

Beyond Cheneyism and Snowdenism

Cass R. Sunstein†

In the domain of national security, many people favor some kind of precautionary principle, insisting that it is far better to be safe than sorry and contending that a range of important safeguards, including widespread surveillance, is amply justified to prevent the loss of life. Those who object to the resulting initiatives, and in particular to widespread surveillance, respond with a precautionary principle of their own, seeking safeguards against what they see as unacceptable risks to privacy and liberty. The problem is that, as in the environmental context, a precautionary principle threatens to create an unduly narrow viewscreen, focusing people on a mere subset of the risks at stake. What is needed is a principle of risk management, typically based on some form of cost-benefit balancing. For some problems in the area of national security, however, it is difficult to specify either costs or benefits, creating a severe epistemic difficulty. Considerable progress can nonetheless be made with the assistance of four ideas, calling for (1) break-even analysis; (2) the avoidance of gratuitous costs (economic or otherwise); (3) a prohibition on the invocation or use of illicit grounds (such as punishment of free speech or prying into people's private lives); and (4) maximin, which counsels in favor of eliminating or reducing the risk of the very worst of the worst-case scenarios. In the face of incommensurable goods, however, the idea of maximin faces particular challenges.

I. TWO TARGETS AND A THESIS

Consider two views:

1. The world has become an unprecedentedly dangerous place. Terrorist threats are omnipresent. As the 9/11 attacks display, numerous people are prepared to engage in

† Robert Walmsley University Professor, Harvard University. I am grateful to Michelle Gemmell, Heidi Liu, and Mary Schnoor for valuable research assistance, and to Eric Posner and Geoffrey R. Stone for helpful comments. I am also grateful to the participants in the Symposium at the University of Chicago Law School for valuable thoughts and suggestions. I served as a member of the President's Review Group on Intelligence and Communications Technologies ("the Review Group"), which operated from 2013 to 2014. For the final report of the Review Group, see generally *Liberty and Security in a Changing World: Report and Recommendations of the President's Review Group on Intelligence and Communications Technologies* (Dec 12, 2013) ("Review Group Report"), archived at <http://perma.cc/FG3M-QE8K>. While fully consistent with the Review Group Report, this Essay does not draw in any way on any information obtained as a result of my participation in the Review Group, and it restricts itself entirely to the public record. It should be unnecessary to say that this Essay has no official status and reflects only the views of the author.

terrorism, and they sometimes succeed. In particular, they want to kill Americans. The first obligation of public officials is to keep the citizenry safe. To do that, the best methods may well involve widespread surveillance, both domestically and abroad. If the result is saved lives, surveillance is worth it. Even when the probability of harm is low, and even if the government is operating in the midst of grave uncertainty, it is appropriate to do whatever must be done, and whatever technology allows, to prevent deaths and to protect the nation, even or perhaps especially from worst-case scenarios.

2. Americans face unprecedented threats from their own government. In the aftermath of the 9/11 attacks, the United States has seen the rise of a massive and (at least until recently) mostly secret security apparatus involving the collection of a vast quantity of data about the communications of ordinary people. Personal privacy is now at serious risk, and the same is true for free speech. “Trust us” is never an adequate response to citizens’ legitimate concerns. We need to create aggressive safeguards to protect civil liberties—not only now but also for periods in which the government is in especially bad hands—and to create precautions against the evident dangers, including worst-case scenarios.

For vividness and ease of exposition, and without ascribing particular views to any particular person, we can describe the first position as “Cheneyism,” in honor of former Vice President Dick Cheney. Consider Cheney’s suggestion:

[S]ooner or later, there’s going to be another attack and they’ll have deadlier weapons than ever before, [and] we’ve got to consider the possibility of a nuclear device or biological agent. . . . And when you consider somebody smuggling a nuclear device into the United States, it becomes very important to gather intelligence on your enemies and stop that attack before it ever gets launched.¹

There is a catastrophic worst-case scenario here, in the form of a nuclear device in the hands of terrorists in the United States.

