

# The Extreme Uncertainty of Extreme Climate Change: An Overview and Some Implications

Martin L. Weitzman\*

October 20, 2009. Preliminary. Comments Appreciated

## Abstract

The possibility of catastrophic climate change is characterized by deep structural uncertainties in the science coupled with an economic inability to evaluate meaningfully the welfare losses from high temperatures. The probability of a disastrous collapse of planetary welfare from global warming seems non-negligible, even if this low probability is very difficult to quantify. Through informal reasoning, elementary examples, and simple numerical exercises, this paper attempts to convey an overview of some of the background uncertainties behind extreme climate change. I argue that the tails of the relevant probability distributions should not be ignored because they are likely to be fat with probability and important. A few implications for climate change analysis and policy are explored.

## 1 Introduction

An important question often asked about climate change is: how bad is it likely to get? The possibility of catastrophic climate change is characterized by deep structural uncertainties in the science combined with severe constraints on the ability to evaluate meaningfully the welfare losses from high temperatures. A key starting point for any cost-benefit analysis (CBA) of climate change must begin by recognizing that values of key future variables – temperatures, climate, comprehensive damages, overall welfare, and so forth – cannot be known now. Instead, they must be conceptualized as random variables drawn from some probability density function (PDF), however fuzzy is that PDF. This paper tries to argue that the probability of a disastrous collapse of planetary welfare from global warming appears

---

\*Department of Economics, Harvard University (mweitzman@harvard.edu). Without tying them to the contents of this paper, I am grateful for critical comments by ...

to be non-negligible, even though this low probability is itself highly uncertain. The tone of the paper is mostly informal and relatively nontechnical, involving discussions, elementary examples, and simple numerical exercises that anyone with a hand calculator can readily reproduce or modify. I think that macro-simplicity and informal discussion may be useful when dealing with highly uncertain climate change extremes, because otherwise it is difficult to see the forest for the trees.

I will argue that the PDFs of distant-future temperatures, comprehensive damages, overall welfare, and so forth, are plausibly fat tailed. By contrast, most traditional CBAs implicitly treat uncertainty in the form of thin-tailed PDFs. A thin-tailed PDF assigns a relatively much lower probability to rare events in the extreme tails than does a fat-tailed PDF.<sup>1</sup> In climate change, it looks to me as if the fat tails of the relevant PDFs might matter a lot. By employing some basic examples and a few simple numerical exercises, the paper attempts to show that the projected confidence levels for excluding catastrophic global warming seem too low for comfort and that these confidence levels are likely to further deteriorate over time at a rapid rate. In other words, it appears as if the “bad” tail of the relevant PDF is already fat and will likely get much fatter in the future as greenhouse gases (GHGs) accumulate in the atmosphere.

The paper provides a broad overview of some of the most critical uncertainties about extreme climate change. The seemingly fat tails of the relevant PDFs, when coupled with nearly unlimited planet-wide liability on the downside, can cause potentially serious challenges for cost-benefit analysis (CBA). Because the (known and unknown) unknowns of severe climate change are so far outside the realm of ordinary experience, probabilistic CBA in the arena of climate change disasters is significantly more fuzzy than probabilistic CBA of more traditional situations – like reinforcing bridges, constructing dikes and levees, or setting building codes for earthquake-prone areas. The paper tries to make sense of the anomaly of catastrophic global warming and explores what might be done in terms of actionable alternatives under such fuzzy circumstances. The welfare economics of catastrophic climate change is highly unsure, in large part because economists are highly unsure about almost all aspects concerning how to appraise distant-future disaster scenarios. This economic uncertainty adds considerably more vagueness to the already formidable scientific uncertainty and

---

<sup>1</sup>As I use the term, a PDF has a “fat” (or “thick” or “heavy”) tail when the tail probability approaches zero *more slowly* than exponentially. The standard example of a fat-tailed PDF is the power law or Pareto distribution, which approaches zero polynomially. A lognormal PDF is also fat-tailed, as is a Student- $t$  or inverted-gamma. By this more or less standard definition, a PDF has a “thin” tail when its tail probability approaches zero *exponentially or faster*. 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 distribution. The ratio of a fat-tailed probability divided by a thin-tailed probability goes to infinity in the limit even though both PDFs approach zero.

may be better treated somewhat informally in an overview paper. I try to explain in broad terms what CBA can and cannot reasonably be expected to accomplish in the presence of extreme structural uncertainty about extreme climate change.

The paper also discusses some policy implications of fat-tailed climate change. One implication suggests some possible directions for future research. Another is the need to compare fat-tailed problems with possibly fat-tailed solutions. Yet another issue concerns the role of learning and the ability to make midcourse corrections. A few other policy implications are touched upon briefly. The paper builds to the idea that we desperately need some semblance of a game plan to be ready for worst-case scenarios that might be coming down the road with low but non-negligible probabilities. One tentative suggestion is that, no matter what else is done realistically to slow GHG buildups, economic analysis lends some support to undertaking serious research into the prospects of “geoengineering preparedness” – as a state-contingent emergency fallback option offering the possibility of buying time by knocking down catastrophic temperatures rapidly.

## 2 Deep Structural Uncertainty about Climate Extremes

In this section I try to make a heuristic-empirical case for there being big structural uncertainties in almost all aspects of extreme climate change. I will argue on intuitive grounds that it would be curious if the way in which this deep structural uncertainty is conceptualized and formalized did not influence the outcomes of a climate-change CBA. My arguments in this section are not intended to be airtight or rigorous. Rather, this is a heuristic presentation based on stylized facts. I try to make this case by citing as examples only four aspects of climate change that do not seem to me to be adequately covered by conventional economic analyses. While many different aspects of structural uncertainty might additionally be cited, I restrict my case to these four examples, which I will call “Exhibits A, B, C and D.” The next section of the paper will attempt to quantify – extremely crudely – some ballpark estimates of the possible significance of these examples.

“Exhibit A” concerns the atmospheric level of (GHGs) over the last 800,000 years. Ice core drilling in Antarctica began in the late 1970s and is still ongoing. The record of carbon dioxide (CO<sub>2</sub>) and methane (CH<sub>4</sub>) trapped in tiny ice-core bubbles currently spans 800,000 years.<sup>2</sup> 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,

---

<sup>2</sup>See Dieter et al (2008), from which my numbers are taken (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](ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2_mm_mlo.txt)).

whereas these ice-core GHG data are directly observed.

The pre-industrial-revolution level of atmospheric CO<sub>2</sub> (about two centuries ago) was 280 parts per million (ppm). The ice-core data show that CO<sub>2</sub> varied gradually during the last 800,000 years within a relatively narrow range roughly between 180 and 280 ppm. Currently, CO<sub>2</sub> is at 385 ppm, and climbing steeply. Methane was never higher than 750 parts per billion (ppb) in 800,000 years, but now this extremely potent GHG, which is 22 times more powerful than CO<sub>2</sub> (per century), is at 1,780 ppb. The sum total of all carbon-dioxide-equivalent (CO<sub>2</sub>e) GHGs is currently at 438 ppm. Even more alarming is the rate of change of GHGs, with increases in carbon dioxide hardly ever exceeding 30 ppm over *any* past thousand-year period, while now CO<sub>2</sub> has risen by 30 ppm in just the last 17 years. Thus, anthropogenic activity has elevated atmospheric CO<sub>2</sub> and CH<sub>4</sub> to levels far outside their natural range at a very rapid rate. 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 even on a time scale of hundreds of millions of years. We seem headed for a unique planetary experiment of subjecting the Earth's system to an unprecedented shock by geologically-instantaneously jolting atmospheric stocks of GHGs far above their highest level over the last several million years. We simply do not know what will happen under such extreme circumstances.

“Exhibit B” concerns the ultimate temperature response to such kind of unprecedented increases in GHGs. “Climate sensitivity” (hereafter denoted  $S$ ) is a key macro-indicator of the *eventual* temperature response to GHG changes. As the Intergovernmental Panel on Climate Change in its IPCC-AR4 (2007) Executive Summary explains the concept: “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 to 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 IPCC defines “likely” as a probability of 66%. A commonly cited guesstimate for the upper 5% confidence level of  $S$  is  $\approx 6^\circ$ , but this guesstimate is accompanied by a great deal of unsureness. The pre-industrial-revolution atmospheric concentration of CO<sub>2</sub> was about 280 parts per million (ppm). If  $P[S > 6^\circ]$  is anywhere in the neighborhood of 5%, it is quite worrisome because, realistically, it seems almost inevitable that CO<sub>2</sub> concentrations will climb well above the doubling level of 560 ppm.

