A long-run target for climate policy: the Stern Review and its critics

Simon Dietz

Grantham Research Institute on Climate Change and the Environment and Department of Geography and Environment, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

s.dietz@lse.ac.uk

Part of a consultancy project for the Committee on Climate Change Secretariat to provide analytical support on the long-term review

2 May 2008

Table of contents

1. Introduction

This report is intended to support the work of the Committee on Climate Change (CCC) Secretariat in its review of the UK's greenhouse gas emissions reduction target for 2050 (defined here as the 'long run'). It first sets out the approach taken in the Stern Review to proposing a range of targets for stabilising the atmospheric concentration of greenhouse gases in 2050. It goes on to survey the vigorous academic debate that followed the Stern Review's publication, with a focus on contributions from academic economists, mainly to the question of how to measure the benefits of emissions reductions.[1] In doing so, I have tried to be as comprehensive as possible, yet it is inevitable that I will not have included all of the relevant contributions. Finally, I will draw conclusions on how the Stern Review and the ensuing debate about it have advanced our knowledge of long-run targets for global greenhouse gas emissions, as well as identifying avenues for future research.

2. The Stern Review's approach to setting climate targets

2.1. Three lines of investigation

The Stern Review's analysis of long-run climate targets essentially spanned the first three parts of the published report, some thirteen chapters (although not all of these thirteen chapters are directly relevant). Table 1 summarises the different sources of evidence. All this evidence was brought together in chapter 13 to make an argument for a target range for stabilisation of the atmospheric stock of greenhouse gases at between 450 and 550 parts per million of carbon dioxide equivalent (ppm CO_2e).

[1] I devote most of the report to difficulties in estimating the benefits of emissions reductions, rather than the costs. Discussing the latter would require a much longer treatment and the contribution of additional experts on energy systems and technology.

Table 1. Direct evidence on the benefits and costs of emissions reductions.

Chapter	*Benefits*	*Chapter*	*Costs*
1, 3, 4, 5	Disaggregated analysis of the physical impacts of climate change on multiple dimensions (e.g. water and food availability, health, and infrastructure)	9	'Bottom-up' estimates of the costs of specific mitigation strategies, based on different portfolios of technologies
6	Integrated assessment (i.e. aggregated analysis) of the economic cost of climate change	10	Macroeconomic modelling of the cost of emissions reductions

Figure 1, reproduced from figure 13.3 of the Stern Review, sets out the logic of the comparison made in chapter 13. It compares the incremental benefits and costs of tightening (relaxing) the eventual level of greenhouse gases in the atmosphere, at which the stock is stabilised. Because neither the costs of emissions reductions nor in particular the benefits are known with certainty, we are presented with overlapping ranges of incremental benefits and costs and hence we can do no better at present than putting forward a range of targets. At the same time, because climate change is a stock externality[2], one must compare benefits and costs along a trajectory of emissions into the long run. That is why the comparison is of the present values of current and future benefits and costs along paths to stabilisation, not just of the benefits and costs of emissions now.[3]

[2] The impacts of climate change, typically indexed by global mean surface air temperature, are a function of the atmospheric concentration of greenhouse gases and consequently cumulative emissions of greenhouse gases over a very long period, rather than the flow of emissions over a period of just a few years.
[3] Note also that there are multiple paths of emissions to any given stabilisation target, a feature of the analysis that figure 1 does not emphasise.

Figure 1. Schematic representation of how to select a stabilisation target.

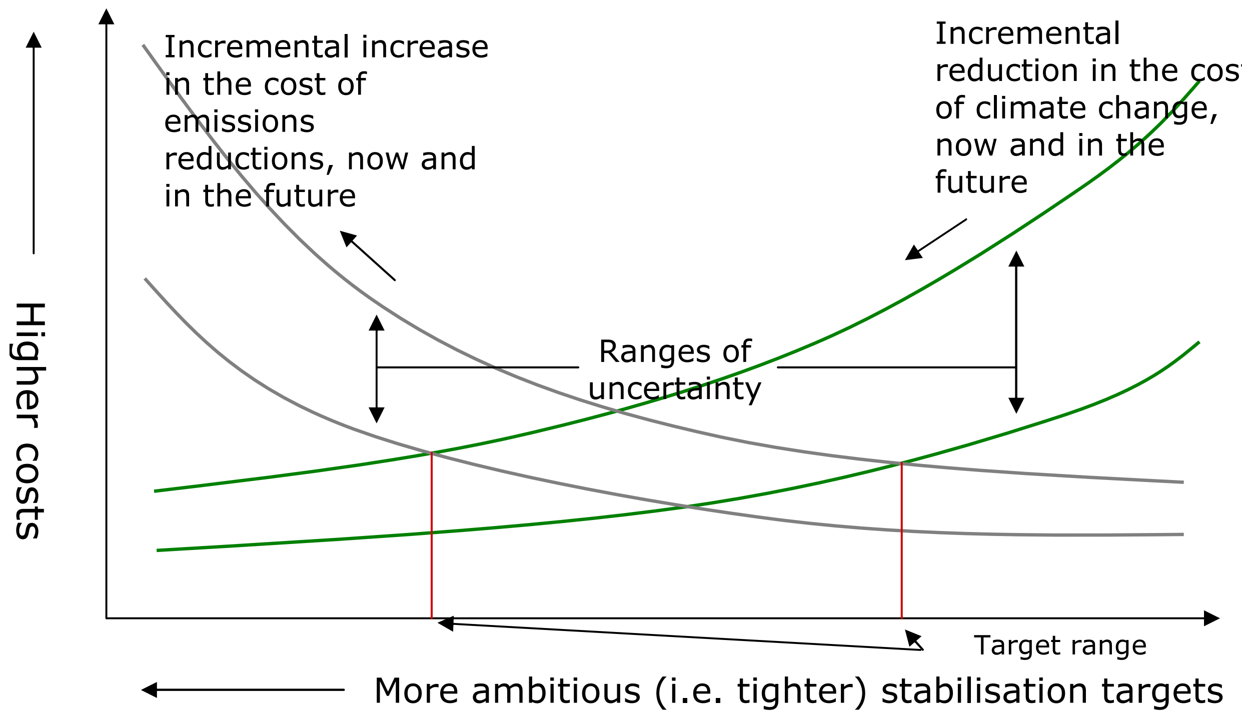

How can the Stern Review's approach to setting climate targets be understood in comparison with other existing analyses? This has in fact been a locus for some of the 'controversy' about the Stern Review since its publication. In a recent report for the CCC Secretariat on setting the first three carbon budgets for the UK, Watkiss *et al*. (2008) create a typology of approaches to setting climate targets and locate the Stern Review on it. While it is difficult to accurately capture all the key features of a diverse range of approaches in one typology, I believe it has pedagogical value. The typology is summarised in figure 2 and organises the extant literature along a spectrum of decision entry-points from formalised economic analysis to what might be described as non-economic or 'precautionary' approaches.

A significant body of economic analysis in the academic literature has used what Hope (2005) describes as 'optimising' Integrated Assessment Models (hereafter IAMs) to calculate the single path of emissions going forward that minimises the present value of the costs and benefits of emissions reductions. This is the basic economic appraisal criterion for policies that maximise social welfare[4] and is what is described in figure 2 as 'formalised cost-benefit analysis'. The literature comprises dozens of papers over the last twenty years or so, but in fact it relies on just a handful of models that are capable of optimising the path of emissions in this way, most notably William Nordhaus' DICE model (and its

[4] By basic I mean that one has to make many simplifying assumptions.

offshoots, like RICE), Richard Tol's FUND model, Alan Manne *et al.*'s MERGE model, and Chris Hope's PAGE model. Because these models strive for a powerful decision criterion, they must simplify radically and thus they are highly aggregated.

What is important and even 'characteristic' of this literature is that it has tended not to make such strong recommendations on emissions reductions as the Stern Review (Tol and Yohe, 2006). To give a very recent example, new analysis by Nordhaus (2008) with the latest version of the DICE model running with standard assumptions puts the optimal concentration of CO_2 in 2050 at 481 ppm.[5] Moreover, the optimal stock of greenhouse gases continues to rise until around 2175, when it reaches nearly 700 ppm CO_2. Because there are differences in assumptions about business-as-usual emissions, the difference between the recommendations of these two studies becomes still clearer if one looks at emissions control rates with respect to baseline emissions (i.e. emissions reductions as a percentage of business-as-usual). Chapter 8 of the Stern Review calculates that the emissions control rate for all six 'Kyoto' greenhouse gases[6] in 2050 will have to be 60-85% (depending in particular on the target), whereas in Nordhaus the optimal emissions control rate for carbon dioxide is around just 25%.

However this is not an immutable feature of the literature on formalised cost-benefit analysis. It is often mistakenly observed that economic analysis of this sort *cannot* support emissions reductions as strong as those put forward in for example the Stern Review. Instead it is more accurate to observe that they *tend not to*, since the analysis depends on many important assumptions, some of which are normative in character. It is possible to support stabilisation targets in the range 450-550 ppm CO_2e using formalised cost-benefit analysis, depending on the structure of the model and assumptions about parameter and variable values. One way to do so is to set a low social discount rate, based on a 'prescriptive' ethical standpoint. By contrast a 'descriptive' approach to discounting, which relies on market data to reveal people's preferences for substituting welfare now for welfare in the future, would set a higher social discount rate. I will return to this issue below. It is further possible to arrive at a very ambitious optimal level of emissions reductions through means other than a low social discount rate. Thus, for instance, a small but emerging number of analyses consider the effect of

---

[5] The comparison is complicated by the fact that Nordhaus only control CO_2, while the Stern Review discusses stabilisation of the basket of six 'Kyoto' gases, which comprises CO_2, methane, nitrous oxide, PFCs, HFCs and SF_6. However, Nordhaus' optimal concentration of all six Kyoto gases would almost certainly exceed 550 ppm CO_2e.

[6] See footnote five.

modelling uncertainties in the science of climate change, in particular including low-probability/high-impact events in the analysis.[7] I will also return to this issue below.[8]

By contrast to formalised cost-benefit analysis and at the other end of the spectrum, a number of studies have considered targets based on a disaggregated analysis of environmental and social impacts of climate change and with little explicit analysis of the cost of emissions reductions. These are described in figure 2 as precautionary/non-economic approaches. Examples include the UK Royal Commission on Environmental Pollution (RCEP) report from 2000 and the analysis underpinning the European Union's target to prevent global mean temperatures exceeding 2°C above pre-industrial levels.[9] These studies tend to recommend very stringent targets indeed. As a number of commentators have noted, the target range recommended by the Stern Review gives us, following the Review's own analysis, at best a fifty-fifty chance of meeting the EU's 2°C target (for stabilisation at 450 ppm CO_2e and interpreting the results of climate modelling studies as probabilities).

In the middle of the spectrum lie a number of approaches that compare the costs and benefits of emissions reductions explicitly, but in ways that do not rest on formal economic appraisal.[10] One example is multi-criteria analysis, which is a formalised method of ranking policy alternatives (i.e. in this case different stabilisation levels) according to their quantitative performance against a list of criteria (i.e. environmental, social and economic), which are themselves quantitatively weighted but not according to welfare-economic theory (i.e. not according to their shadow prices). Cost-effectiveness analysis is also classified in this group. Cost-effectiveness analysis does not by itself have anything to say on the benefits of emissions reductions, but it has tended to be used as an *ex-post* check on the technical and economic feasibility of targets supported by precautionary/non-economic analysis. This is precisely how the 2003 UK Energy White Paper built on the earlier recommendations of the RCEP.

