A R T I C L E

Uncertainty

by Daniel A. Farber

Daniel A. Farber is the Sho Sato Professor of Law and Chair of the Energy and Resources Group, University of California, Berkeley.

Introduction

Our society has sophisticated techniques for analyzing risks that can be modeled and quantified. But other threats—often the most serious ones—do not fit the paradigm. These threats involve what the economist Frank Knight classified as "uncertainty" (where the likelihood of the peril is nonquantifiable) as opposed to "risk" (where the likelihood is quantifiable).[1] Uncertainty is particularly pernicious in situations in which catastrophic outcomes are possible, but conventional decision tools are not equipped to cope with these potentially disastrous results; neither the risk analysis favored in the United States, nor the precautionary principle utilized by Europeans and others, is satisfactory in cases of uncertainty. This Article considers how we can use new advances in economics and decision theory to do better.

Economic modeling and policy analysis are often based on the assumption that extreme harms are highly unlikely, in the technical sense that the "tail" of the probability distributions is "thin"—in other words, that it approaches rapidly to zero. Thin tails allow extreme risks to be given relatively little weight. A growing body of research, however, focuses on the possibility of fat tails, which are common in systems with feedback between different components. As it turns out, determining the precise "fatness" of the tails is often difficult, which causes models involving fat tails to blur from risk into uncertainty.

This Article proposes the "α-precautionary principle" for use when—because of fat tails or otherwise—decisionmakers cannot quantify risks and face Knightian uncertainty. The α-precautionary principle is more nuanced than conventional versions of the precautionary principle though still remaining attentive to possible catastrophic outcomes and simple enough for easy application. For instance, the α-precautionary principle suggests a highly precautionary approach to the uncertainties surrounding climate change but a less precautionary approach to the uncertainties of nanotechnology.

The complete version of this Article was originally published in 99 Geo. L.J. 901 (2011). *It has been excerpted with permission. © 2011 by Georgetown University and The Georgetown Law Journal; Daniel A. Farber. Reprinted with permission.*

1. Frank Knight, Risk, Uncertainty and Profit (1921). Uncertainty also played a central role in the thought of John Maynard Keynes.

The new techniques advanced in this Article occupy a middle space between conventional versions of risk assessment and the precautionary principle, using mathematical tools to help decisionmakers cope with uncertainty, but not requiring the assignment of precise probabilities when doing so would be inappropriate.[2]

I. Current Approaches to Environmental Risks and Uncertainties

The regulatory system often addresses probabilistic harms. Conventional risk assessment—the dominant mode in the United States—is a powerful methodology, but over-reliance on it can lead to a failure to acknowledge any risks that do not lend themselves to the technique. Risk analysis requires that risks be quantified, but it is not always possible to obtain the necessary reliable estimates of probabilities. A focus on conventional risk analysis can therefore lead to disregard of nonquantifiable risks. This can bias decisionmaking and mislead the public about the possible consequences. Indeed, a policy of ignoring all nonquantifiable harms is literally a recipe for disaster.[3]

In contrast, the European Union and other nations are less wedded to quantitative risk assessment than the United States. Instead, the E.U. favors the use of the precautionary principle, which does address unqualified possible harms, but functions more as a source of sound advice than as a method of analysis.

In its most general sense, the precautionary principle advises that lack of certainty is not a justification for inaction in the face of possible risks[4]; more precise statements of the principle focus on situations involving nonquantifiable harms, irreversible harm, or catastrophic harm.[5] This principle has been explained on the basis of risk aversion

2. The fundamental research discussed in this Article is rapidly developing, and work on practical applications is at an even earlier stage. Thus, the conclusions discussed in this Article must be considered preliminary.
3. For example, apparently in the belief that a problem is not significant unless it can be precisely quantified, the Nuclear Regulatory Commission refuses to discuss the possibility of terrorist attacks on nuclear facilities in its environmental impact statements (EIS) because the risk cannot be quantified. *See, e.g.*, Private Fuel Storage, L.L.C., 56 N.R.C. 340, 350-51 (2002). For further discussion, see Daniel A. Farber, *Uncertainty*, 99 Geo. L.J. 901, 909-14 (2011).
4. Jonathan Remy Nash, *Standing and the Precautionary Principle*, 208 Colum. L. Rev. 494, 498-99 (2008).
5. *Id.* at 502-03.

or skepticism about the environment's ability to tolerate damage.[6] The implication of the precautionary principle is that it is better to overregulate than underregulate new technologies—but this can actually result in more harm to public health or welfare under some circumstances.[7]

Despite its broad international acceptance, the precautionary principle is controversial.[8] There seem to be three main criticisms. The first is its vagueness, or "squish[iness]."[9] However, this vagueness critique may be overstated, as a number of efforts have been made to sharpen the precautionary principle in certain settings including where there is uncertainty rather than simply risk and where harm would be "catastrophic."[10] A second criticism of the precautionary principle is that government intervention creates risks of its own.[11] A third criticism connects the precautionary principle with defects in human cognition. Cass Sunstein has argued that when the precautionary principle "seems to offer guidance, it is often because of the operation of probability neglect,"[12] meaning the cognitive incapacity of individuals to attend to the relevant risks.[13] Supporters of the precautionary principle respond that it is actually needed to counter defects in the ways people process probability information. Rather than being part of the problem of limited human rationality, the precautionary principle may be part of the treatment.[14]

