Book Review

A Review of William Nordhaus' *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*

Martin L. Weitzman

We all know the basics of climate change. Human activity has greatly increased the atmospheric stock of greenhouse gases (GHGs), the most important of which, by far, is carbon dioxide (CO_2). Once CO_2 shows up in the atmosphere, it stays there for a long time, on the order of centuries to millennia. The additional CO_2 in the atmosphere is due primarily to the massive burning of fossil fuels that has accompanied the industrial revolution, and it continues to increase at an ever accelerating pace. We are currently at atmospheric concentrations of CO_2 that were last seen over 3 million years ago. Scientists generally agree that increases in atmospheric CO_2 are extremely likely to lead to increased global warming and changes in climate (as well as other effects, such as significant alterations of ocean chemistry). The greater the CO_2 increase, the greater *and more uncertain* the likely warming, climate, and environmental response. If we continue emitting CO_2 at the current (or even a modestly reduced) pace, within a century or so we will likely attain atmospheric CO_2 levels that were last seen about 50 million years ago, when the mean global surface temperature was some 5°C warmer than today. Thus, by increasing atmospheric CO_2 we are performing a human-induced fast-paced experiment on a global scale that is extraordinarily far outside the range of "normal" experience, even by geological standards.

How, by how much, and how fast should we react to this unfolding global warming scenario? Attempts to answer such questions push the discussion into the realm of the economics of climate change.

Economics of Climate Change: A Problem from Hell

The economics of climate change is a problem from hell. Trying to do a benefit-cost analysis (BCA) of climate change policies bends and stretches the capability of our standard economist's toolkit up to, and perhaps beyond, the breaking point. First and foremost, disconcertingly large uncertainties are everywhere, including the most challenging kinds of deep structural

Review of Environmental Economics and Policy, volume 9, issue 1, winter 2015, pp. 145–156
doi:10.1093/reep/reu019

uncertainties. The climate change problem unfolds over centuries and millennia, a long intergenerational human time frame that most people are entirely unaccustomed to thinking about. With such long time frames, discounting becomes ultra-decisive for BCA, and there is much debate and confusion about which long-run discount rate should be chosen. Irreversibilities abound, including the very long residence lifetime of atmospheric CO_2. To add to the challenge, costs of new carbon-free technologies are uncertain. More importantly, for global mean temperature changes much above about 2°C, estimating damages is mostly educated guesswork with a distressingly wide error cone. The evaluation and aggregation of such damages add yet another significant layer of uncertainty; we are even unsure even about what *form* the "damages function" should take. Climate change due to high GHG levels involves nonnegligible tail risks of low-probability catastrophic outcomes, ranging from "known unknown" tipping points to the "unknown unknowns" of black-swan bad-feedback events that we cannot even imagine today.

Climate change is arguably the most serious international public goods problem that the world has ever faced, with global free riding galore and no solution in sight. The temptation to undertake ultra-cheap geoengineering (through the injection of stratospheric reflective aerosols like sulfur dioxide) as a quick fix for too-high global temperatures complicates an already supercomplicated situation. And what are the best *realistic* instruments and policies for dealing with climate change? I could provide many more examples of daunting challenges, but my ultimate point is that doing a BCA of climate change policies is a devilishly iffy proposition. Explaining all of this (and more) to a general/broad audience would seem to be an almost insurmountable hurdle.

Given all of the obstacles I have just described, it is little wonder that there has been no single good book to suggest to an intelligent motivated layperson who wants to get a balanced overview of the main issues in the economics of climate change. No good book until now, that is. William Nordhaus's *The Climate Casino* (2013) is without peer. It covers the basic aspects and issues of the economics of climate change, in about as accessible a form as one might conceivably hope for, especially considering how inherently complicated the subject is. Bill Nordhaus, who more than anyone else founded the economics of climate change and has been a major contributor to the subject over many decades, is a balanced centrist pragmatic observer who avoids extremes of right or left. Although Nordhaus is a gifted expositor, this serious scholarship-based book will not be an easy read for noneconomists. But the journey is well worth it. As Lawrence Summers writes in the book's jacket copy: "There is more insight and good-sense advice in this volume than in many libraries. This book should be as central to climate policy debates as climate change is to humanity's future."

