Climatic Change (2014) 124:21–37
DOI 10.1007/s10584-014-1082-7

Scientific uncertainty and climate change: Part I. Uncertainty and unabated emissions

Stephan Lewandowsky · James S. Risbey · Michael Smithson · Ben R. Newell · John Hunter

Received: 27 September 2012 / Accepted: 7 February 2014 / Published online: 4 April 2014

Abstract Uncertainty forms an integral part of climate science, and it is often used to argue against mitigative action. This article presents an analysis of uncertainty in climate sensitivity that is robust to a range of assumptions. We show that increasing uncertainty is necessarily associated with greater expected damages from warming, provided the function relating warming to damages is convex. This constraint is unaffected by subjective or cultural risk-perception factors, it is unlikely to be overcome by the discount rate, and it is independent of the presumed magnitude of climate sensitivity. The analysis also extends to "second-order" uncertainty; that is, situations in which experts disagree. Greater disagreement among experts increases the likelihood that the risk of exceeding a global temperature threshold is greater. Likewise, increasing uncertainty requires increasingly greater

Part II of this paper is published under doi:10.1007/s10584-014-1083-6.

Electronic supplementary material The online version of this article (doi:10.1007/s10584-014-1082-7) contains supplementary material, which is available to authorized users.

S. Lewandowsky (✉)
University of Western Australia, Perth, WA, Australia
e-mail: stephan.lewandowsky@uwa.edu.au

S. Lewandowsky
University of Bristol, Bristol, UK

J. S. Risbey
CSIRO Marine and Atmospheric Research, Hobart, Tasmania, Australia

M. Smithson
Australian National University, Canberra, Australia

B. R. Newell
University of New South Wales, Sydney, Australia

J. Hunter
Antarctic Climate & Ecosystems Cooperative Research Centre, Hobart, Tasmania, Australia

protective measures against sea level rise. This constraint derives directly from the statistical properties of extreme values. We conclude that any appeal to uncertainty compels a stronger, rather than weaker, concern about unabated warming than in the absence of uncertainty.

1 Introduction

Uncertainty is an unavoidable aspect of scientific endeavors. The IPCC's AR4 of 2007 used the word "uncertain" or its derivatives more than 1200 times in the report of Working Group 1 alone—around 1.2 times per printed page. Although the scientific community has sought to develop ways of dealing with uncertainty (e.g., Intergovernmental Panel on Climate Change 2005; Narita 2012), scientific uncertainty has often been highlighted in public debates to preclude or delay political action on contentious issues (e.g., Freudenburg et al. 2008; Freudenburg and Muselli 2013).

Political appeals to uncertainty tend to follow two streams of argumentation: The first stream asymmetrically focuses on the possibility that the problem (e.g., climate change) may be less serious than anticipated and that any potential surprises are more likely to be benign rather than inauspicious. The second stream paints scientific uncertainty as a "monster" that is not readily managed (Van der Sluijs 2005). Accordingly, some have argued that science makes environmental controversies *worse* (Sarewitz 2004), implying that research is impotent in informing public policy.

We argue that those streams of argumentation misconstrue the impact of scientific uncertainty on policy choices. In the remainder of this article we show that greater uncertainty about climate change implies a greater probability of adverse consequences. In a companion article, we analyze the implications of uncertainty on mitigation (Lewandowsky et al. 2014) and show similarly that greater uncertainty implies a greater, not lesser, impetus for mitigative action.

2 Framing uncertainty: risks and outcomes

Figure 1 presents two possible global "policy responses" (in rows) and two possible states of the world (in columns), with the possible outcomes highlighted in the corresponding cells.

By dichotomizing a range of possible decisions and states of the world, the figure necessarily simplifies. Nonetheless, it highlights several core issues. First, it emphasizes a fact that receives insufficient attention for historical and psychological reasons (Lewandowsky et al. 2013); namely, that the decision *not* to cut emissions is not an inactive "null" default. Instead, withholding emission cuts and continuing with "business as usual" equates to an active decision to add greenhouse gases to the atmosphere. This realization is non-trivial because it reframes the issue from whether or not societies "should do something," to an acknowledgment that we *are* doing something already—adding CO_2 to the atmosphere—and that the consequences of that *action* must be contrasted with the costs of alternative actions.

This raises the second point made by Fig. 1, that the outcomes of actions must be evaluated with respect to the states of the world. There are two main approaches by which the constraints implied by Fig. 1—and the associated uncertainties—are conventionally resolved: One approach relies on economic cost-benefit analysis and the other invokes the precautionary principle.

Policy Response	Actual climate Sensitivity: Sensitivity < 1°C	Actual climate Sensitivity: Sensitivity ≥ 2°C
Rigorous mitigation	Climate change turns out to be harmless. Economic costs of mitigation incurred unnecessarily.	Climate threat turns out to be significant and responses reduce the size of the impact.
Continue to emit	Climate change turns out to be harmless. Economic costs of mitigation avoided.	Climate threat turns out to be significant and is enhanced due to lack of mitigation. Impacts maximal.

Fig. 1 Contingency table relating the (in principle unknowable) true state of the world (*columns*) to possible policy responses (*rows*). Cell entries are the consequences that are likely to arise from each policy responses contingent on the two possible states of the world. For simplicity, the range of possible equilibrium climate sensitivities (i.e., the ultimate response of the climate system to a doubling in CO_2 levels from preindustrial times) is dichotomized into values below 1 °C, which would entail only limited adverse consequences, and values in the vicinity or above 2 °C, which would lead to serious adverse consequences on an unabated emissions path

2.1 Cost-benefit analysis

Cost-benefit analysis (CBA) considers the cost of various emission paths against the associated benefits accrued by a reduction or avoidance of damages from climate change (e.g., Garnaut 2011; Nordhaus 2010; Stern 2007; Tol 2011). Uncertainty can be represented in CBA by probability-weighting the competing costs (cf. Schneider 2002).

Several limitations curtail the utility of CBA: First, notwithstanding its apparent objectivity, CBA is not free of ethical considerations (Aldred 2009; Nolt 2011; Risbey 2006), even if they remain tacit. This problem is most apparent in connection with the monetarization of environmental "goods" such as species diversity (Ressurreição et al. 2011), the existence of songbirds (Funtowicz and Ravetz 1994), or indeed human life (cf. Li et al. 2010). Second, CBA is affected by people's attitudes toward risk (Nordhaus 2011) and is therefore necessarily influenced by cultural and political factors (e.g., Kahan et al. 2006; Slovic 1999). To illustrate, Ackerman et al. (2005) showed retrospectively that CBA would have prevented implementation of a number of regulatory public-health measures, among them limiting workplace exposure to vinyl chloride which likely saved many lives (Michaels 2008).

Third, because climate mitigation costs would be incurred now, whereas damages from climate change will largely come due in the future, CBA must necessarily discount those future costs against present expenditure by applying an interest rate (the discount rate; e.g., Anthoff 2009). Decisions about the discount rate have drastic effects on CBA; if the discount rate is sufficiently high, any future cost, no matter how great, will appear minuscule compared to present-day mitigation costs (e.g., Weitzman 2010a). For example, with a discount rate of 1 % per annum, a presumed damage of $1,000,000 in 300 years is worth around $50,000 today. With a rate of 5 %, the discounted value is less than 50 cents (Sterner and Persson 2008). Relatively small variations in discounting can thus alter the anticipated damage cost by orders of magnitude—thereby undermining the robustness of CBA.

