

Climate Sensitivity Uncertainty: When is Good News Bad?

Mark C. Freeman¹, Gernot Wagner^{2*}, and Richard J. Zeckhauser³

¹ Professor of finance, School of Business and Economics, Loughborough University, Loughborough, Leicestershire LE11 3TU, United Kingdom, 0000-0003-4521-2720. ² Lead senior economist, Environmental Defense Fund, 18 Tremont St. Ste 850, Boston MA 02108, USA, adjunct associate professor, School of International and Public Affairs, Columbia University, 420 W. 188th St #1, New York, NY 10027, USA, and research associate, Harvard Kennedy School, 79 John F. Kennedy St., Cambridge, MA 02138, USA, 0000-0001-6059-0688. ³ Frank P. Ramsey Professor of Political Economy, Harvard Kennedy School, 79 John F. Kennedy St., Cambridge, MA 02138, USA, 0000-0002-4222-9546.

Keywords: Climate sensitivity, climate change, uncertainty, mean preserving spread, risk-return trade-off, willingness to pay

Climate change is real and dangerous. Exactly how bad it will get, however, is uncertain. Uncertainty is particularly relevant for estimates of one of the key parameters: equilibrium climate sensitivity—how eventual temperatures will react as atmospheric carbon dioxide concentrations double. Despite significant advances in climate science and increased confidence in the accuracy of the range itself, the “likely” range has been 1.5–4.5°C for over three decades. In 2007, the Intergovernmental Panel on Climate Change (IPCC) narrowed it to 2–4.5°C, only to reverse its decision in 2013, reinstating the prior range. In addition, the 2013 IPCC report removed prior mention of 3°C as the “best estimate.” We interpret the implications of the 2013 IPCC decision to lower the bottom of the range and excise a best estimate. Intuitively, it might seem that a lower bottom would be good news. Here we ask: When might apparently good news about climate sensitivity in fact be bad news in the sense that it lowers societal wellbeing? The lowered bottom value also implies higher uncertainty about the temperature increase, a definite bad. Under reasonable assumptions, both the lowering of the lower bound and the removal of the “best estimate” may well be bad news.

1. Introduction

What if a study utilizing a new insight on equilibrium climate sensitivity—how temperatures react in the long run as carbon dioxide concentrations double—produced results significantly below what most climate models and scientists now assume? The natural response would be to celebrate, and to conclude that the value of stringent climate policy had diminished.

That celebration would be premature. Countervailing the good news would be the recognition that current climate science, while sound on the big picture, did not understand a fundamental climate uncertainty as well as had been thought, assuming that the new information adds to prior knowledge rather than merely correcting it. If there is one certainty, it is that the new data will not reveal all. Temperature increases predicted by equilibrium climate sensitivity, after all, play out over decades and centuries. Thus, the substantial change in estimates would point to even more and deeper uncertainties than previously recognised.

To put the same matter in statistical terms, our mean estimate on climate sensitivity might have shifted down, but our estimate of its standard deviation would have increased. Deeper uncertainty should magnify concerns, since marginal damages from rising temperatures increase rapidly. Quite possibly, the magnified concerns from variability would outweigh any reassurance from a lowered estimate on the mean (1,2,3,4,5,6). This focus is distinctly different from uncertainty due to high-impact, low-probability events, often referred to as ‘tail risks’ or ‘fat tails’ (2,3). Tail risks magnify concerns significantly; however, they are not the focus of our analysis.

We focus here on climate sensitivity for the simple reason that it’s the most iconic of climate metrics. It is also among the best-studied. Yet, science has not been able to narrow its range in over three decades: double carbon dioxide concentrations and, consensus climate science tells us, expect long-run temperatures to rise by between 1.5 and 4.5°C. That range has stood ever since Jule Charney chaired a National Academy of Sciences Ad Hoc Study Group on Carbon

* Author for correspondence (gernot@gwagner.com).

Dioxide and Climate in the late 1970s (7). In 1990, the Intergovernmental Panel on Climate Change (IPCC) picked up Charney's range for equilibrium climate sensitivity. That verdict held for a further fifteen-plus years of increasingly intense scrutiny and increased recognition of its validity until 2007, when the IPCC decided to cut off the bottom of the range, which became 2°C (8). Apparent bad news: the lowest estimates for climate sensitivity seemed to be ever more out of reach.

In 2013, the IPCC took two steps with respect to the climate sensitivity parameter (9). First, it widened the range once again, keeping the top value of 4.5°C but moving the lower bound of the "likely" range back down to 1.5°C. That step points to deep-seated uncertainties inherent in climate science and, thus, policy. Thirty-five years of amazing advances in most every aspect of climate science apparently have not tightened the range for the eventual realization of this fundamental parameter. Indeed, they have taught us to be more cautious in defining it. See (3) for more on the history and the IPCC's detailed definitions and (3, 10, 11) for discussions on the deep-seated nature of climate sensitivity.

Second, the IPCC (2013) also removed its statement of a "best estimate" for climate sensitivity. Prior mention of that number has always been rather vague, with "best estimate" invariably being interpreted as "mean" or "median" climate sensitivity. Now the IPCC has removed mention of it entirely. No interpretation was given to this move in the report, and we don't venture to second-guess the IPCC's decision. We simply attempt to interpret the implication of removing the value itself by suggesting that any assumed climate sensitivity distribution may now have reduced 'peakedness' as a result. We interpret peakedness using kurtosis, holding the mean and standard deviation fixed. Kurtosis is an indicator of how sharp (high kurtosis) or rounded (low kurtosis) is the peak of the distribution.

