ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

MAXIMIN

Cass R. Sunstein

Discussion Paper No. 1020

11/2019

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center

The Social Science Research Network Electronic Paper Collection:
https://ssrn.com/abstract=3476250

Very preliminary draft 10/19/19
All rights reserved

Maximin

Cass R. Sunstein[*]

Abstract

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 Unites States, but in the context of climate change 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 a terrible 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) The probability distributions may include "fat tails," in which very bad outcomes are more probable than merely bad outcomes; it may make sense to eliminate those outcomes for that reason. (4) In circumstances of Knightian uncertainty, where observers (including regulators) cannot assign probabilities to imaginable outcomes, the maximin rule may make sense. (It may be possible to combine (3) and (4).) 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.

> "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 Knight[1]

> "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?"

[*] Robert Walmsley University Professor, Harvard University. I am grateful to Tyler Cowen, Annie Duke, Eric Posner for superb comments on an earlier draft and to Dinis Cheian for extraordinary research assistance.

[1] Frank H. Knight, Risk, Uncertainty, and Profit (1933).

Electronic copy available at: https://ssrn.com/abstract=3476250

- Jon Elster[2]

> "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-4[3]

I. In Brief

For regulators, what is the appropriate approach to worst-case scenarios? Suppose that genetically modified foods pose a risk of catastrophe – very small, but not zero.[4] Or suppose that some new 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]?

In answering these questions, I am going to be covering a great deal of ground, and while the journey is more important than the destination, it will be useful to set out the basic conclusions at the outset. The first three are straightforward. The remaining four are not.

(1) Regulators should generally focus on expected value and on likely costs and benefits, not on worst cases.[7] They should aim to come up with probability distributions, accompanied by point estimates.[8] When they cannot produce probability distributions, they should try to come up with reasonable ranges of both costs and benefits.

---

[2] *See* JON ELSTER, EXPLAINING TECHNICAL CHANGE: A CASE STUDY IN THE PHILOSOPHY OF SCIENCE 199 (1983).

[3] Available at https://www.transportation.gov/sites/dot.gov/files/docs/OMB%20Circular%20No.%20A-4_0.pdf.

[4] For one view, see Nassim Nicholas Taleb et al., The Precautionary Principle (with Application to the Genetic Modification of Organisms) (2014), available at http://www.fooledbyrandomness.com/pp2.pdf 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."

[5] https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/

[6] https://obamawhitehouse.archives.gov/omb/oira/social-cost-of-carbon

[7] 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).

[8] Point estimates, frequently provided by agencies, can often be understood as reflecting the mean of a probability distributions.

Electronic copy available at: https://ssrn.com/abstract=3476250

(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.[9]
(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.
(4) In some cases, a probability distribution might include "fat tails" on the left-hand side, in which very bad outcomes are more probable than merely bad outcomes; it might make sense to eliminate those very bad outcomes under conventional cost-benefit analysis.
(5) In some circumstances, often described as Knightian uncertainty, observers (including regulators) cannot assign probabilities to imaginable outcomes, and for that reason, the maximin rule is appealing. I will argue that contrary to a vigorously defended view in economics,[10] the problem of uncertainty is real.
(6) Some circumstances may combine (4) with (5), in the sense that fat tails and uncertainty are both present.
(7) With respect to (4), (5), and (6) the problems arise when eliminating dangers also threatens to impose very high costs or to eliminate very large potential gains. When this is so, the analysis is more tractable for (4) than for (5) and (6). There might be fat tails on the right-hand side, suggesting the possibility of wonders or miracles,[11] which might make human life immeasurably better, 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. To do that, they should calculate costs and benefits, with probability distributions as appropriate, and they should proceed if and only if the benefits justify the costs.[12] They should not focus solely or mostly on the worst cases. At the same time, a question remains: Are there any problems that the maximin rule can handle better than welfare maximization? The best answer points to cases of uncertainty, at least 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 weakens as the costs of eliminating the worst-case scenario rise, and as that scenario becomes decreasingly grave.[13]

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

---

[9] 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 Supreme Court Review 93. I bracket that issue here.
[10] An early account is Frank Ramsey, Truth and Probability, in The Foundations of Mathematics and Other Logical Essays, K. Paul, Trench, Trubner and Co., London, 1931.
[11] Arden Rowell, Regulating Best-Case Scenarios, 50 Env. L. (forthcoming 2020), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3157287
[12] 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), and there may be a legitimate role for risk aversion of certain kinds.
[13] 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.

Electronic copy available at: https://ssrn.com/abstract=3476250

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 percent. The likelihood of a significant increase in risk is in the vicinity of one percent, 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, parenthetically, that if you continue to exercise, there is small chance that you will get significant health benefits and thus *reduce* the risk of death.

Now suppose instead that in response to your request that she assign probabilities to the various outcomes, she says, "I can't do that! No doctor can. 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,[14] 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 regulation of new technologies, the same might be true. At least when some technology has a terrible or catastrophic worst-case scenario, the best course might be to avoid it.

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.[15] 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."[16] 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 respect to some problems, including those potentially raised by 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 its benefits justify its costs.[17] 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

[14] FRANK H. KNIGHT, RISK, UNCERTAINTY, AND PROFIT (1933); *see also* R. DUNCAN LUCE AND HOWARD RAIFFA, GAMES AND DECISIONS 275-86 (1957).
[15] Memorandum for the Heads of Executive Departments and Agencies (Mar. 11, 2011), https://obamawhitehouse.archives.gov/sites/default/files/omb/inforeg/for-agencies/Principles-for-Regulation-and-Oversight-of-Emerging-Technologies-new.pdf, included here as Appendix A.
[16] Id.
[17] *See* Executive Order 13563.

Electronic copy available at: https://ssrn.com/abstract=3476250

probability distributions to the extent feasible.[18] But what if it is not feasible to produce probability distributions?