Thus, all in all, *The Climate Casino* is a magnificent achievement. Its contribution is to lay out comprehensively and in a balanced scholarly fashion most major aspects of the current thinking on the economics of climate change. Its limitations derive from the inherent complexity of the subject itself and the alternative nuances and emphases that might be applied to an examination of the problem.

Princeton's Robert Socolow has started many a presentation by asking members of the audience to answer for themselves a version of the following two pop quiz questions. Question 1: "Is climate change an urgent problem?" Question 2: "Is getting the world off fossil fuels difficult?" His point is that most people will answer "yes" to one or the other question but not to both. An environmentalist might tend to answer "yes" to question 1 but "no" to question 2. Perhaps a majority of the public would answer "yes" to question 2 but "no" to question 1 (climate change may be a problem, but it is not an *urgent* problem). I believe that Nordhaus (and Socolow, and I) would answer "yes" to both questions, which might seem like some sort of cognitive dissonance to most people. This means that, right from the start, the main thesis of *The Climate Casino* is going to be a hard sell for most people. In Nordhaus-speak, climate change is a trillion dollar problem that requires a trillion dollar answer. Clearly, this is not a message that many people will want to hear.

Limitations of Integrated Assessment Models

A lot of the scholarly argument in this book is phrased in terms of (or based on) Nordhaus's pioneering integrated assessment model (IAM) called DICE (dynamic integrated climate-economy). An IAM combines a simplified aggregate model of climate change with a simplified aggregate model of the economy. DICE is probably the most famous and widely employed IAM, due in no small part to Nordhaus's public-spirited efforts to make it widely available and user friendly so that critics can readily experiment with the model or even play with their own variant of the equations and parameter values. Nordhaus understands full well that IAMs like DICE have drawbacks, mostly concerning the high level of uncertainty about critical specifications (especially at high GHG concentrations). He also knows that different IAMs, or even different specifications of a single IAM like DICE, can produce different outcomes that may be highly sensitive to some of the underlying assumptions. Nordhaus provides the reader with a fair enumeration of the possible drawbacks and sensitivities of an IAM like DICE.

Nordhaus argues only that having the "fuzzy telescope" of an IAM like DICE is better than having no telescope at all. My concern is that he fails to adequately convey the *relative magnitude* of fuzziness of *this particular* telescope. This is an important point to grasp, because an IAM-inspired BCA of climate change policy is much fuzzier than most ordinary BCAs, having mostly to do with the incredible challenges of stretching economic analysis to its limits, which I alluded to earlier when describing the economics of climate change as the "problem from hell." This is a real issue—one that is especially acute for analyzing the consequences of high concentrations of GHGs. People want to know what will happen and how much it will cost. For something like, say, the social cost of carbon (SCC), the public (and policy makers) want *a number*. Not the *theory* of a number, and not even a *range* of numbers, but an actual number. Thus, it is difficult for an IAM to deliver anything that remotely resembles the degree of certitude that people want. Bill Nordhaus is the kind of constructive guy who would rather light a candle than fumble around in the dark. But the reality, I fear, is that the candle may shed but a feeble light. The public wants precision, and, for expositional purposes, an author like Nordhaus, who in this book is trying to simplify and elucidate the economics of climate change to a general audience, needs something that *looks* like precision. However, the economics of climate change is sufficiently uncertain that the desired precision is just not there, especially for high GHG levels. This is a frustrating situation, and I do not know how we can best resolve it. The public wants answers that are robustly free of uncertainty, but the economics of climate change is largely about, and the range of possible answers can largely depend upon, deep structural uncertainties that are impossible to dismiss. Despite paying homage to uncertainty (titles like "Climate Casino" and "DICE" say it all), IAMs like DICE come up with precise-looking outcomes, even though many of these outcomes are highly sensitive to the assumptions underlying the model. (I would not go nearly as far as the outright rejection of IAMs advocated in Pindyck [2013], but the interested reader might well have a look at his arguments.)