As an alternate approach, we also look at the distribution changing both its kurtosis and standard deviation to hold constant the IPCC's "likely" interval. Many others have attempted to calibrate distribution functions around the IPCC's (2007) climate sensitivity pronouncements (1,2,3,12,13). The 2013 step of removing the "best estimate" might indicate that the chance of hitting close to the peak of any prior probability distribution has decreased. Once again, this additional level of uncertainty is apparent bad news: following Pindyck's (1,12) calibration, decreased peakedness, other factors equal, implies greater Willingness to Pay (WTP) out of current consumption to avoid climate damages in the future. WTP is the maximum amount in dollars someone is willing to pay to avoid something undesirable. More specifically, we define WTP here following Pindyck (1,12), in the sense of how much it is optimal for society to pay to avoid certain degrees of average global warming in order to maximise total societal well-being.

In our subsequent analysis, we ask two sets of questions:

1. When is good news bad? Specifically, under what conditions does a lowering of the lower bound of the "likely" climate sensitivity range lead to increased WTP to avoid global warming?
2. What should we make of the IPCC's removal of its prior "best estimate"? How do we interpret knowing about the climate sensitivity range but not where within that range we might end up? And how does this affect our WTP to avoid global warming?

We focus on the effect of increasing the uncertainty of climate sensitivity on WTP, while leaving the mean unchanged or lowering it, in section 2. Section 3 interprets the second set of questions around 'peakedness', or a lack thereof. Section 4 concludes. An extensive set of appendices presents proofs for our results, and looks for special cases with distinctive properties.

2. When is good news bad? The mean-variance trade-off

A decrease in the mean climate sensitivity, other things equal, is undoubtedly good news for the planet. We could expect eventual global average temperatures to rise less than previously feared. However, when that decrease in mean is due to a widening of the uncertainty range—for example, if it is due to a lowering of the lower bound—this news may not overall be good. In fact, that is what we find may be the case here.

We first develop the general economic framework that underlies our analysis in subsection 2.1. We represent an increase in uncertainty as a Mean Preserving Spread (MPS). Probability distribution B is an MPS of probability distribution A, if it spreads out some portions of A's probability density function, but has the same mean as A. For example, posit that A has temperature increases of 1, 2 and 3 degrees with likelihoods 20%, 60% and 20%, respectively. B has temperature

increases of 1, 1.5, 2, 2.5 and 3 degrees with likelihoods 20%, 20%, 20%, 20% and 20%, respectively. B has spread out A's distribution around 2 degrees, but has the same mean as A. It is thus an MPS of A.

In subsection 2.2, we consider the impact of an MPS on the distribution of future temperatures. We demonstrate that the description of climate sensitivity in Assessment Report 5 (9) could be seen as an MPS of the description in Assessment Report 4 (8). Thus, despite the fact that the IPCC's 2013 position on climate sensitivity might superficially be viewed as 'good' news, our analysis shows that a reasonable interpretation of the results would conclude otherwise: WTP goes up.

In subsection 2.3, we turn to increases in the standard deviation of the distribution of climate sensitivity. Every MPS falls into this category when the increase in variance is mean-preserving, but the converse is not true. Therefore, subsection 2.3 contains a strictly broader category of increased risk than subsection 2.2. We demonstrate that WTP can increase, even when expected climate sensitivity decreases, under a broad range of conditions.

In subsection 2.4, we turn our attention to the situation where there is also uncertainty over future economic conditions and consider the case when consumption growth is correlated with temperature changes.

2.1 The framework

Let y_0 denote current per-capita consumption levels and y_t^* be what Weitzman describes as "*potential* consumption in the complete absence of climate change because it is defined to be what consumption *would be* without any global warming" at some future time t (his emphasis) (13, pp. 58–59). Then realised consumption, y_t at time t in the presence of climate change damage, is given by $y_t = (1 - D(T))y_t^*$, where $D(T) \in [0,1)$, and $D'(T) > 0$ is a multiplicative climate change damage function. We use T , which represents the difference between global average temperatures at time t and pre-industrial levels, as our proxy for climate change. It might, though, just as easily be interpreted as the rise in sea level, extreme weather events, or any other relevant climatic metric. When interpreting our results in terms of climate sensitivity, T is equal to climate sensitivity due to a doubling of the level of carbon dioxide in the atmosphere.

We consider the preferences of a rational social planner in a general economic setting. $U(y, t)$ represents the planner's time-separable utility function with per-capita consumption and time as its elements. In other words, the utility derived in one period does not depend on what happens in other periods. In essence, the discounted levels of wellbeing in different periods are added together. Though we talk about a social planner acting on behalf of all individuals, the formulation could be equally well conducted in per-capita terms for a representative individual within society. This analysis excludes issues associated with the limitations of using expected utility theory in the presence of ambiguity (rather than risk) (14). The amount of consumption, p , that the planner would be prepared to sacrifice today in order to prevent all future climate change damage at time t is:

$$(1) \quad U(y_0, 0) - U(y_0 - p, 0) = E[U(y_t^*, t)] - E[U(y_t, t)].$$

The left-hand side of equation (1) represents the immediate gain in utility from not spending p on mitigation today, while the right-hand side is the expected gain in future utility from mitigating climate change and, thus, consuming y_t^* rather than y_t . For given y_t^* and t , we will define the value function to be $V(T) = U((1 - D(T))y_t^*, t)$.

