



Centre for
Climate Change
Economics and Policy

The Munich Re Programme: *Evaluating the Economics
of Climate Risks and Opportunities in the Insurance Sector*



Grantham Research Institute on
Climate Change and
the Environment

High impact, low probability? An empirical analysis of risk in the economics of climate change

Simon Dietz

September 2009

Centre for Climate Change Economics and Policy
Working Paper No. 10

Munich Re Programme Technical Paper No. 2
Grantham Research Institute on Climate Change and
the Environment

Working Paper No. 9

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council and has five inter-linked research programmes:

1. Developing climate science and economics
2. Climate change governance for a new global deal
3. Adaptation to climate change and human development
4. Governments, markets and climate change mitigation
5. The Munich Re Programme - Evaluating the economics of climate risks and opportunities in the insurance sector (funded by Munich Re)

More information about the Centre for Climate Change Economics and Policy can be found at: <http://www.cccep.ac.uk>.

The Munich Re Programme is evaluating the economics of climate risks and opportunities in the insurance sector. It is a comprehensive research programme that focuses on the assessment of the risks from climate change and on the appropriate responses, to inform decision-making in the private and public sectors. The programme is exploring, from a risk management perspective, the implications of climate change across the world, in terms of both physical impacts and regulatory responses. The programme draws on both science and economics, particularly in interpreting and applying climate and impact information in decision-making for both the short and long term. The programme is also identifying and developing approaches that enable the financial services industries to support effectively climate change adaptation and mitigation, through for example, providing catastrophe insurance against extreme weather events and innovative financial products for carbon markets. This programme is funded by Munich Re and benefits from research collaborations across the industry and public sectors.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training in climate change and the environment. The Institute is funded by the Grantham Foundation for the Protection of the Environment, and has five research programmes:

1. Use of climate science in decision-making
2. Mitigation of climate change (including the roles of carbon markets and low-carbon technologies)
3. Impacts of, and adaptation to, climate change, and its effects on development
4. Governance of climate change
5. Management of forests and ecosystems

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: <http://www.lse.ac.uk/grantham>.

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the author(s) and do not necessarily represent those of the host institutions or funders.

High impact, low probability? An empirical analysis of risk in the economics of climate change

Simon Dietz¹

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

Contact details: e-mail: s.dietz@lse.ac.uk

tel.: +44 (0) 207 955 7589

fax: +44 (0) 207 106 1241

Draft of July 2009

High impact, low probability? An empirical analysis of risk in the economics of climate change

Abstract

To what extent does economic analysis of climate change depend on low-probability, high-impact events? This question has received a great deal of attention lately, with the contention increasingly made that climate damage could be so large that societal willingness to pay to avoid extreme outcomes should overwhelm other seemingly important assumptions, notably on time preference. This paper provides an empirical examination of some key theoretical points, using a probabilistic integrated assessment model. New, fat-tailed distributions are inputted for key parameters representing climate sensitivity and economic costs. It is found that welfare estimates do strongly depend on tail risks, but for a set of plausible assumptions time preference can still matter.

Keywords

Catastrophe; climate change; cost-benefit analysis; discount rate; integrated assessment; risk; uncertainty

JEL codes

Q54

1. Introduction

To what extent does economic analysis of climate change depend on low-probability, high-impact events? This question has received a great deal of attention lately, stimulated by the integrated assessment modelling undertaken by the Stern Review on the Economics of Climate Change (Stern 2007), and in particular by Martin Weitzman's critique of it (Weitzman 2007), which he has subsequently developed into a fuller discussion of how risk and structural uncertainty affect the empirical analysis of climate change (Weitzman 2009). Weitzman shows that, loosely speaking, uncertainty about the economic impacts of climate change could be so large that societal willingness to pay to avoid extreme outcomes should overwhelm other seemingly important assumptions, notably on time preference. If this analysis is correct, it represents a significant departure in the economics of climate change. Without wishing to oversimplify, our understanding of the problem shifts away from one cast largely in terms of a relatively sure intergenerational investment, where the key consideration is equity, towards one cast largely in terms of insuring against very unlikely but catastrophic risks.

This paper is an intuitive empirical analysis of the conceptual issues Weitzman raises, using an integrated assessment model (the PAGE model built by Chris Hope: Hope 2006a). In terms of methodology, the take-home message from Weitzman's (2009) analysis is that estimates of the economic impacts of climate change are better derived from probabilistic studies, in which, crucially, the key relationships linking (i) greenhouse gas emissions with warming (e.g. 'climate sensitivity') and (ii) warming with economic costs (the damage function) are described by distributions with a heavy or 'fat' tail of low-probability, high-impact outcomes.² Thus I approximate one of the distributions for the climate sensitivity reported in the most recent assessment of the Intergovernmental Panel on Climate Change (IPCC: Meehl, Stocker et al. 2007), which has amongst the fattest tails of high-temperature outcomes, based on the largest ever sample of runs from a general circulation model (Stainforth, Aina et al. 2005), and I specify a

² There is some variability in the usage of terms such as heavy and fat tails (see also Nordhaus 2009). In some cases, the terms are used in a very broad sense to describe any probability distribution with a relatively high (though still very small absolute) risk of tail outcomes, especially in comparison with the normal distribution. In other cases, the terms are given specific, statistical meaning, though here too there is some variability. In this paper, I use the terms heavy and fat tail interchangeably to mean any distribution that has heavier tails than the exponential distribution.

damage function linking warming with economic costs that has uncertain curvature, with a probability distribution for that curvature which also has a fat tail.

While a few of the by-now numerous integrated assessments of the cost of climate change have worked with fat tails (e.g. Roughgarden and Schneider 1999; Tol 1999; Tol 2003; Mastrandrea and Schneider 2004), this paper takes an alternative angle, looking more closely at the implications of a potential climate catastrophe for the standard welfare model. The results of the modelling are analysed in terms of the probability of a 'catastrophe', which is taken to be the relative frequency with which arbitrarily large costs (e.g. 90% of consumption) occur, both under a scenario of business-as-usual emissions, and assuming emissions are controlled. Moreover, I explore the properties of welfare estimates as climate damage approaches the equivalent of 100% of consumption. This is a key element of Weitzman's (2009) analysis. In the most widely used empirical model of welfare, with constant and strictly positive relative risk aversion, marginal utility tends to infinity as consumption tends to zero, so we have to find some means of bounding damages. One way to do so, suggested by Weitzman (2009), is to impose an upper bound derived from an analogy to the value of statistical life, but pertaining to whole societies. I show how estimates of expected utility vary as a function of this upper bound.

Overall, I conclude that welfare estimates do strongly depend on tail risks, but for a set of plausible assumptions time preference still matters. With fat tails on the climate sensitivity and on the curvature of the damage function, the PAGE model estimates some positive probability of catastrophe, even when the atmospheric stock of greenhouse gases is stabilised at 550 parts per million (ppm). To be more precise, a very standard damage function produces a risk that climate damages exceed the total consumption of the economy (i.e. completely eliminate it), both in certain regions of the world and globally. Thus the magnitude of the catastrophe has to be bounded at less than 100% of consumption, and the expected welfare cost of climate change is very sensitive to small changes in where the upper bound is placed. However, accepting that there are *a priori* grounds for capping the total economic impact of a climate catastrophe, then time preference can still be an important consideration, because these tail-risk events are by definition unlikely to occur, and are relatively far off. Thus the rate at which utility is discounted, and the elasticity of substitution between consumption at different times, in different regions and in different states of nature can still matter.

