

10

Is Climate Change Cost–Benefit Analysis Defensible? A Critique of the Stern Review

PAUL BAER AND CLIVE L. SPASH

The threat of anthropogenic climate change raises numerous complex problems, but the issue is mainly framed as the need to cut global greenhouse gas (GHG) emissions, with often exclusive emphasis on carbon dioxide (CO₂). Environmental economists reduce the decision further to a monetary cost–benefit analysis (CBA), in which the costs of controlling GHG emissions are balanced against the benefits of avoiding induced climatic-related damages to human welfare. CBA climate professionals then claim an ability to calculate ‘optimal’ long-term policy choices. Such use of CBA, especially for global-scale problems, has been called into question on a variety of grounds by people inside (Vatn and Bromley 1994; Vatn 2000; Spash 2002b) and outside (Sagoff 1988; O’Neill 1993; O’Neill 1997) the economics profession.

Historically, CBA was developed to evaluate well-defined small-scale projects. Even at such a project level there is often scepticism relating to the necessary simplifications and assumptions. In particular, a host of controversial ethical choices are required because of the incommensurability of costs and benefits, the possibility of appropriate compensation, accounting for future generations and non-human species, income inequality, and the distribution of rights. The enormous uncertainties surrounding the relationship between causes of climate change, their potential impact, and valuation raise additional challenges. At the global and multi-century scale, the mismatch between the claims of robust and objective measurement and the realities of subjective and uncertain projection become profound. Yet, despite the considerable range and number of serious critiques,

the CBA approach remains influential and continues to be applied to the debate over preventing human-induced climate change.

A recent example is the so-called Stern Review (SR) (Stern 2006a), an economic analysis commissioned by the United Kingdom's (UK's) ruling Labour Party and chaired by Sir Nicholas Stern, a former chief economist at the World Bank. This report, released in October 2006 with a good deal of fanfare, is known primarily for its headline message that straightforward economic (cost–benefit) analysis justifies ‘prompt and strong action’ to reduce GHG emissions. The SR favours stabilization of GHG concentrations at 450–550 parts per million (ppm) CO₂ equivalent with a target of 500–550 ppm CO₂ equivalent, said to be achievable at a cost of about 1 per cent of gross world product (GWP); business as usual is estimated to cause losses of 5–20 per cent of GWP.¹ Prior to the SR, most climate CBA professionals produced numbers supporting little or no mitigation,² while non-economists and critics of CBA called for stringent mitigation. The significance of the SR is that mainstream economists are found claiming that CBA ‘done properly’ shows that rapid and significant emissions reductions are economically warranted.³ Professional climate economists have then felt the need to defend their own CBAs, and in particular their discounting of future harm (Mendelsohn 2006; Nordhaus 2006; Tol 2006; Yohe 2006). The main claim is that wrong conclusions are drawn due to the making of non-standard assumptions about discounting and so valuing future impacts more highly than those in conventional mainstream economics, that is, in the critics’ own models.

The resulting debate has focused upon whether CBA does warrant a limit of 550 ppm CO₂ equivalent. As a result, the case for a limit at 450 ppm CO₂ equivalent has been neglected. This means effectively accepting global average temperature increases above 2°C, despite this being the limit previously endorsed by both the UK government and the European Commission.⁴ The SR and

¹The scenario was one selected from amongst those of the IPCC (SRES A2).

²Cline (1992) is the most prominent exception.

³A mainstream economic argument for strong mitigation suits a neo-liberal leaning UK Labour government seeking to placate the business class. We return to the political context later.

⁴The SR acknowledges that stabilization at 450 ppm CO₂ equivalent offers at best a roughly even chance of keeping global mean temperature increase below 2°C, with a significant (order of 20 per cent) likelihood of an increase over 3°C. The SR

others then seem happy to use CBA to debate upper but not lower emissions limits. This raises concerns both about the role of CBA in general and the quality of the SR's analysis in particular, for deciding upon GHG control measures.

The SR does express some humility and even scepticism concerning the ability of Integrated Assessment Models (IAM) (linking emissions to economic losses) to produce precise quantified projections, saying at times that such calculations should be considered only as 'indicative'. Indeed, the fact that damages are attributed such a large range (5–20 per cent) makes uncertainty about the impacts of climate change—both their likelihood and their valuation—a central concern. Thus, a logical presumption would be that the choice of limits on emissions and climate was centrally determined by such uncertainties. However, it is shown here that the SR's own standards for addressing uncertainty and value controversy support neither the upper nor the lower bound with justifiable quantitative arguments.

In the remainder of this chapter, we consider how the SR conducts its quantitative analyses. Four subjects will be addressed in turn, namely the treatment of: future generations, risk and uncertainty, extreme and catastrophic impacts, and intra-generational ethics. These areas reflect the ways in which the SR claims to be innovative. We then discuss the issues raised by the reduction of all future climate damages to a single indicator of expected utility. We conclude with an assessment of whether the quantitative results are sufficiently robust to justify the policy conclusions and some interpretation of the role of the SR in the current climate policy context.

One background for our analysis is the framework of 'post-normal science'. Funtowicz and Ravetz (1994) used this approach to critique global CBA estimates of climate change control by Nordhaus (1991a; b). They showed that, in spite of appeals to various tenets of theory and economic estimates to several decimal places of accuracy, Nordhaus produced results on the basis of ad hoc assumptions, educated guesses, and controversial value judgements. Funtowicz and Ravetz (1994) focused on many of the same issues—scientific uncertainty, discounting of future generations, the valuation of impacts—that the SR highlights to differentiate itself from Nordhaus and others. However, the PAGE2002 model used by the SR is directly related

suggests fairly strongly that achieving 450 ppm CO₂ equivalent is already too expensive to be 'worth' the extra risk reduction it accomplishes, but does not rule it out.

to the model developed by Nordhaus and, in addition, the similar basic methodology means identical problems despite the attempted differentiation.

A key aspect of what follows is to show that the SR's argument for stabilization at 550 ppm lacks quantitative economic justification. Subtly different modelling choices allow a case for more stringent mitigation of 450 ppm or even lower. Precisely because the numbers are so pliable, they fail to show that lower targets are economically unwarranted. In addition, non-economic arguments are centrally important and the expression of plural incommensurable values essential in the policy debate. Thus, for many, a highly persuasive argument exists due to the expected physical impacts under 'business as usual' and the resulting inequitable distribution and imposition of harm on the innocent.

THE SR'S CBA ARGUMENT

Two main arguments are made in support of the SR's policy recommendations. First is a justification of targets using a comparison of the marginal costs and benefits of a single ton of CO₂ equivalent emissions. Second is the assertion that the likely damages from business as usual can be equated to a GDP loss of 5–20 per cent, while control costs equate to 1 per cent of GDP to stabilize atmospheric GHGs at 500–550 ppm CO₂ equivalent, which is claimed to avoid 'most of the worst impacts'.

