



HARVARD Kennedy School
JOHN F. KENNEDY SCHOOL OF GOVERNMENT

Harvard Environmental Economics Program

Discussion Paper 09-11

Reactions to the Nordhaus Critique

Martin L. Weitzman

Harvard University

The Harvard Environmental Economics Program

The Harvard Environmental Economics Program develops innovative answers to today's complex environmental issues, by providing a venue to bring together faculty and graduate students from across the University engaged in research, teaching, and outreach in environmental and natural resources economics and related public policy. The program sponsors research projects, develops curricula, and convenes conferences to further understanding of critical issues in environmental and resource economics and policy around the world.

Acknowledgements

The Enel Endowment for Environmental Economics, at Harvard University, provides major support for HEEP. The endowment was established in February 2007 by a generous capital gift of \$5 million from Enel, SpA, a progressive Italian corporation involved in energy production worldwide. HEEP enjoys an institutional home in and support from the Mossavar-Rahmani Center for Business and Government at Harvard Kennedy School. As its name suggests, the Center is concerned with the interaction of the public and private sectors, including with regard to environmental issues.

HEEP is grateful for additional support from: Shell, Christopher P. Kaneb (Harvard AB 1990); the James M. and Cathleen D. Stone Foundation; Paul Josefowitz (Harvard AB 1974, MBA 1977) and Nicholas Josefowitz (Harvard AB 2005); and the Belfer Center for Science and International Affairs at the Harvard Kennedy School. The Harvard Project on International Climate Agreements, co-directed by Robert N. Stavins and closely affiliated with HEEP, is funded primarily by a grant from the Doris Duke Charitable Foundation.

HEEP—along with many other Harvard organizations and individuals concerned with the environment—collaborates closely with the Harvard University Center for the Environment (HUCE). A number of HUCE's Environmental Fellows and Visiting Scholars have made intellectual contributions to HEEP.

Citation Information

Weitzman, Martin L. "Reactions to the Nordhaus Critique" Discussion Paper 2009-11, Cambridge, Mass.: Harvard Environmental Economics Program, April 2009.

The views expressed in the Harvard Environmental Economics Program Discussion Paper Series are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or of Harvard University. Discussion Papers have not undergone formal review and approval. Such papers are included in this series to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Reactions to the Nordhaus Critique

Martin L. Weitzman*

March 17, 2009 – preliminary – comments appreciated

Abstract

In this paper I comment on the reactions of William Nordhaus to a recent article of mine entitled “On Modeling and Interpreting the Economics of Catastrophic Climate Change” that appeared in the February 2009 issue of the *Review of Economics and Statistics*. My target audience here is PhD-level general economists, but this paper could also be viewed as a somewhat less technical supplement to my article, which some interested non-economist readers might conceivably find useful on its own account.

1 Introduction

William Nordhaus has written a thoughtful, considered critique of my article.¹ Professor Nordhaus’s views on the economics of climate change – an area he pioneered with seminal contributions – are worthy of respect. I agree with some of his comments on my article, but respectfully disagree with some others. Naturally, here I focus more on disagreements.

At first I was inclined to debate some of Nordhaus’s criticisms line by line. After all, this is the usual operating procedure of scholars defending their own ideas and their own turf against would-be detractors or encroachers. But on second thought I found myself anxious not to be drawn, by so doing, into having the main focus be on technical details. Instead, I am more keen here to emphasize in fresh language the substantive concepts that, I think, may be more obscured than enlightened by a debate centered on technicalities. I am far more committed to the simple basic ideas that underlie my theory than to the particular mathematical form in which I have chosen to express them. These core concepts could

*Without blaming them for the remaining deficiencies in this paper, I am extremely grateful for the constructive comments of James Annan, Daniel Cole, Stephen DeCanio, Baruch Fischhoff, Don Fullerton, John Harte, David Kelly, Michael Oppenheimer, Robert Pindyck, Joseph Romm, and Richard Tol.

¹Nordhaus (2009) – available online at <http://cowles.econ.yale.edu/P/cd/d16b/d1686.pdf> – reacting to Weitzman (2009a).

have been wrapped in a variety of alternative mathematical shells – and the particular one that I chose is somewhat arbitrary. The implications are roughly similar, irrespective of formalization. Some technical details are unavoidable, but if I can make the underlying concepts acquire greater intuitive plausibility, then I believe that these ideas will become more self-evidently resistant to several of the criticisms that Nordhaus is expressing.

2 Deep Structural Uncertainty about Climate Extremes

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

In this section I try to make a heuristic-empirical case for there being big structural uncertainties in the economics of extreme climate change. I will argue on intuitive grounds that the way in which this deep uncertainty is conceptualized and formalized *should* influence substantially the outcomes of any reasonable CBA of climate change. Further, I will argue that the seeming fact that uncertainty about extremes does *not* substantially influence outcomes from the “standard” CBA is implausible. My arguments in this section are not intended to be airtight or rigorous. Rather, this is an intuitive presentation based on stylized facts.