Also for vividness and ease of exposition, and again without ascribing particular views to any particular person, we can

¹ Chris Wallace, *Former Vice President Dick Cheney Talks NSA Surveillance Program* (Fox News, June 16, 2013), archived at <http://perma.cc/FTS9-LNX4>.

describe the second position as “Snowdenism,” in honor of former NSA contractor Edward Snowden. Consider Snowden’s suggestion:

If we want to live in open and liberal societies, we need to have safe spaces where we can experiment with new thoughts, new ideas, and [where] we can discover what it is we really think and what we really believe in without being judged. If we can’t have the privacy of our bedrooms, if we can’t have the privacy of our notes on our computer, if we can’t have the privacy of our electronic diaries, we can’t have privacy at all.²

There is a catastrophic worst-case scenario here, in the form of a situation in which “we can’t have privacy at all.”³

Both Cheneyism and Snowdenism reflect enthusiasm for aggressive precautions against risks, though they display radically different perspectives on what we have to fear most. My principal goal in this Essay is to reject these two approaches and to link them with a standard, but unhelpful, approach to risks in general.⁴ I sketch a behavioral perspective on why that unhelpful approach has such widespread appeal, perhaps especially in the domain of national security. I suggest that to avoid narrow viewscreens (understood as the limited set of risks to which the analyst attends), a far better approach focuses more broadly on risk management, with a particular focus on cost-benefit analysis. One of the many advantages of cost-benefit analysis is that it reduces (without eliminating) the twin dangers of selective attention and motivated reasoning.

In the face of high levels of uncertainty, however, that approach faces especially serious challenges, above all because we may not know enough to specify either its costs or its benefits. I suggest that it is possible to respond to that uncertainty with four ideas: break-even analysis; the avoidance of gratuitous costs (economic or otherwise); a prohibition on the invocation of certain illicit grounds; and maximin, which requires attention to the worst of the worst-case scenarios. I explore how these ideas might help us move beyond Cheneyism and Snowdenism.

² Alan Rusbridger and Ewen MacAskill, *Edward Snowden Interview - the Edited Transcript* (The Guardian, July 18, 2014), archived at <http://perma.cc/4N7W-AM8M>.

³ Id.

⁴ For a discussion and critique of this approach in broad terms, see generally Cass R. Sunstein, *Laws of Fear: Beyond the Precautionary Principle* (Cambridge 2005). I draw on that discussion here.

II. AN UNHELPFUL PRINCIPLE

A. Precautions and Paralysis

In environmental policy, many people accept the precautionary principle.⁵ The idea takes diverse forms, but its central notion is that regulators should take aggressive action to avoid environmental risks, even if the likelihood of harm is very low.⁶ Suppose, for example, that there is some probability, even a small one, that the genetic modification of food will produce serious environmental harm. 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. Especially if the worst-case scenario is very bad, strong precautions are entirely appropriate. Compare the medical context, in which it is tempting and often sensible to say that even if there is only a small probability that a patient is facing a serious health risk, doctors should take precautions to ensure that those risks do not come to fruition.

In an illuminating account, the precautionary principle is understood as holding that

if an action or policy has a suspected risk of causing severe harm to the public domain (such as general health or the environment), and in the absence of scientific near-certainty about the safety of the action, the burden of proof about absence of harm falls on those proposing the action.⁷

⁵ For a general discussion of the precautionary principle in environmental policy, see *id.* at 64–76; Kerry H. Whiteside, *Precautionary Politics: Principle and Practice in Confronting Environmental Risk* 61–87 (MIT 2006) (comparing the use of the precautionary principle in environmental policy in the United States and Europe). For an especially interesting discussion, see generally Nassim Nicholas Taleb, et al., *The Precautionary Principle (with Application to the Genetic Modification of Organisms)* (NYU School of Engineering Working Paper Series, Sept 4, 2014), archived at <http://perma.cc/SJY5-QDF5> (examining the precautionary principle in the context of genetic modification and the accompanying potential risks).

⁶ See Whiteside, *Precautionary Politics* at viii (cited in note 5) (explaining the traditional applications of the precautionary principle in the context of avoiding environmental disasters).

