# Some Basic Economics of Extreme Climate Change

Martin L. Weitzman

February 19, 2009

**Abstract**

Climate change is characterized by deep structural uncertainty in the science coupled with an economic inability to evaluate meaningfully the welfare losses from high temperature changes. The probability of a disastrous collapse of planetary welfare from too much $CO_2$ is non-negligible, even if this low probability is not objectively knowable. This paper attempts to explain (in not excessively technical language) some of the most basic issues in modeling the economics of catastrophic climate change. The paper builds to a tentative conclusion that, no matter what else is done realistically to slow $CO_2$ buildups, economic analysis lends some support to undertaking serious research now into the prospects of "fast geoengineering preparedness" – as a state-contingent emergency option offering at least the possibility of knocking down catastrophic temperatures rapidly.

## 1 Introduction

Four big questions often asked about climate change are: (1) *how much* global warming and climate change will occur; (2) *how bad* will it get; (3) *when* will all this occur; (4) *what* should be done about it. This paper attempts to explain why science and economics cannot resolve these questions to anywhere near the degree of accuracy that we have come to expect from more traditional applications of cost-benefit analysis (CBA), because there is so much deep structural uncertainty associated with climate change. The "unknown unknowns" of climate change make CBA significantly more fuzzy in this arena than in more traditional applications – like constructing roads, strengthening bridges, or setting building codes in earthquake-prone zones. The paper tries to make sense of this anomalous situation and explores what might be done in terms of actionable alternatives under such fuzzy circumstances.

Climate change is so complicated, and it involves so many sides of so many different disciplines and viewpoints, that no analytically-tractable model or paper can aspire to illuminate more than a few facets of the problem. Because the problem is so complex, economists typically resort to numerical computer simulations. An "Integrated Assessment Model" (hereafter "IAM") for climate change is a multi-equation computerized model linking aggregate economic growth with simple climate dynamics in order to analyze the economic impacts of global warming. An IAM is essentially a dynamic model of an economy with a controllable GHG-driven externality of endogenous greenhouse warming. IAMs have proven themselves useful for understanding several aspects of the economics of climate change – especially in describing outcomes from a complicated interplay of the very long lags and huge inertias involved.

A key starting point for any CBA of climate change should recognize that future temperatures or damages cannot be known exactly and must be expressed as a probability density function (PDF). Yet, most existing IAMs treat central forecasts of temperatures or damages as if they were certain and then do some sensitivity analysis on parameter values. In the rare cases where an IAM formally incorporates uncertainty, it typically uses thin-tailed PDFs including, especially, truncation of PDFs at arbitrary cutoffs. (Often this truncation is more implicit than explicit because a finite discrete-point PDF is used.) What typically emerges from conventional IAM analysis is the so-called "policy ramp" of gradually tightening emissions over time. The underlying rationale of the policy ramp is to postpone pain on climate change prevention, because it is an investment whose payoff comes only in the distant future (by human, if not geological, standards). When the distant-future payoff times are considered, the rate of return on GHG mitigation is lower than the rate of return on education, health, infrastructure, or a variety of other quicker-yielding public investments. As will be explained later, policy-ramp gradualism seems quite sensitive to the functional form of the assumed disutility of high temperature changes, to how the extreme tail probabilities are specified, and to the rate of pure time preference used to discount future utilities and disutilities.

Modeling uncertain catastrophes presents some very strong challenges to economic analysis, the full implications of which have not yet been adequately confronted. Cost-benefit analysis based on expected utility (EU) theory has been applied in practice primarily to cope with uncertainty in the form of a known thin-tailed PDF. I will argue that the PDF of distant-future temperature changes is fat tailed. A thin-tailed PDF assigns a *relatively* much lower probability to rare events in the extreme tails than does a fat-tailed PDF.[1] (Even

---

[1] As I use the term, a PDF has a "fat" (or "thick" or "heavy") tail when its moment generating function (MGF) is infinite – i.e., the tail probability approaches zero *more slowly* than exponentially. The standard

though both limiting probabilities are infinitesimal, the *ratio* of a thick-tailed probability divided by a thin-tailed probability approaches infinity in the limit.) Not much thought has gone into conceptualizing or modeling what happens to EU-based CBA for fat-tailed disasters. A CBA of a situation with known thin tails, even including whatever elements of subjective arbitrariness it might otherwise contain, can at least in principle make comforting statements of the generic form: "if the PDF tails are cut off here, then EU theory will still capture and convey an accurate approximation of what is important." Such accuracy-of-approximation PDF-tail-cutoff statements, alas, do not exist in this generic sense for what in this paper I am calling "fat-tailed CBA."

Fat-tailed CBA has strong implications that have neither been recognized in the literature nor incorporated into formal CBA modeling of disasters like climate-change catastrophes. These implications raise many disturbing yet important questions, which will be dealt with somewhat speculatively in the concluding sections of this paper. Partially answered questions and speculative thoughts aside, I contend that, at least in principle, fat-tailed CBA can change conventional thin-tail-based climate change policy advice. This paper argues that it is quite possible, and even numerically plausible, that the answers to the big policy question of what to do about climate change can hinge on the issue of how the high-temperature damages and tail probabilities are conceptualized and modeled. It is true that some reasonable-looking specifications and plausible parameter values can give rise to a gradualist policy ramp. But I think it is also true that some equally (or even more) reasonable specifications and parameter values can give very different results from a gradualist policy ramp. By implication, the advice coming out of conventional thin-tailed CBAs of climate change should be treated with caution until this low-probability high-impact aspect is addressed seriously and resolved empirically in a true fat-tailed CBA.

This paper explains in non-technical language a connection among the following four basic ideas: 1) the probability distribution of future global temperature has a fat tail at its upper extreme; 2) the disutility-damage of high temperatures is sensitive to the functional form that is assumed; 3) when discounted at an uncertain rate of pure time preference that might be close to zero, the fat tails and temperature-sensitive disutilities can make expected present discounted damages very large; 4) because elevated stocks of $CO_2$ have such a long residence time in the atmosphere, and because it takes so long to learn about irreversible climate changes and to make midcourse corrections, a significant increase in expected welfare

---

example of a fat-tailed PDF is the power law (aka Pareto aka inverted polynomial) distribution, although, for example, a lognormal PDF is also fat-tailed, as is an inverted-normal or inverted-gamma or Student-*t*. By this more or less standard definition, a PDF whose MGF is finite has a "thin" tail – i.e., the tail probability approaches zero *more rapidly* than exponentially. 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.

might be obtained if the upper extremes of the fat tail could be truncated before reaching catastrophic temperatures.

The final parts of the paper concern the welfare and policy implications of coupling fat-tailed uncertainty with high disutility of extreme temperatures. Under any foreseeable technology, elevated stocks of $CO_2$ are committed to persist for a very long time in the atmospheric pipeline. And it also takes a long time to learn about looming realizations of uncertain, but largely irreversible climate changes. Thus, $CO_2$ stock inertia, along with slow learning, makes it difficult to react to unfolding disasters by throttling back $CO_2$ flow emissions in time to avert an impending catastrophe. In this kind of situation, which is akin to trying to turn around an ocean liner in time to get away from a disaster, a large increase in expected welfare might be gained if some relatively benign form of fast geoengineering were deployable as an emergency last-minute response for knocking down rapidly the bad fat tail of temperature change. Even if fast geoengineering is not a replacement for curtailing GHG emissions, because it is too risky to be used as a mainline defence and it has too many other bad consequences, the logic of this paper argues that it still might play an important niche role as an emergency-preparedness fallback component in a balanced portfolio of mixed options for dealing with climate change. The paper highlights the idea that this aspect (the high expected value in this context of being able to truncate the bad fat tail quickly) may constitute a respectable economic underpinning supporting a well-funded research program, undertaken now, to determine the feasibility, environmental side effects, and cost-effectiveness of fast geoengineering preparedness. The paper concludes that, no matter what else is done realistically (within the realm of reason) to slow $CO_2$ buildups, economic analysis lends some support to undertaking serious research now into the prospects of "fast geoengineering preparedness" – as a state-contingent emergency option offering at least the possibility of knocking down catastrophic temperatures rapidly.

