AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

The importance of ‘extremely unlikely’ events: tail risk and the costs of climate change*

John Quiggin[†]

In assessing the risks associated with climate change, ‘tail risks’ (low-probability extreme events) often play a much larger role than their probability alone might indicate. There are three main reasons for this: the linear relationship between sensitivity and warming; the convexity of the damage function; and the concavity of the utility function. Ignoring the upper tail of the distribution of possible outcomes will result in serious underestimates of the social cost of carbon dioxide (CO₂) emissions and of the socially optimal price for emissions.

Key words: Climate change, tail risk, low-probability events.

1. Introduction

The ‘dismal theorem’ of Weitzman (2007, 2014) shows that in problems involving a small risk of catastrophic damage, such as that of climate change, a fat-tailed distribution of possible damages may not have a finite expectation. This result has been the subject of considerable controversy (Karp 2009; Nordhaus 2012; Millner 2013; Pindyck 2013; Horowitz and Lange 2014). Much of this criticism focuses on the problems associated with taking limiting expectations in situations where the potential losses are unboundedly large. However, as Pindyck (2011, p. 269) observes ‘we don’t need a fat-tailed probability distribution to determine that ‘climate insurance’ is economically justified. All we need is a significant (and it can be small) probability of a catastrophe, combined with a large benefit from averting or reducing the impact of a catastrophic outcome’.

In this paper, Pindyck’s point is explored in detail. It is shown that, even in the absence of ‘fat tails’, and without considering ‘exceptionally unlikely’ events (probability 1 per cent or less) or unboundedly large losses, the central policy implication of the ‘dismal theorem’ remains valid as regards climate change; that is, the majority of the social cost of CO₂ emissions is associated with outcomes in the upper tail of the distribution of the key parameter of interest, climate sensitivity (conventionally measured as the equilibrium increase in temperature associated with a doubling of radiative forcing from greenhouse gases). Conversely, the optimal price of CO₂ emissions, calculated

* I thank Nancy Wallace for helpful comments and criticism.

[†] John Quiggin (e-mail: j.quiggin@uq.edu.au) is at School of Economics, The University of Queensland, St Lucia, Brisbane, Queensland, Australia.

over the full probability distribution of the sensitivity parameter, is likely to be substantially larger than an estimate based on a mean or median estimate.

More importantly, in calculations of the optimal level of mitigation, the relevant estimate of the marginal benefits of mitigation, expressed in expected present value terms, will be dominated by outcomes in the right hand (upper tail) of the distribution.

There are three main reasons for this. First, the higher the sensitivity the greater the warming associated with additional emissions. Second, because damage functions are convex, the greater the equilibrium warming, the greater the total and marginal damage associated with an additional degree of warming. Third, because the utility of income is concave, the greater the damage associated with climate change, the greater the risk-adjusted present value of additional damage at the margin.

These effects are mutually reinforcing. So if:

1. the value of sensitivity at the 90th percentile is twice the value at the median; and
2. damage is a quadratic function of warming;

then the 90th percentile of the distribution contributes four times as much to an estimate of the expected damage from climate change as does the 50th percentile. If, in addition,

3. the marginal utility of additional income is twice as high in the case of 95th percentile warming damage as in the median case,

then the 90th percentile of the distribution contributes four times as much to an estimate of the expected marginal damage from additional emissions as does the 50th percentile. Under these circumstances, a simple calculation suggests that the top 10 per cent of the sensitivity distribution will contribute around half the expected damage from warming, and around half of an estimate of the social value of marginal reductions in emissions.

In fact, as will be argued below, each of the estimates (i)–(iii) is likely to be conservative. This suggests that the upper tail of the sensitivity distribution will contribute the majority of the expected present value of mitigation and will largely determine the socially optimal price for CO₂ emissions. Conversely, estimates of total and marginal damage based on median sensitivity will be around half of the correct value, or even less.

The paper is organised as follows: Section 2 deals with the general problem of interpreting low-probability (less than 10 per cent) risks. It is shown that many of the risks with which we are commonly concerned fall into this category. Section 3 contains a formal framework relating climate sensitivity to the social cost of CO₂ emissions. The key result of the paper is that the expected value of the social cost of CO₂ emissions is dominated by costs arising from the right-hand tail of the distribution of values for the climate

sensitivity parameter. In Section 4, this result is illustrated for a variety of parameter values. Some policy implications are discussed in Section 5. Finally, some concluding comments are offered.

2. Understanding tail probabilities

Communication of probabilistic information is always difficult and has posed particular problems for the Intergovernmental Panel on Climate Change (IPCC) in its periodic assessment reports on the state of the global climate. In the Fifth Assessment Report (IPCC 2013), a standard convention was adopted to give verbal equivalents of probability ranges. The equivalences are as follows: virtually certain = 99–100 per cent probability; very likely = 90–100 per cent; likely 66–100 per cent; about as likely as not = 33–66 per cent; unlikely = 0–33 per cent; very unlikely = 0–10 per cent; extremely unlikely 0–5 per cent; exceptionally unlikely = less than 1 per cent.

