

Maximin

Cass R. Sunstein[†]

For regulation, some people argue in favor of the maximin rule, by which public officials seek to eliminate the worst worst-cases. The maximin rule has not played a formal role in regulatory policy in the United States, but in the context of climate change, pandemics, or new and emerging technologies, regulators who are unable to conduct standard cost-benefit analysis might be drawn to it. In general, the maximin rule is not a good idea for regulatory policy, because it is likely to reduce rather than to increase well-being. But under four imaginable conditions, that rule is attractive. (1) The worst-cases are very bad, and not improbable, so that it may make sense to eliminate them under conventional cost-benefit analysis. (2) The worst-case outcomes are highly improbable, but they are so bad that even in terms of expected value, it may make sense to eliminate them under conventional cost-benefit analysis. (3) Observers (including regulators) are in circumstances of Knightian uncertainty, where they cannot assign probabilities to imaginable outcomes. (4) The probability distributions may include “fat tails,” in which very bad outcomes are more probable than is usual; it may make sense to eliminate those outcomes for that reason. With respect to (3) and (4), the challenges arise when eliminating dangers also threatens to impose very high costs or to eliminate very large gains. There are also reasons to be cautious about imposing regulation when technology offers the promise of “moonshots,” or “miracles,” offering a low probability or an uncertain probability of extraordinarily high payoffs. Miracles may present a mirror image of worst-case scenarios.

[†] Robert Walmsley University Professor, Harvard University. I am grateful to Tyler Cowen, Annie Duke, and Eric Posner for superb comments on an earlier draft and to Dinis Cheian for extraordinary research assistance. A few sections of this Article draw on, while also significantly revising and updating, some sections of Cass R. Sunstein, *Irreversible and Catastrophic*, 91 CORNELL L. REV. 841 (2006).

I. In Brief..... 942

II. With and Without Numbers 945

III. Risk and Risk Aversion..... 951

 A. Numbers 951

 B. Precautions and Risk 956

 C. Danger 959

 D. OMB Circular A-4 961

 E. A Note on Loss Aversion 962

IV. Uncertainty and Ignorance..... 965

 A. Strategies of Avoidance 965

 B. Into the Thicket..... 966

 C. Precautions Again..... 968

V. Four Objections..... 970

 A. Triviality..... 970

 B. Maximin Assumes Infinite Risk Aversion 971

 C. Uncertainty Does Not Exist..... 972

 D. Uncertainty is Rare..... 974

VI. A Path Forward..... 976

Uncertainty must be taken in a sense radically distinct from the familiar notion of Risk, from which it has never been properly separated The essential fact is that ‘risk’ means in some cases a quantity susceptible of measurement, while at other times it is something distinctly not of this character; and there are far-reaching and crucial differences in the bearings of the phenomena depending on which of the two is really present and operating.

Frank Knight¹

One could certainly elicit from a political scientist the subjective probability that he attaches to the prediction that Norway in the year 3000 will be a democracy rather than a dictatorship, but would anyone even contemplate *acting* on the basis of this numerical magnitude?

Jon Elster²

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.

OMB Circular A-4³

I. In Brief

For regulators, what is the appropriate approach to worst-case scenarios? In the face of a pandemic, threatening to produce numerous deaths, should costly preventive measures be undertaken, even if the benefits are challenging or speculative to quantify? Or suppose that genetically modified foods pose a risk of catastrophe—very small, but not zero.⁴ Or suppose that some new

-
1. FRANK H. KNIGHT, *RISK, UNCERTAINTY, AND PROFIT* (1933).
 2. See JON ELSTER, *EXPLAINING TECHNICAL CHANGE: A CASE STUDY IN THE PHILOSOPHY OF SCIENCE* 199 (1983).
 3. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, CIRCULAR A-4 (2003), https://www.transportation.gov/sites/dot.gov/files/docs/OMB%20Circular%20No.%20A-4_0.pdf [<https://perma.cc/DE4M-5FBV>] [hereinafter CIRCULAR A-4].
 4. For one view, see Nassim Nicholas Taleb et al., *The Precautionary Principle (with Application to the Genetic Modification of Organisms)* (Sept. 4, 2014) (unpublished manuscript), <http://www.fooledbyrandomness.com/pp2.pdf> [<https://perma.cc/322V-H9ME>], and in particular *id.* at 11:

A lack of observations of explicit harm does not show absence of hidden risks. Current models of complex systems only contain the subset of reality that is accessible to the scientist. Nature is much richer than any model of it. To expose an entire system to something whose potential harm is not understood because extant models do not predict a negative outcome is not justifiable; the relevant variables may not have been adequately identified.

technology poses a catastrophic risk, but that experts cannot say whether it is very small, very large, or somewhere in between.⁵ Should regulators ban that technology? Should the social cost of carbon, designed to capture the damage from a ton of carbon emissions, reflect worst-case scenarios, and if so, exactly how?⁶

With a focus on regulatory policy, my goal here is to answer these questions. I will, above all, be attempting to carve out space for *the maximin rule*, which calls for choosing the approach that eliminates the worst of the worst-case scenarios. That rule has been subject to formidable objections, especially within economics, and I will be acknowledging and attempting to fortify those objections here. Nonetheless, my main aim is to show that the maximin rule deserves a place in regulatory policy. I shall attempt to specify the circumstances in which it deserves that place. Much of the discussion will be abstract, but I shall ultimately suggest a specific addition to OMB Circular A-4,⁷ the general framework for undertaking regulatory impact analysis; the goal of the addition is to codify a potential application of the maximin rule.

In extreme situations, regulators of diverse kinds must decide what kinds of restrictions to put in place against low-probability risks of catastrophe, or against risks that have terrible worst-case scenarios, but to which probabilities cannot readily be assigned. Some people, of course, favor quantitative cost-benefit analysis, whereas others favor some kind of precautionary principle. I am going to be embracing the former here, at least as a general rule, but the claims that deserve emphasis involve the exceptions, which call for precautionary thinking in general and for the maximin rule in particular.

I will be covering a great deal of ground, and while the journey is more important than the destination, it will be useful to specify the basic conclusions at the outset. The first three are straightforward. The remaining three are not.

- (1) Regulators should generally focus on expected value and on likely costs and benefits, not on worst cases.⁸ They should aim to come up with probability distributions, accompanied by point estimates.⁹ When they cannot produce probability distributions, they should try to come up with reasonable ranges of both costs and benefits.

5. Henry A. Kissinger, *How the Enlightenment Ends*, ATLANTIC (June 2018), <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124> [https://perma.cc/YB2F-PBBY].

6. *Social Cost of Greenhouse Gases*, OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, <https://obamawhitehouse.archives.gov/omb/oir/social-cost-of-carbon> [https://perma.cc/NB7S-BRGG].

7. CIRCULAR A-4, *supra* note 3.

8. I am bracketing the various problems with cost-benefit analysis, including the priority of welfare and the relevance of distributional considerations. See MATTHEW ADLER, *MEASURING SOCIAL WELFARE* (2019); CASS R. SUNSTEIN, *THE COST-BENEFIT REVOLUTION* (2017).

9. Point estimates, frequently provided by agencies, can often be understood as reflecting the mean of a probability distribution.

- (2) In some cases, the worst-cases are sufficiently bad, and sufficiently probable, that it may make sense to eliminate them, simply in terms of conventional cost-benefit analysis.¹⁰ (That idea appears to have informed the aggressive responses to the coronavirus pandemic in 2020.¹¹)
- (3) In some cases, the worst-case outcomes are highly improbable, but they are so bad that it may make sense to eliminate them under conventional cost-benefit analysis. (That is a reasonable view about costly efforts to reduce the risk of a financial crisis.¹²)
- (4) In some circumstances, involving what is often described as Knightian uncertainty, observers (including regulators) cannot assign probabilities to imaginable outcomes, and the maximin rule is appealing for that reason. I will argue that, contrary to a vigorously defended view in economics,¹³ the problem of uncertainty is real and sometimes important.
- (5) In some cases, a probability distribution might include “fat tails” on the left-hand side, in which the probability of extreme, very bad events is higher than normal; it might make sense to eliminate those very bad outcomes under conventional cost-benefit analysis. The fact that complex systems are involved might be important here; interactions among people or components of such systems might produce unanticipatedly bad results, as in the case of a pandemic.
- (6) With respect to (4) and (5), the problems arise when efforts to eliminate dangers, including regulation, would also impose very high costs or eliminate very large potential gains. There might be fat tails on the right-hand side, suggesting the possibility of wonders or miracles, which might make human life immeasurably better,¹⁴ and which might be eliminated by aggressive regulation.

This is a long and complicated list, so let us simplify it. In general, agencies should attempt to maximize social welfare (bracketing complex questions about what exactly that means).¹⁵ To do that, they should calculate costs and benefits, with probability distributions as feasible and appropriate,

10. There is also the question of reversibility, which may greatly matter to the cost-benefit analysis. The problem is discussed in Cass R. Sunstein, *Irreparability as Irreversibility*, 2017 SUP. CT. REV. 93. I bracket that issue here.

11. See Michael Greenstone & Vishan Nigam, *Does Social Distancing Matter?* (Univ. of Chi., Becker Friedman Inst. for Econ. Working Paper No. 2020-26, 2020), <https://ssrn.com/abstract=3561244> [<https://perma.cc/LT4K-TCVK>].

12. See Eric A. Posner & E. Glen Weyl, *Cost-Benefit Analysis of Financial Regulations: A Response to Criticisms*, 124 YALE L.J.F. 246 (2015).

13. An early account is FRANK RAMSEY, *Truth and Probability*, in THE FOUNDATIONS OF MATHEMATICS AND OTHER LOGICAL ESSAYS (R. B. Braithwaite ed., 1931). A vigorous defense of the importance and pervasiveness of Knightian uncertainty is JOHN KAY & MERVYN KING, RADICAL UNCERTAINTY: DECISION-MAKING BEYOND THE NUMBERS 35–49, 106–30 (2020). I agree with Kay and King on the question of importance but not quite on the issue of pervasiveness, for reasons explored below.

14. Arden Rowell, *Regulating Best-Case Scenarios*, 50 ENV. L. (forthcoming 2020), <https://ssrn.com/abstract=3157287> [<https://perma.cc/ZDX4-9D7Y>].

15. See ADLER, *supra* note 8.

and they should proceed if and only if the benefits justify the costs.¹⁶ They should not focus solely or mostly on the worst cases; they should not give them more weight than other cases (bracketing for now risk aversion or loss aversion, to which I shall turn in due course). At the same time, calculation of costs and benefits may not be feasible, and an important question remains: are there any problems that the maximin rule can handle better than welfare maximization? The simplest answer points to cases of Knightian uncertainty, where probabilities cannot be assigned.

As we shall also see, the maximin rule is especially appealing when the costs of eliminating the worst-case scenario are not terribly high, and when the worst-case scenario is genuinely grave. For reasons to be explained, we can see the simplest such cases as involving “negative freerolls,” which are best avoided. The argument for use of the maximin rule grows stronger as the badness of the worst-case scenario increases. It grows weaker as the costs of eliminating the worst-case scenario rise, and as that scenario becomes decreasingly grave.¹⁷

II. With and Without Numbers

Imagine that you have a heart condition but that you would like to continue doing strenuous exercise. You ask your doctor for advice, and she says that you probably should not, pointing to the risk of some kind of heart damage, which would in turn increase the risk of a stroke or a heart attack. Suppose that you ask her to assign probabilities to the range of possibilities, from “no adverse health effects at all” to “death.” Suppose that she says, “Okay, you’ve got me. The likelihood of no adverse health effects is very high—maybe 99%. The likelihood of a significant increase in risk is in the vicinity of 1%, probably less. The likelihood of death, as a result of the strenuous exercise that you propose, is trivially small.”

Under such circumstances, you may or may not continue doing strenuous exercise. An important question is how much you like doing it. You might want to weigh the hedonic and other benefits of strenuous exercise against the very small chance of significantly increasing your health risks. The outcome of that weighing will depend on your preferences—on what you care about. If you do not care much about strenuous exercise, you might decide, on precautionary grounds, to stop doing it. If the exercise is something that much matters to you, you might continue. Things might get more complicated if your doctor adds,

16. This claim is meant to be less rigid than it sounds. It should be taken as a presumption rather than a rule. Distributive considerations, or welfarist considerations, might trump the cost-benefit analysis. See MATTHEW ADLER, *WELFARE AND FAIR DISTRIBUTION* (2011). There may also be a legitimate role for risk aversion of certain kinds.