“Exhibit C” concerns possibly disastrous releases over the long run of bad-feedback com-

ponents of the carbon cycle that are currently omitted from most general circulation models. The chief worry here is a significant supplementary component that conceptually should be added on to climate sensitivity  $S$ . This omitted component concerns the potentially powerful self-amplification potential of greenhouse warming 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  $\text{CH}_4$  trapped in the form of hydrates (aka clathrates) – concerning which there is a decidedly non-zero probability over the long run of destabilized methane seeping into the atmosphere if water temperatures over the continental shelves warm just slightly. The amount of  $\text{CH}_4$  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 fossil fuels combined. Over the very long run, a  $\text{CH}_4$  outgassing-amplifier process could potentially precipitate a disastrous strong-positive-feedback warming. Thus, the possibility of a climate meltdown is not just the outcome of a mathematical theory, but it has a real physical basis. Several other examples of an actual physical basis for very bad outcomes could be cited in addition to heat-induced feedbacks from endogenous releases of naturally sequestered GHGs – such as increased respiration of soil microbes, climate-stressed forests, and other weakenings of natural carbon sinks like the oceans. If all indirect carbon cycle feedback effects were included, the PDF of  $S$  would have a tail even heavier with probability.

“Exhibit D” concerns the treatment of damages or disutilities from extreme temperature changes. Most CBAs treat high-temperature damages by whatever specification is assumed to be the low-temperature “damages function.” High-temperature damages extrapolated from a low-temperature damages function seem to be remarkably sensitive to assumed functional forms (and, to a lesser degree, parameter choices). A great variety of functional forms can be made to fit the low-temperature damages assumed by the modeler, even though these functional forms may give enormously different evaluations at higher temperatures. Later I discuss some additional ambiguities and artificialities inherent in the deep uncertainty of appraising severe damages from extreme climate change, prominent among which is the selection and use of an interest rate to discount distant-future events.

To summarize and extend, the economics of climate change consists of a very long chain of tenuous inferences fraught with big uncertainties in every link: beginning with unknown baseline GHG emissions; then compounded by big uncertainties about how available policies and policy levers will transfer into actual GHG emissions; compounded by big uncertainties about how GHG flow emissions accumulate via the carbon cycle into GHG stock concen-

trations; compounded by big uncertainties about how and when GHG stock concentrations translate into global mean temperature changes; compounded by big uncertainties about how global mean temperature changes decompose into regional climate changes; compounded by big uncertainties about how adaptations to, and mitigations of, climate-change damages are translated into utility changes at a regional level; compounded by big uncertainties about how future regional utility changes are aggregated – and then how they are discounted – to convert everything into expected-present-value global welfare changes. 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 mathematically is represented by a PDF that is spread out and heavy with probability in its bad tail.

What I would wish a reader might take away from these four exhibits is the notion that a claim that CBA is robust to such stylized facts seems curious. 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 it is sometimes claimed that CBA is relatively immune to the extraordinarily uncertain probabilities and consequences of catastrophic climate change. A reader should feel intuitively that it goes against the grain of common sense when, in view of such things as the above four examples of structural uncertainty, a climate-change CBA does *not* much depend upon how potential disasters are modeled and incorporated into the CBA.

Throughout the sections of the paper which follow, I want to state clearly at the outset that every number and every formulation is simplistically derived, subject to a myriad of caveats, and open to strong criticisms. This is an unavoidable issue for all aspects of extreme climate change, because we are extrapolating so far outside the realm of ordinary experience that we are left mostly with thought-experimental guesstimates and suggestive numerical examples. The degree of my subjectivity is embarrassing. I lean somewhat towards the side of emphasizing that some not-implausible specifications of the extreme uncertainty of extreme climate change may have big consequences, but I try to keep the examples relatively unbiased and I hope that they are not too unfair. My only excuse is that I do not know a better approach for highlighting what is potentially a very important issue.

In principle, one could plug various scenarios into a big dynamic model and come up with numerical simulations of uncertainty, but when it comes to analyzing extremes I fear that such an approach can miss the forest for the trees. The evaluation of extreme outcomes will depend on the assumptions about extreme climate change being plugged into the simulations, a connection that can get obscured in a big model. This paper takes the opposite approach to the problem of analyzing catastrophic climate change. The emphasis here is on simplicity and transparency. I derive closed form solutions from ultra-simplistic formulations, in which

it is crystal clear what is the core structure and where the results are coming from. One advantage of this approach is to make everything so simple and transparent that anyone can easily plug in alternative numbers or assumptions and come up with their own extremely crude ballpark estimates. A possible disadvantage of my approach is that it may miss subtle interactions that only a bigger multi-equation dynamic model might hope to capture.

### 3 What is the Appropriate “Damages Function”?

I begin with the so-called “damages function.” Let  $T$  represent the *change* in future world-wide average surface temperature, always measured in degrees Centigrade. The initialization time for  $T = 0^\circ$  is not very relevant for this section, so long as it is chosen to be somewhere between the beginning of the industrial revolution and before 1950. Let  $C^*(T)$  represent “welfare equivalent” consumption as a *fraction* of what potential consumption would be if  $T = 0$ . (These concepts themselves have some problematical aspects, which are ignored here, although my intention is that loss of the environment as we know it is included in “welfare equivalent” consumption.) The most popular single formulation in the literature is of the quadratic form

$$C^*(T) = \frac{1}{1 + aT^2}, \quad (1)$$

where  $a$  is a parameter calibrated to give some “reasonable” estimates of  $C^*(T)$ . Most standard estimates of  $a$  in the literature are more or less similar, although I hasten to add in fairness that such calibrations were intended by the authors to capture low-temperature damages and were never intended to be extrapolated to very high temperature changes, which is just what I am doing here. For the sake of having a simple prototype example, I calibrate  $a$  in (1) to conform with the damages function in the well known DICE<sup>3</sup> model of William Nordhaus (2008), which gives  $a = .003$  (approximately). The results in terms of welfare-equivalent consumption are given in the second column of Table 1. I do not find such numbers credible for high temperatures. At  $T = 20^\circ$ , the welfare-equivalent consumption level is  $C^*(20^\circ) = 45\%$ , which means that “only” 55% of consumption is lost at this unimaginably high temperature. The implied consumption losses of 23% for  $T = 10^\circ$  and 16% for  $T = 8^\circ$  also seem far too low for doing a serious analysis of the consequences of catastrophic losses from extreme climate change. To check whether the quadratic *polyno-*

---

<sup>3</sup>Nordhaus’s DICE model is perhaps the most famous integrated assessment model (IAM) in the economics of climate change. My estimate of  $a$  comes from his Figure 3-3 on page 51. Nordhaus wisely does not try to project beyond  $T = 6^\circ$ , which is suitable for his purposes but unsatisfactory for mine.

*mial* form in (1) is responsible for giving such low losses at such extreme temperatures, in the third column of Table 1 is tabulated the corresponding quadratic *exponential* version  $C^*(T) = \exp(-aT^2)$ . This makes some difference, but not very much, as I would say that the high-temperature damages remain far too low for doing a credible analysis of the effects of catastrophic losses from extreme climate change. My tentative conclusion is that the quadratic form (in polynomial or exponential versions), which was never intended to be applied for temperature changes beyond a few degrees centigrade, is not appropriate for assessing the welfare impacts of disastrously high temperature changes. In other words, the “damages functions” enumerated in columns two and three of Table 1 are (undoubtedly unconsciously) preordained to make extreme climate change look empirically negligible, almost no matter what else is assumed, including the fatness or thinness of the extreme tails of the relevant PDFs.

$T$	$C^* = 1/(1 + aT^2)$	$C^* = \exp(-aT^2)$	$C^* = 1/(1 + bT^4)$	$C^* = \exp(-bT^4)$
2°C	98.8%	98.8%	99.2%	99.2%
2.5°C	98.2%	98.2%	98.2%	98.2%
3°C	97.4%	97.4%	96.3%	96.2%
4°C	95.4%	95.3%	89.1%	88.4%
5°C	93.0%	92.8%	76.9%	74.1%
6°C	90.3%	89.8%	61.6%	53.7%
7°C	87.2%	86.3%	46.5%	31.6%
8°C	83.9%	82.5%	33.7%	14.0%
9°C	80.5%	78.4%	24.1%	4.3%
10°C	76.9%	74.1%	17.2%	0.8%
12°C	69.8%	64.9%	9.1%	$5 \times 10^{-5}$
15°C	59.7%	50.9%	4.0%	$3 \times 10^{-11}$
20°C	45.5%	30.1%	1.3%	$4 \times 10^{-34}$

**Table 1:** Welfare-equivalent consumption  $C^*(T)$ . Parameter values  $a=.003$ ,  $b=.00048$ .

Suppose for the sake of argument that I double the exponent so that the quadratic form in the first two columns of Table 1 becomes a quartic form in the last two columns of Table 1. Why double a quadratic exponent to make a quartic? I have no good story except that extrapolating a quartic to high temperatures seems to me no less reasonable offhand than extrapolating a quadratic to high temperatures, and it may be worthwhile to examine what difference it makes. Nobody knows what is the right form for welfare-equivalent

consumption at high temperatures. All anyone can do is to try different forms and pray that the results are “reasonable.” Alas, “reasonableness” in this context must of necessity be a subjective judgement derived from what is at best a fuzzy thought experiment in applied science fiction. The only defence of such exercises is not knowing what else to do about a potentially important issue.