The first computer models of climate change (and the first IAMs) were essentially deterministic. This approach is natural and understandable—they first deal with the case of no uncertainty, do some sensitivity analysis, and leave until later a formal incorporation of uncertainty. In the deterministic formulation, the key issue is the tradeoff between less present consumption (due to greater investment in current CO_2 abatement) and more future consumption (due to less climate change damages because of lower concentrations of future atmospheric CO_2). A deterministic-like IAM is thus addressing a capital-theoretic optimal growth problem with an exotic production function that includes climate change dynamics.

The Key Issue of Tail Risks

Don't get me wrong; the deterministic-based approach to analyzing climate change, pioneered by Nordhaus in early versions of DICE, provides many valuable insights. What concerns me is the nagging sense that, at high GHG concentrations, the main issue in the economics of climate change is more about buying insurance against really bad tail risks (i.e., the small chance of catastrophic climate change) than about deterministic consumption smoothing over time. These tail risks of high-GHG climate change involve issues of deep structural uncertainty that could be decisive in thinking about climate policy, but they are difficult to get a handle on quantitatively because they are so far outside the range of "normal" modeling experience. Once again, Nordhaus is well aware of these tail-risk issues, and he attempts to honestly report the problem (including the results of some rough Monte Carlo simulations). However, I don't think that the possible magnitude of this tail-risk issue or its potential to influence (and, at high GHG concentrations, to possibly dominate) policy outcomes is adequately conveyed to the reader, who might consult Litterman (2013) and/or Weitzman (2014) for a more full discussion and further references.

Let me illustrate some of the issues here with an enormous oversimplification. People would clearly like to have a big-picture sense of what will happen as CO_2 concentrations increase. Thus, a natural question for many people to ask would be along the lines of: "What will be the change to the world's climate, and the resulting damages to welfare, as a function of various concentrations of atmospheric CO_2?" A more precise and easier to answer question is: "What will be the eventual change to global average steady-state surface temperature as a function of various steady-state concentrations of atmospheric CO_2?" It turns out that equilibrium temperature change is proportional to the *logarithm* of equilibrium concentrations of atmospheric CO_2. This means that the question can be simplified to asking: "What will be the eventual change to global average steady-state surface temperature from a steady state *doubling* of atmospheric CO_2?" The science tells us that a doubling of CO_2 from *any* starting position should give rise to approximately the same ultimate temperature-change response. This equilibrium temperature-change response to an equilibrium doubling of CO_2 is called "climate sensitivity."

Climate Sensitivity: A Thought Experiment

Climate sensitivity is critical because it provides a big picture of the eventual macro-level response of global average temperatures to CO_2 concentrations. The underlying thought experiment is to artificially double atmospheric CO_2, keep it perpetually at that doubled level, and then record the (asymptotic) equilibrium global temperature response. This thought experiment misses a *lot* of important details, most notably concerning the timing of climate change. For example, because the oceans can absorb a great deal of heat, very high atmospheric temperature increases, like above 4–5°C, can take a very long time to fully equilibrate, on the order of centuries. Despite its limitations, the concept of climate sensitivity does give people a number that they can wrap their minds around, and it has an almost iconic status in the climate change community. So what is climate sensitivity? Like so much else in this field, climate sensitivity is uncertain. How uncertain?