17. I am bracketing a possible institutional defense of the maximin rule, which is that it is a defense against some systematic bias on the part of regulators, such as undue optimism or short-term thinking. If regulators are systematically biased, the maximin rule might plausibly be a corrective.

parenthetically, that if you continue to exercise, there is a small chance that you will get significant health benefits and thus *reduce* the risk of death.

Now suppose instead that after you ask her to assign probabilities to the various outcomes, she says, “I can’t do that! No doctor can. For you, we just don’t know enough about the likelihood of any of the outcomes, including the bad ones.” What should you do? The doctor might be understood to say that this is a situation of Knightian uncertainty,¹⁸ in which probabilities cannot be assigned to various outcomes. Under such circumstances, some people would be drawn to the maximin rule: an approach that eliminates the worst-case scenario. With respect to pandemics, climate change, and regulation of new technologies, the same might be true. At least when some risk or technology has a terrible or catastrophic worst-case scenario, the best course might be to avoid it.

To understand what the doctor is saying in these cases, and the regulatory problem, we have to understand what it means to assign or to refuse to assign probabilities to future events. If the doctor refuses to do that, the simplest reason is that she lacks enough information. She might have a frequentist understanding of probability, in accordance with which she normally asks: in a large number of cases like this, how many times are there adverse health effects? This is the kind of question that someone might ask in assigning a probability to a fair coin coming up heads on fifty tosses, or a particular baby, born in Jerusalem on August 29, turning out to be female. When a doctor or regulator refuses to assign probabilities, the reason might be that she is a frequentist, and she might not have the kinds of information that frequentists require.¹⁹

An alternative understanding of probability judgments is Bayesian, and it does not depend on knowledge of frequencies.²⁰ It can be used for singular or unique cases.²¹ Bayesian approaches might be used when someone says that the probability of a pandemic five years from now is under 2%, that the probability that the Democratic nominee for president will win is 50%, or that the probability of a particular set of outcomes in 2100, as a result of climate

18. FRANK H. KNIGHT, *RISK, UNCERTAINTY, AND PROFIT* (1933); *see also* R. DUNCAN LUCE & HOWARD RAIFFA, *GAMES AND DECISIONS* 275–86 (1957).

19. For a vigorous defense of frequentism as the only plausible foundation of probability judgments, *see* KAY & KING, *supra* note 13, at 57–68, 110–22. *See also* Gerd Gigerenzer, *How to Make Cognitive Illusions Disappear: Beyond “Heuristics and Biases,”* 2 EUR. REV. SOC. PSYCHOL. 83 (1991); Gerd Gigerenzer, *Why the Distinction Between Single-Event Probabilities and Frequencies Is Important for Psychology (and Vice Versa)*, in SUBJECTIVE PROBABILITY 129, 129–61 (George Wright & Peter Ayton eds., 1994).

20. Eric-Jan Wagenmakers et al., *Bayesian Versus Frequentist Inference*, in BAYESIAN EVALUATION OF INFORMATIVE HYPOTHESES 181 (Herbert Hoijtink et al. eds., 2008).

21. For a brisk, illuminating notation, *see* Daniel Kahneman & Amos Tversky, *On the Reality of Cognitive Illusions*, 103 PSYCHOL. REV. 582, 586 (1996).

Whether or not it is meaningful to assign a definite numerical value to the probability of survival of a specific individual, we submit (a) that this individual is less likely to die within a week than to die within a year and (b) that most people regard the preceding statement as true—not as meaningless—and treat its negation as an error or a fallacy.

change, is over 90%. Bayesians start with a prior probability and then update on the basis of what they learn. Unlike frequentists, they are willing to assign probabilities to singular or nonrepeatable events. At the same time, a Bayesian doctor or regulator might agree that in a particular case, any subjective probability that she assigns to an event is speculative in the extreme; she might acknowledge that she lacks sufficient information to have any confidence in it. For that reason, she might agree that the situation is one of Knightian uncertainty.

It is important to note that frequentists believe that for unique or nonrepeatable events, assignments of probability are essentially meaningless.²² In their view, we have no basis for assigning a probability when we lack a frequency distribution. To say that a particular Democratic nominee has a 50% chance of being president, or that climate change is 90% likely to cause specified damage by 2100, is to speak nonsense, unless either statement can be plausibly justified in frequentist terms. For frequentists, the problem of Knightian uncertainty is therefore pervasive; it exists whenever we are dealing with a unique or nonrepeating problem, and we are doing that much of the time.²³ In my view, frequentists are unconvincing on that count, but it is not necessary to defend that conclusion for present purposes. Bayesians should also be willing to agree that Knightian uncertainty exists (a point to which I will return).

Consider in this regard a document from the White House, Principles for Regulation and Oversight of Emerging Technologies, issued in 2011 and still in effect.²⁴ In general, the document embraces cost-benefit analysis, but in a puzzlingly qualified way: “Benefits and costs: Federal regulation and oversight of emerging technologies should be based on an awareness of the potential benefits and the potential costs of such regulation and oversight, including recognition of the role of limited information and risk in decision making.”²⁵ What, exactly, is the role of “limited information”? What is the role of “risk”? With respect to regulation, the document explicitly calls out the problem of uncertainty: “The benefits of regulation should justify the costs (to the extent permitted by law and recognizing the relevance of uncertainty and the limits of quantification and monetary equivalents).”

The two sentences are different. The first refers to limited information and risk. The second refers to uncertainty and the limits of quantification. But with

22. See KAY & KING, *supra* note 13, at 74-84.

23. *See id.*

24. Memorandum from John P. Holdren, Assistant to the President for Sci. & Tech. Dir., Office of Sci. & Tech. Policy, Cass R. Sunstein, Admin., Office of Info. & Regulatory Affairs, Office of Mgmt. & Budget & Islam A. Siddiqui, Chief Agric. Negotiator, U.S. Trade Representative on Principles for Regulation and Oversight of Emerging Technologies to Heads of Exec. Dept’s & Agencies (Mar. 11, 2011), <https://obamawhitehouse.archives.gov/sites/default/files/omb/infoereg/agencies/Principles-for-Regulation-and-Oversight-of-Emerging-Technologies-new.pdf> [<https://perma.cc/7V3H-9QSU>].

25. *Id.*

respect to some problems, including those potentially raised by pandemics, climate change, and emerging technologies, we should understand the document, taken as a whole, to be emphasizing the epistemic limits of policymakers and regulators, and also to be drawing attention to the problem of Knightian uncertainty. These limits, and that problem, can be seen as qualifications to the general idea, pervasive in federal regulation, that regulators should proceed only if the benefits justify the costs.²⁶ OMB Circular A-4, a kind of Bible for federal regulatory analysis, explicitly recognizes both epistemic limits and Knightian uncertainty, and offers a plea for developing probability distributions to the extent feasible.²⁷ But what if it is not feasible to produce probability distributions, either because we lack frequencies or because Bayesian approaches cannot come up with them?

For a glimpse at the problem, consider a few numbers from cost-benefit reports from the Office of Information and Regulatory Affairs.

- (1) The projected annual benefits from an air pollution rule governing motor vehicles range from \$3.9 billion to \$12.9 billion.²⁸
- (2) The projected annual benefits of an air pollution rule governing particulate matter range from \$3.6 billion to \$9.1 billion.²⁹
- (3) The projected benefits of a regulation governing hazardous air pollutants range from \$28.1 billion to \$76.9 billion.³⁰
- (4) The projected benefits of a regulation governing cross-state air pollution range from \$20.5 billion to \$59.7 billion.³¹

26. See Improving Regulation and Regulatory Review, 3 C.F.R. § 13563 (2020).

27. See CIRCULAR A-4, *supra* note 3, at 41. The relevant passage is worth quoting at length:

Whenever possible, you should use appropriate statistical techniques to determine a probability distribution of the relevant outcomes. For rules that exceed the \$1 billion annual threshold, a formal quantitative analysis of uncertainty is required. For rules with annual benefits and/or costs in the range from 100 million to \$1 billion, you should seek to use more rigorous approaches with higher consequence rules. This is especially the case where net benefits are close to zero. More rigorous uncertainty analysis may not be necessary for rules in this category if simpler techniques are sufficient to show robustness.

28. OFFICE OF INFO. & REGULATORY AFFAIRS, OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, 2015 REPORT TO CONGRESS ON THE BENEFITS AND COSTS OF FEDERAL REGULATIONS AND AGENCY COMPLIANCE WITH THE UNFUNDED MANDATES REFORM ACT 25 (2015), https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2015_cb/2015-cost-benefit-report.pdf [<https://perma.cc/J2RA-687U>] [hereinafter 2015 Report].

29. OFFICE OF INFO. & REGULATORY AFFAIRS, OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, 2014 REPORT TO CONGRESS ON THE BENEFITS AND COSTS OF FEDERAL REGULATIONS AND UNFUNDED MANDATES ON STATE, LOCAL, AND TRIBAL ENTITIES 25 (2014), https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2014_cb/2014-cost-benefit-report.pdf [<https://perma.cc/H5US-8UJS>] [hereinafter 2014 Report].

30. OFFICE OF INFO. & REGULATORY AFFAIRS, OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, 2013 REPORT TO CONGRESS ON THE BENEFITS AND COSTS OF FEDERAL REGULATIONS AND UNFUNDED MANDATES ON STATE, LOCAL, AND TRIBAL ENTITIES 27 (2013), https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2013_cb/2013_cost_benefit_report-updated.pdf [<https://perma.cc/27AJ-EJ9C>].

31. OFFICE OF INFO. & REGULATORY AFFAIRS, OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, 2013 REPORT TO CONGRESS ON THE BENEFITS AND COSTS OF FEDERAL REGULATIONS AND UNFUNDED MANDATES ON STATE, LOCAL, AND TRIBAL ENTITIES 26 (2012),

It is worth pausing over three noteworthy features of these numbers. First, the government does not offer probability estimates to make sense of these ranges. It does not say that the probability at the low end is 1%, or 25%, or 50%. The default implication may be that the probability distribution is normal, so long as it is not specified, which might mean that the point forecast is the mean of the upper and lower bound. But is that what really is meant? Second, the ranges are exceptionally wide. In all four cases, the difference between the floor and the ceiling is much higher than the floor (which is in the billions of dollars)! Third, the wide ranges suggest that the worst-case scenario from government inaction, understood as a refusal to regulate, is massively worse than the best-case scenario. If regulators focus on the worst-case scenario, the relevant regulation is amply justified in all of these cases; there is nothing to discuss. The matter becomes more complicated if regulators focus on the best-case scenario or on the midpoint. But where should they focus?

All of these examples involve air pollution regulation, where projection of health benefits depends on significantly different models, leading to radically different estimates.³² But even outside of that context, relatively standard regulations, not involving new technologies, often project wide ranges in terms of benefits, costs, or both.³³ In terms of monetized costs, the worst case may be *double* the best case.³⁴ In terms of monetized benefits, the best case may be *triple* the worst case.³⁵ For a more general glimpse, consider this table, with particular reference to the wide benefits ranges³⁶:

https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2012_cb/2012_cost_benefit_report.pdf [https://perma.cc/B3U2-MR4L].

32. See 2015 Report, *supra* note 28, at 13-18.

33. See *id.* at 19.

34. See *id.* (food safety rules).

35. See *id.*

36. See 2014 Report, *supra* note 29.

Table 1. Estimates of Annual Benefits and Costs of Non-Environmental Related Health and Safety Rules: October 1, 2003 - September 30, 2013 (billions of 2001 and 2010 dollars)

Area of Safety and Health Regulation	Number of Rules	Estimated Benefits		Estimated Costs	
		2001\$	2010\$	2001\$	2010\$
Safety rules to govern international trade	3	\$0.9 to \$1.2	\$1.0 to \$1.4	\$0.7 to \$0.9	\$0.9 to \$1.1
Food safety	5	\$0.2 to \$9.0	\$0.3 to \$10.9	\$0.2 to \$0.7	\$0.3 to \$0.9
Patient safety	7	\$12.8 to \$21.9	\$12.8 to \$21.9	\$0.9 to \$1.1	\$1.1 to \$1.4
Consumer protection	3	\$8.9 to \$20.7	\$10.7 to \$25.0	\$2.7 to \$5.5	\$3.2 to \$6.6
Worker safety	5	\$0.7 to \$3.0	\$0.9 to \$3.6	\$0.6	\$0.7 to \$0.8
Transportation safety	24	\$13.4 to \$22.7	\$15.4 to \$26.4	\$5.0 to \$9.5	\$6.0 to \$11.4

Some of these gaps are very big, but for pandemics and new technologies, the difference between the worst and the best case might be (much) bigger still.³⁷ It is also important to emphasize that new or emerging technologies may be or include “moonshots,” understood as low-probability (or uncertain probability) outcomes with extraordinarily high benefits; call them miracles. Regulation might prevent those miracles,³⁸ or make them far less likely. In this domain, we may have “catastrophe-miracle” tradeoffs.