This terminology is useful in communicating scientific uncertainty. However, it can be highly misleading in policy evaluation and risk analysis. It might be supposed that an outcome that is regarded as ‘very unlikely’ can be safely disregarded. Much discussion of climate sensitivity appears to incorporate this assumption, at least implicitly. This assumption is even more prevalent with respect to ‘extremely unlikely’ or ‘exceptionally unlikely’ outcomes.

However, many of the risks with which we are concerned in everyday decision-making are realised with probabilities in the range 0–5 per cent or 0–1 per cent, that is, in the IPCC terminology, ‘extremely unlikely’ or ‘exceptionally unlikely’. Most notably, the annual risk of death from any cause in developed countries such as the United States is of the order of 1 per cent (Index Mundi 2015), but it would be absurd to disregard the risk of death in making decisions about, for example, smoking or driving behaviour.

The critical point here is that a small risk of a severe adverse outcome may have a substantial effect on the expected costs of an individual or public policy choice. Unlikely events, including those that may be regarded as extremely unlikely (0–5 per cent) or exceptionally unlikely cannot be disregarded if their outcomes are extremely or exceptionally severe.

In this paper, this point will be illustrated by examining outcomes at or above the 90-th percentile of the probability distribution for climate sensitivity, parameterised by $\tilde{\lambda}$. Notationally, for any $p \in [1, 99]$ λ_p will denote the p -th percentile in the distribution of $\tilde{\lambda}$, with higher values of p corresponding to higher values of $\tilde{\lambda}$. The same convention will apply to other random variables, the values of which will depend on $\tilde{\lambda}$.

2.1 Climate sensitivity and uncertainty

The concept of climate sensitivity is commonly used as a simple guide to the likely extent of climate change (Colman and Braganza 2013). Climate sensitivity refers to the equilibrium adjustment of mean global temperature to

a doubling of the atmospheric CO₂ equivalent concentration¹ As the term is normally used in the literature, climate sensitivity is a property of climate models, informed by data, rather than a directly estimated empirical relationship. An important implication is that feedbacks (positive and negative) that are not included in a given model do not contribute to the model's estimates of sensitivity, or of the uncertainty surrounding that estimate. For example, there is evidence that global warming will contribute to the frequency and severity of forest fires, which will result in emissions of CO₂ from combustion, partially offset by subsequent regrowth. However, these effects are not well understood, so that this feedback is difficult to model accurately.

In general, the inclusion of additional feedbacks and other complexities may either increase or reduce mean and median estimates of sensitivity, but will normally increase the uncertainty surrounding those estimates. This fact explains the seeming paradox that, despite decades of research on climate modelling and the accumulation of vast quantities of data, neither the mean estimate of sensitivity nor the variance associated with that estimate changed significantly between the Second Assessment Report (IPCC 1996) and the Fifth Assessment Report (IPCC 2013).

Annan and Hargreaves (2006) argued that the procedures used by the IPCC overstate the uncertainty surrounding climate sensitivity and that a Bayesian approach to the available evidence would imply a posterior distribution tighter than that proposed by the IPCC. Against this, it must be observed that incomplete models will normally understate uncertainty. Furthermore, Annan and Hargreaves implicitly adopt the position that events that are 'extremely unlikely' or 'exceptionally unlikely' in the IPCC terminology, can safely be ignored. As noted above, from the perspective of risk analysis, this is a fundamental error.

3. Climate sensitivity and the social cost of CO₂ emissions

The relationship between climate sensitivity and the social cost of CO₂ emissions may be derived in a series of steps, each of which involves parameters about which there is significant disagreement and uncertainty. A complete analysis would take account of the joint distribution of climate sensitivity and the vector of model parameters. For simplicity, however, we will take these parameter values as given, and focus on the primary source of uncertainty, namely climate sensitivity. The other crucial variable is the cumulative volume of emissions measured by the concentration of CO₂ equivalents.

The analysis proceeds in four steps:

1. From climate sensitivity to equilibrium warming, expressed in °C;
2. From equilibrium warming to future damage, expressed as a percentage of future global income;

¹ The conversion of a variety of greenhouse gases into CO₂ equivalents raises a number of complex issues, which will not be discussed here.

3. From future damage to current social cost, expressed as a percentage of current global income; and
4. From social cost of warming to the optimal price of emissions, expressed in dollars per tonne of CO₂.

3.1 From climate sensitivity to equilibrium warming

We have

$$\tilde{\Delta} = \tilde{\lambda} * F(E), \quad (1)$$

where:

$\tilde{\Delta}$ is the equilibrium change in mean global temperature relative to the pre-industrial mean;

$\tilde{\lambda}$ is the climate sensitivity parameter;

F is radiative forcing, normalised so that a unit change corresponds to a doubling of atmospheric CO₂ concentrations; and

E is emissions, measured by the concentration of CO₂ equivalents.

