Confronting Deep and Persistent Climate Uncertainty

Gernot Wagner and Richard J. Zeckhauser[1]
20 November 2018

Abstract

Deep-seated, persistent uncertainty is a pernicious feature of climate change. One key parameter, equilibrium climate sensitivity, has eluded almost all attempts to pin down more precisely than a 'likely' range that has stalled at 1.5−4.5°C for over thirty-five years.

The marginal damages due to temperature increase rise rapidly. Thus, uncertainty in climate sensitivity significantly raises the expected costs of climate change above what they would be if the temperature increase was known to be close to a mean value of 3.0°C. The cost of this uncertainty is compounded given that the distribution of possible temperature changes is strongly skewed toward higher values.

Keywords*: Climate sensitivity; uncertainty, risk, fat tails; mitigation, adaptation.
JEL codes*: Q54, D81.

1.

Climate science is not settled.[2]

A statement like this makes good sense. Science is rarely truly 'settled'. But with climate change, the words have taken on a whole new meaning among those trying to distort the science and delay policy (Oreskes and Conway, 2011; Supran and Oreskes, 2017). We focus on a particular aspect of the science not being settled: the deep and persistent uncertainties inherent in climate science. Such deep-seated uncertainties have two critical implications: First, they suggest that there are some possible outcomes due to climate change that we can't even conjecture. Major surprises must be expected. Second, estimates of the fundamental relationships in climate change, notably climate

[2] Slight variations of this statement have appeared repeatedly over time, at least twice as titles of opinion articles in the *Wall Street Journal* alone (Koonin, 2014; Lindzen, 2009).

sensitivity, have stayed persistently uncertain for decades. Despite recent advances in pinning down climate sensitivity, critical tail uncertainties persist.[3]

Climate change has been labelled as "the greatest market failure the world has ever seen" (Stern, 2006) and "the mother of all externalities" (Tol, 2009). Tol (2009) continues by calling it: "larger, more complex, and more uncertain than any other environmental problem." It is. And the uncertainty itself has multiple dimensions.

For one, going from economic activity to climate damages that impact people's lives entails several links in a chain that leads from economic outputs to emissions, from emissions to concentrations, and from concentrations to temperatures for the most discussed climate-related measure humans actually care about. Temperature changes, in turn, lead to climate impacts, which then generate damages. Each of these links has the potential to surprise. A few of the unknowns in this chain represent risks, in the Knightian sense of the term, capable of being captured by a probabilistic function. The vast majority are what Knight (1921) would term uncertainties, where the probabilities themselves are unknown. And some are what we might label 'unknown unknowns' or simply 'ignorance', where even the possible outcomes cannot be identified (Zeckhauser, 2006).

Each of these links brings with it a set of additional complications. For instance, enormous regional variations render already uncertain mean temperature increases even more so. Similarly, local temperature effects are but one of many consequences of climate change. Other prominent effects include increased extreme weather events, increased ocean acidification, disturbances to the global water cycle, and profound effects on biodiversity. Each of these has important localized consequences.

We focus on the link between concentrations and temperatures—specifically what happens to eventual global average temperatures in equilibrium as concentrations of carbon dioxide (CO_2) in the atmosphere double. Arrhenius (1896) first calculated the answer to this all-important magnitude, since named 'climate sensitivity'. His estimate: 5−6°C. That range has by now proven to be too pessimistic. Instead, the current 'likely' range stands at 1.5−4.5°C (IPCC, 2013). That range has persisted for almost four decades.

That consistently wide range is the first and one of the two most significant uncertainties linked to climate sensitivity. Charney *et al.* (1979) first established the 'likely' range of climate sensitivity that remains in use today, despite considerable efforts to refine it. By now, scientists have formulated a more precise idea of what it means to be 'likely'. The Intergovernmental Panel on Climate Change (IPCC) defines it as having a greater-than-66-percent probability (Wagner and Weitzman, 2015, p. 50). That range, in turn, has been 1.5−4.5°C throughout most of IPCC's history (IPCC, 2013, 2001, 1995,

---

[3] Cox et al. (2018a) represents only the latest attempt at specifying an "emergent constraint" on climate sensitivity. This article has sparked a vivid debate and a series of formal responses (Brown et al., 2018; Po-Chedley et al., 2018; Rypdal et al., 2018; Wagner and Weitzman, 2018), followed by a formal reply by Cox et al. (2018b).

1990). The sole exception is the Fourth Assessment Report, which narrowed the range to 2−4.5°C (IPCC, 2007). The range was expanded back to 1.5−4.5°C in the Fifth Assessment Report (IPCC, 2013). Here we do not question the motives behind the IPCC's steps. We simply attempt to interpret the implications.

Meanwhile, the second major uncertainty addresses the skewness of the distribution (Figure 1). Skewness is of secondary interest below the likely range, as science can cut off the lower tail of the climate sensitivity distribution; adding CO_2 is not going to decrease temperatures. That is not the case on the upper end, above the cutoff of the 'likely' range, where skewness stretches out the right tail as it does in our illustrative calibration in Figure 1.