The first argument relates to the theoretical holy grail of CBA. If an analyst could define and equalize the costs with the benefits of reducing a ton of carbon, they would be able to meet the conditions for defining the 'optimal' point for efficient pollution emissions reduction. In the SR, the comparisons are actually never made explicitly. However, in chapter 10 (with estimates based on Grubb et al. 2006), the SR reports that the marginal costs of emissions reductions for a 450 ppm CO₂ stabilization pathway (equivalent to around 500–550 ppm CO₂ equivalent) are around \$27 per ton CO₂ (±50 per cent) in 2030 and around \$15–\$70 per ton CO₂ in 2050.⁵ The benefits of control are reported as being in the order of \$85 per ton of CO₂ equivalent for business as usual, versus about \$30 per

⁵Note that the actual figures reported in the SR (Stern 2006a: 248) are incorrectly converted from tons C to tons CO₂, resulting in numbers that are too high by a factor of 13!

ton if concentrations are stabilized at 550 ppm CO₂ equivalent, and \$25 per ton if concentrations are stabilized at 450 ppm CO₂ equivalent.⁶ Taking these numbers at face value shows ambiguous support for even the 550 ppm upper limit.

The second argument is supposed to provide a relatively self-evident choice in favour of the 550 ppm target. However, as Mendelsohn (2006) has pointed out, this ignores the possibility that stabilization at, say, 650 ppm might also avoid ‘most of the worst impacts’ and have much lower mitigation costs. By the SR’s own admission, picking a stabilization target that is to be defended on grounds of welfare economics still requires a comparison of marginal costs and benefits. Thus, the SR states:

Our work with the PAGE model suggests that, allowing for uncertainty, if the world stabilises at 550ppm CO₂e, climate change impacts could have an effect equivalent to reducing consumption today and forever by about 1.1 per cent. As Chapter 6 showed, this compares with around 11 per cent in the corresponding ‘business as usual’ case—ten times as high. With stabilization at 450ppm CO₂e, the percentage loss would be reduced to 0.6 per cent, so choosing the tougher goal ‘buys’ about 0.5 per cent of consumption now and forever. Choosing 550ppm instead of 650ppm CO₂e ‘buys’ about 0.6 per cent.’ (Stern 2006a: 295).

Note here that the authors claim that the ‘marginal benefits’ of moving from 650 ppm to 550 ppm and 550 ppm to 450 ppm are roughly the same—in both cases, about half a per cent of GDP ‘now and forever’. This implies that the mitigation cost of moving from 650 ppm to 550 ppm would have to be less than half a per cent of GDP for 550 ppm to be clearly warranted on economic efficiency grounds. Yet in one table, the SR shows reductions of approximately this scale (some mismatch occurs because of the conversion of CO₂ to CO₂ equivalent levels), leading to costs at mid-century on the order of 1–4 per cent of GDP (Stern 2006a: 297, Table 13.4), and (in another table based on another meta-analysis), leading to discounted equivalent costs on the order of 0.3–0.8 per cent—(Stern 2006a: 296, Table 13.3). Thus, the marginal benefit of reducing CO₂ equivalent from 650 ppm to 550 ppm is not demonstrably larger than the cost.

⁶The SR reports on the ‘Social Cost of Carbon’, which is highly misleading terminology (on manipulation of cost terminology, see Spash 2002b: 172–77); this is actually referring to the marginal benefits of greenhouse gas reduction, and should not be confused with emissions control costs.

This shows that the quantitative analyses behind the policy recommendations are of questionable robustness, even at the level of the use of the numbers calculated. There is then good reason to be sceptical that monetary conversion and aggregation of impacts, and the appeal to 'state of the art' economic methods, actually can provide justification for policy recommendations. In addition, the case for setting a lower threshold of 450 ppm CO₂ equivalent remains open and may be as, or more, desirable on several grounds.

FUTURE GENERATIONS

The long time scale of human-induced climate change makes the question of our ethical responsibilities to future generations central to the framing of the problem. In economic analysis, this debate is centred on the concept of discounting. This is the practice of reducing the value of future costs and benefits in proportion to their distance in the future, typically through the use of an exponential discount rate (for more detailed discussion in the context of climate change, see Spash 1993, 2002a).

There is a consensus among philosophers—and some economists—that the economic practice of discounting can lead to a dangerous disregard for the well-being of future generations. There is an extensive literature on ethical issues relating to future generations that raises concerns over inter-generational justice, the role of rights and responsibilities, and the standing of future as opposed to present individuals.⁷ No such literature is cited in the SR, making its claim to be a 'review' rather implausible. This merely perpetuates the failure of public discourse to address the ethical implications of multi-generational environmental problems. The fundamental reason for concern is straightforward: for any positive discount rate, a time in the future can be specified at which the effective destruction of civilization would be literally 'not worth preventing'. The high discount rates typically used in climate CBAs mean that the importance of global catastrophe just a few decades in the future is vastly reduced in present decisions, and those in a century or more are effectively written off completely.

The SR recognizes that the choice of a discount rate (or perhaps, more appropriately, a discounting model) is inescapably normative.

⁷See for example, Callahan (1981); d'Arge et al. (1982); Norton (1982); Barry (1983); Page (1983, 1988); Parfit (1983, 1984); Howarth (1997); O'Neill (1999); and Gardiner (2006).

In the mainstream debate within economics, the key question is whether the economic analyst should include a ‘pure rate of time preference’ as a component in the discount rate. In general, mainstream economists accept the validity of a pure time preference and then debate the size of the rate. Despite numerous qualifications, and claims which seem to undermine the practice, the SR does exactly the same. Conventional welfare economics takes for granted that commodity discounting—reducing the importance of future costs or benefits in proportion to the (assumed) increased consumption of future generations—is well justified. Indeed, commodity discounting is based on the assumption of declining marginal utility from consumption, which is itself treated as a self-evident fact requiring no proof.

The SR uses a standard formula for combining the pure rate of time preference and the declining marginal utility of income to define the discount rate r :

$$r = \delta + \eta g \quad (10.1)$$

where δ (delta) is the rate of pure time preference, g is the growth rate of per capita consumption, and η (eta) determines the effect of economic (consumption) growth on the discount rate. The parameter η is also characterized as an ‘inequality aversion’ parameter in the SR, because of the way it is derived from the elasticity of the marginal utility of consumption. The higher the value of η , the greater the weight given to impacts on persons with lower consumption or income levels. Importantly, as we discuss below, η is also characterized as a ‘risk aversion’ parameter.