2.2 Precautionary principle

An alternative to CBA invokes the precautionary principle (e.g., Vlek 2010a, b). Put in its most succinct form, the principle holds that if there is a potential for harm from an activity, and if there is uncertainty about the magnitude of those impacts, then action should be taken to avoid that harm (e.g., Gardiner 2006). Unlike CBA, the precautionary principle is asymmetrical because it focuses on the bottom-right cell in Fig. 1, thereby effectively reversing the "burden of proof": Unless an activity (or product) can be shown to have no harmful consequences, it should cease. Thus, on the precautionary principle, action is triggered *because* there is uncertainty about an outcome, not despite of any uncertainty.

Like CBA, the precautionary principle has been subject to criticism (e.g., Feintuck 2005; John 2010; Kahan et al. 2006; Peterson 2006; Vlek 2010a, b). For example, the principle has been labeled incoherent (Peterson 2006) because if any precautionary action (e.g., cutting CO_2 emissions) itself leads to adverse outcomes (e.g., risks from expanding nuclear power), then the precautionary principle provides no guidance about how to resolve the conundrum. Further criticism cites the ill-specified role of uncertainty, which by an extreme interpretation might suggest that an activity should be banned if there is any possibility, no matter how small, that it might prove harmful (cf. Gardiner 2006).

We therefore suggest that neither CBA nor the precautionary principle are sufficient to resolve the dilemma posed by Fig. 1. We argue instead that the implications of Fig. 1 are best explored by deriving uncertainty-based constraints that are (a) not subject to cultural and personal vagaries of risk perception, (b) are unlikely to be affected by the discount rate, and (c) involve few if any ethical considerations. Our approach rests on developing ordinal constraints—i.e., constraints of the form "greater than" or "lesser than"—that derive from the functional form of the mapping between uncertainty and outcomes.

3 Uncertainty and unabated emissions

In the remainder of this article, we consider the role of uncertainty with respect to the policy decision to continue with unabated emissions (bottom row of Fig. 1). The complementary role of uncertainty with respect to mitigation is addressed in a companion article (Lewandowsky et al. 2014), in which we also show how uncertainty-based constraints might inform policy choices via a simple decision model, such as a "safe-minimum-standard" (i.e., specifying the maximum tolerable risk of exceeding a temperature threshold).

We differentiate and delineate the concepts of risk and various types of uncertainty as follows. We use *risk* to refer to the set of possible consequences of climate change, each with quantifiable probabilities and losses (Schneider 2002). We use *uncertainty*, by contrast, to refer to the imprecision of our knowledge of various crucial climate variables, which is typically captured by the variance of the variable's estimate (cf. Padilla et al. 2011). We are primarily concerned with uncertainty in equilibrium climate sensitivity (ECS); that is, uncertainty about the warming ultimately expected in response to a doubling of CO_2 from preindustrial times. For convenience, we assume that CO_2 will double, from approximately 275 ppm pre-industrially to 550 ppm, such that the expected global temperature increase equals the presumed sensitivity.

The value of climate sensitivity is constrained by several sources of evidence, ranging from paleoclimatology (e.g., Hegerl et al. 2006; Zeebe et al. 2009) to analysis of the observational record in conjunction with climate modeling (e.g., Bender et al. 2010). Those multiple lines of evidence converge on a point estimate of ECS of around 3 °C, with a likely range from 2 °C to 4.5 °C, and a lower bound of 1.5 °C (e.g., Knutti and Hegerl 2008; Meehl et al. 2007). Values substantially higher than 4.5 °C remain subject to debate, with some arguing that it ought to be considered as a firm upper bound (Annan and Hargreaves 2011) and others suggesting the contrary (Roe and Armour 2011).

In light of those multiple converging constraints, we begin by focusing on "first-order" uncertainty (cf. Walley 1991); that is, the error surrounding the (single) estimate of a parameter that is captured by the variance of its inferred probability density function. To prevent confusion, we call this quantity uncertainty$_{\text{ECS}}$ (for *E*quilibrium *C*limate *S*ensitivity) from here on. We introduce other subscripts for other manifestations of uncertainty where necessary. We later extend our analysis to "second-order" uncertainty; where the nature of the probability distribution is uncertain, so that its variance and higher moments are unknown. Second-order uncertainty may arise from lack of relevant information, but also when there are divergent estimates of a parameter such as climate sensitivity (cf. Smithson 1999).

3.1 Asymmetrical tails

Extant climate sensitivity distributions are asymmetric: they are thought to have a "fat" upper tail and a comparatively abruptly truncated lower tail. Thus, the area of the distribution below its mean is more sharply bounded than the area above: values of climate sensitivity far above the central location of the distribution are more likely than values far below (Roe and Baker 2007). The fat upper tail implies that particularly severe consequences arising from CO_2 emissions cannot be ruled out (Weitzman 2011). This concern is supported by episodes of rapid climate change in the geological past (e.g., Bahn et al. 2011; Holmes et al. 2011).

A further consideration concerns the *magnitude* of uncertainty$_{\text{ECS}}$. To examine the impact of the magnitude of uncertainty$_{\text{ECS}}$, its effects must be disentangled from the value of the "best estimate" for ECS. This can be accomplished in several ways, depending on one's choice of "best estimate", or measure of central location of the sensitivity distribution. Figure 2 illustrates two possibilities, one using the mode (panel A) and the other using the mean (panel B), as the measure of location that is to be disentangled from the distribution's spread (i.e., standard deviation). Each panel shows three arbitrary fat-tailed probability density functions of differing spread. In panel A, the mode (labeled *Mo*) is kept constant as the standard deviation (σ) increases. This approach might appear appealing because it creates an increasingly fatter tail while keeping the most likely estimate unchanged. However, one consequence of keeping the mode constant while increasing spread is that the mean (vertical lines μ) increases together with the standard deviation. It follows that the consequences of increasing uncertainty (σ) cannot be disentangled from the contribution of the increasing mean, thereby preventing an unambiguous interpretation. Note that the median also increases with increasing uncertainty.

This problem can be circumvented by keeping the mean of the distributions constant while increasing its spread (panel B in Fig. 2). This approach is known as a mean-preserving examination of uncertainty and is common in economics and finance (e.g., Hartman 1972). The approach is conservative because as uncertainty increases with the mean constant, the mode of the distribution (like its median) *decreases*; hence the effects of increasing uncertainty cannot be attributed to a concomitant increase in any measure of central location. We therefore apply the mean-preserving approach throughout.

We demonstrate the effects of a mean-preserving increase in uncertainty using a simulated lognormal distribution of ECS. The lognormal was chosen for illustrative purposes because it has a fat upper tail, but this choice entails no commitment to the precise shape of the sensitivity distribution. Table 1 summarizes the parameter values and results based on 10,000 samples. The expected value of the distribution, μ_L, was kept constant in all simulation runs, but the uncertainty of its estimate—that is, the standard deviation of the distribution; σ_L—increased from small to considerable across runs.