Figure 1—Climate sensitivity, fitting an illustrative log-normal distribution around the IPCC's (2013) 'likely' range (Wagner and Weitzman, 2015).[4]

Given that the expected costs associated with temperature increase at the margin—i.e., the expected cost curve is convex (Figure 2)—greater skewness magnifies the costs and risks of climate change, all else equal. Here it is important to note the well-known fact that the higher the equilibrium climate sensitivity, the longer it will take for global average temperatures to reach that equilibrium. All else, thus, may not be equal, as discounting of future damages may make the tail of the distribution less severe (Hogan and Wagner, 2018; Roe and Bauman, 2012). Much then, in turn, depends on the appropriate discount rate. There, too, uncertainty plays a role, as uncertainty around the

---

[4] Note that the calibration assumes equal probability mass below and above the likely range. In particular, we follow Wagner and Weitzman (2015) in splitting the difference between the IPCC's 'likely' (66%) and 'very likely' (90%) cut-offs. That implies that the mass between 1.5 and 4.5°C is actually 78%, with 11% probability below 1.5°C and 11% above 4.5°C. That, like all other assumptions taken in our calibrations, is what we describe as a 'conservative' step. Importantly, when we say 'conservative', we mean an assumption that biases our Willingness to Pay (WTP) to avoid climate change downward—unlike 'conservative' how it is commonly used in risk management, which would lead to a much higher WTP in order to minimize climate risk whenever possible.

right rate points to declining rates over time (Arrow et al., 2013, 2014; Cropper et al., 2014; Gollier et al., 2008; Gollier and Weitzman, 2010; Heal, 2017).


Figure 2—Illustrative damage functions with quadratic and exponential extrapolations (Wagner and Weitzman, 2015). [5]

The risks become all the more significant once oft-unquantifiable 'tipping points' and other non-linearities are introduced to the damage function (Hsiang and Kopp, 2018; Kopp et al., 2016; Revesz et al., 2014). In theory, these tipping points could be dealt with the same way any other type of uncertainty is included in decision models: using expected utility theory. A certainty-equivalent[6] level of damages would be computed for each value of temperature increase. Perfectly constructed expected damages curves take account of tipping points. Thus, any jumps in damages get accounted for by the slope of the certainty equivalent curve. While the latest econometric damage function exercises establish ever better estimates for what is known and measurable (Houser et al., 2015; Hsiang et al., 2017), computing a curve that incorporates potential non-linearities beyond the historical record is exceedingly difficult to do, and must be expected to be speculative (Auffhammer, 2018; Convery and Wagner, 2015; Houser et al., 2015; Kopp et al., 2016; Wagner and Weitzman, 2015).

The implications of a heavily right-skewed distribution for climate sensitivity have been studied and interpreted in numerous ways (Convery and Wagner, 2015; Heal, 2017; Wagner and Weitzman, 2015; Weitzman, 2009, 2011, 2012, 2014, 2015), with Roe and Bauman (2012) adding an important counterpoint around the timing element.[7] We will

---

[5] The illustrated exponential damage function in Figure 2 is calibrated to produce damages equal to the quadratic at 4°C.

[6] The certainty equivalent of a lottery is the amount an individual would accept for sure rather than get the prize of the lottery. A risk averse individual might be indifferent between, say, losing $60 for sure and having a ½ chance of losing $0 and a ½ chance of losing $100. Thus −$60 is that individual's certainty equivalent for the lottery [½, $0; ½, −$100].

[7] See Wagner and Weitzman (2015) for a partial response, and Hogan and Wagner (2018) for an attempt at a more complete one.

instead focus on the first kind of uncertainty: the stubbornly persistent wide 'likely' range.

The IPCC's verdict around what constitutes 'likely' is based on a number of different estimates ranging from paleo-climatic data to current temperature records and climate models.[8] It is also important to note that it itself is conservative in yet another sense, namely the equilibrium climate sensitivity parameter estimated here comes under a different name: 'fast climate sensitivity'. 'Fast' applies to what happens over decades or a few centuries. That is distinct from 'Earth system sensitivity', incorporating long-run Earth system feedbacks playing out over many centuries or millennia, which could be more than double the 1.5−4.5°C range (Knutti and Hegerl, 2008).[9]

This definition highlights an important point: whenever asked to make a choice, we tilt toward the more conservative assumption, that is the assumption leading to a lesser estimate of losses. That goes for the 'fat tails' question illustrated in Figure 1, where the actual log-normal calibration used turns out not to be a fat-tailed distribution but rather a 'heavy-tailed' distribution that lies at the cusp of thin- and fat-tailed.[10] Similarly, our focus on fast climate sensitivity does not capture the full extent of warming. It captures what happens in an 'equilibrium' that does not take Earth system feedbacks into account, which in expectation point to much higher, ultimate climate sensitivity.

Positing that the upper tail has been stretched out is clearly bad—in the sense that it ought to magnify concerns about global warming and, rationally, lead to a great willingness to pay (WTP) to avoid global average temperature increases. Uncertainty around the most likely value similarly increases WTP. Marginal damages from rising temperatures increase rapidly. Even in the typically conservative calibration used by the most prominent top-down, climate-economy models, damages increase quadratically (Auffhammer, 2018). More specifically, they follow an assumed quadratic loss function where climate damages $D = (1 - Y)[1 + a \Delta T + b (\Delta T)^2]$. Making the not unreasonable assumption that damages increase more steeply, especially at higher temperatures, would make the results worse. Weitzman (2009), for example, uses an exponential loss function, with $D = (1 - Y)[exp(\beta)(\Delta T)^2]$.[11]

Under these assumptions, it is then clear that an increase in the upper bound of the 'likely' interval would increase WTP to avoid any particular or probabilistic temperature increase. However, what happened between Assessment Reports Four and Five (IPCC,

---

[8] The recent "emergent constraint" work centers on model-based climate sensitivity estimates. See footnote 3.