The dominant convention in CBA has been to use a relatively high discount rate (for example, 5–10 per cent). This includes, explicitly or implicitly, a significant positive pure rate of time preference (for example, 2–3 per cent). The term η is typically set to 1. The SR rejects such a high pure rate of time preference on ethical grounds, but then rather strangely reasserts a very small positive pure rate of time preference (0.1 per cent), based on the probability that human civilization may cease to exist in a century.⁸ The SR uses $\eta=1$, although sensitivity analyses using higher values of η have been added post publication of the main report in an ‘appendix to the postscript’, as a response to critics. Under $\eta=1$, the

⁸In fact, the SR appears to pick the 0.1 per cent number and then use it to estimate what the likelihood of extinction must be! See Stern (2006a: 46–7).

discount rate is equal to 0.1 per cent plus the economic growth rate, which averages 1.3 per cent annually between 2000 and 2200 in the SR's baseline world without climate change (Stern 2006a: 161). Even using this relatively low rate, impacts which occur 200 years in the future have just 6 per cent of their value compared to their occurring today.

Discounting has some appeal as a way of representing certain types of properties in a quantitative way for lay and expert groups. At an intuitive level for lay persons, the more you have the less it is valued (that is, marginal utility of income declines), and people in industrialized economies have been led to expect increases in real income over time. At a more theoretical level for experts, it fits into a family of models within which an ethical judgement (the relative value of consumption to different persons) can be reflected in a single parameter, and then (with a few additional assumptions) 'calibrated' on the basis of 'empirical data'. This ability to extract a rate from observations gives a supposed scientific objectivity that is employed to justify the policy consequences of using the selected parameter. For the analyst, then, all normative aspects are dispelled by a claim that discounting is an empirical fact which can be observed by an objective scientist regardless of any moral implications.

The formula in the SR is standard in welfare economics, and it is a classic example of the way in which economics mixes ethical and empirical claims in the justification of particular calculations. There are a range of problems with this whole approach. First, economists ignore empirical reality which shows individuals can and do hold negative discount rates for some impacts and positive ones for others, for example, bringing forward harms and delaying pleasures (Lowenstein and Prelec 1991). Second, there is no one discount rate in society and there are different rates for different groups, capitals, contexts, and so on. Third, different scenarios imply different rates, even in theory, which makes the rate endogenous to the climate change problem and its policy 'solution'. Fourth, merely observing that something occurs as an empirical fact says nothing of its moral acceptability or repugnance, for example, people murder, rape, torture, and commit genocide. Fifth, adding in risk to discounting conflates separate issues and makes untenable assumptions as to the nature of uncertainty. In brief, the SR fails to seriously address the arguments against discounting and lacks any reasoning as to why, even if one accepts discounting, zero or negative rates are inappropriate.

RISK AND UNCERTAINTY

Uncertainty over future human-induced climate change and impacts is a widely recognized major consideration affecting policy responses. The SR acknowledges this in a variety of places and claims that their approach to the incorporation of risk and uncertainty gives an improved estimate of the overall damages compared to previous climate CBAs. The authors discuss the relationship between risk and uncertainty referring to a variety of debates. However, the fundamental methodology employed reduces strong uncertainty (for example, partial ignorance, social indeterminacy) to known probabilistic events. There are also questionable assumptions about the characteristics of the resulting risk calculations in terms of risk aversion and the treatment of utility.

Economics and the related field of decision theory utilize an idealization of the problem of decision making under uncertainty in which actors—persons, firms, countries—are assumed to behave in such a way as to maximize ‘expected utility’. The approach integrates the probability that specific future states of the world will occur with the ‘utility’ or welfare from the realization of those states. Underlying this is a set of conditions or axioms of assumed ‘rationality’ that impose a very specific model of human behaviour. There has actually been an extensive debate regarding whether persons do, in fact, act ‘rationally’ in this sense, whether the model is fundamentally normative rather than descriptive, and if it is normative whether it is well justified.⁹ As a practical matter, there are clearly many cases where people fail to meet such expectations (for example, Gintis 2000). This brings into question the case for arguing that behaviour can generally be described by simple notions that abstract from the complexity of individual behaviour and empirical reality.

The approach also rapidly becomes messy because of the involvement of an increasing number of countries. Potential futures differ across multiple dimensions and different actors vary in their valuation of alternative states. Even if we assumed a finite number of possible outcomes, there is little reason to assume that there will be well-defined probabilities for those outcomes. Human-induced climate change holds the prospect of large-scale unique changes

⁹For a survey, see Smithson (1989); for a collection of articles, see Gärdenfors and Sahlin (1988).

outside human historical experience. The standard scientific approach of repeated experiments to produce an 'objective' probability distribution is then of no practicable use. Rather, the likelihood of some future state coming to pass is necessarily an opinion, perhaps a well justified opinion, perhaps a consensus opinion among a group, but an opinion nonetheless.

The problem confronting natural and social scientists in their role as policy advisers is how best to address this type of uncertainty. One approach is to take 'opinions' and create probabilities of future events and essentially treat these as if they were derived from empirically observed experiments. These subjective probabilities suffer from numerous problems, not least of which is who has the right to have their opinion determine the weight given to possible future events?

In terms of expected utility analysis, a method is required to incorporate 'loss aversion', that is, the recognized human preference to treat equivalent losses and gains asymmetrically. The SR addresses these concerns about uncertainty and risk aversion using two primary methods. First, a Monte Carlo model is used to create a probability density function (PDF) of climate outcomes and associated economic damages for a specified emissions pathway, based on 1,000 'runs' of the model (on the model, see Hope 2006). Second, a discount rate is employed in each run that varies with the 'realized' rate of economic growth, after climate damages have been subtracted. Since (as discussed above) the effective discount rate increases with economic growth, model runs with higher damages have lower discount rates. As a consequence, those runs with higher damages are weighted more heavily in the aggregation of the multiple Monte Carlo runs, creating loss aversion in a stylized fashion.

There are problems relating to both of these aspects. The model (PAGE2002) requires subjective PDFs for over 30 crucial inputs, everything from the climate sensitivity to the ratio of climate damages in different regions in response to temperature increase. In practice, for only climate sensitivity is there any significant literature on an appropriate PDF; for the remainder, the authors simply use their judgement based on any available evidence, however scanty.¹⁰ Furthermore, the PDFs used are triangular, which means there is zero probability of a value above or below some arbitrarily specified point.

¹⁰Hope (2006: 21) states that 'Most parameter values are taken from the IPCC Third Assessment Report', but it is evident that a great deal of subjective judgement went into converting the numbers into PDFs.