In Nordhaus’s and many other calibrations, the consumption damages at  $T = 2.5^\circ$  are estimated to be in the neighborhood of 1.84%. If one were forced to choose a single dividing line, I think it is fair to regard  $T \approx 2.5^\circ$  as an upper limit on the temperature change for which a plausible approximation of damages might hope to be estimated. As we move beyond  $T \approx 2.5^\circ$ , I believe we are increasingly involved in speculation about welfare-equivalent consumption. In such a spirit, I calibrate the value of  $b$  in column four of Table 1 to give the same welfare-equivalent consumption as column two at  $T = 2.5^\circ$ , so that  $1/(1 + a(2.5)^2) = 1/(1 + b(2.5)^4)$ . Then  $a = .03$  implies  $b = .00048$ . The intent is that the quadratic approximation perhaps applies for  $T \leq 2.5^\circ$ , while the quartic approximation perhaps applies for  $T > 2.5^\circ$ , although this interpretation is not essential. Of course, others are free to make their own calibrations. The only thing I am trying to show is that extrapolation of damages beyond  $T \approx 2.5^\circ$  (much less beyond  $T \approx 4^\circ$ ) is very sensitive to assumptions about functional forms, which are virtually impossible to know a priori.

As is evident, the quartic gives much more pessimistic results than the quadratic at higher temperatures. Furthermore, the distinction between the *polynomial* quartic (column 4) and the *exponential* quartic (column 5) is striking at very high temperatures, which should be contrasted with the minor differences between the polynomial and exponential quadratic formulations of columns 2 and 3.

Now begins a series of questions that have no clear-cut answers. How is one to choose between a quadratic and a quartic at high temperatures? What about a cubic, or a polynomial to the fifth power, or a combination of polynomials to various powers? How does one select between exponential and polynomial formulations, which give similar results for low temperatures but can differ enormously for higher temperatures? What should be the calibration point at which all formulations are made to give the same welfare-equivalent consumption:  $2.5^\circ$  (as in this example), or  $4^\circ$ , or yet some other temperature or temperatures? And so on and so forth.

The damages for  $T = 4^\circ$  in columns 4 and 5 of Table 1 are more than twice as high as the damages for  $T = 4^\circ$  in columns 2 and 3. But the damages for  $T = 4^\circ$  in the more standard formulation of columns 2 and 3 are primarily a quadratic extrapolation of postulated damages from lower temperatures, such as from the base case  $T = 2.5^\circ$ , rather than being an independently established estimate having a more or less objective empirical

foundation. Besides, more sophisticated curve fitting could easily make all of the damages for  $T = 4^\circ$  in Table 1 be much closer together, while emulating the distinctions in Table 1 between quadratic and quartic formulations for much higher temperatures.

A temperature increase of  $4^\circ$  is likely to have some very serious consequences. One recent study<sup>4</sup> asked 52 experts for their subjective probability estimates of triggering a “tipping point of major changes” in each of five possible categories: (1) the Atlantic meridional overturning circulation; (2) the Greenland ice sheet; (3) the West Antarctic Ice Sheet; (4) the Amazon rainforest; (5) the El Niño/Southern Oscillation. For what it is worth, the probability-weighted expected number of such expert-assessment tipping points was *three* (out of a possible five) for a scenario where temperatures increase gradually from  $T \approx 0^\circ$  in year 2000 to  $T \approx 4^\circ - 8^\circ$  in year 2200. There are many other bad things that might conceivably happen over the long run at  $T \approx 4^\circ$ , such as a dieback of boreal forests, releases of methane hydrate from permafrost and, eventually, from the continental shelves, decline of carbon-sink and temperature-absorbing abilities of the oceans, and so forth. I don’t know whether welfare-equivalent consumption losses for  $T \approx 4^\circ$  should be closer to 5% or closer to 10% or closer yet to some third number. For what it is worth, if I were forced to choose between 5% damages and 10% damages at  $T \approx 4^\circ$ , I personally would pick the latter. Others are free to make their own choices and calibrations, of course.

When it comes to  $T \approx 5^\circ$ , some extraordinarily crude comparisons can perhaps be made with a few past geological events. The lowest average temperature of the last ice age was about  $5^\circ$  colder than the present. A new ice age would presumably cause a massive decline in planetary welfare, albeit over a long time period. I don’t know what parallels can be made between  $T \approx -5^\circ$  and  $T \approx +5^\circ$  but I would have to guess that welfare-equivalent losses or damages (including the environment) from  $T \approx +5^\circ$  are closer to the  $\approx 25\%$  in columns 3 and 4 of Table 1 (and maybe even greatly exceed that number) than to the  $\approx 7\%$  of the first two columns.

The last time that the world witnessed temperatures  $\approx +5^\circ$  above the present was during the Eocene period, some 30 to 50 million years ago (mya). In this period the earth was ice free, the land was swampy, and alligators lived near the North Pole. The last time that there was a *geologically rapid* increase in temperatures of magnitude  $\approx +5^\circ$  or so was during the Paleocene-Eocene Thermal Maximum (PETM)<sup>5</sup> about 55 mya, which was accompanied by mass extinctions. This event occurred “geologically rapidly” over a period of about 20K years, which is extremely gradual compared with current worst-case anthropogenically-

---

<sup>4</sup>Kriegler et al, (2009).

<sup>5</sup>My information in this paragraph comes from the Wikipedia entry for PETM and the 42 scientific sources that it cites.

induced trajectories, and was followed by a recovery period that lasted about 150K years. It is unknown what triggered the PETM, but one leading culprit is the strong-feedback release of large amounts of methane hydrates from clathrate formations, mentioned in the previous section. The major point here is that relatively rapid changes of  $\approx +5^\circ$  are extremely rare events very far outside the scope of human experience.

As for huge temperature increases like  $T \approx 10^\circ$ , the planetary effects are almost unimaginable. For what it is worth, which is very little, welfare-equivalent consumption for  $T = 10^\circ$  in Table 1 ranges from about 75% for quadratic formulations to about 1% for the exponential quartic. Once again I am inclined towards the lower number, but anybody’s guess here is undoubtedly as good as mine.

I could go on experimenting with various extrapolations for high-temperature damages, but I think the point has been made already. Many functional forms can be constructed to give about the same welfare-equivalent consumption for low temperature changes, but which give profoundly different results for high temperatures. We simply don’t know what are the appropriate damages for higher temperatures, but some scenarios, which are not totally implausible, seem very bad indeed. So, on intuitive grounds, it does not seem plausible to me that outcomes of climate-change CBAs should not much depend upon what is assumed about damages from high-temperature extremes.

In all of the formulations behind Table 1, the specification is *multiplicative* in some function of temperature change  $T$  times what consumption *would be* if  $T = 0$ . (For example, in equation (1) the consumption-diminishing multiplier is  $1/(1 + aT^2)$ ). This essentially defines a single-attribute utility function, or equivalently, a multi-attribute utility function with strong substitutability between the two attributes. The physical image of a “damages function” of temperature multiplicatively diminishing the ability to produce aggregate consumption, writ large, motivates the popular multiplicative approach to modeling the welfare impacts of global warming. This is 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. In a context of climate change, however, this evaporating-production image might be inappropriate because of the possibility that it can only be extended to a limited subset of other welfare-producing activities or commodities, most of which might better be described by introducing some separate notion of the utility of an “environmental amenity.” Such a formulation would amount to a genuine multi-attribute utility function, and is the more appropriate formulation if the main impact of climate change is on things that are not readily substitutable with material wealth, such as biodiversity and health.

The multiplicative approach implies a high degree of substitutability between output and temperature change, whose limiting elasticity of substitution is one. This accounts

for much of the gradualism seen in models built around this specification, since increased output substitutes so easily for temperature changes. Suppose, though, that the disutility of temperature change is *additively separable* instead of being *multiplicatively separable*. If welfare is instead the additively-separable arithmetic *difference* between a utility function of consumption with constant coefficient of relative risk aversion two minus a quadratic loss function of temperature changes, it is equivalent to having a constant elasticity of substitution (CES) between consumption and temperature change of one half. Empirically, using this additive or CES form – even without any uncertainty – prescribes a significantly much more stringent curtailment of GHG emissions than what emerges from the analogous multiplicative form of the more standard specification<sup>6</sup> – because of the greater difficulty of substituting increased consumption for high temperature changes. So structural uncertainty about the *form* of the utility function, which combines economic output with temperature-dependent environmental amenities, is not ignorable and plausibly can influence significantly climate-change CBA, especially for high temperatures.

Later, I touch briefly upon a few other issues concerning the evaluation of climate-change damages. These include the roles of learning, reversibility, and distant-future discount rates. All of these issues, and many others, seem to be uncertain, to varying degrees. Also highly uncertain is the feasibility and cost of the unfamiliar newly-developed technologies that will be required to diminish significantly atmospheric GHG levels.

## 4 What is the PDF of Long Run Temperature Changes?

Climate sensitivity is the average equilibrium surface warming following a sustained doubling of atmospheric carbon dioxide. Other things being equal, higher values of climate sensitivity raise temperatures in every period by shifting up their dynamic trajectory, but it also takes longer for temperatures to reach any given fraction of their asymptotic limit. Left unanswered by my simplistic treatment here are a great many questions, including whether enough can be learned sufficiently rapidly about high climate sensitivity – relative to tremendous systemic inertias and lags – to be able to undertake realistic midcourse corrections in GHG emissions (more on this later).