It is to this middle group of approaches that the Stern Review can be argued to belong. On the one hand, its approach to setting stabilisation targets comprised explicit consideration of both the benefits

[7] See Weitzman (2008) on theoretical aspects of this. Empirical analysis has been conducted by e.g. Ackerman and Finlayson (2007), Azar and Lindgren (2003), Ceronsky *et al*. (2005) and Roughgarden and Schneider (1999).

[8] Another assumption that would drive downwards the optimal trajectory of emissions reductions is of course cheaper emissions reductions. In certain cases, adding the assumption of endogenous technical change, whereby near-term emissions reductions reduce the costs of future emissions reductions, does just this. See for example Alberth and Hope (2007).

[9] Although the analysis underpinning the EU target is difficult to trace, as Tol (2007) has argued.

[10] For a more detailed discussion of all the approaches set out in figure 2, see Watkiss *et al*. (2008).

and costs of emissions reductions. On the other hand, it is important to appreciate that the Stern Review did not calculate the optimal emissions pathway in a formal sense. Thus it is classified (accurately if somewhat clunkily) as 'non-formalised cost-benefit analysis'.

Figure 2. Decision perspectives and methods for setting climate targets. Note that the Stern Review target range is reported to be 500-550 ppm CO_2e, on the basis that the Review was not able to find any evidence on the costs of achieving a target as low as 450 ppm CO_2e. Nonetheless the Review proposed that 450 ppm CO_2e lay within the recommended target range.

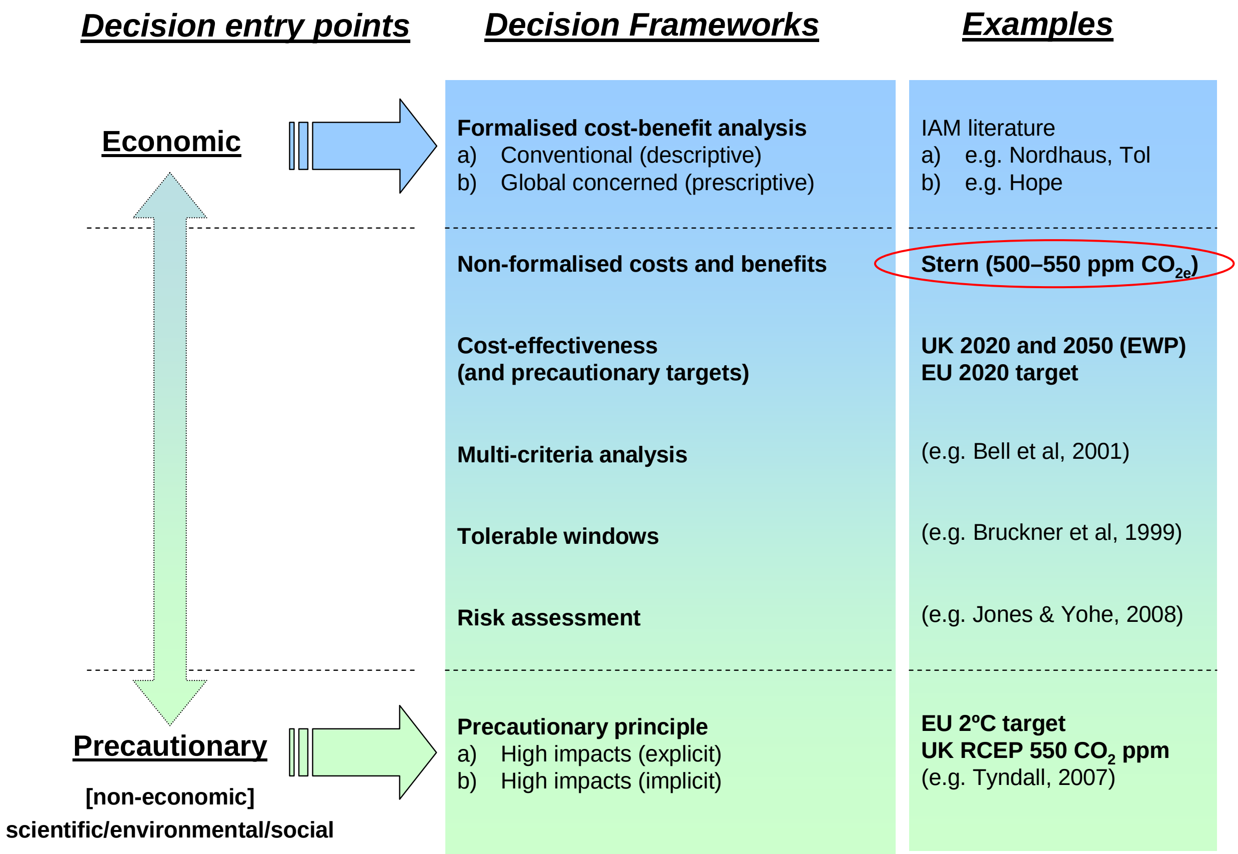

The Stern Review presented a large volume of analysis on the benefits and costs of emissions reductions, as summarised in table 1. In addition, as I have just argued it did not correspond to either of the two 'extreme' approaches to setting targets as classified in figure 2. It is perhaps for both of these reasons that some have found it difficult to understand which elements of the analysis were most important in informing the Stern Review's conclusions. As a result, it is often presumed that the conclusions rested solely on the formal economics of chapters 6, 9 and 10 (see e.g. Mendelsohn, 2008; Neumayer, 2007). An alternative and more unusual perspective is that the Review eschewed mainstream economic analysis altogether (see Barker, forthcoming).

Neither of these views correspond to my understanding. In seeking to clarify any confusion, the Stern Review team itself has argued that the foremost policy question is a disaggregated one: the cost of reducing emissions on any particular pathway to stabilisation is akin to an insurance premium and one should evaluate whether it is worth paying this premium, above all in terms of the disaggregated risks of climate change on for example food and water availability that can be avoided by doing so. As Dietz and Stern (2008, p4) argue: "This question is central, because it presents the basic policy problem as simply and transparently as possible, thereby avoiding the process of aggregating risks and uncertainties across all dimensions (e.g. nations, time, goods etc.), a process for which information and data are extremely thin and which ignores or suppresses so much of what is important".

This comparison must be made sequentially, martialling a range of evidence that cannot be brought together in a single model. The steps are approximately as follows:

1. What is the probability of different increases in global mean temperature accompanying stabilisation of the atmospheric stock of greenhouse gases at a particular level?
2. What are the impacts of climate change for a given increase in global mean temperature?
3. What is the cost of reducing greenhouse gas emissions to this stabilisation level?
4. What is the difference between the impacts of climate change on business-as-usual compared to stabilisation at this level, on multiple metrics (i.e. what are the benefits)?
5. Do the benefits of stabilisation outweigh the costs?
6. Repeat for a range of other stabilisation levels until the costs clearly outweigh the benefits and *vice versa*.

Nevertheless, aggregated economic models, in the form of optimising IAMs and macroeconomic models of the cost of emissions reductions, are useful for exploring particular, stylised aspects of the problem. With respect to the former models, factors can be explored such as the role of attitudes towards intergenerational equity and risk in estimating the cost of climate change. The latter models have more sectoral detail than optimising IAMs and are especially suited to investigating the role of behavioural changes in the economy as a whole in determining the cost of emissions reductions. Thus these tools of mainstream economics also informed the conclusions.

3. Academic responses

3.1. Discounting and intergenerational equity

The most fiercely contested aspect of the Stern Review's analysis has been its choice of social discount rate (this primarily concerns chapter 6, but also chapter 13). Because a large number of papers have already been written in criticism of this choice (notably Beckerman and Hepburn, 2007; Dasgupta, 2006; Nordhaus, 2007a; Weitzman, 2007) and in response[11], I will only provide a rather brief summary here.

The social (otherwise known as the consumption or 'shadow') discount rate determines the present value of flows of future consumption. It should be one of the central factors in formalised cost-benefit analysis of climate change (although see section 3.3), because the external costs of greenhouse gas emissions appear only after a significant lag, but then persist for a very long time.[12] Therefore broadly speaking it is the present generation that must pay the cost of reducing greenhouse gas emissions, but future generations that benefit.

One cannot proceed very far through the vigorous debate about discounting the benefits and costs of greenhouse gas emissions reductions without setting out the relatively simple 'Ramsey' formula for the social discount rate:[13]

$$s = \delta + \eta g \tag{1}$$

Where s is the social discount rate, δ is the rate of pure time preference or utility discount rate, η is the elasticity of marginal utility of consumption and g is the rate of growth of per-capita consumption. The two arguments on the right-hand side of equation (1) play distinctly different roles. The rate of pure time preference or utility discount rate, δ, discounts future utility simply because it accrues in the future, even if future generations are no better off than we are. It is discounting on the basis of time and time alone and it is not the same as the social discount rate; rather it is one element of it. Nevertheless this basic point seems to have confused Varian (2007) in his commentary on the Stern

[11] From the Stern Review team itself by Dietz *et al*. (2007a and 2007b), by Dietz and Stern (2008), by Stern (2008), and by Stern and Taylor (2007). Supportive independent commentaries have also been written by Quiggin (forthcoming) and Sterner and Persson (2008).

[12] In large measure due to the behaviour of CO_2 in the atmosphere. According to Yohe (2003), the marginal increase in temperature from the emission of one tonne of CO_2 is virtually zero at the point of emission, peaks around 50 years after the emission and continues for over 400 years.

[13] From Ramsey (1928).

Review. The product ηg, on the other hand, typically discounts future consumption on the grounds that we will be richer by then (due to positive g), and the marginal utility of an extra unit of consumption falls as consumption rises (this is captured by η). Note that the model underlying (1) tends to assume that the relationship between consumption and utility for an individual also gives the relationship for society, so that we should more accurately describe η as the elasticity of marginal *social* utility of consumption. However we might not be richer in the future, depending on how severe the impacts of climate change turn out to be. Thus in principle at least g could be negative over some period of time and thus the product ηg will correspondingly work towards negative discounting, increasing the present value of future consumption.

In the Stern Review, δ=0.1% per annum and η=1. Because the analysis took into account uncertainty about the impacts of climate change,[14] there was no single growth rate g and consequently no single social discount rate. Rather there was one social discount rate for each possible future state of the world. In addition, in each and every state of the world g was changing over time, so the instantaneous social discount rate was changing too. This point has eluded a number of commentators including Gollier (2006). However for the purposes of understanding to a first approximation what social discount rate these parameter choices give, the Stern Review team tends to report the average growth rate of consumption per capita over the 200-year modelling horizon of the PAGE model in a world without climate change. This is 1.3% and thus gives a social discount rate of 0.1+1*1.3=1.4%.

The source of the controversy lies of course in the fact that this rate is lower than most other social discount rates used in formalised cost-benefit analysis of climate change (with the exception of e.g. Cline, 1992) and lower than many, but not all, discount rates in use in public-sector economic appraisals worldwide. And the empirical evidence that we have suggests that this controversy does matter in setting climate targets. For instance, sensitivity analysis published by the Stern Review team in Dietz *et al*. (2007a) shows that the present-value cost of business-as-usual climate change falls by around 70% in the Review's central scenario[15] if δ is increased from 0.1% to 1.5%. Some well-known studies such as Nordhaus and Boyer (2000) use values of δ as high as 3% at least initially. While this does not directly tell us the effect on the optimal stabilisation target (since the Stern Review undertook no such analysis), we can safely assume it would rise significantly. Similarly, Tol (2007) shows in a

[14] In a Monte Carlo simulation comprising 1000 runs, where each run gives different impacts and consequently different consumption growth rates.
[15] The 'baseline'-climate scenario, with market impacts, non-market impacts and the risk of catastrophe.

meta-analysis of the literature on formalised cost-benefit analysis that the mean estimate of the social cost of carbon (the marginal damage cost of carbon) today increases from \$24/tC to \$317/tC as δ falls from 3% to 0%. Again, this does not directly tell us by how much the optimal stabilisation target would rise, but is highly indicative.