Hence,

$$\frac{\partial \Delta}{\partial E} = \tilde{\lambda} f,$$

where $f = RF'(E)$ is the change in forcing associated with an additional tonne of CO₂ emissions.

Then, we may observe:

$$\frac{\Delta_{95}}{\Delta_{50}} = \frac{\lambda_{95}}{\lambda_{50}} 2.$$

3.2 From equilibrium warming to future damage

The future damage associated with any given level of equilibrium warming may be represented by a function $D(\lambda, E)$ where D is expressed as a proportion of global income in the absence of warming.

D may be evaluated at some future date, such as 2100, or as a present value of a stream of losses. For simplicity, the analysis here will focus on the first approach. A tractable representation of the damage function is:

$$D(\lambda, E) = (\Delta(\lambda, E) - \Delta^0)^\kappa D^1$$

where

Δ^0 is the equilibrium warming that can be sustained with minimal damage; $\kappa > 1$ is a parameter representing the convexity of the damage function; and D^1 is a constant, equal to the damage incurred when $\tilde{\Delta} = \Delta^0 + 1$.

The relationship between uncertainty and expected damage depends crucially on the convexity of damage. By Jensen's inequality, if damage is a

convex function of equilibrium warming, then expected damage will be greater than damage at the mean estimated equilibrium warming.

Standard principles of optimisation (Danskin's theorem and the theorem of the maximum) imply that the maximised value of the objective function in a production problem is convex. Hence, the loss associated with a change in input levels from an initially optimal position must be convex. Since temperature is not a choice variable, the long-run effect of a small change in temperature may be either positive or negative, depending on whether the initial temperature is above or below the climatic optimum for the technology in question. However, once the temperature increase is sufficient to exceed the climatic optimum, the effect of further warming must be negative and the damage function must be convex. Quiggin (2008) surveys estimates of the impacts of climate change on agriculture, all of which display convexity. As argued by Quiggin and Horowitz (2003), an analysis based on the long-term impacts of climate change fails to take account of the costs of adjustment to a change in climate. Such costs are always an increasing and convex function of the rate of warming.

This argument from first principles may be supported by informal and qualitative assessments of the impact of small, moderate and large increases in global temperatures. Consideration of the likely impacts of minimal, moderate and rapid warming shows that the function $D(\Delta)$ is strongly convex.

This point may be illustrated in qualitative terms as follows:

1. Warming already observed between 1880 and 2000 is estimated at 0.8°C and has had a mixture of beneficial and harmful effects, with no clear consensus on the net balance (American Institute of Physics 2017).
2. Warming of 1.5–2°C would have some adverse effects, but is generally accepted as tolerable, as is reflected in the Paris Agreement (United Nations 2015).
3. Warming of 4°C would have a wide range of serious effects (Christoff 2013) and could lead to catastrophic outcomes associated with 'tipping points', although this is inherently uncertain (Glikson 2014).
4. Warming of 8°C would be a catastrophic threat to the survival of all living systems, including human systems. (No assessment has been undertaken, as such a change is outside the current limits of scientific analysis.)

These points were represented graphically by the Intergovernmental Panel on Climate Change (2003) in Figure 1, sometimes referred to as the 'burning embers' diagram illustrates this. Figure 1 illustrates five different kinds of impacts of warming, with the severity represented by the intensity of the red shade. Categories I and II (risks to unique and threatened systems, and extreme climate events) show significant effects even for the modest warming (slightly less than 1°C) that has already taken place. This is consistent with actual experience in Australia, which includes an increase in the severity of droughts and coral bleaching events in the Great Barrier Reef.

Although the convexity of the damage function is generally accepted, there has been some debate about the appropriate representation of convexity. The default practice in optimisation analysis is to assume a quadratic loss function for deviations from the optimum. This practice was followed by Nordhaus (2012) in developing the Dynamic Integrated Climate-Economy (DICE) model, and by Mendelsohn *et al.* (1994) in estimating the long-term impacts of climate change in US agriculture. However, Dietz and Stern (2015) have argued that quadratic functions understate the convexity of the cost function.

The quadratic cost function corresponds to $\kappa = 2$. However, the arguments of Stern and Dietz suggest that this parametric choice is excessively conservative and that the true value of κ is greater than 2. In the illustrative calculations below (Section 4), we will use the value $\kappa = 2.25$, which corresponds to an income loss of 20 per cent for warming of 4°C, relative to the growth that would occur in the absence of damage from climate change.