What appears to bother scientists a lot about the climate sensitivity issue is that, even

---

<sup>6</sup>With coefficient of relative risk aversion two, the described additively-separable specification is mathematically equivalent to the constant elasticity of substitution (CES) specification of Sterner and Persson (2008) having elasticity of substitution 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 more standard CBA.

after some three decades of intensive research, essentially no progress has been made on excluding very high values (substantially higher than  $4.5^{\circ}\text{C}$ ). A total of twenty-two peer-reviewed studies of climate sensitivity published recently in reputable scientific journals and encompassing a variety of methodologies, along with 22 imputed probability density functions (PDFs) of  $S$ , are cited by IPCC-AR4 (2007), which are compiled in Table 9.3 and Box 10.2 of their report.<sup>7</sup> I assume for my purposes here that these 22 reported PDFs might be simplistically aggregated by averaging the 22 PDFs into one “representative” PDF. This form of meta-analysis can be loosely defended as an example of Bayesian model averaging, in which the different PDFs from different studies represent equally credible outcomes from more or less independent models. How best to aggregate different climate sensitivity PDFs from different studies is a controversial subject in climate-change science. At one extreme a classical-frequentist-inspired meta approach might view the PDFs from the different studies as more or less independent draws from the same “correct” PDF of the same “correct” model, which would narrow the distribution of  $S$ . At another extreme it might be argued that most studies have errors in common from hidden biases due to correlated unknowns, including neglect of slow-feedback carbon cycle processes (such as heat-induced feedbacks from endogenous releases of naturally sequestered GHGs, increased respiration of soil microbes, climate-stressed forests, and other weakenings of natural carbon sinks like the oceans) and a great many other neglected slow-feedback effects, which all point in the direction of higher temperatures.<sup>8</sup> These and other omitted unknowns shared in common would presumably broaden and fatten the distribution of  $S$  beyond my simplistic aggregation, perhaps considerably. As a recent *Science* commentary put it: “Once the world has warmed by  $4^{\circ}\text{C}$ , conditions will be so different from anything we can observe today (and still more different from the last ice age) that it is inherently hard to say where the warming will stop.”<sup>9</sup> In any event, it is a fact that the median upper five percent probability level over

---

<sup>7</sup>Eleven of the studies in Table 9.3 overlap with the studies portrayed in Box 10.2. Four of these overlapping studies conflict on the numbers given for the upper 5% level. For three of these differences I chose the Table 9.3 values on the grounds that all of the Box 10.2 values had been modified from the original studies to make them have zero probability mass above  $10^{\circ}\text{C}$ . (The fact that all PDFs in Box 10.2 have been normalized to zero probability above  $10^{\circ}\text{C}$  biases my upper-5% averages here towards the low side.) With the fourth conflict (Gregory et al (2002a)), I substituted  $8.2^{\circ}\text{C}$  from Box 10.2 for the  $\infty$  in Table 9.3 (which arises only because the method of the study itself does not impose any meaningful upper-bound constraint). The only other modification was to average the three reported volcanic-forcing values of Wigley et al (2005a) in Table 9.3 into one upper-5% value of  $6.4^{\circ}\text{C}$ .

<sup>8</sup>In scientific jargon, climate sensitivity would be associated with “fast feedbacks.” What is more relevant to determining long run temperatures would be (very loosely) associated with “slow feedbacks,” which are typically excluded from general circulation models, mostly on the grounds that they are too uncertain and too long run to be included. See Hansen et al (2008), who estimate that the inclusion of all slow feedbacks would raise a more inclusive concept of what might be called “Earth System sensitivity” from a mean of  $\approx 3^{\circ}\text{C}$  to a mean of  $\approx 6^{\circ}\text{C}$ .

<sup>9</sup>Allen and Frame (2007). For a balanced view of climate sensitivity PDFs with references to other

all 22 climate-sensitivity PDFs cited in IPCC-AR4 (2007) is 6.4°C. Even if this upper *five* percent probability level were knocked down to 6.4°C being the median upper *three* percent probability level (which is what I will effectively do in my calibration), such kinds of numbers are disturbing. Besides, I am using climate sensitivity here as a reduced-form metaphor for *all* scientific uncertainty about climate-change extremes, the sum total of which is presumably much more fat tailed than than the PDF of climate sensitivity alone. There are several additional sources of uncertainty in scientific parameters, including all aspects of the absorption and mixing of CO<sub>2</sub> and heat among ocean layers, aerosol radiative forcing, CO<sub>2</sub> fertilization effects, precipitation changes for higher temperatures, and many others. Here climate sensitivity serves as an imperfect proxy for all of these other sources of scientific uncertainty.

As was previously mentioned, the IPCC estimates that  $P[2^\circ \leq S \leq 4.5^\circ] \approx 66\%$ , with a median value of 3°, which implies that  $P[S \geq 4.5^\circ] \approx 17\%$ , while a commonly cited guesstimate for the upper 5% probability level of  $S$  is  $\approx 6^\circ$ , although this number is highly unsure. To be on the conservative side, I make  $P[S \geq 4.5^\circ] = 16\%$  and  $P[S \geq 6^\circ] = 4\%$ . I don't think these kinds of numbers would be considered by most scientists as placing excessive probability mass on high values of climate sensitivity. For the upper-tail PDF of  $S$  I choose a Pareto distribution of the form

$$P[S \geq T] = k T^{-\gamma}, \tag{2}$$

where  $k$  and  $\gamma$  are positive constants to be determined by calibration. Since I am trying here to estimate only the upper tail of the PDF by (2), I calibrate  $k$  and  $\gamma$  so that  $P[S \geq 4.5^\circ] = 16\%$  and  $P[S \geq 6^\circ] = 4\%$ . This determines parameter values  $\gamma \approx 4.819$  and  $k \approx 224.8$ , which describes a Pareto distribution that is fat-tailed but not inordinately so. Of course extrapolating this PDF above the range of  $[4.5^\circ, 6^\circ]$  is an act of faith, but it seems to accord roughly with some of the PDFs depicted in IPCC Box 10.2. The only excuse for such a crude procedure is the intense need to get some kind of a handle on the PDF of high values of climate sensitivity in order to be able to say something about the likelihood of catastrophic temperatures. Numerical values of the resulting distribution of climate sensitivity are given in Table 2. As usual, this is but an illustrative exercise and people who are dissatisfied with this procedure are free to use other parameter values or functional forms if they have other ideas for evaluating high-temperature climate sensitivities.

---

studies, see the recent review article by Knutti and Hegerl (2008)

$T$	4.5°C	5°C	6°C	6.5°C	7°C	8°C	9°C	10°C	12°C	15°C	20°C
$P[S \geq T]$	16.0%	9.6%	4.0%	2.7%	1.9%	1.0%	.57%	.34%	.14%	.05%	.01%

**Table 2:** Climate sensitivity probabilities.  $P[S \geq T] = kT^{-\gamma}$  with  $\gamma = 4.819$ ,  $k = 224.8$ .

My next task is to translate the PDF (2) into confidence levels for temperature change as a function of postulated GHG levels. Carbon dioxide is not the only GHG, although it is by far the most important. Additional anthropogenically induced GHGs include CH<sub>4</sub>, NO<sub>2</sub>, and several others, which can be very roughly translated via “global warming potentials” into “carbon dioxide equivalent” (CO<sub>2e</sub>) by using various conversion factors. Let  $G$  stand for some given stabilized target level of atmospheric GHGs expressed in ppm of CO<sub>2e</sub>. Long run average global surface temperature changes are proportional to the logarithm of  $G$ . Let  $T$  be the equilibrium temperature change corresponding to given values of  $G$  and  $S$ . Then  $T = S \ln(G/280)/\ln 2$ , and the PDF of temperature change as a function of  $G$  corresponding to (2) is given by the formula

$$P[T \geq \hat{T}] = P \left[ S \geq \frac{\ln 2}{\ln(G/280)} \hat{T} \right] = \left[ k \left( \frac{\ln 2}{\ln(G/280)} \right)^{-\gamma} \right] \hat{T}^{-\gamma}. \quad (3)$$

Equation (3) can be inverted to yield

$$\hat{T}(P, G) = \frac{\ln(G/280)}{\ln 2} \left( \frac{k}{P} \right)^{1/\gamma}. \quad (4)$$

Some numerical values calculated from formula (4) are given in Table 3. These are expressed in terms of “confidence levels” (equal to 100(1-P)%) for long-run equilibrium temperatures being *less* than the indicated value.

CO <sub>2e</sub> =	450 ppm	500 ppm	550 ppm	600 ppm	650 ppm	700 ppm	800 ppm
<i>confidence</i>	—————	—————	—————	—————	—————	—————	—————
<b>50%</b>	T<2.1°	T<2.5°	T<2.9°	T<3.3°	T<3.6°	T<4.0°	T<4.5°
<b>75%</b>	T<2.8°	T<3.4°	T<4.0°	T<4.5°	T<4.9°	T<5.4°	T<6.2°
<b>90%</b>	T<3.4°	T<4.1°	T<4.8°	T<5.5°	T<6.0°	T<6.6°	T<7.5°
<b>95%</b>	T<3.9°	T<4.8°	T<5.6°	T<6.3°	T<7.0°	T<7.6°	T<8.7°
<b>99%</b>	T<5.5°	T<6.7°	T<7.8°	T<8.8°	T<9.7°	T<10.6°	T<12.1°
<b>99.9%</b>	T<8.8°	T<10.8°	T<12.6°	T<14.2°	T<15.7°	T<17.1°	T<19.5°

**Table 3:** Confidence levels for long-run equilibrium temperature change as a function of CO<sub>2</sub>e stabilization levels.