In 1979, the first of its kind National Academy of Sciences Ad Hoc Study Group on Carbon Dioxide and Climate concluded that the best estimate of climate sensitivity is 3°C, give or take 1.5°C. The reasoning underlying this estimate was casual by today's standards. Jule Charney, the study's lead author, looked at the two estimates from the two (deterministic) climate models available at that time. One model said 2°C, and the other model said 4°C. Charney averaged these two estimates to get 3°C and added half a degree centigrade to both ends of the two underlying estimates to round out the range to [1.5°C, 4.5°C] because—you guessed it—there is, after all, uncertainty about climate sensitivity.

Fast forward 35 years to today. The IPCC Fifth Assessment reported in 2014 that climate sensitivity is likely between 1.5°C and 4.5°C. Thus, after 35 years of intensive research, the range of climate sensitivity has not been narrowed. It seems that the more we understand and resolve old questions and doubts about climate sensitivity (over the course of the last 35 years), the more new questions and new doubts arise. The IPCC effectively uses the word "likely" to indicate that the probability of climate sensitivity being between 1.5°C and 4.5°C is greater than 66% but less than 90%. However we interpret it, this probability range leaves a worrisome amount of probability in the upper tail of the probability distribution of climate sensitivity (i.e., above 4.5°C).

The Risk of Catastrophic Climate Change

The preindustrial revolution level of atmospheric CO_2 was about 280 parts per million (ppm). The current level of CO_2 is 400 ppm and currently increasing at an accelerating pace of more than 2 ppm per year. If the world goes to 560 ppm (a doubling of preindustrial concentrations), which as of now seems to me to be an optimistic scenario, the "likely" range of *eventual* temperature response is from 1.5°C to 4.5°C. Counting other greenhouse gases, the International Energy Agency (IEA) estimates on the basis of current stated emissions intentions of various countries that the world will end up somewhere around 700 ppm by 2100—two-and-a-half times preindustrial levels—unless major emitters take drastic additional steps over and above their stated intentions. At a level of 700 ppm, the "likely" range of *eventual* temperature response is from 2°C to 6°C. No one is sure what the planet would be like with 6°C (=11°F), or even 4.5°C (=8°F), of global warming, but the consequences of an average planetary surface temperature change of 6°C (or even 4.5°C) is truly frightening and I think qualifies as a global catastrophe. I don't believe that this scary tail-risk aspect of high GHG concentrations is quite adequately conveyed by *The Climate Casino*.

There are lots of legitimate objections to the back-of-the-envelope calculation that I just performed so nonchalantly. A global average surface temperature of 6°C would only be attained after a couple of centuries. And then there are all of the usual issues concerning evaluation of damages, discounting, risk aversion, technological progress, actionable learning over time, and so forth and so on. My simplistic comparison of steady states here is not a careful economic analysis, by any stretch of the imagination. But my own conclusion, based on this crude exercise, is that we do not want to be on a trajectory that gets us anywhere near these kinds of tail risks, even if their full impact may not be felt for a century or two from now. In other words, I believe that high GHG concentrations likely pose a threat to be avoided more because of the *bad-tail* consequences than because of the *most likely* consequence. Thus, I justify this crude climate sensitivity exercise because I think it is the simplest—and starkest—way to illustrate a basic point about long-run tail risks under high GHG concentrations, without getting overwhelmed by the scores of equations in a full blown IAM. (For more details about my estimates and projections here, see Wagner and Weitzman [2015].)

Managed vs. Unmanaged Systems

In his discussion of climate change effects, Nordhaus makes a useful distinction between what he calls "managed systems" and "unmanaged systems." A managed system is one in which human intervention can largely alter the effects of climate change through adaptive reactions. Examples might be indoor living, manufacturing, or even (to some extent) agriculture and public health. By contrast, in unmanaged systems humans are largely unable to alter the impacts of climate change. Examples of these types of impacts include sea level rise, hurricanes, monsoons, ecosystem changes (such as those due to ocean acidification), and loss of biodiversity. Nordhaus is more optimistic about the ability of humans to adjust the effects of climate change in managed than in unmanaged systems. It is also less difficult to evaluate the effects of moderate climate change on managed than on unmanaged systems.