We all know that computer-driven simulations are dependent upon the core assumptions of the model inside the computer. My intuitive examples are frankly aimed at sowing a few seeds of doubt that the “standard” CBA of climate change is fairly representing structural uncertainties, and therefore its conclusions might be more shaky than is commonly acknowledged. I will try to make my case by citing four aspects of the climate science and economics that do not seem to me to be adequately captured by the “standard” CBA. The four examples – which I call “Exhibits 1, 2, 3, and 4” – are limited to structural uncertainty concerning the modeling of climate disasters. While other important aspects of structural uncertainty might also be cited, I restrict my stylized facts to these four examples.

“Exhibit 1” concerns the atmospheric level of greenhouse gases (GHGs) over the last 800,000 years. Ice core drilling in Antarctica began in the late 1970s and is still ongoing. The record of carbon dioxide (CO₂) and methane (CH₄) trapped in tiny ice-core bubbles currently spans 800,000 years.² It is important to recognize that the numbers in this

²See Dieter et al (2008), from which my numbers are taken (supplemented by data from the Keeling curve

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 CO₂ (about two centuries ago) was 280 parts per million (ppm). The ice-core data show that CO₂ varied gradually during the last 800,000 years within a relatively narrow range roughly between 180 and 280 ppm. Currently, CO₂ 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 CO₂ (per century), is at 1,780 ppb. The sum total of all carbon-dioxide-equivalent (CO₂-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 CO₂ were below (and typically well below) 25 ppm within any past sub-period of 1,000 years, while now CO₂ has risen by 25 ppm in just the last 10 years. Thus, anthropogenic activity has elevated atmospheric CO₂ and CH₄ to levels far outside their natural range at a very rapid rate. The unprecedented scale and speed of GHG increases brings us into uncharted territory and makes predictions of future climate change very uncertain. Looking ahead a century or two, the levels of atmospheric GHGs that may ultimately be attained (unless decisive measures are undertaken) have likely not existed for tens of millions of years and the speed of this change may be unique on a time scale of hundreds of millions of years.

Remarkably, the “standard” CBA of climate change takes little account of the magnitude of the uncertainties involved in extrapolating future climate change so far beyond past experience. Perhaps even more surprising, the “policy ramp” of gradually tightening GHG emissions, which emerges as optimal policy from the “standard” CBA, attains stabilization at levels of CO₂-e GHGs that approach 700 ppm. The “standard” CBA thus recommends subjecting the Earth’s system to an unprecedented shock from geologically-instantaneously jolting atmospheric stocks of GHGs up to being two and a half times above their highest level over the last 800,000 years – without mentioning the unprecedented nature of this planetary experiment. This is my Exhibit 1.

“Exhibit 2” concerns the ultimate temperature response to such kind of unprecedented increases in GHGs. “Climate sensitivity” (hereafter denoted S_1) 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

for more recent times, available online at: ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2_mm_mlo.txt).

treatment 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 S_1 , are cited by IPCC-AR4 (2007). This is my “sample.” I assumed for my purposes that these 22 reported PDFs could be simplistically aggregated by averaging the 22 PDFs into one “representative” PDF. This form of meta-analysis can be loosely defended as an example of Bayesian model averaging, in which the different PDFs from different studies represent equally credible outcomes from more or less independent models. In his critique, Nordhaus favors a classical-frequentist-inspired meta approach in which the PDFs from the different studies are treated more like independent draws from the same “correct” PDF of the same “correct” model. In principle, he has a valid point insofar as it is not clear how best to aggregate different climate sensitivity PDFs from different studies. However, the empirical significance of his point seems exaggerated to me because I think it is fair to say that most climate scientists believe it is more appropriate to view most of the 22 studies as different models with different methodologies giving different PDFs, rather than to view the results as different observations from the same model. In any event, it is a fact that the median upper five percent probability level over all 22 climate-sensitivity PDFs cited in IPCC-AR4 (2007) is 6.4°C. Even if Nordhaus’s reasoning were to knock this *five* percent probability level down to the median upper *two* percent probability level being 6.4°C, I find such kinds of numbers disturbing. The actual empirical reason why these upper tails are long and heavy with probability dovetails with the theory of my article: 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.”³ In whichever way I look at the issue of combining different results from different studies, for me the upper tail of the “representative” PDF of climate sensitivity has too much probability not to be disturbing – and this unease is not easily dismissed. This is my Exhibit 2.

“Exhibit 3” concerns possibly disastrous releases over the long run of bad-feedback components of the carbon cycle that are currently omitted from most general circulation models.

³Allen and Frame (2007).