What happens when η varies is more complex. This is because the very standard welfare-economic model behind equation (1) actually asks η to play up to three separate roles, depending on the level of disaggregation. Because η captures the change in marginal utility as consumption changes, it can be applied to differences in consumption over time, as set out above, but also to differences in consumption over regions of the world and over uncertain states of the world (in which case it is the coefficient of relative risk aversion). In the early sensitivity analysis published by the Stern Review team (Dietz *et al.*, 2007a), increases in η from 1 to 2 were found to depress the present value of the cost of business-as-usual climate change in the central scenario by just under 70%. However subsequent work by Dietz *et al*. (2007c) showed that a further increase in η from 2 to 3 could actually drive the present-value cost of business-as-usual climate change back upwards in a hotter 'high'-climate scenario, such that at η=3 the cost is only fractionally lower than at η=1 (reproduced in table 2). Thus what the results tell us in general is that the role of η is ambiguous. At this point it is important to note that the Stern Review only had two dimensions to its analysis, time and uncertainty (or rather risk). The team ran out of time to disaggregate by regions and this is an obvious theme for further work: what would be the combined effect of variation in η if regional 'equity weighting' were also included (see also section 3.2)?

Table 2. Sensitivity of mean total cost of climate change (as a percentage of present global consumption per capita) to rate of pure time preference and elasticity of marginal utility of consumption in high-climate scenario. *Stern Review* estimate in bold. Reproduced from Dietz *et al*. (2007c).

Pure rate of time preference (δ)	*Elasticity of marginal utility of consumption (η)*				
	1.0	*1.5*	*2.0*	*2.5*	*3.0*
0.1	**14.7**	10.2	7.4	8.1	13.2
0.5	10.6	6.5	4.7	5.0	7.8
1.0	6.7	4.0	2.7	2.7	3.9
1.5	4.2	2.5	1.7	1.6	2.1

Returning to the disagreement over the social discount rate, although there are many nuances to this it has often proved helpful to categorise views as being either 'descriptive' or 'prescriptive'.[16] A descriptive approach sets the social discount rate to be consistent with related revealed preferences in the economy. This descriptive approach can focus on s, δ and/or η, because evidence is available directly or indirectly on all of these. Revealed preferences can be of individuals or of society. Examples of the latter include the design of national tax and social-security systems, which produce some evidence on η, and savings rates, which produce some evidence on s. It is usually interpreted to support a social discount rate higher than 1.4%, as both Dasgupta (2006) and Nordhaus (2007a) did in focussing on what the Stern Review's parameter choices would imply for savings rates. By contrast a prescriptive approach focuses on the parameters δ and η and sets them based on ethical first principles. What this implies for s, δ and η is *a priori* ambiguous, since a number of 'ethical' arguments can be made, but the signature argument is usually that δ should be nearly zero so as to be impartial between generations in terms of utility.

The arguments in favour of descriptive and prescriptive approaches have been well rehearsed since publication of the Stern Review and for a long time before, so I will not go into them here. However it is worth noting one further line of argument that has been pursued by Cline (2007) and also noted by Mendelsohn (2007) and by Nordhaus (2007a). Cline in particular notes that the combination of δ=0.1% and η=1 places a great deal of weight in the estimation of optimal policies on what happens very, very far into the future. The Stern Review calculated the cost of business-as-usual climate change over an infinite horizon. To do so it had to extrapolate costs beyond the end of the PAGE modelling horizon, which is 2200. Cline (2007) calculated that with δ=0.1% and η=1 around 93% of the present-value cost of business-as-usual climate change would fall after the end of the modelling horizon, essentially extrapolating. This is somewhat unsettling, because even if it were not necessary to extrapolate (i.e. even if the IAM had a longer horizon), estimates are profoundly uncertain so far into the future. Moreover this raises strongly the issue of learning (see below), which is omitted in simple analyses of benefits and costs when the policy-maker faces an immediate, once-and-for-all decision.

3.2. Equity weighting

[16] Due originally, to the best of my knowledge, to Arrow *et al*. (1996).

As discussed above, one missing element of the Stern Review's integrated modelling in chapter 6 was disaggregation by world regions. Due simply to time constraints, the analysis only considered the impacts of climate change on global welfare.[17] This is likely to be an important omission, because the impacts of climate change are usually estimated by IAM studies to be highest in the poorest regions of the world. If the elasticity of marginal utility of consumption η is positive, then these impacts should be given a higher weight in the analysis. This has come to be known in the literature on cost-benefit analysis of climate policy as 'equity weighting'.

Although the Stern Review did not formally carry out equity weighting, it did include in its 'headline' results a back-of-the-envelope estimate of by how much the present-value cost of business-as-usual climate change would increase if disaggregation by regions had also been possible. This crude estimate was based on the effect of equity weighting in Nordhaus and Boyer's well-known analysis of 2000, which was to increase costs by around one quarter. Applying a factor-of-one-quarter increase to the Stern Review's highest estimate of the present-value cost of business-as-usual climate change (14.4% of global consumption per capita), the Review concluded that the upper end of the range of costs of business-as-usual climate change lay at around 20% of global consumption per capita. Thus, like discounting, equity weighting is also highly consequential for the optimal climate policy.

However, like discounting, equity weighting has also had a controversial track record in the economics of climate change. The Stern Review's argument for making some effort to include equity weighting was characteristically based on first principles: since a positive elasticity of marginal utility of consumption is used in considering the distribution of consumption across time and across states of the world, consistency requires that it also be used in considering the distribution of consumption across regions. Conversely, Maddison (2007) and Mendelsohn (2007), in their commentaries on the Review, have raised a familiar argument against equity weighting. This has in fact been widely deployed in economics against the use of (intragenerational or spatial) equity weighting in cost-benefit analysis more generally (e.g. Mishan, 1981). It is that policies should not be justified on distributional grounds when it is possible to make direct transfers to compensate those who lose out from the policy. As Maddison notes, it may be cheaper at the margin to compensate the losers directly than to do so indirectly through greenhouse gas emissions abatement. For policies within individual nations and

[17] Some clarification is important here. The PAGE model does indeed disaggregate the impacts of climate change to eight regions of the world. However the subsequent analysis of the Stern Review took the aggregate of the impacts over these eight regions as the input into its expected-utility analysis.

generations, this argument may have some traction, since the necessary redistributional apparatus likely exists (e.g. the welfare state). However it is not clear to what extent it exists across regions and even if it were to exist, there is still a problem of consistency in what utility function is assumed for the purposes of discounting, compared with aggregating across regions. Some have looked at redistributive patterns across regions to make a different point against equity weighting. They have argued the pattern of aid giving by high-income countries to low-income countries suggests that aversion to inequality across regions is less than that implied even by η=1. This, as with the discussion on discounting, is a revealed-preference approach. It is vulnerable to the usual objections.

To conclude this section, equity weighting is also an important factor in determining the optimal stabilisation target and remains disputed territory. Since the Stern Review ran out of time to apply equity weights formally, this is an obvious task for further work. However because the effect of equity weighting clearly depends on the size and affluence of the population in different regions of the world, it is likely to bring assumptions about baseline socio-economic change (and adaptation – see section 3.6) back into sharp relief. Some of the assumptions contained in the Stern Review's integrated assessment, notably on the size of the population further into the future, do not appear to be particularly plausible (see section 3.7). Therefore these 'input' assumptions would also need to be revisited.

3.3. Is uncertainty about climate change too large for cost-benefit analysis?

Uncertainty is one of the key features of climate change, due to factors such as futurity and a poor understanding of complex (and sometimes chaotic) systems. In a recent think-piece that has attracted significant attention within economics, Martin Weitzman in effect asks whether the uncertainty is too great for integrated assessment?[18] Here I will sketch out the main features of his thesis and explain how some economists have responded.

Weitzman's paper has in fact the potential to be applied to any policy issue sharing certain key features, but he uses climate change as the motivating example. It takes as its starting point modelling estimates of climate sensitivity reported in the IPCC's *Fourth Assessment Report* (Chapters 9 and 10 of

[18] His arguments are fleshed out in most detail in a working paper, the latest version of which is dated February 8th 2008. I work on a version dated October 29th 2007, but it is not substantively different. His arguments were also published in shorter, preliminary form in the *Journal of Economic Literature* in 2007.

Working Group I). Climate sensitivity is defined as the equilibrium change in global mean surface air temperature for a particular change in radiative forcing and it is commonly estimated in response to a doubling in the (equivalent) atmospheric stock of CO_2. Weitzman calculates on the basis of the IPCC estimates that there is approximately a 4-5% probability of a climate sensitivity greater than 7°C. But in fact he has a broader definition of the climate sensitivity in mind, because he wants to include the effect of uncertain, large-scale positive feedbacks in the carbon cycle as well. The two that he explicitly mentions are the release of methane from melting Arctic permafrost and the release of offshore methane hydrates. Interpreting the paleoclimatic work of Torn and Harte (2006), Weitzman calculates that, if large-scale positive feedbacks are included, then a fuller, more generic definition of the climate sensitivity gives a 2% probability of an increase in global mean temperature of more than 10°C for a doubling of atmospheric CO_2. Since this implies warming unprecedented even in geological time, he postulates that there is a non-negligible probability of 'catastrophe'. What is meant by catastrophe is crucial in understanding his paper and I devote considerable time to it below and in box 1.

Weitzman goes on to explore what the existence of such a possible catastrophe implies for the standard (or 'workhorse') model of welfare economics, which underpins the vast majority of formalised cost-benefit analyses of climate change. The standard model maximises expected utility, where utility is additive across individuals without weighting (except for pure time preference) and utility is an iso-elastic function[19] of aggregate consumption. Uncertainty is dealt with using the standard economic framework for 'risk', where the value of every possible outcome is known along with its probability of occurrence (see box 1 for further explanation). Thus the structure of the utility function implies that the policy maker has constant relative risk aversion.[20] For more details on the standard expected-utility model see chapter 2 of the Stern Review.

How Weitzman postulates the existence of a catastrophe as a result of very high climate sensitivity is , as I have indicated, a critical point in the analysis. Weitzman does not uncover new evidence of the relationship between warming and economic costs at high temperatures. Rather he simply makes the assumption that with so much warming a catastrophe is possible. Thus there is what one might describe as a 'clean break' in his paper between his review of the science and his arguments about the economics. He proceeds directly from some motivating observations about the climate sensitivity and its consequences to a theoretical problem, in which the future probability distribution of economic

[19] That is, where the elasticity of marginal utility of consumption is constant.
[20] Which is also described by the elasticity of marginal utility of consumption.

costs from climate change cannot be given an upper bound. In other words, Weitzman does not rule out an infinite reduction in future consumption. Why he does not is a complex and nuanced point, so I explore it in slightly more detail in box 1 (for which some background knowledge of statistics is likely to be helpful).[21]

This clean break in the story has many consequences, one of which is that, because it is the probability distribution of climate damages he is ultimately interested in, not the probability distribution of the climate sensitivity (even in its generic form), his point would endure even if there were a scientific advance that succeeded in placing a physical upper bound on the generic climate sensitivity. That is to say, even if the generic climate sensitivity were bounded from above at, say, 12°C, Weitzman argues that the cost of climate change could at this point still be unbounded, because of all the additional uncertainties linking temperature changes with economic costs in their various forms.