To see the breadth of the problem, consider a few numbers from recent 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.[19] (2) The projected annual benefits of an air pollution rule governing particulate matter range from $3.6 billion to $9.1 billion.[20] (3) The projected benefits of a regulation governing hazardous air pollutants range from $28.1 billion to $76.9 billion.[21] (4) The projected benefits of a regulation governing cross-state air pollution range from $20.5 billion to $59.7 billion.[22]

It is worth pausing over three noteworthy features of those numbers. First, the government does not offer probability estimates to make sense of those ranges. It does not say that the probability at the low end is 1 percent, or 25 percent, or 50 percent. 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.[23] But even outside of that context, relatively standard regulations, not involving new technologies, often project wide ranges in terms of benefits, costs, or both.[24] In terms of monetized costs, the worst case may be *double* the best case.[25] In terms of monetized benefits, the best case may be *triple* the worst case.[26] For a more general

---

[18] *See* Appendix B. 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. . . ."

[19] https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2015_cb/2015-cost-benefit-report.pdf, at 25.

[20] https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2014_cb/2014-cost-benefit-report.pdf, at 25.

[21] https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2013_cb/2013_cost_benefit_report-updated.pdf, at 27.

[22] https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/2012_cb/2012_cost_benefit_report.pdf, at 26.

[23] *See* 2015 Report, supra note, at 13-18.

[24] *See id.* at 19.

[25] *See* the food safety rules noted at *id.*

[26] *See id.*

Electronic copy available at: https://ssrn.com/abstract=3476250

glimpse, consider this table, with particular reference to the wide benefits ranges[27]:

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 new technologies, the difference between the worst and the best case might be (much) bigger still.[28] 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,[29] 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 a fundamental mistake. 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 really are unable to assign probabilities to outcomes, and when some possible outcomes are catastrophic, the maximin rule may considerable appeal. Climate change is an obvious candidate for this conclusion,[30] and something similar might be said for some new or emerging risks, and others that are not even on the horizon.[31] 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.

---

[27] See 2014 Report, supra note.
[28] As an analogy, consider the social cost of carbon, see note supra, with a range, in 2020 dollars, from $12 to $123 per ton.
https://obamawhitehouse.archives.gov/sites/default/files/omb/inforeg/scc_tsd_final_clean_8_26_16.pdf
[29] Arden Rowell, Regulating Best-Case Scenarios, 50 Env. L. (forthcoming 2020), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3157287
[30] See Stephen M. Gardner, A Perfect Moral Storm: the Ethical Tragedy of Climate Change 411-14 (2011).
[31] 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. ECONS. & POL. 275 (2011); Martin L. Weitzman, *On Modeling and Interpreting the Economics of Catastrophic Climate Change*, 91 REV. ECONS. & STATS. 1 (2009).

Electronic copy available at: https://ssrn.com/abstract=3476250

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.

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. 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.

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). 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. 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.

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

Electronic copy available at: https://ssrn.com/abstract=3476250

almost certainly choose (a). No gambler will do well if she keeps choosing (b).[32] 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 percent of no gain (aside from keeping up with inflation) and a 50 percent chance of an annual gain of two percent. The second, called Risky, creates a 25 percent chance of an annual loss of five percent, a 25 percent chance of no gain (aside from keeping up with inflation), a 25 percent chance of a five percent gain, and a 25 percent chance of a ten percent 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 would a 25 percent loss matter to Smith? What would be the effect of a 25 percent gain? Perhaps a 25 percent loss would be devastating, given Smith's needs and wants, and perhaps a 25 percent 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. 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 this point in the context of climate change.[33] The problem of fat tails is not captured in Problem 4; fat tails consist of unusual probability distributions, when likelihoods tend to increase at the tails, or more particularly, when *the likelihood of terrible outcomes increases on the left-hand side*. Thus:

Problem 5.

Which would you prefer?

(a) A 99% chance of gaining $60, a .01% chance of losing $10, and a .09% chance of losing $100 million; or
(b) A 50% chance of gaining $10, and a 50% chance of losing $10.

---

[32] 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.)
[33] *See* note *supra*.

Electronic copy available at: https://ssrn.com/abstract=3476250

Problem 5 involves 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. Consider Weitzman's suggestion, focusing on climate change[34]:

> 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.

In this passage, Weitzman combines an emphasis on "the unknown unknowns," or uncertainty, with a reference to "the extremes of probability distributions."[35] 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)).

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, 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 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 is a possible example.[36] 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

---

[34] *Fat-Tailed Uncertainty*, supra note, at 275.
[35] *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."
[36] See note supra.

Electronic copy available at: https://ssrn.com/abstract=3476250

$100 million.

This is a problem of "moonshots" or "miracles," understood as low-probability chances of extraordinary returns.[37] 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.)

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[38]? When?

For certain regulatory problems, many people accept the Precautionary Principle.[39] The idea takes multiple forms, 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.[40] 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, 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."[41] 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."[42]

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

---

[37] Arden Rowell, Regulating Best-Case Scenarios, 50 Env. L. (forthcoming 2020), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3157287

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

[39] For general discussion, see CASS R. SUNSTEIN, LAWS OF FEAR (2006).

[40] NASSIM NICHOLAS TALEB ET AL., THE PRECAUTIONARY PRINCIPLE (WITH APPLICATION TO THE GENETIC MODIFICATION OF ORGANISMS) (2014), available at http://www.fooledbyrandomness.com/pp2.pdf. Taleb et al. focus on "propagating impacts resulting in irreversible and widespread damage." 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."

[41] Quoted in BJORN LOMBORG, THE SKEPTICAL ENVIRONMENTALIST 347 (2001).

[42] *See* http://www.monitor.net/rachel/r586.html.