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 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 CH_4 trapped in the form of hydrates (aka clathrates) – which has a decidedly non-zero probability over the long run of destabilized methane seeping into the atmosphere if water temperatures over the continental shelves warm just slightly. The amount of CH_4 involved is huge, although it is not precisely known. Most estimates place the carbon-equivalent content of methane hydrate deposits at about the same order of magnitude as all fossil fuels combined. Over the long run, a CH_4 outgassing-amplifier process could potentially precipitate a disastrous strong-positive-feedback warming. Thus, the possibility of a climate meltdown is not just the outcome of a mathematical theory, but 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 indirect carbon cycle feedback effects that I think should be included in the interpretation of a climate sensitivity-like concept that is relevant here. 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,” the example in the article used S_2 , which (abusing scientific terminology) stands for a more abstract “generalized climate sensitivity-like multiplier” that includes heat-induced feedbacks from endogenous releases of naturally sequestered GHGs, increased respiration of soil microbes, climate-stressed forests, and other weakenings of natural carbon sinks.⁴ The main point here is that the PDF of S_2 has a tail even heavier with probability than the PDF of S_1 . Contrary to what Nordhaus states, my article relied on three recent peer-reviewed scientific studies to estimate roughly the PDF of S_2 . Extraordinarily crude calculations suggested $P[S_2 > 10^\circ\text{C}] \approx 5\%$, which presumably corresponds to a scenario where CH_4 and CO_2 are outgassed on a large scale from degraded permafrost soils, wetlands, and clathrates. The effect of heat-induced GHG releases on the PDF of S_2 is relatively modest at the low end, while being extremely nonlinear at the upper end of the PDF of S_2 because, so to speak, “thick tails conjoined with thick tails beget yet thicker tails.” Should my very rough cal-

⁴In scientific jargon, S_1 would be associated with “fast feedbacks,” while my S_2 would be (very loosely) associated with “slow feedbacks,” which are typically excluded from general circulation models, mostly on the grounds that they are too uncertain to be included. See Hansen et al (2008).

culations be off, and actually $P[S_2 > 10^\circ\text{C}] \approx 2\%$ (say, for example, due in part to Nordhaus’s critique of my meta-analytic approach to combining studies), I think that the substance of Exhibit 3 remains. Even if S_2 could somehow be bounded above by some big number, the value of what might be called “welfare sensitivity” is effectively bounded only by some even bigger number representing something like the value of statistical civilization as we have known it, or maybe even the value of statistical life on earth as we have known it. This is my Exhibit 3.

“Exhibit 4” concerns what I view as a somewhat cavalier treatment of damages or disutilities from extreme temperature changes. The “standard” CBA treats high-temperature damages by a somewhat passive extrapolation of whatever specification is assumed to be the low-temperature “damages function.” High-temperature damages extrapolated from a low-temperature damages function seem to be remarkably sensitive to assumed functional forms (and, to a lesser degree, parameter choices). Almost any function can be made to fit the low-temperature damages assumed by the modeler, even though these functions may give enormously different evaluations at higher temperatures. The “standard” CBA damages function reduces welfare-equivalent consumption by a quadratic-polynomial multiplier, calibrated to some postulated loss for low temperatures, and then does some sensitivity analysis by alternating the exponent. This particular choice of functional form allows the economy to substitute consumption for higher temperatures relatively easily, since the limiting elasticity of substitution between consumption and higher temperatures is one (due to the multiplicative-polynomial assumption). There are several substantive consequences that stem from using this particular high-substitution damages specification. As an example, Nordhaus has argued from his model that serious warming is rendered less important because we would have more output to offset it – due to economic growth being highly correlated with the atmospheric CO_2 stocks that drive the higher temperatures. The strength of this conviction is, at least in part, an artifact of Nordhaus’s choice of a high-substitution multiplicative-quadratic functional form for his damages function.

The “standard” CBA damages specification never had any more compelling rationale than the comfort economists feel from having previously worked with a quadratic-polynomial loss function, along with the ease of interpreting a multiplicative loss because it is directly translatable “as if” into a fraction of lost output. In other words, the multiplicative quadratic-polynomial specification is extrapolated to assess climate-change disutilities at high temperatures for no better reason than casual familiarity and convenience of interpretation. This might be justified as an acceptable approximation for the disutility of small temperature changes, but it is highly questionable when used seriously as an extrapolative device for evaluating the disutility of catastrophic climate changes.

Here is not the place to get involved in all of the details, but suffice it to say that very different optimal policies can be produced when other, in my opinion more plausible, functional forms are used to express the disutility of disastrously high temperatures.⁵ As one example, a multiplicative *exponential* damages specification under structural uncertainty is theoretically capable of inducing a far more stringent curtailment of GHG emissions than the multiplicative *polynomial* specification of the “standard” CBA. Or, to take another example, suppose that the disutility of temperature change is *additively separable* instead of being *multiplicatively separable* (as in the “standard” CBA). In his model, Nordhaus uses a multiplicatively separable quadratic loss embedded within a constant relative risk aversion (CRRA) utility function whose coefficient of relative risk aversion is two. If welfare is instead the analogous additively-separable arithmetic *difference* between a CRRA utility function of consumption (with coefficient of relative risk aversion two) and a quadratic loss function of temperature changes, it implies an elasticity of substitution between consumption and temperature change of one half. Empirically, using this additive form – even without *any* uncertainty – prescribes a significantly more stringent curtailment of GHG emissions than what emerges from the analogous multiplicative form of the “standard” CBA.⁶

