Symposium:
Fat Tails and the Economics of Climate Change

Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change

Martin L. Weitzman*

Introduction

I believe that the most striking feature of the economics of climate change is that its extreme downside is nonnegligible. Deep structural uncertainty about the unknown unknowns of what might go very wrong is coupled with essentially unlimited downside liability on possible planetary damages. This is a recipe for producing what are called "fat tails" in the extremes of critical probability distributions. There is a race being run in the extreme tail between how rapidly probabilities are declining and how rapidly damages are increasing. Who wins this race, and by how much, depends on how fat (with probability mass) the extreme tails are. It is difficult to judge how fat the tail of catastrophic climate change might be because it represents events that are very far outside the realm of ordinary experience.

In this article, which is part of a symposium on Fat Tails and the Economics of Climate Change, I address some criticisms that have been leveled at previous work of mine on fat tails and the so-called "dismal theorem."[1] At first, I was inclined to debate some of the critics and their criticisms more directly. But, on second thought, I found myself anxious not to be drawn into being too defensive and having the main focus be on technical details. Instead, I am more keen here to emphasize anew and in fresh language the substantive concepts that, I think, may be more obscured than enlightened by a debate centered on technicalities. I am far more committed to the simple basic ideas that underlie my approach to fat-tailed uncertainty and the economics of catastrophic climate change than I am to the particular

---

*Department of Economics, Harvard University; e-mail: mweitzman@harvard.edu

Without blaming them for the remaining deficiencies in this article, I am extremely grateful for the constructive comments on an earlier version by James Annan, Daniel Cole, Stephen DeCanio, Baruch Fischoff, Don Fullerton, John Harte, William Hogan, Matthew Kahn, David Kelly, Michael Oppenheimer, Robert Pindyck, Joseph Romm, and Richard Tol.

[1] This symposium also includes articles by Nordhaus (2011) and Pindyck (2011). The "dismal theorem," introduced in Weitzman (2009a), will be discussed later in this article.

Review of Environmental Economics and Policy, volume 5, issue 2, summer 2011, pp. 275–292
doi:10.1093/reep/rer006

mathematical form in which I have chosen to express them. These core concepts could have been wrapped in a variety of alternative mathematical shells—and the particular one that I chose to use previously is somewhat arbitrary. The implications are roughly similar, irrespective of formalization. Some technical details are unavoidable, but if I can give the underlying concepts greater intuitive plausibility, then I believe that this set of ideas will become more self-evident and more self-evidently resistant to several of the criticisms that have been leveled against it.

In the next section, I present an intuitive–empirical argument that deep structural uncertainty lies at the heart of climate change economics. Then, in the following section, I try to explain some of the theory behind fat-tailed extreme events and discuss some possible implications for the analysis of climate change. I offer a few summary remarks in the final section.

Some Empirical Examples of Deep Structural Uncertainty about Climate Extremes

In this section I try to make a heuristic–empirical case (in the form of five "stylized facts") for there being big structural uncertainties in the economics of extreme climate change. I will argue on intuitive grounds that the way in which this deep uncertainty is conceptualized and formalized should influence the outcomes of any reasonable benefit–cost analysis (BCA) of climate change. Furthermore, I will argue that the seeming immunity of the "standard" BCA to the possibility of extreme outcomes is both peculiar and disturbing. My arguments in this section are not intended to be airtight or rigorous. Rather, this is an intuitive presentation based on some very rough stylized facts.

By BCA of climate change, I mean, in the widest sense, some overall economic analysis centered on maximizing (or at least comparing) welfare. My notion of BCA in the present context is so broad that it overlaps with an integrated assessment model (IAM), and here I treat the two as essentially interchangeable. I begin by setting up a straw man that I will label the "standard BCA of climate change." Of course, there is no "standard BCA of climate change," but I think this is an allowable simplification for purposes of exposition here.

We all know that computer-driven simulations are dependent upon the core assumptions of the underlying model. The intuitive examples presented below are frankly aimed at sowing a few seeds of doubt that the "standard BCA of climate change" fairly represents structural uncertainties about extreme events, and that therefore its conclusions might be less robust than is commonly acknowledged. I argue not that the standard model is wrong or even implausible, but rather that it may not be robust with respect to the modeling of catastrophic outcomes. I will try to make my case by citing five aspects of the climate science and economics that do not seem to me to be adequately captured by the standard BCA. The five examples—which I call Exhibits 1, 2, 3, 4 and 5—are limited to structural uncertainty concerning the modeling of climate disasters. While other important aspects of structural uncertainty might also be cited, I restrict my stylized facts to these five examples. In the spirit of performing a kind of "stress test" on the standard BCA, I naturally concentrate on things that might go wrong rather than things that might go right.

Exhibit 1: Past and Present Greenhouse Gas Concentrations

Exhibit 1 concerns the atmospheric level of greenhouse gases (GHGs) over time. Ice-core drilling in Antarctica began in the late 1970s and is still ongoing. The record of carbon dioxide (CO_2) and methane (CH_4) trapped in tiny ice-core bubbles currently spans 800,000 years.[2] It is important to recognize that the numbers in this unparalleled 800,000-year record of GHG levels are among the very best data that exist in the science of paleoclimate. Almost all other data (including past temperatures) are inferred indirectly from proxy variables, whereas these ice-core GHG data are directly observed.

The preindustrial revolution level of atmospheric CO_2 (about two centuries ago) was 280 parts per million (ppm). The ice-core data show that CO_2 varied gradually during the last 800,000 years within a relatively narrow range roughly between 180 and 280 ppm and has never been above 300 ppm. Currently, CO_2 is over 390 ppm and climbing steeply. In 800,000 years, methane has never been higher than 750 parts per billion (ppb), but now this extremely potent GHG, which is twenty-two times more powerful than CO_2 (per century), is at about 1,800 ppb. The sum total of all carbon dioxide equivalent (CO_2e) GHGs is currently at about 450 ppm. An even more startling contrast with the 800,000-year record is the rate of change of GHGs: increases in CO_2 were below (and typically well below) 25 ppm within any past subperiod of 1,000 years, while now CO_2 has risen by over 25 ppm in just the last ten years. Thus, anthropogenic activity has very rapidly elevated atmospheric CO_2 and CH_4 to levels very far outside their natural range. The unprecedented scale and speed of GHG increases brings us into uncharted territory and makes predictions of future climate change very uncertain. Looking ahead a century or two, the levels of atmospheric GHGs that may ultimately be attained (unless decisive measures are undertaken) have likely not existed for tens of millions of years, and the speed of this change may be unique on a time scale of hundreds of millions of years.