Electronic copy available at: https://ssrn.com/abstract=3476250

example) economic, sectoral, or partisan institutional interests."[43] In certain forms, the principle might be taken to reflect the maximin principle: rule out 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*.[44] 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.[45]

Suppose, for example, that steps are taken to regulate or ban genetically modified food on precautionary grounds.[46] Many people believe that any such steps might well result in numerous deaths, and a small probability of many more.[47] The reason is that genetic modification holds out the promise of producing food that is both cheaper and healthier – resulting, for example, in "golden rice," which might have large benefits in developing countries.[48] The point is not that genetic modification will definitely have those benefits, or that the benefits of genetic modification outweigh the risks. The point is only that if the precautionary principle is taken in certain ways, it is offended by regulation as well as by nonregulation. To be sure, the maximin principle might prove helpful here – an issue to which I will return.

Or consider regulation of autonomous vehicles.[49] 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.[50] Use of autonomous vehicles might well increase safety, perhaps dramatically. We are dealing with safety-safety tradeoffs. The example shows again that 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

---

[43] *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).

[44] *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, at 650. The "general concerns" may be on both sides.

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

[46] *See* David Vogel, The Regulation of GMOs in Europe and the United States: A Case-Study of Contemporary European Regulatory Politics (Publication of the Study Group on Trade, Science and Genetically Modified Foods, 2001), available at http://www.cfr.org/pubs/Victor_ModFood_Paper2.html; Symposium, Are the US and Europe Heading for a Food Fight Over Genetically Modified Food? (2001), available at http://pewagbiotech.org/events/1024/; TonyGilland, Precaution, GM Crops, and Farmland Birds, in Rethinking Risk and the Precautionary Principle 84, 84-88 (Julian Morris ed. 2001).

[47] 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).

[48] *Id.*

[49] https://www.transportation.gov/av/3/preparing-future-transportation-automated-vehicles-3.

[50] https://www.zdnet.com/article/how-autonomous-vehicles-could-save-over-350k-lives-in-the-us-and-millions-worldwide/

Electronic copy available at: https://ssrn.com/abstract=3476250

billion, and that the monetized costs will have a range of $200 million to $1.6 billion. (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 $600 million and $3.6 billion.[51]) 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. Should we say that the agency should not proceed, because $1.6 billion is higher than $1.5 billion?

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. 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 his 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 and 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.[52] To be sure, regulators would want to be more disciplined about the 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 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."[53]

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 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

---

[51] See https://www.federalregister.gov/documents/2018/12/21/2018-27283/national-bioengineered-food-disclosure-standard
[52] *See id.*; Richard T. Woodward and Richard C. Bishop, *How to Decide When Experts Disagree: Uncertainty-Based Choice Rules in Environmental Policy*, 73 LAND ECONOMICS 492 (1997).
[53] *See* John C. Harsanyi, *Morality and the Theory of Rational Behavior*, *in* UTILITARIANISM AND BEYOND 40 (Amartya Sen & Bernard Williams eds., 1982).

Electronic copy available at: https://ssrn.com/abstract=3476250

scenario.[54] 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. 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.[55] 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 highly 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.[56] 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 <u>if</u> 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; it is far from irrational to create a margin of safety to protect against disaster. But if the likelihood of a bad outcome is extremely 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-

---

[54] *See* Cass R. Sunstein, *Probability Neglect: Emotions, Worst-cases, and the Law*, 112 YALE L.J. 61, 62-63 (2002).
[55] *See* TODD S. AAGAARD, A FUNCTIONAL APPROACH TO RISKS AND UNCERTAINTIES UNDER NEPA , 1 MICH. J. ENVTL. & ADMIN. L. 87 (2012). .
[56] *See* Richard A. Musgrave, *Maximin, Uncertainty, and the Leisure Trade-Off*, 88 Q. J. ECON. 625, 626-28 (1974).

Electronic copy available at: https://ssrn.com/abstract=3476250

4, finalized in 2003.[57] 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." 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."[58]

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, "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." Optimistically, Circular A-4 concludes: "You should make a special effort to portray the probabilistic results—in graphs and/or tables—clearly and meaningfully."[59]

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.[60] 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 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 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.[61] 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

---

[57] https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/circulars/A4/a-4.pdf. A useful primer can be found at https://www.reginfo.gov/public/jsp/Utilities/circular-a-4_regulatory-impact-analysis-a-primer.pdf
[58] Circular A-4, *supra* note.
[59] *Id.*
[60] *See* notes *supra*.
[61] *See* 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); Colin Camerer, *Individual Decision Making*, in THE HANDBOOK OF EXPERIMENTAL ECONOMICS, 587, 665–670 (John H. Kagel & Alvin E. Roth, eds., 1995).

Electronic copy available at: https://ssrn.com/abstract=3476250

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 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).[62]

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 the introduction of new technologies. 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.[63] 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.[64] 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 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.

---

[62] https://wjh-www.harvard.edu/~dtg/Kermer%20et%20al%202006.pdf.

[63] See William Samuelson and Richard Zeckhauser, Status Quo Bias in Decision Making, 1 J Risk and Uncertainty 7 (1988).

[64] *See* Paul SLOVIC, THE PERCEPTION OF RISK, 140–143 (2000).

Electronic copy available at: https://ssrn.com/abstract=3476250

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 in 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 percent said "yes." When I changed the numbers to 20/30 (for another group), a strong majority (74 percent) 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, with or without uncertainty. 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

In some contexts, risk-related problems involve hazards of ascertainable probability.[65] 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 easily specify even a range of probability.[66] Hence, regulators, and ordinary people, are sometimes acting in a situation of 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).[67] 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.[68]

A. Strategies of Avoidance

[65] 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.
[66] FRANK H. KNIGHT, RISK, UNCERTAINTY, AND PROFIT (1933); KIYOHIKO G. NISHIMURA AND HIROYUKI OZAKI, ECONOMICS OF PESSIMISM AND OPTIMISM: THEORY OF KNIGHTIAN UNCERTAINTY AND ITS APPLICATIONS (2017).
[67] *See id.*; Paul Davidson, *Is Probability Theory Relevant for Uncertainty? A Post-Keynesian Perspective*, 5(1) J. ECON. PERSP. 129 (1991).
[68] On ignorance and precaution, see Poul Harremoes*, Ethical Aspects of Scientific Incertitude in Environmental Analysis and Decision Making*, 11 JOURNAL OF CLEANER PRODUCTION 705 (2003).