The main point of Weitzman's paper is that if the problem is constructed in this way, then the economic cost of climate change is itself unbounded, i.e. infinity. This is an unsettling result, which he terms the 'dismal theorem'. It tells us loosely speaking that uncertainty about climate change has the capacity to overwhelm all other considerations in estimating the cost of climate change and thus optimal climate targets. The parameters of the social discount rate (see above) could be very high indeed and it would not change the result. Indeed, there are no positive values for the pure rate of time preference or the elasticity of marginal utility of consumption that are high enough. Once there is an infinity in the model, nothing else matters. The recommendation that would follow would be to pursue infinitely costly emissions reductions.

[21] It also leans heavily on work by John Geweke (2001).

Box 1. Expected-utility analysis when the probability distribution of the economic cost of climate change is unknown.

In order to perform expected-utility analysis, we have to assume that the decision-maker faces a situation that is mathematically equivalent to what economists would understand as pure 'risk'. In this situation, future outcomes – e.g. the economic cost of climate change – are uncertain, but the full set of possible outcomes is known, along with all the corresponding probabilities of occurrence. But there are very few if any situations in which we can perfectly predict the set of outcomes and probabilities. There may be some examples from the laboratory (Henry, 2006). Thus most analyses adopt a pragmatic position, where subjective predictions of future probabilities are taken to be similar to objective knowledge.

Weitzman works within this framework, in particular following the standard procedures of 'Bayesian' inference. Bayesian inference is a method of statistical inference in which new evidence is used to update one's initial estimate of probability. One starts with an estimate of probability and updates this estimate after new evidence comes to light, where the probabilities in question are usually (and in particular in this case) subjective in nature. The important steps in this process are estimating (1) a 'prior' distribution, before (new) evidence is considered, (2) a 'likelihood function' that loosely speaking sets out what the new evidence tells us about the probabilities, and (3) a 'posterior' distribution, which is the outcome of updating the prior distribution with the evidence.

Weitzman follows this very standard methodology. He begins with the assumption that we know the mean damages from future climate change but, motivated by the large uncertainty about the climate sensitivity and its long right-hand tail, we do not know the distribution. Instead we have a small number of scientific and economic studies (e.g. integrated-assessment studies) to draw upon. This 'sample' of studies gives us the properties of the likelihood function. The assumption that the mean is known could be relaxed, Weitzman claims but does not explicitly prove, without affecting the analysis, but requires greater complexity.

A crucial assumption in Bayesian inference concerns the prior distribution and all of Weitzman's arguments follow from here. Due to our lack of knowledge about the economic cost of climate change, motivated by his reading of the scientific evidence on climate sensitivity, Weitzman assumes a 'non-informative' prior distribution in the form of a generalised power law. Power law distributions have an extremely long tail and are useful in statistics to describe the frequency of occurrence of extremely rare events like stock market crashes. By using such a distribution as his prior, Weitzman is essentially saying that we are practically ignorant of the distribution of economic costs of climate change in the first instance.

Plugging this non-informative prior distribution into the analysis gives a posterior distribution for economic costs that has a so-called 'fat tail'. There are several definitions of fat tails in the literature, but what Weitzman means is that the probability of high consumption losses due to climate change declines less rapidly than the marginal utility of consumption increases. For any utility function with strictly positive relative risk aversion, the marginal cost of climate change is infinity if the marginal utility of consumption wins this 'convergence race', so-to-speak. Expected-utility analysis no longer applies because the marginal benefit of reducing greenhouse gas emissions is infinity.

In fact, this result was also highlighted in an earlier paper by Richard Tol (2003). Tol also pointed out that for cost-benefit analysis to be applicable, the uncertainty about the (marginal) costs and benefits of climate change has to be finite. However in an application of his IAM, FUND, to uncertainty, Tol found a 1/1000 probability of infinitely large marginal utility[22], even though the scenario was of gradual rather than catastrophic climate change. The run in question estimated a catastrophic shortage of water in the former Soviet Union, even though the climate sensitivity was not nearly as high as Weitzman asks us to consider. This is interesting because it constitutes a more direct form of evidence. Weitzman's paper gives us the general framework, by contrast. Similarly, in one of the Stern Review's integrated-assessment scenarios, the PAGE2002 model estimated a 1/500 probability of consumption losses of 100%, which would also produce infinitely high marginal utility.[23] This was also based on a climate scenario that exhibited less rapid warming than Weitzman suggests is possible, although because PAGE2002 is a more aggregated model, there is no direct interpretation of what the cause of these 100% losses was, compared with FUND.

One way to avoid the dismal theorem is to work with a probability distribution of climate damages that rules out unbounded damages *a priori*; i.e. is in some way truncated to less than infinity. It is this option that Weitzman explores conceptually later in his paper, by drawing an analogy with the concept of the value of statistical life. This is fairly intuitive, since our everyday human behaviour clearly betrays the fact that we place a finite value on our lives; otherwise we would avoid running a large number of infinitesimal risks of death (e.g. by never travelling to work).[24] It also means that we have some empirical evidence on the value of statistical life, on which we could draw to set the upper bound (e.g. Viscusi and Aldi, 2003). But there is at the same time a catch, because the concept that Weitzman needs to bound his model is not the value of an individual human's life, but rather the value of humanity or as he puts it "the value of statistical civilization as we know it" (p20). Nevertheless, as a first approximation Weitzman takes the value of statistical life (and thus of statistical civilization) to be one hundred times per-capita consumption. Plugging this into a particular utility function[25] delivers the result that the total economic cost of climate change should be truncated

---

[22] More precisely, one of the thousand Monte Carlo simulation runs produced infinite marginal utility.
[23] In the 'high'-climate scenario with market impacts, non-market impacts and the risk of catastrophe. See chapter 6 of the Stern Review for a description of the scenario. More data can be obtained from the author.
[24] As such the value of statistical life is analytically close to Aumann and Kurz's (1977) concept of 'fear of ruin'.
[25] An iso-elastic function where the elasticity of marginal utility of consumption is equal to two. For the analogy with the value of statistical life to work, this elasticity must be strictly greater than one, so the

at 99% of consumption. This is in fact exactly the point at which the cost of climate change was truncated in the Stern Review, although the latter choice was for convenience, rather than having benefited from Weitzman's work. In a similar vein, Gary Yohe (2003) responded to Richard Tol's (2003) earlier finding by suggesting that the infinity could be removed by implementing an emergency foreign aid programme, designed to maintain economic activity above subsistence levels everywhere. This would be plausible for impacts of a regional nature, since if the catastrophe is global there will be no funding for the aid, and is another way to think about placing an upper bound on costs, this time at the regional level.

However it is important to emphasise that truncating the distribution of climate damages in this way does not make the problems highlighted by Weitzman's analysis disappear. On the contrary it now becomes clear that the overall results of formalised cost-benefit analysis might depend in some important measure on where exactly the rather arbitrary upper bound is set. To see why this is so, consider figure 3, which sets out a standard iso-elastic utility function for someone who is risk-averse.[26] What becomes clear is that utility is very sensitive to changes in consumption for low levels of initial consumption. But it is precisely in this region where the upper-bound on climate damages must be placed. Thus very small changes in the upper bound, from, say, damages of a maximum of 99% of consumption to 99.9%, could have very large consequences indeed.

Weitzman's (2008) arguments have already attracted comment. Bill Nordhaus (2007b) examines what he sees to be the two most important assumptions underpinning Weitzman's dismal theorem, namely (1) that utility should be described by an iso-elastic function and (2) that the (posterior) probability distribution of economic costs of climate change has a fat tail (see box 1). With respect to the first of these assumptions, Nordhaus' view is not that Weitzman's analysis has undermined the whole endeavour of formalised cost-benefit analysis of climate change, but rather that it undermines the validity of the particular utility function that he uses, when consumption is very low. Where the parameter that describes the curvature of this particular class of utility function – the elasticity of marginal utility of consumption or coefficient of relative risk aversion – is greater than or equal to unity, utility tends to minus infinity as consumption tends to zero. But this brings with it the unattractive implication that willingness to pay to prevent an infinitesimal probability is infinite. Yet

---

common natural logarithmic utility function (the special case where the elasticity is equal to unity) cannot be used.

[26] The elasticity of marginal utility of consumption (the coefficient of relative risk aversion) is equal to two.

such willingness to pay is not observed in individual behaviour, nor is it in policy (Nordhaus gives the example of finite – and indeed rather low – spending to track hazardous asteroids, despite the infinitesimal extinction risk). Thus this argument is also in the end about an upper bound on the cost of climate change. Geoffrey Heal (2007) makes similar observations about this first assumption, but his point is more general than that of Nordhaus. Heal asks whether what Weitzman finds suggests that we should not be working from within expected-utility theory at all. After all, there is a peculiar contradiction to Weitzman's analysis, which Heal identifies, namely that it is assumed "we can work within the standard expected-utility framework which assumes known probabilities, even though the essence of [Weitzman's] problem is that the probabilities are unknown" (p14). Consequently Heal points to alternative approaches such as that of Henry and Henry (2002). However which if either of the basic theories is preferable remains very much an open question.

Figure 3. The utility function for a risk-averse policy-maker and an upper bound on the cost of climate change.

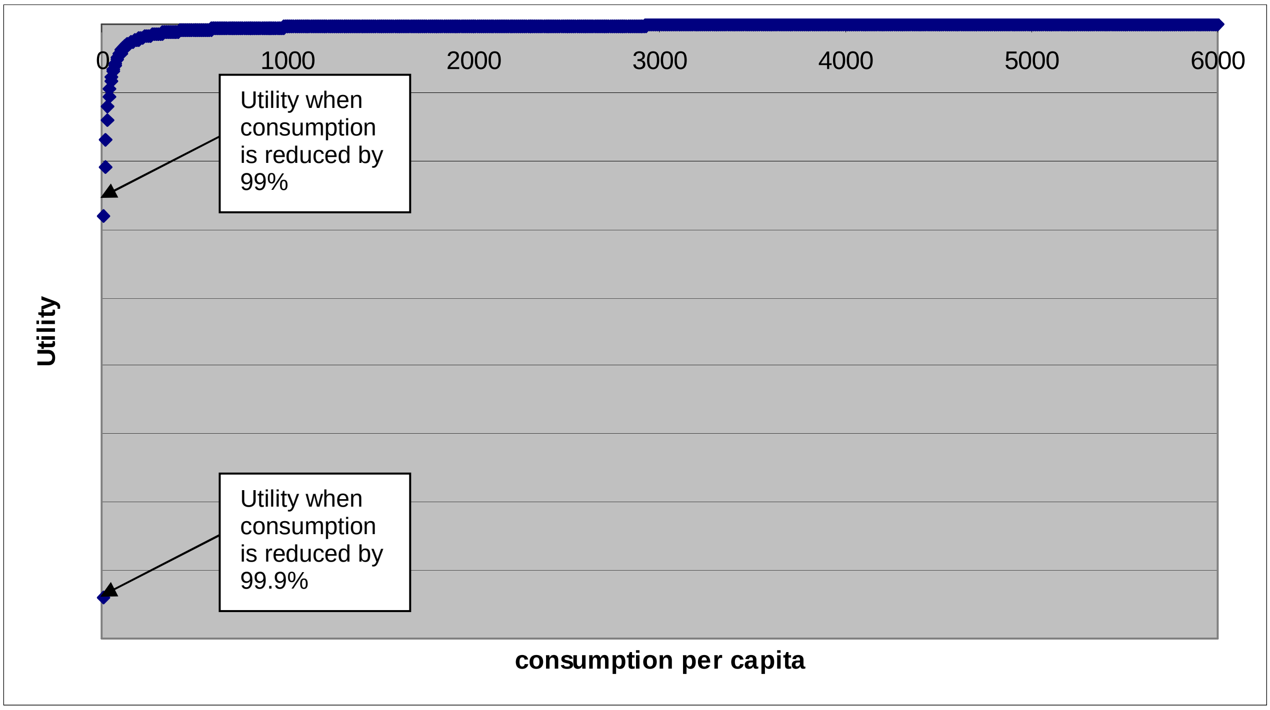

With respect to the second assumption, Nordhaus simply notes with some care that the choice of prior distribution in Bayesian inference (see box 1) is controversial and may assume too much ignorance, in essence. Moreover, Nordhaus raises what I believe to be another important question. That is, if what Weitzman describes as the dismal theorem holds for climate change, does it not also hold for other policy issues such as biotechnology, nuclear proliferation, asteroids and so on? If so then it is not clear what the policy maker should do but "dissolve in a sea of anxiety" (Nordhaus, 2007b, p7).[27]

Gary Yohe and Richard Tol (2007) have also responded. They follow up on Yohe's (2003) earlier arguments about aid to explore the role of learning. They point out that policy makers are not restricted to the form of 'observe and infer' learning that characterises Bayesian inference. Instead, they suggest that policy-makers can invest in scientific research that might provide us with a sort of 'early warning system' against catastrophe. Policy can then focus on avoiding the trigger points for climate catastrophes, if indeed that remains possible (otherwise, as they say and as Clarke and Reed (1994) originally demonstrated, we should make hay while the sun shines).

