group of economists insisted that the report should include it: 'A careful reading of the fine print revealed that they were valuing lives in rich countries at $1,500,000, in middle-income countries at $300,000, and in lowest-income countries at $100,000.'[5] The final report heavily qualified this approach, but not because any consensus was reached. One participant commented: 'The outcome of it all was that the IPCC is very reluctant to engage in that controversy again because the proponents on both sides are still there and obviously still willing to have another fight if the opportunity was given

[5] Using 1995 prices. Frank Ackerman and Lisa Heinzerling, *Priceless* (New Press, 2004), p.73.

to them.'[6] In subsequent reports, the IPCC has adopted a uniform monetary value of life of $1,000,000 to apply in all countries.[7]

Before turning to whether unadjusted values of life taken directly from market data *should* be used in CBA, it is important to see that the issue is broader than the variations in the monetary value of life which arise due to income differences. *Context-specific* valuations of life arise not just because incomes vary, but because the same 'objective' risk – the same probability of harm – is perceived differently in different contexts. (Whether all situations of uncertainty can be treated as cases of risk, involving known probabilities, is a question addressed in a later section). In CBA, the extent of risk is defined solely in terms of its probability, so that two risks with equal probabilities are regarded as 'objectively identical' or 'statistically equivalent'. The probability of harm from having an X-ray in hospital is much higher than that from living near a nuclear power station. But many people regard nuclear power as riskier than hospital X-rays, because they understand risk in a much broader sense. Risk has qualities as well as quantities. A risk is perceived to be much worse if it is involuntary, unfamiliar, unfair, irreversible and uncontrollable; the risks from nuclear power are widely held to embody all these worrying qualities. These concerns are now widely discussed in the literature on the social perception of risk;[8] to the list could be added considerations surrounding the cause of the risk, such as whether it was due to incompetence, negligence or malice – or natural causes. The breadth, complexity and interrelatedness of all these factors affecting risk perception, and their generally qualitative nature, means that there is little hope of capturing them in some

6 Dr. Terry Barker, in oral evidence to the House of Lords Economic Affairs Committee on The Economics of Climate Change, 22 February 2005.

7 The issue is discussed in IPCC 2001. See B. Metz, O. Davidson, et al., (eds.), *Climate Change 2001: Mitigation* (Cambridge University Press, 2001), p. 483. The $1 million value of life is in 1999 prices.

8 See N. Pidgeon, R. Kasperson, and P. Slovic (eds.), *The Social Amplification of Risk*, (Cambridge University Press, 2003).

more sophisticated model of quantitative risk assessment. Instead cost-benefit analysts respond by making *ad hoc* adjustments to the monetary value of the risk; valuations of life become context-specific.[9] So, whether due to income differences or differing perceptions of risk, the same objective risk will be valued differently by different people, yielding differing, context-specific monetary values of life under any market-based valuation method.

When conducting climate change CBA, should lives in different countries be given different, unadjusted, context-specific monetary values, should they be adjusted in some way, or should a uniform value be used? Contrary to the impression sometimes given by government manuals on CBA, the 'measuring rod' of money is not neutral. Most obviously, money is more valuable to the poor. More precisely, an extra unit of money is worth more to the poor than the rich – the marginal value of money is higher for the poor. This follows directly from the utilitarian 'law of diminishing marginal utility', and has long been recognised in the theory underpinning CBA. But the value of an additional unit of money will differ across individuals for other reasons too – 'materialists', who especially enjoy the things money can buy, may value it relatively more highly than 'ascetics'.[10] Thus monetary measures of value emerging from market-based valuation methods will not, as they stand, be *comparable* across individuals. They should be adjusted to reflect the differing marginal values of money for different people. Suppose, for example, that money is three times more valuable to the poor than the rich. Suppose also that in a poor country, the valuation of life emerging from market-based methods is £500,000;

---

[9] A good example of how *ad hoc* is given by Jonathan Wolff, 'Risk, Fear, Blame, Shame and the Regulation of Public Safety', *Economics and Philosophy,* Vol. 22**,** (2006), p. 422. The UK Health and Safety Executive gives all lives a uniform value in its CBA work, except deaths from cancer, which are valued at exactly twice the usual figure.

[10] I am grateful to a referee for reminding me of this point. See K. Brekke, 'The numeraire matters in cost-benefit analysis.' *Journal of Public Economics* Vol. 64, (1997), pp. 117-23.

but in the rich country it is £2,000,000. To make these numbers comparable (allowing them to be aggregated in a CBA), the £500,000 figure should be multiplied by three, a 'distributional weighting'. The adjusted value of life in the poor country is now £1,500,000: we still do not have a uniform value of life, but the range is narrower.

Although appealing to many academic economists, this approach, involving adjusted context-specific valuations, is almost never attempted in practice. There are both theoretical and practical problems involved in obtaining appropriate adjustment weightings.[11] Trying to determine how much more valuable money is to 'the poor' than 'the rich' is hard enough, but this is just an approximation; a comprehensive application of this approach requires that individualized adjustment weightings be determined for each person, reflecting *all* the factors which lead the marginal value of money to vary across people.

There are also practical difficulties in obtaining the (unadjusted) context-specific valuations in the first place. The wage differentials approach is ruled out. Wage differentials are context-specific – but the context is different to the one under consideration. They concern the risk faced by a particular group of workers in a particular job, not the risk being evaluated in the CBA. This leaves the survey method, rendered more difficult and costly to conduct, because a fresh survey is required for each CBA, rather than extrapolating the results from a survey in another context, the so-called 'benefit transfer' method. As well as greatly reducing the cost of conducting a CBA, practitioners often rely on 'benefit transfer' to assist plausibility: they use valuations obtained in a more plausible context (say, willingness

[11] For up-to-date analysis of the theoretical issues, see Olof Johansson-Stenman, 'Distributional weights in cost-benefit analysis - should we forget about them?', *Land Economics* Vol. 81, No. 3 (2005), pp. 337-52.

to pay higher air fares in return for more stringent aircraft safety requirements), rather than attempting to ask directly how much people would pay for policy measures likely to reduce the impact of climate change. Although these concerns are pragmatic, they are deeply entrenched in CBA. Sugden regards the assumption of 'context-independent preferences' as playing an equally important role as standard assumptions about the internal consistency of preferences, in making it 'possible to gather data in one environment and to draw conclusions about economic value in another'.[12] More forcefully:

> The assumed transferability [of valuations] across contexts … is in fact an essential condition without which cost-benefit analyses as we know them today cannot be justified. If all values used in a CBA had to be derived from the precise context of the particular CBA, then the practice of performing CBAs would come close to that of performing opinion polls on the topic to be analysed.[13]

In short, the goal of fully adjusted, truly context-specific valuations is probably unobtainable in practice. Instead, many practitioners adopt a uniform value of life on pragmatic grounds. The idea is that a uniform value will roughly approximate the adjusted context-specific valuations: in the example above, the adjusted valuations of life are more alike than the unadjusted ones, the adjustments acting to 'cancel out' the differences in the value of life obtained from market data. Adopting a uniform value of life assumes, as a short-cut, that this cancelling out is exact.