Because of its relevance to regulation of emerging technologies, I focus throughout on the difference between risk and uncertainty and urge that in the context of risk, adoption of the maximin rule is usually (not always) a fundamental mistake. Everything depends on the particular numbers, but in general, I aim to bury that rule, not to praise it. At the same time, I suggest that it deserves serious attention under identifiable conditions. When regulators

37. As an analogy, consider the social cost of carbon, with a range, in 2020 dollars, from \$12 to \$123 per ton. INTERAGENCY WORKING GROUP ON SOCIAL COST OF GREENHOUSE GASES, U.S. GOV'T, TECHNICAL SUPPORT DOCUMENT: TECHNICAL UPDATE OF THE SOCIAL COST OF CARBON FOR REGULATORY IMPACT ANALYSIS UNDER EXECUTIVE ORDER 12866 (Aug. 2016), https://obamawhitehouse.archives.gov/sites/default/files/omb/inforeg/scc_tsd_final_clean_8_26_16.pdf [<https://perma.cc/7URZ-MBFH>].

38. Rowell, *supra* note 14.

really are unable to assign probabilities to outcomes, and when some possible outcomes are catastrophic, the maximin rule has considerable appeal. Climate change is an obvious candidate for this conclusion,³⁹ and something similar might be said for some pandemics and other new or emerging risks, including some that are not even on the horizon.⁴⁰ But a great deal depends on what is lost by adopting the maximin rule. As we will see, catastrophic risks—of low or uncertain probability—may accompany both regulation and nonregulation. In addition, adoption of the maximin rule may rule out the possibility of miracles.

III. Risk and Risk Aversion

Does it *generally* make sense to eliminate the worst-case scenario? Put the question of uncertainty to one side and begin with numerical examples that involve risk instead. My topic is regulation, of course, but to make conceptual progress on that problem, it will be useful to provide stylized cases involving monetary gambles, which have the advantage of stripping away possible complications.

A. Numbers

Problem 1.

Which would you prefer?

- (a) A 99.9% chance of gaining \$10,000, and a 0.1% chance of losing \$6;
or
(b) A 50% chance of gaining \$5, and a 50% chance of losing \$5.

Under maximin, (b) is preferable, but under standard accounts of rationality, it would be much more sensible to select (a), which has a far higher expected value (outcome multiplied by probability). To choose (b), one would have to show an extraordinary degree of risk aversion.

Problem 2.

Which would you prefer?

- (a) A 70% chance of gaining \$100, and a 30% chance of losing \$30; or
(b) A 50% chance of gaining \$10, and a 50% chance of losing \$10.

39. See STEPHEN M. GARDNER, A PERFECT MORAL STORM: THE ETHICAL TRAGEDY OF CLIMATE CHANGE 411-14 (2011).

40. Broadly related arguments, emphasizing worst-cases and low-probability risks of catastrophe, can be found in Martin L. Weitzman, *Fat Tails and the Social Cost of Carbon*, 104 AM. ECON. REV. 544 (2014); Martin L. Weitzman, *Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change*, 5 REV. ENVTL. ECON. & POL. 275 (2011) [hereinafter Weitzman, *Fat-Tailed Uncertainty*]; and Martin L. Weitzman, *On Modeling and Interpreting the Economics of Catastrophic Climate Change*, 91 REV. ECON. & STAT. 1 (2009).

Under maximin, (b) is again preferable, but under standard accounts of rationality, it would still be much more sensible to select (a), which has a much higher expected value. We could easily proliferate examples in which the magnitude of risk aversion required to justify selection of (b) would be steadily reduced. For example:

Problem 3.

Which would you prefer?

(a) A 60% chance of gaining \$60, and a 40% chance of losing \$40; or

(b) A 50% chance of gaining \$10, and a 50% chance of losing \$10.

Here again, (a) has higher expected value, but it is less obvious that a chooser should choose it, at least if this is the only gamble that she will be offered (a point to which I will return), and at least if the welfare loss of losing \$40 is serious, even though the monetary figure is not so high.⁴¹ Examples of this kind can be mapped onto regulatory problems. For example, a decision to mandate widespread use of some new technology (say, electric cars) might take the form of Problem 2, where (a) is a mandate and (b) is no mandate. This could be so if we are not sure about the social costs and social benefits of such a mandate. Similarly, a decision to allow widespread use of some new technology (say, artificial intelligence in cancer treatment) might take the form of Problem 3, where (a) is widespread use and (b) is nonadoption. This could be so if the reliability of the new technology is not clear.

In life or in public policy, is risk aversion irrational? If one is making a very large number of monetary bets, it certainly is. If you had 10,000 questions like those immediately above, you should almost certainly choose (a). No gambler will do well if she keeps choosing (b).⁴² But in some circumstances, the answer is less obvious. Suppose that a seventy-year-old investor, Smith, is not in the best of health, and is deciding between two strategies for his pension. The first, called Caution, creates a 50% chance of no gain (aside from keeping up with inflation) and a 50% chance of an annual gain of 2%. The second, called Risky, creates a 25% chance of an annual loss of 5%, a 25% chance of no gain (aside from keeping up with inflation), a 25% chance of a 5% annual gain, and a 25% chance of a 10% annual gain.

In terms of expected value, Risky is much better. But without knowing about the effects of these outcomes on the chooser's welfare, it is hard to know which Smith should choose. There is the matter of worry: would Risky cause fear and sleeplessness? Then there is the matter of economics: how much

41. See KAY & KING, *supra* note 13, at 114-16.

42. For a superb discussion, with many implications for policy, see ANNIE DUKE, THINKING IN BETS (2018). I should note that for any gambler, the first bet must be made with an adequate bankroll, which means that a gambler would choose (a) only assuming that she had that. (Thanks to Annie Duke for this qualification.)

would a 5% loss matter to Smith? What would be the effect of a 5% gain? Perhaps a 5% loss would be devastating, given Smith's needs and wants, and perhaps a 5% gain would not much matter. Whether risk aversion is rational depends on the answer to these questions. The monetary figures are insufficient, because they do not tell us about the effects on Smith's welfare. The analysis is similar to the heart disease example with which I began. Something similar might be true in the regulatory context; we need to know what the gains and the losses actually mean, in terms of welfare.

And what happens if the worst cases are catastrophically bad?

Problem 4.

Which would you prefer?

- (a) A 99.99% chance of gaining \$60, and a 0.01% chance of losing \$100 million (resulting in a negative expected value); or
 (b) A 50% chance of gaining \$10, and a 50% chance of losing \$10.

Even if we know everything we need to know, (b) is better, at least in terms of expected value. The example shows that a low-probability risk of catastrophe can drive the outcome of cost-benefit analysis, even if the probability is low indeed, and even if we put risk aversion to one side. Calling attention to "fat tails," Martin Weitzman has emphasized something like this point in the context of climate change.⁴³ The problem of fat tails is not captured in Problem 4; fat tails consist of unusual probability distributions, when the likelihood of bad outcomes is unusually high at the extremes, including cases in which *the likelihood of terrible outcomes is unusually high on the left-hand side*. Thus:

Problem 5.

Which would you prefer?

- (a) A 99% chance of gaining \$60, a 0.01% chance of losing \$10, and a 0.09% chance of losing \$100 million; or
 (b) A 50% chance of gaining \$10, and a 50% chance of losing \$10.

Problem 5 involves a very fat tail (on the left), and (b) is better on cost-benefit grounds. Whether we are dealing with low-probability risks of catastrophe or fat tails, the magnitude of the potential harm can call for serious caution. The point may apply to many problems, including that of pandemics and risky technologies. (In such cases, the fact that we are dealing with complex systems, and unpredictability about how they will work, may be exceptionally important. Observers might not foresee what kinds of outcomes will be produced by interactions among component parts, or among people, as

43. See Weitzman, *Fat-Tailed Uncertainty*, supra note 40.

in the case of exponential growth in illness and death during a pandemic.) Consider Weitzman's important suggestion, focusing on climate change:

Deep structural uncertainty about the unknown unknowns of what might go very wrong is coupled with essentially unlimited downside liability on possible planetary damages. This is a recipe for producing what are called "fat tails" in the extremes of critical probability distributions. There is a race being run in the extreme tail between how rapidly probabilities are declining and how rapidly damages are increasing. Who wins this race, and by how much, depends on how fat (with probability mass) the extreme tails are. It is difficult to judge how fat the tail of catastrophic climate change might be because it represents events that are very far outside the realm of ordinary experience.⁴⁴

In this passage, Weitzman combines an emphasis on "the unknown unknowns," or uncertainty, with a reference to "the extremes of probability distributions."⁴⁵ Problems 4 and 5 do not involve uncertainty. They point only to extreme outcomes, which can be enough to dominate the comparison of expected values. These, then, are cases in which the maximin rule might be justified on the ground that it does not conflict with what would emerge from an analysis of expected value; because of the sheer magnitude of the harm in the worst-case scenario, it has outsized importance in the judgment about what to do. (To be sure, risk-seeking choosers might take their chances with (a).) As I have noted, this might be the right analysis of certain pandemics, especially when we emphasize the possibility (probability?) of exponential growth in infections and deaths.

Note, however, that in some cases, variations on Problem 4 are imaginable and illuminating. For example:

Problem 6.

Which would you prefer?

- (a) A 99.99% chance of gaining \$60, and a 0.01% chance of losing \$100 million; or
- (b) A 49.99% chance of gaining \$10, a 50% chance of losing \$10, and a 0.01% chance of losing \$100 million.

Problem 6 shows that low-probability, high-magnitude outcomes might accompany both options. On one view (with admittedly contested assumptions), climate change is an example. Immediate, very costly steps might be necessary to avert catastrophic risks, but they might themselves impose catastrophic risks, if (for example) they might threaten to create some massive economic downturn and geopolitical instability. (We could easily alter

44. *Id.* at 275.

45. *See also id.* at 285 ("The result of this lengthy cascading of big uncertainties is a reduced form of truly extraordinary uncertainty about the aggregate welfare impacts of catastrophic climate change, which is represented mathematically by a PDF that is spread out and heavy with probability in the tails.").

Problems 5 and 6 so as to include uncertainty.) With respect to new or emerging technologies, of course, there may be potentially massive upsides as well as potentially catastrophic downsides. Artificial intelligence and machine learning are possible examples.⁴⁶ In that regard, consider this:

Problem 7.

- (a) A 51% chance of gaining \$60, and a 49% chance of losing \$1; or
 (b) A 49.99% chance of gaining \$10, a 50% chance of losing \$10, and a 0.01% chance of gaining \$100 million.

This is a problem of “moonshots” or “miracles,” understood as low-probability chances of extraordinary returns.⁴⁷ We can also imagine “fat heads,” parallel to fat tails, or more properly, fat tails on both sides of the probability distribution. Here again, Problem 7 could be altered so as to include uncertainty. If the magnitude of those returns is high enough, they can dwarf the calculation of expected value. On standard grounds, maximax (maximize the best-case scenario) would be the right decision rule. We could also imagine cases in which an option has a negative expected value, but in which the moonshot is nonetheless a reasonable gamble. And if (b) in Problem 7 is combined with (a) in Problem 4, we will face “catastrophe-miracle” tradeoffs, here in circumstances of risk. (With uncertainty, the analytical challenge is even harder, though if catastrophes are bad enough—say, extinction—they may justifiably loom larger than miracles.)

What about the option of inaction? Every one of the foregoing problems could be understood to *include* inaction as one of the two options, producing one of the relevant payoffs, or could be designed so as explicitly to include that option. A simple example, where (b) is understood to mean inaction:

Problem 8.

Which would you prefer?

- (a) A 50% chance of gaining \$1.5 million, and a 50% chance of losing \$1 million; or
 (b) A 50% chance of no change from the status quo, and a 50% chance of losing \$500,000 from the status quo.

In terms of expected value, (a) is better. But in a one-shot gamble, the right choice might not be so clear. One more time: for individuals, a gain of \$1.5 million may produce less welfare than would be lost by a loss of \$500,000. There is a difference between expected value and expected utility (or welfare).

46. See Sendhil Mullainathan and Jann Spiess, *Machine Learning: An Applied Econometric Approach*, 31 J. ECON. PERSP. 87, 98-104 (2017).

