

Catastrophic Uncertainty and Regulatory Impact Analysis

Cost-benefit analysis embodies techniques for the analysis of possible harmful outcomes when the probability of those outcomes can be quantified with reasonable confidence. But when those probabilities cannot be quantified (“deep uncertainty”), the analytic path is more difficult. The problem is especially acute when potentially catastrophic outcomes are involved, because ignoring or marginalizing them could seriously skewing the analysis. Yet the likelihood of catastrophe is often difficult or impossible to quantify because such events may be unprecedented (runaway AI or tipping points for climate change) or extremely rare (global pandemics caused by novel viruses in the modern world). OMB’s current guidance to agencies on unquantifiable risks is now almost twenty years old and in serious need of updating. It correctly points to scenario analysis as an important tool but it fails to give guidance on the development of scenarios. It then calls for a qualitative analysis of the implications of the scenarios, but fails to alert agencies to the potential for using more rigorous analytic techniques.

Decision science does not yet provide consensus solutions to the analysis of uncertain catastrophic outcomes. But it has advanced beyond the vague guidance provided by OMB since 2003, which may not have been state-of-the-art even then. This paper surveys these developments and explains how they might best be incorporated into agency practice. Those developments include a deeper understanding of potential options and issues in constructing scenarios. They also include analytic techniques for dealing with unquantifiable risks that can supplement or replace the qualitative analysis currently suggested by OMB guidance. To provide a standard framework for discussion of uncertainty in regulatory impact analyses, the paper also proposes the use of a structure first developed for environmental impact statements as a way of framing the agency’s discussion of crucial uncertain outcomes.

Proposed Modification to OMB Circular A-4

The following language replaces the discussion of nonquantifiable uncertainty on p. 39 of the current guidance:

When the likelihood of an outcome cannot be quantified with confidence, you should provide (1) a statement that the likelihood of the outcome cannot be quantified with reasonable confidence; (2) a statement of the relevance of the outcome to assessing the proposed regulation; (3) a summary of existing credible scientific evidence which is relevant to evaluating reasonably foreseeable impacts relating on the costs or benefits of a proposed regulation, and (4) the agency's evaluation of such impacts based upon theoretical approaches or research methods generally accepted in the scientific community. For these purposes, "reasonably foreseeable" includes outcomes involving catastrophic consequences, even if their likelihood of occurrence is considered low, provided that the analysis is supported by credible scientific evidence, is not based on pure conjecture, and is within the rule of reason.¹

When an analysis of such an outcome is considered necessary, the agency should consult OMB about potential analytic approaches. The goal of the analysis should be to provide as much guidance as possible to decision makers about the significance of the potential outcomes in question. At a minimum, the analysis should provide a description of alternative scenarios. For instance, in assessing the potential outcomes of an environmental effect, there may be a limited number of scientific studies with strongly divergent results. In such cases, you might present results from a range of plausible scenarios, together with any available information that might help in qualitatively determining which scenario is most likely to occur.²

In developing your analysis, you should consider different methods of scenario development and the potential for applying techniques from decision science in analyzing the import of possible scenarios. You should reference best practices in scenario design and explain the reasons for your choices. You should also discuss how you will assess the outcomes of set of scenarios. Along with qualitative assessments, you may also want to consider techniques including the use of methods such as α -maxmin, robust decision making, and safe minimum standards (among others). Whether you choose to use a more formalized decision method or one that relies on more qualitative judgments and best professional judgment, you should bear in mind whatever guidance the governing statute provides about decision making under uncertainty, such as language requiring an adequate margin of safety.

¹ This language is adapted from the Council on Environmental Quality's guidelines for environmental impact analysis, which is used as a template because it is familiar to many agencies and well understood by courts.

² The preceding sentences in this paragraph are adapted from Circular A-4.

Incorporating Catastrophic Uncertainty in Regulatory Impact Analysis

There are well-developed techniques for incorporating possible harmful outcomes into regulatory impact analysis when the probability of those outcomes can be quantified with reasonable confidence. But when those probabilities cannot be quantified, the analytic path is more difficult.³ The problem is especially acute when potentially catastrophic outcomes are involved, because ignoring or marginally them risks seriously skewing the analysis.⁴ Yet the likelihood of a catastrophic is often difficult to quantify because such events may be unprecedented (runaway AI or climate change) or extremely rare (global pandemics in the modern world).

OMB's current guidance to agencies on this topic is now almost twenty years old and in serious need of updating. It correctly points to scenario analysis as an important tool, but it fails to give guidance on the development of scenarios. It then calls for the a qualitative analysis of the implications of the scenarios, but fails to alert agencies to the potential for using more rigorous analytic techniques.

Decision science does not yet provide consensus solutions to the analysis of uncertain catastrophic outcomes. It has, however, advanced beyond the vague guidance currently provided by OMB, or choosing the alternative with the least-bad possible outcome (maxmin). This paper surveys some of the key developments and describes how they might best be incorporated into agency practice. It is important in doing so to take into account the analytic demands that are already placed on often-overburdened agencies and the variability between agencies in economic expertise and sophistication.

Even at the time OMB's guidance was issued, there was model language that OMB could have used that would have at least provided a general framework for agencies to use. That language is found in the Council on Environmental Quality's guidelines for the treatment of uncertainty in environmental impact statements.⁵ Using that language as a template has several advantages. Many agencies are familiar with the language, as are courts. Judicial interpretation provides additional guidance on the use of that language.⁶ Since agency decisions requiring cost-benefit analysis often

³ Following much of the literature, I will refer to these as situations involving uncertainty, as opposed to risks with quantifiable likelihoods or known probability distributions. For a defense of the view that this is a real and important category, see Cass R. Sunstein, *Maxmin*, 37 Yale J. Reg. 940, 965-966, 976-978 (2020). For a survey of important regulatory situations in which uncertainty arises, see David Weissbach, *Introduction: Legal Decision Making Under Deep Uncertainty*, 44 J. Legal Stud. S319 (2015).