Section two sets out the grounds for including catastrophic climate risks in the analysis, and assesses in more detail the extent to which existing economic studies have taken them into account. Section three discusses the properties of the standard economic model, with constant and strictly positive relative risk aversion, when faced with a climate catastrophe. Section four explains the modelling strategy, section five outlines the results, and section six provides a discussion and some concluding comments.

2. Catastrophic climate risks

Economists' growing interest in catastrophic climate change has been largely stimulated by two concerns, first, that greenhouse gas emissions might lead to very large increases in global temperatures, and, second, that large (perhaps even modest) increases in temperatures, as an index of wider climatic changes, might lead to very large economic costs.

The former concern, about the relationship between greenhouse gases accumulating in the atmosphere and warming, has very often been captured by the notion of climate sensitivity, that is, the equilibrium increase in global mean temperature resulting from a doubling in the atmospheric concentration of carbon dioxide.³ The IPCC's Fourth Assessment Report compiled a number of recent estimates of the climate sensitivity. Inspection of 18 probability density functions (pdfs) and frequency distributions in box 10.2 (Meehl, Stocker et al. 2007, p798) reveals that, in most if not indeed all cases, they have a positive skew, with a long tail of high estimates. These tails can be attributed to uncertainty about feedbacks, related for example to clouds and water vapour, and about the cooling effect of aerosols. This uncertainty is inherently difficult to reduce through measurements and modelling. In summarising this evidence, IPCC (2007) concludes that the best estimate of the climate sensitivity is 3°C, that there is a greater than 66% chance of it falling in the range 2-4.5°C (the IPCC's "likely" range), and a less than 10% chance of it being lower than 1.5°C ("very unlikely"). This leaves around a 17% chance that the climate sensitivity exceeds 4.5°C, and Weitzman (2009) roughly estimates, on the basis of the 18 pdfs, that there is a 5% chance of it exceeding 7°C, and a 1% chance of it exceeding 10°C.

³ This is adequate for a first cut of this sort at the relationship between emissions and warming. However, the rate of warming depends on other factors than just the climate sensitivity.

In comparison with this evidence, previous integrated assessments seem in the main to have been too sanguine. In most cases, the climate sensitivity has been set to a single, best estimate, as part of a deterministic analysis. The best estimate has tended to be placed somewhere in the region 2.5-3°C (e.g. Nordhaus and Boyer 2000; Nordhaus 2008). The small number of probabilistic studies have naturally explored a wider range of values, but while some have considered distributions with a large positive skew (e.g. Tol 1999; Mastrandrea and Schneider 2004; Hope 2006b), others have employed ranges that tend to look unreasonably narrow when set against the evidence presented by IPCC (2007). This particularly concerns the Stern Review (Stern 2007), which used the PAGE model, representing the climate sensitivity with a triangular pdf fitted on a minimum value of 1.5°C, a best estimate of 3°C, and a maximum value of 4.5°C. Thus there was no chance of the climate sensitivity exceeding 4.5°C.

In addition to uncertainty about the climate-sensitivity parameter, the relationship between greenhouse gas emissions and warming is rendered still more uncertain by the risk of a significant positive feedback from the carbon cycle, effectively increasing the atmospheric stock of greenhouse gases for a given pulse of emissions. This feedback could come from a weakening of terrestrial biospheric and oceanic carbon sinks, from the release of methane as permafrost melts and wetland areas undergo drying, and/or even from the release of methane from oceanic stores of gas hydrates. Amongst the prominent economic integrated assessment models, Warren et al. (2009) explain that positive feedbacks have received little attention, with the exception of the PAGE model (Hope 2006a), which has a strong positive feedback calibrated on a weakening in terrestrial biospheric and oceanic sinks.

The latter concern above, about the relationship between warming and economic costs, is the preserve of economic models, and has been captured in a variety of ways. For our purposes, an oft-used and tractable approach is to represent the general relationship between warming and climate damage by an aggregated damage function, expressing economic costs as a function of global mean temperature. This approach is taken in several well-known models (e.g. the PAGE and DICE models: Hope 2006a; Nordhaus 2008). Such a damage function is calibrated first through an estimate of overall impacts at 2.5°C or 3°C warming and second through an estimate of the functional form. A very simple specification of such a function is as follows:

$$d_{j,t,s} = \alpha_{j,s} \left(\frac{W_{j,t,s}}{2.5} \right)^{\beta_s} \quad (1)$$

Where W is warming in region j , at time t , and in state of nature s , in terms of global mean temperature above pre-industrial, d is the per-capita economic cost of climate change, α sets the region-specific cost of 2.5°C warming and β is the damage-function exponent, determining curvature. Clearly other classes of function could be employed here, including exponential functions, but equation (1) has the benefit of closely (e.g. Nordhaus, 2008) or precisely (e.g. Hope, 2006) replicating previous studies.

As there is some evidence to constrain estimates of the costs of 2.5°C warming (i.e. α), the more uncertain issue is what to assume about the functional form, β . Across the range of previous studies, the damage function has tended to be quadratic (which has the convenient analytical property of giving a linear marginal benefit function, but not much more justification than that) or less steep. Peck and Teisberg (1992) explored a cubic function in sensitivity analysis, and following this Hope (2006a) has specified a triangular pdf in the PAGE model, with a minimum value for β of unity (i.e. linear), a best estimate of 1.3, and a maximum value of 3. However, if we assume the global cost of 2.5°C warming is the equivalent of about 1.3% of GDP (which is the best estimate from PAGE)⁴, even a cubic function will lead to costs equivalent to only c. 10% of global GDP for 5°C warming, despite the fact that such warming could well trigger huge environmental changes (being for instance a greater difference in temperature than exists between the present day and the peak of the last ice age, the Last Glacial Maximum).⁵ The general problem in specifying β is that we have almost no evidence on which to base it. While the belief is widely shared that impacts will become disproportionately more severe with rising temperatures, and while there are numerous partial examples of this⁶, we just do not know what a contemporary world that is on average 5°C warmer than pre-industrial times would look like; there is no data point. Nevertheless, it is increasingly recognised that values of β up to and including 3 may not give enough credence to the economic consequences of very large climatic changes, at least in the tail of the distribution.

⁴ Ignoring for the moment the regional weighting on α .

⁵ PAGE includes a separate function representing an economic catastrophe as a result of climate change (Hope, 2006a), but it has been shown to add relatively little (Warren, Hope et al 2006). To avoid double-counting, it is always switched off in the analysis of this paper.

⁶ For example, the relationship between warming and crop productivity at high temperatures (see chapter 3 of Stern, 2007, for a summary). There are of course counter-examples.