While the debate will continue, it may be possible to find consensus on narrower ground, particularly as to a special form of precaution for the uncertainty associated with catastrophic risks. Sunstein, for instance, though a critic of the precautionary principle, nonetheless recognizes that catastrophic risks may be different.[15] He proposes a number of different versions of the catastrophic risk precautionary principle, in increasing order of stringency. Sunstein's observations point helpfully in the right direction but identifying those techniques and clarifying their domain requires further work, and current developments in economics and decision theory allow us to put some flesh on the concept of a catastrophic precautionary principle.

II. Understanding Catastrophic Uncertainty

In many situations, risk falls near the average, such that upside deviations are roughly as likely as downside deviations and extreme deviations are extremely unlikely. These situations are relatively tractable in policy terms, but some issues require much more attention to potential extreme outcomes. One way of understanding the problem begins with the concept of feedback effects.[16] Consider the familiar example of the feedback between a microphone and loudspeakers. If the system is already experiencing a bit of feedback, turning the amplification slightly downward provides only modest benefits, while turning it slightly upward can result in an unnerving shriek from the speakers. Thus, uncertainty about the exact amount of feedback is mostly significant because of the risk that feedback will be higher than expected, resulting in much more noise, rather than the possibility that the feedback will be lower and the noise will be a bit more subdued. The implication is that uncertainty is greatest where it matters most, in terms of extreme events.[17] This section discusses decisionmaking in situations where even rough quantification of probabilities is not feasible.

A. *Fat-Tailed Distributions and Catastrophic Outcomes*

When probabilities form a bell curve (normal distribution), most events are bunched near the average and extreme outcomes fade away quickly.[18] The term *fat tails* is used to describe systems that have a higher likelihood than the normal curve of extreme outcomes—in a graph, the tail of the distribution does not thin out as quickly as the normal distribution.[19]

A common version of fat tails is found in the statistical distribution called a "power law."[20] Rather than following the familiar bell-curve distribution, complex systems often

6. *See* Daniel Farber, Eco-Pragmatism: Making Sensible Environmental Decisions in an Uncertain World 170 (1999).
7. *See* Jonathan H. Adler, *More Sorry Than Safe: Assessing the Precautionary Principle and the Proposed International Biosafety Protocol*, 35 Tex. Int'l L.J. 173, 195-98 (2000).
8. For a recent update on the debate, see Fritz Allhoff, *Risk, Precaution, and Emerging Technologies*, Stud. in Ethics L. & Tech. (Aug. 2009). Allhoff suggests that "precaution supplements cost-benefit analysis given uncertainty." *Id.* at 23.
9. Edward A. Parson, *The Big One: A Review of Richard Posner's Catastrophe: Risk and Response*, 45 J. Econ. Literature 147, 152 (2007) (citing Richard A. Posner, Catastrophe: Risk and Response (2004)).
10. Jonathan Remy Nash, *Standing and the Precautionary Principle*, 208 Colum. L. Rev. 494, 503 (2008) (footnote omitted).
11. *See* Adler, *supra* note 7, at 195; Frank B. Cross, *Paradoxical Perils of the Precautionary Principle*, 53 Wash. & Lee L. Rev. 851, 863-75 (1996) (describing risks created by alternative activities).
12. Cass R. Sunstein, *Probability Neglect: Emotions, Worst Cases, and Law*, 112 Yale L.J. 61, 94 (2002).
13. *Id.* at 62-63. Sunstein further elaborated his critique in Cass R. Sunstein, *Beyond the Precautionary Principle*, 151 U. Pa. L. Rev. 1003 (2003).
14. *See* David A. Dana, *A Behavioral Economic Defense of the Precautionary Principle*, 97 NW. U. L. Rev. 1315, 1327-28 (2003) (arguing that the principle may "result in the generation of more information" and may "provide advocates of regulation with a discursive tool to increase the amount of information generated and the quality of the analysis of that information"). Dana elaborates his position in David A. Dana, *The Contextual Rationality of the Precautionary Principle*, 35 Queen's L.J. 67 (2009) [hereinafter Dana, *Contextual Rationality*].
15. *See* Cass R. Sunstein, *The Catastrophic Harm Precautionary Principle*, Issues in Legal Scholarship (2007), *available at* http://www.bepress.com/ils/iss10/art3.
16. *See generally* Jainguao Liu et al., *Complexity of Coupled Human and Natural Systems*, 317 Science 1513 (2007).
17. For those whose taste runs to equations and numerical examples, this point is mathematically expressed in Daniel A. Farber, *Uncertainty*, 99 Geo. L.J. at 921-22 (2011).
18. This can be seen from the graphs in Eric W. Weisstein, Wolfram Mathwolrd, *Normal Distribution*, http://mathworld.wolfram.com/NormalDistribution.html (last visited June 16, 2012).
19. *See, e.g.*, William Safire, *Fat Tail*, N.Y. Times Mag., Feb. 8, 2009, at 24.
20. For an introduction to power laws, see Mandred Schroeder, Fractals Chaos, Power Laws: Minutes From an Infinite Paradise 103-19 (1991).

at least approximately follow power-law distribution,[21] in which the probability of an event is given by its magnitude taken to a fixed negative exponent.[22] "[T]he distinguishing feature of a power law is not only that there are many small events but that the numerous tiny events coexist with a few very large ones."[23] Such outliers are much less likely when a normal distribution is involved. Power laws conflict with our usual view of the world as consisting of small fluctuations around routine outcomes.