Specifying the "Damages Function"

To do a complete BCA, numbers have to be attached to various damages. In DICE, as in most other IAMs, damages are modeled in terms of a loss of some fraction of potential output, with the fraction increasing as global mean surface temperature increases. So increased temperatures act as if they cause some increased fraction of output to simply evaporate. I am not at all sure that this is the right way to conceptualize or to model climate change damages. This multiplicative-evaporation form makes it relatively easy to substitute increased output for the environmental disamenity of increased temperatures. Moreover, I suspect that the fractional damages function itself is much too low for high temperature changes. In terms of welfare equivalents, the quadratic damages function that Nordhaus employs in DICE makes a temperature change of 6°C correspond to a loss of output of less than 10%. When I think of all the damages that a 6°C global-average temperature change could wreak, this estimate seems much too sanguine. Thus, it is no wonder that simulations based on such a specification are relatively unconcerned about high eventual accumulations of GHGs, as long as the world economy is allowed to grow exponentially at something close to today's rates.

A lot could be riding on the specification of the damages function. If the effect of increased temperatures is on the *growth rate* of output, rather than the *level* of output, which some studies and authors suggest (see, e.g., Dell, Jones and Olken [2014]), then this specification change could have a significant effect on raising the estimated damages from elevated temperatures. Similarly, if environmental amenities and conventional output should have an elasticity of substitution less than one (instead of being essentially equal to one in the conventional multiplicative specification of a damages function), then, as Sterner and Persson (2008) have shown, an optimal policy puts much more of a brake on CO_2 emissions. Or maybe damages are sensitive to the *rate of change* of temperatures, rather than (or in addition to) their absolute level. Or perhaps global average surface temperature is not even the right variable to plug into a damages function for climate change. Although it is unlikely that any of these different specifications will make all that much of a difference for low temperature changes, I can easily imagine them being decisive for the high-GHG, high-tail-risk temperature, and climate changes that concern me most. The simple fact of the matter is that we do not know how to specify or parameterize damages at high GHG or high temperature levels. We simply don't know the right form for the damages function, and, even if we did, we are forced to use the crudest kind of extrapolation for temperatures that are much greater than about 2°C.

Climate Change Catastrophes: Known Unknowns and Unknown Unknowns

What are the climate change "catastrophes" that might materialize at global average temperature changes of 6°C, or 4.5°C, or even, perhaps, 3°C? The list of "known unknown" bad consequences includes ice-free summers in the Arctic Ocean, rapid sea level rise from possible rapid melting of the Greenland and West Antarctica ice sheets, disruption of large-scale weather patterns such as monsoons or desertification, intensification of hurricanes and typhoons, significant alteration of ocean circulation patterns, melting of arctic permafrost with widespread bad-feedback releases of methane and carbon dioxide, and severe diebacks of tropical rainforests (like the Amazon rainforest) and northern taiga (like Siberian forests). This list could easily be extended to include other "known unknowns." Nordhaus covers several of these possibilities in a chapter entitled "Tipping Points in the Climate Casino." The problem is that we don't know how to evaluate the probabilities and damages for these kinds of tail risks. I am also concerned about the possibility, at high GHG concentrations, of "unknown unknowns" (sometimes called "black swan" events), for example, in the form of runaway bad-feedback tail events that we have not yet identified or are currently unaware of (like higher GHG levels begetting higher temperatures, which beget yet higher GHG levels).

Nordhaus is open about these kinds of tail-risk uncertainties. But, again, I don't think that he adequately conveys to the reader the possibility that insurance against these high-GHG catastrophic risks should be a major factor in, and conceivably dominate, the policy analysis of climate change. I believe that it is at least (and maybe even more) important to convey the tail risk consequences of high-GHG scenarios, as it is to convey the most likely outcomes. I think the book's exercises with DICE do not adequately emphasize this tail-risk insurance aspect.