3. Welfare analysis of a climate catastrophe

In cost-benefit analysis of climate change, it is usual to specify the utility of a representative individual as an isoelastic or constant-relative-risk-aversion (CRRA) function:

$$u(c_{j,t,s}) = \begin{cases} c_{j,t,s}^{1-\eta} / (1-\eta) & \text{if } \eta \neq 1 \\ \ln(c_{j,t,s}) & \text{if } \eta = 1 \end{cases} \quad (2)$$

Where u is utility per capita, c is consumption per capita, and η is the elasticity of the marginal utility of consumption, which simultaneously captures aversion to inequality in consumption over regions and time, and relative risk aversion.

Utility is then summed across regions, weighted by their populations, time periods, weighted by the utility discount rate (i.e. the rate of pure time preference or impatience), and states of nature, weighted by their probabilities, to give expected social welfare:

$$E(V) = \sum_{j=1}^J \sum_{t=0}^T \sum_p p_s(u_{j,t,s}) N_{j,t} (1 + \delta)^{-t} \quad (3)$$

Where V is aggregate social welfare, δ is the utility discount rate, and N_j is the size of the population of region j . There are J regions and T time periods.

The CRRA utility function is the mainstay of applied welfare analysis, principally because it is very tractable, but this comes at the price of its limit properties, which are in fact of the utmost relevance when we think of a climate catastrophe that might drive consumption down to very low levels if not indeed eliminate it altogether. As long as η is positive, the CRRA function exhibits diminishing marginal utility of consumption ($u'(c) > 0$, $u''(c) < 0$). In the limit as consumption tends to zero, it has the following properties:

$$u(c_{j,t,s}) \rightarrow \begin{cases} -\infty \\ 0 \\ -\infty \end{cases} \quad \text{as } c \rightarrow \begin{cases} 0 \\ 0 \\ 0^+ \end{cases} \quad \text{if } u(c_{j,t,s}) = \begin{cases} c_{j,t,s}^{1-\eta} / (1-\eta) & \text{and } \eta > 1 \\ c_{j,t,s}^{1-\eta} / (1-\eta) & \text{and } \eta < 1 \\ \ln(c_{j,t,s}) \end{cases} \quad (4)$$

Where $\eta \geq 1$, utility tends to minus infinity as consumption tends to zero, while where $\eta < 1$, utility tends to zero as consumption tends to zero. In all cases of strictly positive relative risk aversion ($\eta > 0$), marginal utility, u' , tends to infinity as consumption tends to zero.

These are of course well-recognised properties of the CRRA utility function, which are at the core of Weitzman's (2009) result, because the implication is that, for strictly positive relative risk aversion, society's marginal willingness to pay to avoid zero consumption is unbounded, and in a probabilistic setting all that is required is the existence of one future state of nature with zero consumption, even if it is very, very unlikely. Put another way, we are willing to pay an infinite amount to avoid the tiny risk of losing everything. In Weitzman (2009), this follows logically from certain assumptions about the structural distribution of consumption, from the assumption of structural uncertainty, and from deriving the parameters of the distribution of consumption from Bayesian learning.⁷ Here I do not replicate Weitzman's (2009) model, which is a generalised theoretical story in which the impact of climate change is simply represented by the distribution of consumption growth. In economic integrated assessments, uncertainty about consumption, net of the cost of climate change, is a much more complex joint distribution of many climatic, environmental and economic processes. Nevertheless, the general point is that model estimates may be very sensitive to the risk of catastrophe, and that is worth exploring.

In an integrated assessment model, all that is required for expected marginal utility to be unbounded is that $\eta > 0$ and:

$$\exists s, \quad s \in [0, S] \quad : \quad d_{j,t} \geq \hat{c}_{j,t} \quad j \in [0, J], t \in [0, T] \quad (5)$$

⁷ Geweke (2001) considers a more general set of cases in which the assumption of CRRA utility can sometimes lead to unbounded welfare estimates, depending on what distribution of consumption is assumed, whether the mean and variance of the distribution are known, whether the parameters of the distribution are derived from Bayesian learning, and on η .

Where d is the per-capita cost of climate change (e.g. from equation 1), and \hat{c} is baseline consumption per capita. That is, there exists a state of nature s , belonging to the set of states of nature S , such that damages in some region j at some time t exceed baseline consumption per capita.

In the catastrophe region j at the time of the catastrophe t , expected marginal utility will be unbounded ($\eta > 0$) and expected total utility will be unbounded or zero depending on η , as explained in equation (4). Expected marginal utility will also be unbounded at the global level if the condition in equation (5) is met for any region j (and $\eta > 0$), and, as long as $\eta \geq 1$, expected total utility will also be unbounded. If $\eta < 1$, then expected total utility still exists and can be positive globally, since while utility in the catastrophe region(s) j is zero, consumption and therefore utility could still be positive in other regions.

Considering evidence on the climate sensitivity and damage function, condition (5) is far from implausible. In standard applications of PAGE, the most vulnerable region of the world to climate change is India and Southeast Asia, in which the best estimate of α from equation (1) (the damage cost of climate change for 2.5°C warming) is 3.1% of regional GDP. Assuming β takes its maximum value in standard PAGE of 3, condition (5) will be met for India and Southeast Asia if regional temperature exceeds around 7.5°C above pre-industrial at some point over the modelling horizon, which is two hundred years. A global catastrophe is not implausible either. Let α and β equal 2.5% of global GDP and 3 respectively.⁸ Then condition (5) will be met with respect to global aggregate consumption if the global mean temperature exceeds about 8.5°C above pre-industrial. Assuming high business-as-usual emissions of greenhouse gases throughout this century (the A1FI SRES scenario: Nakicenovic and Swart 2000), IPCC estimates up to a 17% chance of global mean temperature exceeding 6.4°C by 2100 alone (IPCC 2007). With the more general assumption of stochastic emissions, a recent study by scientists from MIT goes further, estimating a median increase in global mean temperature of 5.2°C by 2100, with a 10% chance of more than 7.4°C warming (Sokolov, Stone et al. 2009). Moreover, the argument was made above that even higher values of β should be considered. Therefore, certainly if our task is to evaluate the cost of business-as-usual emissions over long periods, and even perhaps in a scenario in which emissions are

⁸ $\alpha=2.5\%$ is the maximum assumed value in standard PAGE, ignoring regional weighting.

controlled but climate and ‘welfare’ sensitivity are very high (e.g. in India and Southeast Asia, or in Africa), we need to bound the problem in some way.

Let consumption per capita, net of the cost of climate change, be bounded from below, so that damages never exceed some share ϕ of baseline consumption per capita:

$$c_{j,t,s} = \min[\hat{c}_{j,t} - d_{j,t,s}, \hat{c}_{j,t}(1-\phi)] \quad (6)$$

In effect, the probability distribution on the random variable $\tilde{c}_{j,t}$ is truncated. In principle $0 \leq \phi < 1$, but since the upper bound only exists to create a finite solution where it would not otherwise exist (i.e. where $d_{j,t} \geq \hat{c}_{j,t}$), in practice ϕ is close to, but strictly less than, unity.