## 2 Deep Structural Uncertainty About Extremes

In this section I try to make the case that standard CBAs or IAMs of climate change likely sidestep some important issues concerning improbable but extreme outcomes. I try to make this case by citing three aspects of the climate science that do not seem to be adequately covered by conventional economic analyses. While different aspects of structural uncertainty might additionally be cited, I restrict my case to these three examples, which I will call "Exhibits A, B, and C."

"Exhibit A" concerns the atmospheric level of greenhouse gases 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 ($CO_2$) and methane ($CH_4$) trapped in tiny ice-core bubbles was extended in 2008 to 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) is inferred indirectly by proxy variables, whereas this ice-core GHG data is directly observed.

The pre-industrial-revolution level of atmospheric $CO_2$ (about two centuries ago) was 280 parts per million (ppm). The ice-core data show that carbon dioxide was never outside a range between 180 and 300 ppm during the last 800,000 years, with instances above 280 ppm exceedingly rare (to the point of being almost negligible). Currently, $CO_2$ is at 385 ppm. Methane was never higher than 750 parts per billion (ppb) in 800,000 years, but now this extremely potent GHG, which is 26 times more powerful than $CO_2$, is at 1,780 ppb. Carbon-dioxide-equivalent ($CO_2$-e) GHGs are currently at 435 ppm. Even more alarming is the rate of change of GHGs, with increases in carbon dioxide never exceeding 30 ppm over any past thousand-year period, while now $CO_2$ has risen by 30 ppm in just the last 17 years.

Thus, anthropomorphic activity has elevated $CO_2$ and $CH_4$ to levels very far outside their natural range – and at a stupendously rapid rate. There is no analogue for anything like this happening in the past geological record. Therefore, we do not really know with much confidence what will happen next. The link between GHG levels and temperature change in the ice-core record is not unicausal, and it is not fully understood, but this unsure link just adds more uncertainty to the picture. Any way one looks at it, GHGs are strongly implicated in global warming. Just to stabilize atmospheric $CO_2$ levels at twice pre-industrial-revolution levels would require not just stable but sharply *declining* emissions within a few decades from now. Forecasting ahead a century or two, the levels of atmospheric GHGs that may ultimately be attained (unless drastic measures are undertaken) have likely not existed for at least tens of millions of years and the rate of change will likely be unique on a time scale of hundreds of millions of years.

Astonishingly, conventional CBAs and IAMs take almost no direct account of the magnitude of these unprecedented changes in GHGs – and the enormous uncertainty they create for an economic analysis of climate change. Perhaps even more remarkable is the fact that the gradualist "policy ramp" that emerges from standard CBAs and IAMs attains optimal stabilization at levels of $CO_2$ that are about 650-700 ppm within a century or two. This is my Exhibit A in the case that conventional CBAs and IAMs underplay, or sometimes even disregard, the tremendous structural uncertainties associated with climate change.

"Exhibit B" concerns the ultimate temperature response to such kind of unprecedented increases in GHGs.

---

[2] See Barnola et al (2008)

So-called "climate sensitivity" is a key macro-indicator of the *eventual* temperature response to GHG changes. Let $\Delta \ln CO_2$ be sustained relative change in atmospheric carbon dioxide while $\Delta T$ is equilibrium temperature response. Narrowly defined, climate sensitivity (here denoted $S_1$) converts $\Delta \ln CO_2$ into $\Delta T$ by the formula $\Delta T \approx (S_1/\ln 2) \times \Delta \ln CO_2$. As the Intergovernmental Panel on Climate Change in its IPCC-AR4 (2007) Executive Summary puts it: "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." Climate sensitivity is not the same as temperature change, but, for the benchmark-serving purposes of the simplistic example I will be creating, I assume that the shapes of both PDFs are roughly similar after ≈100-200 years because a doubling of anthropogenically-injected $CO_2$-equivalent ($CO_2$-e) GHGs relative to pre-industrial-revolution levels is essentially unavoidable within the next ≈40-50 years and will plausibly remain well above 2×preindustrial levels for at least ≈100+ years thereafter. Other things being equal, higher values of climate sensitivity will produce higher temperatures at a more remote time in the distant future, which begs the question of whether enough can be learned sufficiently rapidly – relative to the super-long residence time of atmospheric $CO_2$ – to make meaningful mid-course corrections (and whether there would be sufficient political will to do it in time). To fully address these timing issues requires a more complete dynamic model (along with assumptions about the dynamics of information and learning), but I believe the example here is still telling.

In this paper I am mostly concerned with the roughly 15% of those $S_1$ "values substantially higher than 4.5°C" which "cannot be excluded." A grand total of twenty-two peer-reviewed studies of climate sensitivity published recently in reputable scientific journals and encompassing a wide variety of methodologies (along with 22 imputed PDFs of $S_1$) lie indirectly behind the above-quoted IPCC-AR4 (2007) summary statement. These 22 recent scientific studies cited by IPCC-AR4 are compiled in Table 9.3 and Box 10.2. It might be argued that these 22 studies are of uneven reliability and their complicatedly-related PDFs cannot easily be combined, but for the simplistic purposes of this illustrative example I do not perform any kind of formal Bayesian model-averaging or meta-analysis (or even engage in informal cherry picking). Without question, a more sophisticated analysis of how to aggregate scientific data from different sources is required (including a more careful treatment of many aspects I am neglecting, such as possible multiplicative combining of probabilities from overlapping studies). Instead, I just naively assume that all 22 studies have equal credibility and for my purposes here their PDFs can be simplistically aggregated. The upper 5% probability level averaged over all 22 climate-sensitivity studies cited in IPCC-AR4 (2007) is 7°C while the median is 6.4°C,[3] which I take as signifying approximately that $P[S_1 > 7°\text{C}] \approx 5\%$. Glancing at Table 9.3 and Box 10.2 of IPCC-AR4, it is apparent that the upper tails of these 22 PDFs tend to be sufficiently long and fat that one is allowed from a simplistically-aggregated PDF of these 22 studies the rough approximation $P[S_1 > 10°\text{C}] \approx 1\%$. The actual empirical reason why these upper tails are long and fat dovetails beautifully with the theory of this paper: 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 which case one is forced back onto depending more than one might wish upon the prior PDF, which of necessity is largely subjective and relatively diffuse. As a recent *Science* commentary put it: "Once the world has warmed by 4°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." However one looks at it, the long fat tail of climate sensitivity is disturbing. This is Exhibit B in my case that conventional CBAs and IAMs may not adequately cover the deep structural uncertainties associated with climate change.