I think that the numbers in Table 3 speak for themselves. We are currently at 438 ppm of CO<sub>2</sub>e and climbing rapidly. At the median confidence level of 50%, temperature changes are perhaps not so threatening for low CO<sub>2</sub>e stabilization targets. But the confidence levels for excluding dangerously high temperatures are disturbing even for the lower CO<sub>2</sub>e stabilization levels, and they increase at a disturbing rate for the higher CO<sub>2</sub>e stabilization levels.

Again I emphasize that the approach I am taking here is subject to a large number of caveats and can be challenged at every step along the way. My only intention is to demonstrate that the possibility of catastrophic climate change does not appear to be an ignorable event. Quite the contrary, in the light of Table 3 it appears to me on intuitive grounds that it would be curious if a CBA of climate change were not dependent on what is assumed about the PDF of rare high-temperature extremes. We simply do not know what are the appropriate confidence levels for higher temperatures, but some quick and dirty calculations here, which are not totally implausible, indicate some disturbing tendencies.

## 5 Atmospheric Carbon Projections and Persistence

The first thing I am after in this section is some extremely rough sense of how long it might take to reach the CO<sub>2</sub>e levels enumerated in Table 3. Since predictions of atmospheric concentrations of CO<sub>2</sub> and CO<sub>2</sub>e are obviously very important to any analysis of climate change, a great many simulation models have been used to project future stock levels (and temperatures) under various policies. These dynamic models are of varying complexity, but they all are pretty complicated. In keeping with the methodological spirit of this paper, I want to project atmospheric concentrations of CO<sub>2</sub> and CO<sub>2</sub>e under a super-simplistic formula that is reproducible, yet hopefully sufficiently accurate for the rough purposes at hand. Atmospheric CO<sub>2</sub> and CO<sub>2</sub>e levels are subject to a myriad of assumptions (about emissions, the strength of carbon sinks, feedbacks, and so forth and so on), many of which will probably change substantially in the future. Here I just fit a simple not-implausible reduced-form function that predicts well for the *last* 60 post-WWII years, and hopefully gives a decent projection for the *next* 60 years or so. By going right to the reduced form of the variable of central interest (here atmospheric levels of GHGs), I am compressing a very large number of intermediate steps; but since many of these intermediate steps are uncertain anyway, and will be even more so in the future, I am not sure just how much is being missed. The time path I am attempting to project tries to emulate something akin to a “business

as usual” (BAU) scenario. Needless to say, this is one more super-simplistic approach that could readily be challenged, but at least it has the virtue of making it relatively easy for anyone else to use their own favorite projection approach.

Let  $CO_2(t)$  be actual known  $CO_2(t)$  concentrations in ppm for year  $t$ , while  $CO_2(t')$  stands for  $CO_2$  concentrations in ppm at future year  $t' > t$ , as projected from year  $t$  *given known*  $CO_2(t)$  at time  $t$ . To make a long story short, and without introducing too many digressions, the simple formula that I have found works better than any other for predicting future atmospheric  $CO_2$  concentrations over the past 60 years is:

$$CO_2(t') - 280 = [CO_2(t) - 280] \exp(.02(t' - t)). \quad (5)$$

Formula (5) just makes the excess of atmospheric  $CO_2$  over pre-industrial-revolution levels of 280 ppm grow at the rate of 2% per year, which extremely roughly mirrors the average growth rate of emissions over the past 40 years or so and is a number sometimes used for rough projections of future emissions growth. Formula (5) is at best a simple back-of-the-envelope approximation that seems to work tolerably well for the post-war years. There can be no pretence that (5) gives anything but a really crude ballpark estimate of future GHG concentrations under a BAU scenario. The first column of Table 1 specifies decades past and futures. The *past* 60 years of  $CO_2$  data in column 2 of Table 4 is high quality, taken from the Mauna Loa Keeling observations from 1959-2009<sup>10</sup> and from ice core data for 1949.<sup>11</sup> The third column of Table 4 *predicts* what would be the 2009 level of  $CO_2$  on the basis of the then-actual level of  $CO_2$  by using formula (5), while the comparison *actual* level of 2009  $CO_2$  is given in column 4. The fifth column of Table 4 displays future  $CO_2$  projections by decade, based on formula (5), starting from the  $t=2009$  level of  $CO_2=387.6$  ppm. In the fifth column, future  $CO_2e$  levels are ultra-simplistically calculated by using the conversion multiplier  $CO_2e(t')/CO_2(t') = CO_2e(2009)/CO_2(2009) = 1.13$ , which, under a charitable interpretation, has held extremely roughly in the past. (The extra 13% is small enough that I don't think it much matters how it is handled.) I would consider the last two columns of Table 4 to be the most speculative numbers in the paper.

---

<sup>10</sup>[ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2\\_annmean\\_mlo.txt](ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2_annmean_mlo.txt); 2009 from monthly data, same source.

<sup>11</sup><http://cdiac.esd.ornl.gov/ftp/trends/co2/lawdome.smoothed.yr20>

year	actual CO <sub>2</sub>	predicted 2009	actual 2009	future CO <sub>2</sub>	future CO <sub>2</sub> e
-60=1949	310.5 ppm	381.3 ppm	387.6 ppm		
-50=1959	316.0 ppm	377.9 ppm	387.6 ppm		
-40=1969	324.6 ppm	379.3 ppm	387.6 ppm		
-30=1979	336.8 ppm	383.5 ppm	387.6 ppm		
-20=1989	352.9 ppm	388.8 ppm	387.6 ppm		
-10=1999	368.1 ppm	387.6 ppm	387.6 ppm		
0=2009	387.6 ppm	387.6 ppm	387.6 ppm	387.6 ppm	438.0 ppm
+10=2019				411.4 ppm	464.9 ppm
+20=2029				440.5 ppm	497.8 ppm
+30=2039				476.1 ppm	538.0 ppm
+40=2049				519.5 ppm	587.0 ppm
+50=2059				572.5 ppm	646.9 ppm
+60=2069				637.2 ppm	720.0 ppm
+70=2079				716.3 ppm	809.4 ppm
+80=2089				812.9 ppm	918.6 ppm

**Table 4:** Past and projected BAU levels of CO<sub>2</sub> and CO<sub>2</sub>e.

The projections in the last two columns of Table 4 are only suggestive at best. They are extremely uncertain, but they do suggest that CO<sub>2</sub> and CO<sub>2</sub>e trajectories in a BAU scenario might climb rapidly. If nothing much is done to stabilize GHG accumulations before some 70 years or so, the uncertain equilibrium temperatures that may subsequently materialize after that time might acquire the confidence intervals of the last two columns of Table 3. So the combined message from the extraordinarily crude calculations behind Tables 3 and 4 appears to be that the possibility of catastrophic global warming is not a negligible consideration in climate change policy. Much more detailed dynamic models have serious uncertainties of their own, but several of them suggest a similar BAU scenario to what is being conveyed by Tables 3 and 4.<sup>12</sup> Of course these are BAU scenarios. I turn next to what might be accomplished by a very simplistic policy exercise of stopping *all* emissions at some future time. This is followed by a discussion of the feasibility of making “wait and see” midcourse corrections after it has been learned that the planet might be on a bad temperature trajectory.

<sup>12</sup>See, e.g., the recent study of Sokolov et al (2009), which predicts a BAU median CO<sub>2</sub> level of  $\approx 866$  ppm for  $\approx 2095$  with accompanying BAU median temperature change of  $\approx 5.2^\circ\text{C}$ . The latest (2008) projections from the U.K. Met Office Hadley Center predict a BAU median CO<sub>2</sub> level of  $\approx 850$  ppm for  $\approx 2095$  with accompanying BAU median temperature change of  $\approx 5.5^\circ\text{C}$

It is not sufficiently appreciated (or sometimes not even understood) that merely stabilizing GHG flow emissions at some future time will *not* stabilize GHG stock accumulations at the level of that time. To stabilize GHG atmospheric stocks at the level of any given time requires a *very strong decline* in GHG flow emissions during and after that time. This is a critical scientific fact that the public, many economists, and most policy makers do not often take into account.

Atmospheric CO<sub>2</sub>, which is by far the most important anthropogenically-injected GHG, decays over time as a complicated weighted average of several exponential functions, many of which are also fraught with uncertainty because the absorptive power of future sinks is not fully known. In any event, the lifetime of atmospheric CO<sub>2</sub> perturbations has a very long upper tail extending for thousands of years.<sup>13</sup> This extremely important feature is typically ignored in most integrated assessment models (IAMs) of the economic effects of climate change. It means that the widespread idea that we have significant breathing room derived from “wait and see” midcourse corrections is questionable.