⁴ See, e.g., an early criticism of estimates of the social cost of carbon for failing to adequately deal with catastrophic risks. See Michael A. Livermore and Richard L. Revesz, *Retaking Rationality Two Years Later*, 48 Houston L. Rev. 1, 24 (2013). Even today, damage estimates for temperature increases at or above 4 °C are considered "speculative and highly uncertain." Peter H. Howard and Thomas Sterner, *Between Two Worlds: Methodological and Subjective Differences in Climate Impact Meta-Analyses* (Resources for the Future, June 2022), <https://www.rff.org/publications/working-papers/between-two-worlds-methodological-and-subjective-differences-in-climate-impact-meta-analyses/>. Hopefully, such large temperature increases are now of less policy relevance given efforts to date to limit emissions.

⁵ 40 U.S.C. §1502.22.

⁶ For discussion of judicial interpretation of this provision, see Todd S. Aagaard, *A Functional Approach to Risks and Uncertainties Under NEPA*, 1 Mich. J. Env. & Admin. L. 87 (2012). A useful case study is provided by Irene

also require an environmental assessment or impact statement, creating some parallelism between the NEPA requirements and OMB guidance will also allow the relevant documents to mesh more easily.

The paper proposes modifying this template to make specific reference to some leading techniques provided by decision science. There is a plausible argument that decision makers should simply apply their own intuitions to weighing unquantifiable risks of potential catastrophe. This paper takes another tack for two reasons. First, it is difficult to form judgments about scenarios that are likely to be far outside of the decision maker's past experience. Whatever analytic help can be offered in making those judgment calls would be beneficial. Second, to have a chance of adoption, proposals should be geared to the technical, economics-oriented perspective espoused by OIRA. Even for those who do not find formalized decision methods adequate, it is better to move the regulatory process in the direction of fuller consideration of potential catastrophic risks rather than having such possibilities swept to the margins of analysis because of a desire for "rigor."

I. Circular A-4, Uncertainty, and Later Generations

Our starting point is the current OMB guidance on regulatory impact analysis. The section of Circular A-4 begins by discussing uncertainty as a general topic regardless of quantifiability. It then has a short discussion that is specifically focused on unquantifiable uncertainty:

In some cases, the level of scientific uncertainty may be so large that you can only present discrete alternative scenarios without assessing the relative likelihood of each scenario quantitatively. For instance, in assessing the potential outcomes of an environmental effect, there may be a limited number of scientific studies with strongly divergent results. In such cases, you might present results from a range of plausible scenarios, together with any available information that might help in qualitatively determining which scenario is most likely to occur. When uncertainty has significant effects on the final conclusion about net benefits, your agency should consider additional research prior to rulemaking.⁷

The passage then goes on to discuss the possible benefits of obtaining additional information, "especially for cases with irreversible or large upfront investments."⁸ This passage is notably lacking in guidance about how to develop scenarios or what to do with the results.

Circular A-4 also contains a discussion of future generations. It advises:

Special ethical considerations arise when comparing benefits and costs across generations. Although most people demonstrate time preference in their own consumption behavior, it may not be appropriate for society to demonstrate a similar preference when deciding between the well-being of current and future generations. Future citizens who are affected by such choices cannot take part in making them, and today's society must act with some consideration of their interest. One way to do this would be to follow the same discounting techniques described above and supplement the analysis with an explicit

Weintraub, *NEPA and Uncertainty In Low-Risk, High-Impact Scenarios: Nuclear Energy as a Case Study*, 37 *Cardozo L. REV.* 1565 (2018) [student note].

⁷ OMB Circular A-4, "Regulatory Analysis" (2003), at 38-39.

⁸ *Id.* at 39.

discussion of the intergenerational concerns (how future generations will be affected by the regulatory decision).⁹

The remainder of this discussion of Circular A-4 focuses on the choice of discount rates, primarily arguing for discounting on the ground that future generations will be wealthier than the current generation. The literature on discounting over multiple generations has advanced considerably since then, with a consensus in favor of declining discount rates as a hedge against unexpectedly low economic-welfare in future generations.¹⁰

The use of discounting assumes, however, that we can quantify risks and their costs in future time periods. In the nature of things, this becomes increasingly difficult over time, as models are extended further beyond their testable results and as society diverges further from historical experience.

In the nearly 20 years since Circular A-4, the state of the art regarding analysis of unquantified uncertainty has made real progress. OMB's guidance should be improved accordingly.

II. Case Study: Climate Change

The potentially catastrophic multigenerational risks that have been most heavily researched involve global climate change. Climate change is the poster child for the twin problems of uncertainty and impacts on future generations. Impacts on future generations are inherent in the nature of the earth system's response to radiation forcing by greenhouse gases. Uncertainty is due to the immense complexity of the climate system and the consequent obstacles to modeling, combined with the long time spans involved and the unpredictability of human responses. Thus, climate change presses current methods of risk assessment and management to their limits and beyond.

The IPCC uses a standardized vocabulary to characterize quantifiable risks.¹¹

IPCC Probability Designations

Term	<i>Virtually Certain</i>	<i>Extremely Likely</i>	<i>Very Likely</i>	<i>Likely</i>	<i>More Likely than Not</i>	<i>As Likely as Not</i>	<i>Unlikely</i>	<i>Very Unlikely</i>
Likelihood	>99%	>95%	>90%	66-90%	>50%	33-66%	<33%	<10%

Neither lawyers nor economists have evolved anything similar. The efforts of scientists to provide systemized designations of uncertainty is an indication of the attention they have given to both quantifiable and unquantifiable risks.

⁹ *Id.* at 35-36.

¹⁰ For an accessible explanation of the economics, see Daniel A. Farber, *Gambling Over Growth: Economic Uncertainty, Discounting, and Regulatory Policy*, 44 *J. Leg. Studies* S509 (2015).

