Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting

Citation

Weitzman, Martin L. 2009. Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics* 3:2009-39.

Published Version

<http://www.economics-ejournal.org/economics/journalarticles/2009-39>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9639963>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting

Martin L. Weitzman

Harvard University

Abstract

This paper in applied theory argues that there is a loose chain of reasoning connecting the following three basic links in the economics of climate change: 1) additive disutility damages may be appropriate for analyzing some impacts of global warming; 2) an uncertain feedback-forcing coefficient, which might be near one with infinitesimal probability, can cause the distribution of the future time trajectory of global temperatures to have fat tails and a high variance; 3) when high-variance additive damages are discounted at an uncertain rate of pure time preference, which might be near zero with infinitesimal probability, it can make expected present discounted disutility very large. Some possible implications for welfare analysis and climate-change policy are briefly noted.

JEL: O54

Keywords: Climate change; fat tails

Correspondence

Martin L. Weitzman, Department of Economics, Harvard University, e-mail:
mweitzman@harvard.edu.

Without tying them to the contents of this paper, the author is grateful for critical comments on an earlier version by Frank Ackerman, Marcia Baker, Stephen DeCanio, Simon Dietz, David Frame, Olle Häggström, John Harte, Peter Huybers, Michael MacCracken, Torsten Persson, Alan Robock, Michael Schlesinger, Robert Stavins, Jörgen Weibull, and Thomas Wigley.

1 Introduction

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 one or two facets of the problem. This is a paper in applied theory, where the application is to climate change. The paper ranges widely but is primarily about some economic implications of some of the unusually large structural uncertainties surrounding climate-change extremes.

One major structural uncertainty in climate-change economics concerns the appropriate way to represent damages from global warming. The functional form used most frequently in the literature is a nested utility specification, within which consumption is reduced multiplicatively by a quadratic-polynomial function of temperature change. In this paper I argue that an additive form, in which welfare is the difference between the utility of consumption and a quadratic disutility-loss function of temperature change, may make just as much sense for evaluating extreme climate damages. The distinction between multiplicative and additive welfare specifications may seem arcane, but I explain why it can make a surprisingly significant difference in the evaluation of future scenarios involving both high temperatures and high consumption.

There are a great many structural uncertainties about climate change extremes other than the specification of damages, which represents just one part of the economic-welfare side. To represent structural uncertainty on the science side, I also use one specific example, although several others might have served in this capacity. So-called “climate sensitivity” is the equilibrium mean surface temperature response to a doubling of atmospheric CO₂. An oft-asked question is: why has it been so difficult, after some three decades of intensive scientific research, to narrow down the upper-tail probability density function (PDF) of climate sensitivity enough to exclude very high values (say substantially higher than 4.5°C)? A standard answer is that seemingly tiny uncertainties concerning the possibility of a large feedback factor f near one are naturally amplified into broad uncertainties about *very* large values of climate sensitivity λ by a highly nonlinear transformation of the form $\lambda = \lambda_0/(1 - f)$. A detailed examination of the generic analytical mechanism behind such an explanation reveals that the implied upper-tail distribution of climate sensitivity is so “fat” (or “heavy” or “thick” – all synonyms) with probability that its variance is infinite. In other words, essentially the same argument used by most scientists to explain why high values of climate sensitivity cannot be excluded contains within itself the seeds of a generic argument not just for a fat upper tail of the PDF of climate sensitivity, but for a *very* fat tail, which is so spread out that it has infinite variance. With an additive quadratic loss function, this infinite variance translates into infinite expected disutility.

Climate sensitivity is a long run equilibrium concept that abstracts away from the transient dynamics by which it is approached as an asymptote. The next two logical questions concern the transient dynamic phase: (1) what happens to the PDF of temperatures along the uncertain dynamic trajectory that leads to climate sensitivity as an asymptotic limit (and whose variance approaches infinity, but only at an infinitely distant future time)? (2) what is the welfare evaluation of the

uncertain transient trajectory of temperatures? The simplest diagnostic energy-balance model is used to specify quantitatively how the PDF of temperatures varies over time in approaching its infinite-variance limiting PDF. In the case of an additive damages function (here quadratic in temperature changes), welfare evaluation then mainly becomes centered on the issue of how future disutilities should be discounted. I show that when the “rate of pure time preference” or “utility discount rate” is uncertain, but it has a PDF with infinitesimal probability in a neighborhood of zero, then the limiting expected present discounted disutility of additive quadratic temperature damages approaches infinity. The paper closes by mentioning some possible welfare and policy implications of this disturbing theoretical finding.

2 Temperature Damages: Multiplicative or Additive?

This section of the paper argues that it can make a big difference for climate-change policy whether high-temperature damages are specified as entering the overall net utility function multiplicatively or additively with consumption. Most modelers use a multiplicative formulation, perhaps not realizing the degree to which their model’s outcomes depend sensitively on this particular assumption. Here I argue that an additive form might possibly make as much sense as a multiplicative form and indicate why this seemingly obscure distinction might matter a lot, especially at high temperatures conjoined with high consumption.

Although generalizations are possible, suppose for the sake of specificity here that the utility of consumption is isoelastic with coefficient of relative risk aversion two. Let C be consumption, while T stands for temperature change above the pre-warming level. A utility function commonly used in the economics of climate change (for coefficient of relative risk aversion two) is, up to an affine transformation, of the multiplicative form

$$U_M(C, T) = - \left[\left(\frac{1}{C} \right) \times (1 + \gamma_M T^2) \right], \quad (1)$$

where γ_M is a positive coefficient calibrated to some postulated loss for $T \approx 2\text{--}3^\circ\text{C}$.¹ Equation (1) is essentially a single-attribute utility function, or, equivalently, a multi-attribute utility function with strong substitutability between the two attributes. This would be an appropriate formulation if the main impact of climate change is, say, to drive up the price of food and increase the demand for air conditioning.

Instead of the *multiplicative* functional form (1), suppose we now consider, up to an affine transformation, the analogous *additive* functional form (with a quadratic loss function)

$$U_A(C, T) = - \left[\left(\frac{1}{C} \right) + (1 + \gamma_A T^2) \right], \quad (2)$$

¹Such type of calibration is done in Nordhaus (2008) and Sterner and Persson (2008), among others.

where γ_A is a positive coefficient calibrated to some postulated loss for $T \approx 2\text{--}3^\circ\text{C}$. Equation (2) is a genuine multi-attribute utility function. It describes a situation where the main impact of climate change is on things that are not readily substitutable with material wealth, such as biodiversity and health. The 1 in the right side of equation (2) is an inessential constant, intended only to facilitate comparison of (1) with (2). When C is normalized to unity, then $U_M(1, T) = U_A(1, T)$, and the same calibration can be used to fix the same value of $\gamma_A = \gamma_B = \gamma$ in both cases.