The above two examples demonstrate how seemingly minor changes in the specification of high-temperature damages can dramatically alter the gradualist policy ramp outcomes recommended by the “standard” CBA. Such fragility to postulated forms of disutility functions is my Exhibit 4 in the case that “standard” CBA inadequately copes with deep structural uncertainty – here structural uncertainty about the specification of high-temperature damages. Actually, the structural uncertainty of Exhibit 4 is best seen as applying more generally to the *overall* utility function of consumption and high temperatures combined. When Nordhaus challenges a CRRA utility function (which he himself uses in all of his modeling) on the grounds that it can produce extreme results, he is making my case for me that his own results depend non-robustly on, among other things, structural uncertainty about the functional form of the overall utility of consumption and high temperatures.

To summarize, the economics of climate change consists of a very long chain of tenuous inferences fraught with big uncertainties in every link: beginning with unknown base-case GHG emissions; then compounded by big uncertainties about how available policies and

⁵Examples are discussed more elaborately in Weitzman (2009b).

⁶With coefficient of relative risk aversion two, the above additively-separable specification is mathematically equivalent to the constant elasticity of substitution (CES) specification of Sterner and Persson (2008) with elasticity of substitution one half. In their pioneering study, Sterner and Persson showed empirically – by plugging it into Nordhaus’s deterministic DICE model – that their CES (or, equivalently, my additive) welfare specification prescribes a significantly more aggressive policy response to global warming (with a significantly higher carbon tax) than the analogous multiplicative specification of the “standard” CBA. For details on the isomorphism with my additively separable formulation, see Weitzman (2009b).

policy levers will transfer into actual GHG emissions; compounded by big uncertainties about how GHG flow emissions accumulate via the carbon cycle into GHG stock 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 uncertainties about how adaptations to, and mitigations of, climate-change damages are translated into utility changes at a regional level; compounded by big uncertainties about how future regional utility changes are aggregated – and then how they are discounted – to convert everything into expected-present-value global welfare changes. The result of this lengthy cascading of big uncertainties is a reduced form of truly extraordinary uncertainty about the aggregate welfare impacts of catastrophic climate change, which mathematically is represented by a PDF that is spread out and heavy with probability in the tails.

What I would wish a reader might take away from these four exhibits is the notion that the seeming immunity of the “standard” CBA to such stylized facts seems peculiar. An unprecedented and uncontrolled experiment is being performed by subjecting planet Earth to the shock of a geologically-instantaneous injection of massive amounts of GHGs. Yet the “standard” CBA seems almost impervious to the extraordinarily uncertain probabilities and consequences of catastrophic climate change. A reader should feel intuitively that it goes against the grain of common sense when, in view of the above four exhibits of structural uncertainty, a climate-change CBA does *not* much depend upon how potential disasters are modeled and incorporated into the CBA. This uneasy feeling – of a system-wide failure being plausible science fiction that is not adequately represented in the “standard” CBA of climate change – is my opening argument. I turn next to the theory.

3 Infinity, CBA, and the “Dismal Theorem”

I begin this section by asking why is it relevant in the first place to have any supporting theory at all if the four stylized-fact “exhibits” from last section are convincing. Why aren’t these stylized facts alone sufficient evidence that there is a problem with “standard” CBA? My answer is that a combined theoretical plus empirical-intuitive argument delivers a particularly powerful one-two punch at the treatment of structural uncertainty in the “standard” CBA. In this respect I believe that the whole of my argument is bigger than the sum of its two parts. The theoretical part reinforces the empirical part by placing it within a formal mathematical framework. When the intuitive “exhibits” are seen as reflecting some formalized theoretical structure, then it becomes less easy to brush them aside as mere sniping at an established model. In this theoretical section of the paper, as in the last

empirical section, I emphasize the intuitive plausibility of the case I am trying to make – here focusing on the underlying logic driving the theory.

The last section argued that it is only common sense that climate-change policy implications should depend on the treatment of low-probability extreme-impact outcomes. In my article, the main question I attempted to address was whether such intuitive sensitivity is reflecting some deeper principle. My answer was that there is a basic underlying theoretical principle that indeed points in this direction. The logic is simple enough to be grasped intuitively without understanding the fancy math required to state and prove a formal version.