¹¹ IPCC, *Climate Change 2022 — Impacts, Adaptation, and Vulnerability — Summary for Policymakers* 3 n.4 (2022).

As the IPCC explains:

In summary, while high-warming storylines –those associated with global warming levels above the upper bound of the assessed very likely range – are by definition extremely unlikely, they cannot be ruled out. For SSP1-2.6, such a high-warming storyline implies warming well above rather than well below 2°C (high confidence). Irrespective of scenario, high-warming storylines imply changes in many aspects of the climate system that exceed the patterns associated with the best estimate of GSAT [global mean near-surface air temperature] changes by up to more than 50% (high confidence).¹²

High climate sensitivities also increase the potential for warming above the 3 °C level. This, in turn, is linked to possible climate responses whose likelihood is very poorly understood. The IPCC reports a number of potential tipping point possibilities with largely irreversible consequences. Of these, two tipping points are associated with “deep uncertainty” for warming levels above 3°C: collapse of Western Antarctic ice sheets and shelves, and global sea level rise.¹³ Uncertainty is also reported as high regarding the potential for abrupt changes in Antarctic sea ice and in the Southern Ocean Meridional Overturning Circulation.¹⁴

The IPCC also indicates that damage estimates at given levels of warming are also subject to considerable uncertainty:

Projected estimates of global aggregate net economic damages generally increase non-linearly with global warming levels (high confidence). The wide range of global estimates, and the lack of comparability between methodologies, does not allow for identification of a robust range of estimates (high confidence). The existence of higher estimates than assessed in AR5 indicates that global aggregate economic impacts could be higher than previous estimates (low confidence). Significant regional variation in aggregate economic damages from climate change is projected (high confidence) with estimated economic damages per capita for developing countries often higher as a fraction of income (high confidence). Economic damages, including both those represented and those not represented in economic markets, are projected to be lower at 1.5°C than at 3°C or higher global warming levels (high confidence).¹⁵

Economists might view this assessment of the state of the art in modeling climate damages as too harsh. Nevertheless, it seems clear that we cannot assume that present estimates accurately represent the true probability distribution of possible damages.

The possibility of catastrophic tipping points looms large in economic modeling of climate change. Another key issue in terms of the economic analysis is the possibility of unexpectedly bad outcomes, such major melting of ice sheets, releases of large amounts of methane, and halting of

¹² June-Yi Lee et al., *Future Global Climate: Scenario-Based Projections and Near Term Information*, in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* 555 (2021).

¹³ *Id.* at 634.

¹⁴ *Id.* at 635.

¹⁵ *Climate Change 2022, supra* note 13, at 15.

the Gulf Stream.¹⁶ William Nordhaus, who pioneered the economic models of climate change, has explained how these affect the analysis:

[W]e might think of the large-scale risks as a kind of planetary roulette. Every year that we inject more CO₂ into the atmosphere, we spin the planetary roulette wheel. . . .

A sensible strategy would suggest an insurance premium to avoid the roulette wheel in the Climate Casino. . . . We need to incorporate a risk premium not only to cover the known uncertainties such as those involving climate sensitivity and health risks but also ... uncertainties such as tipping points, including ones that are not yet discovered.¹⁷

The difficulty, as Nordhaus admits, is trying to figure out the extent of the premium. Another recent book by two leading climate economists argues that the downside risks are so great that “the appropriate price on carbon is one that will make us comfortable enough to know that we will *never* get to anything close to 6 °C (11 °F) and certain eventual catastrophe.”¹⁸ Although they admit that “never” is a bit of an overstatement – reducing risks to zero is impractical – they clearly think it should be kept as low as feasibly possible. Not all economists would agree with that view, but there seems to be a growing consensus that the possibility of catastrophic outcomes should play a major role in determining the price on carbon.¹⁹

Scientists are beginning to gain greater knowledge about long-term climate impacts after the end of this century. The IPCC reports advances in modeling up through 2300. Under all but the lowest emissions scenarios, global temperatures continue to rise markedly through the 22nd century.²⁰ Moreover, “[s]ea level rise may exceed 2 m on millennial time scales even when warming is limited to 1.5°C–2°C, and tens of metres for higher warming levels.”²¹ Indeed, “physical and biogeochemical impacts of 21st century emissions have a potential committed legacy of at least 10,000 years.”²²

To provide some perspective of the scale of the potential changes:

To place the temperature projections for the end of the 23rd century into the context of paleo temperatures, GSAT [global surface air temperature] under SSP2-4.5 (likely

¹⁶ William Nordhaus, *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World* 51 (2013). Nordhaus observes that “all the systems involved in the analysis of tipping points are perplexing because they involve poorly understood dynamics and nonlinear responses.” *Id.* at 63.

¹⁷ *Id.*

¹⁸ Gernot Wagner and Martin L. Weitzman, *Climate Shock: The Economic Consequences of a Hotter Planet* 78 (2016).

¹⁹ Reaching a similar conclusion after extensive modeling exercises, another group of researchers concludes:

Given the limitations in our ability to model the future climate and climate change damages, in addition to the potential for non-linear and low-probability catastrophic climate responses with large damages, it is inappropriate to rely on “best guess” climate scenarios to dictate optimal mitigation investment pathways in a policy context. . . . [W]orst case scenarios must be considered when crafting policy.

Megan Ceronsky, David Anthoff, Cameron Hepburn, and Richard S.J. Tol, *Checking the Price Tag on Catastrophe: The Social Cost of Carbon Under Non-Linear Response* 17 (2011) (ESRI Working Paper 392). *See also* Robert S. Pindyck, *Climate Change Models: What Do the Models Tell Us?* 14-15 (2010) (NBER working paper 19244).

²⁰ Lee, *supra* note 14, at 631.

²¹ *Id.* at 631.

²² *Id.* at 632.