We will be concerned primarily with the marginal damages from changes in emissions, given by:

$$\frac{\partial D}{\partial E} = \frac{\partial D}{\partial \Delta} \frac{\partial \Delta}{\partial E} = \kappa D^1 (\Delta - \Delta^0)^{\kappa-1} \tilde{\lambda} f.$$

Under these conditions, marginal damages are strongly increasing in λ . In particular, for $\Delta^0 = 0$ and $\kappa = 2$, (quadratic damages)

$$\frac{\partial D}{\partial E}_{|(2\lambda, E)} = 4 \frac{\partial D}{\partial E}_{|(\lambda, E)},$$

for any $\lambda, E > 0$. More generally, $\partial D/\partial E$ increases at least as fast as λ^2 .

As noted above, we may assume $(\Delta_{90}/\Delta_{50}) \geq 2$. This means that, in evaluating the expectation of $\partial D/\partial E$, the 90-th percentile of the distribution of $\tilde{\lambda}$ will contribute at least four times as much as the median (50-th percentile), and the top 10 per cent of the distribution will contribute at least 40 times as much as the 50-th percentile. Hence, the expectation E^* will be around 40 per cent greater than E_{50} , even with conservative choices of parameter values. For more realistic parameter values, the contribution of the upper tail is substantially larger.

3.3 From future damage to current social cost

The analysis so far has dealt with the impact of high values of λ on the expected value of marginal damage from emissions. However, in evaluating uncertain future costs and benefits it is important to take account of risk aversion. In standard economic models, benefits arising in the future are discounted for two reasons. First, a ‘pure’ discount factor may be applied. Second, and more importantly, because consumption levels are normally

Figure 1 Risks and impacts of climate change (IPCC Third Assessment Report 2003, Summary For Policymakers Figure SPM-2) (Burning Embers Diagram). [Colour figure can be viewed at wileyonlinelibrary.com]

expected to be higher in the future, the marginal value of additional consumption in the future is lower than that of additional consumption today, when our consumption needs are more urgent. The associated discount factor may be analysed using a utility function for income to evaluate an objective function

$$U = u(c^1) + \beta E[u(\tilde{c}_2)], \quad (2)$$

where:

U denotes total welfare; c^1 denotes consumption in period 1; $\tilde{c}_2$ denotes (stochastic) consumption in period 2; $\beta \leq 1$ is a discount factor reflecting pure time preference; and u is a von Neumann–Morgenstern utility function.

The most useful functional form for the utility function is given by the constant relative risk aversion (CRRA) family specified by:

$$u(x; \alpha) = \frac{x^{1-\alpha}}{1-\alpha}, \alpha \neq 1$$

$$u(x; \alpha) = \log(x), \alpha = 1,$$

where $\alpha > 0$ is the coefficient of relative risk aversion. The marginal utility of additional consumption is given by:

$$u'(x; \alpha) = x^{-\alpha}.$$

The marginal utility of consumption may be used to express the value of changes in the uncertain future consumption level c^2 in terms of current consumption c^1 . In the absence of uncertainty, 2 shows that a unit increase in c^2 yields the same additional utility as an increase in c^1 by:

$$\rho = \beta \frac{u'(c^2)}{u'(c^1)} = \beta \left(\frac{c^1}{c^2} \right)^\alpha,$$

where ρ is the present value of future income. Under the standard assumption of continued growth in consumption, $c^2 > c^1$ and hence:

$$\left(\frac{c^1}{c^2} \right)^\alpha < 1,$$

with strict inequality except in the risk-neutral case. So, given that $\beta \leq 1$, we have $\rho < 1$ that is, with continued growth, future consumption is discounted even in the absence of pure time preference, simply because consumption is expected to be higher in the future.

However, the assumption of continued growth in consumption may not hold in the case of uncontrolled climate change. If the damage is such as to reduce future living standards below those of today, then the logic of discounting is reversed, and increases in future consumption will be valued more highly than increases in present consumption. In the absence of pure time preference, which in this context amounts to placing less weight on the interests of later-born generations harmed by the actions of the earlier-born generation currently making the relevant decisions, $c^2 < c^1$ implies $\rho > 1$.

3.4 From social cost of warming to the optimal price of emissions

The final step in the analysis is to convert the present value marginal cost of an additional part per million (PPM) of CO₂ into an optimal price on emissions, expressed in dollars per tonne of CO₂ emitted. This is essentially an arithmetic step, requiring:

1. the conversion of damages from percentage points of global income into dollars, using an estimate of current income. The value used here is \$US75 trillion; and
2. the conversion of PPM, taking account of the fact that around half of all emissions are absorbed by sinks, primarily oceans.² From IPCC (2013), 1

² CO₂ absorbed into oceans is not harmless, since it increases the acidity of sea water which tends to dissolve limestone coral reefs. However, since the relationship between atmospheric and oceanic concentrations is monotonic, the damage function in 3.2 is sufficiently flexible to incorporate this additional source of damage.

PPM is equal to 7.8 Gigatonnes (Gt) of CO₂, while the atmospheric fraction (the proportion of emissions that remain in the atmosphere) is approximately 45 per cent. Hence, around 17 Gt of CO₂ emissions generate an increase of 1 PPM in atmospheric concentrations.