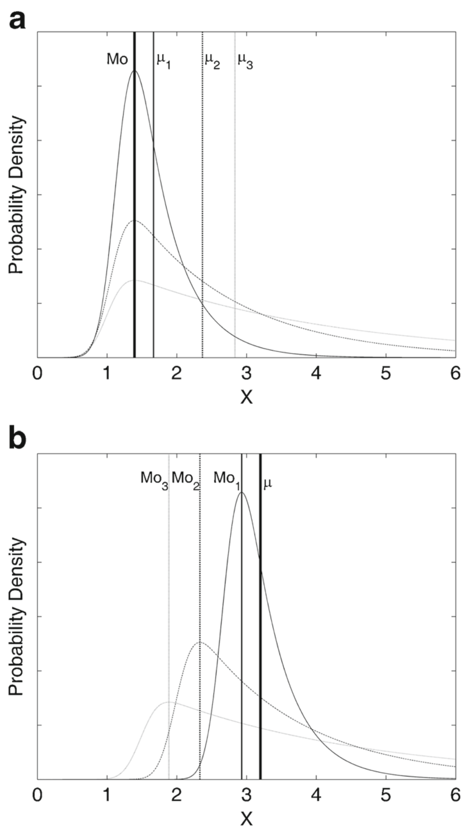

Fig. 2 The effects of increasing the spread (standard deviation, σ) on an arbitrary but fat-tailed probability density function obtained by convolving a Gaussian and an exponential distribution. (The convolution of those two distributions, defined by three parameters, lends itself particularly well to manipulation of mean, standard deviation, and "fatness" of the tail.) Panel **a** shows the effects of keeping the mode (*Mo*) constant at 1.4 while increasing the distribution's standard deviation, σ, from .54 to 1.14 and 1.36, respectively. The accompanying increase in means is denoted by μ_1 through μ_3. Panel **b** shows the effects of keeping the mean (μ) constant at 3.2 while increasing σ from .52 to 1.0 and 1.24, respectively. Mo_1 through Mo_3 denote the accompanying decrease in modes

Table 1 illustrates two principal consequences of increasing uncertainty$_{\mathrm{ECS}}$. First, as the spread of the distribution increases, the probability of a gravely concerning outcome also increases (columns $P(X > T_c)$; i.e., warming in excess of 5 °C or even 7 °C; (Sherwood and Huber 2010)). For a threshold temperature of 5 °C, increasing uncertainty$_{\mathrm{ECS}}$ fivefold, from .5 to 2.5, increases the likelihood of catastrophe by a factor of nearly 250, from .06 to 14.5 %.

This is a necessary consequence of the fact that if two lognormal distributions have an identical expected value (μ_L), then the distribution with the greater standard deviation (call that σ_{L2}) will always have a "fatter tail" than the distribution with the lesser standard deviation (σ_{L1}). Intuitively, this phenomenon arises because the lower bound of all lognormal distributions is zero, and a greater standard deviation must therefore translate into greater probability mass *somewhere* in the upper tail. Formally, there exists a threshold θ, such that for any $x > \theta$, $\{1 - cdf_{\sigma_{L2}}(x)\} > \{1 - cdf_{\sigma_{L1}}(x)\}$, where cdf is the lognormal cumulative density function. The precise location of θ depends on both μ_L and σ_L; here, we are concerned only with the fact that there exists a θ beyond which the distribution with the greater σ_L has more probability mass, not the location of that threshold. We conclude that the greater the uncertainty, the greater the potential for large changes because greater uncertainty$_{\mathrm{ECS}}$ necessarily entails a greater likelihood of extremely high values of sensitivity. Rive and Myhre (2012) recognized the need to communicate the possibility of such extreme temperature changes to the public and policy makers.

Table 1 Parameter values and results for our simulations involving a lognormal climate sensitivity distribution

Simulation	Lognormal		Gaussian [a]		$P(X > T_c)$ [b]		$P(X < T_c)$ [c]	
	μ_L	σ_L	μ_G	σ_G	5 °C	7 °C	μ_L (3 °C)	2 °C (“safe”)
A	3.00	0.5	1.08	0.17	.0006	.0	.536	.009
B	3.00	1.0	1.05	0.32	.044	.003	.568	.044
C	3.00	1.5	0.99	0.47	.101	. 020	.595	.267
D	3.00	2.5	0.83	0.73	.145	.068	.635	.425

[a] Lognormal distributions are often characterized by the parameters of the underlying Gaussian distribution, shown here as μ_G and σ_G. They are related to the parameters of the lognormal distribution via $\mu_L = \exp(\mu_G + .5\sigma_G^2)$ and $\sigma_L = \exp(\mu_G + .5\sigma_G^2) \times \sqrt{\exp(\sigma_G^2) - 1}$.

[b] Proportions of sampled climate sensitivities that exceed the given threshold temperature T_c.

[c] Proportion of sampled climate sensitivities below the given threshold temperature T_c.

The second principal feature of Table 1 (columns $P(X < T_c)$) appears comforting at first glance: As uncertainty$_{\text{ECS}}$ increases, the proportion of lower-end sensitivities increases. This is a necessary consequence of keeping the mean, μ_L, constant while the mode and median shift to the left (cf. Fig. 2b). The proportion of low estimates (< 2 °C) increases to 42.5 % when uncertainty$_{\text{ECS}}$ is greatest.The large proportion of low sensitivities invites a potential alternative interpretation: With a nearly even chance that sensitivity might fall below the ostensibly-safe “guardrail” of 2 °C, perhaps one could legitimately ignore the upper tail, however it might fatten with increasing uncertainty? Further analysis reveals this gamble to be inadvisable for reasons involving the functional form of the relationship between warming and damage.

3.2 The convex damage function

There is pervasive agreement among economic models that further global warming will incur costly damage (e.g., Nordhaus 2010; Tol 2011). There is also agreement that the function that relates warming to cost is convex—that is, damage costs are accelerating with increasing warming. For example, one parameterization of the damage function favored by Nordhaus is:

$$d(t) = -0.0045 \times T(t) + 0.0035 \times T^2(t), \tag{1}$$

where $d(t)$ is damage cost, expressed as a percentage of world GDP, at time t as a function of temperature increase (T; in °C) since pre-industrial times (for a discussion of damage functions, see Wouter Botzen and van den Bergh 2012; Tomassini et al. 2010; Weitzman 2010b).

For our analysis of uncertainty$_{\text{ECS}}$, the precise form of the damage function turns out to be inconsequential: What matters instead are three pervasive attributes of damage functions: First, beyond a minimum at small temperature increases, damage functions are monotonically increasing with further warming. For example, the function in Eq. 1 reaches its minimum at $T = 0.64$ °C. Given that temperatures have already increased approximately 0.8 °C from pre-industrial times, further warming from here on—or from a nearby temperature threshold (Tol 2009)—will incur increasingly greater damages. Second, a pervasive attribute of damage functions is that they are *convex*—i.e., they are upward accelerating. Third, although a damage function cannot be computed without choosing a discount rate, the function retains its convex shape irrespective of the discount rate (e.g., Fig. 5a Tomassini et al. 2010). This convexity turns out to have notable implications.