2.3°C–4.6°C higher than over the period 1850–1900) has not been experienced since the Mid Pliocene, about three million years ago. GSAT projected for the end of the 23rd century under SSP5-8.5 (likely 6.6°C–14.1°C higher than over the period 1850–1900) overlaps with the range estimated for the Miocene Climatic Optimum (5°C–10°C higher) and Early Eocene Climatic Optimum (10°C–18°C higher), about 15 and 50 million years ago, respectively (medium confidence)²³

The current emission pathway seems unlikely to produce the second, more severe scenario (though “unlikely” does not mean impossible). But the first scenario (SSP2-4.5) represents a middle-of-the-road future in which global emissions peak around 2050 and decline through 2100.²⁴ Even that scenario could put the world well out of the range of what homo sapiens has ever experienced, making predictions about future damages inherently uncertain.

III. Decisions Methods under Deep Uncertainty

It is one thing to perceive the significance of potential catastrophic risks; it is another to incorporate them into the analysis in a useful way. The precautionary principle is one effort to do so in a qualitative way, though the principle is controversial, particularly among economists. Use of precaution in the case of catastrophic risks has had some support even from Cass Sunstein, a leading critic of the precautionary principle.²⁵ Sunstein has proposed a number of different versions of the catastrophic risk precautionary principle, in increasing order of stringency.²⁶ The first required only that regulators take into account even highly unlikely catastrophes.²⁷ Another version “asks for a degree of risk aversion, on the theory that people do, and sometimes should, purchase insurance against the worst kinds of harm.”²⁸ Hence, Sunstein said, “a margin of safety is part of the Catastrophic Harm Precautionary Principle – with the degree of the margin depending on the costs of purchasing it.”²⁹ Finally, Sunstein suggested, “it sometimes makes sense to adopt a still more aggressive form of the Catastrophic Harm Precautionary Principle, one “selecting the worst-case scenario and attempting to eliminate it.”³⁰ More recently, Sunstein has endorsed use of maxmin in situations involving catastrophic risks, with some suggestions for possible guardrails

²³ *Id.* at 633.

²⁴ IPCC, 2021: Summary for Policymakers, in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* 13 (2021).

²⁵ Cass R. Sunstein, *Worst Case Scenarios* (2009). As Rowell Arden has pointed out, Sunstein’s view seems to have evolved over time into a greater acceptance of the use of maxmin to avoid catastrophic risks when their probabilities are unknown. See Arden Rowell, *Regulating Best-Case Scenarios*, 50 *Envtl. L.* 1105, 1121-1122 (2020)(notes 54-58 describe his evolving views). Rowell argues for consideration of uncertain but highly beneficial upside outcomes, something that can be accommodated by the α -maxmin approach described later in this paper.

²⁶ Cass R. Sunstein, *The Catastrophic Harm Precautionary Principle*, *ISSUES IN LEGAL SCHOLARSHIP, CATASTROPHIC RISKS: PREVENTION, COMPENSATION, AND RECOVERY* (2007) available at <https://ssrn.com/abstract=2532598>.

²⁷ *Id.* at 28.

²⁸ *Id.*

²⁹ *Id.*

³⁰ *Id.* Sunstein added a caution, however that maximin is “not generally a sensible strategy in the environmental context or elsewhere” because it makes no sense when risks can actually be quantified even roughly and is not attractive when the worst-case scenario is only mildly bad or when the cure inflicts “serious losses of its own.” *Id.* at 28-29.

to ensure that this test is applied sensibly.³¹ As Sunstein’s effort illustrates, it may be possible to clarify the areas of application for the precautionary principle sufficiently to make the principle a workable guide to decisions.

Maximin is a blunt tool, however, particularly in cases where precaution is very expensive in terms of resource utilization or foregone opportunities. There are variations on this principle and other possible tools for dealing with nonquantifiable risks.³² “Ambiguity” is a term that is often used to refer to situations in which the true probability distribution of outcomes is not known.³³ There is strong empirical evidence that people are averse to ambiguity. The classic experiment involves a choice between two urns. One is known to contain half red balls and half blue; the other contains both colors but in unknown proportions. Regardless of which color they are asked to bet on, most individuals prefer to place their bet on the urn with the known composition.³⁴ This is inconsistent with standard theories of rational decision making: if the experimental subjects prefer the known urn when asked to bet on red, that implies that they think that there are fewer than fifty percent red balls in the other urn. Consequently, they should prefer the second urn when asked to bet on blue—if it is less than half red it must be more than half blue—but they do not. Apparently, people prefer not to bet on an urn of uncertain composition.³⁵ Such aversion to ambiguity “appears in a wide variety of contexts.”³⁶ Ambiguity aversion may reflect a sense of lacking competence to evaluate a gamble.³⁷

³¹ Sunstein, *supra* note 5, at 968-970.

³² The discussion of these tools in this paper draws on Daniel A. Farber, *Uncertainty*, 99 *Geo. L.J.* 901 (2011), with significant updates and elaboration. The earlier article also discusses applications to climate change mitigation, climate change adaptation, and nanotechnology regulation. For further thoughts on uncertainty in the climate context, see Daniel A. Farber, *Coping with Uncertainty: Cost-Benefit Analysis, The Precautionary Principle, and Climate Change*, 90 *Wash. L. Rev.* 1659 (2015). On the related problem of fat-tailed probability distributions, see Daniel A. Farber, *Probabilities Behaving Badly: Complexity Theory and Environmental Uncertainty*, 37 *U.C. Davis L. Rev.* 145 (2003).