While the existence of fat tails clearly has relevance to policy, we do not have "a commonly accepted usable economic framework for dealing with these kinds of thick-tailed extreme disasters"—partly because these "probability distributions are inherently difficult to estimate."[24] The reason that the probabilities are difficult to estimate is that data will rarely include instances from the tail (because the events are rare), making it impossible to estimate just how quickly the tail tapers off.

Martin Weitzman has shown on the basis of general considerations of statistical and economic theory that it often "is difficult to infer (or even to model accurately) the probabilities of events far outside the usual range of experience" and that this ultimately leads to a fat-tailed probability distribution of utility losses.[25] Weitzman also shows that even if the "true" probability distribution has a thin tail, the decisionmaker may still be faced with a fat-tailed distribution as a practical matter because it is impossible to get enough evidence to estimate the tail with precision. In effect, estimation errors fatten up the tail. If the parameters of the true distribution are not known with certainty, taking that second-level uncertainty into account leads decisionmakers to act as if they were facing a fat-tailed distribution. These fat tails "represent structural or deep uncertainty about the possibility of rare high-impact disasters that . . . 'scare' any [risk-averse] agent."[26] Thus, an inability to precisely estimate the parameters of a thin-tailed distribution—a form of second-order uncertainty about the first-order probability distribution—may confront the decisionmaker with a fat-tailed distribution in practical terms. Yet, we lack good analytic techniques for quantifying total risk when the distribution has a fat tail.

In sum, there are three connections between fat tails and uncertainty: first, fat tails *contribute* to uncertainty in the sense that they create an epistemic problem of estimation (when we are in a scenario with a fat-tailed distribution, we have difficulty measuring the tail); second, we may encounter *second-order* uncertainty simply because we do not know whether we have a fat tail or not; and third, uncertainty is *more dangerous* if we think we are in a fat-tail scenario because of potential feedback effects. Thus, fat-tailed distributions and uncertainty seem to be connected at a deep level.

B. Uncertainty Models and Worst-Case Scenarios

Unlike situations of pure uncertainty, however, we may have considerable information about the distribution of probabilities for fat-tailed distributions, but just not enough to pin down the fatness of the tail and establish the likelihood of catastrophic outcomes. Nonetheless, there are several approaches to analyzing such situations.

1. Models of Uncertainty and Ambiguity Aversion

"Ambiguity" is a term that is often used to refer to situations in which the true probability distribution of outcomes is not known.[27] There is strong empirical evidence that people are averse to ambiguity,[28] and such aversion "appears in a wide variety of contexts."[29]

There are a number of different approaches to modeling uncertainty about the true probability distribution.[30] I will focus on a particularly tractable approach called α-maxmin models. In these models, α represents the weighting factor between best and worst cases. As Sir Nicholas Stern explains, in these models of uncertainty, "the decisionmaker, who is trying to choose which action to take, does not know which of [several probability] distributions is more or less likely for any given action."[31] In this situation, the decisionmaker would act as if she chooses the action that maximizes a weighted average of the worst expected utility and the best expected utility . . . . The weight placed on the worst outcome would be influenced by concern of the individual about the magnitude of associated threats, or pessimism, and possibly any hunch about which probability might be more or less plausible.[32]

21. It can be difficult to distinguish power laws from other fat-tailed distributions empirically. *See* Aaron Clauset, Cosma Rohilla Shalizi & M.E.J. Newman, Power-Law Distributions in Empirical Data (Feb. 2, 2009) (unpublished manuscript), *available at* http://arxiv.org/pdf/0706.1062 (last visited June 16, 2012).
22. *See* Richard Sole & Brian Goodwin, Signs of Life: How Complexity Pervades Biology 52 (2000) (describing power laws).
23. Albert-La 'Szlo' Barbara 'Si, Linked: The New Science of Networks 67-68 (2002).
24. Martin L. Weitzman, *A Review of* The Stern Review on the Economics of Climate Change, 45 J. Econ. Lit. 703, 723 (2007)
25. Martin L. Weitzman, *On Modeling and Interpreting the Economics of Catastrophic Climate Change*, 91 Rev. Econ. & Stat. 1, 3 n.4 (2009). Indeed, even determining that data exhibits a fat-tailed distribution such as a power law rather than a thinner tailed distribution such as the lognormal distribution can be difficult. *See* M.E.J. Newman, *Power Laws, Pareto Distributions and Zipf's Law*, 46 Contemp. Physics 323, 329-30 (2005).
26. Weitzman, *Catastrophic Climate Change*, *id.* at 9. The distribution that he derives is not a power law but another fat-tailed distribution known for historical reasons as the "Student-t." *Id.* at 8.
27. For other legal applications of ambiguity models, see Daniel A. Farber, *Uncertainty*, 99 Geo. L.J. 901, 928 n.116 (2011) (listing sources).
28. *See* Gideon Keren & Le´onie E.M. Gerritsen, *On the Robustness and Possible Accounts of Ambiguity Aversion*, 103 Acta Psychologica 149, 149 (1999).
29. Nicholas Barberis & Richard Thaler, *A Survey of Behavioral Finance*, *in* Handbook of the Economics of Finance 1053, 1075 (George M. Constantinides, Milton Harris & Rene´ M. Stulz eds., 2003).
30. A good summary can be found in Alessandro Vercelli, *Hard Uncertainty and Environmental Policy*, *in* Sustainability: Dynamics and Uncertainty 191, 196-205 (Graciela Chichilnsky et al. eds., 1998).
31. Nicholas Stern, The Economics of Climate Change: The Stern Review (2007).
32. *Id.*