Note that (1) and (2) are symmetric, with the only difference being the “ $\times$ ” sign in (1) and the “ $+$ ” sign in (2). I think it is fair to say that it is hard to argue strongly for one form over the other from any basic principles, so that, at first glance there might seem to be little basis for choosing between the multiplicative form (1) and the additive form (2). I do not want to take a decisive stand on which of (1) or (2) are “better” formulations of temperature damages. The main purpose of this paper is to point out that a seemingly arcane theoretical distinction between additive and multiplicative disutility damages may have surprisingly strong implications for economic policy. There is not much difference between (1) and (2) for small values of C and T , but when they are large, which is the domain of the utility function about which we are most unsure, I show that the distinction becomes significant.

In an effort to compare and contrast in familiar language the basic properties of the multiplicative form (1) with the additive form (2), I ask the following question. What is the willingness to pay as a fraction of consumption that the representative agent would accept to reduce temperature change to zero? This welfare-equivalent fraction of consumption w must satisfy the equation

$$U((1 - w)C, 0) = U(C, T). \quad (3)$$

Plugging (3) into (1) and (2), one obtains, respectively,

$$w_M = \frac{\gamma T^2}{1 + \gamma T^2} \quad (4)$$

and

$$w_A = \frac{\gamma C T^2}{1 + \gamma C T^2}. \quad (5)$$

If C is normalized to unity, then for all temperature changes the two specifications are identical and $w_M = w_A$. Notice, though, what happens as C increases from its initial value of one. Under the multiplicative specification (4), the fraction of consumption willing to be paid to eliminate temperature change, w_M , is independent of C . This might appear to be odd because one might think that in a rich world the fraction of consumption people would be willing to sacrifice to eliminate a given temperature change would be higher than in a poor world. Note that w_A in (5) has just this property.

Let time be denoted t . Thus, consumption at time t is $C(t)$, while temperature change at time t is $T(t)$. For notational simplicity, normalize so that $C(0) = 1$, $T(0) = 0$. To compare and contrast in familiar language the basic dynamic properties of the multiplicative form (1) with the additive form (2), I now ask the follow-

ing question. What is the willingness to pay as a fraction of *current* consumption $C(0) = 1$ that the representative agent would accept to reduce the temperature $T(t) > 0$ at time $t > 0$ down to $T(t) = 0$? Call this value W . Suppose that the rate of pure time preference or “utility discount rate” is δ . Then W must satisfy the equation

$$U(1, 0) - U((1 - W), 0) = \exp(-\delta t) [U(C(t), 0) - U(C(t), T(t))]. \quad (6)$$

Suppose consumption grows at rate g , so that $C(t) = \exp(gt)$. Then plugging (1) and (2) into (6), after some algebraic rearranging one obtains

$$W_M = \frac{\exp(-(g + \delta)t) \gamma T(t)^2}{1 + \exp(-(g + \delta)t) \gamma T(t)^2} \quad (7)$$

and

$$W_A = \frac{\exp(-\delta t) \gamma T(t)^2}{1 + \exp(-\delta t) \gamma T(t)^2}. \quad (8)$$

The difference between the multiplicative formulation (7) and the additive formulation (8) is that the latter is free of the powerful dampening term $\exp(-gt)$. To give a numerical example emphasizing the significance of this kind of distinction, suppose that $g = 2\%$, $\delta = 0$, $t = 150$. By the Ramsey formula with coefficient of relative risk aversion $\eta = 2$, the corresponding real interest rate in this example is $r = \delta + \eta g = 4\%$. Calibrate γ so that 2% of welfare-equivalent consumption is lost (at $C(0) = 1$) when $T = 2^\circ\text{C}$. Then straightforward calculations show that the willingness to pay at time $t = 0$ to avoid $T(150) = 4^\circ\text{C}$ under the multiplicative specification is $W_M = 0.4\%$, while under the additive specification it is $W_A = 7.5\%$ – a difference of almost twenty times. Another way to see this dramatic difference is ask how much of a welfare-equivalent temperature reduction in 150 years would 7.5% of current consumption buy. With additive utility (8), the answer (from above) is 4°C . With multiplicative utility (7), the answer is 18°C !

In this spirit it might be argued that, relative to the multiplicative form (1), the additive formulation (2) does not trivialize the welfare impacts of large future temperature changes. One lesson to be drawn from this simple numerical example is that a seemingly arcane distinction between an additive and a multiplicative interaction of temperature change with consumption might have big consequences. If so, then it becomes another example of structural uncertainty exerting a decisive influence on climate-change policy (here the structural uncertainty concerns the functional form of utility damages).

In an important article, Sterner and Persson (2008) tested on a leading integrated assessment model (IAM) a utility function of the CES form

$$U(C, E) = \frac{1}{1 - \eta} \left[(1 - b) C^{\frac{\sigma-1}{\sigma}} + b E^{\frac{\sigma-1}{\sigma}} \right]^{\frac{(1-\eta)\sigma}{\sigma-1}}, \quad (9)$$

where the “environmental” good is

$$E = \frac{E_0}{1 + aT^2}. \quad (10)$$

With elasticity of substitution $\sigma = \frac{1}{2}$ and coefficient of relative risk aversion $\eta = 2$, the above formulation (9), (10) is equal (up to an affine transformation) to

$$U_S(C, T) = - \left[\frac{1}{C} + 1 + \gamma T^2 \right], \quad (11)$$

where $\gamma = ab/[(1 - b)E_0]$.

Comparing (11) with (2), $U_S(C, T)$ is identical with $U_A(C, T)$. Thus, for the parameter values $\sigma = \frac{1}{2}$ and $\eta = 2$ chosen by Sterner and Persson, the additive version (2) and the CES version (11) are the same and will therefore give the same results when plugged into any IAM. Consequently, one is free to view this utility function either through the lens of an additive form or through the lens of a CES form, using whichever lens gives more insight for a particular application.

Importantly, Sterner and Persson found empirically that plugging their CES utility function (11) into William Nordhaus’s pathbreaking DICE model² yields a far more stringent emissions policy than Nordhaus found with his multiplicative utility form (1). As an empirical matter, therefore, the seemingly obscure distinction between multiplicative and additive interactions of consumption with temperature change makes a significant difference for optimal climate change policy. This demonstrates 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. I think the underlying reason is more or less transparent from the previous discussion of the comparison of (7) with (8).