---

[12] Robert Sugden, 'Anomalies and stated preference techniques', *Environmental and Resource Economics* Vol. 32, (2005), p.3.

[13] Sven Hansson, 'Philosophical problems in cost-benefit analysis', *Economics and Philosophy,* Vol. 23, No. 2 (2007), p. 179.

Many cost-benefit analysts object to any move away from unadjusted market-based valuations, whether to adjusted valuations, or the short-cut approach of a uniform valuation. They argue that determining policy on this basis may lead to a welfare loss. Suppose again that the unadjusted value of life in poor countries is £500,000, but that economists input the adjusted valuation of £1,500,000 into a CBA; the CBA then recommends immediate costly action to address climate change. If governments in poor countries heed the recommendation, then they may spend more on tackling climate change than their people think it is worth, as revealed by market data. That is, the poor country governments would be adopting policies which may fail an unadjusted cost-benefit test when lives are valued at £500,000. An immediate response to this argument is that the problem can be solved by ensuring poor country governments do not bear the (full) costs of their actions to tackle climate change – for instance, they could receive a subsidy from rich countries. According to this response, the adjusted CBA still identifies the correct policy choice at a global level; the problem is simply about who pays. However, defenders of unadjusted market-based valuations go further. There are likely to be many other policies which would benefit poor countries more than tackling climate change, policies which pass a cost-benefit test in the poor country even when lives are valued at £500,000. The subsidy would be much better spent on supporting (more of) these policies. Subsidising action to tackle climate change distorts the priorities of poor country governments away from other actions which their people value more highly.

These arguments rely on debateable assumptions about how welfare changes are best measured, and whether the *status quo* distribution of global income should be taken as given. However, they raise an important question: is the global perspective of much climate change CBA meaningful, in a world without global government, and

where transfers from rich countries to poor ones are negligible? These are complex issues, but they do not negate the original objection to unadjusted valuations – money is not a neutral measure of value. Therefore aggregate unadjusted valuations do not reflect how we truly value lives: if lives are held to have equal value in different countries this cannot be captured in a CBA by using different unadjusted monetary valuations of life across countries. Nevertheless, once we move away from unadjusted valuations, one of the rationales for putting monetary values on life in the first place appears to have been lost: CBA may no longer track the policy priorities of people, as expressed through market data. Abandoning unadjusted valuations may lead us to abandon the practice of valuing life in monetary terms altogether.

Summing up, it seems that we can reject unadjusted valuations. But adjusted fully context-specific valuations are unworkable in practice and the global perspective implicit in making the adjustments may be irrelevant to policy-making. The short-cut version – a uniform value of life – avoids the practical problems but confronts the question of relevance even more starkly. A uniform value of life such as the IPCC's $1,000,0000 seems abstract and arbitrary, disconnected from the preference data which is supposed to justify it. At the very least, it seems relevant only once some prior questions about global justice and responsibility have been settled. The debate over whether to use context-specific valuations, with or without adjustment weightings, remains far from settled.[14]

[14] For further discussion see Cass Sunstein, 'Are Poor People Worth Less Than Rich People? Disaggregating the Value of Statistical Lives', AEI-Brookings Joint Center for Regulatory Studies Working Paper No. 04-05 (2004). Available at: http://ssrn.com/abstract=506142 ; David Anthoff and Richard Tol, 'On International Equity Weights and National Decision Making on Climate Change' FEEM Working Paper, Milan (2007). Available at: http://www.feem.it/Feem/Pub/Publications/WPapers/default.htm

There is yet another way in which context-specific valuations arise, one which reveals a deeper problem with the very idea of the 'value of statistical life'. Suppose that Ann values a reduction in some particular, well-defined risk of her death of 1 in 1000 at £2000. That is, Ann requires £2000 compensation to undertake the risk. Does this imply that Ann would accept £200000 in return for facing the same type of risk, but now with a 1 in 10 probability of death? Of course not. The relationship between probability and value is not linear. Even if we consider the same type of risk, faced by the same person, it does not follow that valuations are generalisable to new probabilistic contexts in the way that the CBA literature usually assumes in invoking a linearity assumption. It is particularly clear that linearity fails in the extreme case: the amount of money that someone would accept in order to face certain death cannot be calculated by scaling up valuations for very small risks of death. The idea is absurd, and seems to have no relevance to the practice of CBA: defenders of CBA emphasise that it is concerned with 'statistical' lives lost, not real ones. But policies evaluated using CBA often involve real lives being lost, predictably so. Even if the policy imposes only a small increased risk of death on each individual, real lives are highly likely to be lost whenever very large numbers of people are affected, or the risks are repeatedly imposed over a long period. In these cases, by the 'law of large numbers' from probability theory, we can be confident about roughly how many will die. Climate change provides many examples – relatively small increases in the risk of death faced by millions of people.

No argument is being made here against policies which lead to lives being lost. Sometimes such policy choices are correct, because the alternatives are worse. The point is that the CBA treatment of these policies is exactly equivalent to putting a money value on real lives directly, without the 'statistical' smokescreen. For example,

suppose that due to climate change, 60 million people (the UK population) each face a one in one million increase in the risk of death. Each person's increased risk is valued at £2 by the CBA, so the total risk would be valued at £120million. We can confidently predict in advance that on average 60 people will die. So it seems more transparent, and less misleading, to state simply: on average 60 people will die, and the CBA values their lives at £2million each.