One way to understand these models is that we might want to minimize our regret for making the wrong decision, where we regret not only disastrous outcomes that lead to the worst-case scenario, but also we regret having missed the opportunity to achieve the best-case scenario. Alternatively, *a* can be a measure of the balance between our hopes (for the best case) and our fears (of the worst case).

Applying these α-maxmin models as a guide to action leads to what we might call the α-precautionary principle. Unlike most formulations of the precautionary principle, α-precaution is not only aimed at avoiding the worst-case scenario; it also involves precautions against losing the possible benefits of the best-case scenario.[33] In some situations the best-case scenario is more or less neutral, so that α-precaution is not much different from pure loss avoidance, unless the decisionmaker is optimistic and uses an especially low α. But where the best-case scenario is potentially extremely beneficial, unless the decisionmaker's α is very high, α-precaution will suggest a more neutral attitude toward uncertainty in order to take advantage of potential upside gains.

For example, suppose we have two models about what will happen if a certain decision is made. We assume that each one provides us enough information to allow the use of conventional risk assessment techniques *if* we were to assume that the model is correct. For instance, one model might have an expected harm of $1 billion and a variance of $0.2 billion; the other an expected harm of $10 billion and a variance of $3 billion. If we know the degree of risk aversion of the decisionmaker, we can translate outcomes into an expected utility figure for each model. The trouble is that we do not know which model is right, or even the probability of correctness. Hence, the situation is characterized by uncertainty. To assess the consequences associated with the decision, we then use a weighted average of these two figures based on our degree of pessimism and ambiguity aversion. This averaging between models allows us to compare the proposed course of action with other options.

α-maxmin has some important virtues in terms of process. Rather than asking the decisionmaker to assess highly technical probability distributions and modeling, it simply presents the decisionmaker with three questions to consider: (1) What is the best-case outcome that is plausible enough to be worth considering? (2) What is the worst-case scenario that is worth considering? (3) How optimistic or pessimistic should we be in balancing these possibilities? These questions are readily understandable by politicians and members of the public, presenting the key value judgments directly to the officials who should be making them, rather than concealing value judgments in technical analysis by experts.

2. Relating the Models

We seem to be suffering from an embarrassment of riches, in the sense of having too many different models for decisionmaking in situations in which extreme outcomes weigh heavily. At present, it is not clear that any one model will emerge as the most useful for all situations. For that reason, the ambiguity models should be seen as providing decisionmakers with a collection of tools for clarifying their analysis rather than providing a clearly defined path to the "right" decision.

Among this group of tools, what I have been calling α-precaution (utilizing α-maxmin) has a number of attractive features. First, it is complex enough to allow the decisionmaker to continue both the upside and downside possibilities, without requiring detailed probability information that is unlikely to be available. Second, it is transparent. Although the math behind this decision tool is formidable, actually applying the tool requires only simple arithmetic. The user must decide on what parameter value to use for α, but this choice is intuitively graspable as a measure of optimism versus pessimism.[34] Third, α-maxmin can be useful in coordinating government policy. An oversight agency such as OMB can provide benchmark values of α and rules for conducting sensitivity analysis. It can review departures from the benchmarks, where such departures are important, in order to determine that an agency's degree of pessimism or optimism about a problem is consistent with administration policy.

Models of uncertainty and fat-tailed models do not map precisely into each other although they both give us ways of thinking about catastrophic outcomes. Fat-tailed models are technically risk models rather than uncertainty models because the probability distribution is (somewhat) known. The mathematics in fat-tailed models thus looks different from that used in ambiguity models. A heuristic interpretation can link the difficulties of dealing with the dangers incorporated in fat-tailed distributions with the somewhat severe nature of the ambiguity-aversion models. Rather than trying to solve the intractable problem of the potential infinities in fat-tailed distributions, we can cut off the tail at some plausible "worst case"—but then make up for our inability to directly account for the full spectrum of outcomes by giving heavy weight to the chosen bad scenario. In other words, the extremism of maxmin or weighted decisions could be seen as a way of incorporating the fact that we have shunted aside the full range of horrific outcomes. Ambiguity between a finite set of models then functions as a stand-in for the fact that there are multiple alternative models, perhaps only poorly understood, that could lead to worse outcomes.