³³ For other legal applications of ambiguity models, see Uzi Segal & Alex Stein, *Ambiguity Aversion and the Criminal Process*, 81 *Notre Dame L. Rev.* 1495 (2006) (identifying and analyzing higher ambiguity aversion of defendants as opposed to prosecutors results in unbalanced plea negotiations); Eric Talley, *On Uncertainty, Ambiguity, and Contractual Conditions*, 34 *Del. J. Corp. L.* 755 (2009) (concluding that ambiguity explains why the possibility of adverse events sometimes results in the use of a conditions clause in a contract rather than a price adjustment); Joshua C. Teitelbaum, *A Unilateral Accident Model Under Ambiguity*, 36 *J. Legal Stud.* 431 (2007) (arguing that negligence rules may be superior to strict liability under conditions of ambiguity).

³⁴ See Gideon Keren & Léonie E.M. Gerritsen, *On the Robustness and Possible Accounts of Ambiguity Aversion*, 103 *Acta Psychologica* 149, 149 (1999) (“Ambiguity aversion is one of the most robust phenomena documented in the decision making literature....”) The researchers found in one experiment that “ambiguity avoidance is so pervasive that it extends even to situations in which the likelihood of winning in the ambiguity condition is higher than in the risky conditions.” *Id.* at 157.

³⁵ Nicholas Barberis & Richard Thaler, *A Survey of Behavioral Finance*, in *Handbook of the Economics of Finance* 1053, 1074-75 (George M. Constantinides, Milton Harris & René M. Stulz eds., 2003). The basic idea is that decision makers are averse to being faced with multiple possible probability distributions, as opposed to a situation where they are all combined into a single overall distribution. This corresponds to situations where there are multiple models of the world; each model is well understood but we’re unsure which one is correct.

³⁶ *Id.* at 1075.

³⁷ *Id.*

Ambiguity aversion may lead to irrational decision in some settings like laboratory experiments in which subjects know with certainty what scenarios are possible and their contents, it may be a much more reasonable attitude in practice. If we have two models of a situation available with completely different implications, that suggests that we do not understand the dynamics of the situation. That in turn means that the situation could be different from both models in some unknown way, and that whatever process the two models are trying to represent is really unknown to us. In the stylized example, perhaps there are green balls in the urn as well, or maybe the promised payoffs from our gamble won't appear. To take an example from foreign relations, it is one thing to deal with a leader who has a track record of being capricious and unpredictable (a situation involving risk); it is another to have no idea of who is leading the country on any given day (deep uncertainty). Our ability to plan for the future is limited in such situations, and it is reasonable to want to avoid being put in that position.

There are a number of different approaches to modeling uncertainty about the true probability distribution.³⁸ One is the Klibanoff—Marinacci— Mukerji model.³⁹ This approach assumes that decision makers are unsure about the correct probability. Their decision is based on (a) the likelihood that the decision maker attaches to different probability distributions, (b) the degree to which the decision maker is averse to taking chances about which probability distribution is right, and (c) the expected utility of a decision under each of the possible probability distributions. In simpler terms, the decision maker combines the expected outcome under each probability distribution according to the decision maker's beliefs about the distributions and attitude toward uncertainty regarding the true probability distribution. The shape of the function used to create the overall assessment determines in a straightforward way whether the decision maker is uncertainty averse, uncertainty neutral, or uncertainty seeking.⁴⁰

The Klibanoff-Marinacci-Mukerji model has an appealing degree of generality. But this model is not easily applied because the decision maker needs to be able to attach numerical weights to the specific probability distributions, which may not be possible in cases of true uncertainty where the possible distributions are themselves unknown.⁴¹ The model fits best with situations where our uncertainty is fairly tightly bounded: we know all of the possible models and how they behave, although we are unsure which one is correct.

Other models of ambiguity are more tractable and apply in situations where uncertainty may run deeper. As economist Sir Nicholas Stern explained, in these models of uncertainty, “the decision maker, who is trying to choose which action to take, does not know which of [several

³⁸ 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).

³⁹ Peter Kilbanoff, Massimo Marinacci & Sujoy Mukerji, *A Smooth Model of Decision Making Under Ambiguity*, 73 *Econometrica* 1849, 1849 (2005).

⁴⁰ *Id.* at 1869-70.

⁴¹ The model has been extended to dynamic-choice situations in which the decision maker receives additional information over time. See Peter Klibanoff, Massimo Marinacci & Sujoy Mukerji, *Recursive Smooth Ambiguity Preferences*, 144 *J. Econ. Theory* 930 (2009).

probability] distributions is more or less likely for any given action.”⁴² He explains that it can be shown the decision maker would “act as if she chooses the action that maximises 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.”⁴³

These models are sometimes called α -maxmin models, with α representing the weighting factor between best and worst cases.⁴⁴ 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, α 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.⁴⁵ In some situations, the best case scenario is more or less neutral, so that α -maxmin is not much different from pure loss avoidance, unless the decision maker is optimistic and uses an especially low α . But where the best case scenario is potentially extremely beneficial, unless the decision maker'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 decision maker, we can translate each outcome 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.

⁴² Nicholas Stern, *The Economics of Climate Change: The Stern Review* 39 (2007). For a discussion of the debate over the *Stern Review*, see Daniel H. Cole, *The Stern Review and Its Critics: Implications for the Theory and Practice of Benefit—Cost Analysis*, 48 *Nat. Resources J.* 53 (Winter 2008).

⁴³ *Id.*

⁴⁴ For discussion of the so-called α -maxmin model in the context of a more general theory, see Paolo Ghirardato, Fabio Maccheroni & Massimo Marinacci, *Differentiating Ambiguity and Ambiguity Attitude*, 118 *J. Econ. Theory* 133, 153-55 (2004).

⁴⁵ If $\alpha=1$, then α -maxmin becomes ordinary maxmin, in which only the worst case matters. For an axiomatic treatment of maxmin, see Itzhak Gilboa & David Schmeidler, *Maxmin Expected Utility With Non-Unique Prior*, 18 *J. Math. Econ.* 141 (1989).