The argument here is about more than terminology. First, even supporters of CBA are concerned about these kinds of cases. Whether obtained through surveys or revealed in the market through methods such as wage differentials analysis, individual valuations for risks as low as one in one million are likely to be unreliable, due to well known cognitive biases in interpreting small probabilities. Valuations would probably be very low, sometimes zero, implying that CBA might assign a very low total value to the risk, allowing the policy to pass the cost-benefit test, even though on average 60 people will die. This is the average impact; depending on how the risk manifests itself, there could be a catastrophe in which many more would die. Although no worse for *each* of those who die, the overall impact on society of this catastrophe would probably be disproportionately greater. [15]

Second, valuations of small risks should not be used to inform climate change policy when real lives are predictably lost, because linearity is false. The problems here become particularly stark when a third issue, identity, is raised. Consider a hypothetical policy which would lead to the death of an *identifiable* person, John. Clearly the money value which John would place on his own life would be infinite,

15 Richard Posner, *Catastrophe* (Oxford University Press, 2004), p. 165-70; Cass Sunstein, *Laws of Fear* (Cambridge University Press, 2005), p. 160-1. Neither Posner nor Sunstein seem to realise that this example reveals a general flaw with valuing life in terms of risk, not just one applying to small probabilities of catastrophe.

consistent with the theory behind CBA. There is no amount of money which John would be willing to accept as compensation for his death. So because of John's death alone, the costs of this policy are infinite, and it would fail the cost-benefit test. Contrast this case with the policy just described, which might pass a cost-benefit test, even though 60 people die on average. This is possible only because the identity of those who will die is unknown, so there is no particular individual like John. It is absurd that ignorance should make a difference here. What matters is the numbers who will die, not whether we know their identity. The 'value of a statistical life' is a sleight-of-hand device to conceal this absurdity. Presenting the policy as 'causing the death of 60 people on average, but we do not know who' may make it much harder for cost-benefit analysts to persuade their audience that a money value could be put on the risk. Instead, the risk is represented as a loss of statistical lives, or increased probability of death. But it is still the same loss, and we should not be deceived into believing it can be valued in terms of money.[16]

3. Discounting

Discounting concerns the idea that future costs and benefits should have less weight in our decision-making than present ones. It is difficult to overstate the importance of discounting in the economic analysis of climate change. Almost all of the debate among orthodox economists following the Stern Review has concerned its treatment of discounting. It is easy to see why: the results of climate change CBA are acutely sensitive to the discount rate chosen. Stern's effective discount rate is 1.4 per cent, while many economist critics have advocated a rate around 6 per cent. 'The

16 The identity argument first appeared in John Broome, 'Trying to value a life', *Journal of Public Economics* Vol. 9 (1977), pp. 91-100.

present discounted value of a given global-warming loss from a century hence … [at 6 per cent] is *one hundredth* of the present discounted value of the same loss at Stern's annual interest rate of 1.4%.'[17] And yet, for many non-economists, discounting seems obviously inapplicable to climate change policy, because it has an extremely implausible implication: huge impacts in the distant future count for almost nothing in our current decision-making. 'Suppose that Denmark needs to be evacuated due to flooding. Current real estate value in Denmark is estimated at about $238 billion. If a discount rate of 5% is applied, then over 500 years, the same real estate would be worth just $6.'[18] And here is another compounding effect: since human lives are valued in CBA, and CBA discounts all impacts, then future lives are discounted. Discounting at six per cent, one life now is worth 339 lives in 100 years time.

If these examples lead us to adopt a uniform discount rate of zero instead, not only is discounting redundant, but formal climate change CBA too, because the result of a CBA is apparent without any formal analysis. Simply because the costs of adverse climate change are felt by so many generations, extending into the far future, they will substantially outweigh the costs of mitigation borne now. Most economists go further: no discounting implies we should reduce our income to a subsistence level now for the sake of future generations. Any sacrifice now will be outweighed by the benefit to future generations, because investments made now will grow and yield greater benefit in the future, and there are many future generations who will benefit. Economists regard this excessive sacrifice argument as a *reductio ad absurdum* against zero discounting, although there are convincing reasons to doubt the excessive

---

17 Martin Weitzman, 'A Review of *The Stern Review on the Economics of Climate Change*', *Journal of Economic Literature,* Vol. 45, No. 3 (2007), p. 708.

18 Stephen Gardiner, 'Ethics and global climate change', *Ethics,* Vol. 114 (2004), p. 572.

sacrifice argument.[19] But if CBA applied in cases involving large very long-term impacts generates absurd recommendations *both* with zero discount rates and positive ones, then it is tempting to conclude that discounting methods are irrevocably flawed, at least in these cases, and should be abandoned. However, since economists are only convinced of the absurdity in the case of zero discounting, they continue with the practice. Against this background, and given what is at stake, it is worth exploring the core arguments about discounting in some detail.

The current debate over discounting largely concerns two arguments for it:

1. Pure time preference. This is the claim that, all other things being equal, people prefer welfare or satisfaction now to the same amount of it later. The claim is usually justified by reference to individual impatience.

2. Future consumption is worth less. Economic growth means that consumption levels will be higher in the future, so by the law of diminishing marginal utility, the satisfaction from additional future consumption will be lower.

In comparison to orthodox economic theory or past practice in CBA, Stern significantly limited the scope and magnitude of discounting, and attracted substantial flak from many economists for doing so. This occurred despite Stern devoting much more attention to economists' arguments that it discounts too little, than the opposite view of philosophers and others, that it still discounts too much.[20] On ethical grounds,

19 For critique of the excessive sacrifice argument see John Broome, *Counting the Cost of Global Warming* (White Horse Press, 1992); Tyler Cowan and Derek Parfit, 'Against the Social Discount Rate', in Peter Laslett and James Fishkin (eds) *Philsophy, Politics and Society: Future Generations* (Yale University Press, 1992).

20 See especially Stern, *The Economics of Climate Change*, pp. 35-7, 50-4, 59-60.

Stern rightly rejects the argument from pure time preference altogether.[21] That is, Stern is impartial between generations in the sense of attaching equal weight to welfare impacts, whether they are experienced now or by future generations. Stern defends the second argument, although emphasises that it does not apply uniformly across goods and people. It is likely that some future people will be worse off, because increased inequality will outweigh overall economic growth. Also, while some goods will be more abundant in the future, others will be equally scarce or more so, because economic growth cannot generate more of them. Stocks of non-renewable resources such as fossil fuels will decline; the same is arguably true of unique, irreplaceable environmental assets such as beautiful countryside. The upshot is that the discount rate for these goods, or for impacts affecting groups who are worse off in the future, should be zero or negative as appropriate.