Fragility of policy to forms of disutility functions is a disturbing empirical finding because the outcomes of IAMs are then held hostage to basic structural uncertainty about the way in which high temperatures and high consumption interact. Furthermore, this big difference comes from a deterministic IAM (DICE with no numerical simulations of probability distributions) having a relatively high rate of pure time preference $\delta = 1.5\%$ per year. What I show theoretically in the rest of this paper is that if one introduces fat-tailed climate change uncertainty, along with even infinitesimal probabilities of low rates of pure time preference, the difference in optimal policies between additive and multiplicative utilities can become overwhelmingly dominant.

3 Deep Structural Uncertainty about Climate Extremes

In this section I try to make a brief intuitive case for the plausibility of there being big structural uncertainties in the science of extreme climate change. I would interpret

²See Nordhaus (2008).

this as heuristic evidence that an IAM might be missing something important if its results do not much depend on the treatment of these big structural uncertainties.

Ice core drilling in Antarctica began in the late 1970s and is still ongoing. The record of carbon dioxide (CO_2) and methane (CH_4) trapped in tiny ice-core bubbles currently spans 800,000 years.³ The numbers in this unparalleled 800,000-year record of GHG levels are among the very best data that exist in the science of paleoclimate. Almost all other data (including past temperatures) are inferred indirectly from proxy variables, whereas these ice-core GHG data are directly observed.

The pre-industrial-revolution level of atmospheric CO_2 (about two centuries ago) was 280 parts per million (ppm). The ice-core data show that CO_2 varied gradually during the previous 800,000 years within a relatively narrow range roughly between 180 and 280 ppm. Currently, CO_2 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_2 (per century), is at 1,780 ppb. The sum total of all carbon-dioxide-equivalent ($\text{CO}_2\text{-e}$) GHGs is currently at 435 ppm. An even more startling contrast with the 800,000-year record is the rate of change of GHGs: increases in CO_2 were below (and typically well below) 25 ppm within any past sub-period of 1,000 years, while now CO_2 has risen by 25 ppm in just the last 10 years. Thus, anthropogenic activity has elevated atmospheric CO_2 and CH_4 to levels far outside their natural range at an extremely 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 might be unique even on a time scale of hundreds of millions of years.

Another disturbing issue concerns the ultimate temperature response to such kind of unprecedented increases in GHGs. “Climate sensitivity” is a key macro-indicator of the *eventual* temperature response to GHG changes. It is defined as the average global surface warming in equilibrium following a sustained doubling of carbon dioxide concentrations. 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 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 (more on this later).

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 climate sensitivity, are cited by IPCC-AR4 (2007). How to aggregate climate sensitivity PDFs from various studies is currently a serious unresolved issue. The aggregated PDF should have a thinner tail than the individual studies to the extent that the PDFs from the

³See Lüthi 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).

different studies are conceptualized as independent draws from the same “correct” model specification. Against this, the aggregate PDF tail should be fattened to the extent that individual models overlap and are correlated in their mutual omission of important geophysical processes (like ice sheets) or carbon cycle processes (like methane releases). The upper-tail distribution of climate sensitivity remains poorly constrained even after 30 years of research. For what it is worth, the median upper five percent probability level over all 22 climate-sensitivity PDFs cited in IPCC-AR4 (2007) is 6.4°C, which is the number that I use here.

Only so-called “fast feedback” processes are included in the concept of climate sensitivity, narrowly defined. Additionally there are “slow feedback” components that are currently omitted from most general circulation models (mainly on the grounds that they are too uncertain to be included).⁴ A prime 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 CH₄ trapped in the form of hydrates (aka clathrates) – which has a decidedly non-zero probability over the long run of having destabilized methane seep into the atmosphere if water temperatures over the continental shelves warm just slightly. The amount of CH₄ 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 long run, a CH₄ 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 has a real physical basis. Other examples of an actual physical basis for catastrophic outcomes could be cited, but this one will do here.

The above methane-release scenarios are examples of slow carbon cycle feedback effects that I think should be included in the interpretation of a climate-sensitivity-like concept that is relevant for the economics of uncertain extremes. The main point here is that the PDF of fast plus slow feedback processes has a tail much heavier with probability than the PDF of slow feedback processes alone. Extraordinarily crude calculations⁵ suggest that, when slow and fast feedback processes are combined, the probability of eventually exceeding 10°C from anthropogenic doubling of CO₂ is very roughly 5%, which presumably corresponds to a scenario where CH₄ and CO₂ are outgassed on a large scale from degraded permafrost soils, wetlands, and clathrates.

To summarize the major implication for this paper, the economics of climate change consists of a very long chain of tenuous inferences fraught with big uncertainties in every link, of which anthropogenic climate sensitivity (incorporating fast and slow feedbacks) is but one component. The uncertainties begin with unknown base-case GHG emissions; then they are compounded by big uncertainties about how available policies and policy levers will transfer into actual GHG emissions; com-

⁴The distinction between “fast feedbacks” and “slow feedbacks” is explained in Hansen et al (2008).

⁵These calculations are explained in Weitzman (2009a).

pounded further by big uncertainties about how GHG flow emissions accumulate via the carbon cycle into GHG stock concentrations; compounded by big uncertainties about how and when GHG stock concentrations translate into global mean temperature changes; compounded by big uncertainties about how global mean temperature changes decompose into regional climate changes; compounded by big generic uncertainties about the appropriate structure of damage functions and how to discount their disutilities; compounded by big uncertainties about how adaptations to, and mitigations of, climate-change damages are translated into welfare 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 tails. The fat tail of the PDF of overall welfare is the reduced form that concerns economic analysis, not the PDF of climate sensitivity *per se*, which is but one limited illustrative example representing the overall science component of structural uncertainty. In other words, uncertain climate sensitivity serves in this paper as a particular example of uncertain climate science as a whole, which itself is a subset of overall uncertainty.

4 Why is Climate Sensitivity So Unpredictable?

Taking climate sensitivity as a metaphor for climate science, it appears to bother scientists a lot that, even after some three decades of intensive research, essentially no progress has been made on excluding very high values of climate sensitivity (substantially higher than 4.5°C , say). A longstanding informal explanation for this state of affairs – which focuses on the mechanism by which small individual forcing feedbacks are amplified into a large climate-sensitivity multiplier – was formalized recently in a culminating *Science* article by Roe and Baker (henceforth R&B) entitled “Why is climate sensitivity so unpredictable?,” from which the title of this section was taken.⁶ A troubling economic implication of R&B, which I explain later in more detail, is that the asymptotic tail probability of large climate sensitivity appears to be declining toward zero so slowly, relative to its impacts, that this fat-tail aspect seems like it should play a significant role in welfare analysis. Here I reformulate at a high level of abstraction the analytical essence of a R&B-style explanation to emphasize that it implicitly contains a theoretical argument pointing toward *very* fat tails of climate sensitivity. In other words, I am making the simple point that the same (or at least a very similar) argument used by climate scientists to explain why high values of climate sensitivity cannot be excluded contains within itself the seeds of a more general argument not just for fat tails, but for *very* fat tails, which spread the PDF so far apart that its variance is effectively infinite. By

⁶See Roe and Baker (2007). There is a long list of predecessors, ranging from Wigley and Schlesinger (1985) (or maybe even earlier), to Allen et al (2006).

focusing on the feedback-multiplier nub of this argument, I naturally pass over a great many important details of the underlying science.