Remarkably, the "standard BCA of climate change" takes little account of the magnitude of the uncertainties involved in extrapolating future climate change so far beyond past experience. Perhaps even more surprising, the gradual tightening of GHG emissions, which emerges as optimal policy from the "standard" BCA, typically attains stabilization at levels of CO_2 that approach 700 ppm (and levels of CO_2e that are even higher). The "standard" BCA thus recommends subjecting the Earth's system to the unprecedented shock of instantaneously (in geological terms) jolting atmospheric stocks of GHGs up to two-and-a-half times above their highest level over the past 800,000 years—without mentioning the unprecedented nature of this unique planetary experiment. This is my Exhibit 1.

Exhibit 2: The Uncertainty of the Climate Change Response

Exhibit 2 concerns the highly uncertain climate change response to the kind of unprecedented increases in GHGs that were described in Exhibit 1. For specificity, I focus on the uncertainty of so-called "equilibrium climate sensitivity," which is a key macro-indicator of the *eventual* temperature response to GHG changes. This is a good example of a "known unknown."

---

[2] My numbers are taken from Dieter et al. (2008) and supplemented by data from the Keeling curve for more recent times (available online at: ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2_mm_mlo.txt).

However, it should be understood that under the rubric of climate sensitivity, I am trying to lump together an entire suite of other uncertainties, including some nonnegligible unknown unknowns. The insights and results of this Exhibit 2 are not intended to stand or fall on the single narrow issue of accurately modeling the effects of uncertain climate sensitivity. Rather, equilibrium climate sensitivity is to be understood here as a prototype example of uncertainty, or a metaphor, which is being used to illustrate much more generic issues concerning the economics of uncertain climate change.

The Intergovernmental Panel on Climate Change (IPCC-AR4 2007) defines equilibrium climate sensitivity (hereafter S) this way: "The equilibrium climate sensitivity is a measure of the climate system response to sustained radiative forcing. It is not a projection but is defined as the global average surface warming following a doubling of carbon dioxide concentrations. It is *likely* to be in the range 2–4.5°C with a best estimate of 3°C, and is *very unlikely* to be less than 1.5°C. Values substantially higher than 4.5°C cannot be excluded, but agreement of models with observations is not as good for those values." The actual empirical reason for why the upper tails of these probability distributions are long and seem heavy with probability dovetails with Bayesian theory: inductive knowledge is always useful, of course, but simultaneously it is limited in what it can tell us about extreme events outside the range of experience. In such situations, one is forced back onto depending more than one might wish upon prior beliefs in the form of a prior probability distribution, which, of necessity, is largely subjective and relatively diffuse.

Thus, any curve-fitting exercise attempting to attribute probabilities to $S \geq 4.5°\text{C}$, such as I am doing here, is little more than conjectural speculation. My purpose is merely to show that critical results can depend on seemingly casual decisions about how to model tail probabilities. To illustrate some striking implications for the analysis of climate change, I contrast the use of two familiar probability distributions to represent the upper-half tail of climate sensitivity above the median: (a) the Pareto (or Power or Polynomial) distribution, subscripted **P**, which is the prototype example of a fat upper tail; and (b) the Normal distribution, subscripted **N**, which is the prototype of a thin upper tail. (By common definition, probabilities that decline exponentially or faster, like the Normal distribution, are thin tailed, while probabilities that decline polynomially or slower, like the Pareto distribution, are fat tailed). The IPCC defines "likely" as a probability above 66 percent but below 90 percent, which would mean that the probability that climate sensitivity is greater than 4.5°C ($\text{Prob}[S \geq 4.5°\text{C}]$) is between 5 percent and 17 percent. A more recent average estimate of fourteen leading climate scientists is $\text{Prob}[S \geq 4.5°\text{C}] = 23$ percent (Zickfeld et al. 2010). Here I choose $\text{Prob}[S \geq 4.5°\,\text{C}] = 15$ percent and I calibrate parameters of both probability distributions so that $\text{Prob}[S \geq 3°\text{C}] = 0.5$, and $\text{Prob}[S \geq 4.5°\text{C}] = 0.15$.[3]

Table 1 presents cumulative probabilities of climate sensitivity for the Pareto and Normal distributions. Table 1 indicates a tremendous difference in upper tail behavior between the fat-tailed Pareto distribution and the thin-tailed Normal distribution. I think that the Pareto distribution of climate sensitivity has a disturbingly large amount of probability ($\text{Prob}_\mathbf{P}$) in its upper tail. There is no consensus on how to aggregate the results of many different climate

---

[3] I lean more toward $\text{Prob}[S \geq 4.5°\text{C}] \leq 17\%$ than toward $\text{Prob}[S \geq 4.5°\text{C}] \geq 5\%$ because, for a time horizon of a century and a half or so, it is plausible that the more inclusive and larger "earth system sensitivity" (which includes slow feedbacks like methane releases) matters more than the "fast-feedback equilibrium sensitivity" that IPCC-AR4 refers to. For more on the distinction between fast- and slow-feedback climate sensitivities, see, for example, Hansen et al. (2008).

Table 1 Prob$[S \geq \hat{S}]$ for fat-tailed Pareto and thin-tailed Normal distributions

$\hat{S}$ =	**3°C**	**4.5°C**	**6°C**	**8°C**	**10°C**	**12°C**
$\text{Prob}_{\mathbf{P}}[S \geq \hat{S}]$	0.5	0.15	0.06	0.027	0.014	0.008
$\text{Prob}_{\mathbf{N}}[S \geq \hat{S}]$	0.5	0.15	0.02	0.003	7×10^{-7}	3×10^{-10}

sensitivity studies into one overarching probability density function (PDF), and there is much controversy about how it might be done. But for what it is worth (perhaps very little), the median upper 5 percent probability level over all 22 climate sensitivity PDFs cited in IPCC-AR4 is 6.4°C, which fits with the Pareto PDF in Table 1 above.[4]

Table 2 presents some values of probabilities of eventual increased global mean surface temperatures (T) as a function of stationary greenhouse gas concentrations (G).[5] The first row of Table 2 represents levels of G (measured in ppm of CO_2e). The second row indicates the median equilibrium temperature as a function of stabilized GHG stocks. The remaining rows indicate the probabilities ($\text{Prob}_{\mathbf{P}}$ and $\text{Prob}_{\mathbf{N}}$) of achieving at least the steady-state temperature increase represented by the entries in the table (5°C and 10°C) for both of the chosen PDFs (Pareto = **P** = fat tail; Normal = **N** = thin tail).

What is especially striking to me about Table 2 is the reactiveness of high-temperature probabilities to the level of GHGs. The fat-tailed Pareto case seems especially worrisome. One implication is that an optimal policy might be expected to keep GHG levels down and be much less casual about letting CO_2e levels exceed 700 ppm than the "standard" BCA. I believe that Table 2 could be taken as some indirect evidence that the main purpose of keeping GHG concentrations down is effectively to buy insurance against catastrophic global warming. The above examples of the highly uncertain eventual temperature response to unprecedented increases in GHGs constitute my Exhibit 2.