⁷ Taleb, et al., *The Precautionary Principle* at *1 (cited in note 5). Note that Professor Nassim Taleb and his coauthors defend the precautionary principle “only in extreme situations: when the potential harm is systemic (rather than localized) and the consequences can involve total irreversible ruin, such as the extinction of human beings or all life on the planet.” *Id.*

The Wingspread Declaration,⁸ the result of a conference on the precautionary principle, puts it more cautiously: “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.”⁹

The precautionary principle has received prominent attention in other contexts as well. The influential 1992 Rio Declaration¹⁰ states, also with relative caution, that when “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.”¹¹ In Europe, the precautionary principle has sometimes been understood in a still-stronger way, suggesting that it is important to build “a margin of safety into all decision making.”¹² This stronger version, associated with both Cheneyism and Snowdenism, is what I explore here, in the form of a suggestion that when an activity, product, or situation *might* create risks, it is appropriate to take precautions against those risks even if the probability of harm is very low.¹³

In the abstract, these ideas have evident appeal. A clear demonstration of imminent or eventual harm is hardly necessary to justify precautions, not least against the risk of terrorism. But there is a serious, even devastating problem with the

⁸ *Wingspread Consensus Statement on the Precautionary Principle* (Science & Environmental Health Network, Jan 26, 1998), archived at <http://perma.cc/J592-WZ9A>.

⁹ *Id.*

¹⁰ *Rio Declaration on Environment and Development* (UN Environment Programme, June 1992), archived at <http://perma.cc/R8Y4-W7XB>.

¹¹ Bjørn Lomborg, *The Skeptical Environmentalist: Measuring the Real State of the World* 348 (Cambridge 2001).

¹² *Id.* at 349.

¹³ For a valuable and subtle discussion of the precautionary principle, see Daniel Steel, *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy* 44–68 (Cambridge 2015) (examining whether there is such a thing as “the precautionary principle”). For an instructive challenge to my arguments here, at least in the context of genetically modified organisms, see Taleb, et al, *The Precautionary Principle* at *8–11 (cited in note 5). Of course, we could imagine varieties of Cheneyism and Snowdenism that take many different forms. They might, for example, suggest that the danger is real and present rather than conjectural or probabilistic. Even in those forms, however, the analysis here is essentially unaffected. As the interest in national security or in privacy protection begins to focus on the full range of variables at stake—including expected outcomes and probabilities—it begins to converge on the risk management approach that I endorse.

precautionary principle, at least in its crudest forms:¹⁴ risks are on all sides of social situations, and efforts to reduce risks can themselves create risks. For this reason, the precautionary principle forbids the very steps that it requires. If a nation takes aggressive steps against the genetic modification of food, it might deprive people, including poor people, of food that is low in cost and high in nutrition. Precautions themselves can create “a risk of significant health or environmental damage to others or to future generations.”¹⁵

It follows that the very steps commanded by the precautionary principle violate the precautionary principle.¹⁶ The point is that few precautions lack downside risks, however speculative or remote, and if we are concerned enough to build a margin of safety into all decisions, any such margins must apply to precautions, too. Worst-case thinking can be quite dangerous.

For this reason, the precautionary principle turns out to be incoherent or even paralyzing, because it forbids the very measures that it requires. None of this means, of course, that nations should not be concerned about the genetic modification of food, or that they should demand a certainty of harm—or even a probability of harm—before undertaking regulation. If an activity creates a 1 percent risk (or less) of producing catastrophic environmental damage, then it is worthwhile to expend significant resources to eliminate that risk, even if our focus is only on expected value. People buy insurance against low-probability harms, and sensibly so. But reasonable regulators must consider both sides of the equation. Acknowledging the potential difficulty of valuation, they must engage in some form of risk management, and consider whether the costs of precautions are worth the benefits.¹⁷

¹⁴ There are many refinements of the precautionary principle. See, for example, Taleb, et al, *The Precautionary Principle* at *3–5 (cited in note 5); Steel, *Philosophy and the Precautionary Principle* at 9–16 (cited in note 13).