There are likely to be several ways in which ϕ could be derived. Weitzman (2009) does so by specifying it as a function of the value of statistical life. The benefit of this is that the bound on consumption losses can at least be related to a familiar concept in economics, one that has been estimated many times, with the key point being that the evidence (and indeed intuition) suggests it is finite (Viscusi and Aldy 2003). In other words, we do not personally appear to be willing to liquidate all of our wealth to reduce the probability of our own death. The wrinkle in this case is that, since we are dealing with whole societies, if not indeed the whole of human civilisation, the directly relevant construct is “the value of statistical civilization as we know it or maybe even the value of statistical life on Earth as we know it” (Weitzman, 2009, p6). Thus we know very little about ϕ . The purpose of this analysis is not to consider the derivation of ϕ , but rather to look at the role of ϕ in estimating the welfare cost of climate change, for a plausible range of values close to unity.

Substituting (6) into (2) gives:

$$u(c_{j,t,s}) = \begin{cases} \{\min[\hat{c}_{j,t} - d_{j,t,s}, \hat{c}_{j,t}(1-\phi)]\}^{1-\eta} / (1-\eta) & \text{if } \eta \neq 1 \\ \ln\{\min[\hat{c}_{j,t} - d_{j,t,s}, \hat{c}_{j,t}(1-\phi)]\} & \text{if } \eta = 1 \end{cases} \quad (7)$$

From (7) and (4) it follows straightforwardly that, in any state of nature where $d_{j,t} \geq \hat{c}_{j,t}$, for some region j and time t , as ϕ tends in the limit to unity, consumption per capita will tend to zero, utility

per capita will tend either to minus infinity or zero depending on η , and marginal utility will tend to positive infinity when $\eta > 0$.

At this point, we can make three postulates about the results of the empirical analysis. First and foremost, given that marginal utility is very high when consumption per capita is very low, we might anticipate that, if condition (5) is met, then the expected loss in social welfare $E(\Delta V)$ due to climate change will be sensitive to the upper bound on damages, ϕ . Even superficially small differences in ϕ (say, raising it from 0.95 to 0.99) might lead to large differences in $E(\Delta V)$.

Second, if there exists some state of nature s in which $d_{j,t} \geq \hat{c}_{j,t}$, and in the plausible case where $\eta \geq 1$, then in the limit as $\phi \rightarrow 1$, the marginal change in social welfare with respect to the rate of impatience, δ , will tend to zero *in that state*. In fact, it is also true that the marginal change in social welfare with respect to η will tend to zero *in that state*. To see this, take that catastrophic state of nature to be $s=1$. From (7), given condition (5) is satisfied:

iff $d_{j,t,1} \geq \hat{c}_{j,t}$

$$\Rightarrow \lim_{\phi \rightarrow 1} u(c_{j,t,1}) = -\infty \quad \text{iff } \eta \geq 1 \quad (10)$$

$$\Rightarrow V_1 = \sum_{j=1}^J \sum_{t=0}^T (u_{j,t,1}) N_{j,t} (1 + \delta)^{-t} = -\infty \quad (11)$$

$$\Rightarrow \frac{\partial V_1}{\partial \delta} = 0 \quad \text{and} \quad \frac{\partial V_1}{\partial \eta} = 0 \quad (12)$$

This is the postulate that, in the limit, only the risk of catastrophe determines the welfare cost of climate change, and the social discount rate is unimportant.

But, third, as long as $\phi < 1$ and the solution is therefore finite, δ and η will still be important determinants of the expectation of social welfare when integrating over *all states* of nature. To illustrate, consider just two states of nature, $s=1,2$. For $s=1$, as before $d_{j,t} \geq \hat{c}_{j,t}$, so the upper bound is applied as per equation (6). As ϕ becomes vanishingly close to unity, the loss in social welfare in that state, ΔV , becomes vanishingly close to 100%. But for $s=2$, where $d_{j,t} < \hat{c}_{j,t}$, the marginal change in social welfare with respect to ϕ is zero; i.e. the upper bound never kicks in. If $d_{j,t}$ in state

of nature 2 is equal to, say, zero, then the expected welfare loss $E(\Delta V)$ is bounded from above at just under 50% as $\phi \rightarrow 1$.

The key conceptual point here relates to non-uniform convergence. For any given $\phi < 1$, $E(\Delta V)$ is an arbitrarily small function of δ and η . Conversely, for any given δ and η , $E(\Delta V)$ is infinite when $\phi = 1$ and $d_{j,t} \geq \hat{c}_{j,t}$ for some s , and if $\eta \geq 1$. This is a specific formulation of the more general theorem proposed in Stern (2008, p20). In the very simple example above, $d_{j,t} \geq \hat{c}_{j,t}$ for $s=1$ and $d_{j,t}=0$ for $s=2$, so as $\phi \rightarrow 1$, $E(\Delta V)$ is bounded from above by $p[d_{j,t,s} \geq \hat{c}_{j,t}] = 0.5$. This result generalises of course to many states of nature, with different economic losses. The only condition is that the probability of a catastrophe should be less than unity. Then as $\phi \rightarrow 1$, $E(\Delta V)$ will asymptote to some arbitrarily small function of δ and η .

4. Modelling strategy

To reiterate, the purpose of this paper is to test empirically the extent to which economic analysis of climate change depends on low-probability, high-impact events in the climate-economy system. I do so using the PAGE integrated assessment model (Hope 2006a). PAGE is particularly suitable for the present study because it is probabilistic, set up to conveniently run Monte Carlo-type simulations. Furthermore, depending of course on assumptions about key parameters and exogenous variables, PAGE has been demonstrated to provide a good approximation to the results of a range of other models, both of parts of the climate-economy system and of the whole system. As an example of the former, Warren et al. (2009) show that PAGE on the whole provides a good approximation to the projections of atmospheric carbon dioxide concentrations, radiative forcing and warming made by much more complex general circulation models of the climate system, for a consistent set of assumptions about carbon dioxide emissions. As an example of the latter, the mean estimate of the social cost of carbon dioxide made by Hope (2006a) is close to the central estimate from a range of peer-reviewed studies (Tol 2005). In other words, there does not appear to be anything intrinsic to the model itself that would make it an outlier, although clearly the purpose of the present study is precisely to test results from PAGE given certain outlying assumptions about warming and damages.

In common with most applications of integrated assessment models, I avoid an equation-by-equation description of PAGE for the sake of brevity. However, the model is set out in full in Hope (2006a) and Hope (2008) and unless otherwise stated no changes have been made. Here I focus on the parameterisation of two relationships described that are likely to have a particular bearing on my main question; first, the climate sensitivity, and, second, the damage function.

The pdf of the climate sensitivity in most previous applications of PAGE has been triangular. In the Stern Review (Stern 2007), this triangular distribution was fitted on a minimum value of 1.5°C, a best estimate of 3°C, and a maximum value of 4.5°C. Yet in view of the scientific evidence collected by IPCC (Meehl, Stocker et al. 2007) and discussed above, this range looks to be too narrow. Indeed, the triangular distribution is *par excellence* an example of a thin-tailed pdf. Therefore I impute an alternative pdf for the climate sensitivity, fitted on the analysis of Stainforth et al. (2005), using the commercial software package @Risk to make an approximation. @Risk works by fitting a wide range of different kinds of probability distribution to the data and selecting the distribution with the best fit in terms of the lowest root-mean-square error. The resulting pdf in this case is log-logistic⁹, and is reproduced in figure one alongside the Stern Review pdf for comparison. The minimum of the distribution is 1.6°C, the best estimate is 3.2°C, and the 99th percentile is 8.2°C (appropriately, the log-logistic distribution does not have a finite upper support, so the maximum value will vary from one Monte Carlo simulation to another).