Exhibit 3: A Physical Basis for Catastrophic Outcomes

Exhibit 3 concerns possibly disastrous releases over the long run of bad feedback components of the carbon cycle that are currently omitted from most general circulation models. The chief concern here is that there may be a significant supplementary component that conceptually

Table 2 Probabilities of exceeding $T = 5°\text{C}$ and $T = 10°\text{C}$ for given G = ppm of CO_2e

***G*:**	**400**	**500**	**600**	**700**	**800**	**900**
Median T	1.5°	2.5°	3.3°	4.0°	4.5°	5.1°
$\text{Prob}_{\mathbf{P}}[T \geq 5°\text{C}]$	1.5%	6.5%	15%	25%	38%	52%
$\text{Prob}_{\mathbf{N}}[T \geq 5°\text{C}]$	10^{-6}	2.0%	14%	29%	42%	51%
$\text{Prob}_{\mathbf{P}}[T \geq 10°\text{C}]$	0.20%	0.83%	1.9%	3.2%	4.8%	6.6%
$\text{Prob}_{\mathbf{N}}[T \geq 10°\text{C}]$	10^{-30}	10^{-10}	10^{-5}	0.1%	0.64%	2.1%

---

[4] For details, see Weitzman (2009a).

[5] Table 1 indicates probabilities for climate sensitivity, which correspond to GHG levels of 560 ppm of CO_2e (i.e., a doubling of the preindustrial revolution level of 280 ppm). The different values of GHG concentrations in Table 2 reflect the correspondingly different probabilities of temperatures responses, in proportion to the logarithm of GHG concentrations relative to 280 ppm.

should be added on to the so-called "fast feedback" equilibrium concept that IPCC-AR4 works with. This omitted component (which would be part of a more inclusive slow-feedback generalization called "earth system sensitivity") includes the powerful self-amplification potential of greenhouse warming that is due to heat-induced releases of sequestered carbon. One vivid example is the huge volume of GHGs currently trapped in tundra permafrost and other boggy soils (mostly as methane, a particularly potent GHG). A more remote (but even more vivid) possibility, which in principle should also be included, is heat-induced releases of the even vaster offshore deposits of methane trapped in the form of clathrates.[6]

There is a very small and unknown (but decidedly nonzero) probability over the long run of having destabilized methane from these offshore clathrate deposits seep into the atmosphere if the temperature of the waters bathing the continental shelves increases just slightly. The amount of methane involved is huge, although it is not precisely known. Most estimates place the carbon-equivalent content of methane hydrate deposits at about the same order of magnitude as all other fossil fuels combined. Over the long run, a methane outgassing–amplifier process could potentially precipitate a disastrous strong positive feedback warming. If it occurred at all, such an event would likely take centuries to materialize because the presumed initiator would be the slow-acting gradual warming of ocean waters at the depths of the continental shelves. Thus, while it is a low-probability event that might only transpire centuries from now (if at all), the possibility of a climate meltdown is not just the outcome of a mathematical theory but has a real physical basis. Other examples of an actual physical basis for catastrophic outcomes could be cited, but this one will do here. This is my Exhibit 3.

Exhibit 4: Damages of Extreme Climate Change

Exhibit 4 concerns what I view as a somewhat cavalier treatment in the literature of damages or disutilities from extreme temperature changes. The "standard" BCA damages function reduces welfare-equivalent output at mean global temperature change T by a quadratic-polynomial *multiplier* of the form $M(T) = \alpha T^2/(1 + \alpha T^2)$.[7] The results in terms of fractional damages to output are indicated by $M(T)$ in Table 3.

I do not find the numbers in Table 3 convincing for higher temperatures. At an extraordinarily high global average temperature change of $T = 10°\mathrm{C}$, the welfare-equivalent damage in Table 3 is a loss of "only" 19 percent of world output *at the time of impact*. If the annual growth

Table 3 Multiplicative-quadratic damages $M(T)$ (as fraction of output)

T	2°C	4°C	6°C	8°C	10°C	12°C
$M(T)$	1%	4%	8%	13%	19%	26%

---

[6] Clathrates (or hydrates) are methane molecules that are boxed into a semistable state by being surrounded by water molecules under high pressure and low temperatures. For more about methane clathrates, see Archer (2007) and the recent article by Shakhova et al. (2010).

[7] For the sake of concreteness, I take the value of the parameter α used in the latest version of Nordhaus's well-known DICE model (Nordhaus 2008). The DICE model is perhaps the most famous IAM in the economics of climate change. The value $\alpha = 0.002388$ was used to generate his figure 3-3 on page 51. I hasten to add in fairness that this specification was intended only to capture low temperature damages and was never intended to be extrapolated to very high temperature changes.

rate is, say, 2 percent and the time of impact is, say, two centuries from now, then the welfare difference between no temperature change in two hundred years and a temperature change of 10°C in two hundred years is the difference between output that is fifty-five times larger than current levels and output that is forty-four times larger than current levels. This is equivalent to a reduction in projected annual growth rates from 2 percent to 1.9 percent. Such a mild kind of damages function is practically preordained to make extreme climate change look empirically negligible, almost no matter what else is assumed. Conversely, it turns out that fat-tailed temperature PDFs are not by themselves sufficient to make extreme climate change have empirical "bite" without a damages function that is far more immiserizing at high-temperature changes.

So what should the damages function be for very high temperatures? No one knows, of course. Taking an extreme example, suppose for the sake of argument that average global temperatures were to increase by the extraordinary amount of 10°C (with a low probability, and occurring centuries hence). While it is true that people live very well in places where the mean temperature is 10°C higher than in Yakutsk, Siberia, I do not think that this type of analogy justifies using a comparative geography approach for estimating welfare equivalent damages from an average *global* temperature change of 10°C. Global mean temperatures involve a double averaging: across space and over time. A "damages function" is a reduced form representing global welfare losses from global average temperatures, which subsumes a staggering amount of regional, seasonal, and even daily weather heterogeneity. Regional and seasonal climate changes are presumably much more unpredictable than global average surface temperatures. There is just too much structural uncertainty and too much heterogeneity to put trustworthy bounds on the unprecedented, almost unimaginable, changes to planetary welfare that would result from average global temperatures increasing by 10°C. When there is such great uncertainty about catastrophic damages, and when the damages function for high-temperature changes is so conjectural, the relevant degree of risk aversion, yet another important unknown here, will tend to play a significant role in an economic analysis of climate change.