4. Parametric estimates

The parameters of interest in determining the expected marginal damage from emissions are as follows: λ_p , $p \in [1, 100]$ the percentiles of the sensitivity parameter λ ; E , the level of emissions; Δ^0 , the warming consistent with zero/minimal net damage; D^1 , the damage associated with warming of $\Delta^0 + 1$; and κ , the convexity parameter for the damage function.

Consider first λ_{50} , the median value of climate sensitivity. An estimate of 3.5°C for was first put forward by the National Academy of Science (1979) and has been maintained with minor variations ever since. We will assume that λ is normally distributed with mean 3.5. To determine the standard, we will use the statement of IPCC (2014) that climate sensitivity is ‘very unlikely’ to be above 6. As noted above, this term is defined as being associated with probability 0.10, implying that $\lambda_{90} = 6$. This implies a standard deviation of approximately 1.95. The estimates will be truncated (or, more precisely, winsorised) by setting the bottom 5 per cent of the distribution equal to λ_{05} and the top 5 per cent equal to λ_{95} .

We will consider emissions trajectories consistent with stabilisation at 450, 500 and 550 PPM.

Turning to Δ^0 , the warming consistent with zero/minimal net damage, several considerations suggest a value close to 1°C. First, in combination with $\lambda_{50} = 3$, this is consistent with stabilising CO₂ equivalent concentrations at 350 PPM, widely regarded as the maximum safe level. Second, we have already experienced at least 0.5° of warming over the course of the 20th century, without major ill effects. Next, studies of the impact of climate change on agriculture, the human activity most directly affected by climate, find minimal net adverse effects, or even net benefits, for this rate of warming. Further, as a matter of parametric convenience, this choice means that D^1 is the damage associated with 2° of warming, the currently agreed target for international climate agreements. We will set D^1 equal to 2 per cent of $\bar{c}_2$, the future consumption level in the absence of climate change. Finally, as discussed above, we will set $\kappa = 2.5$.

4.1 Results

The results for emissions trajectories consistent with stabilisation at 450 and 550 PPM, respectively. These are presented in Tables 1 and 2. The tables present results for each 5th percentile of the distribution of λ , from 2.5 to 97.5, and also for the 99th percentile. Each of the percentile values from 2.5 to 97.5

is interpreted as representative of the vigintile range³ of which it is the mid-point. For example, $\lambda_{97.5}$ is interpreted as representative of the top 5 per cent of the distribution.

In addition to the percentile values in Column 1, Tables 1 and 2 report the following variables:

Δ the equilibrium change in temperature in °C (Column 2);

D the damage associated with warming, in percentage points of future income (Column 3);

MD the marginal damage, in percentage points of future income per 10 PPM of emissions (Column 4);

DU the risk premium adjustment, a ratio with 1 implying no adjustment (Column 5);

P the implied price for carbon emissions in \$/tonne (Column 6); and Share: the proportion of the mean implied price for carbon emissions contributed by the vigintile associated with λ_p (Column 7).

The final two rows of the tables show, for each of the variables in Columns 1 to 7, the mean value and the ratio of the mean.

4.1.1 Stabilisation at 450 PPM

We will first consider Table 1, associated with the internationally agreed target of stabilisation at 450 PPM. Column 1 shows the distribution of λ , with mean 3.5 and standard deviation 1.95. As shown in Column 2, the median estimate of the warming Δ associated with the mean value of $\lambda = 3.5$ and equilibrium concentration of 450 PPM is 2.1°, which is just above the stated goals of international agreements (2° or less). However, the warming associated with high values of λ is substantially greater, with $\Delta_{97.5} = 4.4$ and $\Delta_{99} = 4.95$. Even if the agreed target for emissions is reached, there is a probability of around 5 per cent of exceeding 4° of warming and a probability of 1 per cent of exceeding 5° of warming. In the IPCC terminology, such low-probability events are ‘extremely unlikely’ or ‘exceptionally unlikely’. However, in the context of risk assessment, catastrophic risks with probabilities of 1 to 5 per cent are commonly central to the analysis.

Column 3 shows the distribution of damage expressed as a percentage reduction in the annualised present value of income. Even with successful stabilisation at 450 PPM, the upper tail of the distribution involves substantial damage from climate change.

Column 4 shows that marginal damage, like total damage, is an increasing function of sensitivity. Around the median, an additional 10 PPM of CO₂ equivalent concentrations would reduce future income by 1 to 1.5 per cent. However, at the 97.5 percentile, $MD_{97.5} = 21.5$.

³ The vigintile associated with λ_p is the set of values from $\lambda_{p-0.025}$ to $\lambda_{p+0.025}$.