[9] See Proistosescu and Huybers (2017) for a further exploration of fast versus slow modes of climate sensitivity.

[10] A 'fat-tailed' distribution is one that approaches zero no faster than a polynomial. A 'thin-tailed' distribution, by contrast, approaches zero at least as fast as an exponential. A log-normal distribution, which we use, is on the knife-edge between thin and fat tailed. Following Wagner and Weitzman (2018, 2015), our calibration employs a log-normal distributions, which approaches zero faster than polynomially but slower than exponentially.

[11] The loss function itself is typically written as an inverse, with the loss equal to $1/[1 + a \Delta T + b (\Delta T)^2]$ and $exp(-\beta)(\Delta T)^2$ for quadratic and exponential loss functions, respectively.

2013, 2007) was that the lower bound of the 'likely' range was lowered. In particular, in 2013, the IPCC lowered the lower bound from 2°C in 2007 back down to 1.5°C, where it had been since 1979.

Superficially, this looks like unambiguous good news; part of the distribution had shifted downward. That indeed would be unambiguous good news if the distribution outside the likely range remained the same. However, this expansion of the likely region reflected greater uncertainty. That, in turn, almost certainly implied that the whole distribution had become broader. Most importantly, uncertainty above the likely range is unambiguously bad. That is also where the marginal damages curve is steepest. This loss might overshadow any reassurance from a lowered estimate on the mean.

If the damages function were determined by merely the mean and the standard deviation, a common but not necessarily realistic assumption, then the traditional mean-standard deviation tradeoff comes into play (Pindyck, 2014). We focus here simply on this mean-standard deviation tradeoff, beginning with section 2. Section 3 provides a more technical discussion, in part since published and expanded upon in Freeman *et al.* (2015). Section 4 discusses some broader implications of persistent climate uncertainties. Section 5 concludes.

2.

Had global warming turned out to be less severe than previously thought, say if the whole distribution or even some portion of the distribution had shifted downward, that would be cause for celebration. But merely reducing the bottom value of the likely distribution in the IPCC report hardly represents such a shift. Rather, it also spreads the overall distribution. Thus, it tells us about the current capabilities of climate science—notably the current understanding of the climate sensitivity parameter—and indicates that the uncertainties are greater than we thought. That alone is disturbing, since greater variability indicates greater expected cost. Thus, the news is bittersweet, a probable reduction in the mean in exchange for greater variability.[12] And the surprise on uncertainty is even more disturbing, since the relationship between CO_2 emissions and global temperatures is perhaps the most studied relationship in the climate debate. Uncertainties elsewhere may be even greater.

Some climate uncertainties are neither 'deep' nor persistent, though the one around climate sensitivity surely is. It is unlike the typical risk situation, where we know the odds we face, though we can't predict the outcome. It also goes beyond traditional uncertainty, where we know the possible outcomes but don't know the probabilities attached to each (Knight, 1921). A situation where even the future states of the world cannot be defined goes well beyond traditional uncertainty, where some outcomes cannot be conjectured. It is what elsewhere has been called "ignorance" (Zeckhauser,

---

[12] Indeed, given right skewness, lowering the lower bound on the 'likely' region may actually increase the mean temperature change given that the tail beyond the 'likely' region will stretch out.

2006), and that we here refer to as deep and persistent uncertainty. This type of uncertainty resides in the realm of unknown unknowns (Morris, 2014; Rumsfeld, 2002). That is the realm where we are with climate sensitivity and, hence, long-run climate change projections in general.

To employ an anthropocentric metaphor, Nature hides her secrets. And on climate sensitivity, she hides them deeply and persistently. A great variety of phenomena have to be understood if they are to be uncovered, and some important phenomena remain unknown today. Some critical phenomena may not have been identified. Other factors equal, the further are today's readings, say on temperature, from what we expected yesterday, the more we should conclude that we do not understand. Or to put the same point in statistical language, the greater today's uncertainties, the greater should be the spread in the subjective distribution of outcomes for what we expect in the future.

This encapsulates our main point: an unexpectedly good low reading today may turn out to be bad news from an expected utility standpoint, given that loss functions are convex, and the distribution of temperature changes is right-skewed. Though it may well shift the mean of what we expect in the future in a favorable direction, it should simultaneously warn us of our unexpectedly weak understanding. And that disappointment about understanding implies that the standard deviation of what we should expect for the future has increased. Given the shape of loss functions, accompanied by right skewness of climate sensitivity, we should be risk averse on variability in climate change, indeed probably strongly so. Thus, the loss from an increased standard deviation may well outweigh any gain from decreased mean.

A 50% chance of a 2°C increase in temperature and a 50% chance of a 4°C increase is far worse than a 3°C increase for sure, and may well be worse than, say, a 2.8°C increase for sure. That is due to the function that maps average temperatures into damage estimates, which increases at an increasing rate.

We wish to make three points clear. First, we recognize that a lesser temperature rise will give more credence to those who have criticized climate 'alarmists' as having significantly exaggerated matters. (However conscientious scientists are, their median estimates of consequences will be too high half the time, and never precisely on target.)