Let welfare W stand for expected present discounted utility, whose theoretical upper bound is B . Let $D \equiv B - W$ be expected present discounted *dis*utility. Here D stands for what might be called the “diswelfare” of climate change. Unless otherwise noted, my default meaning of the term “fat tail”⁷ (or “thin tail”) concerns the upper tail of the PDF of $\ln D$, resulting from whatever combination of probabilistic temperature changes, temperature-sensitive damages, and so forth, by which this comes about. The four intuitive “exhibits” of last section might be interpreted as suggesting that the PDF of $\ln D$ could have a bad tail that is too thick (with probability) for comfort. It may seem arcane, but the tail thickness of the reduced-form PDF of $\ln D$ is the analytical essence of what Nordhaus and I are debating in this interchange. Of course it is extremely difficult to know the thickness of the upper tail of the PDF of $\ln D$, which is my main point.

Because the integral over a nonnegative probability measure is one, the PDF of $\ln D$ must decline asymptotically to zero. In other words, extreme outcomes can happen, but their likelihood diminishes to zero as a function of how extreme the outcome might be. The idea that extreme outcomes cannot be eliminated altogether, but are hypothetically possible with some positive probability, is not at all unique to climate change. Almost nothing in our world has a probability of exactly zero or exactly one. What is worrisome is not the fact that the upper tail of the PDF of $\ln D$ is *long* (reflecting the fact that a meaningful bound on diswelfare does not exist), but that it is *fat* (with probability density). The critical question is *how fast* does the probability of a catastrophe decline relative to the welfare impact of

⁷As 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 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 a Student- t or inverted-gamma. 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. Although both PDFs approach a limit of zero, the ratio of a fat-tailed probability divided by a thin-tailed probability goes to infinity in the limit.

the catastrophe. Other things being equal, a thin-tailed PDF is of less concern because the probability of a bad event declines faster than exponentially. A fat-tailed distribution, where the probability declines polynomially, can be much more worrisome.

My article indicated a theoretical tendency for the PDF of $\ln D$ to have a fat tail. Conceptually, the underlying mechanism is not too difficult to grasp. Structural uncertainty essentially means that the probabilities are unsure. A formal Bayesian translation might be that the structural parameters of the relevant PDFs are themselves uncertain and have their own PDFs. The article expressed this idea in a formal argument that the reduced form “posterior predictive” PDF (in Bayesian jargon) of $\ln D$ tends to be fat tailed because the structural parameters are unknown. Loosely speaking, the driving mechanism is that the operation of taking “expectations of expectations” or “probability distributions of probability distributions” spreads apart and fattens the tails of the compounded posterior-predictive PDF. From past samples alone, it is inherently difficult to learn enough about the probabilities of extreme future events to thin down the bad tail of the PDF, because we don’t have much data about analogous past extreme events. This mechanism provides at least some kind of a generic story about why fat tails might be inherent in many situations.

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

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 of $\ln D$ is thin-tailed because the tail probabilities in the PDF decline faster than exponentially. However, if we do not know the parameters of the normal PDF of $\ln D$, but we have n observations drawn from this normal PDF, then the implied posterior-predictive distribution of $\ln D$ is Student- t with $n-1$ degrees of freedom. A Student- t PDF with $n-1$ degrees of freedom is fat-tailed because it is readily confirmed that the tails decline like one over a polynomial of order $n-1$. The article showed that this example essentially generalizes to uncertainty about the scaling parameter of any distribution. The underlying “true” PDF of $\ln D$ might be thin-tailed, but when there is subjective uncertainty about its structural parameters it can easily turn into a posterior-predictive PDF with a fat tail.

A fat upper tail of the PDF of $\ln D$ makes the willingness to pay (WTP) to avoid extreme climate changes very large, indeed arbitrarily large. The article gave a formal argument within a specific mathematical structure, but this formal argument could have been embedded in alternative mathematical structures – with the same basic message. The particular formal argument I gave in the article came in the form of what I called a “dismal theorem” (DT). In this particular formalization, the limiting expected stochastic discount factor is infinite (or, what I take to be equivalent for purposes here, the limiting WTP to avoid fat-tailed disasters constitutes all of output). Of course, real-world WTPs are not 100% of output. Presumably the PDF in the bad fat tail is thinned, or even truncated, perhaps from considerations akin to what lies behind the value of a statistical life (VSL). (After all, we would not pay an infinite amount to eliminate the fat tail of climate-change catastrophes.) Alas, in whatever way the bad fat tail is thinned or truncated, a climate-change CBA based upon it remains highly sensitive to the details of the thinning or truncation mechanism, because the disutility of extreme climate change has “essentially” unlimited liability. Later I discuss the meaning of this potential lack of robustness in climate-change CBA and speculate on some actionable consequences it might imply regarding what economists do and say.

Disagreements abound concerning how to interpret the infinity symbol that appears in the formulation of DT. There is a natural tendency to sneer at economic models that yield infinite outcomes. This reaction is presumably based on the idea that infinity is a ridiculous result; therefore any model that has an infinity symbol in it is fundamentally mis-specified, and thus dismissable. Critics argue earnestly from their favorite examples that expected disutility from climate change cannot actually be infinite, as if this were a telling indictment of the entire fat-tailed methodology. I believe that, in the particular case of climate change, the infinity is trying to tell us something important. The infinite limit in DT is a formal mathematical way of saying that structural uncertainty in the form of fat tails is, at least in

theory, capable of swamping the outcome of any CBA that disregards this aspect.