47. Rowell, *supra* note 14.

Once we transform money into welfare, (b) might start to look more attractive, even if we put loss aversion (taken up shortly) to one side.

Consider one more case, attempting to broaden the viewscreen:

Problem 9.

Which would you prefer?

(a) A 50% chance of losing \$100 million, and a 50% chance of losing \$200 million; or

(b) An 80% chance of losing \$50 million and a 20% chance of losing \$90 million.

Option (b) is obviously better, though both are bad. (I am understanding the numbers as net losses, compared to the status quo.) We need not speak of the maximin rule in order to reach that conclusion. In 2020, the coronavirus pandemic could easily have been analyzed in terms of Problem 9, with aggressive responses producing (b), and much less aggressive responses producing (a).⁴⁸ Problem 9 is instructive because it shows that when aggressive regulation and nonregulation (understood to include weak regulation) *both* impose significant and even catastrophic net losses, an understanding of standard cost-benefit analysis can call for aggressive regulation.

B. Precautions and Risk

What is the appropriate role of risk aversion in the regulatory context? Should regulators focus on worst-case scenarios? Should they adopt the maximin rule?⁴⁹ When?

For certain regulatory problems, many people accept the Precautionary Principle.⁵⁰ The idea takes multiple forms, some far more cautious and targeted than others,⁵¹ but it is often understood to embody a commitment to risk aversion. The central idea is that regulators should take aggressive action to avoid certain risks, even if they do not know that those risks will come to fruition. Suppose, for example, that there is some probability that genetic modification of food will produce serious environmental harm.⁵²

48. Greenstone & Nigam, *supra* note 11.

49. An influential paper, suggesting the rationality of either maximin or maximax (maximize the best-case scenario), is Kenneth Arrow & L. Hurwicz, *An Optimality Criterion for Decision-Making Under Uncertainty*, in *UNCERTAINTY AND EXPECTATION IN ECONOMICS 1* (C.F. Carter & J.L. Ford eds., 1972).

50. For general discussion, see CASS R. SUNSTEIN, *LAWS OF FEAR* (2006).

51. See Taleb et al., *supra* note 4.

52. *Id.* Taleb et al. focus on “propagating impacts resulting in irreversible and widespread damage.” *Id.* at 1. In their understanding, the Precautionary Principle is designed “to avoid a certain class of what, in probability and insurance, is called ‘ruin’ problems. A ruin problem is one where outcomes of risks have a non-zero probability of resulting in unrecoverable losses.” *Id.* at 2.

For those who embrace the Precautionary Principle, it is important to take precautions against potentially serious hazards, simply because it is better to be safe than sorry. Thus, for example, the 1992 Rio Declaration states, “Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation.”⁵³ The Wingspread Declaration goes somewhat further: “When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically. In this context the proponent of an activity, rather than the public, should bear the burden of proof.”⁵⁴

Whatever the preferred formulation, the Precautionary Principle can be seen as an effort to build in a kind of margin of safety, perhaps because of “a clear normative presumption in favour of particular values or qualities—for instance concerning [the] environment or human health. This is instead of (for example) economic, sectoral, or partisan institutional interests.”⁵⁵ In certain forms, the principle might be taken to reflect the maximin principle: rule out the worst of the worst-case scenarios. But insofar as we are speaking about risk aversion in general, the Precautionary Principle runs into a serious objection: risks may be on all sides of social situations. Regulators are often dealing with *risk-risk tradeoffs* or even *health-health tradeoffs*.⁵⁶ When this is so, it is not helpful to speak of “a clear normative presumption in favour of . . . human health,” because human health is at risk whatever choice regulators make.⁵⁷

Suppose, for example, that steps are taken to regulate or ban genetically modified food on precautionary grounds.⁵⁸ Many people believe that any such steps might well result in numerous deaths, and a small probability of many more.⁵⁹ The reason is that genetic modification holds out the promise of

53. WINGSPREAD CONFERENCE ON THE PRECAUTIONARY PRINCIPLE, SCI. & ENVTL. HEALTH NETWORK, THE WINGSPREAD STATEMENT ON THE PRECAUTIONARY PRINCIPLE (1998) (quoted in BJORN LOMBORG, *THE SKEPTICAL ENVIRONMENTALIST* 347 (2001)).

54. See *The Precautionary Principle*, RACHEL’S ENVT. & HEALTH WKLY. (Envtl. Res. Found., Annapolis, Md.), Feb. 19, 1998.

55. See Andrew Stirling, *Precaution in the Governance of Technology*, in OXFORD HANDBOOK OF LAW, REGULATION, AND TECHNOLOGY 645, 649 (Roger Brownsword et al. eds., 2017).

56. See JOHN GRAHAM AND JONATHAN WIENER, *RISK VS. RISK* (1997). To that extent, it is not right to say that “criticism of the precautionary principle” is necessarily or generally rooted “on the overtly political grounds that it addresses general concerns like environment and human health, rather than more private interests like commercial profit or the fate of a particular kind of technology.” Stirling, *supra* note 55, at 650. The “general concerns” may be on both sides.

57. See Cass R. Sunstein, *Health-Health Tradeoffs*, 63 U. CHI. L. REV. 1533 (1996).

58. See Tony Gilland, *Precaution, GM Crops, and Farmland Birds*, in RETHINKING RISK AND THE PRECAUTIONARY PRINCIPLE 84, 84-88 (Julian Morris ed., 2001); *Are the US and Europe Heading for a Food Fight Over Genetically Modified Food?*, PEW INITIATIVE FOOD & BIOTECHNOLOGY (Oct. 24, 2001), <https://web.archive.org/web/20071011163512/http://pewagbiotech.org/events/1024/> [<https://perma.cc/AAZ4-WLWJ>] (archived from the original).

59. BILL LAMBRECHT, *DINNER AT THE NEW GENE CAFE: HOW GENETIC ENGINEERING IS CHANGING WHAT WE EAT, HOW WE LIVE, AND THE GLOBAL POLITICS OF FOOD* (2001) (tracing but not endorsing the various objections).

producing food that is both cheaper and healthier—resulting, for example, in “golden rice,” which might have large benefits in developing countries.⁶⁰ The point is not that genetic modification will definitely have those benefits, or that the benefits of genetic modification outweigh the risks, or that precautions are a bad idea. The point is only that if the Precautionary Principle is taken in certain ways, it is offended by regulation as well as by non-regulation. To be sure, the maximin principle might prove helpful here, on a certain set of empirical assumptions—an issue to which I will return.

Or consider regulation of autonomous vehicles.⁶¹ There is no question that such vehicles pose risks to public safety. Some of them crash. At the same time, a failure to allow autonomous vehicles, or even to promote them, or perhaps even to *mandate* them, might well be seen to offend the Precautionary Principle, because the result would be, with some probability, to cost lives.⁶² Use of autonomous vehicles might well increase safety, perhaps dramatically. We are dealing with safety-safety tradeoffs. The example shows again that if it is understood in a certain way, the principle seems to forbid the very steps that it requires. To make progress, it would seem necessary, not to speak of precautions or to invoke maximin, but to identify the possible outcomes and to specify the probability that they will occur. That will rapidly move us in the direction of cost-benefit analysis. But what if important information is absent?

To see how hard that question might bite, imagine that technical analysts inform political officials that if they proceed with a regulation, the monetized benefits will have a range of \$300 million to \$1.5 billion, and that the monetized costs will have a range of \$200 million to \$1.6 billion.⁶³ Suppose that the analysts add that they cannot assign probabilities to various points within the range. We seem to have not only a risk-risk tradeoff, in the sense that risks lie on both sides of the problem, but also an uncertainty-uncertainty tradeoff, in the sense that analysts identify outcomes without probabilities on both sides.⁶⁴ Should we say that the agency should not proceed because \$1.6 billion is higher than \$1.5 billion? That is hardly clear.

60. *Id.*

61. U.S. DEP’T OF TRANSP., PREPARING FOR THE FUTURE OF TRANSPORTATION: AUTOMATED VEHICLES 3.0 (AV 3.0) (2018), <https://www.transportation.gov/av/3/preparing-future-transportation-automated-vehicles-3> [<https://perma.cc/NVA5-X39J>].

62. Teena Maddox, *How Autonomous Vehicles Could Save Over 350K Lives in the US and Millions Worldwide*, ZDNET (Feb. 1, 2018), <https://www.zdnet.com/article/how-autonomous-vehicles-could-save-over-350k-lives-in-the-us-and-millions-worldwide> [<https://perma.cc/AA5R-BWSZ>].

63. The example is not so artificial; in the context of genetically modified food, for example, the Department of Agriculture projected first-year costs of between \$569 million and \$3.9 billion. See National Bioengineered Food Disclosure Standard, 83 Fed. Reg. 65,814, 65,869 (2020).

64. Note that the uncertainty is by hypothesis bounded; it is within specific ranges that probabilities cannot be assigned.

C. Danger

Now turn to a mundane illustration of the kinds of decisions in which the maximin rule might seem attractive: a reporter, living in Los Angeles, has been told that she can take one of two assignments. First, she can go to a nation, say Syria, in which conditions are dangerous (perhaps there is a military conflict). Second, she can go to Paris to cover anti-American sentiment in France. The Syria assignment has, in her view, two polar outcomes: a) she might have the most interesting and rewarding experience of her professional life, or b) she might be killed. The Paris assignment has two polar outcomes of its own: a) she might have an interesting experience, one that is also a great deal of fun, or b) she might be lonely and homesick. It might seem tempting for the reporter to choose Paris, on the ground that the worst-case scenario for that choice is so much better than the worst-case scenario for Syria. To know if this is so, she should probably think a bit about probabilities. She might not have numbers, but she might know enough to know, roughly, that the chance of being killed in Syria is quite small, but higher than in Paris, and that she would worry about that risk while in Syria. These points might incline her, reasonably enough, to choose Paris. And if this is correct, the conclusion might bear on regulatory policy, where one or another approach has an identifiably worst worst-case scenario.⁶⁵ To be sure, regulators would want to be more disciplined about both outcomes and probabilities.

But we have seen enough to know that maximin is not always a sensible decision rule. Suppose that the reporter now has the choice of staying in Los Angeles or going to Paris; suppose too that on personal and professional grounds, Paris is far better. It would make little sense for her to invoke maximin in order to stay in Los Angeles on the ground that the plane to Paris might crash. A plane crash is of course extremely unlikely, but it cannot be entirely ruled out. Using an example of this kind, John Harsanyi contends that the maximin rule should be rejected on the ground that it produces irrationality, even madness: "If you took the maximin principle seriously you could not ever cross the street (after all, you might be hit by a car); you could never drive over a bridge (after all, it might collapse); you could never get married (after all, it might end in a disaster), etc. If anybody really acted in this way he would soon end up in a mental institution."⁶⁶

Harsanyi's argument might also be invoked to contest the use of maximin in the choice between Syria and Paris. Perhaps the reporter should attempt to specify the likelihood of being killed in Syria, rather than simply identifying the worst-case scenario and resting content with intuitive assessments. Perhaps

65. See National Bioengineered Food Disclosure Standard, 83 Fed. Reg. at 65,869; Richard T. Woodward & Richard C. Bishop, *How to Decide When Experts Disagree: Uncertainty-Based Choice Rules in Environmental Policy*, 73 LAND ECON. 492 (1997).

66. John C. Harsanyi, *Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls' Theory*, 69 AM. POL. SCI. REV. 594, 595 (1975).

maximin is a way of neglecting probability, and hence a form of irrationality. In some circumstances, people do display probability neglect, in a way that ensures attention to the worst-case scenario.⁶⁷ But if probabilities can actually be assessed, and if that scenario is extremely unlikely to come to fruition, probability neglect is hard to defend even for people who are exceptionally risk-averse. Suppose that the risk of death, in Syria, turns out to be 1/1,000,000, and that the choice of Syria would be much better, personally and professionally, than the choice of Paris. Importantly, it is necessary to know something about the reporter's values and tastes to understand how to resolve this problem, but it is certainly plausible to think that the reporter should choose Syria rather than make the decision by obsessively fixating on the worst that might happen. The Council of Environmental Quality once did but no longer requires worst-case analysis; it refuses to do so on the ground that extremely speculative and improbable outcomes do not deserve attention.⁶⁸ So far, then, Harsanyi's criticism of maximin seems on firm ground.