One of the biggest uncertainties and challenges at high GHG levels is estimating the extent of, and the evaluation of, the loss of the natural environment as we know it. I can imagine a high-GHG scenario in which humans perhaps squeak by in adapting indoor activities to climate change, but the outdoor biosphere is ruined by what some have called the sixth mass extinction. IAMs mostly just ignore this problem, because it is so difficult to quantify. How do we measure the damages from the loss of wildlife and species? Who performs the economic evaluation? Here I suspect there is likely to be enormous variation in estimates depending on who does the evaluation.

Discounting Controversies

Of course if the relevant interest rate is high enough, then even big damages occurring centuries from now can be rendered small in today's present discounted values. There has been a vigorous debate about the proper way to discount climate change damages. Some of the controversy is centered around the issue of whether the discount rate should be derived from a prescriptive (normative) approach, or whether it should rely on a descriptive (opportunity-cost) approach. Nordhaus presents both sides of this debate fairly, but in the end he comes down strongly in favor of the opportunity-cost approach to discounting. He generally uses a real interest rate of around 4 percent per year. While I am not prepared to challenge this opportunity-cost view of discounting head on here, I am perhaps ever so slightly more agnostic than Nordhaus on this issue. Climate change may be sufficiently "different" from other commodities that its intergenerational evaluation over time should also be treated differently. As an economist I am trained to be suspicious about looking at things in this way, but perhaps there is *something* to viewing aspects of the natural environment as a kind of intergenerational merit good to which every generation has some rights.

On the issue of long-term discounting, there is a well-developed argument that if future discount rates are uncertain but likely correlated (e.g., discount or growth rates are mean reverting to a future mean that is itself uncertain), then the effective discount rate should decline over time, possibly to quite low asymptotic values over the course of centuries. This declining discount rate (DDR) argument has become sufficiently mainstream that a treatment of climate-change discounting is generally viewed as incomplete without it. (For a discussion of DDR issues, see Arrow et al. [2014] and the many references cited therein.)

In the context of climate-change discounting there is additional controversy about the role of effects somewhat akin to a climate-change "beta" on the appropriate discount rate. On one side, the argument is that, other things being equal, if investments in CO_2 abatement will help to prevent catastrophic outcomes, then the appropriate discount rate should be lower because potentially worse states of the world are associated with a greater payoff of incremental consumption. The counterargument is that higher economic growth rates, and therefore higher future consumption, are associated with higher damages from correspondingly higher GHG levels, which suggests that a higher discount rate is called for. (This seemingly arcane set of issues is lucidly discussed in Litterman [2013].) The risk-free interest rate is about 1 percent per year, while the economy-wide average rate of return on capital is about 7 percent per year. This leaves a wide range for a climate-change discount rate depending on what part of the probability distribution of damages is being affected under which circumstances, about which we know very little.

The Climate Casino contains a good general discussion of discounting, but given the critical importance of this topic for the economics of climate change, I think that Nordhaus might have addressed more of the issues noted in the above paragraphs, at least to the point of mentioning them. I realize that some of these issues are highly technical, but the counterbalance to such a technical discussion is the importance of the reader realizing that the choice of a discount rate is itself one of the most significant (and controversial) uncertainties in the economics of climate change. The social cost of carbon, for example, is massively affected by the discount rate that is applied.