Of course, I have no objective way to determine the magnitudes of high-temperature damages. The last time the world witnessed periods where global average temperatures were very roughly 10°C or so above the present was during the Eocene epoch ≈55–34 million years ago. During these warming periods, the earth was ice free, while palm trees and alligators lived near the North Pole. The Eocene was also the last epoch in which there were geologically rapid increases in mean global temperatures of magnitude ≈5°C or so above an already warm background. Such hyperthermal events occurred over an average period of very roughly ≈100,000 years or so, which is extremely gradual compared to current worst-case anthropogenically induced trajectories. It is unknown what exactly triggered these Eocene temperature increases, but they were accompanied by equally striking atmospheric carbon increases. One likely culprit is the strong feedback release of large amounts of methane hydrates from clathrate deposits (Exhibit 3), which is a nonnegligible possibility over the next century or two if current GHG emissions trends are extrapolated. The major point here is that relatively rapid changes in global average temperatures of ≈5°C above present values are extremely rare events and are extraordinarily far outside the scope of human experience. For huge temperature increases such as $T \approx 10$°C, the planetary effects are even more difficult to imagine. To find a geologically instantaneous increase in average global temperatures of magnitude $T \approx 10$°C, one would have to go back hundreds of millions of years.

For me, 10°C offers both a vivid image and a reference point, especially in light of a recent study, which estimated that global average temperature increases of $\approx$11–12°C (with, importantly, accompanying humidity in the form of a high wet-bulb temperature[8]) would exceed an absolute thermodynamic limit to metabolic heat dissipation (Sherwood and Huber 2010). Beyond this threshold, represented by a wet-bulb temperature of 35°C, more than half of today's human population would be living in places where, at least once a year, there would be periods when death from heat stress would ensue after about six hours of exposure. (By contrast, the highest wet-bulb temperature anywhere on Earth today is about 30°C). Sherwood and Huber (2010) further emphasize: "This likely overestimates what could practically be tolerated: Our [absolute thermodynamic] limit applies to a person out of the sun, in a gale-force wind, doused with water, wearing no clothing and not working." Even at wet-bulb temperatures, much lower than 35°C, human life would become debilitating and physical labor would be unthinkable. The massive unrest and uncontainable pressures this might bring to bear on the world's human population are almost unimaginable. The Earth's ecology, whose valuation is another big uncertainty, would be upended. Thus, a temperature change of $\approx$10°C would appear to represent an extreme threat to human civilization and global ecology as we now know it, even if it might not necessarily mean the end of *Homo sapiens* as a species.

It must be emphasized strongly that very high atmospheric temperature changes such as $T = 10$°C would likely take several centuries to attain. The higher the limiting temperature, the longer it takes to achieve equilibrium because the oceans will first have to absorb the enormous amounts of heat being generated. Alas, if the oceans are building up enormous amounts of heat it could set in motion irreversible long-term methane clathrate releases from the continental shelves along with some other nasty surprises. Thus, overall damages generated by equilibrium $T = 10$°C are best conceptualized as associated with *being on the trajectory* whose asymptotic limiting atmospheric temperature change is $T = 10$°C.

As noted above, the "standard" BCA damages function reduces welfare-equivalent consumption by a quadratic-polynomial multiplier. This essentially describes a single-attribute utility function, or, equivalently, a multiattribute utility function with strong substitutability between the two attributes of consumption and temperature change. This would be an appropriate formulation if the main impact of climate change is, say, to drive up the price of food and increase the demand for air conditioning. This particular choice of functional form allows the economy to easily substitute higher consumption for higher temperatures, since the limiting elasticity of substitution between consumption and higher temperatures is one (due to the multiplicative-polynomial assumption). However, very different optimal policies can be produced when other functional forms are used to express the disutility of disastrously high temperatures. For example, suppose that instead of being *multiplicatively separable* (as in the "standard" BCA), the disutility of temperature change is instead *additively separable*. This means that welfare is the analogous additively separable *arithmetic difference* between a utility function of consumption and a quadratic loss function of temperature changes. This specification amounts to postulating a genuine

---

[8] Wet-bulb temperature is essentially measured by a thermometer whose bulb is encased in a wet sock being cooled by a very strong wind. Above a wet-bulb temperature of 35°C, humans cannot shed enough core body heat to live—even under ideal circumstances.

multiattribute utility function that describes a situation where the main impact of climate change is on things that are not readily substitutable with material wealth, such as biodiversity and health. If the utility function has a constant coefficient of relative risk aversion of two, it implies an elasticity of substitution between consumption and temperature change of one half. Empirically, using this additive form—even without *any* uncertainty concerning temperatures—prescribes a significantly more stringent curtailment of GHG emissions than what emerges from the analogous multiplicative form of the "standard" BCA.[9] Allowing for the possibility of high temperatures (even with low probabilities) would presumably exaggerate this difference between the additive and multiplicative functional forms yet more. Such fragility to basic functional forms is disturbing because we cannot know with confidence which specification of the utility function is more appropriate.

The above discussions and examples of nonrobustness with respect to a damages function for high temperatures constitutes my Exhibit 4.

Exhibit 5: Discounting the Distant Future

Exhibit 5 concerns the notorious issue of how to discount the distant future. The effects of global warming and climate change will be spread out over centuries and even millennia from now. The logic of compounding a constant positive interest rate forces us to say that what one might conceptualize as monumental—even earth-shaking—events, such as disastrous climate change, do not much matter when they occur in the deep future. Perhaps even more disturbing, when exponential discounting is extended over very long time periods there is a truly extraordinary dependence of BCA on the choice of a discount rate. Seemingly insignificant differences in discount rates can make an enormous difference in the present discounted value of distant future payoffs. In many long-run situations, almost any answer to a BCA question can be defended by one or another particular choice of a discount rate. This is true in general, but it is an especially acute problem when distant future events like climate change (especially catastrophic climate change) are being discounted.

There is a high degree of uncertainty about what should be taken as the appropriate real rate of return on capital in the long run, accompanied by much controversy about its implications for long-run discounting. There is no deep reason or principle that allows us to extrapolate past rates of return on capital into the distant future. The industrial revolution itself began some two centuries ago in Britain and only slowly thereafter permeated throughout the world. The seeming trendlessness of some past rates of return is a purely empirical reduced form observation, which is not based on any underlying theory that would allow us to

---

[9] With a coefficient of relative risk aversion of two, the above additively separable specification is mathematically equivalent to the constant elasticity of substitution (CES) specification of Sterner and Persson (2008) with CES one half. In their pioneering study, Sterner and Persson showed empirically—by plugging it into Nordhaus's deterministic DICE model—that their CES (or, equivalently, my additive) welfare specification prescribes a significantly more aggressive policy response to global warming (with a significantly higher carbon tax) than the analogous multiplicative specification of the "standard" BCA. The disturbingly significant distinction between utility functions that are multiplicatively separable in consumption and temperature change and those that are additively separable in consumption and temperature change is explored in some detail in Weitzman (2009b).

confidently project past numbers far into the future. There are a great many fundamental factors that cannot easily be extrapolated, just one of which is the unknown future rate of technological progress. Even leaving aside the question of how to project future interest rates, additional issues for climate change involve which interest rate to choose out of a multitude of different rates of return that exist in the real world.[10] Furthermore, there is a strong normative element having to do with what is the "right" rate, which brings an ethical dimension to discounting climate change across many future generations that is difficult to evaluate and incorporate into standard BCA. This normative issue is further complicated when the event affecting future generations is a low-probability, high-impact catastrophic outcome. If the utility function is additively separable in damages and consumption, then the appropriate interest rate for discounting damages is not the rate of return on capital but rather the rate of pure time preference, which for a normative evaluation of climate change is arguably near zero.