In summary, where does Weitzman's provocative contribution leave us? I would agree with him that "some sort of a tricky balance is required between being overawed by [the dismal theorem] into abandoning CBA altogether and being underawed by [the dismal theorem] into insisting that it is just another empirical issue to be sorted out" (p30). It seems that for climate economists work on two fronts is urgently required:

1. On theories of economic decision-making other than expected utility;
2. On further experiments with IAMs to explore as far as possible the long right-hand tail of the distribution of climate damages. This includes highly skewed estimated distributions of the climate sensitivity and variations in the upper bound on damages etc.

3.4. The worth of an ice sheet: treatment of 'catastrophic' risks

In an unpublished but interesting note on the Stern Review, Paul Baer (2007) picks up an issue strongly related to Weitzman's arguments. Baer focuses on how the PAGE model used by the Stern Review in chapters 6 and 13 treats catastrophic climate risks, in comparison with the wider scientific literature. He concludes that PAGE underestimates these risks (Weitzman draws this conclusion of the literature more generally). To understand his critique, it is clearly important to set out how PAGE models catastrophic risks.

[27] Nordhaus also points out that new research with the 2007 version of the DICE model never throws up the sorts of problems Weitzman has highlighted. However, since unbounded costs were uncovered by Tol (2003) with FUND and this author with PAGE, we should probably conclude that this is due to the particular features of DICE, rather than holding for other IAMs.

First, what is a climate catastrophe? Richard Posner (2004) defines a catastrophe as "an event that is believed to have a very low probability of materializing but that if it does materialize will produce harm so great and sudden as to seem discontinuous with the flow of events that preceded it". Dietz *et al*. (2007c) similarly offer an economic definition of 'dangerous' climate change that can be taken to be synonymous with catastrophic climate change:[28] "any process that produces losses in consumption-equivalent welfare, which are both rapid and large-scale relative to global trend consumption" (p313). As Dietz *et al*. go on to note, however, "[w]hat constitutes rapid and large-scale is arbitrary: there is no *a priori* basis for drawing the line between dangerous and 'non-dangerous' at any particular rate or level of welfare loss" (p313). Thus they suggest that it is more useful to analyse the pathways to catastrophic climate change and they suggest the following typology:

1. Rapid, large-scale impacts of gradual climate change (such as the 1/1000 probability discovered by Tol, 2003, in his FUND model);
2. Abrupt, discontinuous and large-scale positive natural feedbacks in the climate system that accelerate global warming (such as the melting of the Arctic permafrost);
3. Other abrupt, discontinuous and large-scale changes to the climate system that have more direct economic impacts.

The PAGE model has a separate module for catastrophes of the third type (which Hope calls 'large-scale discontinuities' in his 2006 description of the model) and this is what Baer focuses on. It is worth noting in passing that catastrophes of the first type can also be simulated by PAGE[29], while the Stern Review's 'high'-climate scenario is an example of the second type.

Examples of the third type of catastrophe include a regional or global shutdown of the Thermohaline Circulation, rapid melting of the Greenland and West Antarctic Ice Sheets, transformation of continental monsoons or modification of the El Niño Southern Oscillation. These non-linear changes would have direct economic impacts, in contrast to changes of the second type, which have indirect economic impacts through a major increase in the rate of warming. So far, only two IAMs have built

[28] The frequent interest in defining 'dangerous' climate change in the literature (e.g. Schellnhuber *et al*., 2006) is due to the wording of Article 2 of the United Nations Framework Convention on Climate Change, which states that atmospheric greenhouse gas concentrations should be stabilised such that 'dangerous anthropogenic interference' with the climate system is prevented (United Nations, 1992). Hence we need to find out what is a 'dangerous' stock of greenhouse gases in the atmosphere.

[29] To be more specific, this is achieved by estimating a very non-linear damage function (e.g. cubic) to describe the impacts of gradual warming on market and non-market sectors. With sufficient warming such a function effectively guarantees very high damages as a share of regional and global consumption, satisfying Dietz *et al*.'s definition above.

the possibility of these 'type-three' catastrophes into their core model structure. The pioneering estimates were included in the 1999 version of DICE (Nordhaus and Boyer, 2000). PAGE also includes the risk of catastrophe. When global mean temperature rises to a threshold level (a minimum of 2°C above pre-industrial, mode of 5°C, and maximum of 8°C), the chance of large losses in GDP in the range of 5-20% begins to appear (minimum 5%, mode 10%, maximum 20%). This chance increases by an average of about 10 percentage points per °C rise in global mean temperature (minimum 1, mode 10, maximum 20), so the probability is on average about 10% for a 6°C rise in global mean temperature above the pre-industrial level, c. 20% for a 7°C rise and so on:

$$d(t) = \zeta\theta(T(t) - T_{TRIGGER}) \qquad (2)$$

Where ζ is the loss in GDP if such a catastrophe occurs, θ is the probability of the catastrophe and $T_{TRIGGER}$ is the threshold global mean temperature, above which the catastrophe becomes possible. Thus the risk of direct impacts from these kinds of catastrophe is modelled as a subjective joint probability. The probability distributions were calibrated on information presented in the IPCC *Third Assessment Report* (TAR), although the TAR never addressed the issue directly. In particular, the TAR presented some evidence on $T_{TRIGGER}$, but almost none on θ and ζ, except for some order-of-magnitude quantification. In short, equation (2) is a genuine guesstimate, and this is highlighted by Maddison (2007), who emphasises the thin evidence base underpinning what turns out to be an important component of total climate damage (see e.g. the sensitivity analysis in Dietz *et al*., 2007a and c).

Baer (2007) focuses on the trigger or 'tipping point' for a catastrophe of the third type above and compares the assumptions of the PAGE model with the scientific literature on tipping points reported elsewhere in the Stern Review itself, using the melting of the Greenland Ice Sheet as an example (hence the title of his note and of this section). He notes that the joint probability of a catastrophe seems to be virtually zero in the PAGE model up to at least 3°C (see figure 4), yet the Review itself reports research by Lowe *et al*. (2006), which estimates a trigger temperature range of 2-3°C for the onset of irreversible melting of the Greenland Ice Sheet (reported in table 3.3 of the Review on p96). Other catastrophic changes such as substantial melting of the West Antarctic Ice Sheet are also reported by the Stern Review to be possible at under 3°C.

While Baer's is not a detailed analysis, it does contain a major point, especially given the significant body of scientific research that has recently been published and is now proceeding on tipping points (e.g. Lenton *et al*., 2008). This work can perhaps give us some initial sense of where the tipping points lie. Thus it seems worthwhile exploring the variation in optimal climate policies with changes in the probability distribution describing tipping points (i.e. the parameter $T_{TRIGGER}$). This seems to me to be the most worthwhile parameter of equation (2) to investigate further, because it is the only one for which we might have at least a little empirical evidence. However, a note of caution is necessary. Because of the characteristically aggregated way in which PAGE models catastrophes, $T_{TRIGGER}$ cannot be directly interpreted as a physical tipping point. Rather it describes the trigger point for economic costs as a result of a process like melting of the Greenland Ice Sheet. Thus one would need to make additional assumptions about the relationship between the physical tipping point and the magnitude and timing of economic costs. It is not clear whether we have much evidence to guide these assumptions.

It is also worth noting that the remaining two impact sectors in PAGE are in fact together more important in determining the total economic cost of climate change than are catastrophes. Figure 5 makes this clear as it overlays on a scatter plot of the cost of catastrophes with rising temperatures the cost of 'economic' impacts (i.e. 'market' impacts on e.g. agriculture) and 'non-economic' impacts (i.e. 'non-market' impacts on ecosystems and human health). For these two impact sectors, the form of the damage function is very important, as Hope (2006) and Dietz *et al*. (2007c) have shown in sensitivity analysis. Indeed, when the damage function is very non-linear (e.g. cubic) PAGE produces a climate catastrophe of type one above, which tends to be quantitatively more damaging than a catastrophe of type three. However it is doubtful whether there is a need for further sensitivity analysis based on the shape of the damage function, as we have practically no evidence to guide us on that shape, since we do not know what the economy-wide impacts of large amounts of warming would be.

Figure 4. The cost of a climate catastrophe in PAGE (from Warren *et al*., 2006).

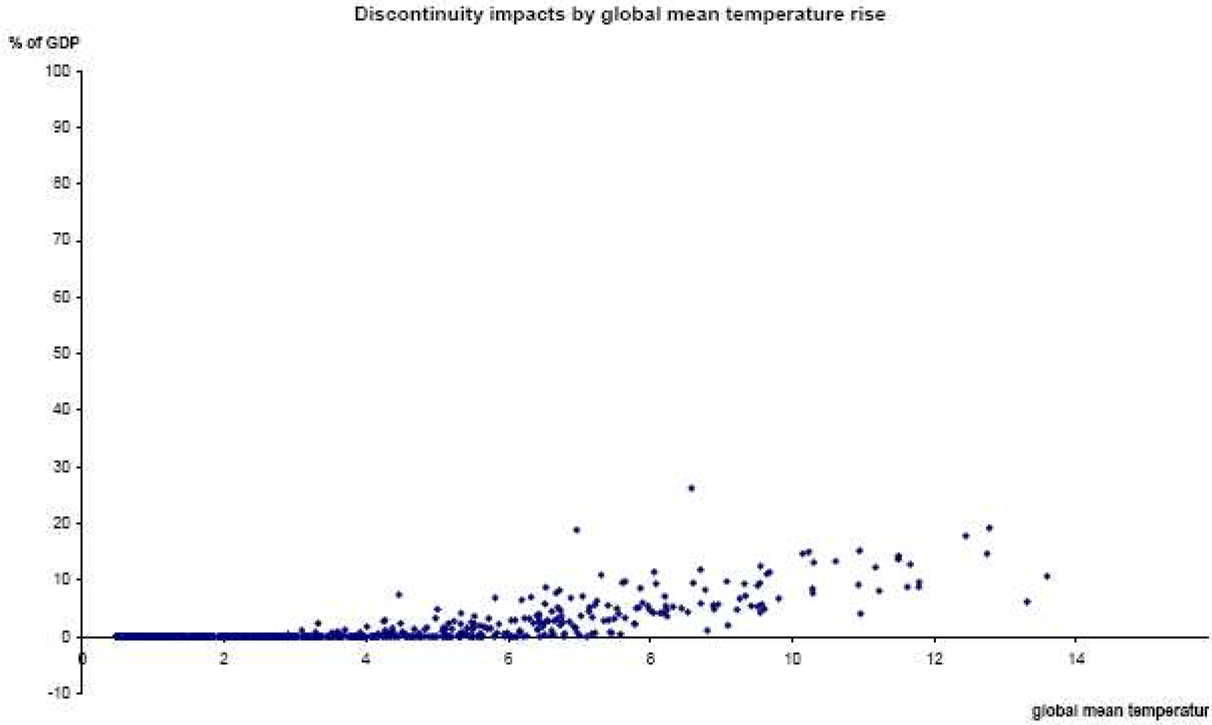

Figure 5. The total economic cost of climate change in PAGE, broken down into its three impact 'sectors' (from Warren *et al*., 2006).