An interesting variant of α -maxmin uses a weighted average that includes not only the best case and worst case scenarios, but also the expected value of the better understood, intermediate part of the probability distribution.⁴⁶ This approach “is a combination between the mathematical expectation of all the possible outcomes and the most extreme ones.”⁴⁷ This tri-factor approach may be “suitable for useful implementations in situations that entangle both more reliable (‘risky’) consequences and less known (‘uncertain’), extreme outcomes.”⁴⁸ However, this approach requires a better understanding of the mid-range outcomes and their probabilities than does α -maxmin.

We seem to be suffering from an embarrassment of riches, in the sense of having too many different method for making decisions in situations in which extreme outcomes weigh heavily. At present, it is not clear that any one method will emerge as the most useful for all situations. For that reason, the ambiguity models should be seen as providing decision makers 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, α -maxmin has a number of attractive features. First, it is complex enough to allow the decision maker to consider both the upside and downside possibilities, without requiring detailed probability information that is unlikely to be available. Second, it is transparent. 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. Third, α -maxmin can be useful in coordinating government policy. It is transparent to higher level decision makers and thus suited to central oversight.

Rather than asking the decision maker to assess highly technical probability distributions and modeling, α -maxmin simply presents the decision maker 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?⁴⁹ These questions are simple enough for politicians and members of

⁴⁶ See Marcello Basili, Alain Chateaufneuf & Fulvio Fontini, *Precautionary Principle as a Rule of Choice with Optimism on Windfall Gains and Pessimism on Catastrophic Losses*, 67 *Ecological Econ.* 485, 486 (2008).

⁴⁷ *Id.* at 490.

⁴⁸ *Id.*

⁴⁹ Adrian Vermeule has argued that any answer to this question is as good as any other:

Substantively, courts sometimes err by assuming that uncertainty demands worst-case reasoning. Courts, that is, assume that under uncertainty some version of worst-case (maximin) or (more generally) highly conservative assumptions are the only rational course. On the contrary, in the face of uncertainty a rational decision maker may set the α -value--the parameter that captures pessimism or optimism--anywhere within a range defined by the worst-case and best-case scenarios; courts should defer to agency choices about how pessimistic to be. There is an inescapable element of arbitrariness in the choice of an α -parameter, yet courts cannot improve the situation by demanding of agencies reasons that they cannot give or by requiring agencies to use maximally pessimistic assumptions--itself an arbitrarily chosen criterion.

Adrian Vermeule, *Rationally Arbitrary Decisions in Administrative Law*, 44 *J. Legal Stud.* S475, S478 (2015). As the term “fog of war” indicates, it is common for decisions made on a battlefield to take place under great uncertainty. Yet, if any decision were as good as any other, there would be no basis for criticisms any action the general might take. Similarly, any action taken by a surgeon who detects internal bleeding from an unknown source would be as good as any other. This seems implausible. Presumably, much of the point of professional training is to inculcate the ability to make reasonable choices under such circumstances. Sunstein, for example, as pointed to a number of qualitative factors that might shape the choice, such as the plausibility of the worse-case scenario and the

the public to understand. More importantly, rather than concealing value judgments in technical analysis by experts, they present the key value judgments directly to the elected or appointed officials who should be making them. Finally, these questions also lend themselves to oversight by outside experts, legislators, and journalists, which is desirable in societal terms even if not always from the agency's perspective.

3. Identifying Robust Solutions

Another approach, which also finds its roots in consideration of worst case scenarios, is to use scenario planning to identify unacceptable courses of action and then choose the most appealing remaining alternatives.⁵⁰ Robustness rather than optimality is the goal. RAND researchers have developed a particularly promising method to use computer assistance in scenario planning.⁵¹ RAND's Robust Decision Making (RDM) technique provides one systematic way of exploring large numbers of possible policies to identify robust solutions.⁵² As one of its primary proponents explains:

RDM rests on a simple concept. Rather than using computer models and data to describe a best-estimate future, RDM runs models on hundreds to thousands of different sets of assumptions to describe how plans perform in a range of plausible futures. Analysts then use visualization and statistical analysis of the resulting large database of model runs to help decisionmakers distinguish future conditions in which their plans will perform well from those in which they will perform poorly. This information can help decisionmakers

reasonableness of the costs of prevention. See text accompanying notes 27 to 34, *supra*. In the case of a regulatory agency, the relevant statute might shape the decision — consider a law requiring the agency choice include an adequate margin of safety. The problem may be that Vermeule's concept of rationality is too narrow to encompass the kinds of considerations that make one choice more reasonable than another. He is probably right, though, that in choosing within this range of opportunities, it may be enough for the decision maker to say: "I was chosen for an office which requires making this choice, and that legitimately authorizes me to make a value judgment about which way to go."

⁵⁰ Computer-based simulations can be used to explore the possibility space in other ways, such as agent-based modeling. See Ian S. Lustick and Philip E. Tetlock, *The Simulation Manifesto: The Limits of Brute-Force Empiricism in Geopolitical Forecasting*, *Futures & Foresight Sci.* 1 (2020), <https://doi.org/10.1002/ffo2.64>

⁵¹ An excellent introduction is provided by R.J. Lempert, *Robust Decision Making (RDM)*, in Marchau et al. (eds), *Decision Making Under Deep Uncertainty* (2019). Robustness can be defined in several different ways, including maxmin, α -maxmin, and minimizing regret (the difference in outcome from the strategy that would have been picked if future development had been known in advance). *Id.* at 34. Lempert illustrates RDM with a study designed to determine whether carbon tax, technological subsidies, or some combination is the best strategy in the face of economic unknowns. *Id.* at 37-43.