In an important recent study<sup>14</sup>, it was calculated how concentrations of CO<sub>2</sub> would be expected to fall off over time if *all* anthropogenic emissions were to cease immediately, following a future 2% annual growth rate of emissions (the same growth rate I used in a previous projection) 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% of the peak enhancement level over the preindustrial level of 280 ppm persevered after 100 years of zero emissions, while approximately 40% of the peak enhancement level over the preindustrial level of 280 ppm persevered after 1,000 years of zero emissions. This means, for example, that if atmospheric CO<sub>2</sub> concentrations were to peak at 800 ppm, followed forever thereafter by *zero* emissions, then atmospheric concentrations would be  $\approx 650$  ppm after 100 years. After 1,000 years of *zero* emissions subsequent to a peak of 800 ppm, CO<sub>2</sub> atmospheric concentrations would come down to  $\approx 500$  ppm. While there may be some role for reactive policies after learning, the above kinds of numbers do not suggest that humanity has significant room for maneuver by using “wait and see” midcourse corrections subsequent to observing unexpectedly bad climate changes. Perhaps more to the point for this paper, I think that the above numerical examples of irreversibility are yet one more demonstration that the possibility of catastrophic

---

<sup>13</sup>For a recent survey of the the implications of the long fat tail and the relevant science behind it, see Archer et al (2009).

<sup>14</sup>Solomon et al (2009).

climate change does not appear to be an ignorable event. And, yet again, it seems to me on intuitive grounds that it would be curious if a CBA of climate change was not much dependent on what is assumed about the structure of rare high-temperature extremes.

## 6 Discounting Under Extreme Uncertainty

In the realm of climate change, there is no more contentious issue dividing economists from most of the rest of society than discounting. At a 4% annual interest rate, \$1 of damages incurred 150 years from now is worth  $\frac{1}{4}$ ¢ today. This effectively trivializes distant-future climate damages, almost no matter how large they are. It tends to infuriate scientists and many others. Maybe it might be more politic for an economic analysis to present information in an undiscounted form (perhaps akin to Tables 1-4 of this paper?), and then leave the public to make up their own minds about what to do – without pouring gasoline on the fire by attempting to discount at market rates of return.

Discounting for distant-future climate change is a very controversial subject that has been much discussed in the literature. Here I limit myself to making one simple point that is often overlooked. This approach begins by recognizing that the most consequential uncertainty in the economic evaluation of climate change is arguably the uncertainty concerning which interest rate to use for discounting.

Under uncertainty about which interest rate to use for discounting, it is not discount *rates* that should be averaged, but discount *factors*. This causes the relevant time-dependent *effective* discount rate to decline over time, approaching asymptotically its lowest possible value.<sup>15</sup>

In what follows next, for the sake of simplicity I divide the population into two sub-populations: hard-hearted economists and soft-hearted naturalists. The hard-hearted economists believe that future costs and benefits should always be discounted at a rate of 4% per annum<sup>16</sup>, so that the relevant discount factor as a function of future time  $t$  is  $A(t) = \exp(-.04t)$ . The soft-hearted naturalists think that there is something special about the nature of climate change that calls for zero discounting and a discount factor of  $A(t) = 1$ . If the soft-hearted naturalists tried to convince the hard-hearted economists, they might employ language like there being no deep underlying reason of generic principle why future

---

<sup>15</sup>This principle was applied in Weitzman (1998), and was recently modeled formally by Gollier (2009). The background setup is that a CBA decision must be made today, before the future discount rate is known. Then the discount rate is drawn tomorrow from the relevant PDF and remains in force over the entire future, during which time the economy adjusts accordingly. This is a crude representation for getting at the idea that distant future discount rates are unknown, but likely to be highly serially correlated.

<sup>16</sup>This is the average discount rate used, e.g., in Nordhaus (2008).

interest rates can be extrapolated from past average values (especially to the distant future), or the idea that catastrophic climate change could cause negative growth rates and low productivity of capital, or that the risk-free real rate of interest is close to zero, or that the modified Pareto criterion is compromised because intergenerational payments cannot be enforced, or that the low elasticity of substitution between environmental amenities and ordinary economic goods turns the terms of trade between the two towards the environment “as if” the relevant discount rate is very low. There are yet some other stories that might be told, for what they are worth.

Suppose that the fraction of hard-hearted economists in the general population is  $\lambda$ , while the fraction of soft-hearted naturalists is  $1 - \lambda$ . Suppose that a democratic leader does not know which position is correct but, akin to the idea of “one person one vote,” believes that the probability of either group being right is proportional to their numbers in the general population. This gives some voice to the beliefs of both groups, and implies that the relevant discount factor is

$$A(t; \lambda) = \lambda \exp(-.04 \times t) + [1 - \lambda] \exp(-0 \times t). \quad (6)$$

In Table 5 below are calculated from formula (6) some values of  $A(t)$  for various values of  $\lambda$ . (The hard-hearted economist’s scenario corresponds to  $\lambda=100\%$ .)

<b>years ahead</b> $\mapsto$	<b>t=50</b>	<b>t=100</b>	<b>t=150</b>	<b>t=200</b>	<b>t=250</b>
$\lambda = 100\%$	A=13.5%	A=1.8%	A=.25%	A=.03%	A=.006%
$\lambda = 95\%$	A=17.9%	A=6.7%	A=5.2%	A=5.0%	A=5.0%
$\lambda = 90\%$	A=22.2%	A=11.6%	A=10.2%	A=10.0%	A=10.0%
$\lambda = 75\%$	A=35.2%	A=26.4%	A=25.2%	A=25.0%	A=25.0%
$\lambda = 50\%$	A=56.8%	A=50.9%	A=50.1%	A=50.0%	A=50.0%

**Table 5:** Discount factors  $A(t | \lambda)$  as a function of  $t$  and  $\lambda$ .

Note how much the discount factor for the  $\lambda = 100\%$  case stands out because it is so extremely much lower for the distant future than the discount factor for the other values of  $\lambda$ .

Another way of making the same point is by considering what would happen if the unknown interest rate for discounting climate change is normally distributed with mean  $\mu=4\%$  per annum and with standard deviation  $\sigma=2\%$ . Under this specification, the probability that the relevant future discount rate is less than zero is (only) .027 yet, in the long run, the

small possibility of negative discount rates will eventually dominate the analysis by reversing the earlier decline of  $A(t)$ . It is readily confirmed that the discount factor applicable for  $t$  years into the future is then lognormally distributed with expected value

$$A(t) = \exp(-.04t + .0002t^2). \tag{7}$$

In Table 6 below are calculated from formula (7) some values of  $A(t)$ .

<b>years ahead:</b>	<b>t=50</b>	<b>t=100</b>	<b>t=150</b>	<b>t=200</b>	<b>t=250</b>
discount factor $A(t)$	22.3%	13.5%	22.3%	100%	122%

**Table 6:** Discount factors  $A(t)$  as a function of  $t$  years ahead.

The above two examples might be considered casual or far fetched, but they do embody an important principle. The very same force of compound interest that makes costs and benefits centuries from now seem insignificant also forces us to recognize the logic that over such long time periods we should be using interest rates and discount factors that decline towards the lower end of the spectrum of possible values drawn from the relevant PDF. As always, others are invited to play with different formulations and different numbers.

Concerning the subject of how to combine environmental amenities with economic output into a single utility function, here is yet another series of controversial structural issues. Should the combining be additively separable, or multiplicatively separable, or something in between? What is the relevant degree of risk aversion for catastrophic outcomes with a low probability? Is there a role for something resembling the statistical value of human civilization or life on earth as we know it, and if so what is its value? Is expected utility the right way to frame an analysis when the low-probability high-impact events are so ambiguous and difficult to quantify? At the end of the day are we in a situation where it might be more politic to provide information in a form roughly akin to Tables 1-4 and allow the public to draw its own conclusions, without yet introducing controversial issues of discounting and utility functions? And so forth and so on. I do not know what is the right procedure here. But the deep structural uncertainties about how to aggregate preferences over low-probability high-impact events, which will occur (if they occur at all) centuries from now, might make it seem curious if a CBA of climate change was not much dependent on what is being assumed about the overall economic evaluation of such rare high-temperature extremes.

I leave aside the issue of what it might cost to research and put into commercial use new and unproven technologies for limiting GHG emissions – or what might happen to climate change if such R&D efforts fail. This is yet another example of tail-fattening extreme uncertainty about extreme outcomes.