Alternatively, we might focus on the uncertainties presented by fat-tailed distributions themselves. In a situation

---

33. If α =1, then α -maxmin becomes ordinary maxmin, in which only the worst case matters.

34. We might be able to narrow the range for α by using empirical evidence showing how individuals approach decisionmaking in situations characterized by ambiguity or through experience over time that might allow officials to develop norms about the appropriate α.

in which a fat-tailed distribution is a possibility, the decisionmaker may face several unknowns: whether the distribution actually does have a fat tail, the type and parameters of the fat-tailed distribution, or whether (and where) to truncate the distribution if there is some possible upper bound on outcomes. Thus, even if a specific fat-tailed distribution (with or without truncation) actually does characterize the situation, the barriers to full knowledge of the distribution may mean that the decisionmaker's problem is more one of uncertainty than risk, making ambiguity models relevant.

III. Applying New Decision Techniques to Regulatory Policy

Of course, the crucial question is whether these various techniques can provide genuine assistance in dealing with important policy issues. This part deploys the approaches economic theory provides in the context of two important current regulatory problems, each of which is characterized by considerable uncertainty: how much society should be willing to pay to mitigate climate change by reducing emissions of greenhouse gases, and whether society should restrict the development of nanotechnology.

A. Climate Change Mitigation

1. Scientific and Economic Confidence and Uncertainty

The primary uncertainty in climate mitigation is the "wide range of possible temperature increases . . . including a five-percent possibility that temperature increases will equal or exceed 6 C° and a two-percent probability of increases equal to or greater than 8 C° within the next 100 to 200 years."[35] Such increases may not sound like much, but a 5° rise is "equivalent to the change in average temperatures from the last ice age to today."[36]

The customary measure for how strongly the climate system responds to changes in the level of greenhouse gases is climate sensitivity. Climate sensitivity is measured as the equilibrium temperature increase caused by a permanent doubling of preindustrial CO_2 concentrations. Studies based on historical climate data find that climate sensitivity is unlikely to be below 1.5° C; the upper bound is more difficult to determine for technical reasons—it could exceed 4.5° C, although such high values are much less likely on the basis of the historical record than those in the 2.0° C to 3.5° C range.[37] A second line of research examines climate sensitivity in models. In each model, the climate sensitivity depends on many processes and feedbacks, and probability distributions can be determined by examining how climate sensitivity tracks variations in various other parameters in the model. Essentially, parameters are subject to variations, and the effect on climate response is measured through many runs of the model. The most frequent sensitivity values are around 3° C, but much higher values cannot be excluded.[38]

Unfortunately, there is no completely satisfactory way of translating these results into a formal probability distribution.[39] If we assume that all current models are equally likely and that they exhaust the possibilities, we can get a probability distribution, but these are somewhat heroic assumptions.[40]

Even when models do agree, there are residual grounds for uncertainty. Models "might share a common error" for example.[41] While there is fairly good evidence that there are no major missing factors, at least in terms of explaining overall 20th-century warming trends,[42] we do know that other factors are relevant and imperfectly modeled for future trends and regional impacts.[43] Some efforts have been made to quantify uncertainty based on various other lines of evidence[44]; new types of computational experiments have been performed to quantify uncertainty about how models respond to external inputs such as changes in solar intensity, for example. Additionally, modelers and other scientists are prone to biases and errors, like the rest of us, despite the strenuous efforts that the scientific enterprise makes to limit the effects of these weaknesses,[45] and this source of error is hard to estimate.

Notwithstanding such concerns, models give us a fair amount of confidence about basic trends. We can be highly confident about the existence of human-caused climate change and the likelihood that it will have serious effects. There is strong residual uncertainty, however, about the scale of climate change impacts, both globally and regionally. This uncertainty might seem to argue against investing in climate change mitigation, but as demonstrated below, the possibility of high-impact scenarios actually provides a further reason to take precautionary steps.

2. Climate Policy, Catastrophic Risks, and α-Maxmin

The more disturbing issues are on the scientific side though, and relate to the possibility that climate change

35. Daniel H. Cole, *The* Stern Review *and Its Critics: Implications for the Theory and Practice of Benefit—Cost Analysis*, 48 Nat. Resources J. 53, 75 (Winter 2008).
36. Stern, *supra* note 31, at xvi.
37. *See* Gerald A. Meehl et al., *Global Climate Projections*, *in* Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Climate Change 2007: The Physical Science Basis 747, 800-01 (Susan Solomon et al. eds., 2007), *available at* http://www.ipcc.ch/pdf/assessment-report/ar4/wg1/ar4-wg1-chapter10.pdf.
38. *Id.* at 799.
39. *Id.*
40. *See id.*
41. Myles Allen et al., *Scientific Challenges in the Attribution of Harm to Human Influence on Climate*, 155 U. Pa. L. Rev. 1353, 1361 (2007).
42. *See id.* at 1375.
43. *See* Meehl et al., *Global Climate Projections*, *supra* note 37, at 797 ("Uncertainty in predictions of anthropogenic climate change arises at all stages of the modeling process . . . .").
44. *Id.* at 754.
45. *See* Myanna Lahsen, *Seductive Simulations? Uncertainty Distribution Around Climate Models*, 35 Soc. Stud. Sci. 895, 904-08 (2005).