The SR acknowledges that the input PDFs are not well constrained, and indeed they address one aspect of this strong uncertainty by running their baseline climate scenario (the SRES A2 scenario) with alternative formulations of the carbon cycle feedback and possible methane releases. This so-called ‘high climate’ scenario leads to an increase in expected damages of about 35 per cent (Stern 2006a: 154–5), and is an important contributor to the 5–20 per cent range of reported damages. They also run the model with a higher climate sensitivity PDF (the baseline PDF has a modal value of 2.5°C and under no possibility can it be higher than 5.0°C), although these results are reported only in a single place (Stern, 2006a: 156), are referred to as ‘particularly speculative’, and play no role in the decision analysis. In their sensitivity analyses, they use alternative PDFs for the primary damage function (Stern, 2006b: 7–10), although later essentially ignore the results.

In a further gesture at the significance of strong uncertainty, the SR discusses, in chapter 2, a specific methodology for dealing with unknown probabilities leading to alternative calculations of expected utility. Citing an unpublished paper by Henry (2006), the authors recommend taking a weighted average of the highest and lowest expected utilities, where the weights ‘would be influenced by concern of the individual about the magnitude of associated threats, or pessimism, and possibly any hunch about which probability might be more or less plausible’. They conclude the discussion: ‘We now have a theory that can describe how to act’ (Stern 2006a: 34). Yet at the heart of the theory are concerns, pessimism, and hunches.

There are also good reasons to doubt the SR’s claim that the differential weighting of damages in an expected utility calculation, based ultimately on a parameter used to describe the declining marginal utility of consumption, is an adequate reflection of loss aversion. This requires the assumption that the monetized valuation of all possible impacts captures everything we care about—impacts must be translated into equivalent monetized losses to count. For example, risks of catastrophic species losses of 25 per cent or even 50 per cent or more would only enter the decision calculus inasmuch as one could put a monetary value on them.

Finally, there is a very powerful normative assumption in the claim that policy should aim to maximize the expected value of a scenario—even if possible losses are weighted higher than possible gains—rather than, say, reducing the risk of crossing some threshold to an acceptably low level. In theory, almost any level of loss aversion could

be 'programmed' into a model like PAGE2002. However, justifying the use of any particular function or parameter requires deciding in advance what risk of catastrophic outcomes should be accepted.

EXTREME AND CATASTROPHIC IMPACTS

The possibility of catastrophic impacts has been discussed in the context of possible states of the world or state changes called irreversible, non-linear, or discontinuous. In most climate CBA models, there is at best a highly stylized inclusion of catastrophic events. For example, Cline (1992) produced a central estimate of damages reaching 6 per cent of GDP with a 10°C warming, and 20 per cent of GDP lost under a pessimistic scenario. He showed that, even with a 5 per cent discount rate, incorporating only a small probability of catastrophe within such economic models is all that is required to justify 'aggressive' action (Cline, 1992: 6). In contrast, Nordhaus and Boyer (2000) estimated the 'willingness to pay' to avoid the risk of catastrophe by using a variety of ad hoc adjustments to an expert survey carried out much earlier (Nordhaus 1994). They used this to justify equating a 2.5°C warming to a 1 per cent loss of GDP and a 6°C warming to a 7 per cent loss of GDP. Even with the ad hoc adjustment (a large component of their estimated damages), 'optimal' global temperature increase is calculated to be 2.44°C above the 1900 level in 2105 (the end of the modelling horizon), just 0.09°C below the business as usual base case, and still rising at 0.20°C per decade.

The SR explicitly lists and represents graphically several of the risks associated with catastrophic impacts—most notably, the melting of ice sheets. However, this is then ignored. Instead, the PAGE2002 model used in the SR broadly follows Nordhaus and Boyer (2000) by including an aggregated probabilistic formulation in which, in every year of each model run, there is some probability (proportional to temperature) of extra GDP losses attributed to unspecified catastrophic impacts. The incorporation of this calculation in the end has the simple effect of raising the expected damages at any particular temperature, and thus at any specified level of emissions. The particular way in which the catastrophic damage function is calculated is necessarily quite arbitrary, as there is no well established basis for any such function or associated PDF.

All the impacts are monetized, and are by assumption presumed to be impossible below a 2°C increase, and never to exceed 20 per

cent of GDP lost in the 'focal region'. A scatter plot of model results reproduced in Warren et al. (2006) suggests that there is essentially a zero possibility of any catastrophic impacts until temperature exceeds 3°C. This is at best inconsistent with the scientific literature. In the SR itself, a finite probability is attributed to the melting of the Greenland ice sheet even below 2°C, with a resulting rise in sea level by several metres. However, an effectively zero probability is then used in the model. This choice, whether conscious or not, shows that contentious scientific and value judgements are embedded within the mathematical analysis of catastrophic impacts. This is a crucial example of how strong uncertainty is converted into weak uncertainty and impacts treated as some quasi-monetized risk.

EQUITY AND DISTRIBUTION

This section considers the distribution of costs and benefits within a generation (although the inter-generational issues are inextricably linked). The premise of welfare economics is that the utility of different individuals can and must be aggregated to calculate the overall ranking of a possible state of the world or outcome. This assumes that in comparing two specific outcomes, the gains to some persons can be directly added to the loss for others. In the ideal world of economic theory, a social welfare function transforms specific gains or losses in utility to particular individuals into cardinal numbers.

In models, such as PAGE2002, one or more 'representative individuals' are used in the calculation. Common formulae assume a declining marginal utility of consumption, which means that the marginal gain or loss from a unit consumed (measured in money) is 'more valuable' to a poor person than to a rich person. Noting this and adjusting calculations to take it into account is known as equity weighting. Depending upon the function and parameters used to model the declining marginal utility of consumption, the relative impact on the poor and the rich of an equal amount of lost consumption can be larger or smaller.

EQUITY WEIGHTING OF THE BENEFITS OF MITIGATION

Most climate CBAs ignore equity weighting and, therefore, implicitly take the distribution of income in society as it stands as being justified. This means that if a person who lives on \$2 a day or less loses \$1 and a millionaire gains \$2, the world is a better place. The few studies

that have included equity weighting have typically shown greater reductions to be warranted since standard damage assessments assert that poor regions will suffer greater proportional harm from anthropogenic climate change. Two such studies that do include equity weighting are cited by the SR (Stern 2006a: 156): Nordhaus and Boyer (2000) claim damages at 5°C warming increase from about 6 per cent to 8 per cent of GDP, while Tol (2002) states damages at 5°C double.¹¹

The SR asserts that equity weighting is appropriate. In fact, the model used is capable of providing regionally disaggregated damage estimates that could straightforwardly be used to calculate equity-weighted aggregate damages. The authors assert that a reasonable estimate of the impact of equity weighting would raise the maximum expected damage estimate associated with business as usual from 14.4 per cent to 20.0 per cent of GDP. This shows that the results of the model are quite sensitive to the use of a stylized incorporation of equity. Obviously, the move to 20.0 per cent of GDP has more to do with picking a nice round number rather than any specific and justifiable parameterization. The SR's authors could just as easily have picked any number—as long as the recommended mitigation remained acceptable to their political sponsors.