## 7 Some Possible Implications for Climate-Change CBA

Perhaps not surprisingly, there is controversy about the implications of fat tails for CBA. My target is not CBA in general, but the particular impression of precision inadvertently conveyed by some CBAs of climate change. I like to think I occupy a middle ground between two extreme positions. An economist does not want to abandon lightly the ideal that CBA should bring independent empirical discipline to any application by being based upon empirically reasonable functional forms and parameter values. Even when fat-tailed logic might apply, climate-change CBA could in principle reveal useful information about whether fat tails are or are not actually relevant for “reasonable” functional forms and parameter values at extreme temperatures. (What “reasonable” means in a context of extreme impacts with uncertain probabilities may not be clear, which in practice can introduce a large gray area into CBA of climate-change catastrophes.) Simultaneously, one does not want to be obtuse by insisting that the extreme uncertainty about extreme climate change behind the fat tails makes no practical difference for climate-change CBA because the parameters just need to be empirically determined and then simply plugged into the analysis. Some sort of a tricky balance is required between being overwhelmed by the deep structural uncertainty about extremes into a Hamlet-like paralysis that leads to abandoning CBA altogether, and being underwhelmed into insisting that it is just another empirical issue to be sorted out by business-as-usual CBA. Economists should, of course, remain open to changing their beliefs on the basis of robust outcomes from well-designed CBAs. One should go ahead and plug into climate-change CBA simulations the PDFs of the uncertain science and the PDFs of the uncertain economic valuations, including, especially, various functional forms – and then respect robust conclusions. On the other hand, I think one should not be especially surprised if outcomes are fragile to specifications concerning catastrophic extremes. There is *some* difference in expected welfare between 550 ppm of CO<sub>2</sub> and 650 ppm of CO<sub>2</sub>, and that difference could be important for policy. But *how much* of a difference may be difficult to extract with reasonable precision from a fat-tailed CBA that is sensitive to seemingly obscure modeling assumptions about climate extremes.

I think that economists would be well advised to convey that not all CBAs are created equal and not all of them should be defended with equal vigor. The public and policy makers need to have more of a sense about which CBAs are more precise than which others. A CBA for constructing or repairing a highway with known technology and known subsoil characteristics is relatively straightforward. Climate change is at the opposite end of the spectrum in being almost unique because deep structural uncertainties about distant future catastrophes are being combined with essentially unlimited downside liability for planetary welfare. This

lends an unusually strong subjective element to climate-change CBA, especially concerning extreme outcomes, which is just not there for more customary CBAs.

The standard packaging of a climate-change CBA appears to offer a constructive ongoing scientific-economic research program for generating ever more precise outputs from ever more precise inputs, although this is not exactly what has happened thus far. By contrast, my main message can seem threatening because it can be painted as anti-scientific and anti-economic. Fat tails and the implied limitations that prevent CBA from reaching robust conclusions are frustrating for economists. After all, we economists make a living from plugging rough numbers into simple models and reaching specific conclusions (more or less) on the basis of these numbers.

It seems plausible to presume that a serious consideration of possible disasters should *qualitatively* move effort in the direction of avoiding or at least minimizing them. But what are we supposed to advise policy makers and politicians *quantitatively* about *how much* effort to spend on averting climate change if conclusions from modeling fat-tailed uncertainties are not so 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 place for even a whiff of fuzziness from two-handed economists. It is threatening for us economists to admit that constructive “can do” climate-change CBA may be up against some limitations on the ability of quantitative analysis to give robust advice. But if this is the way things are with the economics of climate change, then this is the way things are – and non-robustness to subjective assumptions 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 CBA.

In my opinion, economists need to emphasize more openly to the policy makers, the politicians, and the public that, while formal climate-change CBA can be helpful, and it is *always* a good idea to play with functional forms and plug in some numbers, there is a danger of possible overconfidence from undue reliance on subjective judgements about the probabilities and welfare impacts of extreme events when these are an important consideration. What we can do constructively as economists is to explain better the magnitudes of the unprecedented structural uncertainties that are involved in climate change, explain why this feature limits what we can say, and present the best CBAs 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 render clear robust answers. The moral of the story I am trying to tell here is that, under extreme uncertainty about climate-change extremes, it is plausible that seemingly casual decisions about functional forms, parameter values, and tail fatness may dominate CBA. 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 say by effectively marginalizing the very possibilities that make climate change so grave in the first place.

The ideas I am propounding here do not necessarily support the catastrophist view that we are inevitably heading for a meltdown unless super-radical changes are initiated almost immediately. It is important to bear in mind that the ultra-extreme catastrophes are unlikely to ever materialize. That is what a low probability means. Ex post, the world dodged a bullet in the Cuban missile crisis. But were we right to be concerned at that time? Would an ex ante CBA-like analysis of the Cuban missile crisis have given policy advice that might have been especially sensitive to assumptions about the unknowable probabilities and disutilities of atomic war? I think we were right at that time to be very concerned about how to avoid a low-probability extreme-impact situation whose structure is highly uncertain – and I think we are right now to be very concerned about how to avoid a low-probability extreme-impact situation whose structure is highly uncertain.

## 8 Possible Implications for Climate-Change Policy

In situations of potentially unlimited exposure to downside damages, like climate change, it might be appropriate to emphasize a slightly better treatment of the worst-case fat-tail extremes – and what might be done about them, at what cost – relative to refining the calibration of most-likely outcomes or rehashing point estimates of discount rates (or climate sensitivity). A clear implication of this paper is that greater research effort is relatively ineffectual when targeted at estimating central tendencies of what we already know relatively well about the economics of climate change in the more-plausible scenarios. A much more fruitful goal of research might be to aim at understanding even slightly better the deep uncertainties concerning the *less* plausible scenarios located in the bad fat tail. (Alas, the tails are the very part of a PDF that is most difficult to learn, presenting yet another climate-change dilemma.) I also believe that an important complementary research agenda, which stems naturally from the analysis of this paper, is the crying need to comprehend much better *all* of the options for possibly dealing with high-impact climate-change extremes, without trying to pre-censor any of them as socially unacceptable or politically incorrect.

Instead of the existing research emphasis on estimating or simulating economic impacts of the more plausible climate-change scenarios, to at least compensate partially for finite-

sample bias the model of this paper calls for a dramatic oversampling of those stratified climate-change scenarios associated with the most adverse imaginable economic impacts in the bad fat tail. With limited sampling resources, Monte Carlo analysis could perhaps be used more creatively – not necessarily to defend a specific policy result, but to experiment seriously in order to find out more about what happens with fat-tailed uncertainty and significant high-temperature damages in the limit as the grid size and number of runs increase simultaneously. Of course an emphasis on sampling climate-change scenarios in proportion to welfare-weighted probabilities of occurrence forces us to estimate subjective probabilities down to extraordinarily tiny levels and also to put degree-of-devastation weights on disasters with damage impacts up to perhaps being welfare-equivalent to losing 99% (or possibly even more) of consumption – but that is the price we must be willing to pay for having a genuine economic analysis of potentially-catastrophic climate change.

Another suggested research priority concerns the role of “wait and see” midcourse corrections subsequent to observing unexpectedly bad climate changes. We have already seen that the extremely long lifetime of atmospheric CO<sub>2</sub> stocks suggests that there is not a lot of room for maneuver. But still, there might be some scope for learning and reacting. It would appear to be an important research priority to quantify better how much time and how much capability to avert climate change disasters can realistically be provided after filtering out climate change signals from a noisy learning process.

When analyzing the economics of climate change, perhaps it might be possible to make back-of-the-envelope comparisons with empirical probabilities and mitigation costs for extreme events in the insurance industry. One might try to compare numbers on, say, a homeowner buying fire insurance (or buying fire-protection devices, or a young adult purchasing life insurance, or others purchasing flood-insurance plans) with cost-benefit guesstimates of the world buying an insurance policy going some way towards mitigating the extreme high-temperature possibilities. On a U.S. national level, rough comparisons could perhaps be made with the potentially-huge payoffs, small probabilities, and significant costs involved in countering terrorism, building anti-ballistic missile shields, or neutralizing hostile dictatorships possibly harboring weapons of mass destruction. A crude natural metric for calibrating cost estimates of climate-change environmental-insurance policies might be that the U.S. already spends approximately  $2\frac{1}{2}\%$  of national income on the cost of a clean environment.<sup>17</sup> All of this having been said, the bind we find ourselves in now on climate change seems to be characterized by extremely slow convergence of inductive knowledge towards resolving the fat-tailed uncertainties – relative to the lags and irreversibilities from not acting before

---

<sup>17</sup>U.S. Environmental Protection Agency (1990), executive summary projections for 2000, which I updated and extrapolated to 2007.

structure is more fully identified.

Another policy implication concerns the need to compare fat-tailed problems with possibly fat-tailed solutions. There is no sense substituting one low-probability high-impact catastrophe with a “solution” that has comparable probabilities of equally bad potential impacts. Thus, one wants to think seriously about the likelihood and magnitude of worst-case scenarios with expanded nuclear power, carbon capture and storage, or geoengineering, to name but a few options. Are these and other proposed remedies really less risky than the risk of climate change that they are purporting to diminish? My inclination is to answer yes (they are less risky) to all of the above, but it would be consoling to have a more firm foundation for or against such a belief.

## 9 A Portfolio Niche for Fast Geoengineering?

The case for decarbonizing the atmosphere is traditionally made without explicit reference to extreme tail behavior. This argument is presumably strengthened, perhaps considerably, when extreme tail events are considered. Having said this, there is more to say. The fat tails introduce some distinctive issues of their own. Responsible economic analysis of fat tails implies some tolerance for at least considering extreme-sounding proposals that are not normally placed on the policy table for discussion. One consequence of fat-tailed logic might concern the role of fast-acting planetary geoengineering. The opinion that follows might be construed as editorializing, but it seems to me that the thrust of this paper leads logically to a narrowly-defined niche role for a reliable backstop technology that can effectively knock down high planetary temperatures quickly in case of emergency.