Second, we would be happy to accept good news. Thirty years hence, we may find, indeed it is 10% likely, that we will discover that climate change was less serious than our current 10th percentile estimate (assuming that the estimation process is unbiased). That would be fantastic news. By then, we may well know enough from improved science and physical observation to be pretty confident that the good outcome is real.

Third, we recognize that significant catastrophes can happen with parameters that are merely uncertain. If we look at the 2007-8 financial meltdown, the primary terrible outcome was that a broad array of financial assets would tumble like dominoes. This was not an outcome that was beyond imagination. In June 2005, *The Economist* declared "the worldwide rise in house prices" to be "the biggest bubble in history" and exhorted its readers to "prepare for the economic pain when it pops" ("In come the

waves," 2005). But quite apart from such a warning, the collapse that followed was merely a low probability event, no doubt a very low probability event, of the type that we have seen before, assuming that our concern were merely the prices of financial assets. Think of the Asian financial crisis a decade earlier, triggered by the collapsed of the Thai baht. The Dow Jones may plummet to 10,000 or zip up to 30,000, but it will always remain on that scale. It will not turn purple. If there were a chance that financial assets turned into colors, a chance that we could not possibly imagine, then we would be in the world of true ignorance (Zeckhauser, 2006).

Financial markets produce a much easier assessment problem than does climate change. Though they may be prone to bubbles, the price of an asset today reflects what the market thinks it will be worth in the future taking interest rates and intermediate returns into account. In some sense, we get a continual update of the consensus view among vast numbers of investors. And those updates move around responding to daily snippets of information. With climate sensitivity, by contrast, the IPCC only gives us updates every few years. Those updates, in turn, may be highly politicized.

If temperatures in any one year were notably lower than expected, that would be good news. But it might well mean that temperatures were suppressed by a transitory phenomenon we understand, such as volcanic activity, and will soon increase even faster to return to the underlying trend. It surely tells us that the trend itself is more uncertain and our scientific understanding less secure than we previously thought. Conceivably an entirely new process has been at work, one that the scientific literature has yet to consider.

In 2010, the U.S government, deploying the latest models, calculated a social cost of CO_2 at $25 per ton.[13] Re-running the same models in 2013, the number had risen to $40, 60 percent higher.[14] None of the underlying assumptions, such as the discount rate, changed to produce this latter update. The main reason for the shift was due to relatively small updates based on the latest climate models circa 2007, at the time when the fourth IPCC report was published. The major changes involved updates to the carbon cycle representation, and an explicit model of damages from sea-level rise. The latest, 2013, IPCC report would likely lead to further changes, and much progress has been made on a number of dimensions, most prominently perhaps on econometric estimates of damage functions (Auffhammer, 2018; Houser et al., 2015; Hsiang et al., 2017).

The three climate-economic models used in calculating the U.S. social cost of carbon, thus, lag at least half a decade behind incorporating consensus climate science, and much further when it comes to incorporating the latest climate damage estimates. It

---

[13] The exact number, in 2007 US$, was $21 for a ton of CO_2 emitted in 2015, assuming a central 3% discount rate (U.S. Government Interagency Working Group on Social Cost of Carbon, 2010). In 2015 US$, the number is roughly $25.

[14] This is the 'central' estimate for a ton of CO_2 emitted in 2015. The precise number, in 2007 US dollars, is $37, slightly revised down to $36 in 2015 due to a technical fix (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015). Either number is around $40 in 2015 US$.

takes time for the findings of the latest IPCC report to make their way through the peer-review process, and ultimately to feed into policy discussions. The U.S. National Academy of Sciences has since reviewed the process and assumptions leading to the U.S. social cost of carbon and has recommended major revisions based on the latest science (National Academy of Sciences, 2017).[15]

In any case, the $40 per ton of CO_2 externality cost is best viewed as a conservative estimate. Most of the factors that we know are left out would further increase the number. Most of what we don't know would push the number further still (Wagner and Weitzman, 2015). Of course, there could be unforeseen developments that push the number down. One large unknown involves the arena of 'carbon dioxide removal' (CDR). Such interventions would act as a direct backstop technology if its cost, including the costs of the environmental risks it created, would prove to be sufficiently low, or if the social cost of carbon were to climb sufficiently high.[16]

Another large unknown involves the potential use of solar geoengineering—i.e., deliberate attempts to change the Earth's albedo to reflect more sunlight back into space and thereby cool the planet (e.g., Keith, 2000; National Research Council, 2015). It is inconceivable to us that this technology—despite its risks—would not at least be tried in a (perceived) climatic emergency, and long before global average temperatures reach anywhere close to the tails of the currently hypothesized temperature distributions (Figure 1). Despite being quite different than CDR, it, too, could have direct carbon benefits (Keith et al., 2017), which would clearly change the calculation of the optimal carbon price.

3.

Double CO_2 concentrations and, consensus climate science has told us for almost 40 years, long-run temperatures will rise in expectation by around 3°C. That number has stood ever since Jule Charney chaired a National Academy of Sciences Ad Hoc Study Group on Carbon Dioxide and Climate in the late 1970s (Charney et al., 1979). His range around the average number was plus-minus 1.5°C. That range, too, has withstood the test of time.