Assume that the planner's utility function takes the standard constant relative risk aversion form shown in (2). No restrictions are placed on the parameter values ρ and γ , which respectively represent the pure rate of time preference and the coefficient of relative risk aversion, except that the latter must be non-negative ($\gamma \geq 0$):

$$(2) \quad U(y_t, t) = \begin{cases} e^{-\rho t} (y_t^{1-\gamma} - 1) / (1 - \gamma) & \gamma \neq 1 \\ e^{-\rho t} \ln(y_t) & \gamma = 1 \end{cases}.$$

By substituting equation (2) into equation (1), it follows that if g_t , the per-period logarithmic growth rate in consumption in the absence of climate change damage, is defined through the relationship $g_t = t^{-1} \ln(y_t^*/y_0)$, then:

$$(3) \quad \begin{aligned} \left(\frac{y_0 - p}{y_0} \right)^{1-\gamma} &= 1 + e^{-\rho t} E \left[\left((1 - D(T))^{1-\gamma} - 1 \right) e^{t(1-\gamma)g_t} \right] & \gamma \neq 1 \\ \ln \left(\frac{y_0 - p}{y_0} \right) &= e^{-\rho t} E [\ln(1 - D(T))] & \gamma = 1 \end{aligned}$$

We focus on p/y_0 , and call this our "Willingness to Pay" (WTP). The value of preventative action is directly measured by the fraction of consumption that we would willingly spend today to eliminate future climate change, all posited to be

human induced.¹ As the ratio increases, the implied value of policy to avoid climate change becomes stronger. This ratio is closely related to Pindyck's WTP metric, which considers what fraction of consumption society would pay to limit damage to some pre-determined level (1,12). For this analysis, for expositional ease, we set that pre-determined level equal to zero climate change. The qualitative results are the same for other values.

We will denote by f_1 the probability density function that the social planner assigns to T prior to the arrival of news. This density function is then amended to f_2 if the news can be interpreted as an MPS, or f_3 if the news increases variance in a way that may or may not preserve the mean.

2.2 IPCC's lowering of the lower bound as a Mean Preserving Spread

New information leads the social planner to update the probability density function for T from f_1 to f_2 . Let x and v be random variables drawn from f_1 and f_2 respectively. Then we say that f_2 is an MPS of f_1 if and only if we can express $v = x + z(x)$ for a random variable $z(x)$, where $E[z(x)|x] = 0$ for all x (15).²

An MPS results in the variance of f_2 being greater than that of f_1 , with the means of the two distributions being the same. It can also be interpreted as follows: f_2 is an MPS of f_1 if and only if f_1 second-order stochastically dominates f_2 and the two distributions have the same mean. It should be noted, though, that not every mean-preserving increase in variance can be interpreted as an MPS. (See subsection 2.3.)

We now present our main result and an immediate corollary:

Result 1. Assume T and g_t are independent. Then any MPS will increase the social planner's WTP to avoid climate change if and only if $V(T)$ is concave with respect to T .

Proof. See appendix A1.

This is the most restrictive of our results. Its intuition is best explained by looking at a slight extension in the form of a corollary:

Corollary 1. Any weakly risk-averse utility function and any weakly convex damage function, with at least one of these two conditions being strong, provides sufficient conditions for $V(T)$ to be concave with respect to T , and hence, for any MPS to increase the social planner's WTP if T and g_t are independent.

Proof. See appendix A1.

In other words, under most standard assumptions for damage and utility functions, the result holds that WTP increases as uncertainty increases (an MPS) while preserving the mean.

The only additional condition we impose here is that T and g_t be independent. At first, this might seem like a heroic assumption, given that past economic growth and temperature increases have clearly gone hand-in-hand. This positive relationship is also likely to continue in the future if policy makers still pursue economic growth through existing "dirty" technologies. However, there is significant research underway at present in the development of renewable and other non carbon based energy sources. The more successful the innovations in this field, the greater is the potential for strong economic growth at the same time as reduced climate change effects, changing the correlation to negative. Because of this the sign of any correlation between T and g_t cannot be stated with great confidence as it will depend on future technological advances, and indeed this relationship may well be time horizon dependent.³ For this reason, in our baseline

¹ If there is also natural variation in climate due to nature, this could be captured by introducing an additional random factor into the analysis.

² An MPS can be interpreted as follows: At each possible temperature outcome, f_2 introduces a new random gamble, $z(x)$. The distribution of $z(x)$ can vary with x , but the mean must be zero in all cases. Notice that this potentially includes the introduction of trivial gambles ($z(x) = 0$ with probability = 1) for some, but not all, values of x , a property we will use going forward.

³ Climate change damages, which occur over many centuries, can be valued by taking each year separately. Following this course, the methods described here would be applied by constructing a present value for the expected damages in

calibrations we assume that the correlation is zero for parsimony and return to more general correlation structures in Section 2.4 below.

Corollary 1 presents a sufficient condition for any monotonic increasing and non-convex utility function. If we restrict attention to specific utility functions, an analytic condition on the damage function that is weaker than convexity but is both necessary and sufficient can often be identified. We illustrate with the constant relative risk aversion form described in subsection 2.1.⁴

Corollary 2. Assume that T and g_t are independent. Any MPS in the distribution of possible future temperatures will lead a social planner who has constant relative risk aversion utility to have a higher WTP to avoid future climate change if and only if:

$$(4) \quad D''(T) > -\frac{\gamma D'(T)^2}{1-D(T)}$$

In the case when $\gamma = 1$ this result extends to the situation when T and g_t are not independent.

Proof. See appendix A1.

The proof of *Corollary 2* follows from the fact that equation (4) is the necessary and sufficient condition for $V(T)$ to be concave for any constant relative risk aversion utility function. The concave curvature of $V(T)$ comes from the (weakly) concave shape of the constant relative risk aversion utility function, which more than compensates for the maximum permissible concavity of the damage function given condition (4). Given concavity for $V(T)$, from *Corollary 1*, greater uncertainty (an MPS) implies greater WTP to avoid climate change.

We can apply *Result 1* directly to the IPCC's altered descriptions of climate sensitivity. In Assessment Report 4, IPCC stated that climate sensitivity "likely" lies between 2 and 4.5°C, that the most likely outcome is about 3°C, that the probability of it being below 1.5°C is below 10%, and that outcomes substantially higher than 4.5 degrees cannot be excluded (8).