What I mean by “fast geoengineering” is any action having the possibility to lower global temperatures quickly – within decades or even years. Practically, at this time fast geoengineering means albedo enhancement by injecting sunlight-reflective particulates or aerosols, such as sulfur dioxide precursors, into the stratosphere. I do not touch upon the science of fast geoengineering, and even tread lightly upon the economics.<sup>18</sup> My main focus here is on the narrow question of whether the arguments of this paper support a special niche role for fast geoengineering – as one important option in balanced portfolio of global warming strategies and responses. I think the answer is a qualified yes.

The analysis of this paper is suggesting that a significant component of the overall expected damages of climate change may be located in the fat upper tail of the temperature distribution. Cut out the fat upper tail, and you have cut out a major part of the expected

---

<sup>18</sup>Some of the science is reviewed in Rasch et al (2008). The idea of fast geoengineering has been around for a long time, but was recently given much visibility by the influential article of Crutzen (2006).

disutility of global warming, goes the argument. According to this logic, a large increase in expected welfare might be gained if some relatively benign form of fast geoengineering were deployed in readiness to rapidly derail severe greenhouse heating – should this contingency materialize. Because of the largely-irreversible long pipeline commitment of atmospheric CO<sub>2</sub>, this argument might hold even though higher temperatures tend to materialize later and the “emergency” might unfold over a time scale of centuries.

Fast geoengineering seems quite risky, if for no other reason than the law of unintended consequences, and it cannot ward off all of the bad effects of high atmospheric CO<sub>2</sub>, such as ocean acidification. However, to say that fast geoengineering does not now look like a panacea for all the effects of climate and atmosphere changes should not be to prejudge now that it may not have a very important, perhaps even crucial, future role to play in a balanced portfolio of responsible climate-change policies. Even if fast geoengineering (Plan B) is not a replacement for curtailing GHG emissions (Plan A) – because it is too risky to be used as a mainline defence – it might still be critical to have a Plan-B option in reserve to buy more time for Plan A. The crude analysis of this paper formalizes a possibly large potential welfare gain from having the capability to slim down quickly a bad fat global-warming tail during a worst-case emergency. In my opinion, this appears to be a legitimate argument for a well funded Plan-B research program, undertaken now, which might include pilot studies and small-scale field testing. The purpose would be to determine the feasibility, environmental side effects, and cost-effectiveness of responsible geoengineering preparedness – whose intended use is as a state-contingent option giving the ability to respond rapidly to a bad future realization of global-warming uncertainty – in order to buy time. We don’t want to wait for a truck to come barreling down the road before thinking about how it might be stopped, or at least slowed.

A huge issue with fast geoengineering is that, as an externality, it has diametrically opposite cost properties from curtailing emissions of GHGs. For me, the two really inconvenient truths about climate change are: 1) CO<sub>2</sub> abatement is costly; (2) fast geoengineering is cheap. Like it or not, whether it is a panacea or not, whether it lulls the public into a false sense of security that undermines legitimate Plan-A GHG-curtailment strategies or not, the incredible economics of geoengineering is simply not ignorable.<sup>19</sup> The fast geoengineering option currently looks so unbelievably inexpensive as a quick fix for extreme temperature changes that virtually any middle-power developed country might be tempted to implement it unilaterally. For me this means that – as well as there being a strong policy argument that *now* is the time to learn a lot more about fast geoengineering – there is an additional strong

---

<sup>19</sup>Barrett (2008) contains an excellent discussion of some implications of what he has dubbed “the incredible economics of geoengineering.”

policy argument that *now* is also the time to start thinking seriously about an international framework governing the use of this scary option.

## 10 Concluding Comments

Through informal reasoning, elementary examples, and simple numerical exercises, I have tried in this paper to convey an informal and relatively non-technical overview of some of the background uncertainties behind extreme climate change. The emphasis throughout has been on simplicity and transparency. I derived some closed form solutions from ultra-simplistic formulations, which has the advantage of making clear what is the core structure and where the results are coming from. However, these are at best suggestive examples. Every number and every formulation is simple-minded, is based in large part upon subjective thought-experimental guesstimates, is subject to a myriad of caveats, and is open to strong criticisms as well as alternative interpretations.

Alas, a distressingly large element of subjectivity seems unavoidable for extreme climate change, because we are extrapolating so far outside the realm of ordinary experience that we have little to go on. My only excuse has been that I do not know a better approach for highlighting what is potentially a very important issue, and that my methodology allows anyone else to easily plug in alternative numbers or assumptions and come up with their own extremely crude ballpark estimates. There is no doubt that other specifications, equally or perhaps even more plausible, can arrive at different conclusions than those expressed here. However, in a highly uncertain situation where we are not really sure which one is the more plausible of two outcomes with highly asymmetric consequences, I think that the possibility of catastrophic climate change needs to be taken seriously.

The thrust of this paper has been to argue that the bad tails of the relevant probability distributions should not be ignored because they are likely to be fat with probability and play an important role. The overarching theme is that it would seem to be curious if a CBA of climate change was not much dependent on what is assumed about the extreme uncertainty of extreme climate change.

## References

- [1] **Allen, Myles R. and David J. Frame.** “Call Off the Quest” *Science*, 2007 (October 26), 318, pp. 582-583.

- [2] **Barrett, Scott.** “The Incredible Economics of Geoengineering.” *Environmental and Resource Economics*, 2008 (January), 39, pp. 45-54.
- [3] **Crutzen, Paul J.** “Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma?” *Climatic Change*, 2006 (August), 77, pp. 211-220.
- [4] **Dasgupta, Partha.** “Commentary: the Stern Review’s Economics of Climate Change.” *National Institute Economic Review*, 2007, 199, pp. 4-7.
- [5] **James Hansen, Makiko Sato, Puskher Karecha, David Beerling, Valerie Masson-Delmotte, Mark Pagani, Maureen Raymo, Dana L. Royer, James C. Zchos.** “Target Atmospheric CO<sub>2</sub>: Where Should Humanity Aim?” *Open Atmos. Sci. J.*, 2, 217-231.
- [6] **Dieter Lüthi, Martine Le Floch, Bernhard Bereiter, Thomas Blunier, Jean-Marc Barnola, Urs Siegenthaler, Dominique Raynaud, Jean Jouzel, Hubertus Fischer, Kenji Kawamura & Thomas F. Stocker.** “High-resolution carbon dioxide concentration record 650,000–800,000 years before present.” *Nature*, 453, 379-382 (15 May, 2008).
- [7] **Gollier, Christian.** “Expected Net Present Value, Expected Net Future Value, and the Ramsey Rule.” CESifo Working Paper No. 2643, May 2009.
- [8] **IPCC-AR4.** *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge University Press, 2007 (available online at <http://www.ipcc.ch>).
- [9] **Knutti, Reto, and Gabriele C. Hegerl.** “The equilibrium sensitivity of the Earth’s temperature to radiation changes.” *Nature geoscience*, November 2008, 1, pp. 735-438.
- [10] **Elmar Kriegler, Jim W. Hall, Hermann Held, Richard Dawson, Hans Joachim Schellnhuber.** “Imprecise probability assessment of tipping points in the climate system.” *Proceedings of the National Academy of Sciences*, 106, 5041-5046 (31 March, 2009).
- [11] **Matthews, H. Damon, and David W. Keith.** “Carbon-cycle feedbacks increase the likelihood of a warmer future.” *Geophysical Research Letters*, 2007, 34, L09702.
- [12] **Nordhaus, William D.** *A Question of Balance.* Yale University Press, 2008.

- [13] **Rasch, Philip J., S. Tilmes, R. P. Turco, A Robock, L. Oman, C-C Chen, G. L. Stenchikov, and R. R. Garcia.** “An Overview of Geoengineering of Climate using Stratospheric Sulfate Aerosols.” Mimeo, 2008.
- [14] **A.P. Sokolov, P.H. Stone, C.E. Forest, R. Prinn, M.C. Sarofim, M. Webster, S. Paltsev, C.A. Schlosser, D. Kicklighter, S. Dutkiewicz, J. Reilly, C. Wang, B Felzer, H.D. Jacoby.** “Probabilistic forecast for 21st century climate based on uncertainties in emissions (without policy) and climate parameters.” *Journal of Climate*, preprint, 2009: <http://ams.allenpress.com/perlserv/?request=get-abstract&doi=10.1175%2F2009JCLI2863.1>
- [15] **Sterner, Thomas, and U. Martin Persson.** “An Even Sterner Review: Introducing Relative Prices into the Discounting Debate.” *Review of Environmental Economics and Policy*, 2008 (Winter), 2(1), pp. 61-76.
- [16] **U.S. Environmental Protection Agency.** *Environmental Investments: The Cost of a Clean Environment*. Washington, DC: U.S. Government Printing Office, 1990.
- [17] **Weitzman, Martin L.** “Why the Far-Distant Future Should be Discounted at its Lowest Possible Rate.” *Journal of Environmental Economics and Management* , 36, pp. 201-208, (1998).
- [18] **Weitzman, Martin L.** “On Modeling and Interpreting the Economics of Catastrophic Climate Change.” *Review of Economics and Statistics*, February 2009.