In 1990, the IPCC picked up Charney's range and enshrined it as its 'likely' range for equilibrium climate sensitivity of 1.5–4.5°C (IPCC, 1990). That verdict held for a further fifteen-plus years of increasingly intense scrutiny until in 2007, when the IPCC decided to trim the bottom of the range, which became 2–4.5°C (IPCC, 2007). Apparent bad news: the lowest estimates for climate sensitivity seemed to be ever more out of reach.

---

[15] The National Academy of Sciences itself, of course, is not infallible (e.g., Kelleher and Wagner, 2017) and many others have suggested further updates and entirely different approaches (e.g., Heal, 2017; Pindyck, 2017).

[16] See National Academies of Sciences (2018) for the latest comprehensive review.

Instead, and as a corollary of our other analysis here, the narrower 'likely' range was, in fact, probably good news in the sense that it lowered the WTP to avoid climate damages.

That narrower 'likely' range remained in place until 2013 when the IPCC, met with new information, reverted to its prior consensus statement and 1.5°C once again came back as the lower bound. Note that this new old range of 1.5–4.5°C only covers what is by the IPCC's own verdict 66% likely. Obviously, there is a chance that the true value will be even lower than 1.5°C or higher than 4.5°C. In fact, the IPCC ventures more precise guesses as to the likelihood of either possibility. In particular, the 2013 report puts the probability of climate sensitivity below 1°C at 5% or below, what it calls 'extremely unlikely', and the probability of climate sensitivity of above 6°C at 10% or below, what it calls 'very unlikely'. That strongly asymmetric verdict points to the IPCC's assessment that climate sensitivity distribution is clearly skewed toward higher values. That skewness is worrisome, though not our focus here.

Instead, this section focuses on the implications of the mean-standard deviation tradeoff inherent in the IPCC's lowering of the lower climate sensitivity bound. In doing so, we follow perhaps the most direct—if still imperfect—way of looking at the implications of higher temperatures: WTP to avoid climate damages (Pindyck, 2012, 2013).

It is direct because it captures the economic essence of the problem: the worse the (economic) consequences of climate change, society's WTP, appropriately computed, should be higher to avoid those consequences. Importantly, our WTP measure is unlike the WTP of individuals, which are deeply dependent on human nature. Psychologists tell us that an individual's WTP often bears little relation to the magnitude of the problem (Wagner and Zeckhauser, 2011). Framing, for example, typically trumps rational responses: ask people the maximum amount they would pay, $1, $2, or $3, for a particular widget, and responses might tend toward $2. Ask people the same question using $2, 4, or 6, and the most common response might well be twice as high.

WTP here in no way involves asking people their preferred number. It is the direct result of an intertemporal optimization problem looking at climate damages over time. Specifically, we follow Pindyck (2012, 2013), and ask how much society should be willing to pay to avoid global average temperatures exceeding a particular temperature increase by 2100. In most of our calibrations, we look to 2°C as our threshold for eventual global average warming not to be exceeded.[17]

Note the important difference between this overall 2°C threshold of average warming by 2100 and the climate sensitivity parameter itself. The latter estimates what happens to equilibrium temperatures—centuries hence—from a doubling of atmospheric CO_2 concentrations. The former looks to actual temperatures in 2100. The link between the two is given by concentrations of atmospheric CO_2 and other greenhouse gases, and a difference in time scales. The transformation from climate sensitivity to actual

---

[17] Our focus on 2°C stands in contrast to Pindyck (2012), who focuses on 3°C for most of the analysis and also 0°C in some sensitivity tests (Freeman et al., 2015).

temperatures in 2100 requires addressing further uncertainties beyond the climate sensitivity parameter itself.

We discuss two such uncertainties here: predicting concentrations of greenhouse gases in the atmosphere, and time scales for global average temperatures to reach equilibrium. Without a massive global decarbonization effort, global atmospheric CO_2 concentrations are expected to pass 560 ppm, double their pre-industrial levels of 280 ppm, well before the end of the century. But we cannot simply multiply projected concentrations in 2100 by the climate sensitivity, as the latter shows warming in equilibrium. In fact, projected global average temperature increases by 2100 under business-as-usual conditions, based on four wildly different IPCC scenarios, range from a low of 1°C to a high of 5.5°C above pre-industrial levels (IPCC, 2013). Taking a rough average of the outer bounds of this range yields an average of slightly above 3°C of average warming by 2100, assuming no further climate policy—either in the form of massive decarbonization or also direct changes to radiative forcing via solar geoengineering. Nevertheless, we follow Pindyck (2012, 2013) in ignoring the potential of solar geoengineering and the difference in time scales and instead conflate the two distributions of climate sensitivity and temperatures by 2100.

Specifically, we follow Pindyck's (2012, 2013) calibration, but depart in two ways: First, instead of relying on Pindyck's displaced gamma distribution for temperature change, we rely on a log-normal calibration most prominently employed by Weitzman (2009).[18] To do so, we replace the temperature probability distributions plotted in Figure 1 of Pindyck (2013) with a standard log-normal distribution calibrated to a 66% probability of being within the IPCC's 'likely' range (both 2-4.5°C in one instance and then 1.5-4.5°C in the post-2013 assessment).

Our second departure from Pindyck's calibration is that we focus solely on uncertainty in the climate sensitivity parameter. Thus, rather than fit a further distribution around likely economic damages (another displaced gamma distribution in Pindyck's case), we employ instead point estimates for economic damages.[19] That simplification underplays the full impact of known uncertainties and therefore dampens the adverse consequences of temperature uncertainties. It, thus, biases our results toward less climate action. Meanwhile, it does not affect the qualitative implications of our results.