But return in this light to the Precautionary Principle and notice that something important is missing from Harsanyi's argument and even from the reporter's analysis of the choice between Los Angeles and Paris. Risks, and equally bad worst-case scenarios, are on all sides of the hypothesized situations. If the reporter stayed in Los Angeles, she might be killed in one way or another, and hence the use of maximin does not by itself justify the decision to stay in the United States. And contrary to Harsanyi's argument, the maximin rule does not really mean that people should not cross streets, drive over bridges, and refuse to marry. The reason is that failing to do those three things has worst-case scenarios of its own (including death and disaster). To implement the maximin rule, or an injunction to take precautions, it is necessary to identify all relevant risks (including both outcomes and probabilities), not a subset.

Nonetheless, the more general objection to the maximin rule holds under circumstances of risk. If probabilities can be assigned to the various outcomes, it usually does not make sense to follow maximin when the worst case is exceptionally improbable and when the alternative option is both much better and much more likely. As noted, many people are risk-averse, or averse to particular risks, and on welfare grounds, some kinds of risk aversion, or aversion to particular risks, might be a good idea for individuals and societies. But when probabilities can be assigned, the maximin rule, imposed rigorously, seems to require infinite risk aversion.⁶⁹

67. See Cass R. Sunstein, *Probability Neglect: Emotions, Worst-cases, and the Law*, 112 YALE L.J. 61, 62-63 (2002).

68. See Todd S. Aagaard, *A Functional Approach to Risks and Uncertainties Under NEPA*, 1 MICH. J. ENVTL. & ADMIN. L. 87 (2012).

69. See Richard A. Musgrave, *Maximin, Uncertainty, and the Leisure Trade-Off*, 88 Q.J. ECON. 625, 626-28 (1974).

Compare this choice: (1) have a one-week family vacation in Florida, where it would be a great deal of fun, but where there is a 0.0001% chance of being killed by a Burmese python, or (2) stay home in Boston, where it would be relatively boring. Option (2) avoids (let us stipulate) the worst-case scenario, but does it really make sense to reject it for that reason? It follows that the reporter would do well to reject maximin and to go to Paris, even if the worst-case scenario for Paris is worse than that for Los Angeles, if the realistically likely outcomes are so much better in Paris.

These points are not meant to suggest that in order to be rational, the reporter must calculate expected values, multiplying imaginable outcomes by probability and deciding accordingly. Life is short; people are busy and occasionally risk-averse; anxiety and worry are themselves harms, and may cause harms; important information might be missing or unavailable; it is far from irrational to create a margin of safety to protect against disaster. But if the likelihood of a bad outcome is exceptionally small, and if much is to be gained by deciding in accordance with expected values, maximin is foolish. It does not make sense, as a general rule, to identify the worst-case scenario and to attempt to eliminate it. But the problem of uncertainty raises distinctive questions.

D. OMB Circular A-4

For regulatory impact analysis in the U.S. government, the key document is OMB Circular A-4, finalized in 2003.⁷⁰ That document offers a detailed discussion of how to proceed in the absence of complete information. It recognizes that

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.⁷¹

It adds that “whenever possible, you should use appropriate statistical techniques to determine a probability distribution of the relevant outcomes. For rules that exceed the \$1 billion annual threshold, a formal quantitative analysis of uncertainty is required.”⁷²

But that analysis might leave gaps, simply because insufficient information is available to produce specific numbers. In such cases, Circular A-4 offers guidance about how to proceed, calling for a “formal probabilistic analysis of the relevant uncertainties, possibly using simulation models and/or expert judgment.” In such assessments,

70. Circular A-4, *supra* note 3. A useful primer can be found at OFFICE OF INFO. & REGULATORY AFFAIRS, OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, REGULATORY IMPACT ANALYSIS: A PRIMER, https://www.reginfo.gov/public/jsp/Utilities/circular-a-4_regulatory-impact-analysis-a-primer.pdf [<https://perma.cc/9YT8-83FZ>] (last visited May 30, 2020).

71. Circular A-4, *supra* note 3, at 39.

72. *Id.* at 41.

expert solicitation is a useful way to fill key gaps in your ability to assess uncertainty. In general, experts can be used to quantify the probability distributions of key parameters and relationships. These solicitations, combined with other sources of data, can be combined in Monte Carlo simulations to derive a probability distribution of benefits and costs.⁷³

Optimistically, Circular A-4 concludes: “You should make a special effort to portray the probabilistic results—in graphs and/or tables—clearly and meaningfully.”⁷⁴

It is safe to say that the ambition of this discussion has not been fulfilled. In the context of air pollution rules, which sometimes cost at least \$1 billion, a formal probabilistic analysis is not usually offered. Instead agencies tend to report ranges.⁷⁵ There might be some pragmatic judgments in the background here. Agencies might be thinking that the analysis suggested by Circular A-4 is quite demanding, and if the benefits of a rule exceed the costs on any reasonable assumptions, the costs of undertaking the analysis might exceed the benefits. But without investigating particular problems in detail, we cannot know whether that is true. And in some cases, involving new risks and emerging technologies, the approach suggested by Circular A-4 might well be the right way to go.

Suppose, for example, that the technical analysis converges on these conclusions: *The cost of a regulation is \$1 billion. The benefits range from \$800 million to \$1.3 billion.* The first step would be to see if the benefits range could be turned into some kind of point estimate. The second would be to see if probabilities could be assigned to various points along the range, perhaps with the use of the approaches outlined in OMB Circular A-4. Under the Circular, the agency should be pressed to do exactly that.

E. A Note on Loss Aversion

People tend to be loss-averse, which means that they view a loss from the status quo as more undesirable than an equivalent gain is seen as desirable.⁷⁶ When we anticipate a loss of what we now have, we can become genuinely afraid, in a way that greatly exceeds our feelings of pleasure when we anticipate some (equivalent) supplement to what we now have. So far, perhaps, so good. The problem comes when individual and social decisions downplay potential gains from the status quo, and fixate on potential losses, in such a way as to produce overall increases in risks and overall decreases in well-being. The

73. *Id.*

74. *Id.* at 42.

75. See *supra* notes 28-31 and accompanying text.

76. See Colin Camerer, *Individual Decision Making*, in THE HANDBOOK OF EXPERIMENTAL ECONOMICS, 587, 665-670 (John H. Kagel & Alvin E. Roth, eds., 1995); Richard H. Thaler, *The Psychology of Choice and The Assumptions of Economics*, in QUASI-RATIONAL ECONOMICS 137, 143 (1991) (arguing that “losses loom larger than gains”); Daniel Kahneman, Jack L. Knetsch & Richard H. Thaler, *Experimental Tests of the Endowment Effect and the Coase Theorem*, 98 J. POL. ECON. 1325, 1328 (1990).

problem is heightened by the possibility that loss aversion is an “affective forecasting error”—that is, people might think (at the time of decision) that losses will have a much greater effect on their well-being than they actually do (in experience).⁷⁷

In the context of risk regulation, there is a clear implication: people will be closely attuned to the losses produced by any newly introduced risk, or by any aggravation of existing risks, but far less concerned with the benefits that are foregone as a result of regulation. The point very much bears on decisions of the Food and Drug Administration, where the risks of allowing unsafe or ineffective drugs on the market may be quite visible, while the risks of not allowing potentially safe and effective drugs on the market may be hidden. The point bears on the introduction of new technologies more generally, where regulators might be highly attuned to the risks of allowing them (and imposing losses), and less attuned to the risks of forbidding them (and failing to obtain gains). More generally, loss aversion often helps to explain what makes the Precautionary Principle operational. The opportunity costs of regulation may register little or not at all, whereas the threats posed by the activity or substance in question may be visible. In fact, this is a form of status-quo bias.⁷⁸ The status quo marks the baseline against which gains and losses are measured, and a loss from the status quo seems much worse than a gain from the status quo seems good.

If loss aversion is at work, we would predict that the Precautionary Principle would place a spotlight on the losses introduced by some risk and downplay the benefits foregone as a result of controls on that risk. Recall the emphasis, in the United States, on the risks of insufficient testing of medicines as compared with the risks of delaying the availability of those medicines. If the “opportunity benefits” are offscreen, the Precautionary Principle will appear to give guidance notwithstanding the objections I have made. At the same time, the neglected opportunity benefits sometimes present a serious problem with the use of the Precautionary Principle.

Loss aversion is closely associated with another cognitive finding: people are far more willing to tolerate familiar risks than unfamiliar ones, even if they are statistically equivalent.⁷⁹ For example, the risks associated with driving do not usually occasion a great deal of concern, even though in the United States alone, tens of thousands of people die from motor vehicle accidents each year. The relevant risks are simply seen as part of life. By contrast, many people are quite concerned about risks that appear newer, such as the risks associated with genetically modified foods, recently introduced chemicals, and terrorism. Part of the reason for the difference may be a belief that with new risks, we are in

77. Deborah A. Kermer et al., *Loss Aversion Is an Affective Forecasting Error*, 17 *PSYCHOL. SCI.* 649 (2006).

78. See William Samuelson and Richard Zeckhauser, *Status Quo Bias in Decision Making*, 1 *J. RISK & UNCERTAINTY* 7 (1988).

79. See PAUL SLOVIC, *THE PERCEPTION OF RISK* 140-43 (2000).

the domain of uncertainty (meaning that we cannot assign probabilities to bad outcomes) rather than risk (where probabilities can be assigned), and perhaps it makes sense to be cautious when we are not able to measure probabilities. But the individual and social propensity to focus on new risks outruns that sensible propensity. It makes the Precautionary Principle operational by emphasizing a subset of the hazards actually involved.

At first glance, it is tempting to think that if regulators fall prey to loss aversion, they will blunder. Consider a situation in which automated vehicles will produce twenty-five deaths that would not have occurred, but prevent fifty deaths that would have occurred. Unless those numbers conceal other factors, it seems clear that automated vehicles should be allowed. That is indeed the right result, but if people are loss averse, they might not weight a loss from a new technology in the same way that they would weight a loss from the status quo. Because loss aversion bears on public reactions, and because the public might be outraged or frightened by deaths that would not otherwise have occurred, regulators might have to work carefully to prevent beneficial new technologies from being discredited.

To test these questions, I conducted a survey on Amazon's Mechanical Turk, asking about 400 people to assume that in a city in their state, officials were deciding whether to go forward with a pilot project allowing automated vehicles on the road. Then I asked respondents this:

Imagine that the experts project that if automated vehicles are allowed, they would be responsible for 15 accidents that would not have otherwise occurred, during the next six months—but that automated vehicles would also prevent 50 accidents that would otherwise have occurred, in those next six months.

The question was whether the project should go forward. Fully 84% said "yes." When I changed the numbers to 20/30 (for another group), a strong majority (74%) again said "yes." A strong majority appears not to be loss averse, at least in the sense that they think that fewer overall accidents is the right test.

In general, the majority is correct on that point. But there is a countervailing consideration. Suppose that we are dealing with fat tails on both sides. (Recall that fat tails mean that at the extremes, probabilities are unusually high.) If things go very badly, we might have a catastrophe. If things go very well, we might have a miracle. Reasonable regulators might prevent a possible catastrophe, even if the price is to prevent a possible miracle. The downside risk of (say) extinction might reasonably be seen to deserve more attention than the upside potential of (say) immortality.

IV. Uncertainty and Ignorance

Now let us turn to what are, in a sense, the largest issues. In some contexts, risk-related problems involve hazards of ascertainable probability.⁸⁰ It may well be possible to say that the risk of death, from a certain activity, is 1/100,000, or at least that it ranges from (say) 1/20,000 to 1/500,000, with an exposed population of (say) 10 million. Or it may be possible to say that the risk of catastrophic harm from some activity is under 10% but above 1%. But as we have seen, it is possible to imagine instances in which analysts cannot specify even a range of probability, easily or at all,⁸¹ perhaps because they are frequentists who cannot find relevant frequencies, or perhaps because they are Bayesians who lack necessary information. Hence, regulators, and ordinary people, are sometimes acting in a situation of Knightian uncertainty (where outcomes can be identified but no probabilities can be assigned) rather than risk (where outcomes can be identified and probabilities assigned to various outcomes).⁸² And they are sometimes acting under conditions of ignorance, in which they are unable to specify either the probability of bad outcomes or their nature—where regulators do not even know the magnitude of the harms that they are facing.⁸³ One reason might be that they are dealing with a unique or nonrepeatable event. Another reason might be that they are dealing with a problem involving interacting components of a system, in which regulators cannot know much about how components of the system are likely to interact with each other.⁸⁴

A. Strategies of Avoidance

Of course, it is also true that over time, some problems that involve ignorance might shift to problems of uncertainty, and that problems of uncertainty might shift to problems of risk—a point that may counsel in favor of delay while new information is received. OMB Circular A-4 emphasizes this point: “For example, when the uncertainty is due to a lack of data, you might consider deferring the decision, as an explicit regulatory alternative, pending further study to obtain sufficient data.”⁸⁵ But as the circular notes, “Delaying a decision will also have costs, as will further efforts at data gathering and

80. In the remainder of this Article, I draw heavily on a section of Cass R. Sunstein, *Irreversible and Catastrophic*, 91 CORNELL L. REV. 841 (2006), while also revising and updating the discussion in significant ways.