"Exhibit C" 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 of climate change. The chief worry here is a significant supplementary component that conceptually should be added on to climate sensitivity $S_1$. This omitted component concerns the powerful self-amplification potential of greenhouse warming due to heat-induced releases of sequestered carbon. One example is the huge volume of GHGs currently sequestered in arctic permafrost and other boggy soils (mostly as methane, a particularly potent GHG). A yet-more-remote possibility, which in principle should also be included, is heat-induced releases of the even-vaster offshore deposits of $CH_4$ trapped in the form of hydrates (aka clathrates) – for which there is a decidedly non-zero probability of destabilized methane seeping into the atmosphere if water temperatures over the continental shelves warm just slightly. The amount of methane involved is huge, although it is not precisely known.[4]

---

[3] Details of this calculation are available upon request. 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°C. (The fact that all PDFs in Box 10.2 have been normalized to zero probability above 10°C biases my upper-5% averages here towards the low side.) With the fourth conflict (Gregory et al (2002a)), I substituted 8.2°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°C.

[4] IPCC4 contains some discussion of methane releases.

Most estimates place the carbon content of methane hydrate deposits at about the same order of magnitude as the sum total of all of traditional fossil fuel deposits that will ever be extracted and burned by humans. A $CH_4$ outgassing-amplifier process could potentially precipitate (over the very long run, to be sure) a cataclysmic high-positive-feedback warming. This real physical basis for a catastrophe scenario is my exhibit C in the case that conventional CBAs and IAMs do not adequately cover the structural uncertainties associated with possible climate-change disasters. Other examples of a real physical basis for a disastrous outcome could be cited, but this one will do here.

The real physical possibility of endogenous heat-triggered releases at high temperatures of the enormous amounts of naturally-sequestered GHGs is a good example of indirect carbon-cycle feedback-forcing effects that I think should be included in the abstract interpretation of a concept of "climate sensitivity" that is relevant for this paper. What matters for the economics of climate change is the reduced-form relationship between atmospheric stocks of *anthropogenically-injected* $CO_2$-e GHGs and temperature change. Instead of $S_1$, which stands for "climate sensitivity narrowly defined," I work throughout the rest of this paper with $S_2$, which (abusing scientific terminology somewhat here) stands for a more abstract "generalized climate-sensitivity-like multiplier-parameter" that includes heat-induced feedbacks on the forcing from the above-mentioned releases of naturally-sequestered GHGs, increased respiration of soil microbes, climate-stressed forests, and other weakenings of natural carbon sinks. The transfer from $\Delta \ln$[anthropogenically-injected $CO_2$-e GHGs] to eventual $\Delta T$ is not linear (and is not even a true long-run equilibrium relationship), but for the purposes of this highly-aggregated example the linear approximation is good enough. This suggests that a doubling of anthropogenically-injected $CO_2$-e GHGs causes (very approximately) ultimate temperature change $\Delta T \approx S_2$.

The main point here is that the PDF of $S_2$ has an even-longer even-fatter tail than the PDF of $S_1$. A recent study by Torn and Harte (2006) can be used to give some very rough idea of the relationship of the PDF of $S_2$ to the PDF of $S_1$. It is universally accepted that in the absence of any feedback gain, $S_1$=1.2°C. If $g_1$ is the conventional feedback gain parameter associated with $S_1$, then $S_1=1.2 / [1-g_1]$, whose inverse is $g_1=[S_1-1.2] / S_1$. Torn and Harte estimated that heat-induced GHG releases add about .067 of gain to the conventional feedback factor, so that (expressed in my language) $S_2=1.2 / [1-g_2]$, where $g_2=g_1+.067$. (The .067 is only an estimate in a linearized formula, but it is unclear in which direction higher order terms would pull the formula and even if this .067-coefficient were considerably lower my point would remain.) Doing the calculations, $P[S_1>7°\text{C}]=5\%=P[g_1>.828]=P[g_2>.895]$ implies $P[S_2>11.5°\text{C}]=5\%$. Likewise, $P[S_1>10°\text{C}]=1\%=P[g_1>.88]=P[g_2>.947]$ implies $P[S_2>22.6°\text{C}]=1\%$ and presumably corresponds to a scenario where $CH_4$ and $CO_2$ are outgassed on a large scale from degraded permafrost soils, wetlands, and clathrates.[5] The effect of heat-induced GHG releases on the PDF of $S_2$ is *extremely* nonlinear at the upper end of the PDF of $S_2$ because, so to speak, "fat tails conjoined with fat tails beget yet-fatter tails."

Of course my calculations and the numbers above can be criticized, but (quibbles and terminology aside) I don't think most climate scientists would say that these calculations are fundamentally wrong in principle or that there exists a clearly superior method for generating rough estimates of extreme-impact tail probabilities. Without further ado I just assume for purposes of this simplistic example that $P[S_2 > 10°\mathrm{C}] \approx 5\%$ and $P[S_2 > 20°\mathrm{C}] \approx 1\%$, implying that anthropogenic doubling of $CO_2$-e eventually causes $P[\Delta T > 10°\mathrm{C}] \approx 5\%$ and $P[\Delta T > 20°\mathrm{C}] \approx 1\%$, which I take as my base-case tail estimates in what follows. These small probabilities of what amounts to huge climate impacts occurring at some indefinite time in the remote future are wildly-uncertain unbelievably-crude ballpark estimates – most definitely *not* based on hard science. But the subject matter of this paper concerns just such kind of situations and my overly simplistic example in this case does not depend at all on precise numbers or specifications. To the contrary, the major point I am trying to make is that such numbers and specifications *must* be imprecise and that this is a significant part of the climate-change economic-analysis problem, whose strong implications have thus far largely been ignored.

Stabilizing anthropogenically-injected $CO_2$-e GHG stocks at anything like twice pre-industrial-revolution levels looks now like an extremely ambitious goal, which would require sharply declining GHG emissions within a few decades. Given current trends in emissions, we will attain such a doubling of anthropogenically-injected $CO_2$-e GHG levels around the middle of this century and will then go far beyond that amount unless drastic measures are taken starting soon. Projecting current trends in business-as-usual GHG emissions, a *tripling* of anthropogenically-injected $CO_2$-e GHG concentrations would be attained relative to pre-industrial-revolution levels by early in the 22nd century. Countering this effect is the idea that we just might begin someday to seriously cut back on GHG emissions (especially if we learn that a high-$S_2$ catastrophe is looming – although the extraordinarily long inertial lags in the commitment pipeline converting $CO_2$ emissions into temperature increases might severely limit this option). On the other hand, maybe currently-underdeveloped countries

---

[5] I am grateful to John Harte for guiding me through these calculations, although he should not be blamed for how I am interpreting or using the numbers in what follows. The Torn and Harte study is based upon an examination of the 420,000-year record from Antarctic ice cores of temperatures along with associated levels of $CO_2$ and $CH_4$. While based on different data and a different methodology, the study of Sheffer, Brovkin, and Cox (2006) supports essentially the same conclusions as Torn and Harte (2006). A completely independent study from simulating an interactive coupled climate-carbon model of intermediate complexity in Matthews and Keith (2007) confirms the existence of a strong carbon-cycle feedback effect with especially powerful temperature amplifications at high climate sensitivities.

like China and India will develop and industrialize at a blistering pace in the future with even more GHG emissions and even less GHG emissions controls than have thus far been projected. Or, who knows, we might someday discover a revolutionary new carbon-free energy source or make a carbon-fixing technological breakthrough. Perhaps natural carbon-sink sequestration processes will turn out to be weaker (or stronger) than we thought. There is also the unknown role of climate engineering. The recent scientific studies behind my crude ballpark numbers could turn out to too optimistic or too pessimistic – or I might simply be misapplying these numbers by inappropriately using values that are either too high or too low. And so forth and so on. For the purposes of this very crude example (aimed at conveying some very rough empirical sense of the fatness of global-warming tails), I cut through the overwhelming enormity of climate-change uncertainty and the lack of hard science about tail probabilities by sticking with the overly simplistic story that $P[S_2{>}10°\text{C}] \approx P[\Delta\text{T}{>}10°\text{C}] \approx 5\%$ and $P[S_2{>}20°\text{C}] \approx P[\Delta\text{T}{>}20°\text{C}] \approx 1\%$. I can't know precisely what these tail probabilities are, of course, but no one can – and *that* is the point here. To paraphrase again the overarching theme of this example: the moral of the story does not depend on the exact numbers or specifications in this drastic oversimplification, and if anything it is enhanced by the fantastic uncertainty of such estimates.