Stainforth et al.'s (2005) estimates of the climate sensitivity are one result from the huge *climateprediction.net* experiment, in which distributed computing resources were used to run the first 'grand ensemble' of simulations through a general circulation model of the climate system. A grand ensemble is essentially a Monte Carlo simulation, in which model parameters are varied (a so-called 'perturbed physics' ensemble), alongside initial climate conditions (e.g. initial temperatures and humidity levels) and carbon dioxide emissions or concentrations. This gives one of the most comprehensive available estimates of climate-model uncertainty about equilibrium temperature change, based on the largest ever set of model runs. Other pdfs reported by IPCC (Meehl, Stocker et al. 2007, box 10.2) could also be used, but the Stainforth et al. distribution is a good proxy of their typical features.

⁹ i.e. the probability distribution of a random variable (the climate sensitivity) whose logarithm has a logistic distribution.

FIGURE ONE HERE

Another key element in the climatic response to greenhouse gas emissions is the carbon cycle, and in particular the risk of a large-scale positive feedback. This is captured in the basic structure of PAGE by a parameter representing natural stimulation of carbon dioxide from weakening terrestrial biospheric and oceanic carbon sinks, which is as usual triangular distributed. Future work could attempt to substitute this by an alternative pdf, perhaps including the likelihood of natural stimulation of methane from melting permafrost, from drying of wetlands, and even from the oceans. However, there is currently very little evidence upon which such a pdf could be based, compared with the climate sensitivity. Therefore the triangular distribution of natural carbon dioxide stimulation is retained.¹⁰

In estimating the relationship between warming and economic losses, section two argued that the key uncertainty can be captured in the curvature of a generalised damage function (the parameter β in equation 1). Once again, PAGE has in past studies been run with a triangular distribution for β , where the minimum value is unity, the best estimate is 1.3, and the maximum value is 3. Yet section two argued an intuitive case for exploring values higher than 3, as Weitzman (2009) has too. Thus @Risk is once again used to fit a revised pdf, but, given the lack of evidence on β , this time I make a minimal set of assumptions about damages. First I assume that if there is no warming there is no economic cost, all else equal. Second, I assume that the best estimate of β is 2, giving a quadratic damage function. This has the merit of replicating a very common assumption in the literature, although, as mentioned above, the assumption itself has little grounding except in analytical convenience. Third and finally, I make the arbitrary assumption that the 90th percentile is 3, so that there is a 10% likelihood of β exceeding 3. The best fitting distribution (figure two) is lognormal, with a minimum value of 0.6, a best estimate of 1.7 and a 99th percentile of 3.8 (the lognormal distribution does not have a finite upper support either).

FIGURE TWO HERE

¹⁰ In simulations where the atmospheric concentration of greenhouse gases is stabilised at 550 parts per million, the natural stimulation of carbon dioxide is halved relative to its baseline value (Hope, 2006a), following a suggestion by Chris Hope and based on research by van Vuuren et al. (2009).

Economic growth, population growth and greenhouse gas emissions are exogenous to the PAGE model. In this paper they are taken from the IPCC SRES A2 scenario (Nakicenovic and Swart 2000), converted to purchasing power parity and extrapolated from 2100 to 2200 by Hope (2006a). The Stern Review (Stern 2007) also used the A2 scenario, so it provides a point of contrast. The overall modelling horizon is 2000 to 2200. GDP from the A2 scenario is converted to consumption for estimation of welfare using a constant exogenous saving rate of 20%. Business-as-usual emissions are considered, alongside a scenario that sets out to stabilise the atmospheric concentration of greenhouse gases at 550 parts per million (ppm), although the carbon-cycle feedback results in an eventual concentration somewhat higher than 550 ppm. This provides estimates of the reduction in the risk of catastrophe brought about by emissions reductions.

In performing Monte Carlo simulations, a Latin Hypercube Sample is taken, with 1000 simulation draws. Latin Hypercube Sampling improves coverage of the full range of input values relative to a simple, non-stratified Monte Carlo procedure, an especially advantageous feature when investigating probability distributions with a large skew.

5. Results

The first step in the empirical analysis is to quantify the probability of a high-impact event in the climate-economy system, a catastrophe. Since there is no canonical meaning of ‘catastrophe’ in economics (or indeed elsewhere), I simply investigate the relative frequency with which arbitrarily large economic costs occur. In particular, table one begins by estimating the relative frequency with which the global cost of climate change, at any point in time between 2000 and 2200, exceeds three thresholds, 75% of global consumption, 90% of global consumption, and 100% of global consumption. In terms of the discussion above, the last threshold of 100% of consumption is especially important, as its breaching will directly lead to infinite or zero social welfare, and infinite marginal utility for strictly positive relative risk aversion. Hence we will gain an initial insight into the relevance of the upper bound on damages, ϕ .

Table one presents estimates of the probability of a catastrophe conditional on two sets of assumptions about the climate sensitivity and the damage-function exponent (β). The first set comprises the triangular pdf for the climate sensitivity and for β from the Stern Review (Stern, 2007: labelled ‘Stern’ for convenience). The second set comprises the log-logistic pdf for the climate

sensitivity fitted on data from Stainforth et al. (2005), and the lognormal pdf for β outlined above. Estimates are presented both for business-as-usual emissions and for the stabilisation scenario.

TABLE ONE HERE

Looking first at business-as-usual emissions, the probability that the global cost of climate change exceeds 75% of global consumption is 0.1% (i.e. 1 in 1000 simulation draws), assuming the triangular pdf for the climate sensitivity from the Stern Review, and the triangular pdf for β . The probability of costs greater than 90% of consumption, and in turn greater than 100% of consumption, is zero, in terms of zero draws from the Latin Hypercube Sample. However, for the more plausible assumption that the climate sensitivity is log-logistically distributed, while β is lognormally distributed, the probability of a global climate catastrophe is positive across all three thresholds. The probability of costs greater than 75% of global consumption is 6.9%, while the probability of costs greater than 90% of consumption is 5.5% and the probability of costs reaching 100% of consumption is 5.0%. This is worrying, of course. It also gives immediate empirical backing to the relevance of the issues raised by Weitzman (2007; 2009), including ϕ .

Turning to the stabilisation scenario, we can see that emissions reductions buy down the probability of catastrophe. With the ‘Stern’ climate sensitivity and damage-function exponent, damages in the sample never exceed the three threshold levels of consumption. On the other hand, with the fat-tailed distributions, the probability of catastrophe is not eliminated in the sample. Global damages exceed the 75% threshold in 0.7% of draws, the 90% threshold in 0.6% of draws and reach the 100% threshold in 0.5% of draws.