81. KIYOHICO G. NISHIMURA & HIROYUKI OZAKI, *ECONOMICS OF PESSIMISM AND OPTIMISM: THEORY OF KNIGHTIAN UNCERTAINTY AND ITS APPLICATIONS* (2017); KNIGHT, *supra* note 1.

82. *See id.*; Paul Davidson, *Is Probability Theory Relevant for Uncertainty? A Post-Keynesian Perspective*, 5 J. ECON. PERSP. 129 (1991).

83. On ignorance and precaution, see Poul Harremoës, *Ethical Aspects of Scientific Incertitude in Environmental Analysis and Decision Making*, 11 J. CLEANER PRODUCTION 705 (2003).

84. *See* Taleb et al., *supra* note 4.

85. *See* Circular A-4, *supra* note 3.

analysis.”⁸⁶ Delay of regulation may mean serious harm (including large numbers of deaths; consider the coronavirus pandemic of 2020). In principle, agencies would calculate the costs and benefits of delay. But because of the very problem that counsels in favor of delay (lack of information), that calculation is not possible.

It is also true that agencies might use *breakeven analysis* to make progress in the face of uncertainty (at least if it is bounded).⁸⁷ Suppose, for example, that the costs of regulation are \$100 million, that the benefits range from \$150 million to \$5 billion, and that technical analysts state that at the present time, they cannot assign probabilities to the lower or upper bound, or to points along the range. Even so, it is clear that the regulation should go forward. Or suppose that the monetized costs of some new technology (say, a variation on fracking) are \$500 million, but that the monetized benefits range from \$600 million to \$10 billion. A regulatory ban would not be a good idea. We could easily imagine variations on these numbers. Breakeven analysis can enable regulators to identify reasonable paths forward even in the midst of uncertainty.

The Principle of Insufficient Reason says that when people lack information about probabilities (say, 1% to 40%), they should act as if each probability is equally likely.⁸⁸ But why is it rational to do so? By hypothesis, there is no reason to believe that each probability is equally likely. Making that assumption is no better than making some other, very different assumption. The Principle of Insufficient Reason is essentially arbitrary.⁸⁹

B. Into the Thicket

When strategies of avoidance are unappealing or unsuccessful, regulators might be drawn to the maximin rule: *Choose the policy with the best worst-case outcome*.⁹⁰ In the context of regulation of pandemics or new technologies, for example, perhaps elaborate precautions can be justified by reference to the maximin rule, asking officials to identify the worst case among the various options, and to select that option whose worst-case is least bad. Perhaps the maximin rule would lead to a Catastrophic Harm Precautionary Principle, by, for example, urging elaborate steps to combat potential risks. It follows that if aggressive measures are justified to reduce the risks associated with emerging technologies, one reason is that those risks are potentially catastrophic and existing science does not enable us to assign probabilities to the worst-case scenarios. The same analysis might be applied to many problems, including the

86. *Id.*

87. See Cass R. Sunstein, *The Limits of Quantification*, 102 CAL. L. REV. 1369 (2014).

88. See LUCE & RAIFFA, *supra* note 18, at 284; JOHN RAWLS, A THEORY OF JUSTICE 146 (revised ed. 1999) (“When we have no evidence at all, the possible cases are stipulated to be equally probable”).

89. See KAY & KING, *supra* note 13, at 63–64.

90. For a technical treatment of the possible rationality of maximin, see Arrow & Hurwicz, *supra* note 49; for a non-technical overview, see ELSTER, *supra* note 2, at 185–207.

risks associated with genetically modified food,⁹¹ nuclear energy,⁹² pandemics, and terrorism.

To understand these claims, we need to back up a bit. I have suggested that maximin has sometimes been recommended under circumstances of uncertainty rather than risk.⁹³ In an influential discussion, John Rawls, focusing on justice, offers a justification for a rule that “directs our attention to the worst that can happen.”⁹⁴ As he puts it, “this unusual rule” is plausible in light of “three chief features of situations.”⁹⁵ The first is that we cannot assign probabilities to outcomes, or at least we are extremely uncertain of them. The second is that the chooser “has a conception of the good such that he cares very little, if anything, for what he might gain above the minimum stipend that he can, in fact, be sure of by following the maximin rule.”⁹⁶ For that reason, it “is not worthwhile for him to take a chance for the sake of further advantage.” The third is that “the rejected alternatives have outcomes that one can hardly accept.” In other words, they involve “grave risks.” Under the stated conditions, the gains are limited from running a catastrophic risk, which means that choosers do not much value them, and it is worthwhile giving them up to protect against a downside outcome that choosers deplore.

Rawls emphasizes that the three “features work most effectively in combination,” which means that the “paradigm situation for following the maximin rule is when all three features are realized to the highest degree.”⁹⁷ That means that the rule does not “generally apply, nor of course is it self-evident.”⁹⁸ It is “a maxim, a rule of thumb, that comes in its own in special circumstances,” and “its application depends upon the qualitative structure of the possible gains and losses in its relation to one’s conception of the good, all this against a background in which it is reasonable to discount conjectural estimates of likelihoods.”⁹⁹

Rawls’ own argument is that for purposes of justice, the original position, as he understands it, is “defined so that it is a situation in which the maximin rule applies”¹⁰⁰—which helps to justify his principles of justice. It is

91. TALEB ET AL., *supra* note 4.

92. See Jon Elster, Explaining Technical Change: A Case Study in the Philosophy of Science 188-205 (1979) (unpublished manuscript) (on file with author).

93. See, e.g., ELSTER, *supra* note 2, 188-205 (1983).

94. See RAWLS, *supra* note 88, at 132-39. Rawls draws on but adapts WILLIAM FELLNER, PROBABILITY AND PROFIT 140-42 (1965).

95. RAWLS, *supra* note 88, at 134.

96. *Id.*

97. *Id.*

98. I am cheating a little bit here, referring to the original rather than the revised version of Rawls’ book. See JOHN RAWLS, A THEORY OF JUSTICE 155 (1971). (Sometimes the original is best.) It should be noted that in later work in particular, Rawls emphasized the Kantian foundations of the Veil of Ignorance, see JOHN RAWLS, POLITICAL LIBERALISM (1993), and those ideas could also be connected with the difference principle. I am bracketing that discussion for my purposes here.

99. JOHN RAWLS, A THEORY OF JUSTICE 155 (1971).

100. *Id.* (Note: This is only in the original, again.)

worthwhile noting that the same argument can help to identify situations in which maximax applies. Assume, first, that people are acting under conditions of uncertainty, or close to it. Assume, second, that the chooser “has a conception of the good such that he cares greatly for what he might gain by following the maximax rule.” Assume, finally, that grave or even significant risks are not involved, which is to say that if things go sour, and the chooser does not end up with the best possible outcome, he is nonetheless well enough off, given his conception of the good.

We can think of these cases as involving something akin to a “negative freeroll”: a choice in which one can incur losses but obtain no (real) gains.¹⁰¹ Who wants that? In such cases, applying maximin seems quite rational.

C. Precautions Again

These points bear on regulatory policy, where Rawls’ defense of maximin has inspired a defense and reconstruction of the Precautionary Principle in an important essay by Stephen Gardiner.¹⁰² To make the underlying intuition clear, Gardiner begins with the problem of choosing between two options, A and B:

If you choose A, then there are two possible outcomes: either (A1) you will receive \$100, or (A2) you will be shot. If you choose B, there are also two possible outcomes: either (B1) you will receive \$50, or (B2) you will receive a slap on the wrist. According to a maximin strategy, one should choose B. This is because: (A2) (getting shot) is the worst outcome on option A and (B2) (getting a slap on the wrist) is the worst option on plan B; and (A2) is worse than (B2).¹⁰³

It should be immediately apparent that if we can assign probabilities to outcomes, A might turn out to be the better choice. Suppose that if you choose A, there is a 99.99999% chance of A1, and that if you choose B, there is a 99.99999% chance of (B2). If so, A might seem better. But let us stipulate that assignment of probabilities is not possible. In Gardiner’s view, this conclusion helps support what he calls the Rawlsian Core Precautionary Principle in the regulatory setting: when Rawls’ three conditions are met, precautions, understood as efforts to avoid the worst-case scenario, should be adopted. As he puts it: “If one really were faced with the genuine possibility of disaster, cared little for the potential gains to be made by avoiding disaster and had no reliable information about how likely the disaster was to occur, then, other things being equal, choosing to run the risk might well seem like a foolhardy and thereby extreme option.”¹⁰⁴

101. I am grateful to Annie Duke for this point.
 102. See Stephen Gardiner, *The Core Precautionary Principle*, 14 J. POL. PHIL. 33 (2006).
 103. *Id.* at 46.
 104. *Id.* at 49.

Gardiner adds, importantly, that to justify the maximin rule, the threat posed by the worst-case scenario must satisfy some minimal threshold of plausibility. In his view, “the range of outcomes considered are in some appropriate sense ‘realistic,’ so that, for example, only credible threats are considered.”¹⁰⁵ If they can be dismissed as unrealistic, then maximin should not be followed. Gardiner believes that the problem of climate change, and also that of genetically modified organisms, can be usefully analyzed in these terms and that it presents a good case for the application of the maximin rule:

The RCPP [Rawlsian Core Precautionary Principle] appears to work well with those global environmental issues often said to constitute paradigm cases for the precautionary principle, such as climate change and genetically-modified crops. For reasonable cases can be made that the Rawlsian conditions are satisfied in these instances. For example, standard thinking about climate change provides strong reasons for thinking that it satisfies the Rawlsian criteria. First, the “absence of reliable probabilities” condition is satisfied because the inherent complexity of the climate system produces uncertainty about the size, distribution and timing of the costs of climate change. Second, the “unacceptable outcomes” condition is met because it is reasonable to believe that the costs of climate change are likely to be high, and may possibly be catastrophic. Third, the “care little for gains” condition is met because the costs of stabilizing emissions, though large in an absolute sense, are said to be manageable within the global economic system, especially in relation to the potential costs of climate change.¹⁰⁶

Gardiner adds, sensibly, that to justify maximin, the threats that are potentially catastrophic must satisfy some minimal threshold of plausibility.¹⁰⁷ Gardiner believes that the problem of climate change can be usefully analyzed in these terms and that it presents a good case for the application of maximin.¹⁰⁸ In a similar vein, Jon Elster, speaking of nuclear power, contends that maximin is the appropriate choice when it is possible to identify the worst-case scenario and when the alternatives have the same best consequences.¹⁰⁹ A related argument, ventured by Nassim Nicholas Taleb et al. in an illuminating discussion and specification of the Precautionary Principle, is that genetically modified crops pose a “ruin” problem, involving a low probability of catastrophically high costs.¹¹⁰

Taleb et al. contend that for such problems, it is best to take strong precautions—in this case, placing “severe limits” on genetically modified food. The discussion is technical, but let us bracket the science and suppose that it is correct. If so, the question is whether genetically modified crops really do

105. *Id.* at 51.

106. *Id.* at 55.

107. *See id.* at 51-52. There are some conceptual puzzles here. If an outcome can be dismissed as unrealistic, then we are able to assign at least some probabilities. Gardiner’s argument must be that in some cases we might know that the likelihood that a bad outcome will occur really is trivial.

108. *See id.* at 55.

109. *See* ELSTER, *supra* note 2, at 203.

110. TALEB ET AL., *supra* note 4.

create ruin problems. Perhaps they do, but it is also possible to read the most recent science to suggest that they do not; if the probability of catastrophic harm is vanishingly low and essentially zero, rather than merely very low, we can fairly ask whether Taleb's argument applies. If they can be dismissed as unrealistic, then maximin should not be followed.

But the larger point is that in identifiable circumstances, the argument for the maximin rule seems plausible. Taken seriously, this conclusion would have real consequences for regulatory policy, perhaps especially in the context of new risks or emerging technologies.

V. Four Objections

A. Triviality

An evident problem with this argument is that it risks triviality.¹¹¹ If individuals and societies can eliminate an uncertain danger of catastrophe for essentially no cost, then of course they should eliminate that risk. If people are asked to pay \$1 to avoid a potentially catastrophic risk to which probabilities cannot be assigned, they might as well pay \$1. And if two options have the same best-case scenario, and if the first has a far better worst-case scenario, people should of course choose the first option.