It is difficult to imagine what $\Delta T \approx 10°\text{C-}20°\text{C}$ might mean for life on earth, but such high temperatures have not been seen for hundreds of millions of years and such a rate of change over a few centuries would be unprecedented even on a time scale of billions of years. Global average warming of 10°C-20°C masks tremendous local and seasonal variation, which can be expected to produce temperature increases *much* greater than this at particular times in particular places. Because these hypothetical temperature changes would be geologically instantaneous, they would effectively destroy planet Earth as we know it. At a minimum such temperatures would trigger mass species extinctions and biosphere ecosystem disintegration matching or exceeding the immense planetary die-offs associated in Earth's history with a handful of previous geo-environmental mega-catastrophes. There exist some truly terrifying consequences of mean temperature increases ≈10°C-20°C, such as: disintegration of the Greenland and at least the Western part of the Antarctic ice sheets with dramatic raising of sea level by perhaps 30 meters or so, critically-important changes in ocean heat transport systems associated with thermohaline circulations, complete disruption of weather, moisture and precipitation patterns at every planetary scale, highly consequential geographic changes in freshwater availability, regional desertification – and so forth and so on.

All of the above-mentioned horrifying examples of climate-change mega-disasters are incontrovertibly possible on a time scale of centuries. They were purposely selected to come across as being especially lurid in order to drive home a valid point. The tiny probabilities of nightmare impacts of climate change are all such crude ballpark estimates (and they would occur so far in the future) that there is a tendency in the literature to dismiss altogether these highly uncertain forecasts on the "scientific" grounds that they are much too speculative to be taken seriously. In a classical-frequentist mindset, the tiny probabilities of nightmare catastrophes are so close to zero that they are highly statistically insignificant at any standard confidence level and one's first impulse can understandably be to just ignore them or wait for them to become more precise. My main theme contrasts sharply with the conventional wisdom of *not* taking seriously extreme-temperature-change probabilities *because* such probability estimates aren't based on hard science and are statistically insignificant. The exact opposite logic holds because there is a Bayesian sense in which, other things being equal, the more speculative and fuzzy are the tiny subjective tail probabilities of extreme events, the less ignorable and the more serious is the impact on present discounted expected utility for a risk-averse agent.

When fed into an economic analysis, the great open-ended uncertainty about eventual mean planetary temperature change cascades into yet-much-greater yet-much-more-open-ended uncertainty about eventual changes in welfare. There exists here a very long chain of tenuous inferences fraught with huge uncertainties in every link beginning with unknown base-case GHG emissions; then compounded by huge uncertainties about how available policies and policy levers will transfer into actual GHG emissions; compounded by huge uncertainties about how GHG-flow emissions accumulate via the carbon cycle into GHG-stock concentrations; compounded by huge uncertainties about how and when GHG-stock concentrations translate into global mean temperature changes; compounded by huge uncertainties about how global mean temperature changes decompose into regional temperature and climate changes; compounded by huge uncertainties about how adaptations to, and mitigations of, climate-change damages are translated into utility changes – especially at a regional level; compounded by huge 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 immense cascading of huge uncertainties is a reduced form of truly stupendous uncertainty about the aggregate expected-present-discounted utility impacts of catastrophic climate change, which mathematically is represented by a very-spread-out very-fat-tailed PDF of what might be called "welfare sensitivity."

Even if a generalized climate-sensitivity-like scaling parameter such as $S_2$ could be bounded above by some big number, the value of "welfare sensitivity" is effectively bounded only by some *very* big number representing something like the value of statistical civilization as we know it or maybe even the value of statistical life on earth as we know it. *This* is the essential point of this simplistic motivating example. Suppose it were granted for the sake of argument that an abstract climate-sensitivity-like scaling parameter such as $S_2$ might somehow be constrained at the upper end by some fundamental law of physics that assigns a probability of exactly zero to temperature change being above some critical physical constant instead of continuously higher temperatures occurring with continuously lower probabilities trailing off asymptotically to zero. Even granted such an upper bound on $S_2$, the essential point here is that the enormous unsureness about (and enormous sensitivity of CBA to) an arbitrarily-imposed "damages function" for high temperature changes makes the relevant reduced-form criterion of welfare sensitivity to a fat-tailed generalized scaling parameter seem almost unbelievably uncertain at high temperatures – to the point of being essentially unbounded for practical purposes. This is my Exhibit C.

## 3 Fat Tails, High-Temp Disutilities, and Discounting

Because the integral of a PDF is one, the PDF of catastrophic temperatures must decline to an asymptote of zero probability. Thus, extreme outcomes can happen, but their likelihood diminishes to zero as a function of how extreme is the output. The fact that extreme outcomes cannot be eliminated altogether, but are theoretically possible with some positive probability, is not unique to climate change. What is worrisome is not the fact that extreme tails are *long* per se (reflecting the fact that a meaningful upper bound does not exist), but that they are *fat* (with probability density). The critical question is *how fast* does the probability of a catastrophe decline relative to the scope and impact of the catastrophe. Other things being equal, a thin-tailed PDF is of less concern because the probability of the bad event declines exponentially (or faster). A fat-tailed distribution, where the probability declines polynomially in the temperature, can be much more worrisome.

A variety of mechanisms will produce a fat-tailed distribution of long-run temperature changes. One mechanism concerns the fact that climate sensitivity is of the form $S = 1/(1-g)$ where $g < 1$ is a feedback-gain coefficient. If the PDF of $g$ allows values near one, even though with very low probability, then the PDF of $S$ tends to be long and fat-tailed.[6]

Another mechanism for generating fat tails is structural uncertainty.[7] The basic idea behind this mechanism can be illustrated by a specific example. Oversimplifying enormously, how warm the climate ultimately gets is approximately a product of two factors – anthropogenically-injected $CO_2$-e GHGs and a critical climate-sensitivity-like scaling multiplier. Both factors are uncertain, but the scaling parameter is more open-ended on the high side with a longer and fatter upper tail. This critical scale parameter reflecting huge scien-

---

[6] This idea is developed in an influential article by Roe and Baker (2007).

[7] This idea is developed extensively in Weitzman (2009).

tific uncertainty is then used as a multiplier for converting aggregated GHG emissions – an input mostly reflecting economic uncertainty – into eventual temperature changes. Suppose the true value of this scaling parameter is unknown because of limited past experience, a situation that can be modeled *as if* inferences must be made inductively from a finite number of data observations. At a sufficiently high level of abstraction, each data point might be interpreted as representing an outcome from a particular scientific or economic study. Having an uncertain scale parameter in such a setup can add a significant tail-fattening effect to posterior-predictive PDFs, even when Bayesian learning takes place with arbitrarily large (but finite) amounts of data. Loosely speaking, the driving mechanism here is that the operation of taking "expectations of expectations" or "probability distributions of probability distributions" spreads apart and fattens the tails of the reduced-form compounded posterior-predictive PDF. It is inherently difficult to learn from finite samples alone enough about the probabilities of extreme events to thin down the bad tail of the PDF because, by definition, we don't get many data-point observations of such catastrophes. This mechanism provides some kind of a generic argument why fat tails are almost inherent in many situations.