To be more precise, the IPCC uses "likely" for any probability above 66% and "very likely" for any probability above 90%. It therefore stands to reason—taking the IPCC's language at face value—that the probability the IPCC meant to convey lies between 66% and 90%. In our subsequent discussion, we pick 66% for expositional ease. (Others have split the difference between 66 and 90%, using 78% instead in order to have a more conservative interpretation of the IPCC's language when it comes to analyzing the probabilities associated with tails outside that "likely" range (3). Picking 66%, 78% or any other number (or range) makes no material difference to our analysis, as long as we are consistent across scenarios.)

While others treat climate sensitivity as a continuous, log-normal calibration (2,3), for expositional ease, we illustrate our argument using a discrete distribution, f_1 . The possible outcomes are temperature increases of {1.40, 1.75, 3.00, 4.75, 6.25} °C, with associated probabilities of {7%, 10%, 66%, 13%, 4% }.

Assessment Report 5, five years later, changed the assessment of climate sensitivity primarily by lowering the lower

each of these years and then summing them to get the total cost of greenhouse gas emissions. As each cash flow is valued separately, there is no difficulty in applying different correlation structures as years stretch forward. We would also note that, in our analysis g_t is economic growth *without* climate damages, and therefore should not be confused with discussions in, for example, (3,16), which include damage effects.

⁴ Equation (A5) in the appendix provides the less analytically tractable necessary and sufficient conditions for other utility functions.

bound from 2°C to 1.5°C, expanding the 66% likely range to 1.5 to 4.5°C (9). Moreover, the IPCC assessed the probability of climate sensitivity less than 1°C at less than 5% (“extremely unlikely”) and the probability that it is greater than 6°C at less than 10% (“very unlikely”). For expositional ease, we can once again capture this assessment with a discrete distribution, f_2 , with six possible outcomes: temperature changes can equal {0.90, 1.40, 1.75, 3.3214, 4.75, 6.25} degrees, with associated probabilities {4%, 13%, 10%, 56%, 13%, 4%}.

Figure 1 plots the cumulative distribution functions of these two distributions.

[Figure 1]

Notice that f_2 has increased the mass in the left-hand tail. The offsetting rightward distribution of mass, so as to preserve the mean, is in the centre of the distribution. There is no increase in mass in the right-hand tail.

This example is ready made for *Result 1*. Define $z(x) = 0$ with probability = 1 for all values of x except $x = 3^\circ\text{C}$. Then set $z(3) = \{-2.1, -1.6, +0.3214\}$ with associated probabilities {6.06%, 9.09%, 84.85%}. Think of the gamble z as being equivalent to scientists discovering additional uncertainties attached to the previous scenario leading to a 3°C temperature increase. Thus, it represents an MPS that added no density to either tail and indeed was all concentrated on one side of the median. First, $E[z(x)|x] = 0$ for all x in this case. That’s the “mean preserving” part. Second, when we add this $z(x)$ onto f_1 , f_2 emerges. In other words, f_2 is an MPS of f_1 .

Recall how f_1 is a representation of Assessment Report 4 and f_2 is a representation of Assessment Report 5’s description of climate sensitivity. *Result 1* then tells us that the IPCC’s 2013 step to reduce the lower bound in a mean-preserving way—what superficially might have appeared to be ‘good’ news for the planet—might well have been instead bad news. Posing independence of T and g_t , this holds true as long as $V(T)$ is concave, which, in turn, holds true for any convex damage function and any weakly risk averse utility function.

2.3 Increase in the variance of the distribution

What if rather than an MPS, we interpreted the IPCC’s step of lowering the lower bound as an increase in the standard deviation? Assume that the original distribution f_1 changed to f_3 , where $\sigma_3 > \sigma_1$. At this point, we place no other restrictions on the two means (μ_1 and μ_3) nor on the specific distribution of f_1 or f_3 .

The situation now becomes more complex. Knowing the properties of the mean and standard deviation of the distribution of future temperatures is generally insufficient to determine the social planner’s WTP. We return to this point in detail below and in the technical appendices.

2.3.1 Logarithmic utility

Though our initial analysis was in terms of an MPS, WTP can increase even when the expected value of climate sensitivity diminishes by a significant amount. An elegant illustration emerges if we assume a logarithmic utility and a negative quadratic exponential damage function, $D(T) = 1 - e^{-\beta T^2}$, of the type employed by Weitzman (2) and Pindyck (1,12), among others. Given these functional forms, even if the IPCC reduced the expected value of climate sensitivity by a significant amount ($>0.05^\circ\text{C}$) in Assessment Report 5, when compared to Assessment Report 4, the WTP could still have increased (8,9).

Result 2. Assume that the social planner has logarithmic utility, and the damage function is defined as $D(T) = 1 - e^{-\beta T^2}$.⁵ Then moving from f_1 to f_3 increases the social planner’s WTP if and only if $\sigma_3^2 - \sigma_1^2 > \mu_1^2 - \mu_3^2$.

Proof. In the case of logarithmic utility, substituting the damage function into the value function gives $e^{-\rho t} E[\ln(1 - D(T))] = -\beta e^{-\rho t} E[T^2]$. This is a quadratic function with expectation determined solely by the mean and variance.

⁵ The damage function is convex if and only if $T < \sqrt{(2\beta)^{-1}}$. It also satisfies the inequality of equation (4) for all T when $\gamma \geq 1$, which includes log utility. For $\gamma < 1$, the inequality is met if and only if $T < \sqrt{(2\beta(1 - \gamma))^{-1}}$. Therefore complete convexity in the damage function is not required either for *Result 2* (except in the case of risk neutrality) or *Result 3* below.