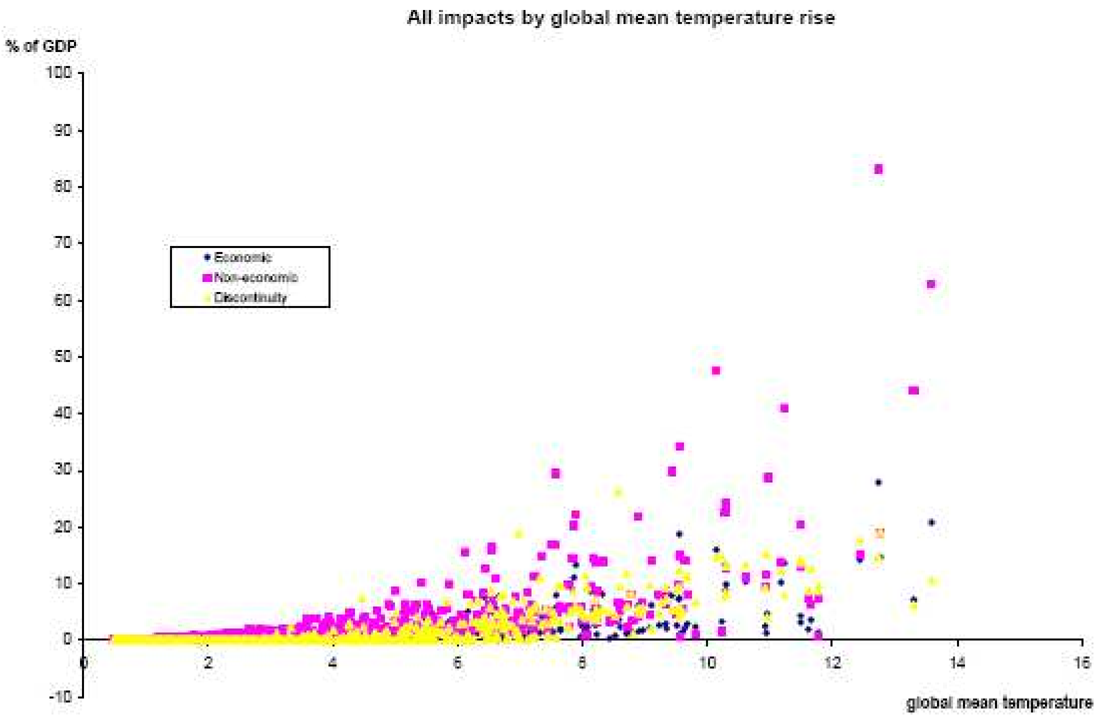

3.5. An even 'Sterner' Review: the role of the relative price of the environment

The standard welfare-economic model underpinning formalised cost-benefit analysis has a restrictive structure. I have already discussed this in relation to the triple role that must be played by the elasticity of marginal utility of consumption, η. However Sterner and Persson (2008) and Neumayer (2007) have reminded us of another way in which it is restrictive, essentially concerning the extent to which continually rising output of material goods and services can substitute for ongoing environmental degradation as a result of climate change.

Existing IAMs are based on a very simple and highly aggregated model of human welfare, in which utility depends only on the aggregate consumption of all goods and services. An estimate of material consumption is fairly easy to extract from the national accounts, being a portion of total output (the remainder is investment). The monetary value of environmental damage in all its forms due to climate change is calculated with an IAM and deducted from output. This output, net of the impact of climate change, is divided between consumption and investment. The former portion allows the cost of climate change to be calculated for a given year in terms of utility and welfare.[30] It can be seen from this that environmental damage is treated as if it is an additively separable component of total output. That is, one pound of environmental damage from climate change can be precisely offset by one pound of material consumption.

However this may well miss aspects of the problem. In an article amusingly titled "An even Sterner Review", Thomas Sterner and Martin Persson (2008) consider the implications of changing relative prices, as output of material goods and services (they give the contemporary example of mobile telephones) grows but environmental quality (e.g. biodiversity) becomes more scarce. The relative price of environmental quality can be expected to rise. This is another way of saying that material goods and services may not be *perfectly substitutable* with environmental goods and services, so that one pound of material consumption does not offset one pound of environmental damage from climate change. This not only concerns *final* consumption; it is also an issue for how production depends on inputs of material and environmental goods and services. All else equal, it increases the value in mitigating climate change and reduces the optimal atmospheric stock of greenhouse gases.

[30] The latter portion affects welfare in future years by eroding the accumulation of capital through investment.

Nevertheless as Neumayer has long argued[31] and as he (2007) and Sterner and Persson (2008) reaffirm, IAMs assume perfect substitutability. I will return to this issue as I believe it is more complicated than this bald statement suggests.

Sterner and Persson conduct an empirical analysis with the DICE model to test the implications of an increase in the relative price of the environment. They find that emissions reductions of the order of those recommended by the Stern Review can be supported even if the social discount rate is higher than in the Stern Review. In particular, optimal emissions are initially higher than in their approximation of the Stern Review analysis with DICE, but by the end of the century they are lower. Cumulative emissions over the 21st century appear to be only slightly higher in their optimal scenario with a higher social discount rate but with relative price changes (figure 6). Clearly the combination of Stern Review discounting and relative price changes would indeed be an 'even Sterner' stabilisation target.

Figure 6. Optimal carbon emissions with imperfect substitutability between material goods and services and the environment (from Sterner and Persson, 2008, p70). 'Nordhaus discounting' runs DICE (2007 version) in conventional mode, where the rate of pure time preference δ is 1.5% and the elasticity of marginal utility of consumption η is 2. 'Stern discounting' runs DICE where δ=0.1% and η=1. 'Relative price effects' is Nordhaus discounting but with a utility function that distinguishes between material consumption and consumption of environmental quality. The relative price of environmental quality rises.

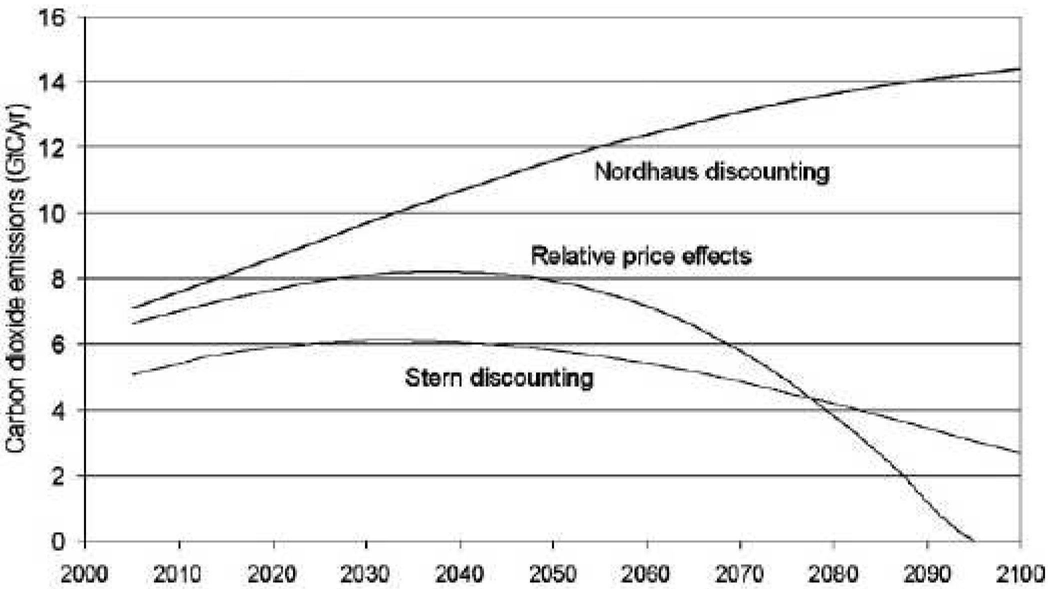

[31] Since *Weak Versus Strong Sustainability*, first published in 1999.

One or two notes of caution are again in order. First, as I hinted above, I do not think it is possible to claim without caveats that IAMs assume perfect substitutability. The problem is rather that they have nothing direct to say about substitutability at all, either in production (i.e. in the production function) or in consumption (i.e. in the utility function). It is true that, once estimated, one pound of extra material consumption exactly offsets the effect of one pound of environmental damage on human welfare. But the devil lies in the method of estimating how much environmental damage is caused by greenhouse gas emissions; in other words how the damage function is estimated. My point is that in models such as PAGE the damage function can be extremely non-linear, so that with enough warming the cost of climate change approaches 100% of global consumption. It is difficult in my view to interpret this in any other way than as being the result of depletion and degradation of non-substitutable environmental assets. One implication of this is that a model of relative price changes should not be grafted on to an existing IAM that estimates highly non-linear damages, otherwise the damages could in effect be partly double-counted. It is an open question whether this affects DICE and therefore Sterner and Persson's analysis.

Second, estimating relative price changes introduces further parameters into an IAM that are as difficult to calibrate as any other. In the Sterner and Persson (2008) paper, two parameters of significance are the contribution to current utility of environmental quality and the elasticity of substitution between material goods and services on the one hand and environmental goods and services on the other. It is very difficult to generate evidence on these parameters.

Nonetheless this is another important contribution to the economics of climate change. It is difficult to see what grounds one could have in principle to object to the model used by Sterner and Persson[32], because clearly at some level we cannot do without the environment altogether. In other words in the limit the environment is non-substitutable (see Dasgupta and Heal, 1979). But where does the limit lie? This is the empirical challenge of estimating a model with relative price changes. One could look at Sterner and Persson (2008) as another contribution to shifting the burden of proof towards those who doubt the case for strong action (I return to this point in section 4). In the context of IAMs, given that relative price changes are likely to be important, the case for working with non-linear damage

[32] And set out more formally in Hoel and Sterner (2007).

functions (i.e. quadratic and upwards) is reinforced. On the other hand if relative price changes are explicitly modelled then one has to reconsider the shape of the damage function.

3.6. Adaptation

A number of commentaries on the Stern Review have focused on the treatment of adaptation to climate change (notably Maddison, 2007; Mendelsohn, 2007; Tol and Yohe, 2006). Mendelsohn (2007) considers that the assumptions about adaptation made in chapters 3, 4 and 5 are too pessimistic and Maddison (2007) appears to make a similar point. However of more interest to this discussion are the specific criticisms that Tol and Yohe (2006) make about the treatment of adaptation in PAGE, which as usual applies to chapters 6 and 13.

Adaptation is explicitly modelled in PAGE, which sets it apart from other IAMs like DICE. In these other IAMs, adaptation is an implicit, inseparable component of the damage function, so it is difficult to establish – and even more difficult to vary – assumptions about adaptation. It is not even clear to what extent adaptive costs are measured and included in the cost-benefit analysis (Dietz *et al*., 2007a), even though such costs are certainly a cost of climate change and, insofar as they can be avoided, they are a benefit of mitigation.