If the SR had actually calculated equity weights, this might have stimulated a debate over the numbers employed. The SR might have shown, for example, that equity weighting increased business as usual damages from 14.4 per cent to 19.7 per cent. Or perhaps, using different estimates for the relevant parameters, to anywhere from 17.3 per cent to 26.2 per cent. Or the equity weight could have followed Tol's analysis, which would have meant doubling the SR's damage factors. Such a debate would seem likely to have revealed the excessive precision being claimed on the basis of ad hoc assumptions—wrongly precise, and precisely wrong.

There is then a direct parallel with the estimate of GHG control benefits presented by Nordhaus (1991a; b) and criticized by Funtowicz and Ravetz (1994). After presenting a table of numbers with as many as three significant digits on some figures and not

¹¹In fact, these numbers appear to be taken from a graph in the IPCC's Third Assessment Report (Smith et al. 2001: Table 19.4) which is reproduced in the SR (Stern 2006a: 147). The numbers on the equivalent graph in Nordhaus and Boyer (2000: 95, Figure 4.3) appear to be closer to 7 per cent and 8 per cent, while the numbers from Tol's study do not appear in the cited 2002 paper, but appear to be from a working-paper version of the same study used in the TAR.

even a clear sign on others, Nordhaus simply increased the total to 2 per cent GDP loss (for a doubling of CO₂) to account for his intuitions. The SR follows Nordhaus in producing a figure with a calculated deceptive precision, which is simply arbitrarily rounded off to another number.

The idea of inequality aversion, described by a function or parameter, suggests that a collective social attitude towards inequality can be modelled, and appropriate functional forms and parameter values inferred from observable data. Hidden in this debate is a question about exactly what this weighting is supposed to mean, and how it is supposed to justify policy choices. The methodology of welfare economics is prone to a relatively frequent slipping back and forth between ostensibly normative and empirical concepts. The fact that an 'equity' parameter can be set by the preferences of the modeller seems to imply that the aggregate value of the outcome is simply an expression of the CBA climate experts' concern, or lack of concern, over poverty.

EQUITY WEIGHING OF THE COSTS OF MITIGATION

In economic analysis of pollution control, income inequality is generally taken as given, as if a natural consequence of life, while the distribution of mitigation costs is presumed to be a consequence of policy. This would seem to make its analysis a self-evident requirement of policy design. Strangely then, the distributional impacts of mitigation costs are rarely modelled or even discussed.

In any model assuming a declining marginal utility of consumption, the aggregate pollution control cost will be affected as much by distributional considerations as will the aggregate benefits from pollution control (for example, climate change avoidance). Mitigation costs can be made to appear arbitrarily small by distributing them to ever smaller and wealthier fractions of the population. The welfare impacts of a policy that can be modelled as a tax (as climate mitigation can) tend toward zero as the tax is shifted towards the wealthiest fraction of the population.

Countries like USA and Australia, which have opposed GHG emission control, tend to reflect a view that the wealthy should refuse a distribution of mitigation costs that burdens them disproportionately, even if it demonstrably minimizes global welfare losses. Climate change economists are generally inconsistent in their analysis of this position. They, as in the SR, assume that global

welfare maximization can be an effective justification for the choice of a stabilization target, but not for the distribution of mitigation costs. Yet, compared to control benefits, equity weighting might have a similar (or even greater) impact on the estimation of control costs to that for the estimation of control benefits. As a result, the actual emissions target being recommended would necessarily be different based upon the specific assumptions about the distribution of mitigation costs.

CBA LEGITIMACY AND EQUIVALENT CONSUMPTION LOSSES

Key to understanding the SR's argument for policy targets is to be aware of what the quantities measured as projected costs and benefits are intended to describe. The central figures fail to represent a range of possible impacts, but rather give a range of 'expected values' where a possible future in quasi-monetary terms is weighted according to its estimated likelihood.¹² We refer to quasi-monetary terms because both consumption (proportional to future GDP) and welfare losses (from climate harm) are transformed by a mathematical operation into utility. This move, which, as the SR authors note, is standard practice in applied economics, plays a variety of important roles in their analysis. Utility is then further aggregated, discounted, and compared at the margin, in order to allow comparison of control benefits (avoided damages) with the costs of reducing emissions.

There are several points that need to be made about this idea of expected utility. Crucially, there is no straightforward link to anything real in the world. Rather, it represents a hypothetical valuation of possible future worlds associated with some policy scenario, integrating the perceived likelihood of different possibilities with the presumed desirability of those possibilities. As such, it is a kind of judgement that can reasonably be expected to differ among different persons, and indeed the SR discusses the kinds of disagreements that might be expected to lead to different estimations of the value (expected utility) of a particular scenario. The idea that even a single individual could have a well-defined view of the expected utility of an uncertain future is open to serious question. As discussed, such projections involve addressing not merely processes that are well understood

¹²Indeed, if non-market impacts, catastrophic risks, and high feedbacks are taken into account, the SR's model calculates at least a 5 per cent likelihood of impacts exceeding 32.5 per cent of GDP (Stern 2006a:158).

but uncertain (in a probabilistic sense), but processes about which we are at least partially ignorant, or which are indeterminate due to human choice. The expected utility approach requires a world of ‘weak uncertainty’ in which the range of possible outcomes and their respective probabilities are well bounded (Spash 2002d), as opposed to one of ‘strong uncertainty’ (Spash 2002c). Thus, strong uncertainty must be reduced to weak uncertainty but such a move simultaneously undercuts the robustness of the resulting calculations.

A fundamental justification for such reductionism is the claim that choice amongst alternatives requires a single scalar index of ‘value’ to achieve a ranking. This conversion of all aspects of a scenario (from loss of life to the melting of the Greenland and West Antarctic ice sheets) into quantitatively commensurable objects is extremely controversial. Even supporters of such global CBA, like the SR authors, note that this is ‘problematic’ (Stern, 2006a: 145–146). There is no account of the involuntary imposition of physical harm and threat of harm to people spread across countries and generations. Framing the policy question as a trade-off between fewer commodities and greater risk of harm to the innocent is an ethical decision.