Table 1 Distribution of warming cost measures by percentile of climate sensitivity (vigintile midpoints), with stabilization at 450 ppm

Percentile	$\Delta\ddagger$	Damage $\ddagger$	Marginal damage $\S$	Risk premium $\P$	Implied CO ₂ price (\$/t) $\dagger\dagger$	Share of implied price $\ddagger\ddagger$
2.5	0.00	0.0	0.0	1.00	\$0.00	0.0
7.5	0.38	0.0	0.0	1.00	\$0.00	0.0
12.5	0.73	0.0	0.0	1.00	\$0.00	0.0
17.5	0.99	0.0	0.0	1.00	\$0.00	0.0
22.5	1.21	0.0	0.0	1.00	\$1.86	0.1
27.5	1.40	0.1	0.1	1.00	\$5.73	0.2
32.5	1.57	0.2	0.1	1.01	\$11.16	0.5
37.5	1.74	0.5	0.2	1.01	\$18.05	0.8
42.5	1.90	0.8	0.2	1.02	\$26.45	1.1
47.5	2.05	1.1	0.3	1.02	\$36.52	1.6
52.5	2.20	1.6	0.4	1.03	\$48.52	2.1
57.5	2.35	2.1	0.5	1.04	\$62.88	2.7
62.5	2.51	2.8	0.7	1.06	\$80.19	3.5
67.5	2.68	3.6	0.9	1.08	\$101.42	4.4
72.5	2.85	4.7	1.1	1.10	\$128.10	5.5
77.5	3.04	6.0	1.3	1.13	\$162.86	7.0
82.5	3.26	7.7	1.6	1.17	\$210.75	9.1
87.5	3.52	10.1	2.1	1.24	\$283.07	12.2
92.5	3.87	14.0	2.8	1.35	\$413.55	17.8
97.5	4.41	21.5	4.1	1.62	\$729.54	31.4
Mean values						
50.0	2.133	0.038	0.008	1.094	116.032	5.0

$\dagger$ Equilibrium change in global mean temperature. $\ddagger$ As percentage of income. $\S$ As percentage of income. $\P$ Ratio of mean income to certainty equivalent. $\dagger\dagger$ Socially optimal price of carbon, derived as risk-adjusted expected marginal damage. $\ddagger\ddagger$ Proportion of price associated with vigintile (5 percentiles) of sensitivity distribution.

Column 5 shows the relative risk premium arising from the effects of climate change on the marginal utility of future income derived as risk-adjusted expected marginal damage. For values of λ up to and including the median this value is near 1, indicating that the loss of income associated with climate change is sufficiently small that its effect on the marginal utility of future income may be safely disregarded. By contrast, at the 95th percentile, the value is 1.62, indicating that the marginal utility of future income will be 60 per cent higher than in the median case.

Column 6 shows the implied socially optimal price of carbon. The median value is around \$40/tonne, close to the \$38/tonne value often used in benefit-cost analysis. However, for sensitivities in the upper tail of the distribution, the implied carbon price is well above \$300/tonne. The mean value is \$116/tonne, more than twice as much as the median.

Column 7 shows the share of the estimated mean price contributed by each vigintile. The top two vigintiles, corresponding to values of λ above the 90th percentile, contribute nearly 50 per cent of the total.

The analysis here is undertaken using the mid-point of each vigintile, so that the highest value of sensitivity considered is the 97.5 percentile value

Table 2 Distribution of warming cost measures by climate sensitivity (vigintile midpoints) with stabilization at 550 ppm

Percentile	$\Delta^\dagger$	Damage ‡	Marginal damage §	Risk premium ¶	Implied CO ₂ price (\$/t) ††	Share of implied price ‡†
2.5	0.00	0.0	0.0	1.00	\$0.00	0.0
7.5	0.60	0.0	0.0	1.00	\$0.00	0.0
12.5	1.16	0.0	0.0	1.00	\$0.00	0.0
17.5	1.57	0.2	0.0	1.00	\$0.00	0.0
22.5	1.92	0.8	0.2	1.02	\$17.52	0.0
27.5	2.22	1.7	0.3	1.03	\$31.70	0.0
32.5	2.50	2.8	0.4	1.06	\$49.59	0.1
37.5	2.76	4.1	0.6	1.09	\$71.62	0.1
42.5	3.01	5.7	0.8	1.13	\$98.60	0.1
47.5	3.25	7.6	1.0	1.17	\$131.81	0.2
52.5	3.50	9.8	1.3	1.23	\$173.20	0.2
57.5	3.74	12.4	1.6	1.30	\$225.83	0.3
62.5	3.99	15.5	1.9	1.40	\$294.63	0.3
67.5	4.25	19.0	2.3	1.53	\$388.06	0.4
72.5	4.53	23.4	2.8	1.70	\$521.88	0.6
77.5	4.83	28.7	3.4	1.97	\$729.00	0.8
82.5	5.18	35.7	4.1	2.42	\$1,090.85	1.3
87.5	5.59	45.2	5.1	3.33	\$1,874.45	2.2
92.5	6.15	60.2	6.7	6.32	\$4,642.69	5.4
97.5	7.00	88.3	9.5	72.61	\$76,353.71	88.1
Mean values						
50.000	3.39	0.18	0.02	5.22	4,334.76	5.0

† Equilibrium change in global mean temperature. ‡ As percentage of income. § As percentage of income. ¶ Ratio of mean income to certainty equivalent. †† Socially optimal price of carbon, derived as risk-adjusted expected marginal damage. ‡† Proportion of price associated with vigintile (5 percentiles) of sensitivity distribution.