There is nothing wrong with this argument, but the real world rarely presents problems of this form. Where policy and law are disputed, the elimination of uncertain dangers of catastrophe imposes both costs and risks. In the context of climate change, for example, it is implausible to say that regulatory choosers can or should care "very little, if anything," for what might be lost by following maximin. If nations followed maximin for climate change, they would spend a great deal to reduce greenhouse gas emissions.¹¹² The result would almost certainly be higher prices for gasoline and energy, probably producing increases in unemployment and poverty.

Something similar can be said about genetic modification of food, because elimination of the worst-case scenario, through aggressive regulation, might well eliminate an inexpensive source of nutrition that would have exceptionally valuable effects on countless people who live under circumstances of extreme

111. Cf. David Kelsey, *Choice Under Partial Uncertainty*, 34 INT'L ECON. REV. 297, 305 (1993):

It is often argued that lexicographic decision rules such as maximin are irrational, since in economics we would not expect an individual to be prepared to make a small improvement in one of his objectives at the expense of large sacrifices in all of his other objectives. This criticism is less powerful in the current context since we have assumed that the decision maker has a weak order rather than a cardinal utility function on the space of outcomes. Given this assumption the terms "large" and "small" used in the above argument are not meaningful.

In many contexts, however, decision makers do have a cardinal utility function, not merely a weak order.

112. See WILLIAM D. NORDHAUS & JOSEPH BOYER, *WARMING THE WORLD: ECONOMIC MODELS OF CLIMATE CHANGE* 168 (2000).

deprivation.¹¹³ If we eliminate the worst-case scenarios for all pandemic risks, people will simply be required to stay at home, today, tomorrow, and the day after. While that might be the right approach, the fact that a very bad worst-case scenario is associated with the pandemic (worse, let us stipulate, than the worst-case associated with the mandate) cannot easily be taken to justify that mandate without trying to know more about probabilities.

The real question, then, is whether regulators should embrace maximin in real-world cases in which doing so is costly or extremely costly. If they should, it is because condition (3) is too stringent and should be abandoned. Even if the costs of following the maximin rule are significant, and even if regulators care a great deal about incurring those costs, the question is whether it makes sense to follow the maximin rule when they face uncertain dangers of catastrophe. In the environmental context, some people have so claimed.¹¹⁴ This claim takes us directly to the next objection to maximin.

B. Maximin Assumes Infinite Risk Aversion

Rawls' arguments in favor of adopting maximin, for purposes of distributive justice, were subject to withering critiques from economists—critiques that many economists accept to this day.¹¹⁵ The central challenge was that the maximin principle would be chosen only if choosers showed *infinite risk aversion*. In the words of one of Rawls' most influential critics, infinite risk aversion "is unlikely. Even though the stakes are great, people may well wish to trade a reduction in the assured floor against the provision of larger gains. But if risk aversion is less than infinite, the outcome will not be maximin."¹¹⁶ To be more specific: suppose that you have a choice between two options. Option A has a 99.9999% likelihood of great wealth and welfare and a 0.0001% likelihood of a terrible outcome. Option B has a 60% chance of a very bad outcome and a 40% chance of a just-short-of-terrible outcome. Would it really make sense to choose Option B?

To adapt this objection to the environmental context: it is plausible to assume a bounded degree of risk aversion with respect to catastrophic harms, to support some modest forms of a Catastrophic Harm Precautionary Principle. But even under circumstances of uncertainty—the argument goes—maximin is senseless unless societies are to show infinite risk aversion.

113. See Kym Anderson & Chantal Pohl Nielsen, *Golden Rice and the Looming GMO Debate: Implications for the Poor* 7-8 (Centre for Economic Policy Research, Discussion Paper No. 4195, 2004), <https://ssrn.com/abstract=508463> [<https://perma.cc/Q8NP-JDV4>].

114. See Richard T. Woodward and Richard C. Bishop, *How to Decide When Experts Disagree: Uncertainty-Based Choice Rules in Environmental Policy*, 73 *LAND ECON.* 492, 505 (1997).

115. See, e.g., Kenneth J. Arrow, *Some Ordinalist-Utilitarian Notes on Rawls' Theory of Justice*, 70 *J. PHIL.* 245 (1973); Harsanyi, *supra* note 66.

116. Musgrave, *supra* note 69, at 627.

This is a standard challenge, but it is wrong, because maximin does not assume infinite risk aversion.¹¹⁷ By stipulation, *we are dealing with situations in which probabilities cannot plausibly be assigned to various outcomes.*¹¹⁸ Perhaps that is rare in the regulatory context. But in principle, the objection that maximin assumes infinite risk aversion depends on a denial that uncertainty exists; it assumes that subjective choices will be made and that they will reveal subjective probabilities. It is true that subjective choices will be made. But such choices do not establish that objective uncertainty does not exist. To see why, it is necessary to engage that question directly.

C. Uncertainty Does Not Exist

Many economists have denied the existence of uncertainty.¹¹⁹ Milton Friedman, for example, writes of the risk-uncertainty distinction that “I have not referred to this distinction because I do not believe it is valid. I follow L.J. Savage in his view of *personal probability*, which denies any valid distinction along these lines. We may treat people as if they assigned numerical probabilities to every conceivable event.”¹²⁰ Friedman and other skeptics are correct to insist that people’s choices suggest that they assign probabilities to events. On a widespread view, an understanding of people’s choices can be taken as evidence of subjective probabilities. People’s decisions about whether to fly or instead to drive, whether to go to a store during a pandemic, whether to walk in certain neighborhoods at night, and whether to take risky jobs can be understood as an implicit assignment of probabilities to events. Indeed, regulators themselves make decisions, including decisions about climate change, from which subjective probabilities can be calculated.

But none of this makes for anything like a good objection to Knight, who was concerned with objective probabilities rather than subjective choices.¹²¹

117. See C.Y. Cyrus Chu & Wen-Fang Liu, *A Dynamic Characterization of Rawls’s Maximin Principle: Theory and Implications*, 12 CONST. POL. ECON. 255, 268 (2001).

118. See *id.* at 264-65.

119. For an account and a lament, see KAY & KING, *supra* note 13, at 106-54.

120. See MILTON FRIEDMAN, *PRICE THEORY* 282 (1976); see also JACK HIRSHLEIFER & JOHN G. RILEY, *THE ANALYTICS OF UNCERTAINTY AND INFORMATION* 10 (1992):

In this book we disregard Knight’s distinction, which has proved to be a sterile one. For our purposes risk and uncertainty mean the same thing. It does not matter, we contend, whether an ‘objective’ classification is or is not possible. For, we will be dealing throughout with a ‘subjective’ probability concept (as developed especially by Savage, 1954): probability is simply *degree of belief*. . . . [Because we never know true objective probabilities, d]ecision-makers are . . . never in Knight’s world of risk but instead always in his world of uncertainty. That the alternative approach, assigning probabilities on the basis of subjective degree of belief, is a workable and fruitful procedure will be shown constructively throughout this book.

For the purposes of the analysis by Hirschleifer and Riley, the assignment of subjective probabilities may well be the best approach. But the distinction between risk and uncertainty is not sterile when regulators are considering what to do but lack information about the probabilities associated with various outcomes.

121. See Stephen F. LeRoy & Larry D. Singell, Jr., *Knight on Risk and Uncertainty*, 95 J. POL. ECON. 394 (1987) (arguing that, against many critics, that Knight’s work supported the idea

Animals, no less than human beings, make choices from which subjective probabilities can be assigned. But the existence of subjective probabilities—from dogs, horses, and elephants—does not mean that animals do not ever face (objective) uncertainty.

Suppose that the question is the likelihood that at least one hundred million human beings will be alive in 10,000 years. For most people, equipped with the knowledge that they have, no probability can sensibly be assigned. Perhaps uncertainty is not unbounded; the likelihood can reasonably be described as above 0% and below 100%. (I think.) But beyond that point, there is little to say. Or suppose that I present you with an urn, containing 250 balls, and ask you to pick one; if you pick a blue ball, you receive \$1000, but if you pick a green ball, you have to pay me \$1000. Suppose that I refuse to disclose the proportion of blue and green balls in the urn—or suppose that the proportion has been determined by a computer, which has been programmed by someone that neither you nor I know. You can make a pick, but what does that tell us about actual probabilities? Regulators may be in a similar position at the early stage of a pandemic or when dealing with a new technology. These examples suggest that it is wrong to deny the possible existence of uncertainty, signaled by the absence of objective probabilities.¹²²

For Friedman and other skeptics about uncertainty, there is an additional problem. When necessary, human beings do assign subjective probabilities to future events. So what? The assignment can be a function of how the situation is described, and formally identical descriptions can produce radically different judgments. There is reason to believe, for example, that people will not give the same answer to the question, “What is the likelihood that 80% of people will suffer an adverse effect from a certain risk?” and to the question, “What is the likelihood that 20% of people will not suffer an adverse effect from a certain risk?”¹²³ The merely semantic reframing may well affect probability judgments.¹²⁴

In any case, probability judgments are notoriously unreliable because they are frequently based on heuristics and biases that lead to severe and systematic errors.¹²⁵ Suppose that subjective probability estimates are rooted in the

of subjective probabilities). For a vigorous and sustained argument on behalf of the pervasiveness of uncertainty, see KAY & KING, *supra* note 13, at 35–49. For a clear explanation of why uncertainty exists, see ELSTER, *supra* note 2, at 193–99, 199 (“One could certainly elicit from a political scientist the subjective probability that he attaches to the prediction that Norway in the year 3000 will be a democracy rather than a dictatorship, but would anyone even contemplate *acting* on the basis of this numerical magnitude?”).

122. See ELSTER, *supra* note 2, at 195–99.

123. See *id.*

124. *Id.*

125. For a good overview of this topic, see JONATHAN BARON, THINKING AND DECIDING 125–47 (3d ed. 2000). Elster briefly notes how this point relates to the debate over uncertainty: “There are too many well-known mechanisms that distort our judgment, from wishful thinking to rigid cognitive structures, for us to be able to attach much weight to the numerical magnitudes that can be elicited by the standard method of asking subjects to choose between hypothetical options.” ELSTER, *supra* note 2, at 199 (internal citations omitted).

availability heuristic, leading people to exaggerate risks for which examples readily come to mind (“availability bias”) and also to underestimate risks for which examples are cognitively unavailable (“unavailability bias”).¹²⁶ Why should regulators believe that subjective estimates, subject as they are to framing, heuristics, and biases, have any standing in the face of the objective difficulty or impossibility of making probability judgments?

Even if individuals and governments assign subjective probabilities, do their assignments bear on what ought to be done? As Elster puts it, speaking of scientists and bureaucrats: “There are too many well-known mechanisms that distort our judgment, from wishful thinking to rigid cognitive structures, for us to attach much weight to the numerical magnitudes that can be elicited by the standard method of asking subjects to choose between hypothetical options.”¹²⁷ Even if this account is too pessimistic (as I think it is), there are some problems for which merely subjective probabilities cannot plausibly be taken to show that we are operating in circumstances of risk rather than uncertainty. In any case, recall the benefits ranges reported above, in which officials declined to offer probability estimates, evidently on the ground that no adequate evidence was thought to support them.

Writing in 1937, Keynes, often taken to be a critic of the idea of uncertainty, clearly saw the distinction between objective probabilities and actual behavior: “The sense in which I am using the term [‘uncertain’ knowledge] is that in which the prospect of a European war is uncertain About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.”¹²⁸ This is so even if, as Keynes immediately added, we act “exactly as we should if we had behind us a good Benthamite calculation of a series of prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed.”¹²⁹ Even if subjective expected utilities can be assigned on the basis of behavior, regulators (like everyone else) may well be operating in circumstances of genuine uncertainty.

D. Uncertainty is Rare

Notwithstanding these points, regulatory problems do not typically involve genuine uncertainty. Using frequentist strategies, regulators are often able to assign probabilities to outcomes, and Bayesian approaches can also be used. When they cannot, perhaps they can instead assign probabilities to probabilities (or even, where this proves impossible, probabilities to

126. See Timur Kuran & Cass R. Sunstein, *Availability Cascades and Risk Regulation*, 51 STAN. L. REV. 683 (1999); Amos Tversky & Daniel Kahneman, *Judgment under Uncertainty: Heuristics and Biases*, in JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES 3, 11 (Daniel Kahneman ed., 1982).