In PAGE, investment in adaptive measures can increase the rate and level of warming before costs are incurred (essentially the amount of 'free' climate change) and can also reduce costs above this level, both in market and non-market sectors (but not due to a climate catastrophe[33]). Adaptive capacity is thus expressed first as a 'tolerable' rate and level of warming and second in terms of the percentage of market and non-market costs that can be reduced by adaptation above the tolerable level. In industrialised regions of the world, the assumption in PAGE is that adaptation allows a tolerable rate of warming in market sectors of the economy of 1°C per decade up to a tolerable level of 2°C. Beyond this, adaptation reduces market impacts by a constant 90%[34], irrespective of warming. In lower-income regions – Africa, India and Southeast Asia, and Latin America – there is no tolerable level of climate change (i.e. all warming is costly). Adaptation reduces market impacts by 50%[35], irrespective of

[33] i.e. a large-scale discontinuity as set out in section 3.4.
[34] In fact the percentage of costs avoided rises from 18% to 90% over the first ten years of the model, from 2001 to 2010. Thereafter it is constant at 90% until the end of the modelling horizon in 2200.
[35] As above, this percentage rises from 10% in 2001 to 50% in 2010 and is then constant over the remaining 190 years to 2200.

warming. There is no tolerable level of warming for non-market impacts. Worldwide, they are reduced by 25% in value terms through adaptation, again at all levels of warming. This lower non-market adaptive capacity partly reflects the insuperable challenges facing natural systems in adapting to rapid shifts in habitat conditions. Adaptation is highly cost-effective in PAGE, coming at a fraction of the cost of avoided impacts as global mean temperatures rise to higher levels.

Tol and Yohe's (2006) contention is that vulnerability to climate change – roughly the opposite to adaptive capacity – is proportional to development and therefore on this basis adaptive capacity should increase in developing regions of the world as they develop, rather than remaining at a constant 50%. It follows that in their view the cost of business-as-usual climate change in the Stern Review should be lower, because more adaptation should be assumed, especially in the medium to long term. However this criticism perhaps misses to some extent the point that in PAGE adaptive capacity is measured in relative terms. The percentage impacts of climate change that we can adapt to are constant, but since the impacts of climate change rise under business-as-usual, so does the absolute contribution of adaptation. At most there does seem to be a case for sensitivity analysis to allow adaptive capacity in currently developing regions to converge to that in developed regions, and Dietz *et al*. (2007a) do precisely that. This change in assumption has a small effect in relation to similar sensitivity analysis around, for example, the discount rate. Thus there does not seem to be a case for focusing scarce resources on further work in this particular area.

3.7. Business-as-usual socio-economic change and emissions

A set of assumptions made across the Stern Review's analysis of climate targets that has rightly generated interest is about the business-as-usual path of population, economic growth, technology and emissions of greenhouse gases. These assumptions were made jointly to construct internally consistent scenarios of future socio-economic change by the IPCC in its *Special Report on Emissions Scenarios* from 2000 (SRES: Nakicenovic and Swart, 2000). The SRES is a very common source of assumptions for analyses of climate targets and the Stern Review drew upon it heavily. In chapter 3, which set out the disaggregated impacts of climate change worldwide, results were reported for a number of SRES scenarios (because this was the strategy pursued by the 'FastTrack' project on which the chapter drew heavily). However chapter 6 relied upon the A2 scenario in particular and assumptions in other parts of the Review, notably in chapter 9 on the 'bottom-up' cost of emissions reductions, are also consistent with A2.

Compared with the full set of SRES scenarios, A2 gives the second highest atmospheric concentration of CO_2 in 2100. Since the SRES stated that all of its scenarios should be considered equally likely, this could be interpreted as a source of potential bias. That it has tended not to be, at least in terms of emissions, is indicative of the fact that recent data on actual emissions has been even higher than the A2 scenario predicted (see the work of the Garnaut Review in Australia). Indeed, emissions growth this decade has outstripped even the prediction of the A1FI scenario, a high-growth, fossil-fuel intensive scenario that gives the highest emissions of all the SRES scenarios. Thus A2 looks at present to be if anything slightly conservative.

However there are other aspects of the A2 scenario that are more concerning. Of all the SRES scenarios, A2 predicts the highest population growth (a global population of roughly 15 billion by 2100) and GDP growth that is not especially high in comparison with population growth. Thus per-capita GDP growth is relatively low. There are a number of consequences of this. First, baseline assumptions about population and economic growth have been shown to greatly affect vulnerability to climate change.[36] Roughly speaking, the more low-income people there are living in a region geographically vulnerable to climate change, the higher will be the impacts. Thus A2 looks particularly pessimistic about the cost of climate change in this respect. Note, however, that this feature does not directly affect adaptive capacity in the formal integrated modelling of chapter 6, because in PAGE assumptions about adaptive capacity are independent of the baseline scenario of socio-economic change (see the previous section). Second, the very high population in A2 will affect the formal estimation of economic welfare, because utility losses due to climate change are added across all individuals affected. The higher the population, the more individuals are affected, the higher are the utility losses (i.e. this is essentially weighting by population). This is a particular concern in PAGE, because SRES estimates of socio-economic change, which end in 2100, must be extrapolated to 2200. Chris Hope (2006) does this by taking the global average growth rate of, in this case, population between 2080 and 2100 and extrapolating it to 2200. For population in the A2 scenario, this gives a global population in 2150 of 17.6 billion (table 3) and rising[37], which is much higher than the most recent medium-scenario estimate issued by the UN[38] (9.7 billion), though lower than the high-scenario estimate (24.8 billion). Thus there is a case for reconsidering the formal analysis of chapter 6 with a

[36] The results of the FastTrack project, used extensively in chapter 3 of the Stern Review, show this very clearly.

[37] In 2200 it is 21.5 billion. However the UN does not issue comparable projections for 2200.

[38] http://www.un.org/esa/population/publications/longrange/longrange.htm

different scenario of business-as-usual socio-economic change. Alberth and Hope (2007), for instance, use the post-SRES 'Common Poles Image' scenario of den Elzen *et al*. (2003).

Table 3. Population (billions) in the A2 scenario, extrapolated from 2100 to 2200 in Hope (2006), compared with UN Long-Range Projections.

Region	*2001*	*2002*	*2010*	*2020*	*2040*	*2060*	*2080*	*2100*	*2150*	*2200*
IPCC SRES A2										
EU	0.40	0.40	0.42	0.43	0.45	0.46	0.49	0.54	0.65	0.80
EE	0.60	0.61	0.63	0.65	0.70	0.78	0.89	1.01	1.23	1.50
US	0.30	0.31	0.33	0.36	0.42	0.48	0.56	0.67	0.82	1.00
CA	1.31	1.33	1.45	1.60	1.93	2.26	2.64	2.99	3.66	4.46
IA	2.03	2.07	2.37	2.74	3.43	3.97	4.28	4.33	5.29	6.46
AF	0.82	0.84	1.05	1.33	1.92	2.48	2.87	3.04	3.71	4.54
LA	0.51	0.52	0.60	0.70	0.92	1.15	1.38	1.62	1.98	2.42
OT	0.20	0.20	0.21	0.21	0.22	0.22	0.23	0.24	0.30	0.36
Total	6.19	6.28	7.05	8.01	9.98	11.80	13.35	14.45	17.64	21.54
UN										
Low	6.03	-	-	-	-	-	-	5.15	3.23	-
Medium	6.06	-	-	-	-	-	-	9.46	9.75	-
High	6.09	-	-	-	-	-	-	16.18	24.83	-

4. Conclusions: towards a set of sufficient conditions for tough targets

Recommendations on long-term climate targets depend on assumptions. This report has set out the assumptions that underpinned the Stern Review's analysis and has compared them with other assumptions that have been made, notably in the long-standing literature on formalised cost-benefit analysis. That literature has tended not to support targets for the atmospheric stock of greenhouse gases as low as 550 ppm CO_2e (although there have been exceptions) and the strong criticism of the Stern Review, in particular by Robert Mendelsohn, William Nordhaus and Richard Tol, shows that there remains disagreement on long-term targets and by implication (but with additional complications) on near-term targets.

In this context, the contribution of the Stern Review has been cast in terms of shifting the burden of proof on to those who deny the need for immediate and ambitious reductions in greenhouse gases (e.g. Heal, 2007). And while scholars such as Kenneth Arrow (Arrow, 2007) and Martin Weitzman continue to disagree with some of the assumptions made in the Stern Review (notably on the social discount rate), they have concluded that the Review is 'right for the wrong reasons'.[39] Thus, in his review of the literature, Geoffrey Heal puts forward a set of sufficient conditions to support ambitious long-term climate targets. I believe this is a very good way to summarise the economics of climate change eighteen months on from publication of the Stern Review and I present my own version in table 4, which lays out five key assumptions determining optimal emissions reductions. The table should be interpreted with care, as it attempts to compare empirical studies in some cases with commentaries in others. Therefore some of the classifications are imprecise. In addition, some economists have themselves shown how different sets of assumptions can support different conclusions, so it will no doubt be possible to show that my interpretation contradicts some of their work. On this point, I have tried to summarise what I take to be each economist's overall view.

The first assumption is on vulnerability; essentially one's expectations of human and environmental capacity to adapt to future climate change. The second is the rate of pure time preference, δ, explained in section 3.1. The third is the elasticity of marginal utility of consumption, η, explained in sections 3.1 and 3.2. The fourth is explicit modelling of the change in the relative price of environmental goods and services, as they become relatively more scarce due to continuing emissions of greenhouse gases and resulting climate change. This can also be interpreted as a view on the importance of the environment to production in the economy and to human welfare. Hence I interpret some economists to place significant emphasis on this, without having modelled the process explicitly. The fifth is uncertainty and captures the extent to which focus is placed on the whole distribution of costs of climate change, including the tails.

As for example Dietz *et al*. (2007b) have argued, the case for strong action in the Stern Review was based on low discounting and a framework that placed strong emphasis on uncertainty. In Weitzman (2008), the case for strong action is built on uncertainty, coupled with the possibility of catastrophic outcomes. In Sterner and Persson (2008) and in Neumayer (2007), it is built on high vulnerability and the relative scarcity of environmental quality. While it would be unfair to accuse them of denying the importance of uncertainty, they present their assumptions with the confidence of a best guess. By contrast, Nordhaus (2008) and Mendelsohn (2007) do not support such strong action and their cases are built on high discounting and a framework that places less emphasis on uncertainty. In addition,

[39] This moniker has become a popular criticism of the Stern Review but it is important to distinguish between those such as Weitzman who think the Review's proposed climate targets are about right, and those such as Richard Tol who think the Review is right only insofar as it proposes to reduce greenhouse gas emissions by some amount.

Mendelsohn is optimistic about adaptation to climate change, estimating low vulnerability in particular in market sectors. Tol's work has tended to assume high adaptive capacity and he has argued strongly against the Stern Review's assumptions about the rate of pure time preference. However he has also published work exploring uncertainty and his analyses occasionally support strong action.

Table 1. Sufficient conditions for tough targets.