Specifically, $E[T^2] = E^2[T] + Var[T]$, and from equation (3):

$$(5) \quad \ln\left(\frac{y_0 - p}{y_0}\right) = -\beta e^{-\rho t} (E^2[T] + Var[T]).$$

Hence, p is monotonically increasing in $E^2[T] + Var[T]$. Therefore, the social planner's WTP is greater under f_3 than f_1 if and only if $\mu_3^2 + \sigma_3^2 > \mu_1^2 + \sigma_1^2$. *QED*.

Crucially, under the conditions of *Result 2*, any mean-preserving increase in variance will increase WTP. For logarithmic utility and exponential quadratic damage functions, this is a stronger finding than *Result 1*, since not all mean-preserving increases in variance are also an MPS. Also note that *Result 2* does not depend on the independence of T and g_t .

Focusing on the mean climate sensitivity alone then could once again lead us astray. Consider the implications of *Result 2* in the context of the discrete distribution functions for climate sensitivity given in the previous subsection. Keeping all other values the same, we can reduce the most likely value of f_2 from 3.3214°C to 3.25°C to produce f_3 . The density function f_3 represents an unambiguous improvement over f_2 . The mean of f_3 is less than f_1 : 3.08°C compared to 3.12°C. Yet, as $(\sigma_3^2 - \sigma_1^2) - (\mu_1^2 - \mu_3^2) = 0.125 > 0$, we would still prefer to confront f_1 rather than f_3 . The WTP has gone up, even though the new distribution has a lower mean value of climate sensitivity, more mass in the left-hand tail, and the same amount of mass in the right-hand tail. That is because outcomes got worse in the middle. What at first glance might have appeared to be 'good' news is, in fact, 'bad' news. This result holds if we change the most likely value of f_3 to any value higher than 3.2155°C, while leaving all other values unchanged, when $\mu_3 = 3.06^\circ\text{C}$.

2.3.2 Non-logarithmic utility

For non-logarithmic utility—values of $\gamma \neq 1$ —the impact on the social planner's WTP of moving from f_1 to f_3 is more complex. This is because no natural damage function combines with non-logarithmic utility to give a quadratic value function. We are, therefore, no longer in the comfortable world of mean-variance decision-making, where *Result 2* holds. However, we conjecture that if both (i) f_1 and f_3 belong to the same major traditional families of distributions, and (ii) T and g_t are independent, then a mean-preserving increase in variance will generally increase the social planner's WTP.⁶

While it is beyond the scope of this paper to prove this conjecture for all possible families of distributions, all possible damage functions, and all potential utility functions, we can illustrate it for a particularly relevant case. Restrict both f_1 and f_3 to be drawn from a generalised gamma distribution (alternatively known as a generalised normal distribution (16)) where the second shape parameter is set to be 2; $f(T): GG(\alpha, 2, \lambda)$:

$$(6) \quad f(T) = \frac{2}{\lambda \Gamma(\alpha)} \left(\frac{T}{\lambda}\right)^{2\alpha-1} e^{-(T/\lambda)^2},$$

with $\alpha, \lambda > 0$ and $T \geq 0$. Figure 2 illustrates the probability density function (6) for four different choices of parameter values. All assume that $\mu = 3^\circ\text{C}$; but they vary in the value of σ .⁷

[Figure 2]

In this parameterisation, the probability density function, $g(T^2)$, of the square of future temperature is given by a gamma distribution $g(T^2) = \Gamma(\alpha, \lambda^2)$, where the second parameter is the scale parameter. This distribution has a well-known moment generating function.

Result 3. If $D(T) = 1 - e^{-\beta T^2}$, $f(T): GG(\alpha, 2, \lambda)$, $\gamma \neq 1$, and g_t and T are independent, then any mean-preserving increase in variance raises the WTP.

⁶ If T is drawn from a family of probability density functions whose members differ by location and scale only. ("Two cumulative distribution functions $G1(\cdot)$ and $G2(\cdot)$ are said to differ only by location parameters α and β if $G1(x) = G2(\alpha + \beta x)$ with $\beta > 0$ "; 17 p. 422.) This family includes normal and uniform distributions, then *Result 2* will generalize to non-logarithmic utility functions (17).

⁷ If we centre climate change damage around $\mu = 3^\circ\text{C}$ of warming, with a standard deviation of $\sigma = 1.5^\circ\text{C}$, then this gives parameter values of $\alpha = 1.0876$ and $\lambda = 3.2162$.

Proof. We provide the full proof of *Result 3* in appendix A3. The analytical derivation relies on the assumption that, for values of α that are likely to be relevant for describing potential future climate change damage, $(\Gamma(\alpha + 0.5)/\Gamma(\alpha))^2$ is well approximated by $-A + \alpha$ for some positive constant A . The appendix also describes a series of empirical tests of this result which do not rely on this approximation. Under a wide range of plausible parameter values, in no case is *Result 3* found to be violated.

Result 3 also implies that *Result 2* extends to non-logarithmic utility for most plausible functional forms and parameter values. Figure 3 illustrates the empirical magnitude of this result. Let p_c denote the value of p if $T = \mu$ with certainty. Figure 3 presents values for p/y_0 and p_c/y_0 under the conditions of *Result 3* and when $g_t = 0$ for different choices of γ (Panel A), σ (Panel B), and β (Panel C). The baseline parameter values used in these figures are: $\mu = 3^\circ\text{C}$, $\sigma = 1.5^\circ\text{C}$, $\gamma = 3$ (Panel B also includes $\gamma = 0.5$) and $\beta = 0.006585$.