¹⁵ Sunstein, *Laws of Fear* at 19 (cited in note 4), quoting *Cloning, 2002: Hearings before a Subcommittee of the Committee on Appropriations, United States Senate, 107th Cong., 2d Sess 19 (2002)* (statement of Dr. Brent Blackwelder, President of Friends of the Earth).

¹⁶ For an interesting refinement and counterargument in the case of genetic modification, focused on the risk of truly catastrophic harm, see Taleb, et al, *The Precautionary Principle* at *1 (cited in note 5). Even if Taleb and his coauthors’ argument is taken as convincing, it is explicitly limited to unusual contexts and hence does not bear on the general points made in this Essay. See *id.*

¹⁷ In this Essay, I am largely bracketing the question of how to specify costs and benefits, as well as questions about distribution and equity. For a superb discussion of

B. The Appeal of Precautions

These points raise a genuine puzzle: Why do reasonable people accept forms of the precautionary principle that do not make much sense? The answers bear directly on environmental policy, but as we shall see, they help to account for the appeal of Cheneyism and Snowdenism as well. The most general point is that the precautionary principle seems appealing and workable because (and when) people use narrow viewscreens, focusing on a subset of the risks at stake rather than the whole. (There are close analogues in the domain of investor behavior.)¹⁸

Narrow viewscreens can also produce *motivated reasoning*. Suppose that we are focused above all on the risks associated with terrorism. If so, we might be motivated to discount and treat as trivial the privacy and liberty risks said to be associated with certain measures that are designed to reduce the risks of terrorism. Or suppose that we are focused above all on privacy and liberty. If so, we might be motivated to discount and treat as trivial the risks said to be associated with certain measures that are designed to protect against risks to privacy and liberty. In my view, both forms of motivated reasoning have been playing a significant role in this domain.

Three more particular factors seem especially important. The first is the *availability heuristic*. A risk that is familiar, like the risk associated with nuclear power, will be seen as more serious than a risk that is less familiar, like the risk associated with heat during the summer.¹⁹ So too will recent events have a greater impact than earlier ones. This point helps explain much risk-related behavior, including decisions to take or urge precautions. In the words of Professors Amos Tversky and Daniel Kahneman, “[A] class whose instances are easily retrieved will appear more numerous than a class of equal frequency whose instances are less retrievable.”²⁰

these topics, see generally Matthew D. Adler, *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis* (Oxford 2012).

¹⁸ See Hilary J. Allen, *A New Philosophy for Financial Stability Regulation*, 45 *Loyola U Chi L J* 173, 191–95 (2013) (identifying similarities between environmental and financial systems).

¹⁹ See Eric Klinenberg, *Heat Wave: A Social Autopsy of Disaster in Chicago* 10 (Chicago 2002) (comparing the death toll of the 1995 Chicago heat wave to that of other, better-known disasters).

²⁰ Amos Tversky and Daniel Kahneman, *Judgment under Uncertainty: Heuristics and Biases*, 185 *Science* 1124, 1127 (1974).

The central point is that for those who embrace the precautionary principle, some risks are cognitively available and others are not. Because the focus is on the former, the principle seems far more coherent than it is. Suppose, for example, that the precautionary principle has appeal in the context of nuclear power. The appeal might have a great deal to do with highly salient incidents in which the risks associated with nuclear power came to fruition, or close to it—as in the cases of Three Mile Island and Fukushima.²¹ Or suppose that the principle seems to suggest the importance of a new initiative to reduce the risk of train accidents. It would not be surprising if those who were motivated by the principle were alert to a recent train accident, which would appear to justify precautions.

The second factor involves *loss aversion*.²² Behavioral scientists have emphasized that people much dislike losses from the status quo.²³ In fact, they dislike losses about twice as much as they like corresponding gains.²⁴ The precautionary principle often seems coherent only because losses (or particular kinds of losses) are salient, while forgone gains (or other kinds of losses) are not. In the context of genetically modified food, for example, the environmental risks seem to many to be salient and “on screen” because they are self-evident losses, while the various costs of regulation might not be seen this way because they prevent potential gains.²⁵ And in the context of privacy, loss aversion can be especially important, as people strongly resist a loss of the privacy that they have come to expect.²⁶ (The idea of “reasonable expectation of privacy” may, in fact, encode some form of loss aversion.)