⁵² See David G. Groves, *New Methods for Identifying Robust Long-Term Water Resources Management Strategies for California* 12 (2006), available at http://www.rand.org/pubs/rgs_dissertations/2006/RAND_RGSD196.pdf For an effort to test the usability of this approach for water agencies, see David G. Groves et al., *Presenting Uncertainty About Climate Change to Water-Resource Managers: A Summary of Workshops with the Inland Empire Utilities Agency* 73-74 (2008), available at www.rand.org/pubs/technical_reports/2008/RAND_TR505.pdf. For a more recent application of this technique, this time to strategies for reopening after COVID, see De Lima et al., *Reopening California: Seeking Robust, Non-Dominated COVID-19 Exit Strategies*, 16 *PLoS One* (2021), <https://doi.org/10.1371/journal.pone.0259166>.

identify, evaluate, and choose robust strategies — ones that perform well over a wide range of futures and that better manage surprise.⁵³

During each stage of the analysis, RDM uses statistical analysis to identify policies that perform well over many possible situations.⁵⁴ It then uses datamining techniques to identify the future conditions under which such policies fail. New policies are then designed to cope with those weaknesses, and the process is repeated for the revised set of policies. As the process continues, policies become robust under an increasing range of circumstances, and the remaining vulnerabilities are pinpointed for decision makers.⁵⁵ More specifically, “RDM uses computer models to estimate the performance of policies for individually quantified futures, where futures are distinguished by unique sets of plausible input parameter values.”⁵⁶ Then, “RDM evaluates policy models once for each combination of candidate policy and plausible future state of the world to create large ensembles of futures.”⁵⁷ The analysis “may include a few hundred to hundreds of thousands of cases.”⁵⁸

The process could be compared with stress-testing various strategies to see how they perform under a range of circumstances. There are differences, however. RDM may also consider how strategies perform under favorable circumstances as well as stressful ones; it is able to test a large number of strategies or combination of strategies; and strategies are often modified (for instance by incorporating adaptive learning) as a result of the analysis.

A related concept, which discards strategies known to be dangerous, is known as the safe minimum standards (SMS) approach.⁵⁹ This approach may apply in situations in which there are discontinuities or threshold effects, but there is considerable controversy about its validity.⁶⁰ A related variant is to impose a reliability constraint, requiring that the odds of specified bad

⁵³ Richard J. Lempert et al., *Making Good Decisions Without Predictions: Robust Decision Making for Planning Under Deep Uncertainty*, https://www.rand.org/pubs/research_briefs/RB9701.html.

⁵⁴ See Groves, *New Methods*, *supra* note, at 123-48.

⁵⁵ See *id.* at 132-33 (describing the RDM robustness criterion).

⁵⁶ *Id.* at 124-25 (describing how the method achieves “robust” policies that are “relatively insensitive to the key uncertainties and different preferences held by decision makers”); see also David G. Groves & Robert J. Lempert, *A New Analytic Method for Finding Policy-Relevant Scenarios*, 17 *Global Environ. Change* 73, 75 (2007) (stating that “[t]he central idea is to use multiple runs of computer simulation models to identify those scenarios most important to the choices facing decision makers,” based on the foundation of RDM).

⁵⁷ Groves, *New Methods*, *supra* note 54, at 125.

⁵⁸ *Id.*

⁵⁹ See Michael Margolis & Eric Naevdal, *Safe Minimum Standards in Dynamic Resource Problems— Conditions for Living on the Edge of Risk* (Res. for the Future, Discussion Paper No. 04-03, 2004), available at www.rff.org/RFF/Documents/RFF-DP-04-03.pdf. Margolis and Naevdal show that “SMS is optimal policy if managers can put lower bounds on two parameters: the seriousness of the catastrophe and a parameter that determines how the magnitude of risk varies with the state-variable's position in state space.” *Id.* at 3.

⁶⁰ For further discussion and critique of this approach, see Michael C. Farmer & Alan Randall, *The Rationality of a Safe Minimum Standard*, 74 *Land Econ.* 287, 288 (1998) (arguing in favor of a harder version of this approach where a consensus exists that a resource is a human necessity); J.C. Rolfe, *Ulysses Revisited—A Closer Look at the Safe Minimum Standard Rule*, 39 *Aus. J. Agric. Econ.* 55, 66-68 (1995) (arguing for a softer version in which SMS is understood as merely a switching rule triggering more intensive scrutiny of costs and benefits).

outcomes be kept below a set level.⁶¹ The existence of threshold effects makes information about the location of thresholds quite valuable. For instance, in the case of climate change, a recent paper estimates that the value of early information about climate thresholds could be as high as three percent of gross world product.⁶²

3. Scenario Construction

The RDM methodology and α -maxmin are technical overlays on the basic idea of scenario analysis. Robert Verchick has emphasized the importance of scenario analysis—and of the act of imagination required to construct and consider these scenarios—in the face of nonquantifiable uncertainty.⁶³ As he explained, scenario analysis avoids the pitfall of projecting a single probable future when vastly different outcomes are possible; broadens knowledge by requiring more holistic projections; forces planners to consider changes within society as well as outside circumstances; and, equally importantly, “forces decision-makers to use their imaginations.”⁶⁴ He added that the “very process of constructing scenarios stimulates creativity among planners, helping them to break out of established assumptions and patterns of thinking.”⁶⁵ In situations in which it is impossible to give confident odds on the outcomes, scenario planning may be the most fruitful approach.⁶⁶

There is now a well-developed scholarly literature about the construction and uses of scenarios.⁶⁷ By 2010, it could be said that “combining rich qualitative storylines with quantitative modeling techniques, known as the storyline and simulation approach, has become the accepted method for integrated environmental assessment and has been used in all major assessments including those by the IPCC, Millennium Ecosystem Assessment, and many regional and national studies.”⁶⁸ Storylines “describe plausible, but alternative socioeconomic development pathways that allow scenario analysts to compare across a range of different situations, generally from 20 to 100 years into the future.”⁶⁹ Because our limited ability to predict the future evolution of society on a multi-decadal basis (let alone multi-century), scenarios are a particularly useful technique in dealing with problems involving long time spans and future generations. Best practices for constructing scenarios have also emerged. They include “pooling independently elicited

⁶¹ See David McInerney & Klaus Keller, *Economically Optimal Risk Reduction Strategies in the Face of Uncertain Climate Thresholds*, 91 *Climatic Change* 29, 30 (2007).