Electronic copy available at: https://ssrn.com/abstract=3476250

Of course, it is also true that over time, problems that seem to 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."[69] But as the circular notes, " Delaying a decision will also have costs, as will further efforts at data gathering and analysis."[70] Delay of regulation may mean serious harm (including large numbers of deaths). 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).[71] 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.[72] There is some evidence that people follow that principle, at least in surveys.[73] 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.

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.*[74] In the context of regulation of new technologies, 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 risks

---

[69] See note supra.
[70] Id.
[71] See Cass R. Sunstein, The Limits of Quantification, 102 Cal. L. Rev. 1369 (2014).
[72] *See* Rawls, supra note, at 146 ("When we have no evidence at all, the possible cases are stipulated to be equally probable"); DUNCAN LUCE & HOWARD RAIFFA, GAMES AND DECISIONS 284 (1957).
[73] See Sunstein, supra note.
[74] For a technical treatment of the possible rationality of maximin, see Kenneth Arrow and Leonid Hurwicz, *An Optimality Criterion for Decision-Making Under Ignorance*, *in* UNCERTAINTY AND EXPECTATIONS IN ECONOMICS: ESSAYS IN HONOR OF G.L.S. SHACKLE (1972); for a non-technical overview, see JON ELSTER, EXPLAINING TECHNICAL CHANGE 185–207 (1983).

Electronic copy available at: https://ssrn.com/abstract=3476250

associated with genetically modified food,[75] nuclear energy,[76] 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.[77] 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."[78] As it puts it, "this unusual rule" is plausible in light of "three chief features of situations."[79] 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."[80] 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."[81] That means that the rule does not "generally apply, nor of course is it self-evident."[82] 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."[83]

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"[84] – which helps to justify his principles of justice. It is 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.[85] Who wants that? In such cases, applying maximin seems quite rational.

---

[75] NASSIM NICHOLAS TALEB ET AL., THE PRECAUTIONARY PRINCIPLE (WITH APPLICATION TO THE GENETIC MODIFICATION OF ORGANISMS) (2014), available at http://www.fooledbyrandomness.com/pp2.pdf.
[76] *See* JON ELSTER, EXPLAINING TECHNICAL CHANGE 188–205 (1979).
[77] *See, e.g.*, JON ELSTER, EXPLAINING TECHNICAL CHANGE 188-205 (1983).
[78] *See* JOHN RAWLS, A THEORY OF JUSTICE 132-39 (revised ed. 1999). Rawls draws on but adapts William Fellner, Probability and Profit 140-42 (1965).
[79] Rawls*, supra note,* at 134.
[80] *Id.*
[81] *Id.*
[82] 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.)
[83] *Id.*
[84] *Id.* (Note: This is only in the original, again.)
[85] I am grateful to Annie Duke for this point.

Electronic copy available at: https://ssrn.com/abstract=3476250

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.[86] To make the underlying intuition clear, Gardiner begins with the problem of choosing between two options, A and B[87]:

> 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).

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 percent 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."[88]

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."[89] 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[90]:

> 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.

---

[86] *See* Stephen Gardiner, *The Core Precautionary Principle*, 14 J. POLIT. PHIL. 33 (2006).
[87] *Id.* at 46.
[88] *Id.* at 49.
[89] *Id.* at 51.
[90] *Id.* at 55.

Electronic copy available at: https://ssrn.com/abstract=3476250

Gardiner adds, sensibly, that to justify maximin, the threats that are potentially catastrophic must satisfy some minimal threshold of plausibility.[91] 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.[92] 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.[93] A related argument, ventured by Nassim Nicholas Taleb et al. in an illuminating discussion of the precautionary principle, is that genetically modified crops pose a "ruin" problem, involving a low probability of catastrophically high costs.[94] 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 create ruin problems. Perhaps they do, but it is certainly 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 or emerging technologies.

V. Four Objections

A. Triviality

An evident problem with this argument is that it risks triviality, above all because of condition (3).[95] 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

---

[91] *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 some probabilities, at least. Gardiner's argument must be that in some cases, we might know that the likelihood that a bad outcome would occur really is trivial.

[92] *See id.* at 55.

[93] *See* JON ELSTER, EXPLAINING TECHNICAL CHANGE at 203.

[94] NASSIM NICHOLAS TALEB ET AL., THE PRECAUTIONARY PRINCIPLE (WITH APPLICATION TO THE GENETIC MODIFICATION OF ORGANISMS) (2014), available at http://www.fooledbyrandomness.com/pp2.pdf.

[95] *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.

Electronic copy available at: https://ssrn.com/abstract=3476250

greenhouse gas emissions.[96] 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 lives under circumstances of extreme deprivation.[97]

The real question, then, is whether regulators should embrace maximin in real-world cases in which doing so is extremely costly. If they should, it is because condition (3) is too stringent and should be abandoned. Even if the costs of following maximin 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.[98] 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.[99] 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."[100] 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 the Catastrophic Harm Precautionary Principle. But even under circumstances of uncertainty—the argument goes—maximin is senseless unless societies are to show infinite risk aversion.

This is a standard challenge, but it is wrong, because maximin does not assume infinite risk aversion.[101] By stipulation, we are dealing with situations in which probabilities cannot plausibly be assigned to various outcomes.[102] 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.