The constant interest rates used for discounting in the "standard" BCA would be viewed by many people as severely biasing BCA toward minimizing into near nothingness the present discounted value of distant future events, such as climate change. This kind of exponential discounting, perhaps more than anything else, makes scientists and the general public suspicious of the economist's "standard" BCA of climate change, since it trivializes even truly enormous distant future impacts. Honestly, I think that there are few economists who do not feel uneasy about evaluating distant future climate change impacts this way. One line of research, in which I have been involved, shows that when the discount rate itself is uncertain, it implies that the "effective" discount rate declines over time to its lowest possible value (see Weitzman 1998). Empirically, this effect can be quite powerful (see Weitzman 2010). The driving force is a "fear factor" that derives from risk aversion to permanent productivity shocks representing bad future states of the world. Whatever its source, the unknown discount rate (and the extraordinary sensitivity of policy to its choice) is yet another big structural uncertainty in the economic analysis of climate change, especially when evaluating possible catastrophes. This is my Exhibit 5.

A Long Chain of Structural Uncertainties

The above five exhibits could readily be extended to incorporate yet more examples of structural uncertainty, but enough is enough. To summarize, the economics of climate change consists of a very long chain of tenuous inferences fraught with big uncertainties in every link: beginning with unknown base-case GHG emissions; compounded by big uncertainties about how available policies and policy levers will affect actual GHG emissions; compounded by big uncertainties about how GHG flow emissions accumulate via the carbon cycle into GHG stock concentrations; compounded by big uncertainties about how and when GHG stock concentrations translate into global average temperature changes; compounded by big uncertainties about how global average temperature changes decompose into specific changes in regional weather patterns; compounded by big uncertainties about how adaptations to, and mitigations of, climate change damages at a regional level are translated into regional utility changes via an appropriate "damages function"; compounded by big

---

[10] For more on this, see Weitzman (2007).

uncertainties about how future regional utility changes are aggregated into a worldwide utility function and what its overall degree of risk aversion should be; compounded by big uncertainties about what discount rate should be used to convert everything into expected present discounted values. The result of this lengthy cascading of big uncertainties is a reduced form of truly extraordinary uncertainty about the aggregate welfare impacts of catastrophic climate change, which is represented mathematically by a PDF that is spread out and heavy with probability in the tails.

What I would wish the reader to take away from these five exhibits is the notion that the seeming immunity of the "standard" BCA to such stylized facts is both peculiar and disturbing. An unprecedented and uncontrolled experiment is being performed by subjecting planet Earth to the shock of a geologically instantaneous injection of massive amounts of GHGs. Yet, the standard BCA seems almost impervious to the extraordinarily uncertain probabilities and consequences of catastrophic climate change. In light of the above five exhibits of structural uncertainty, the reader should feel intuitively that it goes against common sense that a climate change BCA does *not* much depend upon how potential future disasters are modeled and incorporated into the BCA. This uneasy intuitive feeling based on stylized empirical facts is my opening argument. I turn next to the theory.

The Dismal Theorem, Infinity, and BCA: A Theoretical Framework

I begin this section by asking why it is relevant in the first place to have any supporting theory if the five stylized fact exhibits from the previous section are convincing. Why aren't these stylized facts alone sufficient evidence that there is a problem with the "standard BCA of climate change"? My answer is that a combined theoretical plus empirical–intuitive argument delivers a particularly powerful one–two punch at the treatment of structural uncertainty in the standard BCA. In this respect, I believe that the whole of my argument is bigger than the sum of its two parts. The theoretical part reinforces the empirical part by placing it within a formal mathematical framework. When the intuitive exhibits are seen as reflecting some formalized theoretical structure, then it becomes less easy to brush them aside as mere sniping at an established model. In this theoretical section, as in the empirical section above, I emphasize the intuitive plausibility of the case I am trying to make—focusing here on the underlying logic that drives the theory.

In the previous section I argued that it is only common sense that climate change policy implications should depend on the treatment of low-probability, extreme-impact outcomes. The main question I attempted to address in Weitzman (2009a) was whether such intuitive dependence is reflecting some deeper principle. My answer was that there is indeed a basic underlying theoretical principle (the "dismal theorem," hereafter DT) that points in this direction. The simple logic can be grasped intuitively without using (or understanding) the fancy math required to state and prove a formal version of the principle. In this section I restate the theoretical arguments underlying the DT in what is hopefully a more intuitive form.

Let welfare W stand for expected present discounted utility, whose theoretical upper bound is B. Let $D \equiv B - W$ be expected present discounted *dis* utility. Here D stands for what might be called the "diswelfare" of climate change. Assume for the sake of argument that D is "essentially" unbounded in the particular case of climate change because global liability is "essentially" unlimited in a worst-case scenario. (More later on what happens when D is, technically, bounded.) Because the integral over a nonnegative probability measure is one, the PDF of $\ln D$ must decline to zero. In other words, extreme outcomes can happen, but their likelihood diminishes to zero as a function of how extreme the outcome might be. The idea that extreme outcomes cannot be eliminated altogether but are hypothetically possible with some positive probability is not at all unique to climate change. Almost nothing in our world has a probability of exactly zero or exactly one. What is worrisome is not the fact that the upper tail of the PDF of $\ln D$ is *long* (reflecting the fact that a meaningful bound on diswelfare does not exist), but that it might be *fat* (reflecting the fact that the probability of a catastrophic outcome is not sufficiently small to give comfort). The critical question, which tail fatness quantifies, is *how fast* does the probability of a catastrophe decline relative to the welfare impact of the catastrophe.