The Role of Geoengineering

The Climate Casino discusses the role of geoengineering, specifically solar radiation management (SRM) through the injection of aerosols, such as sulfur dioxide, into the stratosphere. A geoengineered "sunshade" has a long list of things going against it. Some of the negatives include continued ocean acidification, possible lowering of stratospheric ozone, dependency effects, changed regional weather patterns, and a possible weakening of resolve to cut GHG emissions. However, geoengineering is estimated to be so cheap (probably less than $10 billion per year to offset the global temperature effects of a doubling of CO_2), and it lowers temperatures so fast that essentially any determined country (with even a medium-sized economy) could, if unopposed, put up a geoengineered sunshade on its own, to respond to its own perceived need to lower global temperatures and change its own climate quickly. This is a true "twin" externality to the conventional externality of emitting greenhouse gases. The conventional CO_2 emissions externality has a serious "free rider" problem because severely cutting GHG emissions is sufficiently expensive that it makes it difficult to attain meaningful global agreement on apportioning compliance costs. In contrast, a geoengineered sunshade could be called a "free driver" problem because lowering global average temperatures via this route is so cheap and so quick that, in principle, many countries could do it unilaterally to meet their own particular perceived needs, thereby imposing a potentially dangerous "public bad" on some other nations. Thus, the world faces not one, but two, global externalities from climate change. Nordhaus treats SRM geoengineering more as a social planning problem. Perhaps greater emphasis should have been placed on this double externality problem of climate change (free riding *and* free driving), how this double trouble creates even more tail-risk uncertainty, and what might be done about it. (For more details on the view of geoengineering presented here, see Wagner and Weitzman [2015].)

My Overall Assessment

Let me try to summarize my assessment of *The Climate Casino*. It is a great book in a class of its own. There is no other comparable textbook-like but relatively simple book on the economics of climate change. Nordhaus covers almost every important aspect of the economics of climate change in a scholarly, balanced, and intellectually honest way. But it will not be an easy read for someone without at least some minimal previous exposure to economics and/or a high degree of motivation. I learned a lot from this book, and I think that any professional economist interested in climate change would profit greatly from reading it.

My main concerns with the book are related to differences in nuance and emphasis on aspects of climate change where the debate among economists is already alive and well.

Much of Nordhaus's analysis is based on DICE. My main criticism here has to do with the way IAMs treat potentially catastrophic tail risks with big impacts and small probabilities. I think that the main purpose of keeping GHG levels under control is to buy insurance against extreme bad outcomes. At high GHG levels, the structural uncertainties associated with IAMs are great enough to make the models highly sensitive to some seemingly arbitrary assumptions about high-temperature damages, long-term discounting, and several other important features about which we are unsure and sometimes highly unsure. Nordhaus clearly understands that major uncertainties and disagreements remain about the results produced by computer models. Nevertheless, he believes that we have no choice but to rely on a "fuzzy telescope" in making our best guesses and wisest choices. I do not necessarily disagree, but I think we may need to emphasize more the deep fuzziness that arises from deep structural uncertainties, especially concerning the evaluation of tail risks from high GHG concentrations.

The "standard" BCA approach to climate change appears to offer a constructive ongoing scientific-economic research program for generating ever more precise outputs from ever more precise inputs. By contrast, my emphasis may seem off-putting because it can be painted as antiscientific and antieconomic. Tail risks and the resulting limitations they impose on the ability of BCA to reach robust conclusions are frustrating for economists. After all, we make a living from plugging rough numbers into simple models and reaching specific conclusions (more or less) on the basis of those numbers. What quantitative advice are we supposed to provide to policy makers and politicians about how much to spend on averting climate change if the conclusions from modeling tail-risk uncertainties are not clear-cut?

Men and women of action have a low tolerance for vagueness and crave some kind of an answer. Thus, they have little patience for even a whiff of fuzziness from two-handed economists. It is threatening for us economists to admit that constructive "can do" climate-change BCAs and IAMs may be up against some basic limitations on the ability of quantitative analysis to yield robust policy advice. But if this is the way things are with the economics of climate change, then this is the way things are. Nonrobustness to subjective assumptions about high-GHG damages, discounting, and other aspects of potentially disastrous outcomes is an inconvenient truth to be lived with rather than a fact to be denied or evaded just because it looks less scientifically objective in BCA. Alas, this limits the ability to give detailed and concrete answers to an impatient public. What remains is a somewhat generic warning to keep well away from high GHG concentrations.