B. Uncertainty Does Not Exist

[96] *See* WILLIAM D. NORDHAUS & JOSEPH BOYER, WARMING THE WORLD: ECONOMIC MODELS OF CLIMATE CHANGE 168 (2000).
[97] *See* Kym Anderson and 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), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=508463.
[98] *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).
[99] *See, e.g.*, Kenneth J. Arrow, *Some Ordinalist-Utilitarian Notes on Rawls' Theory of Justice*, 70 J. PHIL. 245 (1973); J.C. Harsanyi, *Can the Maximin Principe Serve As a Basis for Morality? A Critique of John Rawls' Theory*, 69 AM. POL. SCI. REV. 594 (1975).
[100] Musgrave, supra note, at 627.
[101] *See* C.Y. Cyrus Chu and Wen-Fang Liu, *A Dynamic Characterization of Rawls's Maximin Principle: Theory and Implications*, 12 CONST. POL. ECON. 255, 268 (2001).
[102] *See id.* at 264-65.

Electronic copy available at: https://ssrn.com/abstract=3476250

Many economists have denied the existence of uncertainty. 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."[103] 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 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 a good objection to Knight, who was concerned with objective probabilities rather than subjective choices.[104] 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 genuine 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%. 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. These examples suggest that it is wrong to deny the possible existence of uncertainty, signaled by the absence of objective probabilities.[105]

For Friedman and other skeptics about uncertainty, there is an additional problem. When necessary, human beings do assign subjective probabilities to future events. But the assignment is a function of how the situation is described, and formally identical descriptions can produce radically different judgments. There is every 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

---

[103] *See* MILTON FRIEDMAN, PRICE THEORY 282 (1976); *see also* JACK HIRSHLEIFER AND 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 Hirshleifer 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.

[104] *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 of subjective probabilities). For a clear explanation of why uncertainty exists, see JON ELSTER, EXPLAINING TECHNICAL CHANGE: A CASE STUDY IN THE PHILOSOPHY OF SCIENCE 193–99, 199 (1983) ("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?").

[105] *See* ELSTER, *supra* note, at 195–99.

Electronic copy available at: https://ssrn.com/abstract=3476250

effect from a certain risk?"[106] The merely semantic reframing will almost certainly affect probability judgments.[107] In any case, probability judgments are notoriously unreliable because they are frequently based on heuristics and biases that lead to severe and systematic errors.[108] Suppose that subjective probability estimates are rooted in the 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 cognitive unavailable ("unavailability bias").[109] Why should regulators believe that subjective estimates, subject as they are 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."[110] Even if this account is too pessimistic, there are some problems for which subjective probabilities cannot plausibly 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."[111] 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."[112] 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

Perhaps regulatory problems rarely involve genuine uncertainty. Perhaps regulators are usually able to assign probabilities to outcomes; and where they cannot, perhaps they can instead assign probabilities to probabilities (or even, where this proves impossible, probabilities to 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 climate change is above 2% but below 30%. Many scientists and economists believe that climate change is not likely to create catastrophic harm, and that the real costs, human and economic, will be high but not intolerable. In their view, the

---

106 *See id.*
107 *Id.*
108 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, at 199 (internal citations omitted).
109 *See* Amos Tversky and Daniel Kahneman, *Judgment under Uncertainty: Heuristics and Biases*, *in* JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES 3, 11 (Daniel Kahneman ed., 1982); Timur Kuran and Cass R. Sunstein, *Availability Cascades and Risk Regulation*, 51 STAN. L. REV. 683 (1999).
110 *See* JON ELSTER, EXPLAINING TECHNICAL CHANGE 199 (1983).
111 JOHN MAYNARD KEYNES, A TREATISE ON PROBABILITY 214 (1921).
112 *Id.*

Electronic copy available at: https://ssrn.com/abstract=3476250

worst-case scenarios can be responsibly described as improbable.

Perhaps we can agree that pure uncertainty is rare. 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.

VI. A Path Forward

A great deal of work explores the question whether people should follow maximin under circumstances of uncertainty.[113] 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.[114] Other work is highly formal,[115] adopting certain axioms and seeing whether maximin violates them. The results of this work are not conclusive.[116] Certainly, maximin cannot be been ruled out as a candidate for rational choice under uncertainty.

I will rest content with a general suggestion. In deciding whether to follow the maximin rule in the regulatory context, 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 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; 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.

These points have the important implication of suggesting the possibility of a (rough) cost-benefit analysis of maximin 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

---

[113] *See, e.g.*, Kenneth J. Arrow and Leonid Hurwicz, *An Optimality Criterion for Decision-Making Under Ignorance*, *in* UNCERTAINTY AND EXPECTATIONS IN ECONOMICS (C.F. Carter & J.L. Ford eds., 1972) (suggesting the rationality of either maximin or maximax).

[114] *See* Harsanyi, *supra* note.

[115] *See, e.g.*, R. DUNCAN LUCE AND HOWARD RAIFFA, GAMES AND DECISIONS: INTRODUCTION AND CRITICAL SURVEY 286–97 (1957).

[116] *See id.*

Electronic copy available at: https://ssrn.com/abstract=3476250

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 climate change with the loss from huge expenditures on reductions of greenhouse gas emissions. 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 an 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 a 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 hundred 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, maximin 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).[117] 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.[118] 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 that time, available information makes it possible to answer those questions at least in general terms.

My goal here has not been to suggest an amendment to Circular A-4, but we could easily imagine one of the following form[119]:

---

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

[118] *Id.*

[119] The last two sentences of the first paragraph are largely drawn from the current version of Circular A-4. See Appendix B.

Electronic copy available at: https://ssrn.com/abstract=3476250

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.

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. In special circumstances, you might consider avoiding the worst-case scenario and thus following the maximin rule. The strongest cases for following that rule would involve three factors: (1) uncertainty, understood as an inability to assign probabilities to various options; (2) catastrophic or grave consequences from one option, but not from other options; (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.

Nothing here is meant as a proof that maximin is forbidden, or even not required, by rationality.[120] My claim is instead 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. At the same time, it is important for prudent regulators to focus as well on the best-case scenarios, which may promise miracles[121]; that possibility may provide an important cautionary note about efforts to eliminate risks.