Unless otherwise noted, my default meaning of the term "fat tail" (or "thin tail")[11] henceforth concerns the upper tail of the PDF of $\ln D$, resulting from whatever combination of probabilistic temperature changes, temperature-sensitive damages, utility functions, discounting, and so forth, makes this come about. *This* is the PDF that ultimately matters. Empirically, it is not the fatness of the tail of the climate sensitivity PDF *alone* or the reactivity of the damages function to high temperatures *alone*, or the degree of relative risk aversion *alone*, or the rate of pure time preference *alone*, or any other factor *alone*, that counts, but rather the *interaction* of all such factors in determining the upper-tail fatness of the PDF of $\ln D$. For example, other things being equal, the PDF of $\ln D$ will have a relatively fatter tail the larger is the probability of high temperatures and the greater is the reactivity of the damages function to high temperatures, but neither condition alone implies a fat-tailed PDF of $\ln D$. It may seem arcane, but the upper-tail fatness of the reduced form PDF of $\ln D$ is the core issue in the economics of catastrophic climate change. Of course, it is extremely difficult to know the fatness of the upper tail of the PDF of $\ln D$, which is precisely the main point of this line of research and this article.

In Weitzman (2009a), I indicated a theoretical tendency for the PDF of $\ln D$ to have a fat tail. Conceptually, the underlying mechanism is fairly straightforward. Structural uncertainty essentially means that the probabilities are unsure. A formal Bayesian translation might be that the structural parameters of the relevant PDFs are themselves uncertain and have their own PDFs. Weitzman (2009a) expressed this idea in a formal argument that the reduced form "posterior predictive" PDF (in Bayesian jargon) of $\ln D$ tends to be fat tailed because the structural parameters are unknown. Loosely speaking, the driving mechanism is that the operation of taking "expectations of expectations" or "probability distributions of probability

---

[11] There is some wiggle room in the definition of what constitutes a fat-tailed PDF or a thin-tailed PDF, but everyone agrees that probabilities declining exponentially or faster (like the Normal) are thin tailed, while probabilities declining polynomially or slower (like the Pareto) are fat tailed. The standard example of a fat-tailed PDF is the power law (aka Pareto aka inverted polynomial) distribution, although, for example, a Student-t or inverted-gamma PDF is also fat tailed. A normal or a gamma are examples of thin-tailed PDFs, as is any PDF having finite supports, like a uniform distribution or a discrete-point finite distribution. Although both PDFs must approach a limit of zero, the ratio of a fat-tailed probability divided by a thin-tailed probability goes to infinity in the limit.

distributions" tends to spread apart and fatten the tails of the compounded posterior predictive PDF. From past samples alone, it is inherently difficult to learn enough about the probabilities of extreme future events to thin down the "bad" tail of the PDF because we do not have much data about analogous past extreme events. This mechanism provides at least some kind of a generic story about why fat tails might be inherent in some situations.

The part of the distribution of possible future outcomes that we might know now (from inductive information of a form conveyed by past data) concerns the relatively more likely outcomes in the middle of the probability distribution. From past observations, plausible interpolations or extrapolations, and the law of large numbers, there may be at least some modicum of confidence in being able to construct a reasonable picture of the central regions of the posterior-predictive PDF. As we move toward probabilities in the periphery of the distribution, however, we are increasingly moving into the unknown territory of subjective uncertainty, where our probability estimates of the probability distributions themselves become increasingly diffuse because the frequencies of rare events in the tails cannot be pinned down by previous experiences. From past data alone, it is not possible to know enough now about the frequencies of future extreme tail events to make the outcomes of a BCA be independent from artificially imposed limitations on the extent of possibly catastrophic outcomes. Climate change economics generally and the fatness of climate change tails specifically are prototypical examples of this principle, because we are trying to extrapolate inductive knowledge far outside the range of limited past experience. To put a sharp point on this seemingly abstract issue, the thin-tailed PDFs that implicitly support gradualist conclusions have at least some theoretical tendency to morph into fat-tailed PDFs when it is admitted that we are unsure about the functional forms or structural parameters behind these implicitly assumed thin-tailed PDFs—at least where high temperatures are concerned.

A fat upper tail of the PDF of $\ln D$ makes the willingness to pay (WTP) to avoid extreme climate changes very large, indeed arbitrarily large if the coefficient of relative risk aversion is bounded above one. In Weitzman (2009a), I presented a formal argument within a specific mathematical structure, but this formal argument could have been embedded in alternative mathematical structures—with the same basic message. The particular formal argument I presented was in the form of what I called a "dismal theorem" (DT). In this particular formalization, the limiting expected stochastic discount factor is infinite (or, what I take to be equivalent for purposes here, the limiting WTP to avoid fat-tailed disasters constitutes all of output). Of course, in the real world, WTP is not 100 percent of output. Presumably the PDF in the bad fat tail is thinned, or even truncated, perhaps due to considerations akin to what lies behind the "value of a statistical life" (VSL)—after all, we would not pay an infinite amount to eliminate the fat upper tail of climate change catastrophes. Alas, in whatever way the bad fat tail is thinned or truncated, a climate change BCA based upon it might remain sensitive to the details of the thinning or truncation mechanism because the disutility of extreme climate change is "essentially" unbounded.[12] Later, I discuss the meaning of this potential lack of robustness in climate change BCA and speculate on some of its actionable consequences.

---

[12] There is "essentially" unlimited liability here because global stakeholders cannot short the planet as a hedge against catastrophic climate change.

Interpreting Infinity

Disagreements abound concerning how to interpret the infinity symbol that appears in the formulation of the DT. There is a natural tendency to scoff at economic models that yield infinite outcomes. This reaction is presumably inspired by the idea that infinity is a ridiculous result, and that, therefore, an economic model that yields an infinity symbol as an outcome is fundamentally misspecified, and thus dismissable. Critics cite examples to argue earnestly that expected disutility from climate change cannot actually be infinite, as if this were an indictment of the entire fat-tailed methodology. I believe that, in the particular case of climate change, the infinity symbol is trying to tell us something important. That is, the infinite limit in the DT is a formal mathematical way of saying that structural uncertainty in the form of fat tails is, at least in theory, capable of swamping the outcome of any BCA that disregards this uncertainty.

The key issue here is *not* a mathematically illegitimate use of an infinite limit in the DT. Infinity is a side show that has unfortunately diverted attention from the main issue, which is nonrobustness. It is easy to modify utility functions, to add on VSL-like restrictions, to truncate probability distributions arbitrarily, or to introduce ad hoc priors that cut off or otherwise severely dampen low values of welfare-equivalent consumption. Introducing any of these (or many other attenuating mechanisms) formally replaces the infinity symbol by some uncomfortably large, but finite, number. Unfortunately, removing the infinite limit in these or other ways does not eliminate the underlying problem because it then comes back to haunt us in the form of a WTP that is arbitrarily large in order to erase the structural uncertainty. Just how large the WTP is can depend greatly upon obscure details about how the upper tail of the PDF of $\ln D$ has been thinned or bounded.