[Figure 3]

Even when the social planner is risk-neutral ($\gamma = 0$), the predominantly convex damage function results in $p > p_c$. As γ increases, p/y_0 rises: the more risk averse the planner is, the more the social planner is willing to pay today to tackle climate change when future temperature is uncertain. This particular illustration assumes $g_t = 0$; we discuss different growth assumptions in Section 2.4 below. Moreover, as temperature uncertainty increases, so does p/y_0 , consistent with *Result 3*. Finally, and as expected, the greater are the actual damages from temperature change, as indicated by β , the higher is the WTP, and the greater is the difference between p and p_c .

In the technical appendices, we provide a counter-example to the general principle that WTP increases with a mean-preserving increase in the variance of the distribution that describes climate sensitivity. This occurs when f_1 and f_3 are strongly different probability density functions. For non-logarithmic utility, the value function will often result in a dislike of kurtosis as well as variance. It is therefore possible to construct examples with high variance (but low kurtosis) that give the social planner a lower WTP than under an alternate distribution with lower variance (but higher kurtosis). See Technical Appendix TA1. We also show in Section 2.4 below that the WTP may not always increase with uncertainty when growth and temperature changes are positively correlated.

While these counter-examples are of academic interest, their relevance for real world policy making is uncertain. For example, it seems unlikely that news about climate sensitivity would significantly reduce kurtosis while simultaneously raising the variance of our prevailing climate sensitivity distribution.

2.4 Economic growth

So far we have focused on how changes in the distribution of climate sensitivity would affect WTP. Under a wide range of conditions, both an MPS and a potentially mean-reducing increase in variance increases WTP.

We could imagine a similar analysis with respect to uncertainty regarding economic growth. For example:

Result 4. Assume that g_t and T are independent and that $\gamma \neq 1$. Then for any damage function $D(T) \in [0,1)$ and for any probability density function describing future temperature uncertainty, WTP increases following an MPS in the distribution of future consumption growth. For logarithmic utility ($\gamma = 1$), WTP to avoid climate change is independent of the properties of g_t even if there is dependence with T .

Proof. See appendix A4.

This result draws parallels with the well-known finding that the social discount rate should decline with time if future growth rates are uncertain (19,20,21).

When g_t and T are correlated, we turn to numerical analysis when evaluating the WTP. As before, let $D(T) = 1 - e^{-\beta T^2}$ with $\beta = 0.006585$, and $f(T) = GG(\alpha, 2, \lambda)$, with $\alpha = 1.0876$ and $\lambda = 3.2162$. Further, as is standard in the economics literature, assume that per-period logarithmic consumption growth is independently and identically normally distributed with constant mean $M = 1.9\%$, and standard deviation, $S = 3.0\%$. These parameter values are broadly consistent with long-run historical real per-capita annual consumption growth in the US.

We then run 25,000 simulations where g_t and T are drawn from their respective distributions and with correlation between them of ρ . The value of p/y_0 is calculated numerically. This is then compared to a “certainty equivalent” value, p^* , that captures the case when temperature is non-stochastic, but consumption growth remains uncertain. Results are presented in Figure 4 for the case when $t = 50$ years and either $\gamma = 5$ (left-hand axis) or $\gamma = 0.5$ (right-hand axis) for $\rho \in [-1, +1]$.

[Insert Figure 4 about here]

Consistent with earlier results, when the correlation is zero, greater uncertainty over future temperatures increases the willingness to pay, $p > p^*$. However, for high γ and strongly positive correlation (greater than approximately 0.4), then $p < p^*$, meaning that higher temperature uncertainty does not necessarily lead to a stronger economic rationale for acting now. In addition, our willingness to pay is much higher when the value of γ is lower.

Both these effects have the same underlying cause. Parameter γ indicates risk aversion over consumption between different periods. In contrast to the examples in the previous section, where $g_t = 0$ with certainty, consumption is now very likely to grow over the long term implying that future generations will be wealthier than our own. The higher is γ , the lower is the incentive to spend money now to help with future problems. In addition, when g_t and T are positively correlated, the largest temperature changes occur when society is wealthiest, which again is of less current concern, particularly when γ is high.

3. Removal of “best estimate” as decreased peakedness

The previous section focused on how the increase in uncertainty that is associated with the IPCC’s step, in its Fifth Assessment Report, to widen the “likely” range of climate sensitivity may well have increased the WTP. In this section, we turn to the other important recent change in the way it reports estimates of this parameter. The IPCC’s 2007 report included a “best estimate” for climate sensitivity of 3°C. No longer. In its 2013 report the IPCC abandoned its statement identifying a “best estimate.” There is a mention of a “mean” climate sensitivity parameter of 3.2°C in Chapter 12 of the full IPCC report. However, in a break from the 2007 IPCC report, neither the Summary for Policy Makers nor the chapter summaries themselves include a statement of the “best estimate” (3,9).

We interpret the resulting change as one that seems to be best captured by a look at the distribution’s decreased ‘peakedness’. To isolate the implications, we shift our analysis of climate sensitivity to a more general class of distributions while keeping mean and standard deviation constant. In particular, we use a probability-density function considered by Zeckhauser and Thompson (22):

$$(7) \quad f(z; \mu, \sigma, \theta) = [2\sigma\Gamma(1 + 1/\theta)]^{-1} \exp\left\{-\left|\frac{z-\mu}{\sigma}\right|^\theta\right\}$$

with $\sigma > 0$ and $\theta > 0$, where θ defines the distribution’s kurtosis: the higher is θ , the lower is its kurtosis, and vice versa. Trimming the peak in effect decreases the kurtosis of the distribution. **Figure 5** shows the effects: the higher is θ , the lower is the distribution’s peakedness. Importantly, Zeckhauser and Thompson’s θ is not equal to kurtosis. It is rather a parameter that directly affects kurtosis (what we call peakedness throughout the text), even though it operates in the opposite direction. For a normal distribution, $\theta = 2$, while kurtosis = 3. θ increases with a decrease in peakedness. The kurtosis of a uniform distribution equals 1.8, while $\theta = \infty$ (22).