Although the basic idea is more general, it can be illustrated concretely by the relationship between the normal distribution and the Student-$t$. A normal distribution is thin tailed because the tail probabilities in the PDF decline faster than exponentially. If we do not know the parameters of the normal distribution (the mean and, more importantly, the standard deviation), but we have $n$ observations drawn from the normal distribution, the implied posterior-predictive distribution is Student-$t$ with $n$ degrees of freedom. A Student-$t$ PDF with $n$ degrees of freedom is thick tailed because it is readily confirmed that the tails are polynomial of order $n$.

A Student-$t$ "child" posterior-predictive PDF from a large number of observations looks almost exactly like its bell-shaped normal "parent" except that the probabilities are somewhat more stretched out, making the tails appear relatively fatter at the expense of a slightly-flatter center. In the limit, the ratio of the fat Student-$t$ tail probability divided by the thin normal tail probability approaches infinity, even while both tail probabilities are approaching zero. Intuitively, a normal density "becomes" a Student-$t$ from a tail-fattening spreading-apart of probabilities caused by the variance of the normal having itself a (inverted gamma) probability distribution. It is then no surprise that people are more averse qualitatively to a relatively fat-tailed Student-$t$ posterior-predictive child distribution than they are to the relatively thin-tailed normal parent which begets it. Perhaps more surprising is the quantitative strength of this endogenously-derived aversion to the effects of unknown tail-structure. The story behind this quantitative strength is that fattened posterior-predictive bad tails represent structural or deep uncertainty about the possibility of rare high-impact disasters that – using colorful language here – "scare" any agent having a utility function with relative risk aversion everywhere bounded above zero.

Uncertain structural parameters (coupled with finite data, under conditions of everywhere-positive relative risk aversion) can have strong consequences for CBA when catastrophes are theoretically possible, because in such circumstances it can drive CBA much more than anything else, including discounting. When fat-tailed temperature PDFs are combined with a utility function that is sensitive to high temperatures, it can make quite a difference on the outcomes of CBA. To see why the functional form of damages from high temperatures can be critical in this context, even when limited to a quadratic expression, consider the following formulation. Let $U(C, A)$ stand for utility as a function of consumption $C$ and environmental amenities $A$. Let $A(\Delta T)$ stand for environmental amenities as a function of temperature change. Suppose in what follows the base case

$$A(\Delta T) = \frac{1}{1+\gamma(\Delta T)^2} \tag{1}$$

for some positive constant $\gamma$.

All existing IAMs treat high-temperature damages by a rather casual extrapolation of whatever specification is (typically arbitrarily) 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 parameter choices. Almost any function can be made to fit the low-temperature damages assumed by the modeler, even though these functions can give enormously different values at higher temperatures. Most IAM damages functions reduce welfare-equivalent consumption by a quadratic-polynomial multiplier equivalent to (1), with $\gamma$ calibrated to some postulated loss for $\Delta T \approx$ 2-3°C. The standard conventional functional form combining consumption and environmental amenities is

$$U(C, A) = \frac{(CA)^{1-\eta}}{1-\eta}, \tag{2}$$

where $\eta$ is the coefficient of relative risk aversion. The particular (and, in this context, perhaps peculiar) choice of functional form (1), (2) allows the economy to substitute consumption for high temperatures relatively easily, since the elasticity of substitution between $C$ and $A$ in this particular formulation is one.

There was never any more compelling rationale for the particular functional form (1), (2) than the comfort that economists feel from having worked with it before. In other words, the multiplicative quadratic-polynomial specification is extrapolated to assess climate-change damages at high temperatures for no better reason than casual familiarity with this particular form from other cost-of-adjustment dynamic economic models, where it has been used primarily for analytical simplicity in a situation that, at best, approximates reality for small changes. I would argue that if, for some unclear reason, climate-change economists want dependence of damages to be a multiplicative function of $(\Delta T)^2$ of form (2), then a far better function at high temperatures for a consumption-reducing welfare-equivalent quadratic-based multiplier is the exponential form $A(\Delta T) = \exp(-\gamma(\Delta T)^2)$. Why? Look at the specification choice abstractly.

With isoelastic utility, the exponential specification is equivalent to $dU/U \propto dA$, while the polynomial specification is equivalent to $dU/U \propto dA/A$. For me it is obvious that, between the two, the former is much superior to the latter. When temperatures are already high in the latter case, why should the impact of $dA$ on $dU/U$ be artificially and unaccountably diluted via dividing $dA$ by high values of $A$? The same argument applies to any polynomial in $\Delta T$. I cannot prove that my favored choice is the more reasonable of the two functional forms for high $\Delta T$ (although I truly believe that it is), but no one can disprove it either – and *this* is the point here.

The value of $\gamma$ required for calibrating welfare-equivalent consumption at $\Delta T \approx$2-3°C to be (say) ≈97-98% of consumption at $\Delta T = 0$°C is so miniscule that both the polynomial-quadratic multiplier $1/[1+\gamma(\Delta T)^2]$ and the exponential-quadratic multiplier $\exp(-\gamma(\Delta T)^2)$ give virtually identical outcomes for relatively small values of $\Delta T$, but at ever higher temperatures they gradually, yet ever increasingly, diverge. With a fat-tailed PDF of $\Delta T$, there can be a big difference between these two functional forms in the implied willingness to pay (WTP) to avoid or reduce uncertainty in $\Delta T$. When the consumption-reducing welfare-equivalent damages multiplier has the exponential form $\exp(-\gamma(\Delta T)^2)$, then with fat tails the WTP to avoid (or even reduce) fat-tailed uncertainty can be a very large fraction of consumption.

In a recent article, Sterner and Persson (2008) tested on a leading IAM a utility function of the constant-elasticity-of-substitution (CES) form

$$U(C, A) = \frac{1}{1-\eta}\left[C^{\frac{\sigma-1}{\sigma}} + bA^{\frac{\sigma-1}{\sigma}}\right]^{\frac{(1-\eta)\sigma}{\sigma-1}}, \qquad (3)$$

where $b$ and $\sigma$ are positive constants, with $\sigma$ being the the (constant) elasticity of substitution between $C$ and $A$. The particular multiplicative form (2) is a special case (for $\sigma = 1$) of the more general CES form (3). Sterner and Persson chose as elasticity of substitution what they argue is a more appropriate (than $\sigma = 1$) value of $\sigma = \frac{1}{2}$, and show that this change can make a big difference on the economic policy recommended by an IAM. When $\sigma = \frac{1}{2}$, the policy-ramp tightening of GHG emissions is much stronger and steeper than for the conventional IAM case $\sigma = 1$.

When $\sigma = \frac{1}{2}$ and $\eta = 2$, which are the numerical values chosen by Sterner and Persson in their example, then (3) (with (1)) becomes

$$U = -\left[\frac{1}{C} + \gamma(\Delta T)^2\right] - 1. \tag{4}$$

Exactly the same policy implications that are shown by Sterner and Persson would come out of the simpler specification of additively-separable utility of the form

$$U = \frac{C^{1-\eta}}{1-\eta} - \gamma(\Delta T)^2 \tag{5}$$

for the plausible coefficient of relative risk aversion value $\eta = 2$. Form (5) is a standard isoelastic utility of consumption minus a conventional quadratic loss function in temperature changes. Because (4) and (5) are identical (except for a multiplicative constant), the numerical experiment of Sterner and Persson can be interpreted as showing that there is a big difference in policy implications between the standard (in the literature) multiplicative form (2) and the less standard (but no less familiar) additive form (5). In an optimal policy, the additive form (5) induces a much more stringent curtailment of GHG emissions than the multiplicative form (2). This demonstrates clearly how seemingly minor changes in the specification of high-temperature damages (here from multiplicative to additive) can dramatically change the climate-change policies recommended by an IAM. Such fragility of policy to forms of disutility functions might be considered a fourth exhibit – "Exhibit D" if you will – in making the case that conventional CBAs and IAMs do not adequately cope with structural uncertainty – here uncertainty about the specification of damages.