In estimating climate sensitivity (denoted λ), feedbacks are everything. The ultimate temperature response of a climate system to imposed GHG shocks is unknown mostly because the exogenous initial forcing is amplified by uncertain endogenous feedback factors like albedo, water vapor, clouds, and so forth. These feedback factors have complex, nonlinear, and even chaotic features. Overarching this messiness, R&B argue, feedbacks still combine additively and linear systems analysis is still a useful way of seeing the forest for the trees. Climate sensitivity in this R&B-style view is a derived concept able to be portrayed abstractly as an amplifier (or multiplier) for a forcing impulse to a linear feedback process. More basic in this process than the multiplier-amplifier λ is the aggregate feedback factor or coefficient f . Not only does f act on the original CO₂ forcing, but it also acts on the results of its own forcing action, and so forth, ultimately causing an infinite series of feedback loops as described by the differential equation in the next section.

The aggregate feedback coefficient f has a critical additivity property in its components:

$$f = \sum_{j=1}^n \tilde{f}_j, \quad (12)$$

where each of the n primitive $\tilde{f}_j$ represents a feedback sub-factor, such as the albedo, water vapor, or clouds previously alluded to. In the R&B-style worldview, the total feedback-forcing factor f is considered more fundamental than the climate-sensitivity multiplier λ because f scales additively in its primitive sub-components $\tilde{f}_j$ (whereas λ does not scale additively in $\lambda_j = \lambda_0/(1 - \tilde{f}_j)$) and because each $\tilde{f}_j$ is (at least in principle) directly measurable, in the laboratory or in the field. By contrast, λ is observable only indirectly, as the equilibrium limit of an iterative multiplier process which requires (at least in principle) the passage of an infinite number of multiplier rounds over real time by the formula

$$\lambda = \lambda_0 \sum_{i=1}^{\infty} (f)^i = \frac{\lambda_0}{1 - f}. \quad (13)$$

If each primitive $\tilde{f}_j$ is an independently distributed random variable (RV) and n is large, then from (12) and the central limit theorem a case could be made that the RV f might be approximately normally distributed. A normal PDF for f is the prototype case considered in the R&B-style modeling tradition. To draw out the generic implications of a R&B-style explanation for the derived fatness of the upper tail of λ requires showing formally how the argument generalizes from the normal to essentially any reasonable PDF of f . The base-case normal PDF in the R&B-style tradition is presumably truncated from above at $f = 1$, or else it could be argued that the implied unstable dynamics would have produced a runaway feedback amplification at some time in the past. However, it is far from clear how exactly this truncation at $f = 1$ (of the normal or any other PDF) is to be carried

out in practice, or even how it is to be conceptualized – and the formal R&B-style argument for $\lambda = \lambda_0/(1 - f)$ having a PDF with a fat upper tail is left somewhat dangling on this point.

Let the PDF of the RV f be $\phi(f)$ with upper support at $f = 1$. The spirit of a R&B-style explanation is that high values of $f \approx 1^-$ having “small but non-negligible” probability get nonlinearly skewed upwards into a fat-upper-tail PDF of $\lambda = \lambda_0/(1 - f)$. Without further ado I assume that a fair translation of the idea that high values of $f \approx 1^-$ have “small but non-negligible” probability is that the PDF $\phi(f)$ has the properties

$$\phi(1) = 0, \phi'(1) < 0. \quad (14)$$

It is important here to understand that I am not assuming a positive point probability of occurrence for a feedback value $f = 1$, in which case the results to follow would be trivial. I am not even assuming that the probability *density* at $f = 1$ is positive. I am only assuming in (14) that in the limit the probability density of f changes linearly (from an initial value of zero) within an arbitrarily small neighborhood of $f \approx 1^-$.

With the usual Jacobian change-of-variable transformation $\phi(f) df \rightarrow \psi(\lambda) d\lambda = \phi(f(\lambda)) f'(\lambda) d\lambda$, applied to $\lambda = \lambda_0/(1 - f)$, the derived PDF of λ is

$$\psi(\lambda) = \lambda_0 \phi\left(1 - \frac{\lambda_0}{\lambda}\right) \frac{1}{\lambda^2} \quad (15)$$

for all $\lambda > 0$. The mean of λ is given by the expression

$$E[\lambda] \equiv \lim_{M \rightarrow \infty} \int_0^M \lambda \psi(\lambda) d\lambda < \infty, \quad (16)$$

while its variance is

$$V[\lambda] \equiv \lim_{M \rightarrow \infty} \int_0^M (\lambda - E[\lambda])^2 \psi(\lambda) d\lambda = \infty, \quad (17)$$

where the integral in (17) is blowing up essentially because, from (15), the integrand inside of (17) is proportional to $1/\lambda$ as $\lambda \rightarrow \infty$, making the integral (17) approach $-\phi'(1) \ln M$ as $M \rightarrow \infty$.

The significance of (17) for economic policy is not subtle. As climate sensitivity goes, so goes the *eventual* mean planetary temperature response to increased GHGs. While global warming is just one example of a fat-tailed $f \rightarrow \lambda$ feedback-multiplier process, it is special because of the enormous potential damages to worldwide welfare associated with very large values of λ . If additive economic damages increase at least as fast as quadratically in temperatures, as in (2), then (17) indicates that the probability-weighted *expected* value of climate damages is infinite.

A standard criticism of my (or any) oversimplified reliance on the time-independent long run equilibrium concept of climate sensitivity is that the catastrophically high

temperature values will materialize (if they materialize at all) only in the remote future. If one brought back the time element by focusing more on the transient dynamics and less on the stationary limit, this line of argument goes, short to medium term concerns would dominate – and the climate sensitivity issue might recede. Some have even interpreted the R&B-style explanation of fat-tailed climate sensitivity as signifying that the concept itself has run into diminishing returns, and the scientific community should essentially “call off the quest” of trying to make more precise estimates of λ in favor of concentrating greater effort on analyzing short-to medium-term temperature dynamics and implications. In the rest of this paper I will show that when one formalizes the uncertain trajectory of temperature dynamics, along with analyzing carefully the issue of discounting utility (or disutility) under uncertain rates of pure time preference, then the long run behavior of the system can in principle continue to play a significant role in economic analysis and policy discussion.