will *not* be moderate. Based on an analysis of reported studies, Weitzman estimates that a "'best guess' estimate of the extreme bad tail" places the odds at about 5% of a temperature increase over 10° C (18° F) and a 1% chance of an increase of 20° C (36° F).[46] It is hard to improve on his explanation of the gravity of these findings:

> [s]ocieties and ecosystems in a world whose average temperature has changed in the geologically instantaneous time of two centuries or so by 10° C-20° C . . . are located in *terra incognita*, since such high temperatures have not existed for hundreds of millions of years and such a rate of global temperature change might be unprecedented even on a timescale of billions of years.[47]

Hence, "the planetary welfare effect of climate changes [from such increases] . . . implies a nonnegligible probability of worldwide catastrophe."[48]

As Weitzman says, the normative implication is clearly a higher degree of precaution, making "insurance" against catastrophe a critical factor in climate policy.[49] It is difficult to extract more specific guidance from his approach,[50] and we might instead turn to ambiguity-based models for guidance.

Ambiguity theory suggests that we weigh the best-case scenario (unimpeded economic growth combined with modest investment in climate adaptation) and the worst-case scenario (catastrophic climate outcomes), perhaps also including as a mid-case the standard economic models of climate change (which, as it happens, are not too far away from the best case).[51]

The implication of this analysis would be a high degree of precautionary catastrophe insurance, as Weitzman suggests. This argument can be seen as an application of Sunstein's "catastrophic harm precautionary principle."[52] If we think in terms of α-maxmin models, the worst-case scenario is grim, perhaps on the order of the end of civilization; the best-case scenario is that harm from climate change is modest. Unless we are inclined to be optimistic and place extraordinarily weight on the best-case scenario, business as usual does not seem to be an appealing strategy—in fact, we should be willing to make major investments to reduce climate change. This conclusion is α robust under a variety of assumptions, as shown below.

Specifically, if H_w is the harm in the worst-case scenario and H_B is the harm in the best-case scenario, we would attribute a cost of $\alpha H_w + (1—\alpha)H_B$l to the strategy of doing nothing. Even if H_B is zero (no net harm from climate change), the no-action option will not be appealing. The reason is that, because H_w is so large, αH_w will be a large number unless α is very small indeed. For example, suppose we are equally balanced between optimism and pessimism (α = 0.5) and that we take the worst case as being a temperature change equivalent to at least a trillion dollars in value. Then we would be willing to spend $500 billion or more to avoid this risk.

If we take into account more catastrophic outcomes, the case for doing nothing evaporates even if we are optimistic about avoiding the worst-case scenario. As we have seen, Weitzman suggests that the most extreme outcomes could result in the end of civilization. If we interpret that as a complete collapse of world GDP, we would get an estimated loss of $\$10^{16}$, or $1 quadrillion (or in more familiar terms, $1000 trillion).[53] In order to reflect optimism about climate change, assume that the best-case scenario is actually a $1 trillion benefit from warming, and take α = 0.01 (meaning that we put 99 times as much emphasis on the best case as on the worst case). With some simple arithmetic, we come up with a loss figure of .01(1000 trillion)-0.99(1 trillion), or approximately $9 trillion. Therefore, even if we are highly optimistic about the best-case scenario, a serious investment in climate mitigation would still be warranted if the downside risk is as severe as Weitzman suggests.

Thus, the α-precautionary principle would warrant a high degree of precaution to avoid the negative uncertainties of climate change. Based on reasoning of this type, the *Stern Review* suggests that the cost of climate change should be assessed at between 13% and 20% of current global consumption, with the weight used to average the figures being based on "crude judgments about likelihoods of different kinds of probability distributions, on judgments about the severity of losses in this context, and on the basic degree of cautiousness on the part of the policymaker."[54] The World Bank estimates world GDP in 2008 at about $60.5 trillion,[55] so the value of eliminating climate change would be roughly $6-$12 trillion. Because Stern's is only