In fact, the estimates in table one may still underplay the importance of high-impact events, because they may not pick up the existence of a regional catastrophe, where climate damage exceeds 75%, 90%, or reaches 100% of regional consumption, but consumption in other regions is sufficient to compensate. Therefore table two takes the same set of simulation draws and this time counts every draw in which there is a catastrophe in any one or more of PAGE’s eight major world-regions.¹¹

¹¹ (i) Western Europe (i.e. the old EU-15); (ii) the former Soviet Union and Eastern Europe; (iii) the USA; (iv) China and centrally planned Asia; (v) India and Southeast Asia; (vi) Africa and the Middle East; (vii) Latin America; (viii) other OECD.

TABLE TWO HERE

As expected, the probability of a regional catastrophe is higher than that of a global catastrophe. Under business-as-usual and assuming the ‘Stern’ climate sensitivity and damage-function exponent, the probability that costs exceed 75% of regional consumption in any one or more regions is 3.8%, while there is a 2.8% probability of costs exceeding 90% of consumption, and a 2.3% chance of costs reaching 100% of consumption. Substituting the more plausible, fat-tailed distributions for warming and damages, the probabilities rise to an alarming 17.7%, 14.4% and 12.9% respectively. The regions most likely to suffer a catastrophe are India and Southeast Asia, Africa and the Middle East, and Latin America, although the only region to escape a catastrophe altogether is the former Soviet Union and Eastern Europe.¹² The risk of catastrophe is highest in the period from 2150 to 2200, although, in the scenario with fat-tailed distributions, there is a 0.8% probability of costs reaching 100% of consumption between 2060 and 2100.¹³ Once again, stabilising greenhouse gas concentrations reduces the risk of a catastrophe, but assuming the fat-tailed distributions, it remains considerable.

Given the risk of a catastrophe, tables one and two clearly show that the introduction of an upper bound on climate damages, ϕ , could be important. Section three also speculated that the precise value of ϕ may be important, even to within a very narrow range in the region of unity. Figure three thus presents estimates of the variation in the welfare cost of climate change, measured as the expected change in social welfare $E(\Delta V)$ normalised to present GDP per capita¹⁴, for a range of values of ϕ between 0.75 and very close to unity (0.9999999). As we are now considering the sensitivity of social welfare (equation 3), three curves are plotted, for different combinations of the utility discount rate, δ , and the elasticity of the marginal utility of consumption, η . The climate-sensitivity pdf is from Stainforth et al. (2005), while β is lognormally distributed.

FIGURE THREE HERE

¹² Clearly this depends on the assumptions of the model. Indeed, in Tol (2003), the FUND model generated its only regional catastrophe in the former Soviet Union.

¹³ Further data on the incidence of a ‘catastrophe’ by region and/or by time period can be obtained from the author on request. PAGE has ten uneven time steps: 2000; 2001; 2002; 2010; 2020; 2040; 2060; 2080; 2100; 2150; 2200.

¹⁴ Formally, the change in the balanced growth equivalent (BGE), as used by the Stern Review (Stern, 2007) and fully elaborated in Anthoff and Tol (2009).

The bottom-most curve traces the relationship between the expected loss in social welfare and ϕ , for the Stern Review's (Stern 2007) disputed assumptions that $\eta=1$ and $\delta=0.1\%$ per annum. The shape of the curve is striking, showing not only that the welfare cost of climate change is sensitive to ϕ , increasing from 11.8% of present GDP per capita when ϕ is 0.75 to 18.6% when ϕ is very close to unity (0.9999999), but also that the welfare cost is increasing more than proportionately in ϕ . As $\phi \rightarrow 1$, the marginal change in the welfare cost with respect to ϕ tends to infinity. Because the CRRA utility function exhibits diminishing marginal utility for $\eta>0$, this is what we would expect. Indeed, the curve looks like the mirror-image of a CRRA utility function ($\eta>0$) plotted against rising consumption per capita. Thus, as we might have anticipated, the analysis is particularly sensitive to very small increases in the value of ϕ in the region of (strictly less than) unity.

However, just as important is to observe the cluster of observations as ϕ tends in the limit to unity. What this shows is that, while the incremental expected loss in social welfare increases, the incremental increase in ϕ itself becomes vanishingly small. Indeed, the latter effect exceeds the former, so that the welfare cost of climate change asymptotes to a particular value, in this case roughly 20% of present GDP per capita. This reflects the third postulate in section three: for any given $\phi<1$, the welfare cost of climate change is an arbitrarily small function of δ and η . Here, the Latin Hypercube Sample gives a 5% probability that climate damages reach 100% of global consumption (table one). As ϕ tends in the limit to unity, the expected loss in social welfare will tend to 100% for these 5% in 1000 simulation draws. But for the remaining 95% of draws, climate damages are below ϕ , and hence the increase in ϕ is irrelevant. It is in this sense that the welfare cost of climate change is bounded from above by the probability of a catastrophe.

Turning now to the other two sets of assumptions about δ and η , the top-most curve plots the relationship for $\delta=3\%$. η remains at unity. What can be seen is that increasing the utility discount rate lowers the welfare cost of climate change for all values of ϕ (postulate three again). Indeed, the difference is substantial. Setting ϕ to, for example, 0.99, the welfare cost of climate change falls from 15.3% of present GDP per capita to just 2.8%, when δ is increased from 0.1% to 3%. This is unsurprising; the importance of the utility discount rate to economic analysis of climate change is well known. What can also be seen is the same shape of relationship between the expected loss in social welfare and ϕ , namely a more than proportionate increase in cost as ϕ is increased.

The middle curve plots the relationship for $\eta=3$, while holding δ at 0.1%. Again, the same relationship is traced between the expected loss in social welfare and ϕ . And, again, the expected loss in social welfare is reduced for all values of ϕ , when η is increased. Setting ϕ to 0.99, the welfare cost of climate change falls from 15.3% of present GDP per capita when $\eta=1$ to 14.1% when $\eta=3$. This could not have been confirmed *a priori*, because in the standard welfare model η simultaneously captures aversion to inequality in consumption over regions and over time, and relative risk aversion. With increasing η , more weight is placed on consumption losses in poor regions, and more weight is placed on the tail risks. Since poor regions are expected to suffer the worst relative impacts of climate change, and PAGE represents this in its parameterisation of α_r (Hope 2006a), both of these would tend to increase the expected loss in social welfare. However, assuming consumption growth, an increase in η also effectively increases the social discount rate. Thus what figure three indicates is that the temporal effect outweighs the spatial and risk effects in this particular instance.

What figure three also shows is the relative difference in the welfare cost of climate change as ϕ is varied. Consider first the utility discount rate, δ . We can see from inspection of figure three that the difference in the expected loss in social welfare between $\delta=0.1\%$ and $\delta=3\%$ actually *increases* in ϕ . That is to say, the effect of the utility discount rate on the welfare cost of climate change rises as the upper bound on climate damages rises. This is actually quite intuitive: increasing the impact of these events in the extreme tail, which are far off, increases the relevance of δ , and higher δ reduces the weight placed on these impacts. The expected loss in social welfare has a different asymptote when $\delta=0.1\%$ (about 20% of present GDP per capita) compared with $\delta=3\%$ (about 6% of present GDP per capita), since the probability that $d_{j,t} \geq \hat{c}_{j,t}$ is quite low, and so the loss of social welfare in states of nature where $d_{j,t} < \phi \hat{c}_{j,t}$, which is a function of δ , determines the asymptote. To reiterate, for any given $\phi < 1$, the welfare cost of climate change is an arbitrarily small function of δ .