The main result reveals the clear tradeoff between mean and standard deviation, as highlighted by Pindyck (2014) and others. We zero in on the 2013 IPCC decision to extend the climate sensitivity range by lowering the lower bound from 2°C to 1.5°C. In our optimization problem, presented in detail in Freeman *et al.* (2015), that simple move increases WTP from 0.4% to 0.5% of GDP.

---

[18] See also Wagner and Weitzman (2015) for further explanation and calibration, presented here in Figure 1.

[19] Specifically, we set Pindyck's (2012) loss function parameter γ to 0.0001, slightly below the mean estimate used as input in Pindyck's displaced gamma distribution (Figure 3 in Pindyck, 2012).

The mean-variance discussion in this section is distilled from Freeman *et al.* (2015). That paper extends this discussion to a much broader set of questions around mean-variance tradeoffs, compares it to a mean-preserving spread, and tests the implications when various conditions are relaxed. It also addresses the question of the IPCC removing the 'most likely' or 'mean' value altogether. Whereas IPCC (2007) included a 'most likely' value of 3°C, IPCC (2013) did not. Interpreting that move as one that lowers the kurtosis of the distribution, as Freeman *et al.* (2015) shows, would also imply a higher WTP.

4.

Uncertainty around the seemingly all-important climate sensitivity addresses but one dimension. We haven't yet even considered the human and policy dimensions, which are themselves beset with major uncertainties. Deep and persistent uncertainty is a constant companion of policymaking for posterity (Summers and Zeckhauser, 2008).

In particular, climate policy also questions standard thinking around 'saving' for presumably wealthier future generations. The standard story goes something like this: Future generations will be much richer than we are, so why sacrifice portions of our measly incomes today when future generations will be much more easily able to pay larger sums? Such thinking would make sense if only wealth were involved, such as the question of how much we should invest to boost economic productivity, where most of the eventual productivity increases cannot be claimed by those who pay for them, and when those who receive most of the increases will be much richer.

Thinking around climate policy differs in three important respects. The first is irreversibility. Since temperatures and sea levels will rise for centuries due to actions (not) taken today, it becomes almost irrelevant to argue that future generations would have an easier time decreasing their CO_2 emissions because they will be richer. *Our* choices—not theirs—define *their* future. An analogy might be drawn to taking actions that yield pleasure today but promote damage to the genes we pass on to our descendants. No matter their wealth, or the care they take in protecting the genes they inherit, barring unforeseen scientific advances, they will have little ability to spend monies to reverse the damages that we impose.

Second, since future generations can indeed be expected to be richer, they would be willing and able to pay more for a stable climate. This is not merely a calculation of how much future generations would be willing to pay to avoid the worst. On a pure utilitarian basis, if climate and wealth are complements as we might expect, future generations will get more utility out of a superior climate due to their wealth (Summers and Zeckhauser, 2008). This observation, combined with elements of irreversibility, implies the need for more climate action today for the benefit of future generations, precisely because they will be richer (Sterner and Persson, 2008).

The third difference returns to the theme of persistent uncertainty: We often just don't know. Almost four decades of climate science have produced amazing advances in a host

of areas. But climate science has not enabled us to pin down the range of equilibrium climate sensitivity more finely. There are problems that lurk, ones that have received little or no notice to date.

The implications of that third uncertainty may be the most profound, especially because they are virtually impossible to quantify. The traditional focus has been on 'fat tails' and extreme events, an important concept in itself. Yet considering the 'likely' climate sensitivity values themselves can have similarly profound implications, with the conclusion that greater uncertainty often increases the case for climate action (Brock and Hansen, 2017; Freeman et al., 2015; Pindyck, 2014). It may, thus, not even be necessary to look to the low-probability, high-impact tail events. The not-so-low-probability, not-quite-as-high-impact events closer to the mean may carry even more punch. In either case, unknown and possibly unknowable areas lead us to our framing of the problem as deep and persistent uncertainty.

5.

Climate uncertainty comes in three flavors: stochastic uncertainty, measurement uncertainty, and model uncertainty (Brock and Hansen, 2017). The first is ever-present and hardly distinctive here. Measurement uncertainty is the simple fact that we know too little about fundamental parameters of the phenomena we do understand. That type of uncertainty has indeed diminished over time. We now know more about many climatic phenomena than we did four decades ago. Model uncertainty is the crux of the issue that concerns us. Four decades of impressive advances in climate science, due to the efforts of thousands of individuals, have gotten us no closer to pinning down the true value of climate sensitivity. If anything, the latest IPCC report takes a modest step back in that regard.

Extending the 'likely' range of climate sensitivity to include lower values, 1.5°C instead of 2°C as the bottom of the range, is at first glance *good* news. There is the potential that climate change is not as bad as has been feared. Sadly, increasing the 'likely' range while keeping the probability of what it means to be within that range constant at 66%—much like removing the concept of 'most likely' value of 3°C entirely—is *bad* news.

Despite important advances in other areas of climate science, we have discovered new uncertainties that make us even less confident about the range of equilibrium climate sensitivity than we were before the latest IPCC report was published. Given the increasing marginal costs of global warming, greater uncertainty, other factors equal, raises the returns from curbing greenhouse gases. On an expected-value basis, the massive uncertainties afflicting climate change, thus, *should* be a prod to policy action.