There will be a *really* big impact on making optimal GHG-emissions policy be more stringent if the discount rate used for discounting future climate-change disutilities is small (combined with fat tails of high temperatures and a disutility function that is sensitive to high-temperature damages). It is critical to bear in mind that the number being discussed here for discounting is the so-called "rate of pure time preference" or "utility discount rate," which is a subjective taste-like parameter that is difficult to pin down. It is *much* harder to argue that this utility discount rate should *not* be almost zero than it is to make such an argument for the so-called "goods interest rate," which is far more directly tied to observed market rates of return on capital that are unquestionably significantly positive. (Any goods interest rate can be made compatible with a zero rate of pure time preference just by adjusting the elasticity of marginal utility appropriately.[8]) If a near-zero rate of pure time

---

[8] This is explained, e.g., in Dasgupta (2007).

preference is used to discount disutilities of temperature damages, the optimal policy may curtail GHG emissions very severely in formulation (5) (or (4)). Even when the rate of pure time preference is positive, but it is not known, there can still be a big impact if there is a possibility of the rate of time preference being near zero. When this "utility discount rate" itself has a PDF with non-negligible probability density in a neighborhood of zero, then the expected present discounted disutility (from additive quadratic temperature damages) can be very large. In order to avoid such a probability-weighted bad possibility, the optimal policy will curtail GHG emissions, typically severely.

Reasonable attempts to constrict the length or the fatness of the "bad" tail (or to modify the utility function) can still leave us with uncomfortably big numbers whose exact value depends non-robustly upon artificial constraints or parameters that we really do not understand. The only legitimate way to avoid this potential problem is when there exists strong *a priori* knowledge that restrains the extent of total damages. If a particular type of idiosyncratic uncertainty affects only one small part of an individual's or a society's overall portfolio of assets, exposure is naturally limited to that specific component and bad-tail fatness is not such a paramount concern. However, some very few but very important real-world situations have potentially *unlimited* exposure due to structural uncertainty about their potentially open-ended catastrophic reach. Climate change potentially affects the whole worldwide portfolio of utility by threatening to drive all of planetary welfare to disastrously low levels in the most extreme scenarios.

The part of the distribution of possible future outcomes that can most readily be learned (from inductive information of a form as if conveyed by data) concerns the relatively more likely outcomes in the middle of the distribution. From previous experience, 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 towards probabilities in the periphery of the distribution, however, we are increasingly moving into the unknown territory of subjective uncertainty where our probability estimate of the probability distributions themselves becomes increasingly diffuse because the frequencies of rare events in the tails cannot be pinned down by previous experiences or past observations. It is not possible to learn enough about the frequency of extreme tail events from finite samples alone to make the outcome of a CBA independent of artificially-imposed bounds on the extent of possibly-ruinous disasters. Climate-change economics generally – and the fatness of climate-sensitivity tails specifically – are prototype examples of this principle, because we are trying to extrapolate inductive knowledge far outside the range of limited past experience.

# 4 Some Implications of "Fat-Tailed Logic"

By "fat-tailed logic" I mean a combination of fat tails and temperature-sensitive disutilities, along with low rates of pure time preference. A common reaction to the conundrum for CBA implied by fat-tailed logic is to acknowledge its mathematical foundation but to wonder how it is to be used constructively for deciding what to do in practice. Is this fat-tailed logic an economics version of an impossibility theorem which signifies that there are fat-tailed situations where economic analysis is up against a strong constraint on the ability of any quantitative analysis to inform us without committing to an empirical CBA framework that is based upon some explicit numerical estimates of the miniscule probabilities of all levels of catastrophic impacts down to absolute disaster? Even if it were true that this logic represents a valid economic-statistical precautionary-like principle which, at least theoretically, might dominate decision making, would not putting into practice this "generalized precautionary principle" freeze all progress if taken too literally? Considering the enormous inertias that are involved in the buildup of GHGs, and the warming consequences, is the possibility of learning and mid-course corrections a plausible counterweight to this fat-tailed logic, or, at the opposite extreme, has the commitment of GHG stocks in the ultra-long pipeline already fattened the bad tail so much that it doesn't make much difference what is done in the near future about GHG emissions? How should the bad fat tail of climate uncertainty be compared with the bad fat tails of various proposed solutions such as nuclear power, geoengineering, or carbon sequestration in the ocean floor? Other things being equal, this fat-tailed logic suggests as a policy response to climate change a relatively more cautious approach to GHG emissions, but *how much* more caution is warranted?

I simply don't know the full answers to the extraordinarily wide range of legitimate questions that fat-tailed logic raises. I don't think anyone does. But I also don't think that such questions can be allowed in good conscience to be simply shunted aside by arguing, in effect, that when probabilities are small and *im*precise, then they should be set *precisely* to zero. To the extent that uncertainty is formally considered at all in the economics of climate change, the artificial practice of using thin-tailed PDFs – especially the usual practice of imposing *de minimis* low-probability-threshold cutoffs that casually dictate what part of the high-impact bad tail is to be truncated and discarded from CBA – seems arbitrary and problematic.[9] In the spirit that the unsettling questions raised by fat-tailed CBA for the economics of climate change must be addressed seriously, even while admitting that we don't know all of the answers, I offer here some speculative thoughts on what it all means. Even if

---

[9] Adler (2007) sketches out in some detail the many ways in which *de minimis* low-probability-threshold cutoffs are arbitrary and problematic in more-ordinary regulatory settings.

the quantitative magnitude of what fat-tailed logic implies for climate-change policy seems somewhat hazy, the qualitative direction of the policy advice is nevertheless quite clear.

In ordinary run-of-the-mill limited exposure or thin-tailed situations, there is at least the underlying theoretical reassurance that finite-cutoff-based CBA might (at least in principle) be an arbitrarily-close *approximation* to something that is accurate and objective. In fat-tailed unlimited-exposure situations, by contrast, there is no such theoretical assurance underpinning the arbitrary cutoffs – and CBA outcomes are not robust to fragile assumptions about the likelihood of extreme impacts and how much disutility they cause.

One does not want to abandon lightly the ideal that CBA should bring independent empirical discipline to any application by being based upon empirically-reasonable parameter values. Even when fat-tailed logic applies, CBA based upon empirically-reasonable functional forms and parameter values might reveal useful information. Simultaneously one does not want to be obtuse by insisting that the catastrophe logic behind fat tails makes no practical difference for CBA because the parameters just need to be determined empirically and then simply plugged into the analysis along with some extrapolative guesses about the form of the "damages function" for high-temperature catastrophes (combined with speculative extreme-tail probabilities). So some sort of a tricky balance is required between being overawed by fat-tailed catastrophe logic into abandoning CBA altogether and being underawed into insisting that it is just another empirical issue to be sorted out by business-as-usual CBA. By all means plug in tail probabilities, disutilities of high impacts, rates of pure time preference, and so forth, and then see what emerges empirically – but do not be surprised when CBA outcomes are very sensitive to specifications and parameter values.