While plainly recognizing some of the issues, the SR is inconsistent in its treatment of the critical question about the ethical legitimacy of the conclusions reached by CBA. There are a variety of gestures towards questions about the priority of rights, the idea of stewardship, and other non-utilitarian approaches to justifying climate policy, and in these sections the SR appears humble about the role of economic considerations in such decisions. Yet the policy conclusions of the report fail to evidence this humility; rather they assert that economic analysis has set the upper and lower bound on reasonable policy objectives, and that ethical disputes about uncertainty, distribution, and fairness can only make adjustments within this range. As the SR states:

There will always be disagreements about the size of the risks being run, the appropriate policy stance towards risk, and the valuation of social, economic and ecological impacts into the far future. But the range suggested here provides room for negotiation and debate about these. And we would argue that agreement on the range stated does not require signing up to all of the judgements specified above. In presenting the arguments, for example, we have omitted a number of important factors that are likely to point to still higher costs of climate change and thus still higher benefits of lower emissions and a lower stabilisation goal. (Stern 2006a: 299–300).

In the section of the SR from which this is taken, the authors demonstrate the relationship between their argument, the ethical considerations that policy must address, and the many relevant kinds of uncertainty. This brings into question the ignoring of the case for a 2°C limit on allowable warming. The SR authors are asserting that people with differing values may differ on a stabilization target, but can only do so within the 450–550 ppm CO₂ equivalent range. They nonetheless admit to biasing their argument by omitting ‘a number of important factors’ that would have supported lower targets.

There is no specific reason given why people who reject the role of CBA in determining policy should accept the upper and lower ranges. On the contrary, people who support the 2°C target are simply classified as holding that the target should be met ‘whatever the cost’ and, hence, economically irrational. We suggest that the structure of the policy problem at hand includes an obvious asymmetry that has specific implications. Where, as a consequence of self-interested action, costs are imposed on others—a negative externality, in mainstream economic terms—we suggest that the burden of proof should be greater for arguing that a limit on pollution is too strict. If our emissions limits are too lax, more people will die from climate harm; if they are too strict, our economies will grow more slowly and some will have to consume a bit less than otherwise.¹³

Climate sceptics argue that the risks are in fact symmetrical. They argue that given the levels of global poverty—its relationship to preventable death, health risks, and indeed even vulnerability to climate extremes—reducing economic growth today will cause harm of the same moral consequence as humans-induced climate change. That is to say, millions of people may die as a consequence of climate change mitigation, due to slower economic growth in poor countries. This is an argument that must be taken seriously; clearly there are tens of millions of preventable deaths annually from causes related fundamentally to poverty, a number that greatly exceeds estimates of likely deaths from human-induced climate change in the near future. Put simply, in a world in which premature death is ubiquitous, there are opportunity costs to investing resources in any one approach to reducing it. This is the heart of the argument made by economists such as Schelling (1997) and Tol (2006) as well as by political lobbyists such as Lomborg (2006) and other professional environmental sceptics:

¹³This asymmetry could be a primary justification for the precautionary principle, though it is rarely articulated this way.

many more lives in poor countries could be saved by other ways of investing the same money than will be saved by emissions mitigation.

This is a relevant argument with regard to setting an emissions objective. However, turning the argument around slightly shows it is not decisive in favour of the sceptics. The structure of the case can be explained as follows: Group A is carrying out an activity (call it polluting) that causes X deaths to group B, and it would cost Group A \$Y to eliminate those X deaths; but if Group A can reduce X deaths in Group B for \$Z, which is less than \$Y but does not address the pollution problem, that would be morally preferable and the two groups can negotiate how to divide the surplus; this potentially allows more lives to be saved than by eliminating the pollution. The point to this example is not that one solution is *a priori* right or wrong, but rather that neither solution is free of moral judgement or dilemma. On the one hand, we can save more lives and on the other, we choose that some people will die due to preventable pollution so that others might live. We have simply rediscovered the basic conflict between the utilitarian intuition, that the sum of all harm matters, and the deontological intuition, that some categories of harm should just be avoided. We might go further and raise some assessment of the democratic legitimacy of the process of making such a decision and the problems of who specifically dies being different and so on.

This gets to the heart of the debate over the applicability of CBA to climate change and other risk assessment problems. One obvious issue is that the policy choice at hand—how much to reduce greenhouse pollution—is not in fact being debated in the context of the question ‘what would be the best way to save lives in developing countries?’ No one is saying ‘Instead of reducing emissions by such-and-such per cent, we will invest in sanitation, or malaria reduction, or whatever’. Furthermore, since those who are most at risk are the poorest people—who are not participating in policy debate—and also poor (as well as wealthy) people in the future, the procedural legitimacy of any decision to sacrifice the specific interests of those at risk from human-induced climate change faces a substantial challenge. This is magnified by the fact that those who benefit most from GHG emissions are today’s wealthy.

The controversy over using CBA in this context is precisely about such issues as assuming we can legitimately trade lost lives for consumer goods—a dilemma which symbolizes debates about commensurability. The structure imposed by standard economic

analysis makes inevitable the reduction of lost lives to their equivalent in lost consumption, a move that is in many contexts and to many people morally indefensible, and that is indeterminate even if it is accepted as necessary. In the end, the numbers produced by the SR are only meaningful if one accepts that the prospective human deaths (plus extinction of species and other losses) due to human-induced climate change can be defensibly converted into equivalent amounts of consumption today.

CONCLUDING REMARKS

Mainstream economics addresses all of the major areas of impact under the enhanced greenhouse effect (GHE)—future generations, risk and uncertainty, extreme and catastrophic impacts, distributional equity—through the unjustifiable reduction of complexity and ethical controversy into a single scalar value. Reasonable differences about choices in each of these areas lead to very wide variance in the possible valuation of alternative policy scenarios. There cannot be a ‘correct’ value associated with any specific scenario nor ‘correct’ selection of a limited set of future scenarios. Moreover, the claim of such reductionism to any authority at all depends upon the assertion that, in fact, the core problems of CBA—commensurability, compensation, and the distribution of impacts—can all be adequately incorporated in a scientifically objective framing by an elite group of professional climate CBA experts. Such authority is clearly unwarranted.

The SR’s authors have plainly expressed their desire to be persuasive in the policy debate as they perceive it. To achieve this they have used a variety of methods of rhetorical and quantitative argument, but the SR’s persuasiveness also depends upon the social and political context and the broader credibility of its authors. The credibility of mainstream economic analysis in general rests in part on three crucial factors: the disciplinary authority of mainstream economics in the elite academic world, the apparently robust quantitative measures it produces, and its flexibility in supporting policies desired by economic elites. We suggest that the apparent insensitivity, or lip service, of standard economic analyses to relevant alternative considerations (for example, justice) is not simply a regrettable flaw, but rather a critical failure undermining the justification for giving economic analyses such a great weight in policy making.

The SR makes some effort to point out that action is supported

even without relying upon the aggregation of all mitigation costs and benefits into a single comparable figure. Nonetheless, at the heart of the analysis is a model which reduces uncertainty to risk and all climate impacts to a single quasi-monetary value, to be compared with an equivalent quasi-monetized pollution control cost. The SR recognizes the fallacy of a single number approach but proceeds regardless. Some carefully crafted arguments regarding the uncertainty of the results are then meant to justify the numbers calculated as upper and lower bounds on 'reasonable' stabilization targets.