Author	Stern	Nordhaus	Weitzman	Sterner and Persson	Tol	Mendel-sohn	Neumayer
Vulnerability	Medium	Medium-high	Very high	High	Low	Low	High
δ	Low	High	High	High[40]	High	High	High
η	Low	High	High	High	Low	High	High
Relative prices	Not directly	Not directly	Not directly	Yes	No	No	Yes
Uncertainty	High	Low	Very high	Low	High	Low	Low
Strong action	Yes	No	Yes	Yes	Perhaps	No	Yes

5. Building on the Stern Review: research priorities for setting long-run targets

In this section I offer some suggestions on how we can build on the Stern Review's assessment of long-run climate targets. First, it is useful to take stock of the integrated-assessment literature and set out important new modelling avenues. In this context, perhaps the most important recent contribution has been that of Weitzman (2008). As he argues and as Yohe and Tol (2007) also note, the nature of the uncertainty about the future costs of climate change dictates that much more modelling effort is invested into exploring the tails of the distribution. Thus far, most of the effort has been on the body of the distribution, around our 'best guess' of future events. There are a number of practical steps that can be taken to do so:

- Investigate the sensitivity of the optimal trajectory to small shifts in the upper bound on total costs of climate change;

[40] Note that Sterner and Persson actually support the Stern Review assumptions on δ and η, but offer an alternative way to support similar targets.

- Investigate scenarios, grounded in scientific evidence to the greatest possible extent[41], that emphasise high impacts with low probabilities. Examples include probability distributions for the climate sensitivity with a long right-hand tail and tipping points for catastrophic events;
- Consider theories of decision-making other than expected-utility maximisation, such as that set out by Henry and Henry (2002).

In addition, it may be worth exploring the role of the changing relative price of environmental goods and services, although for any IAM that has not already been adjusted to do so, this may involve significant work. Finally, all future work should take into account the differential impacts of climate change across regions of the world and conduct sensitivity analysis to different degrees of equity weighting.

Second, the work of the CCC Secretariat should not be and indeed is not confined to formalised cost-benefit analysis and therefore there is a broader question of whether the Stern Review's analysis can be advanced. The structure of this analysis was set out in section 2. Aside from updating the analysis of the various relevant chapters with new data[42], can the CCC Secretariat narrow down the Stern Review's target range? I am not optimistic that this can be done, in general because (i) the overall uncertainties have not reduced and (ii) value judgements remain disputed. Therefore recommendations on targets will for the time being continue to be predicated on assumptions that can be contested.

There is one way in which new economic and scientific evidence could narrow the range, and this lies in identifying non-linearities in the incremental costs and benefits of changes in the long-run target. The difficulties in comparing incremental changes in costs and benefits are acute if the magnitudes are small, because such a comparison requires aggregation, whether implicit or explicit, and aggregation requires assumptions to be made about poorly understood or fiercely contested parameters. Such difficulties diminish, however, if a strong non-linearity can be found. If, for example, it could be ascertained that a number of tipping points for globally catastrophic climatic processes are crossed at a certain temperature, then the case for attempting to avoid crossing this threshold is much

---

[41] It is easy to generate catastrophic impacts in an IAM like PAGE, for example by manipulating the damage function. What is much more difficult is adducing the basis for doing so.

[42] For instance, chapters 9, 10 and 13 reveal the paucity of data on the cost of stabilising the atmospheric stock of greenhouse gases at below 550 ppm CO_2e.

strengthened. Similarly, if the costs of emissions reductions can be shown to escalate very rapidly at some level of ambition, this level may be easier to rule out.

References

Ackerman, F. and I. J. Finlayson (2006). "The economics of inaction on climate change: a sensitivity analysis." Climate Policy **6**.

Alberth, S. and C. Hope (2007). "Climate modelling with endogenous technical change: Stochastic learning and optimal greenhouse gas abatement in the PAGE2002 model " Energy Policy **35**(3): 1795-1807.

Arrow, K. J. (2007). "Global climate change: a challenge to policy." Economists' Voice **June**: 1-5.

Aumann, R. J. and M. Kurz (1977). "Power and taxes." Econometrica **199**: 1137-1161.

Azar, C. and K. Lindgren (2003). "Catastrophic events and stochastic cost-benefit analysis of climate change." Climatic Change **56**: 245-255.

Baer, P. (2007). The worth of an ice sheet: a critique of the treatment of catastrophic impacts in the Stern Review. Ecoequity discussion paper, Ecoequity.

Barker, T. (Forthcoming). "The economics of avoiding dangerous climate change." Climatic Change.

Beckerman, W. and C. J. Hepburn (2007). "Ethics of the discount rate in the Stern Review." World Economics **8**(1): 187-210.

Ceronsky, M., D. Anthoff, C. Hepburn and R. S. J. Tol (2005). Checking the price tag on catastrophe: the social cost of carbon under non-linear climate response. Working Paper FNU-87. Hamburg, Hamburg University and Centre for Marine and Atmospheric Science, Research Unit Sustainability and Global Change.

Cline, W. (2007). Evaluating the Stern Review on the Economics of Climate Change, Center for Global Development and Peterson Institute for International Economics.

Cline, W. R. (1992). The Economics of Global Warming. Washington, Institute for International Economics.

Dasgupta, P. (2006). Comments on the Stern Review's Economics of Climate Change, University of Cambridge.

Dasgupta, P. and G. Heal (1979). Economic Theory and Exhaustible Resources. Cambridge, Cambridge University Press.

den Elzen, M. G. J., M. M. Berk, P. Lucas, B. Eickhout and D. P. van Vuuren (2003). Exploring climate regimes for differentiation of commitments to achieve the EU climate target. RIVM report 728001023/2003. Amsterdam, RIVM.

Dietz, S., C. Hope, N. Stern and D. Zenghelis (2007a). "Reflections on the Stern Review (1): a robust case for strong action to reduce the risks of climate change." World Economics **8**(1): 121-168.

Dietz, S., D. Anderson, N. Stern, C. Taylor and D. Zenghelis (2007b). "Right for the right reasons: a final rejoinder on the Stern Review." World Economics **8**(2): 229-258.

Dietz, S., C. Hope and N. Patmore (2007c). "Some economics of 'dangerous' climate change: reflections on the Stern Review." Global Environmental Change **17**(3-4): 311-325.

Dietz, S. and N. Stern (2008). "Why economic analysis supports strong action on climate change: a response to the Stern Review's critics." Review of Environmental Economics and Policy **2**(1): 94-113.

Geweke, J. (2001). "A note on some of the limitations of CRRA utility." Economics Letters **71**: 341-345.

Gollier, C. (2006). Institute Outlook: Climate Change and Insurance: An Evaluation of the Stern Report on the Economics of Climate Change, Barbon Institute.

Heal, G. (2007). Climate change economics: a meta-review. Mimeo.

Henry, C. and M. Henry (2002). Formalization and applications of the precautionary principle. Discussion Paper 2002009. Louvain, Belgium, Université catholique de Louvain, Institut de Recherches Economiques et Sociales.

Hope, C. (2005). The economics of climate change. House of Lords Select Committee on Economic Affairs, Volume II: evidence. London, The Stationery Office.

Hope, C. (2006). "The marginal impact of CO2 from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern." Integrated Assessment **6**(1): 19-56.

Lenton, T. M., H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf and H. J. Schellnhuber (2008). "Tipping elements in the Earth's climate system." Proceedings of the National Academy of Sciences **105**(6): 1786-1793.

Lowe, J., J. M. Gregory, J. Ridley, P. Huybrechts, R. J. Nicholls and M. Collins (2006). The role of sea-level rise and the Greenland ice sheet in dangerous climate change: implications for the stabilisation of climate. Avoiding Dangerous Climate Change. H. J. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley and G. W. Yohe. Cambridge, Cambridge University Press: 29-36.

Maddison, D. J. (2007). Further comments on the Stern Review. Mimeo. Birmingham, UK, University of Birmingham.

Mendelsohn, R. O. (2006). "A critique of the Stern Report." Regulation **Winter**: 42-46.

Mendelsohn, R. O. (2008). "Is the Stern Review an economic analysis." Review of Environmental Economics and Policy **2**(1): 45-60.

Mishan, E. J. (1981). Introduction to Normative Economics. New York, Oxford University Press.

Nakicenovic, N. and R. Swart, Eds. (2000). Special Report on Emissions Scenarios: a Special Report of Working Group III of the Intergovernmental Panel on Climate Change. Cambridge, Cambridge University Press.

Neumayer, E. (2007). "A missed opportunity: the Stern Review on climate change fails to tackle the issue of non-substitutable loss of natural capital." Global Environmental Change **17**(3-4): 297-301.

Nordhaus, W. D. (2007a). "A Review of The Stern Review on the Economics of Climate Change." Journal of Economic Literature **45**(3): 686-702.

Nordhaus, W. D. (2007b). The real meaning of Weitzman's dismal theorem. Mimeo. New Haven, CT, Yale University.

Nordhaus, W. (2008). A Question of Balance: Weighing the Options on Global Warming Policies. New Haven and London, Yale University Press.

Nordhaus, W. D. and J. Boyer (2000). Warming the World: Economic Models of Global Warming. Cambridge, Mass., MIT Press.

Posner, R. A. (2004). Catastrophe: Risk and Response. Oxford, Oxford University Press.

Quiggin, J. (Forthcoming). "Stern and the critics on discounting." Climatic Change.

Ramsey, F. P. (1928). "A mathematical theory of saving." Economic Journal **38**(543-559).

Roughgarden, T. and S. H. Schneider (1999). "Climate change policy: quantifying uncertainties for damages and optimal carbon taxes." Energy Policy **27**: 415-429.

Schellnhuber, H. J., W. Cramer, N. Nakicenovic, T. Wigley and G. W. Yohe (2006). Avoiding Dangerous Climate Change. Cambridge, UK, Cambridge University Press.

Stern, N. and C. Taylor (2007). "Climate change: risk, ethics and the Stern Review." Science **317**: 203-204.

Sterner, T. and U. M. Persson (in press). "An even Sterner Review: introducing relative prices into the discounting debate." Review of Environmental Economics and Policy.

Tol, R. S. J. (2003). "Is the uncertainty about climate change too large for expected cost-benefit analysis." Climatic Change **56**: 265-289.

Tol, R. S. J. (2007). "Europe's long term climate target: a critical evaluation." Energy Policy **35**(1): 424-434.

Tol, R. S. J. (2007). The social cost of carbon: trends, outliers and catastrophes. E-conomics Discussion Papers 2007-44.

Tol, R. S. J. and G. W. Yohe (2006). "A review of the Stern Review." World Economics **7**(4): 233-250.

Torn, M. S. and J. Harte (2006). "Missing feedbacks, asymmetric uncertainties, and underestimation of future warming." Geophysical Research Letters **33**: L10703.

United Nations (1992). United Nations Framework Convention on Climate Change. Bonn, Germany, United Nations.

Varian, H. R. (2007). Recalculating the Costs of Global Climate Change. New York Times, December 14th.

Viscusi, W. K. and J. E. Aldy (2003). "The value of a statistical life: a critical review of market estimates throughout the world." Journal of Risk and Uncertainty **5**: 5-76.

Warren, R., C. Hope, M. Mastrandrea, R. Tol, N. Adger and I. Lorenzoni (2006). Spotlighting impacts functions in integrated assessment: research report prepared for the Stern Review on the Economics of Climate Change. Working Paper 91, Tyndall Centre for Climate Change Research.

Watkiss, P., W. Blyth, S. Dietz, T. E. Downing, A. Hunt and K. Fletcher (2008). Review of the Methodological Approaches Available for Setting UK Carbon Budgets. Oxford, Paul Watkiss Associates.

Weitzman, M. L. (2008). The role of uncertainty in the economics of catastrophic climate change. Mimeo. Cambridge, MA, Harvard University, Department of Economics.

Yohe, G. W. (2003). "More trouble for cost-benefit analysis." Climatic Change **56**(3): 235-244.

Yohe, G. W. and R. S. J. Tol (2007). Precaution and a dismal theorem: implications for climate policy and climate research. Working Paper FNU-145. Hamburg, Hamburg University and Centre for Marine and Atmospheric Science.