With a convex function, we can rely on a theorem known as Jensen's inequality (Jensen 1906) to derive ordinal predictions about the damage expected from climate change. Jensen's inequality can be stated as: "If X is a (non-degenerate) random variable taking values in an interval (r, s), and if $u(X)$ is a strictly convex function on (r, s), then $mean[u(X)] > u(mean[X])$, providing that $mean[X]$ and $mean[u(X)]$ exist and are finite" (Brewster et al. 2005, p. 394). We explore the implications of Jensen's inequality, namely that increasing variance in X elevates the response $mean(u[X])$ if u is convex (Smallwood 1996), by simulation.

Figure 3 shows the relationship between simulated lognormal distributions of ECS and the total risk from climate change—that is, the distribution of likely damages. This illustrative simulation used a simple quadratic cost function, $d(T) = T^2$, where T was the value drawn from the climate-sensitivity distribution (bottom horizontal graphs in each panel) and d was the corresponding "reflected" observation in the damage-cost distribution (vertical graphs on left). The damage function is shown in the large plot within each panel.

We again fixed the mean of the sensitivity distributions for the reasons noted earlier, and changed only uncertainty$_{\text{ECS}}$. The figure reveals that increased uncertainty$_{\text{ECS}}$—reflected in the increased spread of the climate sensitivity distribution in panel B compared to panel A—translates into greater expected cost: The mean of the damage distribution is greater in panel B than in panel A (compare the length of vertical double-headed arrows between panels). This result is a necessary outcome provided the function relating temperature increases (i.e., sensitivity) to cost is convex (e.g., Smallwood 1996) and it holds—by Jensen's inequality—irrespective of the shape of the uncertainty$_{\text{ECS}}$ distribution. We show in the Supplementary material (Section S1) that this conclusion also holds when risk is measured in ways other than mean expected damage; for example, by examining the consequences for certain thresholds of sensitivity.

Some additional points apply to Fig. 3: First, not only does greater uncertainty$_{\text{ECS}}$ increase the mean expected damage, but it also increases the uncertainty around that expected damage (call that uncertainty$_{\text{D}}$). The increase in uncertainty$_{\text{D}}$ is relevant because it must be taken into account when assessing total risk. Second, we show in the supplementary material that this result holds irrespective of the absolute value of mean sensitivity. Although increasing mean sensitivity increases the expected mean damage, there always is an additional contribution from an increase in uncertainty$_{\text{ECS}}$ (see Table S1).

3.3 Speed of warming and the discount rate

Finally, we examine the relationship between climate sensitivity and the rate of future warming. There is considerable evidence that greater sensitivity translates into more rapid warming, all other things being equal (e.g., Bahn et al. 2011; Raupach et al. 2011; Ross et al. 2012; Winton et al. 2013). It follows that greater uncertainty about climate sensitivity translates not only into greater expected damage, but it also implies that greater damages are likely to arrive sooner rather than later.[1] We provide a formal analysis of the

[1] Notwithstanding the greater speed of warming, the climate system will take longer overall to reach equilibrium if climate sensitivity is greater (Hansen et al. 1985).

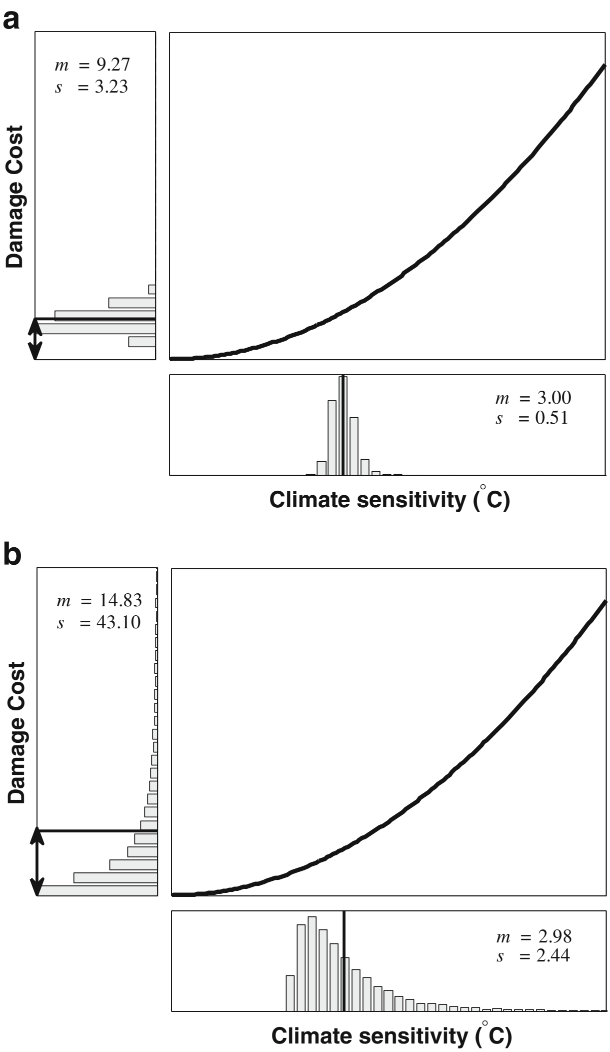

Fig. 3 Illustration of the consequences of Jensen's inequality across two levels of uncertainty$_{\text{ECS}}$. Standard deviations of the lognormal climate sensitivity distributions (σ_L), shown at the bottom of each panel, are .5 and 2.5 in panels A and B, respectively, with μ_L fixed at 3. Distributions of expected damage costs are shown at the left in each panel, using units that are arbitrary but with a scale that is aligned between panels. Mean (m) and standard deviation (s) of the simulated distributions are also shown. For all distributions the means are represented by *thick solid lines*. *Vertical arrows* highlight distance of the mean of the damage distribution from 0. The damage function in the large plot within each panel is arbitrary and shown for illustration only

relationship between speed of warming and climate sensitivity in the Supplementary material (Section S2).

This point has an important implication because it prevents the discount rate from altering the functional relationship between uncertainty$_{\text{ECS}}$ and the magnitude of damages. To clarify why that is the case, it is helpful to consider the consequences of the opposite outcome, viz. if greater damages were delayed, rather than accelerated, as a function of greater uncertainty$_{ECS}$. If that opposite outcome were the case, then an expected damage cost of, say, \$100,000,000 under greater uncertainty$_{ECS}$ would occur some time later than a cost of, say, \$80,000,000 if uncertainty$_{ECS}$ were lower. When a greater cost is delayed relative to a lesser one, its value can always be discounted below the lesser cost by choosing a convenient discount rate. We noted earlier that the choice of discount rate can alter discounted amounts by orders of magnitude: For our analysis it is therefore important that greater uncertainty$_{ECS}$ translates into acceleration as well as amplification of climate-related damages. Because discounting, by definition, only affects future outcomes, any cost that occurs *sooner* than another one cannot be discounted below the other one: \$100,000,000 due now cannot be discounted to be less than \$80,000,000 that is due at some later time. This argument is again ordinal in nature and does not depend on how *much* quicker warming is if climate sensitivity is greater than anticipated: Provided greater sensitivity does not delay a given extent of warming—and it does not (Bahn et al. 2011; Raupach et al. 2011; Ross et al. 2012; Winton et al. 2013)—the relationship between uncertainty$_{ECS}$ and damages just uncovered holds irrespective of the discount rate.