Three points then need to be raised in drawing conclusions about the SR. First, addressing human-induced climate change has created a complex political debate, in which there is a vast distance between parties who see effectively no mitigation to be warranted, and those who see extremely stringent mitigation as warranted. The extreme positions can be caricatured as those who see a global average temperature increase of 4°C or more as no problem or even beneficial, and those who see a temperature increase of only 2°C as an unmitigated human disaster. Evidence supporting a middle path then has an air of respectability and political rationality, whether produced and paid for by an economic and political elite with vested interests of its own, whether right or wrong.

Second, in this debate, those who oppose stringent mitigation typically speak in the language of economics, and oppose mitigation on the basis of projected financial costs. Typically, the costs highlighted are aggregated at the national level, especially in the USA, although some reference may be made to the global economy and/or particular economic sectors. The fundamental method of reducing GHG emissions is perceived to be via the reduced usage of fossil fuel energy. Modern economies are heavily dependent on fossil fuels and stored energy in general. The idea of controlling consumption via demand management is outside the political frame. Short-term costs are then intuitively accepted to be high for any stringent mitigation effort.

This political economic battleground was staked out by the energy industry and the trenches dug some time ago. In this respect, perhaps there should be no surprise that the SR's major 'innovations' are not particularly innovative, and reflect work by other modellers. Similar problems to these are also evident in the treatment of catastrophes and the distribution of impacts. Rather than innovation, the SR delivers only highly subjective and scientifically questionable PDFs (in the case of catastrophic impacts), and an arbitrary multiplication factor and a gesture at 'further research' (in the case

of distributional equity). A more thorough approach would have made the analysis appear less rather than more robust. In addition, the whole economic framing of the problem would have been brought into question.

Third, these debates take place across communities within a political economy. There is a more-or-less academic community, in which there is a presumption of commitment to reasoned and disinterested argument. Then there is a political community in which parties use arguments strategically, attempting to win support for the policies they prefer by selecting favourable evidence and attempting to discredit evidence which opposes their vested interests. The scientific foundations of human-induced climate change mean that political actors legitimize their policy arguments on the *prima facie* credibility of the academic community, and deploy a wide range of 'scientific evidence'. This is plainly not the conduct of a disinterested truth-seeking exercise. The approach assumes that the best process for seeking truth is to have zealous advocates make their case and weaken those of their opponents in the 'if you are not with us then you are against us' school of thought. In this regard, the primary focus on justifying the higher 550 ppm CO₂ equivalent upper limit is clearly a political statement.

The discounting in the SR is still substantial, and the justification is open to question on a variety of grounds. The reduction of strong uncertainty to expected utility with a particular function is methodologically flawed, and even putting this to one side, the treatment of weak uncertainty could easily justify more serious risk aversion. The treatment of catastrophic risk has implausibly low damages as temperature increases between 2°–3°C. The sources cited by the SR for calibrating equity weighting justify higher possible damage adjustments. All of these would argue for greater mitigation. Yet in the end, the SR chose to place a minimum stabilization level at a threshold which the authors themselves claim has at least even odds of exceeding a 2°C warming.

A close look shows many reasons why the critical issues concerning the enhanced GHE cannot be decisively resolved in any quantitative exercise. Each area of the modelling process requires subjective judgements about likelihood and valuation, which lead to large changes in the results. Metaphorically, the model has a bunch of control knobs that can be turned to different settings to represent different views of particular concepts, mixing (for example) views about

ethical responsibilities to future generations with views about the risk of exceeding some climatic threshold. Despite the mathematical formalism and air of objectivity (employed by all global CBAs), no purely scientific determination is possible for the settings of these knobs, and there are plausible settings of the control knobs which would warrant even more stringent mitigation.

Among those most opposed to GHG regulation are the industries (notably oil, coal, electricity, and transportation) which suspect that the greatest impact will be on their power and profits. These industries include many of the world's largest multinational corporations and also corporations with enormous influence in particular countries and over ruling governments. Greenhouse gas regulation must literally be imposed against the will of many of these corporations, which can in turn count on popular support from politicians, consumers, and workers who expect to see prices increase and jobs lost. For those who most vehemently oppose mitigation on 'economic' grounds, the fact that 550 ppm has been shown to be economically warranted will not be convincing—since they are not interested in being convinced—and they will continue to use the opinions of Mendelsohn, Nordhaus, Tol, and their like to defend themselves. The point of the SR is to enlist the prestige of economics to persuade the uncommitted rather than to persuade the committed opponents of mitigation. In this regard, the fact that the analysis is not robust is of minor importance. No one whom it intends to persuade is expected to read or understand it any more than those appealing to the ad hoc numbers produced by Nordhaus for 20 years have ever paid any attention to their fallibility.

REFERENCES

- Barry, B., (1983), 'Intergenerational Justice in Energy Policy', in D. MacLean and P.G. Brown. (eds), *Energy and the Future*; Totowa, New Jersey: Rowman and Littlefield, pp. 15–30.
- Callahan, D., (1981), 'What Obligations do We have to Future Generations?', in E. Partridge (ed.), *Responsibilities to Future Generations: Environmental Ethics*, Buffalo, New York: Prometheus Books, pp. 73–85.
- Cline, W.R., (1992), *The Economics of Global Warming*. Harlow, Essex: Longman.
- d'Arge, R.C., W.D. Schulze and D.S. Brookshire, (1982), 'Carbon Dioxide and Intergenerational Choice', *American Economic Association Papers and Proceedings*, Vol. 72, No. 2, pp. 251–56.