As a case in point of just how fuzzy an actual upper bound might be, consider imposing VSL-like restrictions on the WTP to avoid planet ruining climate change. Individuals make choices all the time that involve trading off some amount of monetary wealth for some change in the probability of exposure to fatal risks. The VSL is typically defined as the monetary premium $\triangle M$ a person would be willing to pay to avoid exposure to a tiny increased probability of death $\triangle q$, per increment of probability change. A large number of studies have been used to estimate the VSL. Some time period is chosen, over which $\triangle M$ and $\triangle q$ are measured, and then the ratio $\text{VSL} = \triangle M/\triangle q$ is calculated. Very rough VSL values for the United States have typically been estimated at about \$10 million.[13] With average per capita income in the United States at ≈\$50,000 per year, the VSL represents some two hundred years of income per unit change in mortality probability. In the case of climate change, GHG concentrations of 800 ppm of CO_2e imply that $\text{Prob}_{\mathbf{P}}\,[T \geq 10°\text{C}] \approx 1/20$ (see Table 2). Suppose for the sake of argument that $T \geq 10°\text{C}$ represents something like the "death of humanity." Then a naively calculated per capita WTP to avoid this scenario would be about ten years' worth of income, a big number. The fact that this scenario would occur centuries from now, if at all, would lower this WTP. The fact that this event represents the "death of humanity," rather than the death of a single human, would raise this WTP, perhaps considerably. Numerous other underlying considerations would also affect the calculation of this WTP. My numerical example here has many serious flaws, is easy to criticize, and is extraordinarily nonrobust. But I think it

---

[13] See, for example, the numbers cited in Viscusi and Aldi (2003), updated into 2010 dollars.

illustrates how difficult it can be in practice to place an upper bound on the WTP to avoid catastrophic climate change.

Thus, although one can easily remove the infinity symbol from the DT, one cannot so easily "remove" the basic underlying economic problem of extreme sensitivity to fat tails and the resulting conundrum of deciding policy under such circumstances. The overwhelming majority of real-world BCAs have thin upper tails in $\ln D$ from limited exposure to globally catastrophic risks. However, a few very important real-world situations have effectively unlimited exposure due to structural uncertainty about their potentially open-ended catastrophic reach. Climate change is unusual in potentially affecting the entire worldwide portfolio of utility by threatening to drive all of planetary welfare to disastrously low levels in the most extreme scenarios.

Some Implications of the DT for BCA and Policy

The "standard" BCA approach appears to offer a constructive ongoing scientific–economic research program for generating ever more precise outputs from ever more precise inputs. By contrast, my main message can seem off-putting because it can be painted as antiscientific and antieconomic. Fat tails and the resulting limitations they impose on the ability of BCA to reach robust conclusions are frustrating for economists. After all, we make a living from plugging rough numbers into simple models and reaching specific conclusions (more or less) on the basis of those numbers. What quantitative advice are we supposed to provide to policy makers and politicians about *how much* effort to spend on averting climate change if the conclusions from modeling fat-tailed uncertainties are not clear cut? Practical men and women of action have a low tolerance for vagueness and crave some kind of an answer, so they have little patience for even a whiff of fuzziness from two-handed economists. It is threatening for us economists to admit that constructive "can do" climate change BCA may be up against some basic limitations on the ability of quantitative analysis to yield robust policy advice. But if this is the way things are with the economics of climate change, then this is the way things are. Nonrobustness to subjective assumptions about catastrophic outcomes is an inconvenient truth to be lived with rather than a fact to be denied or evaded just because it looks less scientifically objective in BCA. If this limits the ability to give fine-grained and concrete answers to an impatient public, then so be it.

BCA is valuable, even indispensable, as a disciplined framework for organizing information and keeping score. But all BCAs are not created equal. In rare situations with effectively unlimited downside liability, like climate change, BCAs can be fragile to the specifications of extreme tail events. Perhaps economists need to emphasize more openly to the policy makers, the politicians, and the public that, while formal BCA can be helpful, in the particular case of climate change there is a danger of possible overconfidence from undue reliance on subjective judgments about the probabilities and welfare impacts of extreme events. What we can do constructively as economists is to better explain both the magnitudes of the unprecedented structural uncertainties involved and why this feature limits what we can say, and then present the best BCAs and the most honest sensitivity analyses that we can under fat-tailed circumstances, including many different functional forms for extremes. At the end of the day, policy makers must decide what to do on the basis of an admittedly sketchy economic analysis of a gray area that just cannot be forced to render clear robust answers. The moral of the dismal theorem is that, under extreme tail uncertainty, seemingly casual decisions about functional forms, parameter values, and tail fatness can dominate BCA. Economists should not pursue a narrow, superficially crisp, analysis by blowing away the low-probability, high-impact catastrophic scenarios as if this is a necessary price we must pay for the worthy goal of giving answers and advice to policy makers. An artificial infatuation with crispness is likely to make our analyses go seriously askew and undermine the credibility of what we have to offer by effectively marginalizing the very possibilities that make climate change so grave in the first place.

Role of Systemic Inertia and Learning

The issue of how to deal with the deep structural uncertainties of climate change would be completely different and immensely simpler if systemic inertias, such as the time required for the system to naturally remove extra atmospheric CO_2, were short, as is the case for many airborne pollutants like sulfur dioxide, nitrous oxides, and particulates. Then an important component of an optimal strategy might be along the lines of "wait and see." With strong reversibility, an optimal climate change policy would logically involve (among other elements) waiting to learn how far out on the bad fat tail the planet will end up, followed by midcourse corrections if we seem to be headed for a disaster. Alas, the problem of climate change seems bedeviled at almost every turn by significant stock-accumulation inertias—in atmospheric CO_2, in the absorption of heat or CO_2 by the oceans, and in many other relevant physical and biological processes—that are slow to respond to attempts at reversal.

Take atmospheric CO_2 as a prime example. Solomon et al [2009] calculated how concentrations of CO_2 would be expected to fall off over time if *all* anthropogenic emissions were to cease immediately, following a future 2 percent annual growth rate of emissions up to peak concentrations of 450, 550, 650, 750, 850, and 1,200 ppm. As the authors state: "The example of a sudden cessation of emissions provides an upper bound to how much reversibility is possible, if, for example, unexpectedly damaging climate changes were to be observed." Results differed for different trajectories and scenarios, but a crude rule of thumb seemed to be that approximately 70 percent of the peak concentration level over the preindustrial level of 280 ppm persevered after one hundred years of *zero* emissions, while approximately 40 percent of the peak concentration level over the preindustrial level of 280 ppm persevered after one thousand years of *zero* emissions. Based on these estimates, were atmospheric CO_2 concentrations to peak at 800 ppm, followed forever after by *zero* emissions, then atmospheric concentrations would be $\approx$650 ppm after one hundred years and $\approx$500 ppm after one thousand years. A different recent study reached essentially the same conclusions (Archer et al. 2009). These numbers do not look to me like evidence supporting "wait and see" policies. The capacity of the oceans to take up atmospheric heat, the saturation of carbon sinks, the loss of albedo, and many, many other relevant mechanisms tell a similar story of long stock-accumulation irreversibilities relative to the time it takes to filter out and act upon meaningful signals of impending disasters.