Our analysis thus far rested on three assumptions. First, we assumed an asymmetric "fat-tailed" distribution of ECS. Second, by focusing on a mean-invariant transformation of a presumed probability density function, our analysis has been limited to what is called "first-order" uncertainty; namely, uncertainty associated with the estimate of a single parameter. Third, we assumed that the function relating warming to damages is convex, as is pervasively assumed in economic modeling. We now show that our analysis remains unaffected by a relaxation of those assumptions.

3.4 Symmetrical climate sensitivity

We show in the Supplementary material (Section S3) that the implications of Jensen's inequality generalize to situations in which the distribution of climate sensitivities is entirely symmetrical. Although the assumption of symmetry is likely incorrect (Roe and Baker 2007, 2011), the removal of a fat upper tails is also the most conservative assumption one can make about the evolution of the climate. The fact that increasing uncertainty nonetheless leads to increasing expected damages attests to the robustness of our approach.

3.5 Ambiguity among estimates of climate sensitivity

We argued earlier that enough is known about climate sensitivity to consider its estimation to involve mainly "first-order" uncertainty, or uncertainty intrinsic to the estimate of a single value. What are the implications of introducing "second-order" uncertainty, for example by considering competing and divergent estimates of ECS?

Figure 4 (Panel A) illustrates second-order uncertainty with 3 hypothetical climate-sensitivity distributions that differ in mean and variance. Each distribution can be taken to represent the knowledge of a single expert obtained by an expert-elicitation methodology (e.g., Morgan and Keith 1995). Although the hypothetical distributions share a common (lognormal) shape and overlap considerably, there is also some heterogeneity among the expert judgments. This is best illustrated by considering the areas of the curves that exceed a cutoff (T_c), which in Fig. 4 was arbitrarily set to 4 °C. Note that the same cutoff applies to all three distributions, and that the area exceeding the cutoff is a function of both mean and variance of each distribution.

Given that we assume a doubling of atmospheric CO_2 throughout this article, the shaded areas also represent the probabilities π_i that each expert i assigns to the risk that global

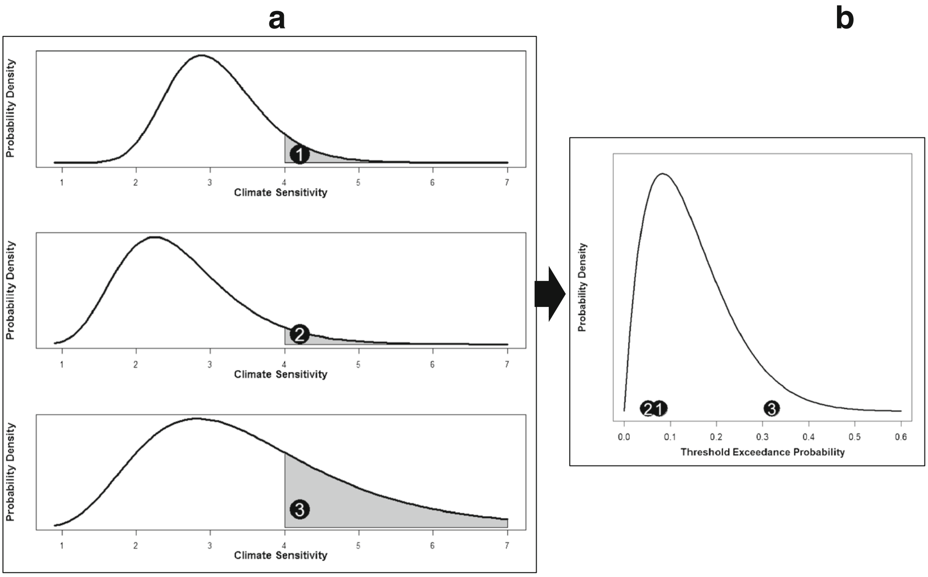

Fig. 4 **a** Three hypothetical lognormal probability density functions (PDF) of climate sensitivity with different means and variances ($\mu_G = 1.1, 0.9$, and 1.2, respectively, from top to bottom; corresponding $\sigma_G = 0.2, 0.3$, and 0.4). Each distribution can be taken to represent the subjective PDF of an individual expert. The shaded areas, numbered 1 through 3, identify the probabilities with which a temperature threshold (T_c; in this case 4 °C) would be exceeded if CO_2 levels doubled from preindustrial levels. **b** Hypothetical Beta distribution (parameters $a = 2$; $b = 12$) that characterizes the exceedance probabilities (*shaded areas* in Panel A; π_i's) across individuals. Numbered points refer to shaded areas in Panel A. See text for more details

temperatures may increase by more than T_c (4 °C). Second-order uncertainty, then, is captured by the distribution of those exceedance probabilities, π_i's, across experts. Panel B in Fig. 4 shows a hypothetical Beta distribution that characterizes those exceedance probabilities across a presumed population of experts. (Panel B also shows the three exceedance probabilities from Panel A for illustrative purposes.)

It is intuitively obvious that as disagreement among experts increases—thereby introducing greater heterogeneity into their estimated exceedance probabilities—the variance of the Beta distribution that characterizes those probabilities also increases. This in turn necessarily implies that there is greater probability mass over the larger values of π_i, thereby raising the likelihood of greater risks of exceeding T_c (4 °C).

This intuition is formalized in Fig. 5, which plots a family of complementary cumulative distribution functions (i.e., the tail distribution $\bar{F}(x) = \mathrm{P}(X > x) = 1 - F(x)$) for Beta distributions of constant mean ($\mu_B = 0.2$) but differing variance. The weight of the lines in Fig. 5 maps into the variance of the underlying distribution: It is clear that as the variance increases, the probability mass over the larger exceedance probabilities increases. Conceptually, this means that as disagreement among experts increases, greater risks of exceeding a global temperature threshold also become more likely.