5 A Dynamic Aggregative Model of Global Warming

This section compresses into a single differential equation what is arguably the simplest meaningful deterministic-dynamic model of the physical process of global warming.⁷ Of course this particular one-differential-equation model cannot possibly capture the full complexity of climate change. However, I think that the highly aggregated approach taken here is realistic enough to serve as a springboard for meaningful discussions of some basic climate-change issues, which, for the purposes of this paper, may actually be clarified when tightly framed in such stark simplicity.

Factors that affect climate change are standardly segregated into “forcings” and “feedbacks.” A climate *forcing* is a direct, primary, or exogenous energy imbalance imposed on the climate system, either naturally or by human activities. Examples include changes in solar irradiation (the prototype, in whose units all other forcings may be expressed), volcanic emissions, deliberate land modification, or anthropogenic changes in atmospheric stocks of greenhouse gases, aerosols, and their precursors. (The radiative forcing from CO₂ happens to be proportional to the logarithm of its atmospheric concentration, but this is not true in general for all GHGs.) A climate *feedback* is an indirect, secondary, or endogenous radiative imbalance that amplifies or dampens the climate response to an initial forcing. An example is the increase in atmospheric water vapor that is induced by an initial warming due to rising CO₂ concentrations, which then acts to amplify the warming through the greenhouse properties of water vapor, further accelerating the process.

Suppose, for simplicity, that in pre-industrial-revolution times ($t \leq 0$) the planetary climate system had been in a state of (relative) equilibrium at a constant

⁷This super-simple diagnostic energy-balance model is sprinkled throughout the scientific literature and appears formally in, e.g., Andrews and Allen (2007), or Roe (2007), both of which contain further references to it, including who created it and more realistic extensions of it. My only possible originality here is in expositing this basic one-differential-equation model to a broader audience, primarily economists.

temperature with constant radiative forcing and no radiative imbalance. Let $F(t)$ stand for radiative forcing at time t . Normalize $F(t) = T(t) = 0$ for $t \leq 0$. Imagine, in a thought experiment, that a sustained radiative imbalance of constant magnitude R_Σ has been *additionally* imposed. (Whether this constant additional radiative imbalance R_Σ is itself exogenous or endogenous is not relevant in this context because only the reduced-form total imbalance matters for its expository role here.) Let $T(t)$ be the temperature response at time $t > 0$. If the earth were a blackbody planet, with no atmosphere and no further feedbacks, the long-run “reference” non-feedback temperature response as $t \rightarrow \infty$ would be $T(t) \rightarrow \lambda_0 R_\Sigma$, where λ_0 is the feedback-free constant defined by the fundamental physics of a blackbody reference system as described by the Stefan-Boltzman law. Even in richer more-realistic situations with atmospheric feedbacks and complicated dynamics, other things being equal it is not a terrible approximation that at any time the temperature moves with an instantaneous velocity roughly proportional to the reference imbalance. This means that the linearized differential equation of temperature motion is

$$\dot{T}(t) = \frac{1}{k} [R_\Sigma(t) - \frac{T(t)}{\lambda_0}], \quad (18)$$

where the positive coefficient k in (18) represents the aggregate thermal inertia or effective capacity of the system as a whole to absorb heat. In this application, k essentially stands for the overall planetary ability of the oceans to take up heat.

The full temperature dynamics of an idealized non-blackbody planetary system can now most simply be described as follows. Count time in the conventional modeling format where the present corresponds to $t = 0$. At any time $t > 0$, suppose that the system is subjected to an exogenously imposed additional radiative forcing of $F(t)$ (relative to its pre-industrial-revolution equilibrium rest state of zero). In the application here, the exogenously imposed additional radiative forcing is essentially the logarithm of the relative increase of atmospheric CO_2 over pre-industrial-revolution levels. Without loss of generality, it is convenient throughout this paper to normalize the unit of forcing to correspond to a doubling of CO_2 . If $G(t)$ is the concentration of atmospheric carbon-dioxide-equivalent (CO_2 -e) GHGs at time t (in parts per million (ppm)), and $\underline{G}$ ($\approx 280\text{ppm}$) is the pre-industrial-revolution CO_2 -e concentration of atmospheric GHGs, then

$$F(t) = \frac{1}{\ln 2} \ln \left(\frac{G(t)}{\underline{G}} \right). \quad (19)$$

The trajectory of *exogenous* (or primary) radiative forcings $\{F(s)\}$ for $0 < s < t$ (here essentially standing for past anthropogenic increases in atmospheric GHG stocks) causes temperatures to rise over time, which induces feedback-like changes in secondary radiative imbalances (such as cloud formation, water vapor, ice albedo, lapse rates, and so forth). Lumped together, these “secondary” radiative imbalances are typically more powerful in ultimate magnitude than their “primary” inducers. Let the endogenously *induced* overall radiative imbalance at time t be denoted $R_I(t)$.

Let the *total* change in radiative imbalances at time t be denoted $R_{\Sigma}(t)$. Then

$$R_{\Sigma}(t) = F(t) + R_I(t). \quad (20)$$

In the problem at hand, the temperature change $T(t)$ induces a (comparatively-fast-acting, relative to (18)) endogenous radiative imbalance $R_I(t)$ according to the formula

$$R_I(t) = \frac{f}{\lambda_0} T(t), \quad (21)$$

where the (linear) *feedback factor* f is a basic parameter of the system. Not only does f act on the original CO₂ forcing, but it also acts on the results of its own forcing action, and so forth, ultimately causing an infinite series of feedback loops. As mentioned, the relevant feedback factors in climate change involve cloud formation, water vapor, albedo, and many other effects. A key property of linear feedback factors is that (as with radiative forcings or radiative imbalances) the various components and subcomponents can be aggregated simply by adding them all up because they combine additively.

Plugging (21) and (20) into (18) then yields, after simplification, the basic differential equation

$$\dot{T}(t) = \frac{1}{k} \left[F(t) - \frac{1-f}{\lambda_0} T(t) \right] \quad (22)$$

with the initial conditions $F(0) = T(0) = 0$. The closed-form solution of (22) is

$$T(t) = \frac{1}{k} \int_0^t F(s) \exp \left[(s-t) \left(\frac{1-f}{\lambda_0} \right) \right] ds. \quad (23)$$

The oversimplifications of physical reality that have gone into the one-differential-equation temperature change trajectory (22) are numerous. As just one example, the parameters that appear in (22) or (23) are not true constants because they might co-vary over time in complicated ways that this simplistic formulation is incapable of expressing. The only defence of this ultra-macro approach is a desperate need for analytical simplicity in order to see the forest for the trees. It seems fair to say that (22) captures the dynamic interplay of forces along a global-warming path decently enough for the purposes at hand – and almost surely better than any alternative formula based on one simple linear differential equation.