Turning to the elasticity of the marginal utility of consumption, η , figure three shows that as ϕ increases, the difference in the welfare cost of climate change initially *decreases*. As discussed above, by the time $\phi=0.99$ the difference is only 1.2 percentage points. But as ϕ becomes very close to unity, the difference begins to increase again. What is happening is that, with higher η , the expected loss in social welfare reaches its asymptote more quickly, as more weight is placed on

catastrophic outcomes with very low consumption. This initially serves to reduce the difference in $E(\Delta V)$ due to η , but then serves to increase it again (when $\eta=3$, the asymptote is reached by the time $\phi=0.999$, whereas when $\eta=1$ it takes longer).

Finally, table three gives estimates of the social cost of carbon (specifically the marginal damage cost of carbon dioxide) accompanying these experiments. The estimates are conditional on different values of the disputed welfare parameters δ and η , comparing the estimate for business-as-usual emissions with the stabilisation scenario. ϕ is set to 0.99. The striking feature is of course that the estimates are very high, conditional on assumptions about δ and η . For the Stern Review's (Stern, 2007) assumptions that $\eta=1$ and $\delta=0.1\%$, the mean social cost of carbon along business-as-usual is a dizzying US\$430.93/tCO₂, within a 90% confidence interval of over US\$1700/tCO₂. For the stabilisation scenario, the equivalent mean price is US\$291.59/tCO₂. However, when $\eta=3$ and $\delta=3\%$, the estimates are much lower. Thus the table makes two points. First, replacing thin-tailed pdfs for the climate sensitivity and for the curvature of the damage function with more plausible fat-tailed pdfs, the social cost of carbon rises significantly. This has strong implications for any current policy that attempts to calibrate the actual price of carbon (e.g. the carbon tax rate, or the price of a tradable permit) on earlier, thin-tailed assessments. Second, the table demonstrates the ongoing importance of assumptions about social discounting, even when catastrophic climate change is possible.

TABLE THREE HERE

6. Discussion and conclusions

To what extent does economic analysis of climate change depend on low-probability, high-impact events? The short answer is a great deal (see also Ackerman, Stanton et al. 2009), but not to the exclusion of other factors that we already know to be very important, in particular the discount rate. Inputting fat-tailed distributions for two key parameters in the analysis, namely the climate sensitivity and the curvature of the damage function, I find that a catastrophe is possible in the integrated assessment model PAGE, both under business-as-usual emissions and even stabilising the atmospheric stock of greenhouse gases at nearly 550 ppm. This brings to the fore the issues recently raised by Martin Weitzman (2007; 2009).

In order to then undertake welfare analysis, it is necessary to bound the catastrophe in some way. Failure to do so in the standard welfare model with risk aversion implies we are willing to pay an infinite amount at the margin to reduce greenhouse gas emissions, making it impossible to come to useful conclusions about, for example, optimal rates of carbon taxation. This paper has explored one way to do so, introducing an upper bound on climate damages that could in principle be derived from the concept of the (finite) value of statistical life, albeit essentially scaled up to the level of whole societies, or even to the ultimate level of human life itself (Weitzman 2009). Unsurprisingly, the valuation of climate impacts turns out to be sensitive to this upper bound, in particular to very small differences in the upper bound towards its limiting value. This is concerning, since there is very little evidence on which to base it.

Nevertheless, if you subscribe to three assumptions made in this paper, there is cause to believe that the importance of catastrophes is bounded, and long-standing controversies over the discount rate continue to matter. The first assumption is that the probability of a catastrophe is, roughly speaking, small. The second is that the catastrophe will not occur for decades to come, if not indeed centuries. The third is that there are *a priori* grounds to 'cap' the total impact of the catastrophe (e.g. in reasoning by analogy to the value of statistical life). If all three assumptions hold, the welfare cost of climate change will asymptote to some value much less than the total value of the global economy, and moreover the two parameters of the social discount rate, the utility discount rate and the elasticity of the marginal utility of consumption, will be strong determinants of where that asymptote lies.

Besides interpretation of the results, this paper raises a number of open questions and issues for further research. Immediately, there will be suggestions to tweak the empirical approach, trying for instance an alternative probability distribution for the climate sensitivity, and in particular introducing a fat-tailed distribution for the carbon-cycle feedback, or for the curvature of the damage function. These would of course be interesting extensions. They are likely to be important in taking very much that is quantitative from the paper, but on the other hand would be unlikely in my view to qualitatively alter its main conclusions, unless for some reason a catastrophe can be ruled out through these changes.

Instead, I would like to draw attention to two issues. The first is the interpretation of the upper bound on damages, and indeed alternative ways in which the problem could be bounded. More

rigorous thinking on these issues is surely warranted, given the sensitivity of the analysis to where (and quite probably how) it is bounded. As mentioned, Weitzman (2009) draws an analogy with the value of statistical life, since individuals make trade-offs on a daily basis between decreasing present consumption and reducing the probability of their own death, at least tacitly. This, he shows with some ballpark estimates, might lead to an upper bound of 99% of welfare-equivalent consumption. Unfortunately, the analysis in this paper has confirmed the suspicion that it matters whether it is 99%, or 99.9%, and so on.

Nordhaus (2009), on the other hand, reasons by analogy with the actions of public institutions towards mitigating tiny risks of other catastrophes, such as the collision of a very large asteroid with the Earth, or unwanted, catastrophic consequences of research into biotechnology and nanotechnology. We spend just a few million dollars per year researching asteroids, while research into bio- and nanotechnology continues, despite the risks. This, it seems to me, is a promising area for further research, because the decision context – i.e. public institutions governing societal risks – is more directly comparable to the problem of mitigating catastrophic risks from climate change. We do not, however, have to choose. We need more research on both. And there are still other ways forward. For one, reasoning by analogy with the value of statistical life is very similar to reasoning by analogy with the concept of ‘fear of ruin’ (Aumann and Kurz 1977). That is, how much would an individual be willing to pay to be fully insured against the possibility of ruin? Fear of ruin can be applied to a wide range of choice situations. For another, the problems discussed in this paper are in large part a consequence of the reliance on a particular form of utility function. Other utility functions do not possess the same, troublesome limiting properties as the CRRA utility function. And when it comes to regional catastrophes, Yohe (2003) has argued that the upper bound can be interpreted as the receipt of aid in the affected region, from other regions, to maintain living standards above subsistence levels.

The second issue to highlight is the analysis of probabilities. This is a paper about risk, in the sense that it is assumed we know the mean and the variance of the joint probability distribution of climate impacts. Weitzman (2009) also considers this distribution, but his paper concerns how structural uncertainty about the mean and variance propagates through to give a ‘fat-tailed’ distribution of impacts. The question Heal (2008) raises is whether, if we cannot describe the impacts of climate change probabilistically, we should be working within standard expected-utility analysis at all? There are reasonably close alternatives within economics, in the shape of, for

example, ambiguity theory (Henry and Henry 2002). And of course analysis of climate change from outside the economics profession has long applied notions of policies that are 'robust' in some sense to various, non-probabilistic scenarios. While such approaches have their own fair share of weaknesses, this seems to be a potentially rich vein of enquiry as well.