Note that this result is based entirely on the distribution of exceedance probabilities: It is therefore unimportant whether experts differ with respect to their estimates of mean climate sensitivity, or their estimates of uncertainty surrounding that mean estimate, or both. The only constraint in our analysis of second-order uncertainty is, once again, that the mean

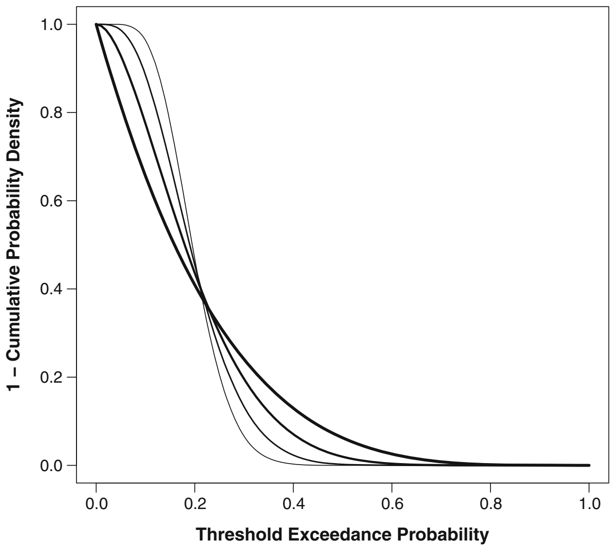

Fig. 5 Family of complementary cumulative distribution functions for Beta distributions of constant mean ($\mu_B = 0.2$) but differing variance. The standard deviations of the functions range from $\sigma_B = 0.062$ for the *lightest line*, to 0.087, 0.121, and 0.163 in increasing order of *line thickness*

exceedance probability (i.e., μ_B) remains constant—it is only by keeping the mean constant that the effects of varying degrees of disagreement among experts can be modeled.

In the Supplementary material (Section S4), we apply the preceding analysis to an expert-elicitation study (Morgan and Keith 1995). The data of Morgan and Keith (1995) are particularly suitable because they contain one outlying observation; namely, an expert who provided a very low but highly certain estimate of climate sensitivity. We show that removal of this outlier, which reduces ambiguity among experts, tends to reduce the likelihood of extreme risks of exceeding a global temperature threshold. Conversely, we show that if this outlying expert is replicated and replaces some other observations at random, the likelihood of exceeding a tolerable risk of excessive warming increases in most cases. For that analysis, both mean and variance of the distribution of exceedance probabilities were allowed to vary freely: We nonetheless found that greater second-order uncertainty entailed a greater likelihood of threshold exceedance in the majority of cases.

3.6 Sea level rise and inundation risk

Sea level rise (SLR) is one of the more costly consequences of climate change: Coping with SLR requires that existing coastal structures must be relocated, raised, or protected by levees or dikes. To prevent an increase in the risk from flooding as sea levels rise, it is insufficient to focus on projected mean SLR: Instead, an additional allowance for extreme events (e.g., storm surges) must be made.

The extra allowance for extreme events is a function of the uncertainty in the estimated mean SLR, called uncertainty$_{\mathrm{SLR}}$ from here on. That is, the total height of protective measures required to keep the risk from flooding constant is a function of the projected increase (call that X) as well as its uncertainty$_{\mathrm{SLR}}$ (i.e., the standard deviation of X). Hunter (2012) computed the required magnitude of the constant-risk protective response to SLR. Risk here is measured in units of flooding events whose magnitude is assumed to remain constant. In line with recent recommendations (Katz et al. 2013), Hunter (2012) focused on modeling of extreme events. Specifically, Hunter (2012) used a Gumbel distribution, the simplest of the set of generalised extreme-value distributions, which is known to characterize sea-level extremes (Van den Brink and Können 2011). The implications of Hunter (2012)'s analysis are shown in Fig. 6.

The figure plots the total protective response necessary to cope with SLR as a function of uncertainty$_{SLR}$, using three different distributions to characterize uncertainty$_{SLR}$. Note that unlike the distribution of climate sensitivity, the three modeled distributions for uncertainty$_{SLR}$—uniform, Gaussian, and raised cosine—are all symmetrical. For the purposes of this illustration, SLR was expected to be 0.5 m, represented by the horizontal dashed line. Accordingly, when there is no uncertainty$_{SLR}$ in the estimate, the total protective response is equal to the expected SLR, namely 0.5 m.

When uncertainty$_{SLR}$ is non-zero, then irrespective of what assumptions are made about the distribution of SLR, the required protective response increases and deviates rapidly and in an accelerating manner from the anticipated mean SLR. For example, under a Gaussian assumption, if uncertainty$_{SLR}$ is around 0.36 m, this raises the required protective response to around 1 m. That is, an expected SLR of 0.5 m requires that dikes and levees be raised by *twice* that amount in order to keep the risk of flooding constant in light of uncertainty$_{SLR}$. If other distributional assumptions are made, the values change but the in-principle conclusion remains the same: Greater uncertainty$_{SLR}$ translates into a greater required protective response. Note that the value of 0.36 m was not chosen arbitrarily but represents a current estimate of the uncertainty$_{SLR}$ in SLR projections (Hunter 2012; Nicholls et al. 2011). Likewise, the expectation of 0.5 m global mean SLR is consonant with IPCC projections for

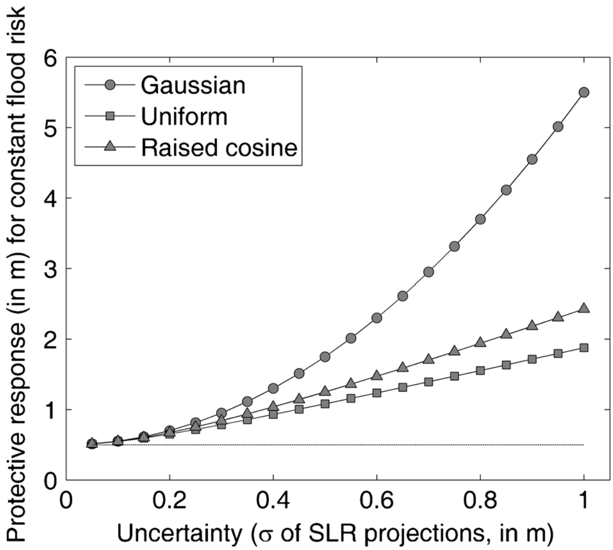

Fig. 6 The effects of uncertainty$_{SLR}$ in future sea level rise (SLR) on the protective response required to keep the risk from flooding constant. Protective response is expressed in *m* (e.g., of raising dikes, levees, or buildings) and uncertainty$_{SLR}$ is expressed as the standard deviation (σ) of the expected SLR. For this illustration, sea level is expected to rise by 0.5 m (*dotted horizontal line*). The three lines represent different distributional assumptions about future SLR. Figure produced from the equations provided by Hunter (2012)

century's end on the current emissions trajectory (Meehl et al. 2007), although that estimate is now considered a lower bound on likely SLR (e.g., Nicholls et al. 2011).

We underscore two points in connection with Fig. 6: First, the increase in the required protective response results from greater uncertainty$_{SLR}$ *only*—the mean expected SLR is constant for all data points in the figure. Second, the effects of uncertainty$_{SLR}$ derive from the mathematical properties of extreme values (for details, see Hunter 2012), without any auxiliary assumptions such as the convexity of the economic damage function, and are therefore in little doubt at least at an ordinal level.

4 Conclusions: uncertainty and unabated emissions

4.1 Potential objections and limitations

Unlike related precedents (e.g., Tomassini et al. 2010; Webster et al. 2003; Weitzman 2009, 2011), our analysis ignored sources of uncertainty other than those associated with ECS or SLR. Several other sources of uncertainty exist, such as potential amplifying loops in the carbon cycle itself: Because warming can accelerate respiration, soil may turn from a carbon sink to a carbon source (Cox et al. 2004), thereby shortening the time until CO_2 has doubled. The considerable uncertainty about the global policy response likewise affects the time of CO_2 doubling. One may therefore question whether our analysis extends to those other sources of uncertainty.