²¹ See John P. Christodouleas, et al, *Short-Term and Long-Term Health Risks of Nuclear-Power-Plant Accidents*, 364 *New Eng J Med* 2334, 2334–35 (2011) (comparing three major nuclear accidents and their consequences).

²² See Eyal Zamir, *Law, Psychology, and Morality: The Role of Loss Aversion* 119–65 (Oxford 2015) (explaining the theory of loss aversion and the implications of this phenomenon in the context of law).

²³ See *id.* at 4–5.

²⁴ *Id.* at 6 (“Tversky and Kahneman estimated that monetary losses loom larger than gains by a factor of 2.25.”).

²⁵ For a discussion of the importance of what is “on screen” and what is not, see Howard Margolis, *Dealing with Risk: Why the Public and the Experts Disagree on Environmental Issues* 76–77 (Chicago 1996).

²⁶ See Alessandro Acquisti, Leslie K. John, and George Loewenstein, *What Is Privacy Worth?*, 42 *J Legal Stud* 249, 255 (2013) (describing the loss-aversion heuristic as the most supported explanation of a gap in privacy valuation).

The third factor, and perhaps the most important, involves *probability neglect*.²⁷ The largest point is that if a bad outcome is emotionally gripping, people might well be inclined to eliminate it even if it has a low probability of coming to fruition.²⁸ The emotionally gripping outcome crowds out an assessment of the question of probability. And in fact, both Cheneyism and Snowdenism seem to derive a significant amount of their attraction from probability neglect. Suppose that you are asked how much you would pay to eliminate a small risk of a gruesome death from cancer, of a terrorist attack, or of a small child's death. You might well focus on the tragic outcome and not so much on the question of probability. A great deal of evidence confirms the phenomenon of probability neglect.²⁹ The precautionary principle often has appeal, and seems sensible, because some subset of risks appears emotionally gripping and the bad outcomes associated with those risks serve to crowd out other considerations.

Consider, in this regard, the finding that when people are asked how much they will pay for flight insurance for losses resulting from terrorism, they will pay more than if they are asked how much they will pay for flight insurance from all causes.³⁰ The evident explanation for this peculiar result, fitting with a form of Cheneyism, is that the word "terrorism" evokes vivid images of disaster, thus crowding out probability judgments. Note also that when people discuss a low-probability risk, their concerns rise even if the discussion consists mostly of apparently trustworthy assurances that the likelihood of harm really is infinitesimal.³¹ With these points in mind, both the appeal and the apparent administrability of the precautionary principle should be clear. The principle seems to work, and to be attractive, because of identifiable features of human cognition.

²⁷ See Cass R. Sunstein, *Probability Neglect: Emotions, Worst Cases, and Law*, 112 *Yale L J* 61, 62–63 (2002).

²⁸ See *id.* at 62 (providing examples in which people seek to avoid a negative outcome, even if the probability of its occurrence is very low).

²⁹ See *id.* at 64–68 (summarizing a number of studies documenting the cognitive biases related to probability neglect).

³⁰ See Eric J. Johnson, et al, *Framing, Probability Distortions, and Insurance Decisions*, 7 *J Risk & Uncertainty* 35, 39 (1993).

³¹ See, for example, James Flynn, Paul Slovic, and C.K. Mertz, *The Nevada Initiative: A Risk Communication Fiasco*, 13 *Risk Analysis* 497, 498–99 (1993) (describing the impact of a "counterproductive" advertising campaign that emphasized the safety of a nuclear waste repository site but changed almost nothing regarding the public's opposition).