BCA is valuable, even indispensable, as a disciplined framework for organizing information and keeping score. But not all BCAs are created equal. In rare situations with effectively unlimited downside liability, like high-GHG climate change, BCAs can be fragile to the specifications of extreme tail events. Perhaps economists need to emphasize more openly to policy makers, politicians, and the public that while formal BCA can be helpful, in the particular case of climate change there is a danger of possible overconfidence from an over-reliance on subjective judgments about the probabilities and welfare impacts of extreme events. What we can do constructively as economists is to better explain both the magnitudes of the unprecedented structural uncertainties involved in high-GHG scenarios and why this feature limits what we can say, and then present the best BCAs and the most honest sensitivity analyses that we can under tail-risk circumstances, including trying many different functional forms for extremes.

At the end of the day, policy makers must decide what to do on the basis of an admittedly sketchy economic analysis of a gray area that just cannot be forced to render clear robust answers. The bottom line is that, under the extreme tail uncertainty associated with high GHG levels, seemingly casual decisions about functional forms, parameter values, discount rates, tail fatness, and so forth can dominate BCA. Economists should not pursue a narrow, superficially crisp analysis by blowing off the low-probability high-impact catastrophic scenarios associated with high GHG levels, as if this were a necessary price we must pay for the worthy goal of giving answers and advice to policy makers. In fact, this artificial infatuation with crispness can make our analyses go seriously awry and undermine the credibility of what we have to offer by effectively marginalizing the very possibilities that make high-GHG climate change so grave in the first place.

Qualitatively, considerations of tail risk favor more aggressive policies to lower GHG concentrations than when the more “standard” BCAs of climate change are used. I believe that this issue of tail-risk insurance is an important part of the reason why we want to keep GHG levels from ballooning. Alas, the quantitative implications are less clear. In addition to other uncertainties that I have mentioned, much can depend on the specification of welfare – whether it is expressed in the form of the workhorse of expected present discounted utility or more exotic preferences that can easily amplify the role of risk aversion independent of time, or even involve the thorny issue of Knightian-uncertainty aversion. And of course, no matter what form we choose for welfare evaluation, the parameter values matter when evaluating welfare. The natural consequences of tail-risk uncertainty should be to make economists less confident about climate-change BCA and to make them adopt a more modest precautionary-insurance tone that befits less robust policy advice. My own sense is that the sheer magnitude of the deep structural uncertainties concerning high-GHG catastrophic outcomes, and the way we express this in our models, is likely to influence plausible applications of BCA to the economics of climate change.

So let us use the fuzzy telescope described in *The Climate Casino*. But let us also be aware of the need to purchase extra mitigation insurance against the extreme tail-risk uncertainty of high-GHG scenarios, even though it is inherently difficult for us to say exactly how much insurance we need.

Martin L. Weitzman
Professor of Economics, Harvard University
October 18, 2014

References

Arrow, Kenneth *et al.* 2014. Should governments use a declining discount rate in project analysis? *Review of Environmental Economics and Policy* 8 (2): 145–63.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2014. What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature* 52 (3): 740–98.

Litterman, Robert. 2013. What is the right price for carbon emissions? *Regulation* 36 (2): 38–43.

Nordhaus, William. 2013. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*. New Haven, CT: Yale University Press.

Pindyck, Robert S. 2013. Climate change policy: What do the models tell us? *Journal of Economic Literature* 51 (3): 860–72.

Sterner, Thomas, and U. Martin Persson. 2008. An even sterner review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy* 2 (1): 61–76.

Wagner, Gernot, and Martin L. Weitzman. 2015. *Climate Shock*. Princeton, NJ: Princeton University Press (forthcoming).

Weitzman, Martin L. 2014. Fat tails and the social cost of carbon. *American Economic Review Papers and Proceedings* 104 (5): 544–46.