(sensitivity of 7.2°). Consideration of the 99th percentile values would yield an even larger divergence between mean and median values. At the 99th percentile of the sensitivity distribution (not reported), the marginal social cost of emissions is over \$1300/tonne.

4.1.2 Stabilisation at 550 PPM

Table 2 presents results for stabilisation at 550 PPM, which is the concentration level implied by commitments made prior to the Paris Conference in 2015. For this atmospheric concentration, the top 10 per cent of the sensitivity distribution is associated with warming of 6° or more, which would take the global climate far outside the range of human experience. Estimates of damage associated with such unprecedented warming are necessarily speculative. Such an outcome would certainly be associated with mass extinctions and reductions in food production on a massive scale, but we have, at present, no basis on which to form accurate estimates. Moreover, estimates of the monetary value of such damage have little meaning if, as would inevitably be the case, relative prices change so radically as to make comparison of present and future consumption bundles inappropriate. The results derived

in Table 2 for the top 10 per cent of the sensitivity distribution should not be regarded as model estimates; rather, they should be taken as an illustration of the outcomes when a model is pushed beyond its range of applicability.

With these qualifications, it should be observed that the model results indicate that calculations of total damage and marginal social cost are entirely dominated by the top vigintile, that is by the worst 5 per cent of outcomes, which account for 88 per cent of the marginal social cost of carbon. While the quantitative estimates are unreliable, the qualitative conclusion that 'extremely unlikely' values of sensitivity account for a large share of the expected costs of climate change is robust.

For concentrations of 600 PPM or higher, commonly associated with 'Business As Usual' projections of carbon emissions, the top 20 per cent of the sensitivity distribution yields more than 6° of warming, and the median warming is around 4°. Using the IPCC terminology, a catastrophic outcome is as likely as not, rendering a focus on the upper tail of the distribution largely irrelevant. For this reason, results are not reported for concentrations of 600 PPM or higher.

5. Policy implications

The analysis above shows that a large proportion, and possibly the majority, of the expected damage from global warming is associated with low-probability (less than 10 per cent) events in the upper tail of the distribution of climate sensitivity. One way of responding to this observation is to partition the expected damage into two components: a 'best estimate' component associated with the median value for sensitivity and an 'insurance' component associated with the upper tail of the sensitivity distribution. For some policy variables, such as the optimal social price of carbon, this distinction is, in practical terms, irrelevant. The optimal price is the expected value of marginal damage associated with emissions, and is the same however the probability distribution is partitioned.

For other aspects of policy, however, the distinction assumes greater significance. In particular, there are significant implications for the targets agreed at the Paris Conference of Parties in 2015, and the associated emissions trajectories. The parties agreed on the need for aggregate emission pathways consistent with holding the increase in the global average temperature to well below 2°C above pre-industrial levels and pursuing efforts to limit the temperature increase to 1.5°C.

For median values of sensitivity, a stabilisation target of 450 PPM would be consistent with equilibrium warming of 2°C. However, for sensitivity values in the upper decile of the distribution, the implied equilibrium warming is above 4°C. Hence, the agreed target for equilibrium warming implies a conditional target for emissions reductions, with more rapid reductions in emissions required for higher values of sensitivity.

For sufficiently high values of sensitivity, the implied stabilisation target is below the current concentration of 400 PPM, implying the need for negative net emissions. Only preliminary consideration has been given so far to the options for achieving negative net emissions. The options include the following:

- negative net carbon emissions achieved through a large net increase in global forest cover;
- reductions in emissions of methane, which has a relatively short residence time;
- direct extraction of CO₂ from the atmosphere, through technologies yet to be developed and tested; and
- 'geoengineering' technologies, such as measures to increase absorption of CO₂ by oceans.

All of these options are problematic in one way or another. Perhaps the most problematic is geoengineering. This term is used to encompass a variety of strategies involving radical human-induced changes in atmospheric and oceanic systems, aimed at offsetting the warming effects of greenhouse gas emissions.

Uncertainty about sensitivity values will be resolved over time. The dependence of cost estimates on sensitivity values in the tail of the distribution implies both a high value for information and a high value for flexibility.

With more information, uncertainty about sensitivity can be resolved earlier. If the value of sensitivity is found to be low, high-cost and high-risk options such as geoengineering can be ruled out. On the other hand, if the value of sensitivity is found to be high, early information will permit a rapid acceleration in mitigation efforts.