Stern also responds to a particular objection to discounting when future generations are affected. It is not a new objection, but until climate change issues became prominent, most cost-benefit analyses assumed that all the major impacts of CBA-based decisions are felt within thirty years or so – that is, by the current generation.[22] Climate change means the objection can no longer be ignored. The problem is that future generations are disenfranchised.[23] In CBA, the value assigned to the discount rate is influenced by market interest rates. (The influence varies according to differing views about how discount rates should be derived, as discussed

---

21 Stern rejects the argument for pure time preference based on impatience, but argues that there is a non-trivial probability that the human race may not recognisably survive into the distant future, because of catastrophes like a meteorite impact or devastating nuclear war. This leads Stern to assume a small degree of *effective* pure time preference: pure time preference equivalent to the human race facing an almost 1 in 10 chance of extinction within 100 years, which seems implausibly large if this is really the only justification. See Stern, *The Economics of Climate Change*, pp. 35-6, 53.

22 For instance, guidance for the conduct of CBAs in previous editions of the UK Treasury *Green Book* recommended time horizons of around 30 years.

23 This version of the argument that future generations are disenfranchised draws heavily on John Broome, *Ethics out of Economics* (Cambridge University Press, 1999), Chapter 4.

in the next section.) At least in principle, market-determined rates reflect preferences – specifically, how impatient we are. But only the preferences of current generations are reflected in markets; the preferences of future generations are represented only insofar as current generations care about them. Future generations are disenfranchised just in the way that women are disenfranchised in societies where only the views of men have any impact on the decision. This may seem like the familiar point about intergenerational justice requiring impartiality between generations. But the present argument is different: even if such ethical considerations turn out to sanction the practice of discounting in principle, in practice we should not let current interest rates influence calculation of the discount rate, because only current generations have any sway over these interest rates. Obviously, it is inevitable that only current generations make decisions – in *that* sense future generations are unavoidably disenfranchised – but it is less obvious, and not inevitable, that this bias is reinforced by relying on market interest rates. Clearly the argument here goes beyond interest rates. An identical objection can be made to using observed saving rates to determine the rate of pure time preference (even if a positive rate of pure time preference is otherwise regarded as legitimate).

In response to this argument, rather than abandon discounting, economists have characteristically embraced a technical modification with great enthusiasm. This new work in economic theory is seen as answering the sceptics, and has swiftly turned into UK Government policy, now reflected in Stern.[24] The proposal is that discount rates should decline over time, rather than remain constant as in standard CBA. With

---

24 *The Green Book*, 2003. The UK Office of the Deputy Prime Minister commissioned a report on recent research on discounting, which explicitly recommended declining discount rates as a response to the philosophical objections: *A social time preference rate for use in long-term discounting* (Office of the Deputy Prime Minister, 2002), p. 13. The Stern Review shares this conclusion, although for different reasons. See Stern, *The Economics of Climate Change*, pp. 35-7, 56-7.

declining rates, the distant future counts for *much* more than under conventional discounting, so that impacts affecting future generations are discounted less. But this more favourable treatment of future generations is a by-product of declining discount rates, not their rationale, which concerns uncertainty. The rationale runs as follows. Until very recently, if they did not know what discount rate to pick, economists simply adopted an average (possibly weighted) of the plausible rates, constant over time. It is now known that this strategy reflects a mathematical error. The costs and benefits are actually multiplied by discount *factors*, so it is these factors which should be averaged in response to uncertainty, not the rates themselves; averaged discount factors imply, counter-intuitively, a much lower effective discount rate, and one which declines over time.[25]

How does this new approach claim to answer the ethical objections to discounting? Ethical doubts about the discount rate are interpreted as a form of uncertainty. For example, a 'materialist' might favour a high discount rate, while a 'conservationist' would favour a low one.[26] Then we should adopt declining discount rates over time in order to 'average' the weight given to the two ethical positions. While sceptics about discounting may welcome the belated recognition given to competing ethical views, the overall approach is a bizarre one. Trying to 'average' the implications of 'materialist' and 'conservationist' views is akin to the person who, because they are unsure whether to become a Christian or a Buddhist, decides to be a

---

25 The discount factor in period t, δ_t, is related to the discount rate r using the standard net present value formula $\delta_t = 1/(1+r)^t$. A pioneering theoretical contribution on declining discount rates is Martin Weitzman, 'Why the far distant future should be discounted at its lowest possible rate', *Journal of Environmental Economics and Management,* Vol. 36 (1998), p. 201-8. For a simple example illustrating the mathematical relationships which give rise to declining discount rates, see D. Pearce, B. Groom, C . Hepburn and P. Koundouri, 'Valuing the Future', *World Economics,* Vol. 4, No. 2 (2003), pp. 128-9.

26 C. Li and K. Lofgren, 'Renewable resources and economic sustainability', *Journal of Environmental Economics and Management*, Vol. 40 (2000), pp. 236-50.

Buddhist on Mondays, Wednesdays and Fridays, and a Christian on Tuesdays, Thursdays and Sundays.

To sum up, CBA aims to provide guidance on what we should do, so our obligations to future generations matter. Discounting disenfranchises future generations, and the technical fix of declining discount rates does not address the problem. Debates over discounting bring ethical issues to the fore, but discounting may not offer the appropriate framework for the resolution of difficult ethical dilemmas raised by problems of intergenerational justice. At best, discount rates express the resolution in simplified form; they do not help to reach it.

4. Risk and uncertainty

It is worth stressing the challenge which uncertainty poses to climate change policy-making, because some suggest that the uncertainty is no greater than in many other policy problems. Climate change uncertainty is driven by distinctive features. There are complex feedback effects which cut back and forth between the social and natural worlds. For example the scale and rate of climate change will affect economic growth and the speed of technological innovation, both of which in turn influence emissions levels, feeding back to shape climate change. Uncertainty about the causes of climate change is not just *ex ante* but *ex post*. That is, there is no event which resolves the uncertainty. There may be many more hurricanes than previously, but there is no way of knowing whether any particular hurricane is due to climate change. Uncertainty is only partly a result of incomplete knowledge or lack of research. Rather, research has shown that uncertainty is intrinsic to the problem. One reason is that increased uncertainty about the climate is itself a *consequence* of climate change,

not just an exogenous constraint. Another is that climate change manifests itself in a highly non-linear, chaotic system which is inherently unpredictable. For example, a characteristic of chaotic models is their extreme sensitivity to parameter values – popularly illustrated by the 'butterfly effect'. In sum, it is hard to see what other policy issue shares all these features of uncertainty.[27]

Stern represents an improvement on past practice in CBA in terms of its treatment of uncertainty. The Keynesian/Knightian distinction between risk and uncertainty is explicitly acknowledged: under risk, the full range of possible outcomes is known, along with their probabilities; under uncertainty, either the range is unknown, or the probabilities, or both. (This meaning of uncertainty will be termed *pure uncertainty* hereafter). As already indicated, in climate change analysis, pure uncertainty does not arise only when there is an absolute blank in our knowledge, but as an intrinsic phenomenon, or because scientists fundamentally disagree about probabilities or impacts.