III. PRECAUTIONS: NATIONAL SECURITY AND PRIVACY

A. Cheneyism

Some people embrace a version of the precautionary principle that no one rejects—one that grows out of the self-evidently correct claim that it is exceedingly important to counteract serious threats to the nation, including terrorist attacks.³² Vice President Cheney himself offered the core of the principle:

We have to deal with this new type of threat in a way we haven't yet defined . . . [w]ith a low-probability, high-impact event like this.

...

If there's a one percent chance that Pakistani scientists are helping al Qaeda build or develop a nuclear weapon, we have to treat it as a certainty in terms of our response.³³

In terms of standard decision theory, of course, it seems preposterous to treat a 1 percent risk the same way that one would treat a certainty. People should not, and ordinarily do not, live their lives that way. But as the stakes grow higher, the expected cost of a 1 percent risk becomes higher as well, and a precautionary approach to a 1 percent risk of catastrophe has a great deal of appeal.

For the purposes of illustration, let us focus on the question of surveillance. Even if some kinds of surveillance sweep up an immense amount of material, including much that has no interest from the standpoint of national security, some people think: surely it is better to be safe than sorry. It is tempting to emphasize the great difficulty of ruling out the possibility that, if the intelligence community obtains as much information as technology permits, it will find some information that is ultimately helpful for national-security purposes. “Helpful” here is not mere abstraction; it may mean “saves lives” or “prevents catastrophes.” Perhaps surveillance could prevent another 9/11; perhaps some forms of surveillance have not proved indispensable in the recent past but could prove indispensable in the future. A

³² See, for example, Genevieve Lennon, *Precautionary Tales: Suspicionless Counter-Terrorism Stop and Search*, 15 *Crimin & Crim Just* 44, 46–51 (2015) (describing the use of the precautionary principle in the context of “suspicionless counter-terrorism stop and search” programs in the United Kingdom).

³³ Ron Suskind, *The One Percent Doctrine: Deep inside America's Pursuit of Its Enemies since 9/11* 61–62 (Simon & Schuster 2006).

precautionary measure in ordinary life—say, the purchase of safety equipment for a car—is not valueless just because it has not proved necessary over the initial years of ownership. The measure might well be worthwhile if it avoids just one incident at some time during the life of the vehicle.

This claim could be elaborated in different ways, emphasizing diverse consequences from a successful terrorist attack. Whenever such an attack occurs, it has a series of proliferating costs, economic and otherwise. And if a future attack occurs, it might well lead to a demand for further restrictions on civil liberties, meaning that aggressive steps that are designed to protect against attacks—and that are, in the eyes of some, objectionable from the standpoint of civil liberties—might ultimately be justified or even necessary *as a means of protecting civil liberties*. With these points in view, it seems plausible to argue that, at least in the context of national security, a precautionary principle makes a great deal of sense.

In light of that point, it is similarly tempting to think: if we *can* obtain information, we *should* obtain information. This thought is especially tempting to those whose mission is to protect the nation from harm. If your job is to reduce the risk of terrorist attacks—and if you will be responsible, at least in part, for any such attacks if they occur—you might well want every available tool to increase the likelihood that no attacks will occur on your watch. That attitude might even seem indispensable to the successful performance of your most important task.

Here as elsewhere, however, the problem is that multiple risks are involved. The point may be simplest to see when the question involves standard warmaking. Any effort to use military force will create obvious risks, including risks to life and limb. What is required is a balance of risks, including probabilistic ones, rather than an abstract invocation of the idea of precaution.

The same point holds true for widespread surveillance, which creates multiple risks of its own.³⁴ Of these, perhaps the most obvious risks involve personal privacy. If the government holds a great deal of information, there is at least a risk of abuse—perhaps now or soon, but if not, then potentially in the future. We could imagine a range of possible fears and threats.

³⁴ For a discussion of the risks of widespread surveillance, see *Liberty and Security in a Changing World: Report and Recommendations of the President's Review Group on Intelligence and Communications Technologies* *46–49 (Dec 12, 2013) (“Review Group Report”), archived at <http://perma.cc/FG3M-QE8K>.