Bibliography

Arrhenius, S., 1896. On the influence of carbonic acid in the air upon the temperature of the ground. Lond. Edinb. Dublin Philos. Mag. J. Sci. 41, 237–276.

Arrow, K., Cropper, M., Gollier, C., Groom, B., Heal, G., Newell, R., Nordhaus, W., Pindyck, R., Pizer, W., Portney, P., Sterner, T., Tol, R.S.J., Weitzman, M., 2013. Determining Benefits and Costs for Future Generations. Science 341, 349–350. https://doi.org/10.1126/science.1235665

Arrow, K.J., Cropper, M.L., Gollier, C., Groom, B., Heal, G.M., Newell, R.G., Nordhaus, W.D., Pindyck, R.S., Pizer, W.A., Portney, P.R., Sterner, T., Tol, R.S.J., Weitzman, M.L., 2014. Should Governments Use a Declining Discount Rate in Project Analysis? Rev. Environ. Econ. Policy reu008. https://doi.org/10.1093/reep/reu008

Auffhammer, M., 2018. Quantifying Economic Damages from Climate Change. J. Econ. Perspect. 32, 33–52. https://doi.org/10.1257/jep.32.4.33

Brock, W.A., Hansen, L.P., 2017. Wrestling with Uncertainty in Climate Economic Models.

Brown, P.T., Stolpe, M.B., Caldeira, K., 2018. Assumptions for emergent constraints. Nature 563, E1. https://doi.org/10.1038/s41586-018-0638-5

Charney, J.G., Arakawa, A., Baker, D.J., Bolin, B., Dickinson, R.E., Goody, R.M., Leith, C.E., Stommel, H.M., Wunsch, C.I., 1979. Carbon dioxide and climate: a scientific assessment. National Academy of Sciences, Washington, DC.

Convery, F.J., Wagner, G., 2015. Reflections–Managing Uncertain Climates: Some Guidance for Policy Makers and Researchers. Rev. Environ. Econ. Policy 9, 304–320. https://doi.org/10.1093/reep/rev003

Cox, P.M., Huntingford, C., Williamson, M.S., 2018a. Emergent constraint on equilibrium climate sensitivity from global temperature variability. Nature 553, 319–322. https://doi.org/10.1038/nature25450

Cox, P.M., Williamson, M.S., Nijsse, F.J.M.M., Huntingford, C., 2018b. Cox et al. reply. Nature 563, E10. https://doi.org/10.1038/s41586-018-0641-x

Cropper, M.L., Freeman, M.C., Groom, B., Pizer, W.A., 2014. Declining Discount Rates. Am. Econ. Rev. 104, 538–543. https://doi.org/10.1257/aer.104.5.538

Freeman, M.C., Wagner, G., Zeckhauser, R.J., 2015. Climate sensitivity uncertainty: when is good news bad? Phil Trans R Soc A 373, 20150092. https://doi.org/10.1098/rsta.2015.0092

Gollier, C., Koundouri, P., Pantelidis, T., 2008. Declining discount rates: Economic justifications and implications for long-run policy. Econ. Policy 23, 758–795. https://doi.org/10.1111/j.1468-0327.2008.00211.x

Gollier, C., Weitzman, M.L., 2010. How should the distant future be discounted when discount rates are uncertain? Econ. Lett. 107, 350–353. https://doi.org/10.1016/j.econlet.2010.03.001

Heal, G., 2017. The Economics of the Climate. J. Econ. Lit. 55, 1046–1063. https://doi.org/10.1257/jel.20151335

Hogan, W.W., Wagner, G., 2018. Carbon prices, preferences, and the timing of uncertainty.

Houser, T., Hsiang, S., Kopp, R., Larsen, K., 2015. Economic Risks of Climate Change: An American Prospectus. Columbia University Press.

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D.J., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larsen, K., Houser, T., 2017. Estimating economic damage from climate change in the United States. Science 356, 1362–1369. https://doi.org/10.1126/science.aal4369

Hsiang, S., Kopp, R.E., 2018. An Economist's Guide to Climate Change Science. J. Econ. Perspect. 32, 3–32. https://doi.org/10.1257/jep.32.4.3

In come the waves, 2005. . The Economist.

IPCC, 2013. Fifth Assessment Report.

IPCC, 2007. Fourth Assessment Report.

IPCC, 2001. Third Assessment Report: Climate Change.

IPCC, 1995. Second Assessment Report: Climate Change.

IPCC, 1990. Report prepared for Intergovernmental Panel on Climate Change by Working Group I. Camb. U. K. N. Y. NY USA.

Keith, D.W., 2000. Geoengineering the climate: history and prospect. Annu. Rev. Energy Environ. 25, 245–284. https://doi.org/10.1146/annurev.energy.25.1.245

Keith, D.W., Wagner, G., Zabel, C.L., 2017. Solar geoengineering reduces atmospheric carbon burden. Nat. Clim. Change 7, 617–619. https://doi.org/10.1038/nclimate3376

Kelleher, J.P., Wagner, G., 2017. Ramsey Discounting Calls for Subtracting Climate Damages from Economic Growth Rates (SSRN Scholarly Paper No. ID 3003848). Social Science Research Network, Rochester, NY.

Knight, F.H., 1921. Risk, uncertainty and profit. N. Y. Hart Schaffner Marx.