127. See ELSTER, *supra* note 2, at 199.

128. JOHN MAYNARD KEYNES, A TREATISE ON PROBABILITY 214 (1921).

129. *Id.*

probabilities of probabilities). In many cases, regulators might be able to specify a range of probabilities saying, for example, that the probability of catastrophic outcomes from a pandemic or climate change is above 2% but below 30%.¹³⁰ At least some scientists and economists believe that climate change is not likely to create truly catastrophic harm, and that the real costs, human and economic, will be very high but not intolerable. In their view, the worst of the worst-case scenarios can be responsibly described as improbable.

Whatever we think of that example, perhaps we can agree that pure uncertainty is rare.¹³¹ Perhaps we can agree that at worst, regulatory problems involve problems of “bounded uncertainty,” in which we cannot assign probabilities within specified bands. It is possible to think, for example, that the risk of a catastrophic outcome is above 1% but below 10%, without being able to assign probabilities within that band. The pervasiveness of uncertainty depends on what is actually known. If uncertainty is rare, then Rawls’ argument, or variations on it, do not apply outside of exotic cases. Fair enough. But even if this is so, exotic cases may turn out to be important.

130. I am bracketing here frequentist claims about the pervasiveness of uncertainty. See KAY & KING, *supra* note 13, at 35-49. Even if we are frequentists, regulators are often dealing with repeated cases for which frequentist assignments of probability are perfectly feasible; consider food safety, occupational safety, and air pollution.

131. *But see* KAY & KING, *supra* note 13. In their provocative and spirited book, Kay and King are very hard on the idea of maximizing expected value, emphasizing that we often do not know enough to do anything like that. Instead of generating numbers, they urge that regulators, officials, and others should ask, “What is going on here?” *See id.* at 10. They also ask for close attention to “narratives.” *Id.* at 178-95. This is not the space to explore their analysis and their proposals, but in brief, the “What is going on here?” question cannot easily yield sensible answers. How can regulators possibly know how to handle (say) food safety, nanotechnology, genetic modification of food, if that is their question? The analysis is best disciplined at least through a rough sense of both probabilities and outcomes, which is often obtainable; lacking those, maximin is a candidate solution.

Revealingly, Kay and King defend the “What is going on here?” question in part by reference to President Barack Obama’s decision to kill Osama Bin Laden without knowing that Bin Laden was actually present in the relevant location. *Id.* at 8-9. In my view (and I was in the White House at the time, though not involved in any way with the decision), this is not a helpful example; it counts strongly against the central argument offered by Kay and King. Obama’s decision was Bayesian, and it involved a careful assessment of costs and benefits (and hence expected value). Roughly: The benefits of killing Bin Laden would be very high; the costs of failing would be high but manageable; the likelihood that he was present fell within an ascertainable range (in the vicinity of 50%, *id.* at 8); and importantly, the likelihood that he would be found, in the future, was relatively low. My experience is that public officials approach many non-repeatable events in this way, not by asking, “What is going on here?”

In addition, Kay & King rightly draw attention to the importance of resilience and robustness, as ways of handling uncertain risks. *Id.* at 423-25. (Consider the risks associated with climate change and pandemics.) But how resilient, and how robust? Resilience and robustness can be very costly indeed. We would want to spend infinite costs, today, to create resilience against a pandemic in a decade. Under conditions of risk, or of bounded uncertainty, calculation of expected value can be more than helpful, and for reasons discussed in text, maximin has its place. “Narrative,” by contrast, is not of much use. I am acutely aware that these are complex topics and that what I have said here is inadequate; it should be taken as a kind of a promissory note.

VI. A Path Forward

A great deal of work asks whether people really should follow maximin under circumstances of uncertainty.¹³² Some of this work draws on people's intuitions, in a way that illuminates actual beliefs but may tell us little about what rationality requires.¹³³ Other work is highly formal,¹³⁴ adopting certain axioms and seeing whether maximin violates them. The results of this work are not conclusive.¹³⁵ Certainly, maximin cannot be ruled out as a candidate for rational choice under uncertainty.

I will rest content with three general suggestions. *First:* As we have seen, the maximin rule is sometimes justified by standard cost-benefit analysis. If some potential outcomes are genuinely catastrophic and not highly improbable, eliminating them might be the approach that maximizes net benefits. Even if such outcomes are highly improbable (say, 1 in 100,000), the same conclusion might be the right one, if the expected value of precautions outweighs their expected costs.

Second: In the face of fat tails on the left-hand side (suggesting a higher-than-normal risk of catastrophe, as in "ruin problems"), there may be a good argument for the maximin rule, again depending on the numbers (and on what is known and what is unknown).

Third: Uncertainty is real; sometimes regulators lack information about probabilities. In deciding whether to follow the maximin rule under circumstances of Knightian uncertainty, or something close to it (such as bounded uncertainty), a great deal should turn on two questions: (a) How bad is the worst-case scenario, compared to other bad outcomes? (b) What, exactly, is lost by choosing the maximin rule? Of course, it is possible that choosers, including regulators, will lack the information that would enable them to answer these questions. But (and this is the central point) in the regulatory context, answers to both (a) and (b) may well be possible *even if it is not possible to assign probabilities to the various outcomes with any confidence*. By emphasizing the relative badness of the worst-case scenario, and the extent of the loss from attending to it, I am attempting to build on the Rawls/Gardiner suggestion that maximin is the preferred decision rule when little is lost from following it.

To see the relevance of the two questions, suppose that you are choosing between two options. The first has a best-case outcome of 10 and a worst-case outcome of -5. The second has a best-case outcome of 15 and a worst-case outcome of -6. It is impossible to assign probabilities to the various outcomes. Maximin would favor the first option, to avoid the worse worst-case (which is -

132. See, e.g., Arrow & Hurwicz, *supra* note 49 (suggesting the rationality of either maximin or maximax).

133. See Harsanyi, *supra* note 66.

134. See, e.g., LUCE & RAIFFA, *supra* note 18, at 286-97 (1957).

135. See *id.*

6); but to justify that choice, we have to know something about the *meaning* of the differences between 10 and 15 on the one hand and –5 and –6 on the other. If 15 is much better than 10, and if the difference between –5 and –6 is a matter of relative indifference, then the choice of the first option is hardly mandated. But if the difference between –5 and –6 greatly matters—if it is a matter of life and death—then the maximin rule is much more attractive.

Consider a regulatory analogue. Suppose that as compared with a ban, allowing automated vehicles would have a best-case outcome of \$2 billion in annual net benefits and a worst-case outcome of \$10 million in annual net losses. Suppose that we cannot assign probabilities to the various outcomes. Under the maximin rule, we should ban automated vehicles. But if the net loss of \$10 million is not a big deal, we might reject the maximin rule on something like the Rawls/Gardiner theory. Of course we could vary the numbers in such a way as to make the maximin rule much more attractive.

These points have the important implication of suggesting the possibility of a (rough) cost-benefit analysis of whether to follow the maximin rule under conditions of both risk and uncertainty. Sometimes the worst-case is the worst by far, and sometimes we lose relatively little by choosing the maximin rule. It is typically thought necessary to assign probabilities in order to engage in cost-benefit balancing; without an understanding of probabilities, such balancing might not seem able to get off the ground. But a crude version of cost-benefit balancing is possible even without reliable information about probability. For the balancing exercise to work, of course, it must be possible to produce cardinal rankings among the outcomes—that is, it must be possible to rank them not merely in terms of their badness but also in at least rough terms of how much worse each is than the less-bad others. That approach will not work if cardinal rankings are not feasible—as might be the case if (for example) it is not easy to compare the catastrophic loss from a pandemic with the loss from huge expenditures on efforts to control a pandemic. Much of the time, however, cardinal rankings *are* possible in the regulatory context.

Here is a simpler way to put the point. It is often assumed that in order to undertake cost-benefit analysis, it is necessary to assign probabilities, with the understanding that point estimates represent the average or most probable case. But in some cases, a sensible rule-of-thumb can be adopted without assigning probabilities. An understanding of the magnitude of the relevant payoffs can help regulators to navigate difficult situations. If one option has a large downside but no substantial upside, it can be rejected in favor of one that lacks that downside but that has a roughly equivalent upside.

To appreciate the need for some kind of analysis of the effects of following the maximin rule, imagine an individual or society lacking the information that would permit the assignment of probabilities to a series of hazards with catastrophic outcomes; suppose that the number of hazards is ten, or twenty, or a thousand. Suppose too that such an individual or society is able to assign probabilities (ranging from 1% to 90%) to an equivalent number of

other hazards, with outcomes that range from bad to extremely bad, but never catastrophic. Suppose, finally, that every one of these hazards can be eliminated at a cost—a cost that is high, but that does not, once incurred in individual cases, inflict harms that count as extremely bad or catastrophic. The maximin rule suggests that our individual or society should spend a great deal to eliminate each of the ten, or twenty, or thousand potentially catastrophic hazards. But once that amount is spent on even one of those hazards, there might be nothing left to combat the extremely bad hazards, even those with a 90% chance of occurring. We could even imagine that a poorly informed individual or society would be condemned to real poverty and distress, or even worse, merely by virtue of following maximin. In these circumstances, the maximin rule should be rejected.

This suggestion derives indirect support from the empirical finding that when asked to decide on the distribution of goods and services, most people reject the two most widely discussed principles in the philosophical literature: average utility, favored by Harsanyi, and Rawls' difference principle (allowing inequalities only if they work to the advantage to the least well-off).¹³⁶ Instead, people choose average utility with a floor constraint—that is, they favor an approach that maximizes overall well-being, but subject to the constraint that no member of society may fall below a decent minimum.¹³⁷ Insisting on an absolute welfare minimum to all, they maximize over that floor. Their aversion to especially bad outcomes leads them to a pragmatic threshold in the form of the floor. So too, very plausibly, in the context of precautions against risks. A sensible individual, or society, would not always choose maximin under circumstances of risk or uncertainty. Everything depends on what is lost, and what is gained, by eliminating the worst-case scenario; and much of the time, available information makes it possible to answer those questions at least in general terms.

If we apply these various points, we can easily imagine an amendment to OMB Circular A-4 that takes the following form¹³⁸:

In general, it is appropriate to focus on costs and benefits, calculated by reference to the expected value of various options. Thus, your analysis should include two fundamental components: a quantitative analysis characterizing the probabilities of the relevant outcomes and an assignment of economic value to the projected outcomes. It is essential that both parts be conceptually consistent. In particular, the quantitative analysis should be conducted in a way that permits it to be applied within the more general analytical framework of benefit-cost analysis. If one or another outcome is potentially catastrophic (a "worst case"), it might make sense to eliminate it, if the analysis shows that doing that maximizes net benefits. In considering potential catastrophe, you should

136. NORMAN FROHLICH & JOE A. OPPENHEIMER, CHOOSING JUSTICE: AN EXPERIMENTAL APPROACH TO ETHICAL THEORY (1992).

137. *Id.*

138. Significant parts of the italicized text are drawn from the current version of Circular A-4, *supra* note 3.

consider the possibility of “fat tails,” which arise when the probability of extreme negative outcomes is unusually high. Complex systems may be especially prone to fat tails.

In some cases, it may not be feasible to come up with probability distributions. If so, your analysis should be as complete as the available evidence permits. For example, it might include a specification of lower and upper bounds, with a qualitative analysis of their respective likelihoods (to the extent possible). In special circumstances, you might consider avoiding the worst-case scenario and thus following the maximin rule, which calls for eliminating the worst of the worst-cases. The strongest cases for following that rule would involve three factors: (1) Knightian uncertainty, understood as an inability to assign probabilities to various options; (2) catastrophic or grave consequences from one option, but not from other options; and (3) low or relatively low costs, or low or relatively low benefits foregone, as a result of choosing the option that avoids the worst-case scenario. Again in cases of uncertainty, more difficult cases, in which (for example) the costs of avoiding the worst-case scenario are very high, might also justify use of the maximin rule if (for example) the worst-case scenario is genuinely catastrophic.

My modest claim here is that for prudent regulators, attempting to proceed in the midst of important epistemic gaps, the maximin rule makes most sense when the worst-case scenario, under one course of action, is much worse than the worst-case scenario under the alternative course of action, when there are no huge disparities in gains from either option, and when the choice of maximin does not result in extremely significant losses. Variations on this basic case will present harder challenges, but in some situations, they too will allow room for maximin. At the same time, it is important for prudent regulators to focus as well on the best-case scenarios, which may promise miracles;¹³⁹ that possibility may provide an important cautionary note about efforts to eliminate risks, including those posed by new technologies.

139. See Rowell, *supra* note 14. Rowell’s illuminating discussion refers to “wonders.”