Webster et al. (2003) differentiated between climate and emissions uncertainty, suggesting that either on its own is just over half the uncertainty of both combined. Specifically, the standard deviation of expected temperature increases by 2100 was 0.69 °C for climate uncertainty alone, 0.76 °C for emissions uncertainty, and 1.18 °C for both combined. Emissions uncertainty thus adds to the uncertainty$_{ECS}$ we have considered here, suggesting that this additional source of uncertainty amplifies the impact of our analysis.

4.2 Implications

Our analyses of the role of uncertainty$_{ECS}$ in relation to damage costs from climate change yielded a fairly clear conclusion: The greater the uncertainty, the greater the expected cost of unmitigated global warming. When damages are considered at a macro-economic level, this relationship arises directly from the convexity of the damage function, an assumption shared by all extant economic models. When the consequences of climate change are considered with respect to sea level rise, the relationship between greater uncertainty$_{SLR}$ and greater adaptation costs emerges without making any assumptions about the cost function, simply from mathematical examination of the behavior of extreme values.

Our conclusion is unlikely to be affected by the discount rate because increasing uncertainty$_{ECS}$ likely also accelerates damages in addition to increasing their magnitude, thereby preventing any application of discounting. Similarly, because our analysis presents ordinal constraints (i.e., of the form "greater than") that do not depend on absolute estimates of uncertainty$_{ECS}$ our conclusions also are not subject to cultural and personal risk perception variables. Our conclusion is also quite robust to "second-order" uncertainty, such as disagreement among experts about the climate-sensitivity distribution.

Our analysis permits at least one optimistic conclusion: The converse of our argument is that any reduction in uncertainty$_{ECS}$ arising from further research will translate into lesser expected damages from unmitigated climate change (cf. Webster et al. 2003).

At the outset, we identified two misconstruals of climatic uncertainty; viz. that surprises will likely be favorable and that uncertainty prevents policy decisions from being scientifically informed. This article has shown the first claim to be flawed: Any appeal to uncertainty about the evolution of the climate implies a stronger, rather than weaker, reason to be concerned about unmitigated climate change than in the absence of uncertainty. To complete the picture, we must confront the second misconstrual by analyzing whether uncertainty can prevent science from informing policy. This is addressed by the companion article (Lewandowsky et al. 2014).

Acknowledgments Preparation of this paper was facilitated by a Discovery Grant and a Discovery Outstanding Researcher Award from the Australian Research Council, and a Wolfson Research Merit Award from the Royal Society, to the first author, by a Future Fellowship from the Australian Research Council to Ben Newell, and funding from the Australian Research Council Centre of Excellence in Climate Systems Science. The work was also supported by a Linkage Grant from the Australian Research Council, by the National Climate Change Adaptation Research Facility, and the CSIRO Climate Adaptation Flagship. We thank four reviewers for their incisive critique and helpful comments. Correspondence to the first author at the School of Experimental Psychology, University of Bristol, 12A Priory Road, Bristol BS8 1TU, United Kingdom (stephan.lewandowsky@bristol.ac.uk). Personal web page: http://www.cogsciwa.com.

References

Ackerman F, Heinzerling L, Massey R (2005) Applying cost-benefit to past decisions: was environmental protection ever a good idea? Adm Law Rev 57:155–192

Aldred J (2009) The skeptical economist: revealing the ethics inside economics. Earthscan

Annan JD, Hargreaves JC (2011) On the generation and interpretation of probabilistic estimates of climate sensitivity. Clim Chang 104:423–436

Anthoff D, Tol RSJ, Yohe GW (2009) Discounting for climate change. Economics: the open access open assessment E-Journal 3

Bahn O, Edwards NR, Knutti R, Stocker TF (2011) Energy policies avoiding a tipping point in the climate system. Energy Policy 39:334–348. doi:10.1016/j.enpol.2010.10.002

Bender FA-M, Ekman AM, Rodhe H (2010) Response to the eruption of mount pinatubo in relation to climate sensitivity in the CMIP3 models. Clim Dyn 35:875–886

Brewster JF, Graham MR, Mutch AC (2005) Convexity, Jensen's inequality and benefits of noisy mechanical ventilation. J R Soc Interface 2:393–396

Cox PM, Betts RA, Collins M, Harris PP, Huntingford C, Jones CD (2004) Amazonian forest dieback under climate-carbon cycle projections for the 21st century. Theor Appl Climatol 78:137–156

Feintuck M (2005) Precautionary maybe, but what's the principle? The precautionary principle, the regulation of risk, and the public domain. J Law Soc 32:371–398. doi:10.1111/j.1467-6478.2005.00329.x

Freudenburg WR, Gramling R, Davidson DJ (2008) Scientific certainty argumentation methods (SCAMs): science and the politics of doubt. Sociol Inq 78:2–38

Freudenburg WR, Muselli V (2013) Reexamining climate change debates: scientific disagreement or scientific certainty argumentation methods (SCAMs)? Am Behav Sci 57:777–795. doi:10.1177/0002764212458274

Funtowicz SO, Ravetz JR (1994) The worth of a songbird: ecological economics as a post-normal science. Ecol Econ 10:197–207

Gardiner S (2006) A core precautionary principle. J Political Philos 14:33–60

Garnaut R (2011) The Garnaut review 2011: Australia in the global response to climate change. Cambridge University Press, Cambridge

Hansen J, Russell G, Lacis A, Fung I, Rind D, Stone P (1985) Climate response times: dependence on climate sensitivity and ocean mixing. Science 229:857–859

Hartman R (1972) The effects of price and cost uncertainty on investment. J Econ Theory 5:258–266

Hegerl GC, Crowley TJ, Hyde WT, Frame DJ (2006) Climate sensitivity constrained by temperature reconstructions over the past seven centuries. Nature 440:1029–1032

Holmes J, Lowe J, Wolff E, Srokosz M (2011) Rapid climate change: lessons from the recent geological past. Glob Planet Chang 79:157–162

Hunter J. (2012) A simple technique for estimating an allowance for uncertain sea-level rise. Clim Chang 113:239–252. doi:10.1007/s10584-011-0332-1

Intergovernmental Panel on Climate Change (2005) Guidance notes for lead authors of the IPCC Fourth Assessment Report on addressing uncertainties. (Technical Report)

Jensen JLWV (1906) Sur les fonctions convexes et les inégalités entre les valeurs moyennes. Acta Mathematica 30:175–193

John S (2010) In defence of bad science and irrational policies: an alternative account of the precautionary principle. Ethical Theory Moral Pract 13:3–18. doi:10.1007/s10677-009-9169-3

Kahan DM, Slovic P, Braman D, Gastil J (2006) Laws of fear: beyond the precautionary principle. Harv Law Rev 119:1071–1109

Katz RW, Craigmile PF, Guttorp P, Haran M, Sansto B, Stein ML (2013) Uncertainty analysis in climate change assessments. Nat Clim Change 3:769–771

Knutti R, Hegerl GC (2008) The equilibrium sensitivity of the Earth's temperature to radiation changes. Nat Geosci 1:735–743