[120] *See* Luce & Raiffa, *supra* note, at 286-97.
[121] See Arden Rowell, Regulating Best-Case Scenarios, 50 Env. L. (forthcoming 2020), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3157287. Rowell's illuminating discussion refers to "wonders."

Electronic copy available at: https://ssrn.com/abstract=3476250

Appendix A

Electronic copy available at: https://ssrn.com/abstract=3476250

March 11, 2011

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM: John P. Holdren
Assistant to the President for Science and Technology
Director, Office of Science andTechnology Policy

Cass R. Sunstein
Administrator, Office of Information and Regulatory Affairs
Office of Management and Budget

Islam A. Siddiqui
Chief Agricultural
Negotiator
United States Trade Representative

SUBJECT: Principles for Regulation and Oversight of Emerging Technologies

Innovation with respect to emerging technologies -- such as nanotechnology, synthetic biology, and genetic engineering, among others -- requires not only coordinated research and development but also appropriate and balanced oversight. The White House Emerging Technologies Interagency Policy Coordination Committee (ETIPC) has developed the following broad principles, consistent with Executive Order 13563, to guide the development and implementation of policies for oversight of emerging technologies at the agency level.

We share a fundamental desire for regulation and oversight that ensure the fulfillment of legitimate objectives such as the protection of safety, health, and the environment. Regulation and oversight should avoid unjustifiably inhibiting innovation, stigmatizing new technologies, or creating trade barriers.

To advance these goals, the following principles, consistent with Executive Order 13563 and discussed and approved by the ETIPC, should be respected to the extent permitted by law:

Scientific Integrity: Federal regulation and oversight of emerging technologies should be based on the best available scientific evidence. Adequate information should be sought and developed, and new knowledge should be taken into account when it becomes available. To the extent feasible, purely scientific judgments should be separated from judgments of policy.

Public Participation: To the extent feasible and subject to valid constraints (involving, for example, national security and confidential business information),

Electronic copy available at: https://ssrn.com/abstract=3476250

relevant information should be developed with ample opportunities for stakeholder involvement and public participation. Public participation is important for promoting accountability, for improving decisions, for increasing trust, and for ensuring that officials have access to widely dispersed information.

Communication: The Federal Government should actively communicate information to the public regarding the potential benefits and risks associated with new technologies.

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.

Flexibility: To the extent practicable, Federal regulation and oversight should provide sufficient flexibility to accommodate new evidence and learning and to take into account the evolving nature of information related to emerging technologies and their applications.

Risk Assessment and Risk Management: Risk assessment should be distinguished from risk management. The Federal Government should strive to reach an appropriate level of consistency in risk assessment and risk management across various agencies and offices and across various technologies. Federally mandated risk management actions should be appropriate to, and commensurate with, the degree of risk identified in an assessment.

Coordination: Federal agencies should seek to coordinate with one another, with state authorities, and with stakeholders to address the breadth of issues, including health and safety, economic, environmental, and ethical issues (where applicable) associated with the commercialization of an emerging technology, in an effort to craft a coherent approach. There should be a clear recognition of the statutory limitations of each Federal and state agency and an effort to defer to appropriate entities when attempting to address the breadth of issues.

International Cooperation: The Federal Government should encourage coordinated and collaborative research across the international community. It should clearly communicate the regulatory approaches and understanding of the United States to other nations. It should promote informed choices and both sharing and development of relevant data, particularly with respect to the benefits and costs of regulation and oversight. The Federal Government should participate in the development of international standards, consistent with U.S. law and guidance (e.g., the National Technology Transfer and Advancement Act and OMB Circular A-119). When appropriate, international approaches should be coordinated as far in advance as possible, to help ensure that such approaches are consistent with these principles.

Electronic copy available at: https://ssrn.com/abstract=3476250

Regulation: The Federal Government should adhere to Executive Order 13563 and, consistent with that Executive Order, the following principles, to the extent permitted by law, when regulating emerging technologies:

- Decisions should be based on the best reasonably obtainable scientific, technical, economic, and other information, within the boundaries of the authorities and mandates of each agency;

- Regulations should be developed with a firm commitment to fair notice and to public participation;

- 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);

- Where possible, regulatory approaches should promote innovation while also advancing regulatory objectives, such as protection of health, the environment, and safety;

- When no significant oversight issue based on a sufficiently distinguishing attribute of the technology or the relevant application can be identified, agencies should consider the option not to regulate;

- Where possible, regulatory approaches should be performance-based and provide predictability and flexibility in the face of fresh evidence and evolving information; and

- Regulatory approaches shall comply with established requirements and guidance such as the following:

    - Executive Order 13563 – Improving Regulation and Regulatory Review. Federal Register, Vol. 76, No. 14, Friday, January 21, 2011, 3821-3823, available at http://www.gpo.gov/fdsys/pkg/FR-2011-01-21/pdf/2011- 1385.pdf;

    - Executive Order 12866 – Regulatory Planning and Review. Federal Register Vol. 58, No. 190, Monday, October 4, 1993, 51735-51744, available at http://www.whitehouse.gov/omb/inforeg/eol2866.pdf;

    - Information Quality Act (Sec. 515 of the Treasury and General Government Appropriations Act for FY 2001, Pub. L. No. 106-554); Information Quality Guidelines: OMB (2002) Guidelines for Ensuring and Maximizing the Quality, Objectivity, Utility, and Integrity of Information Disseminated by Federal Agencies (2002), 67 Fed. Reg.