Even accepting the enormous oversimplifications of reality that go into an equation like (22) or (23), there remain massive uncertainties concerning the appropriate values of the structural parameters. For simplicity, the critical feedback parameter f is chosen to be the only uncertainty, but it should be appreciated that the relevant values of k and of forcings $\{F(t)\}$ are also very uncertain and co-vary in ways that are not fully represented here, even leaving aside the model itself being an uncertain

simplification of a much more complicated reality. Generally speaking, additional uncertainty strengthens the message of this paper.

Just glancing at equation (23) is suggestive of why it is so difficult in practice to infer f directly from data. The record of past forcing histories is extremely noisy and such components as aerosol concentrations are notoriously difficult to identify. Furthermore, it is readily shown that the first-order response of a system like (23) to a change in forcings does not involve long-run parameters like f at all, but more centrally concerns the overall ability of the oceans to take up heat as embodied in the thermal inertia coefficient k . The parameter k itself is not very well known in this aggregative context and can be interpreted as having time-varying values for different stages of different heat-absorbing processes. It is statistically very difficult to distinguish between a high- f low- k world and a low- f high- k world. To be able to infer f at all precisely would require a long and fairly accurate time series of past natural forcings along with a decent knowledge of the relevant thermal inertias – none of which are readily available. From this, from the difficult politics of the situation, and from the very long pipeline commitment of atmospheric CO₂ stocks, it follows that prospects for a meaningful “wait and see” reactive policy for GHG flow emissions may be quite limited.

For notational convenience and analytical sharpness, I restrict the situation here to the most basic case of the dynamic temperature reaction to a step-function forcing represented by an instantaneous doubling of CO₂-e GHGs: $F(t) \equiv 1$ for $t \geq 0$ and $F(t') \equiv 0$ for $t' < 0$. Then (23) simplifies down to

$$T(t | f) = \frac{\lambda_0}{1-f} \left[1 - \exp \left(- \left(\frac{1-f}{k\lambda_0} \right) t \right) \right]. \quad (24)$$

Note that the right hand side of (24) approaches t/k when going to the limit as $f \rightarrow 1$. This implies from (24) that $T(t | f) < T(t | 1) = t/k$ for $f < 1$, implying that the bounded RV $T(t | f)$ must have finite (but increasing) variance no matter what is the PDF of f satisfying (14).

Equilibrium climate sensitivity is defined as

$$\lambda \equiv \lim_{t \rightarrow \infty} T(t) \quad (25)$$

and it is apparent from applying (25) to (24) that

$$\lambda = \frac{\lambda_0}{1-f}, \quad (26)$$

which is one of the most basic relationships of climate change. Conventional as-if-deterministic point estimates might be $\lambda_0 \approx 1.2$, $f \approx .65$, $\lambda \approx 3$, with the standard deviation of f approximately $\sigma_f \approx .13$. Relevant values of k might vary widely in this aggregate context, depending on the time scale of the heat absorption process. Other things being equal, higher values of f shift up temperatures T all along the trajectory (24), but with higher values of f it also takes a longer response time to reach any given fraction of the asymptotic value λ represented by (25), (26).

The third link in my chain of reasoning concerns the welfare and policy implications of the infinite limiting variance of λ from the second link. Under any foreseeable technology, elevated stocks of CO₂ are committed to persist for a very long time in the atmospheric pipeline. Ballpark estimates imply that, for every unit of CO₂ anthropogenically added to the atmosphere, $\approx 70\%$ remains after 10 years, $\approx 35\%$ remains after 100 years, $\approx 20\%$ remains after 1,000 years, $\approx 10\%$ remains after 10,000 years, and $\approx 5\%$ remains after 100,000 years.⁸ It can also take a long time to learn about looming realizations of uncertain, but irreversible, climate changes. Thus, the CO₂ stock inertia, along with slow learning, makes it unreliable to react to unfolding disasters by throttling back CO₂ flow emissions in time to avert an impending catastrophe. Here I just simplistically assume that the planet will never be able to react to bad future scenarios by stabilizing atmospheric concentrations below a doubling of CO₂-e GHG concentrations relative to pre-industrial-revolution levels. Such a stark approach may be an acceptable proxy for reality in the context of the message I am trying to convey, because it seems to me, alas, that CO₂-e GHGs ≈ 560 ppm are essentially unavoidable within the next half century or so, and will plausibly remain well above this level for one or two centuries thereafter, no matter what new information is received in the meantime. This represents an extreme and perhaps unrealistic interpretation, but the modeling strategy of this paper is to lay out the essential structure of my argument as simply as possible, leaving more realistic refinements for later work.

6 How Should Climate-Change Disutilities be Discounted?

The analysis of last section showed that pushing a R&B-style explanation of fat-tailed climate sensitivity all the way to its logical conclusion implies a PDF having infinite variance. With an additive quadratic disutility of temperature change, this implies an infinite loss of expected welfare, but this infinite loss is occurring at an infinitely remote future time. The obvious next question is: what happens to expected present discounted welfare when the disutility damages of high temperatures are discounted at the appropriate rate of pure time preference?

Suppose that the damages of temperature changes are quadratic in T and of the additively separable form $U_A(C, T) = -[1/C + 1 + \gamma T^2]$ from (2), which has the analytical convenience of allowing one to calculate separately the disutility impact of temperature changes, irrespective of the time trajectory of consumption. The super-strong result that follows depends critically upon additively separable utility and would not hold (in such strength) for the multiplicative form (1): $U_M(C, T) = -(1 + \gamma T^2)/C$.

⁸See Archer (2007), pages 122-124, and the further references he cites.

The temperature change at time t conditional on the realization f is $T(t | f)$ given by formula (24). The expected disutility damages at time t are then given (up to an inessential multiplicative constant) by the expression

$$D_E(t) = \int_{-\infty}^1 [T(t | f)]^2 \phi(f) df, \quad (27)$$

and it is straightforward to show that $D_E(t)$ increases monotonically over time, approaching a limit of ∞ , which is consistent with climate sensitivity having an infinite variance in (17). If the rate of pure time preference used for discounting future utiles or disutilities is $\delta > 0$, and if the future is artificially truncated at time horizon H , then expected present discounted disutility is

$$D^*(\delta; H) = \int_0^H D_E(t) e^{-\delta t} dt. \quad (28)$$

It is essential to realize that the number δ being discussed here for discounting future disutilities is the so-called “rate of pure time preference” or “utility discount rate,” an elusive concept that is subjective and not directly observable. The “utility discount rate” δ is *not* the much-more-familiar number that is used to discount ordinary goods ($r = \delta + \eta g$ by the Ramsey formula), and which is identified with the everyday concept of an interest rate. It is *much* harder to argue that this utility discount rate δ should be significantly above zero than it is to make such an argument for the “goods interest rate” r , which is far more directly tied to observed market rates of return on capital that are significantly positive.⁹