46. Weitzman, *supra* note 25, at 1.
47. *Id*. A leading critic of Weitzman concurs that "[m]any people would agree that a 5% chance of a 10° change, or a 1% chance of a 20° change, would be a catastrophic prospect for human societies." William D. Nordhaus, *An Analysis of the Dismal Theorem* 10 (Cowles Found., Discussion Paper No. 1686, 2009), *available at* http//ssrn.com/abstracts= 1330454.
48. Weitzman, *supra* note 25, at 1.
49. *Id*. at 18. The fat-tail aspect of Weitzman's analysis seems to be crucial. Using a thin-tail analysis while still taking into account possible extreme outcomes, Pindyck finds a case for moderate climate mitigation but nothing more. *See* Robert S. Pindyck, *Uncertain Outcomes and Climate Change Policy* 22 (MIT Sloan Sch., Working Paper No. 4742-09, 2009), *available at* http://ssrn.com/abstract = 1448683. Pindyck provides an important caveat:
    > We have no historical or experimental data from which to assess the likelihood of a ΔT [change in temperature] above 5° C, never mind its economic impact, and one could argue á la Weitzman (2009) that we will never have sufficient data because the distributions are fat-tailed, implying a WTP [willingness to pay] of 100% [of consumption] (or at least something much larger than 2%). *Id*.
50. It is hard to quarrel, however, with Weitzman's statement that "[e]ven just acknowledging more openly the incredible magnitude of the deep structural uncertainties that are involved in climate-change analysis . . . might go a long way toward elevating the level of public discourse concerning what to do about global warming." Weitzman, *supra* note 25, at 18.
51. *See supra* notes 116-33 and accompanying text.
52. *See* Cass R. Sunstein, *The Catastrophic Harm Precautionary Principle*, Isues Legal Scholarship, 2007, *available at* http://www.bepress.com/ils/iss10/art3.
53. Nordhaus, *supra* note 47, at 14 (stating that "the discounted value of world consumption is in the order of $\$10^{16}$).
54. Stern, *supra* note 31, at 187. As Cole, *supra* note 35, at 62, explains, these numbers are controversial, but they are at least illustrative.
55. *See Key Development Data and Statistics*, The World Bank, http://econ.worldbank.org/WBSITE/EXTERNAL/DATASTATISTICS/0,,contentMDK:20535285% A0menuPK:1192694%A0pagePK:64133150%CA0piPK:64133175%A0theSitePK:239419,00.html (last visited Dec. 20, 2010).

one model, the actual range of estimates is wider, making the choice of the weighting factor (α) even more important. It seems clear, however, that it would be worth investing a large amount of money in climate mitigation.

It is tempting to seek a higher degree of precision in this recommendation, but in practical terms, the precision is probably irrelevant. If we take seriously that there is even a small possibility that climate change could wipe out our present society,[56] the indicated amount of precaution is probably higher than anything we could plausibly expect from the political system. So, the identity of the "correct" policy is this: the most stringent policy that is politically feasible[57] (though unfortunately that policy still probably runs a haunting risk catastrophe).

The basic lesson here is quite simple and does not depend on the details of the analysis. Climate policy cannot be based simply on the outcomes we consider most likely. The full range of possible consequences must be considered. Given the possibility of dire consequences from climate change, corrective measures should be supported even if some people believe that climate change most likely will not occur or that it will be beneficial.

B. *Nanotechnology*

Nanotechnology presents different sorts of unknowns and therefore a different context for investigating regulatory uncertainty. As a technology in its early stages of development, it presents the possibility of extraordinary benefits as well as serious risks. We have little ability to attach probabilities to any of the outcomes, making this a case of true uncertainty.[58]

Nanotechnology is the domain of the remarkably small. One nanometer (nm) is equal to one-billionth of a meter (or about 0.00000004 inches), an incredibly tiny length. Importantly, nanoparticles can have properties quite different from larger amounts of the same substance—for example, opaque particles can become transparent to visible light but reflective of ultraviolet light at nano size.[59]

Anticipated applications of nanotechnology in the relatively near term include cosmetics, materials for remediating hazardous waste sites, fuel cells, video displays, batteries, and fuel additives[60]; longer-term projects may involve revolutionary developments rather than incremental evolution, including new tests and treatments for cancer, greatly improved renewable energy, universal access to clean water, and higher crop yields through use of nanosensors to detect plant diseases.[61]

But the same properties that make nanotech appealing, such as high surface reactivity and ability to cross cell membranes, may also pose risks—risks that are still poorly understood.[62] A study by the Royal Society indicated that "there is a lack of information about [nanoparticles'] health, safety and environmental impacts," requiring reliance on research results regarding other small particles from pollution and occupational research.[63] Given the uncertainties, the Royal Society recommended a ban on use of free nanoparticles for cleaning up toxic sites,[64] and it put a high priority on investigation by regulators of the safety of nanoparticles in consumer products.[65]

The Congressional Research Service (CRS) also recently surveyed the risks and potential benefits of nanotechnology,[66] viewing the long-run picture as potentially involving revolutionary developments but also recognizing risks as scientists already know that some nanomaterials (carbon nanotubes and fullerenes) can cause lung damage in mice, brain damage in fish, and DNA damage.[67]

Environmental advocates call for a moratorium on commercial release of food and agricultural materials containing manufactured nanomaterials until a new legal structure is in place.[68] Public interest groups "have invoked the Precautionary Principle in advocating a more draconian regulatory approach to address potential risks from nanomaterials."[69] Others argue that the precautionary principle "freezes us in place," because "[n]o technology at its inception can satisfy the precautionary principle, so the principle becomes a formula for doing nothing."[70] Thus, further study and investment in liability insurance are arguably better approaches.[71] Another possibility would be to impose a substantial bond requirement for