The value of flexibility (closely related to the concept of option value) is associated with strategies that allow for such a rapid acceleration in mitigation efforts, at relatively low cost, in the event that sensitivity is found to be high. For example, investment in the development of a diverse range of emissions reductions technologies increases the likelihood that a technology capable of rapid and large-scale implementation will be available if needed. Such diversity is unlikely to arise from the market incentives associated with a moderate carbon price, unless the set of policies incorporates measures to finance the development of backstop options that may never be used.

6. Concluding comments

In a properly formulated risk analysis, 'tail risks' (low-probability extreme events) often play a much larger role than their probability alone might indicate. As has been shown in this paper, the problem of tail risk is significant in the context of climate change. There are three main reasons for

this: the linear relationship between sensitivity and warming; the convexity of the damage function; and the concavity of the utility function. These factors combine in a multiplicative fashion. Ignoring the upper tail of the distribution of possible outcomes will result in estimates of the social cost of CO₂ emissions that are less than half the correct value, and sometimes much less than half. A focus on the median value for sensitivity implies a serious underestimate of the severity of the climate change problem and of the magnitude of the measures needed to deal with it.

References

- American Institute of Physics (2017). The Discovery of Global Warming: Impacts of Climate Change. Available from URL: <http://history.aip.org/climate/impacts.htm> [Accessed 30 May 2017].
- Annan, J. and Hargreaves, J. (2006). Can we believe in high climate sensitivity? Arxiv preprint physics/0612094.
- Christoff, P. (2013) *Four Degrees of Global Warming: Australia in a Hot World*. Routledge, London and New York.
- Colman, R. and Braganza, K. (2013). 'Explainer: what is climate sensitivity?' <https://theconversation.com/explainer-what-is-climate-sensitivity-18815> [Accessed 30 May 2017].
- Dietz, S. and Stern, N. (2015). Endogenous growth, convexity of damage and climate risk: How Nordhaus' framework supports deep cuts in carbon emissions, *Economic Journal* 125, 574–620.
- Glikson, A. (2014). What climate tipping points should we be looking out for? *The Conversation*, <https://theconversation.com/what-climate-tipping-points-should-we-be-looking-out-for-27108>
- Horowitz, J. and Lange, A. (2014). Cost-benefit analysis under uncertainty: a note on Weitzman's dismal theorem, *Energy Economics* 42, 201–203.
- Index Mundi (2015). United States Death rate. Available from URL: http://www.indexmundi.com/united_states/death_rate.html [accessed 25 May 2016].
- Intergovernmental Panel on Climate Change (1996). *Climate Change Third Assessment Report*. IPCC, Geneva.
- Intergovernmental Panel on Climate Change (2003). *Third Assessment Report, Working Group I Report (WGI): Climate Change: Summary for Policymakers*. IPCC, Geneva.
- Intergovernmental Panel on Climate Change (2013). *Climate Change Fifth Assessment Report*. IPCC, Geneva.
- Intergovernmental Panel on Climate Change (2014). *Climate Change 2014 Synthesis Report Summary for Policymakers*. IPCC, Geneva.
- Karp, L.S. (2009). Sacrifice, discounting and climate policy: five questions, CESifo Working Paper Series.
- Mendelsohn, R., Nordhaus, W. and Shaw, D. (1994). The impact of global warming on agriculture: A Ricardian analysis, *American Economic Review* 84, 753–771.
- Millner, A. (2013). On welfare frameworks and catastrophic climate risks, *Journal of Environmental Economics and Management* 65, 310–325.
- National Academy of Science (1979). *Carbon Dioxide and Climate: A Scientific Assessment*. NAS, Washington, DC.
- Nordhaus, W.D. (2012). Economic policy in the face of severe tail events, *Journal of Public Economic Theory* 14, 197–219.
- Pindyck, R.S. (2011). Fat tails, thin tails, and climate change policy, *Review of Environmental Economics and Policy* 5, 258–274.

- Pindyck, R. (2013). Climate Change Policy: What Do the Models Tell Us?, *Journal of Economic Literature* 51, 860–872.
- Quiggin, J. (2008). Counting the cost of climate change at an agricultural level, *CAB Reviews: Perspectives in Agriculture, Veterinary Science, Nutrition and Natural Resources* 2(092), 1–9.
- Quiggin, J. and Horowitz, J. (2003). Costs of adjustment to climate change, *Australian Journal of Agricultural and Resource Economics* 47, 429–446.
- United Nations (2015). Paris Agreement, http://unfccc.int/files/essential/_background/convention/application/pdf/english/_paris_agreement.pdf.
- Weitzman, M. (2007). Structural uncertainty and the value of statistical life in the economics of catastrophic climate change. *National Bureau of Economic Research*, <https://doi.org/citesee-rx.ist.psu.edu/viewdoc/download?doi=10.1.1.317.9499&rep=rep1&type=pdf>
- Weitzman, M.L. (2014). Fat tails and the social cost of carbon, *The American Economic Review* 104, 544–546.