The next obvious question is: what are appropriate values of δ and H to use in evaluating (28)? For H the answer is relatively easy: by longstanding economic logic and practice, in principle the horizon ought to be infinite and (28) should be evaluated by taking the limit as $H \rightarrow \infty$. The more difficult and more controversial issue concerns the appropriate rate of pure time preference to be used for discounting intergenerational disutility damages from future climate change. The question here is: what is δ ? I think an honest direct answer is: somewhere between zero and very roughly about 1% per year – some people might have opinions, but nobody really knows. I think it is also fair to point out that a notable minority of some very distinguished economists believe that the appropriate rate of pure time preference for discounting intergenerational utilities generally, and disutility damages from future climate change particularly, should be arbitrarily close to zero. Without taking sides directly on this issue, I approach the problem indirectly by postulating some given distribution of subjective probabilities representing overall “degrees of belief” in the appropriate value of δ to be plugged into (28).¹⁰

Let the subjective PDF of the RV δ be $h(\delta)$, with a lower support at $\delta = 0$. The formal treatment of the RV δ that follows in this section parallels the formal treatment of the RV f from the last section. In a spirit of giving at least

⁹Dasgupta (2007) has an insightful discussion of some of the main issues here.

¹⁰The logic of this position is spelled out further in Weitzman (2001).

some limited voice to the opinion that the rate of pure time preference for inter-generational discounting might be arbitrarily close to zero, I assume that low values of $\delta \approx 0^+$ have “small but non-negligible” probability in the sense that the PDF $h(\delta)$ obeys

$$h(0) = 0, \quad h'(0) > 0, \quad (29)$$

which is the analogue here of condition (14). Essentially, uncertainty concerning the possibility of a small rate of pure time preference δ near zero is naturally amplified into uncertainty about very large values of the present discounted value π of a unit flow by a highly nonlinear transformation of the form $\pi = 1/\delta$.

It is critical to understand here that I am not assuming a positive point probability of occurrence for a zero rate of pure time preference, in which case the main result of this paper would be trivial. I am not even assuming that the probability *density* of a zero rate of time preference is positive. I am only assuming in (29) that the PDF of δ increases linearly (from an initial value of zero) within an arbitrarily small neighborhood of $\delta \approx 0^+$.

There are now two RVs: f and δ . What might be called the “expected expected” present discounted disutility of temperature change is

$$D^{**}(H) = \int_0^H D^*(\delta; H) h(\delta) d\delta. \quad (30)$$

The final question to be addressed here is: what happens to $D^{**}(H)$ in the limit as $H \rightarrow \infty$? The answer is not obvious. Other things being equal, as the time horizon recedes the expected disutility damage from more-variable future temperatures is increasingly dominated by the limiting infinite variance of climate sensitivity. However, other things also being equal, discounting at a positive rate counteracts this infinite-variance asymptote. Yet a third wild card here is that pure time preference itself is a legitimately-unknown RV in this context, with some “small but non-negligible” probability of being close to zero, which tends to favor lower effective discount rates at longer horizons – again other things being equal.¹¹ The value of the “expected expected” present discounted disutility of temperature change in (30), which emerges from this pulling of different forces in different directions, is the main result of the paper that is explicated in the next section.

7 A Dismal Proposition

The following “dismal proposition” hints that fat-tailed infinite-variance climate sensitivity may have economic ramifications conceivably impacting current policy analysis.

Theorem 1 *In the above model with “small but non-negligible” probabilities,*

$$\lim_{H \rightarrow \infty} D^{**}(H) = \infty. \quad (31)$$

¹¹This effect is described in Weitzman (1998).

Proof. Consider the expression

$$\int_0^{\infty} [T(t; f)]^2 \exp(-\delta t) dt = \left(\frac{\lambda_0}{1-f} \right)^2 \int_0^{\infty} \left[1 - \exp \left(- \left(\frac{1-f}{k\lambda_0} \right) t \right) \right]^2 \exp(-\delta t) dt. \quad (32)$$

By brute-force integration, the right hand side of (32) is shown to be proportional to

$$1/[\delta(\delta + (1-f)/k\lambda_0)(\delta + 2(1-f)/k\lambda_0)]. \quad (33)$$

If it were true that

$$\lim_{\delta \rightarrow 0, f \rightarrow 1} \left\{ \frac{[\phi(f)] [h(\delta)]}{\delta(\delta + (1-f)/k\lambda_0)(\delta + 2(1-f)/k\lambda_0)} \right\} = \infty, \quad (34)$$

then it would follow that

$$E_{\delta, f} \left[\int_0^{\infty} [T(t; f)]^2 \exp(-\delta t) dt \right] = \infty, \quad (35)$$

which in turn would imply (31). But expression (34) must hold because

$$\lim_{\delta \rightarrow 0, f \rightarrow 1} \left\{ \frac{[-(1-f)\phi'(1)] [\delta h'(0)]}{\delta(\delta + (1-f)/k\lambda_0)(\delta + 2(1-f)/k\lambda_0)} \right\} = \infty, \quad (36)$$

which concludes this streamlined proof. ■

Were quadratic disutility to be discounted at a zero rate of time preference, then it would be straightforward that an infinite-variance tail of eventual temperature change would have a big impact on present discounted expected welfare. What is perhaps surprising is that this high-impact result can continue to hold even when the probability density at a zero rate of time preference is zero and is merely increasing linearly (in the small) with time preference. The infinite limit of the theorem is coming in main part from a perhaps counterintuitive implication of (29) that, with uncertain rates of pure time preference, utilities in the far-distant future are discounted at the lowest possible rate of pure time preference, here zero.¹² This holds even when the probability of $\delta \approx 0^+$ is infinitesimal. In a sense, the model has been reverse engineered via (29) to put weight on the limiting PDF of temperature changes, whose variance approaches the infinite variance of the PDF of climate sensitivity. With economic damages quadratic in temperatures, expected present discounted disutility then approaches infinity as the horizon recedes.

Of course any interpretation must be based on an assessment of the model's overall assumptions. The model of this paper is really more of a suggestive example than a fully general formulation, and an example that has been reverse engineered at that. The theorem depends on a conjunction of several basic assumptions, none

¹²For an exposition of this logic, see Weitzman (1998).

of which is beyond criticism. Analyzing what happens for $H \rightarrow \infty$ stretches the logic even further. Still, taken as a whole (and even admitting that it has been somewhat rigged), I think this “dismal proposition” makes it somewhat less easy to dismiss the significance of unpredictable climate sensitivity on the grounds that high values will have impact only in the distant future.