Many orthodox economists endorse some form of Bayesianism. That is, there is no such thing as pure uncertainty in the sense I have defined it. Decision-makers always use probabilities, consciously or otherwise, when outcomes are not certain. Most orthodox economists who reject these claims nevertheless adhere to a weaker, normative version of Bayesianism: *rational* choice requires that we must attach probabilities, explicitly or implicitly – so it is better to be explicit about it. Although distinct, Stern's 'subjective preference' approach remains close to this view, despite

[27] See Lisa Heinzerling, and Frank Ackerman, 'Law and Economics for a warming world', *Harvard Law and Policy Review*, Vol. 1, No. 2**,** (2007), pp. 331-62.

its acceptance of the possibility of pure uncertainty.[28] Both of the techniques which Stern uses in response to this possibility involve quantifying the uncertainties.

First, Stern uses an econometric technique ('Monte Carlo simulation') to show how the results of his core model are highly uncertain: they vary dramatically depending on the assumed parameter values. However, for each of the 31 parameters in the model, Stern still assumes we know both the exact range of possible values, and the probability distribution of the parameter.[29] Second, in cases of pure uncertainty, Stern does not demand probabilities for all the unknowns, but *does* assume we know the exact range of possible probabilities, *and* can put numerical weights on these possible probabilities according to factors such as 'the magnitude of associated threats, or pessimism, and possibly any hunch about which probability might be more or less plausible.'[30] The result of these assumptions is a kind of quantitative pseudo-probability attached to each uncertainty. Stern's approach here suffers from the same defects as Bayesianism applied to CBA. The numerical weights to be attached to probabilities must be invented, forcing inherently qualitative views about catastrophic risks and the like into an unhelpful quantitative framework. With both techniques, the outcome of the CBA will have a misleadingly reassuring precision, implying that we know more than we do.

Stern also claims that the numerical weights which form part of the quantitative pseudo-probabilities enable the precautionary principle to be incorporated into CBA. The idea appears to be that if some possible outcome is particularly catastrophic, then a greater weight can be attached to the probabilities associated with

---

28 Stern, *The Economics of Climate Change*, pp. 37-9.
29 *Ibid.*, p. 173.
30 *Ibid.*, p. 39.

it. This attempt to give catastrophe risks special treatment in the analysis is welcome, and it is accompanied by various 'special mentions' throughout the Review. Nevertheless, it is unconvincing from several perspectives. To begin with, Stern's approach advises the decision-maker to 'act as if she chooses the action that maximises a weighted average of the worst expected utility and the best expected utility.'[31] This does not capture the underlying intuition behind the precautionary principle, that when the worst outcome could be severe, we should take all reasonable steps to avoid it. No version of the precautionary principle involves trading off the best and worst outcomes against each other. We return to the precautionary principle shortly.

Support for the view that Stern understates the severity of uncertainty also comes from a very different source. Harvard economist Martin Weitzman is a leading figure in orthodox economics who has made innovative theoretical contributions, such as his work on declining discount rates outlined above. Weitzman begins by noting that IPCC 2007 considers six 'equally sound' scenarios for global average temperature increase. There is pure uncertainty, so IPCC does not attach probabilities, but *if* these 'equally sound' scenarios were 'equally likely' then it follows that there is a probability of around 3 per cent of global warming of 6ºC or more within 100 years.[32] Since uncertainty means such non-trivial probabilities cannot be ruled out, and given the horrific effects of 6ºC warming, Weitzman argues that catastrophes should be central to any analysis. Consistent with his view that CBA parameters should be derived from market data, Weitzman looks to behaviour in markets for clues on how to treat catastrophe risks. Specifically, he argues that

---

31 *Ibid.*, p. 39.

32 Weitzman, 'A Review of *The Stern Review*', p. 716.

various puzzling patterns of behaviour in financial markets – so-called 'asset return puzzles' – only make sense if 'investors are disproportionately afraid of rare disasters'[33]; and he adds a theoretical econometric argument suggesting that this disproportionate fear is legitimate in a world involving pure uncertainty. Weitzman concludes that the damages in case of catastrophe will generally be decisive in determining the recommendation of any kind of CBA, largely regardless of its other parameters such as the discount rate. Applying this framework to climate change, two important points emerge. First, even economists persuaded of the wisdom of attempting global CBA of *average* impacts of climate change must recognise that impacts in a world with 6ºC warming 'are located in the *terra incognita* of what any honest economic modeller would have to admit is a planet Earth reconfigured as science fiction.'[34] It is simply absurd to attempt to measure these impacts in monetary terms. Second, no such attempt is necessary, even from a narrow CBA perspective, because there are good reasons in economic theory to justify very substantial costs to eliminate the risk of uncertain disasters. Weitzman summarizes these arguments with the following contrast. Standard climate change CBA is about 'consumption smoothing' – trading-off sacrifices in consumption now for the sake of benefits in the future. Instead we should think in terms of 'catastrophe insurance'.

To non-economists, and many economists, Weitzman's analysis may seem idiosyncratic. But it may prove crucial in shifting the terms of debate in orthodox economics. At the very least, it makes one fundamental claim. Economic analysis of climate change is preoccupied with average outcomes: attempting to assess the average effects, if the Earth experiences the average (i.e. most probable) amount of

33 *Ibid.*, p. 715.
34 *Ibid.*, p. 716.

climate change. But if averages cannot be calculated, then this analysis is misleading and futile. Instead we must consider *all* possible outcomes, rightly and inevitably paying the most attention to possible catastrophes.