56. A caveat is that we could downplay the potential catastrophic possibilities if, as Nordhaus argues, we could learn that catastrophe is impending fast enough to make a sufficiently quick and vigorous global response to head off the possibility. *See* Nordhaus, *supra* note 47, at 20. In my view, Nordhaus is excessively optimistic about this last-minute policy response, in part because of the potential for "climate surprises" involving abrupt climate change that might not leave a great deal of time for a response. *See* John D. Cox, Climate Crash: Abrupt Climate Change and What It Means for Our Future 189 (2005). Nevertheless, the potential for detecting and heading off catastrophic climate change does need to be considered as part of the analysis.
57. *See, e.g.*, Robert W. Hahn, *Climate Policy: Separating Fact From Fantasy*, 33 Harv. Envtl. L. Rev. 557, 577 (2009).
58. For a recent discussion that emphasizes the importance of these uncertainties, see Douglas A. Kysar, *Ecologic: Nanotechnology, Environmental Assurance Bonding, and Symmetric Humility*, 28 UCLA J. Envtl. L. & Pol'y 201 (2011).
59. The Royal Soc'y & the Royal Acad. of Eng'g, Nanoscience and Nanotechnologies: Opportunities and Uncertainties 9 (2004).
60. *Id.* at 10-12.
61. John F. Sargent Jr., Cong. Research Serv., RL 34511, Nanontechnology: A Policy Primer 1, 3-4 (2009).
62. Royal Society, *supra* note 59, at 35.
63. *Id.* at 47. As of 2004, according to the Royal Society, "very few studies have been published on the potential adverse effects that nanoparticles or nanotubes may have on humans, and only one to our knowledge on environmental effects." *Id.* at 75.
64. *Id.* at 47.
65. *Id.* at 74.
66. Sargent, *supra* note 61.
67. *Id.* at 9.
68. *See* Georgia Miller & Rye Senjen, Friends of the Earth, Out of the Laboratory and Onto Our Plates 3 (2008).
69. David B. Fischer, *Nanotechnology—Scientific and Regulatory Challenges*, 19 Vill. Envtl. L.J. 315, 330 (2008). Dana, *Contextual Rationality*, *supra* note 14, at 18-29, argues that the precautionary principle may correct market incentives to avoid investigating possible environmental and health risks.
70. Robin Fretwell Wilson, *Nanotechnology: The Challenge of Regulating Known Unknowns*, 34 J.L. Med. & Ethics 704, 710 (2006).
71. *Id.* at 711.

substances that are allowed on the market after passing screening tests.[72]

Because nanotechnology has potential large upsides as well as downsides, an attitude of pure precaution seems inappropriate. Instead, we would do better to use ambiguity models that balance upside and downside outcomes, such as α-maxmin.[73] The α-precautionary principle would probably not justify efforts to forestall research and development of nanotechnology given its high upside potential. It would, however, justify a degree of caution.

An appropriate strategy could involve sustained research into health and safety issues of current uses of nanomaterials,[74] restrictions on uses involving potential public exposure until further risk information is available, and sensitivity to potential large downside risks in R & D for longer term, nonevolutionary nanotechnologies. Given the unknown hazards associated with nanomaterials, it is surprising that regulatory authorities have failed to treat them as new substances for regulatory purposes but have instead given them the more favorable treatment available to existing products.[75] That said, on balance nanomaterials do not require a more precautionary approach than new chemicals in general.

Conclusion

It is sometimes tempting to ignore the imperfectly understood dimensions of hazards as speculative. That is clearly the wrong response. Just because you do not know exactly how big a number is, there is no reason to assume it to be zero.

As we have seen, such uncertainties can be associated with fat-tailed distributions, while in other situations, we may simply have no good idea of how to assign probabilities in the first place or of what the probability distribution might look like. Ambiguity theory helps address these situations, and the most easily applied models advise assessing decisions based on a combination of the best-case and worst-case scenarios. This leads to the α-precautionary principle, which weighs the best and worst potential outcomes in assessing a course of action. Although there is no easy recipe for divining the right solution to problems the parameters of which involve so much uncertainty, but we can gain some much-needed clarity with the tools discussed in this Article.

72. This proposal is made in Albert C. Lin, *Size Matters: Regulating Nanotechnology*, 31 Harv. Envtl. L. Rev. 349, 396-404 (2007). Kysar, *supra* note 58, at 208-09, presents an alternative bonding proposal that emphasizes the role played by worst-case outcomes in establishing bond amounts.

73. *See supra* notes 41-50 and accompanying text.

74. EPA has embraced such a research program, but if past practice is a guide, it could take a decade or more before the work even begins.

75. *See* Diana M. Bowman & Graeme A. Hodge, *A Small Matter of Regulation: An International Review of Nanotechnology Regulation*, 8 Colum. Sci. & Tech. L. Rev. 1, 36 (2007). An EPA advisory is now considering whether to recommend that nanosilver products be treated as new pesticides requiring a new pesticide registration. *See* Lynn L. Bergeson, *FIFRA Scientific Advisory Panel Considers Nanosilver*, 39 ELR 11143, 11143-44 (Dec. 2009).