Theorem 1 indicates that the willingness to pay to avoid climate change is unbounded. There are several possible ways to escape this disturbing paradox of infinity. The troubling infinite limit is technically eliminated by imposing *ad hoc* inequality constraints like $H \leq 200$ years, or $T \leq 7^\circ\text{C}$, or $f \leq .99$, or $\delta \geq .001$, or so forth. However, removing the ∞ symbol in this way does not truly eliminate the underlying problem, because it then comes back to haunt in the form of an arbitrarily large expected-present-discounted disutility, whose exact value depends sensitively upon obscure bounds, truncations, severely-dampened or cut-off prior PDFs, or whatever other formal mechanisms have been used to banish the ∞ symbol. The take-away message here is that reasonable attempts to constrict bad-tail fatness can leave us with uncomfortably big numbers whose exact value may depend non-robustly upon artificial constraints or parameters, the significance of which we do not honestly comprehend. Theorem 1 should therefore be taken only figuratively as holding for some “uncomfortably big number” – but not for infinity. A reader interested in understanding more about how the infinite limit of the “dismal proposition” is to be interpreted and applied in a finite world should consult the much fuller discussion of this set of issues in Weitzman (2009a) and Weitzman (2009b). Here I restrict myself just to commenting on the widespread notion that the extreme realizations being described in the tails are so improbable that they can effectively be ignored.

One may not legitimately discard the bad tail of the PDF of a disaster on the grounds alone that the probabilities are “too small to matter.” Such *de minimis* truncation requires some asymptotic argument along the lines that the limiting *product* of the decreasing probability *times* the increasing disutility is “too small to matter.” The most practical way to avoid this vexing tail-evaluation issue is when there exists strong *a priori* knowledge that meaningfully 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 welfare, which is the usual case, 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. In these unusual situations, there is no choice but to evaluate somehow or other the limiting product of probability times disutility. 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. This tail-evaluation feature, which is essentially ignored by most conventional IAMs, can understandably dominate the economics of climate change. Such a feature makes an economic analysis of climate change look and feel uncomfortably subjective, but, at least in the formulation of this paper, it is the way things can be with fat tails and unlimited liability.

Because climate-change catastrophes develop slower than some other potential catastrophes, there is perhaps more chance for learning and midcourse corrections with global warming, relative to some other catastrophic scenarios. The possibility of “learning by doing” may well be a more distinctive feature of global-warming disasters than of some other disasters, and in that sense deserves to be part of an optimal climate-change policy. The other horn of this dilemma, however, is that the ultimate temperature responses to CO₂ stocks have tremendous inertial pipeline-commitment lags that are very difficult to reverse once they are in place. This nasty fact can be brutal on illusions about the easy corrective potential of “wait and see” reactive policies. As I already noted, it seems implausible to me that ultimate stabilized values of GHGs will end up being much less than twice pre-industrial-revolution values, no matter what realistic future responses to global warming are undertaken. Reacting to an impending climate disaster by changing a CO₂ emissions-flow instrument (to control the CO₂ stock accumulation inducing the disaster) seems offhand like using an outboard motor to maneuver an ocean liner away from an impending collision with an iceberg. The role of learning and midcourse corrections is a subject worthy of further detailed study, the outcome of which could potentially soften all of my conclusions. However, a conservative position, at least for the time being, might be to consider that by the time we learn that a climate-change disaster is impending it may be too late to do much about it.

8 Concluding Comments

Issues of uncertainty and discounting are fundamental to any economic analysis of climate change. This paper combines together three forms of structural uncertainty: how to formulate damages, how to discount these damages, and how to express future temperature dynamics. The paper shows that the single most widespread scientific explanation of why climate sensitivity is so uncertain at the upper end contains within itself a generic argument in favor of a very fat upper tail of temperature changes. When this is merged with an additively separable damages function and a rate of pure time preference that is unknown but might conceivably be close to zero, the combination can in principle dominate an economic analysis of climate change. Such a message is not intended to cause despair for the economics of climate change, nor to negate the need for further study and numerical simulations to guide policy. The message is just a cautionary note that this *particular* application of cost-benefit analysis to climate change seems more inherently prone to being dependent on subjective judgements about structural uncertainties than most other, more ordinary, applications.

References

- Allen, M., N. Andronova, B. Booth, S. Dessai, D. Frame, C. Forest, J. Gregory, G. Hegerl, R. Knutti, C. Piani, D. Sexton, and D. Stainforth (2006). "Observational Constraints on Climate Sensitivity." Chapter 9 in *Avoiding Dangerous Climate Change*, ed. H. J. Schellnhuber, Cambridge University Press, pp. 281–289.
- Allen, Myles R. and David J. Frame (2007). "Call Off the Quest." *Science* 318 (October 26): 582–583.
- Andrews, David G. and Myles R. Allen (2007). "Diagnosis of Climate Models in Terms of Transient Climate Response and Feedback Response Time." Mimeo dated July 19.
- Archer, David (2007). *Global Warming*. Blackwell Publishing.
- Dasgupta, Partha (2007). "Commentary: the Stern Review's Economics of Climate Change." *National Economic Review* 199: 4–7.
- James Hansen, Makiko Sato, Puskher Karecha, David Beerling, Valerie Masson-Delmotte, Mark Pagani, Maureen Raymo, Dana L. Royer, James C. Zachos (2008). "Target Atmospheric CO₂: Where Should Humanity Aim?" *Open Atmos. Sci. J.* 2: 217–231.
- 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 (2008). "High-Resolution Carbon Dioxide Concentration Record 650,000–800,000 Years before Present." *Nature* 453: 379–382 (15 May).
- Nordhaus, William (2008). *A Question of Balance*. Yale University Press.
- Roe, Gerard (2007). "Feedbacks, Timescales, and Seeing Red." Mimeo dated September 2.
- Roe, Gerard H., and Marcia B. Baker (2007). "Why is Climate Sensitivity So Unpredictable?" *Science* 318 (October 26): 629–632.
- Stern, Thomas, and U. Martin Persson (2008). "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate." *Review of Environmental Economics and Policy* 2 (1): 61–76.
- Weitzman, Martin L. (1998). "Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate." *Journal of Environmental Economics and Management*. 36 (3, November): 201–208.
- Weitzman, Martin L. (2001). "Gamma Discounting." *American Economic Review* 91 (1, March): 260–271.
- Weitzman, Martin L. (2009a). "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91 (1): 1–19.

- Weitzman, Martin L. (2009b). “Reactions to Nordhaus.”
<http://www.economics.harvard.edu/faculty/weitzman/files/ReactionsCritique.pdf>.
- Wigley, Thomas M. L., and Michael E. Schlesinger (1985). “Analytical Solution for the Effect of Increasing CO₂ on Global Mean Temperature.” *Nature* 318: 649–652.