These considerations return us to the precautionary principle. There has been a voluminous literature on the PP, much of it critical. But some of the harshest critics have recently been persuaded that a more coherent PP can be recovered, one which escapes their objections.[35] This 'core PP' has a much narrower range of application, inspired by Rawls' criteria to determine the applicability of the difference principle.[36] The precautionary principle is analogous to Rawlsian 'maximin' because in both approaches, priority is given to the worst off/case rather than trading off this priority against others. Under the core PP, precaution is justified when three conditions are met: (i) some plausible possible outcomes are catastrophic; (ii) probabilities cannot be attached to these outcomes; (iii) the loss from following maximin is relatively small. It is clear that climate change meets the first two conditions; regarding the third, much hangs on the word *relatively*. Essentially, precaution is justified if the costs of aggressive action to stabilize carbon emissions are small relative to the size of the potential catastrophe. Many would argue that this is true, if the costs can be understood as around 2 per cent of global GDP. But of course such cost estimates are still vulnerable to the criticisms of climate change CBA made here and elsewhere. Plausible versions of the precautionary principle do not, on their own, allow us to escape from standard cost-benefit thinking. On the other hand, if we make the ethical judgement that the effects of a climate change catastrophe would be qualitatively

35 Cass Sunstein is one prominent critic who now supports a limited precautionary principle. See Sunstein, *Laws of Fear*, chapter 5.

36 See especially Stephen Gardiner, 'A Core Precautionary Principle', *Journal of Political Philosophy,* Vol. 14, No. 1, (2006), pp. 33-60.

different from, and worse than, the consumption losses involved in taking action now – immeasurably worse – then aggressive action can be justified. But as the next section argues, most orthodox economists are deeply reluctant to admit explicit ethical judgements, despite their ubiquity in CBA.[37]

5. Revealed preference, ethics and cost-benefit science

In light of the problems discussed in the previous three sections, CBA appears to be a seriously flawed tool for guiding decisions about climate change policy. There is a large and vitally important question about the extent to which governments should continue to rely on CBA for these purposes. But given its enormity, I will approach just one aspect of this question indirectly, examining the way in which arguments about climate change CBA are conducted.

'Revealed preference' is the conceptual ghost which haunts all aspects of the analysis. Individual choices in markets supposedly reveal people's preferences, which are then used by economists to deduce important parameter values in their models. These parameters can have a decisive influence on the recommendations of CBA. In the above discussion of the Stern Review and related work, there are many examples of the attempt to derive parameter values from observed market behaviour, including deriving the discount rate from observed saving rates and interest rates, and determining monetary valuations, such as the value of life, using revealed preference methods such as the wage differentials approach. The practice of inferring

[37] For further discussion of the reluctance of economists to acknowledge ethics, see Daniel Hausman and Michael McPherson, *Economic Analysis, Moral Philosophy and Public Policy*, (Cambridge University Press, 2006); Philippe Mongin, 'Value Judgments and Value Neutrality in Economics', *Economica*, Vol. 73 (2006) pp. 257-86.

probabilities in Bayesian fashion from market choices is another example of revealed preference reasoning. Although discussion in economics often proceeds as though there is no alternative to revealed preference approaches, debates about discounting show more awareness of the issues.

At least since influential work by Arrow and Stiglitz for IPCC1995, there has been explicit recognition of 'descriptive' versus 'normative' approaches to discounting.[38] The descriptive approach seeks to derive discount rates from market data; the normative one appeals to ethical principles to determine the rate. Arrow emphasises the merits of the normative approach, but nevertheless takes what might be termed a 'weak descriptive' view: even if market data alone cannot determine the discount rate, our ethical choices must be 'consistent' with observed market behaviour. Perhaps the most notable feature of the debate following Stern, and surrounding climate change CBA more generally, is the conflict over whether to take a descriptive or normative approach to discounting. The debate has been characterised, by those on both sides, as reflecting a fundamental difference of view between UK and other economists.[39] The language of the debate is sometimes intemperate: Nordhaus lambasts Stern for adopting 'the lofty vantage point of the world social planner, perhaps stoking the dying embers of the British empire'.[40] On the UK side, Dasgupta has endorsed Stern's approach, which in its willingness to reject market data on ethical grounds, can be traced back to Ramsey, via Sen.[41]

---

38 K. Arrow, W. Cline, K.-G. Maler, M. Munasinghe, and J. Stiglitz, 'Intertemporal equity, discounting, and economic efficiency', in J. Bruce, H. Lee and E. Haites (eds) *Climate Change 1995: economic and social dimensions of climate change*, (Cambridge University Press, 1996), p. 125.

39 See for instance Weitzman, 'A Review of *The Stern Review*', p. 703-24; Angus Deaton, 'On transatlantic vices, or Stern in America', *Royal Economic Society Newsletter,* Vol. 139 (2007), p. 3-4.

40 William Nordhaus, 'A Review of *The Stern Review on the Economics of Climate Change*', *Journal of Economic Literature,* Vol. 45, No. 3, (2007), p. 691.

41 Partha Dasgupta, 'The Stern Review's Economics of Climate Change', *National Institute Economic Review,* Vol. 199, (2007) p. 4-7.

Elsewhere, the majority revealed preference view among economists was adopted by Lomborg's 'Copenhagen Consensus' panel, and endorsed again in the aftermath of Stern by Arrow, Nordhaus, and Weitzman, among others.[42]

Despite their persisting popularity among economists, revealed preference arguments are fundamentally flawed; the 'weak descriptive' view of Arrow hardly fares better. To begin with, choice cannot reveal preference because choice depends on both belief and preference.[43] Only with comprehensive information about beliefs can we hope to infer a preference from an observed choice. For example, wage differentials cannot alone form the basis for valuing life. The observing economist would also need to know how workers think about risks. Specifically, the economists would need to know, or assume, that workers consciously weigh up risks against wages, have accurate probability information about the risks they face, construct their beliefs about risk on the basis of this probability information (rather than 'trusting to fate'), and do not suffer from common cognitive biases in their apprehension of probabilities. Similarly, observed saving rates reflect habit, inertia, trends in house prices, and beliefs about future inflation and interest rates, as much or more than they reveal pure time preferences between present and future consumption. These examples raise a second worry: even if the observing economist somehow knows enough about the individual's beliefs to derive her preferences from her choices, the standard 'consumer sovereignty' justification for relying on these preferences breaks down if the preferences are based on false beliefs, incomplete information or cognitive errors. That is, policy based on a correct reading of preferences from

---

42 Kenneth Arrow, 'Global Climate Change: A Challenge to Policy', *The Economists' Voice*, (2007); Lomborg, *Global Crises, Global Solutions*; Nordhaus, 'A Review of *The Stern Review*'.