The key issue here is *not* a mathematically illegitimate use of an infinite limit in DT. It is easy to modify utility functions, to add on VSL-like restrictions, to truncate probability distributions arbitrarily, or to introduce ad hoc priors that cut off or otherwise severely dampen low values of welfare-equivalent consumption. Introducing any of these (or many other attenuating mechanisms) formally replaces the infinity symbol by some uncomfortably large, but finite, number. Unfortunately, removing the infinite limit in these or other ways does not eliminate the underlying problem because it then comes back to haunt in the form of a WTP to erase the structural uncertainty that is arbitrarily large. How large depends sensitively upon obscure details about how the upper tail of the PDF of $\ln D$ has been thinned. One can easily remove the infinity symbol from DT, but one cannot so easily “remove” the underlying substantive economic problem of extreme sensitivity to fat tails and the resulting conundrum of deciding policy under such circumstances. The overwhelming majority of real-world CBAs have thin upper tails in $\ln D$ from limited exposure to system-wide catastrophic risk. However, a few very important real-world situations have effectively unlimited exposure due to structural uncertainty about their potentially open-ended catastrophic reach. Climate change is unusual in potentially affecting the entire worldwide portfolio of utility by threatening to drive all of planetary welfare to disastrously low levels in the most extreme scenarios. This has policy implications, some of which are discussed later.

Perhaps not surprisingly, there is controversy about the implications of fat tails for CBA. My target is not CBA in general, but the particular impression of precision inadvertently conveyed by the “standard” CBA of climate change. I like to think I occupy a middle ground between two extreme positions. An economist does not want to abandon lightly the ideal that CBA should bring independent empirical discipline to any application by being based upon empirically reasonable functional forms and parameter values. Even when fat-tailed logic might apply, climate-change CBA could in principle reveal useful information about whether fat tails are or are not actually relevant for “reasonable” functional forms and parameter values at extreme temperatures. (What “reasonable” means in a context of extreme impacts with uncertain probabilities may not be clear, which in practice can introduce a large gray area into CBA of climate-change catastrophes.) Simultaneously, one does not want to be obtuse by insisting that the logic behind fat tails makes no practical difference for climate-change CBA because the parameters just need to be empirically determined and then simply plugged into the analysis. Some sort of a tricky balance is required between being overwhelmed by fat-tailed logic into a Hamlet-like paralysis that leads to abandoning CBA altogether, and being underwhelmed into insisting that it is just another empirical issue to be sorted out by business-as-usual CBA. Economists should, of course, remain open

to changing their beliefs on the basis of robust outcomes from well-designed CBAs. One should go ahead and plug into climate-change CBAs tail probabilities, plug in disutilities of disastrous impacts, plug in rates of pure time preference or coefficients of relative risk aversion, plug in various functional forms – and then respect robust conclusions. On the other hand, one should not be especially surprised if outcomes are fragile to specifications concerning catastrophic extremes. There is *some* difference in expected welfare between 500 ppm of CO₂ and 600 ppm of CO₂, and that difference could be important for policy. But *how much* of a difference may be difficult to extract with reasonable precision from a fat-tailed CBA that is sensitive to obscure modeling assumptions about climate extremes.

The “standard” CBA appears to offer a constructive ongoing scientific-economic research program for generating ever more precise outputs from ever more precise inputs. By contrast, my main message can seem threatening because it can be painted as anti-scientific and anti-economic. Fat tails and the implied limitations that prevent CBA from reaching robust conclusions are frustrating for economists. After all, we make a living from plugging rough numbers into simple models and reaching specific conclusions (more or less) on the basis of these numbers. What are we supposed to advise policy makers and politicians quantitatively about *how much* effort to spend on averting climate change if conclusions from modeling fat-tailed uncertainties are not clear-cut? Practical men and women of action have a low tolerance for vagueness and crave some kind of an answer, so they have little place for even a whiff of fuzziness from two-handed economists. It is threatening for us economists to admit that constructive “can do” climate-change CBA may be up against some limitations on the ability of quantitative analysis to give robust advice. But if this is the way things are with the economics of climate change, then this is the way things are – and non-robustness to subjective assumptions is an inconvenient truth to be lived with rather than a fact to be denied or evaded just because it looks less scientifically objective in CBA.