Acknowledgements

I am grateful to Chris Hope for enabling me to use the PAGE model and for invaluable advice. I would also like to thank Alex Bowen, Christian Gollier, Robert Hahn, Cameron Hepburn, Nicola Ranger, Leonard Smith, Nicholas Stern, Martin Weitzman and Dimitri Zenghelis for comments on earlier drafts, and to participants in the session "The Role Of Integrated Assessment Models In Handling Climate Change Part I" at the IARU International Scientific Congress in Copenhagen on 11th March 2009. This research has been supported by the Munich Re programme of the Centre for Climate Change Economics and Policy; "Evaluating the Economics of Climate Risks and Opportunities in the Insurance Sector". All errors are my own.

References

- Ackerman, F., E. A. Stanton, et al. (2009). Fat tails, exponents, and extreme uncertainty: simulating catastrophe in DICE. Stockholm Environment Institute Working Paper WP-US-0901.
- Aumann, R. J. and M. Kurz (1977). "Power and taxes." Econometrica **199**: 1137-1161.
- Geweke, J. (2001). "A note on some of the limitations of CRRA utility." Economics Letters **71**: 341-345.
- Heal, G. (2008). "Climate economics: a meta-review and some suggestions for future research." Review of Environmental Economics and Policy **3**(1): 4-21.
- 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. (2006a). "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.
- Hope, C. (2006b). "The social cost of carbon: what does it actually depend on?" Climate Policy **6**(5): 566-572.
- Hope, C. (2008). "Discount rates, equity weights and the social cost of carbon." Energy Economics **30**(3): 1011-1019.
- IPCC (2007). Summary for Policymakers. Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. S. Solomon, D. Qin, M. Manning et al. Cambridge, United Kingdom and New York, NY, USA, Cambridge University Press.
- Mastrandrea, M. D. and S. H. Schneider (2004). "Probabilistic integrated assessment of "dangerous" climate change." Science **304**: 571-575.
- Meehl, G. A., T. F. Stocker, et al. (2007). Global climate projections. Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. S. Solomon, D. Qin, M. Manning et al. Cambridge, UK, Cambridge 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.
- Nordhaus, W. D. (2008). A Question of Balance: Weighing the Options on Global Warming Policies. New Haven and London, Yale University Press.
- Nordhaus, W. D. (2009). An analysis of the dismal theorem. Cowles Foundation Discussion Paper No. 1686, Yale University.
- Nordhaus, W. D. and J. Boyer (2000). Warming the World: Economic Models of Global Warming. Cambridge, Mass., MIT Press.
- Peck, S. C. and T. J. Teisberg (1992). "CETA: a model for carbon emissions trajectory assessment." Energy Journal **13**(1): 55-77.
- Roughgarden, T. and S. H. Schneider (1999). "Climate change policy: quantifying uncertainties for damages and optimal carbon taxes." Energy Policy **27**: 415-429.
- Sokolov, A. P., P. H. Stone, et al. (2009). "Probabilistic forecast for 21st century climate based on uncertainties in emissions (without policy) and climate parameters." Journal of Climate: doi: 10.1175/2009JCLI2863.1.
- Stainforth, D. A., T. Aina, et al. (2005). "Uncertainty in predictions of the climate response to rising levels of greenhouse gases." Nature **433**(7024): 403-406.
- Stern, N. (2007). The Economics of Climate Change: The Stern Review. Cambridge, UK, Cambridge University Press.

- Stern, N. (2008). "The economics of climate change: Richard T. Ely lecture." American Economic Review: Papers and Proceedings 98(2): 1-37.
- Tol, R. S. J. (1999). "Safe policies in an uncertain climate: an application of FUND." Global Environmental Change 9: 221-232.
- 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. (2005). "The marginal damage costs of carbon dioxide emissions: an assessment of the uncertainties." Energy Policy 33(16): 2064-2074.
- 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 27(1): 5-76.
- Warren, R., C. Hope, et al. (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.
- Warren, R., M. Mastrandrea, et al. (2009). Variation in the climatic response to SRES emissions scenarios in integrated assessment models. Mimeo. Norwich, Tyndall Centre, University of East Anglia.
- Weitzman, M. L. (2007). "A review of the Stern Review on the Economics of Climate Change." Journal of Economic Literature 45(3): 703-724.
- Weitzman, M. L. (2009). "On modeling and interpreting the economics of catastrophic climate change." Review of Economics and Statistics 91(1): 1-19.
- Yohe, G. W. (2003). "More trouble for cost-benefit analysis." Climatic Change 56(3): 235-244.

Figure 1. Climate sensitivities compared; Stern Review on the left, Stainforth et al. on the right.




Figure 2. The damage-function exponent (β) compared; standard PAGE on the left, fat-tailed distribution on the right.




Figure 3. Expected loss in social welfare as a function of the upper bound on damages (ϕ).




Table 1. Probability of a global ‘catastrophe’.

Climate sensitivity	Damages	$p[d_{j,t} \geq 0.75\hat{c}_{j,t}]$	$p[d_{j,t} \geq 0.9\hat{c}_{j,t}]$	$p[d_{j,t} \geq \hat{c}_{j,t}]$
Business-as-usual emissions (IPCC A2)				
Stern	Stern	0.1% (i.e. 1 in 1000)	0%	0%
Stainforth et al.	Lognormal	6.9%	5.5%	5.0%
550 ppm stabilisation (IPCC A2)				
Stern	Stern	0% (i.e. 0 in 1000)	0%	0%
Stainforth et al.	Lognormal	0.7%	0.6%	0.5%

Table 2. Probability of a regional ‘catastrophe’.

Climate sensitivity	Damages	$p[d_{j,t} \geq 0.75\hat{c}_{j,t}]$	$p[d_{j,t} \geq 0.9\hat{c}_{j,t}]$	$p[d_{j,t} \geq \hat{c}_{j,t}]$
Business-as-usual emissions (IPCC A2)				
Stern	Stern	3.8% (i.e. 38 in 1000)	2.8%	2.3%
Stainforth et al.	Lognormal	17.7%	14.4%	12.9%
550 ppm stabilisation (IPCC A2)				
Stern	Stern	0% (i.e. 0 in 1000)	0%	0%
Stainforth et al.	Lognormal	2.9%	2.3%	1.6%

Table 3. The marginal damage cost of carbon dioxide ($\phi = 0.99$).

Utility discount rate, δ	Elasticity of marginal utility of consumption, η	Marginal damage cost (year 2008 \$US/tCO ₂)		
		5%	mean	95%
Business-as-usual emissions (IPCC A2)				
0.1	1	30.69	430.93	1732.94
3	3	2.95	105.71	121.61
550 ppm stabilisation (IPCC A2)				
0.1	1	35.29	291.59	993.66
3	3	2.82	55.06	34.07