43 For seminal contributions in economics and philosophy see Amartya Sen, 'Behaviour and the Concept of Preference', *Economica,* Vol. 40, (1973), pp. 241-59; Donald Davidson, 'Actions, reasons and causes', *Journal of Philosophy,* Vol. 60, (1963), pp. 685-700.

choices may still fail to be welfare-maximising. More generally, recent work in behavioural economics provides overwhelming evidence that individuals are not good judges of their own interests, so welfarist appeals to consumer sovereignty are spurious. A third problem is the context-specificity of preferences. We have seen this in the case of valuing life but the problem is more general. Preferences revealed through market data may be entirely different from those the economist is actually seeking to uncover – the preferences in the non-market context of climate change policy.

Regardless of the status of revealed preference arguments, it does not follow that the preferences so revealed should guide policy. Individual preferences expressed in markets reflect the narrower concerns of individuals as consumers, whereas climate change policy demands the preferences – or rather, judgements – of individuals as citizens who consider what society should do. For example, while individuals may exhibit impatience (pure time preference) in their behaviour as consumers in markets, preferring consumption now to the same amount later, as citizens they may judge pure time preference to be wholly inappropriate to societal decision-making. So even if we decide that discount rates should be determined by aggregating the judgements of members of society, these judgements cannot be derived by reading off preferences from market behaviour. Climate change policy also raises particular ethical considerations which take us beyond market data: this data can only reflect the preferences of current generations, so future generations are disenfranchised. This point was noted above with regard to using market interest rates and savings rates to guide the setting of discount rates, but again, the objection to market data is more general.

In light of all these objections to using market data to determine valuations and model parameters, it is worth asking why most orthodox economists still do so. An important part of the explanation appears to be a widespread determination to deny the role of ethics, in favour of economics as 'science'. Nordhaus finds 'ethical reasoning on discount rates [in Stern] largely irrelevant'.[44] In an acerbic passage, Weitzman is equally contemptuous, stating that Stern (and Cline, a precursor also adopting a normative approach) rely:

> mostly on a priori philosopher-king ethical judgements about the immorality of treating future generations differently from the current generation—instead of trying to back out what possibly more representative members of society than either Cline or Stern might be revealing from their behavior is *their* implicit rate of pure time preference. An enormously important part of the 'discipline' of economics is supposed to be that economists understand the difference between their own personal preferences for apples over oranges and the preferences of others for apples over oranges. Inferring society's revealed preference value of [rate of pure time preference] is not an easy task … but at least a good-faith effort at such an inference might have gone some way towards convincing the public that the economists doing the studies are not drawing conclusions primarily from imposing their own value judgements on the rest of the world.[45]

These comments are all the more remarkable because Weitzman appears to take the more nuanced 'weak descriptive' approach of Arrow; Weitzman is certainly aware that actual market behaviour often appears inconsistent with economic theory (for instance in his 'asset return puzzles'), so that revealed preference methods cannot be applied straightforwardly. In short, Weitzman appears much more concerned to exclude ethics than to insist on using market data alone to derive discount rates.

---

44 Nordhaus, 'A Review of *The Stern Review*', p. 692.
45 Weitzman, 'A Review of *The Stern Review*', p. 712.

Among economists who favour a normative approach, there is still unease about acknowledging the role of ethics. Although Dasgupta has written on some philosophical issues around discounting, he appears to regard such discussion as inappropriate within economics: he has devoted an entire paper to arguing that economists argue about facts, not values.[46]

More generally, both Nordhaus and Weitzman dismiss the Stern Review as being much more a 'political document' than a scientific analysis, making repeated references to economics being a science.[47] We have seen the scientific pretensions of climate change CBA elsewhere, in the focus on quantification and consistency. Quantification goes far beyond the realm of the market, making progressively more heroic assumptions to extract numbers, as in the treatment of pure uncertainty and the monetary valuation of life. The practice of valuing life is also dominated by the drive for consistency, as is the practice of replacing subjective judgements about risk with objective probabilities whenever those are available.

At this point, there is no need to confront difficult questions about the nature of science, because the project to present CBA as an ethics-free science is *internally* flawed. Firstly, CBA clearly rests on a comprehensive ethical framework, perhaps best labelled 'welfarism'. This may appear obvious, but it is notable that Nordhaus, in appealing to alternative ethical systems to support his belief in pure time preference, does not acknowledge that such systems would typically abandon CBA altogether as a

---

46 Partha Dasgupta, 'Three conceptions of intergenerational justice', in H. Lillehammer and D.H. Mellor (eds), *Ramsey's Legacy* (Oxford University Press, 2005); Partha Dasgupta, 'What do economists analyze and why: values or facts?' *Economics and Philosophy,* Vol. 21, No. 2 (2005)**,** pp. 221-78.

47 See for instance Nordhaus, 'A Review of *The Stern Review*', p. 688; Weitzman, 'A Review of *The Stern Review*', p. 723.

guide to social decision-making.[48] It is as though CBA *per se* (as opposed to normative approaches to discounting) can be conceived in ethics-free terms, without appeal to its welfarist foundations. Recall that many economists still believe CBA can be justified in terms of Pareto efficiency alone, and 'efficiency' is sometimes seen as a matter of fact, not value.[49]

This raises a second, more fundamental difficulty. There is little recognition in orthodox economics that even if first order ethical judgements are sidestepped, second order judgements are made about the institutional setting for those first order judgements. Effectively, a kind of meta-ethical judgement replaces the first order ethical judgement. For example, the decision to determine discount rates and other parameters on the basis of market data is in this sense a meta-ethical judgement to place the ethical choice in the hands of the market. Some economists appear barely to notice making *this* judgement because they do not acknowledge there is a choice to be made: there is no alternative. That is, there is almost no recognition that legitimate policy-making can proceed on a different basis to private decision-making, that factors relevant to one may not be relevant to the other, so that public and private decision-making can be inconsistent in the sense of 'revealing' different implicit parameters (such as different discount rates) without either decision process being flawed.

Third, at the level of practice, CBA is unavoidably riddled with *ad hoc* ethical assumptions if monetary values are to be assigned. This is less widely acknowledged than welfarism: recognition of the role of ethics at this level does not usually extend beyond the choice of distributional weights and the discount rate. It is not just the key

---

48 Nordhaus, 'A Review of *The Stern Review*', p. 692.

49 For careful exposition of the flaws in this view, see Hausman and McPherson, *Economic Analysis*.

parameters which depend on ethical assumptions, but the minor assumptions too, which are acknowledged, if at all, in a footnote. Two examples must suffice here. Stern makes modelling assumptions which rule out 'consumption externalities' such as envy.[50] 'Envy' may appear undesirable, but Stern's assumption rules out the possibility that the satisfaction from material consumption may depend on the consumption levels of others, a possibility backed by substantial psychological evidence. In providing guidance for CBAs of transport projects, the U.K. Treasury *Green Book* assumes that time spent sitting in a traffic jam is less unpleasant (i.e. less costly) than time spent walking or waiting at a bus stop.[51] While these kinds of hidden assumptions may not have much impact on the overall analysis individually, the cumulative effect may be significant. This is more likely if the assumptions, although apparently independent, are informed by a common normative perspective, perhaps concerning the preferences people ought to have, or the belief that the market is usually the best institution for resource allocation.