Lewandowsky S, Oreskes N, Newell BR, Smithson M, Risbey J (2013) Uncertainty as a barrier to mitigation: historical and psychological factors. (Manuscript in preparation)

Lewandowsky S, Risbey JS, Smithson M, Newell BR (2014) Scientific uncertainty and climate change: part II. Uncertainty and mitigation. Clim Chang. doi:10.1007/s10584-014-1083-6

Li M, Vietri J, Galvani AP, Chapman GB (2010) How do people value life? Psychol Sci 21:163–167. doi:10.1177/0956797609357707

Meehl GA, Stocker TF, Collins WD, Friedlingstein P, Gaye AT, Gregory JM, Zhao Z-C (2007) Global climate projections. I. In: S. Solomon, et al. (eds) Climate change 2007: the physical science basis. Contribution of Working Group I to the Fourth assessment report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, pp 748–845

Michaels D (2008) Doubt is their product: how industry's assault on science threatens your health. Oxford University Press, New York

Morgan MG, Keith DW (1995) Subjective judgments by climate experts. Environ Sci Tech 29:468A–476A

Narita D (2012) Managing uncertainties: the making of the IPCCs special report on carbon dioxide capture and storage. Public Underst Sci 21:84–100

Nicholls RJ, Marinova N, Lowe JA, Brown S, Vellinga P, De Gusmao D, Tol RSJ (2011) Sea-level rise and its possible impacts given a 'beyond 4 °C' world in the twenty-first century. Philos Trans R Soc A 369:161–181

Nolt J (2011) How harmful are the average American's greenhouse gas emissions? Ethics. Policy Environ 14:3–10

Nordhaus WD (2010) Economic aspects of global warming in a post-Copenhagen environment. Proc Natl Acad Sci 107:11721–11726

Nordhaus WD (2011) The economics of tail events with an application to climate change. Rev Environ Econ Policy 5:240–257

Padilla LE, Vallis GK, Rowley CW (2011) Probabilistic estimates of transient climate sensitivity subject to uncertainty in forcing and natural variability. J Clim 24:5521–5537. doi:10.1175/2011JCL13989.1

Peterson M (2006) The precautionary principle is incoherent. Risk Anal 26:595–601. doi:10.1111/j.1539-6924.2006.00781.x

Raupach MR, Canadell JG, Ciais P, Friedlingstein P, Rayner PJ, Trudinger CM (2011) The relationship between peak warming and cumulative CO2 emissions, and its use to quantify vulnerabilities in the carbon-climate-human system. Tellus 63B:145–164

Ressurreição A, Gibbons J, Dentinho TP, Kaiser M, Santos RS, Edwards-Jones G (2011) Economic valuation of species loss in the open sea. Ecol Econ 70:729–739

Risbey JS (2006) Some dangers of dangerous climate change. Clim Pol 6:527–536

Rive N, Myhre G (2012) Communicating the probabilities of extreme surface temperature outcomes. Atmos Clim Sci 2:538–545

Roe GH, Armour KC (2011) How sensitive is climate sensitivity? Geophys Res Lett 38:L14708. doi:10.1029/2011GL047913

Roe GH, Baker MB (2007) Why is climate sensitivity so unpredictable? Science 318:629–632

Roe GH, Baker MB (2011) Comment on Another look at climate sensitivity by Zaliapin and Ghil (2010). Nonlinear Process Geophys 18:125–127

Ross A, Matthews HD, Schmittner A, Kothavala Z (2012) Assessing the effects of ocean diffusivity and climate sensitivity on the rate of global climate change. Tellus B 64:17733

Sarewitz D (2004) How science makes environmental controversies worse. Environ Sci Policy 7:385–403

Schneider SH (2002) Can we estimate the likelihood of climatic changes at 2100? Clim Chang 52:441–451

Sherwood SC, Huber M (2010) An adaptability limit to climate change due to heat stress. Proc Natl Acad Sci 107:9552–9555
Slovic P (1999) Trust, emotion, sex, politics, and science: surveying the risk assessment battlefield. Risk Anal 19:689–701
Smallwood PD (1996) An introduction to risk sensitivity: the use of Jensen's inequality to clarify evolutionary arguments of adaptation and constraint. Am Zool 36:392–401
Smithson M (1999) Conflict aversion: preference for ambiguity vs conflict in sources and evidence. Organ Behav Hum Decis Processs 79:179–198
Stern N (2007) The economics of climate change: the Stern review. Cambridge University Press, Cambridge
Sterner T, Persson UM (2008) An even Sterner review: introducing relative prices into the discounting debate. Rev Environ Econ Policy 2:61–76
Tol RSJ (2009) The economic effects of climate change. J Econ Perspect 23:29–51
Tol RSJ (2011) The social cost of carbon. Ann Rev Res Econ 3:419–443
Tomassini L, Knutti R, Plattner G-K, van Vuuren DP, Stocker TF, Howarth RB, Borsuk ME (2010) Uncertainty and risk in climate projections for the 21st century: comparing mitigation to non-intervention scenarios. Clim Chang 103:399–422. doi:10.1007/s10584.009.9763.3
Van der Sluijs JP (2005) Uncertainty as a monster in the science policy interface: four coping strategies. Water Sci Tech 52(6):87–92
Van den Brink HW, Können GP (2011) Estimating 10000-year return values from short time series. Intl J Clim 31:115–126
Vlek C (2010a) Judicious management of uncertain risks: I. Developments and criticisms of risk analysis and precautionary reasoning. J Risk Res 13:517–543
Vlek C (2010b) Judicious management of uncertain risks: II. Simple rules and more intricate models for precautionary decision-making. J Risk Res 13:545–569. doi:10.1080/13669871003629903
Walley P (1991) Statistical reasoning with imprecise probabilities. Chapman and Hall, London
Webster M, Forest C, Reilly J, Babiker M, Kicklighter D, Mayer M, Wang C (2003) Uncertainty analysis of climate change and policy response. Clim Chang 61:295–320. doi:10.1023/B:CLIM.0000004564.09961.9f
Weitzman ML (2009) On modeling and interpreting the economics of catastrophic climate change. Rev Econ Stat 91:1–19
Weitzman ML (2010a) How should the distant future be discounted when discount rates are uncertain? Econ Lett 107:350–353
Weitzman ML (2010b) What is the damages function for global warming and what difference might it make? Clim Chang Econ 1:57–69
Weitzman ML (2011) Fat-tailed uncertainty in the economics of catastrophic climate change. Rev Environ Econ Policy 5:275–292
Winton M, Adcroft A, Griffes SM, Hallberg RW, Horowitz LW, Stouffer RJ (2013) Influence of ocean and atmosphere components on simulated climate sensitivities. J Clim 26:231–245
Wouter Botzen W, van den Bergh JC (2012) How sensitive is Nordhaus to Weitzman? climate policy in DICE with an alternative damage function. Econ Lett 117:372–374
Zeebe RE, Zachos JC, Dickens GR (2009) Carbon dioxide forcing alone insuffcient to explain Palaeocene Eocene Thermal Maximum warming. Nat Geosci 2:576–580