Electronic copy available at: https://ssrn.com/abstract=3476250

8452 (Feb. 22, 2002), available at
http://www.whitehouse.gov/omb/fedreg/reproducible2.pdf;

- National Technology Transfer and Advancement Act of 1995 ("NTTAA"). Public Law 104-113, available at http://standards.gov/standards_gov/nttaa.cfm;

- Office of Management and Budget (OMB) Circular A-119, Transmittal Memorandum, Federal Participation in the Development and Use of Voluntary Standards (02/10/1998), available at http://www.whitehouse.gov/omb/circulars/al 19/al 19.html;

- OMB Final Information Quality Bulletin for Peer Review (December 16, 2004), available at http://www.whitehouse.gov/omb/memoranda/fy2005/m05-03.pdf;

- OMB Bulletin No. 07-02 (M-07-07), Issuance of OMB's "Final Bulletin for Agency Good Guidance Practices" (January 18, 2007), available at http://www.whitehouse.gov/omb/memoranda/fy2007/m07-07.pdf

- OMB/OSTP Memorandum: M-07-24, Updated Principles for Risk Analysis (September 19, 2007), available at http://www.whitehouse.gov/omb/memoranda/fy2007/m07-24.pdf;

- The Trade Agreements Act of 1979, as amended (Pub.L. 96-39, 93 Stat. 144, enacted July 26, 1979, codified at 19 U.S.C. ch.13 (19 U.S.C. § 2501-2581);

- "A Strategy for American Innovation: Driving Towards Sustainable Growth and Quality Jobs" (September 2009), available at: http://www.whitehouse. gov/assets/documents/SEPT_20__Innovation Whitepaper FINAL.pdf; and

- Office of Information and Regulatory Affairs, Disclosure and Information As Regulatory Tools (June 18, 2010), available at http://www.whitehouse.gov/sites/default/files/omb/assets/inforeg/discl osur e principles.pdf

Electronic copy available at: https://ssrn.com/abstract=3476250

Appendix B

Electronic copy available at: https://ssrn.com/abstract=3476250

Circular A-4

September 17, 2003

. . .

Treatment of Uncertainty

The precise consequences (benefits and costs) of regulatory options are not always known for certain, but the probability of their occurrence can often be developed. The important uncertainties connected with your regulatory decisions need to be analyzed and presented as part of the overall regulatory analysis. You should begin your analysis of uncertainty at the earliest possible stage in developing your analysis. You should consider both the statistical variability of key elements underlying the estimates of benefits and costs (for example, the expected change in the distribution of automobile accidents that might result from a change in automobile safety standards) and the incomplete knowledge about the relevant relationships (for example, the uncertain knowledge of how some economic activities might affect future climate change).[25] By assessing the sources of uncertainty and the way in which benefit and cost estimates may be affected under plausible assumptions, you can shape your analysis to inform decision makers and the public about the effects and the uncertainties of alternative regulatory actions.

---

25 In some contexts, the word "variability" is used as a synonym for statistical variation that can be described by a theoretically valid distribution function, whereas "uncertainty" refers to a more fundamental lack of knowledge. Throughout this discussion, we use the term "uncertainty" to refer to both concepts.

Electronic copy available at: https://ssrn.com/abstract=3476250

The treatment of uncertainty must be guided by the same principles of full disclosure and transparency that apply to other elements of your regulatory analysis. Your analysis should be credible, objective, realistic, and scientifically balanced.[26] Any data and models that you use to analyze uncertainty should be fully identified. You should also discuss the quality of the available data used. Inferences and assumptions used in your analysis should be identified, and your analytical choices should be explicitly evaluated and adequately justified. In your presentation, you should delineate the strengths of your analysis along with any uncertainties about its conclusions. Your presentation should also explain how your analytical choices have affected your results.

In some cases, the level of scientific uncertainty may be so large that you can only present discrete alternative scenarios without assessing the relative likelihood of each scenario quantitatively. For instance, in assessing the potential outcomes of an environmental effect, there may be a limited number of scientific studies with strongly divergent results. In such cases, you might present results from a range of plausible scenarios, together with any available information that might help in qualitatively determining which scenario is most likely to occur.

When uncertainty has significant effects on the final conclusion about net benefits, your agency should consider additional research prior to rulemaking. The costs of being wrong may outweigh the benefits of a faster decision. This is true especially for cases with irreversible or large upfront investments. If your agency decides to proceed with rulemaking, you should explain why the costs of developing additional information—including any harm from delay in public protection—exceed the value of that information.

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.[27] Delaying a decision will also have costs, as will further efforts at data gathering and analysis. You will need to weigh the benefits of delay against these costs in making your decision. Formal tools for assessing the value of additional information are now well developed in the applied decision sciences and can be used to help resolve this type of complex regulatory question.

"Real options" methods have also formalized the valuation of the added flexibility inherent in delaying a decision. As long as taking time will lower uncertainty, either passively or actively through an investment in information gathering, and some costs are irreversible, such as the potential costs of a sunk investment, a benefit can be assigned to the option to delay a decision. That benefit should be considered a cost of taking immediate action versus the alternative of delaying that action pending more information. However, the burdens of delay—including any harm to public health, safety, and the environment—need to be analyzed carefully.

1. Quantitative Analysis of Uncertainty

[26] When disseminating information, agencies should follow their own information quality guidelines, issued in conformance with the OMB government-wide guidelines (67 FR 8452, February 22, 2002).

[27] Clemen RT (1996), *Making Hard Decisions: An Introduction to Decision Analysis*, second edition, Duxbury Press, Pacific Grove.

Electronic copy available at: https://ssrn.com/abstract=3476250

Examples of quantitative analysis, broadly defined, would include formal estimates of the probabilities of environmental damage to soil or water, the possible loss of habitat, or risks to endangered species as well as probabilities of harm to human health and safety. There are also uncertainties associated with estimates of economic benefits and costs, such as the cost savings associated with increased energy efficiency. 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 a more general analytical framework, such as benefit-cost analysis. Similarly, the general framework needs to be flexible enough to incorporate the quantitative analysis without oversimplifying the results. For example, you should address explicitly the implications for benefits and costs of any probability distributions developed in your analysis.