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

Koonin, S.E., 2014. Climate Science Is Not Settled. Wall Str. J.

Kopp, R.E., Shwom, R., Wagner, G., Yuan, J., 2016. Tipping elements and climate-economic shocks: Pathways toward integrated assessment. Earths Future 4, 346–372. https://doi.org/10.1002/2016EF000362

Lindzen, R.S., 2009. The Climate Science Isn't Settled. Wall Str. J.

Morris, E., 2014. The Certainty of Donald Rumsfeld (Part 1). N. Y. Times.

National Academies of Sciences, 2018. Negative Emissions Technologies and Reliable Sequestration: A Research Agenda. https://doi.org/10.17226/25259

National Academy of Sciences, 2017. Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide.

National Research Council, 2015. Climate Intervention: Reflecting Sunlight to Cool Earth. National Academies Press.

Oreskes, N., Conway, E.M., 2011. Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming. Bloomsbury Publishing USA.

Pindyck, R.S., 2017. Taxes, Targets, and the Social Cost of Carbon. Economica.

Pindyck, R.S., 2014. Risk and Return in the Design of Environmental Policy. J. Assoc. Environ. Resour. Econ. 1, 395–418. https://doi.org/10.1086/677949

Pindyck, R.S., 2013. The Climate Policy Dilemma. Rev. Environ. Econ. Policy 7, 219–237.

Pindyck, R.S., 2012. Uncertain outcomes and climate change policy. J. Environ. Econ. Manag. 63, 289–303.

Po-Chedley, S., Proistosescu, C., Armour, K.C., Santer, B.D., 2018. Climate constraint reflects forced signal. Nature 563, E6. https://doi.org/10.1038/s41586-018-0640-y

Proistosescu, C., Huybers, P.J., 2017. Slow climate mode reconciles historical and model-based estimates of climate sensitivity. Sci. Adv. 3, e1602821. https://doi.org/10.1126/sciadv.1602821

Revesz, R.L., Howard, P.H., Arrow, K., Goulder, L.H., Kopp, R.E., Livermore, M.A., Oppenheimer, M., Sterner, T., 2014. Global warming: Improve economic models of climate change. Nature 508, 173–175.

Roe, G.H., Bauman, Y., 2012. Climate sensitivity: should the climate tail wag the policy dog? Clim. Change 117, 647–662. https://doi.org/10.1007/s10584-012-0582-6

Rumsfeld, D., 2002. Press Conference at NATO Headquarters.

Rypdal, M., Fredriksen, H.-B., Rypdal, K., Steene, R.J., 2018. Emergent constraints on climate sensitivity. Nature 563, E4. https://doi.org/10.1038/s41586-018-0639-4

Stern, N.H., 2006. Stern Review: The economics of climate change. HM treasury London.

Sterner, T., Persson, U.M., 2008. An Even Sterner Review: Introducing Relative Prices into the Discounting Debate. Rev. Environ. Econ. Policy 2, 61–76. https://doi.org/10.1093/reep/rem024

Summers, L., Zeckhauser, R., 2008. Policymaking for posterity. J. Risk Uncertain. 37, 115–140. https://doi.org/10.1007/s11166-008-9052-y

Supran, G., Oreskes, N., 2017. Assessing ExxonMobil's climate change communications (1977–2014). Environ. Res. Lett. 12, 084019. https://doi.org/10.1088/1748-9326/aa815f

Tol, R.S., 2009. The economic effects of climate change. J. Econ. Perspect. 23, 29–51.

U.S. Government Interagency Working Group on Social Cost of Carbon, 2015. Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.

U.S. Government Interagency Working Group on Social Cost of Carbon, 2010. Appendix 15a. Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866, in: Final Rule Technical Support Document (TSD): Energy Efficiency Program for Commercial and Industrial Equipment: Small Electric Motors. U.S. Department of Energy.

Wagner, G., Weitzman, M.L., 2018. Potentially large equilibrium climate sensitivity tail uncertainty. Econ. Lett. 168, 144–146. https://doi.org/10.1016/j.econlet.2018.04.036

Wagner, G., Weitzman, M.L., 2015. Climate Shock. Princeton University Press.

Wagner, G., Zeckhauser, R.J., 2011. Climate policy: hard problem, soft thinking. Clim. Change 110, 507–521. https://doi.org/10.1007/s10584-011-0067-z

Weitzman, M.L., 2015. Book Review—A Review of William Nordhaus' The Climate Casino: Risk, Uncertainty, and Economics for a Warming World. Rev. Environ. Econ. Policy 9, 145–156.

Weitzman, M.L., 2014. Fat tails and the social cost of carbon. Am. Econ. Rev. 104, 544–546.

Weitzman, M.L., 2012. GHG Targets as Insurance Against Catastrophic Climate Damages. J. Public Econ. Theory 14, 221–244. https://doi.org/10.1111/j.1467-9779.2011.01539.x

Weitzman, M.L., 2011. Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change. Rev. Environ. Econ. Policy 5, 275–292. https://doi.org/10.1093/reep/rer006

Weitzman, M.L., 2009. On Modeling and Interpreting the Economics of Catastrophic Climate Change. Rev. Econ. Stat. 91, 1–19. https://doi.org/10.1162/rest.91.1.1

Zeckhauser, R.J., 2006. Investing in the Unknown and Unknowable. Capital. Soc. 1/2.