[Figure 5]

The assumed distribution is symmetric, and is not cut off at zero. These properties make it far from perfect to describe climate sensitivity. Indeed, they tilt our results toward conservatism in the sense that we are clearly underestimating the true uncertainties involved, in particular on the upper tail of what might more accurately be captured by a skewed distribution.

Result 5. Removing the “best estimate” for climate sensitivity, a step we interpret as decreasing peakedness of the distribution, increases WTP.

We rely on Pindyck’s (1,12) model to analyse the implications of varying kurtosis, employing the Zeckhauser-Thompson distribution from equation (7) in lieu of Pindyck’s displaced gamma distribution. The lower is peakedness (the higher is

θ), the higher is the WTP to avoid damages from climate change, in particular for a constant 1.5-4.5°C “likely” range (Figure 6).

[Figure 6]

“Constant σ ” in Figure 6 shows the results of varying θ , while keeping everything else constant. That shows the cleanest possible trade-off of various levels of peakedness. However, it also changes the probabilities of climate sensitivity between 1.5 and 4.5°C. In particular, increasing θ without other adjustments increases the likelihood of being between 1.5 and 4.5°C, cutting off tails (on both ends) even further.

By adjusting the standard deviation, σ , in addition to θ , we hold constant the IPCC’s “likely” probability. The difference is small but important in itself. A constant “likely” probability guarantees that WTP is a strictly increasing function for reasonable values of θ . Any amount of increased uncertainty *within* the 66% likely range—a decrease in peakedness—leads to an increased WTP to avoid climate damages.

Importantly, this result plays out entirely within the “likely” climate sensitivity range. Higher θ (lower peakedness) implies less density in the tails, which typically drive the results (2,3). Hence, here WTP increases with decreased peakedness *despite* decreased mass in the upper tail, not because of it. That decreased density in the upper tail is also the reason why the WTP-line for constant standard deviation tilts downward slightly with θ above 3, and why the WTP for a constant 66% “likely” line tapers off (23).⁸

Given that, it is important to add a warning: Figure highlights the importance of relative differences across WTP levels with different levels of peakedness; the absolute WTP levels are largely irrelevant, as the uncertainty of climate sensitivity here does not play out in the all-important (fat) upper tail of the distribution. It operates solely within the 1.5 to 4.5°C IPCC “likely” range.

4. Conclusion

Climate sensitivity—the eventual temperature outcome based on a doubling of carbon dioxide concentrations in the atmosphere—is the key parameter determining long-run global average temperature outcomes and, thus, the costs of failing to take significant steps to mitigate future climate change. Doubling of carbon dioxide concentrations in the atmosphere is not a hypothetical. Current concentrations of around 400 parts per million (ppm) are already up by 40% over pre-industrial levels of roughly 280 ppm. At the current rate of increase, pre-industrial levels will double well before the end of the century.

The title question is also not a hypothetical. Some apparently new facts have led IPCC (2013) to reconsider the full implications for long-term equilibrium warming and have led it to widen the “likely” range, once again, to 1.5 to 4.5°C. In addition, the IPCC removed its prior “best estimate” of 3°C.

Unambiguous conclusions are hard to reach in this arena where uncertainties are large. Bearing this caution in mind, our analysis strongly suggests that, other things equal, the IPCC’s recent widening of the “likely” range of temperature change by reducing its bottom value may well increase appropriately calibrated WTP to avoid such change in the future. The same holds for the removal of the “best estimate.” Given the increasing marginal costs of global warming, greater uncertainty raises the return for taking action to curb greenhouse emissions. Recent steps in the IPCC’s assessments reflect greater uncertainty.

Additional Information

Acknowledgments

⁸ If we did not have the result of diminished peakedness in the updated IPCC report, we might think that $\theta < 2$, implying a higher peak but much more density in the tails as compared with a normal distribution. Whether this would be better or worse than the normal case would depend on both the form of the damage and utility functions.

We thank Katherine Rittenhouse for excellent research assistance and Michael Aziz, Frank Convery, Howard Kunreuther, Chuck Mason, Ilissa Ocko, Michael Oppenheimer, Daniel Schrag, Katheline Schubert, Thomas Sterner, Martin Weitzman, Matthew Zaragoza-Watkins, and seminar participants at Duke, Harvard, the University of Gothenburg, the University of Minnesota, the 2014 American Economic Association meetings, and the 2014 World Congress of Environmental and Resource Economists for comments and discussions. All remaining errors are our own.

Funding Statement

None specified.

Competing Interests

We have no competing interests.

Authors' Contributions

All authors contributed equally to this paper.