As with other elements of regulatory analysis, you will need to balance thoroughness with the practical limits on your analytical capabilities. Your analysis does not have to be exhaustive, nor is it necessary to evaluate each alternative at every step. Attention should be devoted to first resolving or studying the uncertainties that have the largest potential effect on decision making. Many times these will be the largest sources of uncertainties. In the absence of adequate data, you will need to make assumptions. These should be clearly identified and consistent with the relevant science. Your analysis should provide sufficient information for decision makers to grasp the degree of scientific uncertainty and the robustness of estimated probabilities, benefits, and costs to changes in key assumptions.

For major rules involving annual economic effects of $1 billion or more, you should present a formal quantitative analysis of the relevant uncertainties about benefits and costs. In other words, you should try to provide some estimate of the probability distribution of regulatory benefits and costs. In summarizing the probability distributions, you should provide some estimates of the central tendency (e.g., mean and median) along with any other information you think will be useful such as ranges, variances, specified low-end and high-end percentile estimates, and other characteristics of the distribution.

Your estimates cannot be more precise than their most uncertain component. Thus, your analysis should report estimates in a way that reflects the degree of uncertainty and not create a false sense of precision. Worst-case or conservative analyses are not usually adequate because they do not convey the complete probability distribution of outcomes, and they do not permit calculation of an expected value of net benefits. In many health and safety rules, economists conducting benefit-cost analyses must rely on formal risk assessments that address a variety of risk management questions such as the baseline risk for the affected population, the safe level of exposure or, the amount of risk to be reduced by various interventions. Because the answers to some of these questions are directly used in benefits analyses, the risk assessment methodology must allow for the determination of expected benefits in order to be comparable to expected costs. This means that conservative assumptions and defaults (whether motivated by science policy or by precautionary instincts), will be incompatible with benefit analyses as they will result in benefit estimates that exceed the expected value. Whenever it is possible to characterize quantitatively the probability distributions, some estimates of expected value (e.g., mean and

Electronic copy available at: https://ssrn.com/abstract=3476250

median) must be provided in addition to ranges, variances, specified low-end and high-end percentile estimates, and other characteristics of the distribution.

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. You may consider the following analytical approaches that entail increasing levels of complexity:

- Disclose qualitatively the main uncertainties in each important input to the calculation of benefits and costs. These disclosures should address the uncertainties in the data as well as in the analytical results. However, major rules above the $1 billion annual threshold require a formal treatment.
- Use a numerical sensitivity analysis to examine how the results of your analysis vary with plausible changes in assumptions, choices of input data, and alternative analytical approaches. Sensitivity analysis is especially valuable when the information is lacking to carry out a formal probabilistic simulation. Sensitivity analysis can be used to find "switch points" -- critical parameter values at which estimated net benefits change sign or the low cost alternative switches. Sensitivity analysis usually proceeds by changing one variable or assumption at a time, but it can also be done by varying a combination of variables simultaneously to learn more about the robustness of your results to widespread changes. Again, however, major rules above the $1 billion annual threshold require a formal treatment.
- Apply a formal probabilistic analysis of the relevant uncertainties B possibly using simulation models and/or expert judgment as revealed, for example, through Delphi methods.[28] Such a formal analytical approach is appropriate for complex rules where there are large, multiple uncertainties whose analysis raises technical challenges, or where the effects cascade; it is required for rules that exceed the $1 billion annual threshold. For example, in the analysis of regulations addressing air pollution, there is uncertainty about the effects of the rule on future emissions, uncertainty about how the change in emissions will affect air quality, uncertainty about how changes in air quality will affect health, and finally uncertainty about the economic and social value of the change in health outcomes. In formal probabilistic assessments, expert solicitation is a useful way to fill key gaps in your ability to assess uncertainty.[29] 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. You should

[28] The purpose of Delphi methods is to generate suitable information for decision making by eliciting expect judgment. The elicitation is conducted through a survey process which eliminates the interactions between experts. See Morgan MG and Henrion M (1990), *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Riskand Policy Analysis*, Cambridge University Press.
[29] Cooke RM (1991), *Experts in Uncertainty: Opinion and Subjective Probability in Science*, Oxford University Press.

Electronic copy available at: https://ssrn.com/abstract=3476250

pay attention to correlated inputs. Often times, the standard defaults in Monte Carlo and other similar simulation packages assume independence across distributions. Failing to correctly account for correlated distributions of inputs can cause the resultant output uncertainty intervals to be too large, although in many cases the overall effect is ambiguous. You should make a special effort to portray the probabilistic results—in graphs and/or tables—clearly and meaningfully.

New methods may become available in the future. This document is not intended to discourage or inhibit their use, but rather to encourage and stimulate their development.

2. Economic Values of Uncertain Outcomes

In developing benefit and cost estimates, you may find that there are probability distributions of values as well for each of the outcomes. Where this is the case, you will need to combine these probability distributions to provide estimated benefits and costs.

Where there is a distribution of outcomes, you will often find it useful to emphasize summary statistics or figures that can be readily understood and compared to achieve the broadest public understanding of your findings. It is a common practice to compare the "best estimates" of both benefits and costs with those of competing alternatives. These "best estimates" are usually the average or the expected value of benefits and costs. Emphasis on these expected values is appropriate as long as society is "risk neutral" with respect to the regulatory alternatives. While this may not always be the case, you should in general assume "risk neutrality" in your analysis. If you adopt a different assumption on risk preference, you should explain your reasons for doing so.

3. Alternative Assumptions

If benefit or cost estimates depend heavily on certain assumptions, you should make those assumptions explicit and carry out sensitivity analyses using plausible alternative assumptions. If the value of net benefits changes from positive to negative (or vice versa) or if the relative ranking of regulatory options changes with alternative plausible assumptions, you should conduct further analysis to determine which of the alternative assumptions is more appropriate. Because different estimation methods may have hidden assumptions, you should analyze estimation methods carefully to make any hidden assumptions explicit.

. . .

Electronic copy available at: https://ssrn.com/abstract=3476250

Electronic copy available at: https://ssrn.com/abstract=3476250