The degree to which the kind of "generalized precautionary principle" that comes out of fat-tailed reasoning is relevant for a particular application must be decided on a case-by-case "rule of reason" basis. In the particular application to the economics of climate change, with so obviously limited data and limited experience about the catastrophic reach of climate extremes, to ignore or suppress the significance of rare fat-tailed disasters is to ignore or suppress what economic-statistical decision theory seems to be telling us here is potentially the most important part of the analysis.

Where does global warming stand in the portfolio of extreme risks currently facing us? There exist maybe half a dozen or so serious "nightmare scenarios" of environmental disasters perhaps comparable in conceivable worst-case impact to catastrophic climate change. These might include: biotechnology, nanotechnology, asteroids, strangelets, pandemics, runaway computer systems, nuclear proliferation.[10] It may well be that each of these possibilities of environmental catastrophe deserves its own CBA application of fat-tailed logic along with

---

[10] Many of these are discussed in Posner (2004), Sunstein (2007), and Parson (2007).

its own empirical assessment of how much probability measure is in the extreme tails. Even if this were true, however, it would not lessen the need to reckon with the strong potential implications of fat-tailed logic for CBA in the particular case of climate change.

Perhaps it is little more than raw intuition, but for what it is worth I do not feel that the handful of other conceivable environmental catastrophes are nearly as critical as climate change. I illustrate with two specific examples. The first is widespread cultivation of crops based on genetically-modified organisms (GMOs). At casual glance, climate-change catastrophes and bioengineering disasters might look similar. In both cases, there is deep unease about artificial tinkering with the natural environment, which can generate frightening tales of a planet ruined by human hubris. Suppose for specificity that with GMOs the overarching fear of disaster is that widespread cultivation of so-called "Frankenfood" might somehow allow bioengineered genes to escape into the wild and wreak havoc on delicate ecosystems and native populations (including, perhaps, humans), which have been fine-tuned by millions of years of natural selection. At the end of the day I think that the potential for environmental disaster with Frankenfood is much less than the potential for environmental disaster with climate change – along the lines of the following loose and oversimplified reasoning.

In the case of Frankenfoods interfering with wild organisms that have evolved by natural selection, there is at least *some* basic underlying principle that plausibly dampens catastrophic jumping of artificial DNA from cultivars to landraces. After all, nature herself has already tried endless combinations of mutated DNA and genes over countless millions of years, and what has evolved in the fierce battle for survival is only an infinitesimal subset of the very fittest permutations. In this regard there exists at least some inkling of a prior argument making it fundamentally implausible that Frankenfood artificially selected for traits that humans find desirable will compete with or genetically alter the wild types that nature has selected via Darwinian survival of the fittest. Wild types have already experienced innumerable small-step genetic mutations, which are perhaps comparable to large-step human-induced artificial modifications and which have not demonstrated survival value in the wild. Analogous arguments may also apply for invasive "superweeds," which so far represent a minor cultivation problem lacking ability to displace either landraces or cultivars. Besides all this, safeguards in the form of so-called "terminator genes" can be inserted into the DNA of GMOs, which directly prevent GMO genes from reproducing themselves.

A second possibly-relevant example of comparing climate change with another potential catastrophe concerns the possibility of a large asteroid hitting Earth. In the asteroid case it seems plausible to presume there is much more inductive knowledge (from knowing something about asteroid orbits and past collision frequencies) pinning down the probabilities to very small "almost known" values. If we use $P[\Delta T > 20°\mathrm{C}] \approx 1\%$ as the very rough probability of a climate-change cataclysm occurring within the next two centuries, then this is roughly ten thousand times larger than the probability of a large asteroid impact (of a one-in-a-hundred-million-years size) occurring within the same time period.

Contrast the above discussion about plausible magnitudes or probabilities of disaster for genetic engineering or asteroid collisions with possibly-catastrophic climate change. The climate-change "experiment," whose eventual outcome we are trying to infer now, "tests" the planet's response to a geologically-instantaneous exogenous injection of GHGs. An exogenous injection of this much GHGs this fast seems unprecedented in Earth's history stretching back perhaps billions of years. Can anyone honestly say now, from very limited prior information and very limited empirical experience, what are reasonable upper bounds on the eventual global warming or climate change that we are currently trying to infer will be the outcome of such a first-ever planetary experiment? What we *do* know about climate science and extreme tail probabilities is that the rate of change of GHGs seems almost unprecedented in geological history, planet Earth hovers in an unstable trigger-prone "whipsaw" ocean-atmosphere system[11], chaotic dynamic responses to geologically-instantaneous GHG shocks are possible, and all twenty-two recently published studies of climate sensitivity cited by IPCC-AR4 (2007), when mechanically aggregated together, estimate on average that $P[S_1 > 7°\mathrm{C}] \approx 5\%$. To my mind this open-ended aspect with a way-too-high subjective probability of a catastrophe makes GHG-induced global climate change vastly more worrisome than cultivating Frankenfood or colliding with large asteroids.

These two examples hint at making a few meaningful distinctions among the handful of situations where fat-tailed logic might reasonably apply. My discussion here is hardly conclusive, so we cannot rule out a biotech or asteroid disaster. However, I would say on the basis of this line of argument that such disasters seem *very very* unlikely, whereas a climate disaster seems "only" *very* unlikely. In the language of this paper, synthetic biology or large asteroids feel more like high-knowledge situations that we know a lot more about relative to climate change, which by comparison feels more like a low-knowledge situation about which we know relatively little. Whether my argument here is convincing or not, the overarching principle is this: the mere fact that my logic might also apply to a few other environmental catastrophes does *not* constitute a valid reason for excluding it from applying to climate change.

The simplistic story I am telling here represses the real-option value of waiting and learning. Concerning this aspect, however, with climate change we are on the four horns of two dilemmas. The horns of the first dilemma are the twin facts that built-up stocks

---

[11] On the nature of this unstable "whipsaw" climate equilibrium, see Hansen et al (2007).

of GHGs might end up *ex post* representing a hugely-expensive irreversible accumulation, but so too might massive investments in non-carbon technologies that are at least partly unnecessary.

The second dilemma is the following. Because climate-change catastrophes develop slower than some other potential catastrophes, there is ostensibly somewhat more chance for learning and mid-course corrections with global warming relative to, say, biotechnology (but not necessarily relative to asteroids when a good tracking system is in place). The possibility of "learning by doing" may well be a more distinctive feature of global-warming disasters than some other disasters, and in that sense deserves to be part of an optimal climate-change policy. The other horn of this second dilemma, however, is the nasty fact that the ultimate climate response to GHGs has tremendous inertial pipeline-commitment lags of several centuries up to millennia (via the very long atmospheric residence time of $CO_2$). When all is said and done, I don't think there is a smoking gun in the biotechnology, asteroid, or any other catastrophe scenario quite like the idea that a crude amalgamation of numbers from the most recent peer-reviewed published scientific articles is suggesting something like $P[S_2{>}10°\text{C}]{\approx}5\%$ and $P[S_2{>}20°\text{C}]{\approx}1\%$.

The logic of catastrophic climate change seems to be suggesting here that the debate about what interest rate to use for discounting goods and services, which has dominated the discussion so far, may be secondary to a debate about the open-ended catastrophic reach of climate disasters. While it is always fair game to challenge the assumptions of a model, when theory provides a generic result (like "free trade is Pareto optimal" or "steady growth eventually outstrips one-time change") the burden of proof is commonly taken as being upon whoever wants to over-rule the theorem in a particular application. The burden of proof in climate-change CBA might be upon whomever calculates expected discounted utilities and disutilities without considering that structural uncertainty might matter more than discounting or pure objective risk.