In my opinion, economists need to emphasize more openly to the policy makers, the politicians, and the public that, while formal climate-change CBA may be helpful, there is a danger of possible overconfidence from undue reliance on subjective judgements about the probabilities and welfare impacts of extreme events. What we can do constructively as economists is to explain better the magnitudes of the unprecedented structural uncertainties that are involved, explain why this feature limits what we can say, and present the best CBAs and the most honest sensitivity analyses that we can under fat-tailed circumstances, including many different functional forms for extremes. At the end of the day, policy makers must decide what to do on the basis of an admittedly sketchy economic analysis of a gray area that just cannot render clear robust answers. The moral of the dismal theorem is that, under extreme tail uncertainty, seemingly casual decisions about functional forms, parameter

values, and tail fatness can dominate CBA. Economists should not pursue a narrow, superficially crisp, analysis by blowing away the low-probability high-impact catastrophic scenarios as if this is a necessary price we must pay for the worthy goal of giving answers and advice to policy makers. An artificial infatuation with crispness is likely to make our analyses go seriously askew and undermine the credibility of what we say by effectively marginalizing the very possibilities that make climate change so grave in the first place.

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

4 Whom or What Should a Person Believe?

The issue of how to deal with the deep structural uncertainties in climate change would be completely different and immensely simpler if systemic inertias, like the time required for the system to naturally remove extra atmospheric CO₂, were short, as is the case for many airborne pollutants like ozone, sulfur dioxide, and particulates. Then an important component of an optimal strategy might be along the lines of “wait and see.” With strong reversibility, an optimal climate-change policy should logically involve (among other elements) waiting to learn how far out on the bad fat tail the planet will end up, followed by midcourse corrections if we seem to be headed for a disaster. This is the ultimate backstop rebuttal of DT given by some critics of fat-tailed reasoning, including Nordhaus. Alas, the problem of climate change seems bedeviled almost everywhere by significant stock-accumulation inertias – in atmospheric CO₂, in the absorption of heat or CO₂ by the oceans, and in many other relevant physical and biological processes – that are slow to respond to attempts at reversal. It is a legitimate open question whether or not we can learn enough in sufficient time to make feasible midcourse corrections. Right now, we don’t know any technologies for rapidly removing existing atmospheric CO₂ that are even remotely viable on a large scale. When the critics are gambling on a midcourse-correction learning mechanism to undercut

the message of DT, they are relying more on an article of faith than on an evidence-based scientific argument.

Take atmospheric carbon dioxide as one specific example. The time it takes for an excess bulge of CO₂ to get purged from the atmosphere can be approximated by a weighted sum of exponential decay terms, each term representing a different CO₂ absorption process on a different time scale. Ballpark estimates imply that, for every unit of CO₂ anthropogenically added to the atmosphere, ≈70% remains after 10 years, ≈35% remains after 100 years, ≈20% remains after 1,000 years, ≈10% remains after 10,000 years, and ≈5% remains after 100,000 years.⁸ These numbers do not look to me like evidence supporting “wait and see” policies. The capacity of the oceans to take up atmospheric heat, and several other relevant mechanisms, tell a similar story of long stock-accumulation irreversibilities relative to the time it takes to extract and act upon meaningful signals of impending disasters.

The examples of possible midcourse-correction strategies cited by Nordhaus strike me as not sufficiently reliable in a context of imposing on the planet such an unprecedented GHG shock. Are we really so sure that over time we will understand sufficiently accurately future carbon-cycle dynamics, future oceanic uptake dynamics (for heat or CO₂), future atmospheric temperature dynamics and so forth, and that we will be able to invert the dynamics and extract clear enough signals in time to act? Is it technologically or politically possible to shut down all GHG emissions in a short period of time? Will we have viable technologies for removing existing stocks of CO₂ from the atmosphere? Even if an impending systemic-failure meltdown becomes known at some future time, and no expense is spared to avert it, would you want to rely as a backup on the reversibility strategies that Nordhaus mentions in his critique? The relevance of learning over future time is an important unresolved issue, difficult to model and not formally treated in any CBA, which in principle could decide this debate but needs to be researched much more thoroughly. Of course things could change, but for me now the built-in pipeline inertias and irreversibilities are sufficiently large that, if and when we detect that we are heading for disastrous climate change, there is an uncomfortable non-zero probability of being too late to do a lot about it (except, possibly, for lowering temperatures by geoengineering the atmosphere to reflect back incoming solar radiation, which has its own fat-tailed problems).

Nordhaus states that there are so many low-probability catastrophic-impact scenarios around that “if we accept the Dismal Theorem, we would probably dissolve in a sea of anxiety at the prospect of the infinity of infinitely bad outcomes.” This is rhetorical excess and, more to the point here, it is not convincing. In my article I listed what I consider to be the half-dozen or so serious contenders with climate change for potentially catastrophic im-

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

pacts with non-negligible probabilities: biotechnology, nanotechnology, asteroids, strangelets, pandemics, runaway computer systems, nuclear proliferation – and went on to give a few tentative reasons why I think that climate change is especially worrisome. It may well be that each of the other half-dozen or so serious candidates for fat-tailed disasters deserves its own ballpark estimates of tail probabilities along with extremely crude calculations of policy implications, which is about the best we can do with potential catastrophes. Even if this were true, however, it would not lessen the need to reckon with the strong potential implications of DT for CBA-like calculations in the particular case of climate change. The flaw in Nordhaus’s position here is in his trying to argue his case via guilt by association.