- Funtowicz, S.O. and J.R. Ravetz, (1994), 'The Worth of a Songbird: Ecological Economics as a Post-normal Science', *Ecological Economics*, Vol. 10, No. 3, pp. 197–207.
- Gärdenfors, P. and N.E. Sahlin, (1998), *Decision, Probability, and Utility: Selected Readings*, Cambridge: Cambridge University Press.
- Gardiner, S.M., (2006), 'A Perfect Moral Storm: Climate Change, Intergenerational Ethics and the Problem of Moral Corruption', *Environmental Values*, Vol. 15, No. 3, pp. 397–414.
- Gintis, H., (2000), 'Beyond Homo Economicus: Evidence from Experimental Economics', *Ecological Economics*, Vol. 35, No. 3, pp. 311–22.
- Grubb, M., C. Carraro and J. Schellnhuber, (2006), 'Technological Change for Atmospheric Stabilization: Introductory Overview to the Innovation Modelling Comparison Project', *Energy Journal*, Vol. 27, Special Issue 1, pp. 1–16.
- Henry, C., (2006), 'Decision-making under Scientific, Political and Economics Uncertainty', *Chaire Développement Durable*, DDX-06-12 Paris: Laboratoire d'Econometrie de l'École Polytechnique.
- Hope, C., (2006), 'The Marginal Impact of CO₂ from PAGE2002: An Integrated Assessment Model Incorporating the IPCC's Five Reasons for Concern', *The Integrated Assessment Journal*, Vol. 6, No. 1, pp. 19–56.
- Howarth, R.B., (1997), 'Sustainability as Opportunity', *Land Economics*, Vol. 73, No. 4, pp. 569–79.
- Lomborg, B., (2006), (ed.), *How to Spend \$50 Billion to Make the World a Better Place*, Cambridge: Cambridge University Press.
- Lowenstein, G. and D. Prelec, (1991), 'Negative Time Preference', *American Economic Review*, Vol. 81, pp. 347–52.
- Mendelsohn, R., (2006), 'A Critique of the Stern Review', *Regulation*, Winter, pp. 42–6.
- Nordhaus, W.D., (1991a), 'A Sketch of the Economics of the Greenhouse Effect', *American Economic Review*, Vol. 81, No. 2, pp. 146–50.
- , (1991b), 'To Slow or Not to Slow: The Economics of the Greenhouse Effect', *Economic Journal*, Vol. 101, pp. 920–38.
- , (1994), 'Expert Opinion on Climate Change', *American Scientist*, Vol. 82, No. 1, pp. 45–51.
- , (2006), 'The Stern Review on the Economics of Climate Change', 16 November, <http://nordhaus.econ.yale.edu/SternReviewD2.pdf>.
- Nordhaus, W.D. and J. Boyer, (2000), *Warming the World: Economic Models of Global Warming*, Cambridge, MA: The MIT Press.
- Norton, B.G., (1982), 'Environmental Ethics and the Rights of Future Generations', *Environmental Ethics*, Vol. 4, No. 4, pp. 319–37.
- O'Neill, J., (1993), *Ecology, Policy and Politics: Human Well-being and the Natural World*, London: Routledge.
- , (1997), 'Managing without Prices: On the Monetary Valuation of Biodiversity', *Ambio*, Vol. 26, pp. 546–50.

- , (1999), 'Self, Time and Separability', in T.-C. Kim and R. Harrison (eds), *Self and Future Generations: An Intercultural Conversation*, Cambridge: White Horse Press, pp. 91–106.
- Page, T., (1983), 'Intergenerational Justice as Opportunity', in D. Maclean and P.G. Brown (eds), *Energy and the Future*, Totowa, New Jersey: Rowman and Allanheld, pp. 33–58.
- , (1988), 'Intergenerational Equity and the Social Rate of Discount', in V.K. Smith (ed.), *Environmental Resources and Applied Welfare Economics*, Washington, DC: Resources for the Future, pp. 71–89.
- Parfit, D., (1983) 'Energy Policy and the Further Future: The Social Discount Rate', in D. Maclean and P.G. Brown (eds), *Energy and the Future*, Totowa, New Jersey: Rowman and Allanheld, pp. 31–37.
- Parfit, D., (1984), *Reasons and Persons*, Oxford, England: Clarendon Press.
- Sagoff, M., (1988), *The Economy of the Earth: Philosophy, Law, and the Environment*, Cambridge: Cambridge University Press.
- Schelling, T.C., (1997), 'The Costs of Combating Global Warming: Facing the Trade-offs', *Foreign Affairs*, Vol. 76, No. 6, pp. 8–14.
- Smith, J.B., H.J. Schellnhuber, M.Q. Mirza, S. Fankhauser, R. Leemans, L. Erda, L. Ogallo, B. Pittock, R. Richels, C. Rosenzweig, U. Safriel, R.S.J. Tol, J. Weyant, and G. Yohe, (2001), 'Vulnerability to Climate Change and Reasons for Concern: A Synthesis', in J.J. McCarthy, O.F. Canziani, N.A. Leary, D.J. Dokken, and K.S. White (eds), *Climate Change 2001: Impacts, Adaptation, and Vulnerability*, Cambridge: Cambridge University Press, pp. 913–67.
- Smithson, M., (1989), *Ignorance and Uncertainty: Emerging Paradigms*, New York: Springer-Verlag.
- Spash, C.L., (1993), 'Economics, Ethics, and Long-term Environmental Damages', *Environmental Ethics*, Vol. 15, No. 2, pp. 117–32.
- , (2002a), 'Dividing Time and Discounting the Future', in C.L. Spash, *Greenhouse Economics: Value and Ethics*, London: Routledge, pp. 201–20.
- , (2002b) (ed.), *Greenhouse Economics: Value and Ethics*, London: Routledge.
- , (2002c), 'Strong Uncertainty: Ignorance and Indeterminacy', in C.L. Spash (ed.), *Greenhouse Economics: Value and Ethics*, London: Routledge, pp. 120–52.
- , (2002d), 'Weak Uncertainty: Risk and Imperfect Information', in C.L. Spash (ed.), *Greenhouse Economics: Value and Ethics*, London: Routledge, pp. 97–119.
- Stern, N., (2006a), *Stern Review on the Economics of Climate Change*, UK Government Economic Service, London, available at www.sternreview.org.uk.
- , (2006b), 'Technical Annex to Postscript', *The Stern Review on the Economics of Climate Change*, London, available at [http://www.hm-treasury.gov.uk/media/1/8/Technical annex to the postscript P1-6.pdf](http://www.hm-treasury.gov.uk/media/1/8/Technical_annex_to_the_postscript_P1-6.pdf).

- Tol, R.S.J., (2002), 'Estimates of the Damage Costs of Climate Change: Part II Dynamic Estimates', *Environmental and Resource Economics*, Vol. 21, pp. 135–60.
- , (2006), 'The Stern Review on the Economics of Climate Change: A Comment', 2 November, available at www.fnu.zmaw.de/fileadmin/fnu-files/reporst/sternreview.pdf.
- Vatn, A., (2000), 'The Environment as Commodity', *Environmental Values*, Vol. 9, No. 4, pp.493–509.
- Vatn, A. and D.W. Bromley, (1994), 'Choices without Prices without Apologies', *Journal of Environmental Economics and Management*, Vol. 26, No. 2, pp. 129–48.
- Warren, R., C. Hope, M. Mastrandrea, R.S.J. Tol, N. Adger, and I. Lorenzoni, (2006), 'Spotlighting Impacts Functions in Integrated Assessment', *Tyndall Centre on Climate Change*, Working Paper No. 91, Norwich.
- Yohe, G., (2006), 'Some Thoughts on the Damage Estimates Presented in the Stern Review: An Editorial', *The Integrated Assessment Journal*, Vol. 6, No. 3, pp. 65–72.