# 5 Possible Implications for Climate-Change Policy

Instead of the existing IAM 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 for the big IAMs, Monte Carlo analysis could be used much 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 utility-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.

In situations of potentially unlimited damage exposure 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 policy 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.

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.[12] All of this having been said, the bind we find ourselves in now on climate change starts from

---

[12] U.S. Environmental Protection Agency (1990), executive summary projections for 2000, which I updated and extrapolated to 2007.

a diffuse prior situation to begin with, and is characterized by extremely slow convergence of inductive knowledge towards resolving the tail uncertainties – relative to the lags and irreversibilities from not acting before structure is more fully identified.

The point of all of this is that economic analysis is not completely helpless in the presence of deep structural uncertainty and potentially unlimited exposure. We can say a few important things about the relevance of fat-tailed CBA to the economics of climate change. The analysis is much more frustrating and much more subjective – and it looks much less conclusive – because it requires some form of speculation (masquerading as an "assessment") about the extreme bad-fat-tail probabilities and utilities. Compared with the thin-tailed case, CBA of fat-tailed potential catastrophes is inclined to favor paying a lot more attention to learning how fat the bad tail might be and – if the tail is discovered to be too heavy for comfort after the learning process – is a lot more open to at least considering undertaking serious mitigation measures (including, perhaps, geoengineering in the case of climate change) to slim it down fast. This paying attention to the feasibility of slimming down overweight tails is likely to be a perennial theme in the economic analysis of catastrophes. The key economic questions here are: what is the overall cost of such a tail-slimming weight-loss program and how much of the bad fat does it remove from the overweight tail?

## 6 An Analytical Foundation for Fast Geoengineering?

The economist's case for a carbon tax is traditionally made without explicit reference to extreme tail behavior. This argument is presumably strengthened when extreme tail events are considered. Whatever value it happens to be, the "uncomfortably big number" for expected disutility that tends to emerge from fat-tailed logic can be reduced by imposing carbon taxes. So the very first thing to say here is that the fat upper tail of the PDF of possible temperature changes lends even greater urgency to reducing GHG emissions by levying a substantial tax on the burning of fossil fuels. 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 analysis 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.[13] My main focus here is on the narrow question of whether the analytical argument of this paper supports 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 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 $CO_2$, 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 the bad effects of high atmospheric $CO_2$, 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. The 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.

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) $CO_2$ abatement is really costly; (2) fast geoengineering

---

[13] 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).

is really 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.[14] 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 policy argument that *now* is also the time to start thinking seriously about an international framework governing the use of this scary option.

# 7 Conclusion

Heroic attempts at constructive suggestions notwithstanding, it is painfully apparent that fat-tailed logic makes economic analysis trickier and more open-ended in the presence of deep structural uncertainty. The economics of fat-tailed catastrophes raises difficult conceptual issues which cause the analysis to appear less scientifically conclusive and to look more contentiously subjective than what comes out of an empirical CBA of more usual thin-tailed situations. But if this is the way things are with fat tails, then this is the way things are, and it is an inconvenient truth to be lived with rather than a fact to be evaded just because it looks less scientifically objective in cost-benefit applications.

Perhaps in the end the climate-change economist can help most by not presenting a cost-benefit estimate for what is inherently a fat-tailed situation with potentially unlimited downside exposure as if it is accurate and objective – and perhaps not even presenting the analysis as if it is an approximation to something that is accurate and objective – but instead by stressing somewhat more openly the fact that such an estimate might conceivably be arbitrarily *in*accurate depending upon what is subjectively assumed about the high-temperature damages function along with assumptions about the fatness of the tails and/or where they have been cut off. Even just acknowledging more openly the incredible magnitude of the deep structural uncertainties that are involved in climate-change analysis – and explaining better to policy makers that the artificial crispness conveyed by conventional IAM-based CBAs here is especially and unusually misleading compared with more-ordinary non-climate-change CBA situations – might elevate the level of public discourse concerning what to do about global warming. All of this is naturally unsatisfying, frustrating, and not what economists are used to doing – but in rare situations like climate change, where fat-tailed logic

---

[14] Barrett (2008) contains an excellent discussion of some implications of what he has dubbed "the incredible economics of geoengineering."

applies, we may be deluding ourselves and others with misplaced concreteness if we think that we are able to deliver anything much more precise than this with even the biggest and most-detailed climate-change IAMs as currently constructed and deployed.

This paper has presented a basic theoretical principle that holds under temperature-sensitive disutilities and potentially unlimited exposure. In principle, what might be called the catastrophe-insurance aspect of such a fat-tailed unlimited-exposure situation, which can never be fully learned away, can dominate discounting, objective-probability risk, and consumption smoothing. Even if this principle in and of itself does not provide an easy answer to questions about how much catastrophe insurance to buy (or even an easy answer in practical terms to the question of what exactly *is* catastrophe insurance buying for climate change or other applications), I believe it still might provide a useful way of framing the economic analysis of catastrophes.

# References

[1] **Adler, Matthew D.** "Why De Minimis?" AEI-Brookings Joint Center Related Publication 07-17, dated June 2007.

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

[3] **Barrett, Scott.** "The Incredible Economics of Geoengineering." *Environmental and Resource Economics*, 2008 (January), 39, pp. 45-54.

[4] **Crutzen, Paul J.** "Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma?" *Climatic Change*, 2006 (August), 77, pp. 211-220.

[5] **Dasgupta, Partha.** "Commentary: the Stern Review's Economics of Climate Change." *National Institute Economic Review*, 2007, 199, pp. 4-7.

[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] **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).

[8] **Matthews, H. Damon, and David W. Keith.** "Carbon-cycle feedbacks increase the likelihood of a warmer future." *Geophysical Research Letters*, 2007, 34, L09702.

[9] **Nordhaus, William D.** "The Stern Review on the Economics of Climate Change." *Journal of Economic Literature*, 2007, 45 (3), pp. 686-702.

[10] **Parson, Edward A.** "The Big One: A Review of Richard Posner's *Catastrophe: Risk and Response.*" *Journal of Economic Literature*, 2007, XLV (March), pp. 147-164.

[11] **Posner, Richard A.** *Catastrophe: Risk and Response.* Oxford University Press, 2004.

[12] **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.

[13] **Roe, Gerard H. and Marcia B. Baker.** "Why is Climate Sensitivity So Unpredictable?" *Science*, 2007 (October 26), 318, pp. 629-632.

[14] **Sunstein, Cass R.** *Worst-Case Scenarios.* Harvard University Press, 2007.

[15] **Sheffer, Martin, Victor Brovkin, and Peter M. Cox.** "Positive feedback between global warming and atmospheric $CO_2$ concentration inferred from past climate change." *Geophysical Research Letters*, 2006 (May), 33 (10), L10702.

[16] **Stern, Nicholas et al.** *The Economics of Climate Change.* Cambridge University Press, 2007.

[17] **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.

[18] **Torn, Margaret S. and John Harte.** "Missing feedbacks, asymmetric uncertainties, and the underestimation of future warming." *Geophysical Research Letters*, 2006 (May), 33 (10), L10703.

[19] **U.S. Environmental Protection Agency.** *Environmental Investments: The Cost of a Clean Environment.* Washington, DC: U.S. Government Printing Office, 1990.

[20] **Weitzman, Martin L.** "The Stern Review of the Economics of Climate Change." *Journal of Economic Literature*, 45 (3), pp. 703-724, (2007).

[21] **Weitzman, Martin L.** "On Modeling and Interpreting the Economics of Catastrophic Climate Change." Forthcoming in *Review of Economics and Statistics*, February 2009.