Some critics of DT promote alternative thin-tailed specifications that do not imply nearly such extreme expected outcomes as do my specifications. Their thin-tailed reduced-form specifications appear superficially to be plausible, and my fat-tailed reduced-form specifications (I hope) appear superficially to be plausible. They have credentials and fans of their approach, but so do I. Whom or what is a reader to believe? Of course the reader should weigh the plausibility of the arguments and the reasonableness of the specifications on their own merits. But it is difficult to form opinions about probabilities of climate-change extremes, or about disutility functions for extreme temperatures, or about lots of other relevant things for deciding the tail fatness of the PDF of $\ln D$. Suppose, for the sake of argument, that a policy maker believes the probability is 50% that my fat-tailed specification is correct and 50% that the thin-tailed specification of some “representative critic” is correct. Then rational policy should lean more in the direction of my fat-tailed conclusions than in the direction of the representative critic’s thin-tailed conclusions because of the highly asymmetric consequences of fat tails vs. thin tails. In this sense, whether it is fair or unfair, the playing field is not level between me and the “representative critic.” If one expert advises you that a fire insurance policy protecting your house against extreme losses is unnecessary because so few houses of your kind burn to the ground, while another expert advises you that a complete fire insurance policy is necessary in your case, should you flip a coin in deciding what to do just because both expert advisers seem equally credible?

As for Nordhaus’s framing of the issue that a combination of three (implicitly unlikely, in his mind) conditions must simultaneously be fulfilled in order to buy into what I am calling fat-tailed logic, I think it is a subjective judgement as to where the burden of proof lies. His point is essentially correct, but the issue is how to interpret it. To cut to the analytical core, the reduced form that Nordhaus must assume to justify his gradualist “policy ramp” is a thin-tailed PDF of $\ln D$ (and, very importantly, his reduced-form PDF must be thin tailed *after* integrating out the uncertainty in functional forms and structural parameters). Can anyone claim to know the extreme tail probabilities of the logarithm of the diswelfare

of high temperatures? Whether a fat-tailed or a thin-tailed PDF of $\ln D$ emerges, or does not emerge, from some particular combination of temperature-reactive disutilities, climate sensitivity, or wait and see policies, is essentially secondary. The primary issue that a reader must decide, from all of the evidence taken together, is which of us assumes more restrictive conditions and less plausible specifications overall than the other – in light of the welfare implications of fat vs. thin tails for the PDF of $\ln D$.

5 Concluding Comments

Taking fat tails into account has implications for climate-change policy. Qualitatively, fat tails clearly favor more aggressive policies to lower GHGs than the “standard” CBA. Alas, the quantitative implications are less clear. As this paper has stressed, the natural consequence of fat-tailed CBA is to exude less confidence and to convey an appearance of less robust policy advice than the “standard” thin-tailed CBA – especially concerning *how much* extra effort should be expended to lower GHGs.

Nordhaus summarizes his critique with the idea that deep uncertainties surround virtually every aspect of the natural and social sciences of climate change – but these uncertainties can only be resolved by continued careful analysis of data and theories. I heartily endorse his constructive attitude about the necessity of a research program targeted toward a goal of resolving as much of the uncertainty as is humanly possible. Future learning might well narrow the uncertainties faster than they expand, but we will not know this without an ongoing research effort. I would just add that we should also recognize the reality that, for now and perhaps for some time to come, the sheer magnitude of the deep structural uncertainties, and the way we express them in our models, will likely dominate plausible applications of CBA to the economics of climate change.

References

- [1] **Allen, Myles R. and David J. Frame.** “Call Off the Quest.” *Science*, 2007 (October 26), 318, pp. 582-583.
- [2] **Archer, David.** *Global Warming*. Blackwell Publishing, 2007.
- [3] **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).

- [4] **James Hansen, Makiko Sato, Puskher Karecha, David Beerling, Valerie Masson-Delmotte, Mark Pagani, Maureen Raymo, Dana L. Royer, James C. Zachos.** “Target Atmospheric CO₂: Where Should Humanity Aim?” *Open Atmos. Sci. J.*, 2, 217-231.
- [5] **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>).
- [6] **Nordhaus, William D.** “An Analysis of the Dismal Theorem.” Cowles Foundation Discussion Paper No. 1686, Yale University, January 2009 (available online at <http://cowles.econ.yale.edu/P/cd/d16b/d1686.pdf>).
- [7] **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.
- [8] **Weitzman, Martin L.** “On Modeling and Interpreting the Economics of Catastrophic Climate Change.” *Review of Economics and Statistics*, February 2009 (2009a).
- [9] **Weitzman, Martin L.** “Some Basic Economics of Extreme Climate Change.” Mimeo, February 2009 (2009b).