Alternative Specifications and Catastrophic Outcomes

Many researchers promote alternative specifications that imply outcomes that are much less extreme than those implied by my specifications. I am not arguing that these alternative formulations are wrong or even implausible. I am merely pointing out that they are not likely to be robust with respect to assumptions about extreme catastrophic climate change and that they therefore fail an important "stress test." Of course the reader should weigh the plausibility of the arguments and the reasonableness of the various specifications on their own merits. But it is difficult to form opinions about probabilities of climate change extremes, or about disutility functions for extreme temperatures, or about lots of other things that are relevant for deciding the tail fatness of the PDF of $\ln D$. Suppose, for the sake of argument, that a policy maker believes the probability is 50 percent that my fat-tailed specification is correct and 50 percent that the thin-tailed specification of someone else is correct. Then, other things being equal, rational policy should lean more in the direction of my fat-tailed conclusions than in the direction of someone else's thin-tailed conclusions because of the highly asymmetric consequences of fat tails versus thin tails. In this sense, whether it is fair or unfair, the playing field is not level between me and someone else. To further illustrate this point, suppose one person advises you that a fire insurance policy protecting your house against extreme losses is unnecessary because so few houses of your kind burn to the ground, while another person advises you that a complete fire insurance policy is necessary in your case. Other things being equal, should you flip a coin to decide what to do just because both advisers seem to be giving equally credible guidance?

Climate change is not the only possible catastrophic threat to humanity. In Weitzman (2009a), I listed what I consider to be the half-dozen or so serious contenders with climate change for potentially catastrophic global impacts with nonnegligible probabilities—biotechnology, nanotechnology, asteroids, "strangelets," pandemics, runaway rogue computers, nuclear proliferation—and went on to give a few tentative reasons why I think that climate change is especially worrisome. It may well be that each of the other half-dozen or so serious candidates for fat-tailed disasters deserves its own ballpark estimate of tail probabilities along with extremely crude calculations of policy implications, which is about the best we can do with potential catastrophes. Even if this were true, however, it would not lessen the need to reckon with the strong potential implications of fat tails for BCA-like calculations in the particular case of climate change.

Conclusions

Taking fat tails into account has implications for climate change research and policy. For example, perhaps more emphasis should be placed on research about the extreme tails of relevant PDFs rather than on research about central tendencies. As another example, the fatness of the bad fat tail of proposed solutions (such as, perhaps, the possibility that buried CO_2 might escape) needs to be weighed against the fatness of the tail of the climate change problem itself. With fat tails generally, we might want more explicit contingency planning for bad outcomes, including, perhaps, a niche role for last-resort portfolio options like geoengineering.

Qualitatively, fat tails favor more aggressive policies to lower GHGs than the "standard" BCA. Alas, the quantitative implications are less clear. As this article has stressed, the natural consequence of fat-tailed uncertainty should be to make economists less confident about climate change BCA and to make them adopt a more modest tone that befits less robust policy advice. My own conclusion is that the sheer magnitude of the deep structural uncertainties concerning catastrophic outcomes, and the way we express this in our models, is likely to influence plausible applications of BCA to the economics of climate change for the foreseeable future.

References

Archer, David. 2007. *Global Warming*. Oxford, UK: Blackwell Publishing.

Archer, David, M. Eby, V. Brovkin, A. Ridgwell, L. Cao, U. Mikolajewicz, K. Caldeira, K. Matsumoto, G. Munhoven, A. Montenegro, and K. Tokos. 2009. Atmospheric lifetime of fossil fuel carbon dioxide. *Annual Review of Earth and Planetary Sciences* 37: 117–134.

Dieter Lüthi, Martine Le Floch, Bernhard Bereiter, Thomas Blunier, Jean-Marc Barnola, Urs Siegenthaler, Dominique Raynaud, Jean Jouzel, Hubertus Fischer, Kenji Kawamura, and Thomas F. Stocker. 2008. High-resolution carbon dioxide concentration record 650,000–800,000 years before present. *Nature* 453: 379–382.

IPCC-AR4. 2007. *Climate change 2007: The physical science basis.* Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK: Cambridge University Press. (http://www.ipcc.ch).

Hansen, James, Makiko Sato, Puskher Karecha, David Beerling, Valerie Masson-Delmotte, Mark Pagani, Maureen Raymo, Dana L. Royer, and James C. Zachos. 2008. Target atmospheric CO_2: Where should humanity aim? *Open Atmospheric Science Journal* 2: 217–31.

Nordhaus, William D. 2008. *A question of balance*. New Haven, CT: : Yale University Press.

———. 2011. The economics of tail events with an application to climate change. *Review of Environmental Economics and Policy,* 5(2):240–57. 10.1093/reep/rer004.

Natalia Shakhova, Igor Semiletov, Anatoly Salyuk, Vladimir Yusupov, Denis Kosmach, and Örjan Gustafsson. (2010). Extensive methane venting to the atmosphere from sediments of the East Siberian Arctic Shelf. *Science* 327: 1246–50.

Pindyck, Robert S. 2011. Fat tails, thin tails, and climate change policy. *Review of Environmental Economics and Policy,* 5 (2): 258–74. 10.1093/reep/rer005.

Sherwood, Steven C., and Matthew Huber. 2010. An adaptability limit to climate change due to heat stress. *Proceedings of the National Academy of Sciences* 107: 9552–5.

Solomon, Susan, Gian-Kasper Plattner, Reto Knutti, and Pierre Friedlingststein. 2009. Irreversible climate change due to carbon dioxide emissions. *Proceedings of the National Academy of Sciences* 106: 1704–9.

Sterner, Thomas, and U. Martin Persson. 2008. An even Sterner Review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy* 2: 61–76.

Viscusi, W. Kip, and Joseph E. Aldy. 2003. The value of a statistical life: A critical review of market estimates throughout the world. *Journal of Risk and Uncertainty* 27: 5–76.

Weitzman, Martin L. 1998. Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management* 36: 201–8.

———. 2007. The Stern Review of the economics of climate change. *Journal of Economic Literature* 45: 703–24.

———. 2009a. On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 91: 1–19.

———. 2009b. Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics — The Open-Access, Open-Assessment E-Journal* 3: 1–29.

———. 2010. Risk-Adjusted Gamma Discounting. *Journal of Environmental Economics and Management* 60: 1–13.

Zickfeld, Kirsten, M. Granger Morgan, David J. Frame, and David W. Keith. 2010. Expert judgments about transient climate response to alternative future trajectories of radiative forcing. *Proceedings of the National Academy of Sciences* 106: 16129–35.