References

1. Pindyck, Robert S. 2013. "The Climate Policy Dilemma." *Rev. Environ. Econ. Policy* **7**(2), 219-237. (doi: 10.1093/reep/ret007)
2. Weitzman, Martin L. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Rev. Econ. Stat.* **91**, no. 1, 1-19. (doi:10.1162/rest.91.1.1)
3. Wagner, Gernot and Martin L. Weitzman. 2015. *Climate Shock: the Economic Consequences of a Hotter Planet*. Princeton University Press.
4. Lewandowsky S, Risbey JS, Smithson M, Newell BR. 2014. "Scientific uncertainty and climate change: Part I. Uncertainty and unabated emissions." *Clim. Change* **124**, no. 1-2, 21-37. (doi: 10.1007/s10584-014-1082-7)
5. Lewandowsky S, Risbey JS, Smithson M, Newell BR. 2014. "Scientific uncertainty and climate change: Part II. Uncertainty and mitigation." *Clim. Change* **124**, no. 1-2, 39-52. (doi: 10.1007/s10584-014-1083-6)
6. Pindyck, Robert S. 2014. "Risk and return in the design of environmental policy." *JAERE* **1**, no. 3, 395-418. (doi:10.1086/677949)
7. Charney, Jule G., Akio Arakawa, D. James Baker, Bert Bolin, Robert E. Dickinson, Richard M. Goody, Cecil E. Leith, Henry M. Stommel, and Carl I. Wunsch. 1979. "Carbon dioxide and climate: a scientific assessment." *National Academy of Sciences*.
8. Intergovernmental Panel on Climate Change (IPCC). 2007. *Fourth Assessment Report, Working Group I*.
9. Intergovernmental Panel on Climate Change (IPCC). 2013. *Fifth Assessment Report, Working Group I*.
10. Roe, G., & Baker, M. (2007). Why Is climate sensitivity so unpredictable? *Science* **318**, no. 5850, 629-632.
11. Stainforth, D. A., Allen, M. R., Tredger, E. R., & Smith, L. A. (2007). Confidence, uncertainty and decision-support relevance in climate predictions. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **365**, no. 1857, 2145.
12. Pindyck, Robert S. 2012. "Uncertain outcomes and climate change policy." *J. Environ. Econ. Manage.* **63**, no. 3, 289-303. (doi: 10.1016/j.jeem.2011.12.001)
13. Weitzman, Martin L. 2010. "What Is The "Damages Function" For Global Warming—And What Difference Might It Make?." *Clim. Change Econ.* **1**, no. 01, 57-69. (doi: 10.1142/S2010007810000042)
14. Millner A, Dietz S, Heal G. 2013. "Scientific ambiguity and climate policy." *Environ. Resource Econ.* **55**, 21-46. (doi: 10.1007/s10640-012-9612-0)
15. Rothschild, Michael, and Joseph E. Stiglitz. 1970. "Increasing risk: I. A definition." *JET* **2**, no. 3, 225-243. (doi: 10.1016/0022-0531(70)90038-4)
16. Dietz, Simon, Christian Gollier, and Louise Kessler. 2015. "The climate beta." Grantham Research Institute on Climate Change and the Environment Working Paper no. 190.
17. Meyer, J. 1987. "Two-moment decision models and expected utility maximization." *Am. Econ. Rev.* **77**, 421-430.
18. Khodabina M, Ahmadabadib A. 2010. "Some properties of generalized gamma distribution." *Mathematical Sciences* **4**, 9-28.
19. Arrow, K., M. Cropper, C. Gollier, B. Groom, G. Heal, R. Newell, W. Nordhaus, Robert S. Pindyck, William A. Pizer, Paul R. Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman. 2013. "Determining Benefits and Costs for Future Generations." *Science* **341**, no. 6144, 349-350. (doi: 10.1126/science.1235665)

-
20. Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, Robert S. Pindyck, William A. Pizer, Paul R. Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman. 2014. "Should Governments Use a Declining Discount Rate in Project Analysis?" *Rev. Environ. Econ. Policy* **8**, 145-163 (doi: 10.1093/reep/reu008).
 21. Cropper ML, Freeman MC, Groom B, Pizer WA. 2014. "Declining discount rates." *Am. Econ. Rev.* 104:538–543. (doi: 10.1257/aer.104.5.538)
 22. Zeckhauser, Richard, and Mark Thompson. 1970. "Linear regression with non-normal error terms." *The Rev. Econ. Stat.* **52**, no. 3, 280-286.
 23. *The Economist*. 2013. "Climate science: a sensitive matter." March 30th.
 24. Nuccitelli, Dana and Michael E. Mann. 2013. "How The Economist got it wrong." Australian Broadcasting Corporation Environment, 12 April.
 25. Cowtan, Kevin, and Robert G. Way. 2013. "Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends." *Q. J. Roy. Meteor. Soc.* **140**, no 683, 1935-1944 (doi: 10.1002/qj.2297)
 26. *The Economist*. 2014. "Global Warming: Who Pressed the Pause Button." March 8th.




Figure 1—Cumulative distribution functions for climate sensitivity under IPCC Assessment Reports 4 (AR4) and 5 (AR5). The distributions are represented by functions f_1 and f_2 , respectively.







Figure 2—The probability density function of future temperature changes, which are assumed to be generalised gamma distributed with second shape parameter equal to 2: $f(T): GG(\alpha, 2, \lambda)$. The values of α and λ are chosen so that the mean of T , $\mu = 3^\circ\text{C}$ and the standard deviation of T , $\sigma \in \{0.5^\circ\text{C}, 1.0^\circ\text{C}, 1.5^\circ\text{C}, 2^\circ\text{C}\}$.



Panel A: Risk aversion



Panel B: Standard deviation of temperature



Panel C: Temperature damages

Figure 3—This figure illustrates the amount that a rational social planner would spend as a proportion of current consumption to prevent climate change both in the presence, \mathbf{p} , and absence, $\mathbf{p_c}$, of future temperature uncertainty. Panels A, B and C respectively vary the coefficient of relative risk aversion, γ , the standard deviation of future temperature, σ , and the extent of the damages caused by climate change β .




Figure 4—This figure illustrates the amount, p/y_0 , that a rational social planner would spend as a proportion of current consumption to prevent climate change when there is uncertainty concerning both potential consumption growth and future temperature levels. Results are shown for a range of different correlations between these two variables. The figure also shows the equivalent amounts when temperature change is known but consumption is stochastic (p^*/y_0). Results are shown for $\gamma = 5$, to be read against the left-hand axis, and $\gamma = 0.5$, which should be read against the right-hand axis.




Figure 5—Climate sensitivity distribution, calibrating a standard normal distribution to the IPCC’s “likely” range of 1.5 to 4.5°C.




Figure 6—Willingness to Pay (WTP) to avoid climate damages at various levels of peakedness (inversely related to θ).