Fourth, another reason for the ubiquity of ethical assumptions is that economists sometimes reject market data even when it is available and unambiguous. As already noted, there is a tension between economists' loyalty to market data – consumer sovereignty – and the desire for scientific consistency and objectivity. A related conflict arises between market data and economic theories of rationality, which have a normative element. If the market data reveals 'irrational' preferences, then some economists side with the raw preferences: pure time preferences and

50 Stern, *The Economics of Climate Change*, p. 56.

51 *The Green Book*, 2003, Annex 2.

preferences about subjective risks are treated like any other preferences.[52] Others favour the theoretical account of rationality, and accordingly abandon, modify or otherwise launder the data so that it more closely reflects the preferences market participants would have, if they were 'fully rational'.

Once the full range of value judgements is acknowledged, then climate change CBA becomes recognised as a quantitative framework for structuring ethical decision-making – applied moral mathematics rather than science. Many facts in CBA are inextricably entangled with values. These facts can only be stated in terms which are partly evaluative: description and evaluation cannot be separated.[53]

It is worth stressing what is *not* being argued here. There is no suggestion that facts and values are indistinguishable. Or even that cost-benefit analysts cannot make value-free statements (although they may be hard to identify in practice).[54] Although ethics cannot be stripped out of CBA, it does not follow that CBA is redundant. Rational argument about ethical judgements – whether within the framework of CBA or not – is not only possible, but central to policy-making.

6. Concluding remarks

What role, if any, should CBA play in informing climate change policy? One reply is that applications of moral mathematics such as CBA have a distinctive role to play in clarifying the nature of ethical disagreements, facilitating the kind of rational

---

52 An influential defence of this view is John Harsanyi, 'Morality and the theory of rational behaviour', in Amartya Sen and Bernard Williams (eds) *Utilitarianism and Beyond* (Cambridge University Press, 1982).

53 Hilary Putnam, *The Collapse of the Fact/Value Dichotomy,* (Harvard University Press, 2002).

54 Mongin, 'Value Judgments'.

argument just mentioned. Another response points to a major objection to CBA which has not been mentioned so far – incommensurability, the view that expressing all climate change impacts in monetary terms is at best unhelpful, usually misleading, and in some cases impossible. Taking incommensurability seriously has powerful implications. Returning to an earlier example, the incommensurability judgement there – that the impacts of a climate change catastrophe are incommensurably worse than the consumption losses involved in taking action now – arguably did more work than the precautionary principle in justifying immediate aggressive action.[55] For many non-economists, the incommensurability objection to CBA is overwhelming in the case of climate change. Stern sometimes appears to agree:

> Our preference is to consider the multiple dimensions of the cost of climate change separately, examining each on its own terms. A toll in terms of lives lost gains little in eloquence when it is converted into dollars; but it loses something, from an ethical perspective, by distancing us from the human cost of climate change.[56]

Nevertheless, Stern does 'convert into dollars', in presenting a full CBA which expresses all impacts on human health, lives lost, and the environment in monetary terms. (Admittedly, Stern reports *both* monetary and non-monetary assessments of the impacts). A discussion of incommensurability, and Stern's apparently mixed views on the issue, is far beyond the scope of this paper.[57] Leaving incommensurability aside,

---

[55] Some readers may baulk at the phrase 'incommensurably worse'. There are two possible interpretations. First, it is a shorthand for 'the alternatives are not directly comparable, but on other ethical grounds one alternative is worse'. Second, the alternatives are comparable (one is worse), but not commensurable – there is no cardinal measure for ranking them. For an influential statement of the comparability/commensurability distinction, see

[56] Stern, *The Economics of Climate Change*, pp. 163.

[57] The standard arguments used by economists to reject incommensurability problems in CBA are analysed in Jonathan Aldred, 'Incommensurability and Monetary Valuation', *Land Economics,* Vol. 82, No. 2 (2006), pp. 141-61.

any assessment of the role CBA should play in informing climate change policy will be crucially incomplete.

However, this paper can be interpreted as an exercise in ground-clearing, before we begin that assessment in earnest. The nature of climate change CBA has been clarified, showing how it is more an application of applied moral mathematics than science. With the ground cleared, we can approach the assessment via questions which will be familiar to institutional political economists: who benefits from different climate change policies and why? Relatedly, what interests do the cost-benefit analysts themselves have? Apart from an institutional analysis of the policy formation process – different interest groups, their power relations and the impacts of climate change upon them – political psychology may offer important insights too. Are arguments couched in cost-benefit terms more or less persuasive than the overtly 'moral' arguments which have in the past been used to encourage population-wide behavioural change?

More prosaically, some key issues in the Stern Review have been examined in detail. Although it represents a major improvement on past climate change CBA, stubborn difficulties remain. CBA claims just to value 'statistical' lives, but if enough people are affected, statistical lives add up to real lives lost. The justification for giving lives a monetary value is to achieve consistent policy-making, which requires a uniform monetary value. But a uniform value is hard to defend. Ethical disagreements lie at the heart of debates about the discount rate but the practice of discounting may itself obscure these disagreements. In particular, use of a single discount rate across different people and goods is illegitimate, and discounting disenfranchises future generations. Declining discount rates over time do not address these problems. Pure

uncertainty is a fundamental challenge for climate change policy-making, and it is hard to see how any purely quantitative approach can address it. In the meantime, the lesson of Weitzman's analysis for contemporary political debate is that instead of focusing on estimating the average impact of climate change, we must pay much more attention to the possibility of truly horrific catastrophe.

All these difficulties arguably suggest not that we must find better numbers to feed into climate change CBA, but that we need a different framework entirely, one which allows explicit ethical debate. If there are persuasive arguments for incommensurability, this conclusion would be strongly reinforced. But if we persist with CBA, whether for reasons of political expediency or otherwise, a first step would be to release the analysis from the meta-ethical straitjacket of attempting to base all parameters and monetary valuations on market data.