⁶² Klaus Keller et al., *What Is the Economic Value of Information About Climate Thresholds?*, in *Human-Induced Climate Change: An Interdisciplinary Assessment* 343, 349 fig. 28.5 (Michael E. Schlesinger et al. eds., 2007).

⁶³ Robert R. M. Verchick, *Facing Catastrophe: Environmental Action for a Post-Katrina World* 239-49 (2010). For an example of the use of scenario analysis to deal with a situation of high uncertainty, see Gary E. Machlis & Marcia K. McNutt, *Scenario-Building for the Deepwater Horizon Oil Spill*, 329 *Science* 1018 (2010).

⁶⁴ Verchick, *supra* note 65, at 242-43.

⁶⁵ *Id.* at 243.

⁶⁶ This is not to say that the process is without pitfalls. See Lustick and Tetlock, *supra* note 52, at 5.

⁶⁷ For a review of the literature, see Mark D.A. Rounsevell and Mac J. Metzger, *Developing Qualitative Scenario Storylines for Environmental Change Assessment*, 1 *WIREs Climate Change* 606 (2010).

⁶⁸ *Id.* at 608.

⁶⁹ *Id.*

judgments, extrapolating trend lines, and looking through data through the lenses of alternative assumptions.”⁷⁰

Within the scenario family, however, are several different subtypes, varying in the process for developing the scenario, whether the scenario is exploratory or designed to exemplify the pathway to a given outcome (such as achieving the Paris Agreements goals), or informal as opposed to probabilistic.⁷¹

Use of scenarios in the context of climate change has been a particular subject of attention.⁷² The IPCC has developed one set of scenarios (the SSP scenarios⁷³) for future pathways of societal development. Roughly speaking, these scenarios differ in the amount of international cooperation and whether society is stressing economic growth or environmental sustainability.⁷⁴ The SSP scenarios include detailed assumptions about population, health, education, economic growth, inequality and other factors.⁷⁵ A different set of scenarios, the RCPs,⁷⁶ are used to model climate impacts based on different future trajectories of GHG concentrations. Thus, “the RCPs generate climate projections that are not interpreted as corresponding to specific societal pathways, while the SSPs are alternative futures in which no climate impacts occur nor climate policies implemented.”⁷⁷ Integrated models are then used to combine the socioeconomic storylines, emissions trajectories, and climate impacts into a single simulation.⁷⁸ One significant finding is that only some scenarios are compatible with achieving the Paris Agreements’ aim of keeping warming to 1.5°C.⁷⁹ One advantage of the IPCC approach is that it provides standardized scenarios that can be used by many researchers looking at many different aspects of climate change.

There seems to be no one “right” way to construct scenarios. As with the assessment methods discussed above, the most important thing in practice may be for the agency to explain the options and the reasons for choosing one approach over the others. It might also be useful for agencies to standardize their scenarios where possible, which would create economies of scale in scenario development, allow comparison of results across different regulations, and provide focal points for researchers.

⁷⁰ Lustick and Tetlock, *supra* note 52, at 6.

⁷¹ Table 1 in Rounswell et al. provides a useful typography of the varieties of scenarios. See Rounswell et al., *supra* note 69, at 613.

⁷²For a review of the literature, see Brian C. O’Neill et al., *Achievements and Needs for the Climate Change Scenario*, 10 *Nature Climate Change* 1074 (2020).

⁷³ SSP stands for Shared Socioeconomic Development Pathways.

⁷⁴ See Rounsevell, *supra* note 69, at 610.

⁷⁵ O’Neill, *supra* note 74, at 1075.

⁷⁶ RCP stands for Representative Concentration Pathways.

⁷⁷ *Id.* at 1074.

⁷⁸ See Bas J. van Ruijven et al., *Enhancing the Relevance of Shared Socioeconomic Pathways for Climate Change Impacts, Adaptation and Vulnerability Research*, 122 *Climate Change* 481, 483 (2014).

⁷⁹ Joeri Rogelj et al., *Scenarios Towards Limiting Global Mean Temperature Increase Below 1.5 °C*, 8 *Nature Climate Change* 325 (April 2018). The research also revealed factors that enabled or blocked achievement of this target in various scenarios. *Id.* at 329-330.

Conclusion

Current guidance to agencies treats situations of deep uncertainty — those situations where we cannot reliably quantify risks — as an afterthought. This is particularly unfortunate in terms of catastrophic risks. Such risks are generally rare, meaning a sparse history of prior events (which may be completely lacking when novel threats are involved.) They often involve complex dynamics that make prediction difficult. The more catastrophic the event, the more likely it is to have impacts across multiple generations, changing the future in ways that are hard to predict. Yet planning only for better understood, more predictable risks may blind regulators to crucial issues. Catastrophic risks, even when they seem remote before the fact, may matter far more than the more routine aspects of a situation, in the spirit of “Except for that, how did you like the play, Mrs. Lincoln?” Fortunately, devastating pandemics, economic collapses, environmental tipping points, and disastrous runaway technologies are not the ordinary stuff of regulation. But where such catastrophes are relevant, marginalizing their consideration is dangerous.

There is no clear-cut answer to how we should integrate consideration of potential catastrophic outcomes into regulatory analysis. However, two decades after the last OMB guidance, we have a clearer understanding of the available analytic tools and better frameworks for applying them. If we cannot expect agencies to get all the answers right, we can at least expect them to carefully map out the analytic techniques they have used and why they have chosen them. This paper has explained the range of tools available and called for a revamping of OMB guidance to give clearer direction to agencies on their use. Even imperfect tools are better than relegating the risk of catastrophe to the shadows in regulatory decision making.