Perhaps it is the mere fact of collection that is objectionable. Perhaps public officials are learning, or would learn, about interactions or relationships for which people have a reasonable expectation of privacy. Perhaps people could be threatened or punished for their political commitments or their religion. Perhaps their conversations or relevant metadata could be released to the public, thus endangering domains that are or have become central to private life. Perhaps officials will see such conversations or such metadata, thus producing a degree of intrusion into the private domain.

There is also a risk to civil liberties, including the freedom of speech; if the government acquires metadata, there might well be (or perhaps there now is) a chilling effect on free discussion, on journalists, and on journalists' sources. Extensive forms of surveillance also create risks to commercial and economic interests and to relationships with foreign nations.

Each of these risks could be elaborated in great detail. For now, we need not undertake the relevant elaboration; the underlying risks have received a great deal of attention and have helped animate not merely proposals for reform³⁵ but also significant legislative changes.³⁶ The central point is that a form of Cheneyism, focused reasonably but solely on risks associated with terrorism, artificially truncates an appropriately wide viewscreen.³⁷

B. Snowdenism

Focusing on an important subset of risks, some people embrace a privacy precautionary principle. In their view, the risk to personal privacy requires political reforms that reduce the risk that an incompetent or ill-motivated government might, now or at some future time, jeopardize personal privacy.³⁸ In one form, associated with Snowdenism, the objection is that some invasions of privacy have already occurred and are unacceptable in a

³⁵ See generally, for example, *id.*

³⁶ See Tribune Wire Reports, *Obama Signs Bill Remaking NSA Phone Records Program* (Chi Trib, June 2, 2015), archived at <http://perma.cc/UM4C-RZSE>.

³⁷ It is true, of course, that one might endorse policies that are designed above all to reduce the risks of terrorism and that give little attention to privacy, civil liberties, and related values—not because of a limited viewscreen, but on the theory that a sensible approach to risk management, taking full account of the relevant values, justifies those policies. The discussion below is meant to address this conclusion.

³⁸ See, for example, Giovanna De Minico, *Privacy-Security: A Precautionary Principle Is Needed* (EurActiv, Jan 30, 2014), archived at <http://perma.cc/97CL-STAE>.

free society.³⁹ In another form, also associated with Snowdenism, the claim is that more egregious invasions are possible or likely if corrective steps are not taken.⁴⁰

An evident source of the privacy precautionary principle is the availability heuristic: to some people, certain highly publicized cases of abuse, at some point in the past, are very salient, not least in the United States and Europe, and they make the risk of future abuse seem far from speculative. Another underpinning is loss aversion: people are used to certain safeguards against invasions into what they see as their private domain, and widespread surveillance threatens to impose significant losses to core interests in freedom, dignity, and civic respect. A final underpinning is probability neglect. It is easy to imagine (and in the view of some, to identify) privacy violations of an extreme or intolerable sort;⁴¹ because those violations call up strong emotions, the very possibility that they will occur in the future stirs strong emotions.

At the same time, and to return to my general theme, a privacy precautionary principle, taken by itself and for all that it is worth, would not make a great deal of sense, if only because it would give rise to national-security risks—and potentially serious ones at that. The problem is that if our only or central goal were to eliminate any and all risks to privacy, we would abandon forms of surveillance that might turn out to save lives. Safeguards for privacy are of course exceedingly important, but at the conceptual level, the question remains: Why should a nation adopt a form of precautionary thinking in the context of privacy while repudiating it in the context of national security?

This question suggests that the relevant inquiries are best understood as involving a form of risk management. Here as elsewhere, risks of many kinds are on both sides of the ledger, and the task is to manage the full set, not to focus on one or a few. But the concept of risk management remains to be specified,

³⁹ See, for example, Shami Chakrabarti, *Let Me Be Clear – Edward Snowden Is a Hero* (The Guardian, June 14, 2015), archived at <http://perma.cc/X6JL-Z4L7> (“For years, UK and US governments broke the law. . . . [S]urely we can have an open and balanced discussion about how we adapt to new threats while safeguarding the intimacy and dignity rightly craved by human beings.”).

⁴⁰ See, for example, John Cassidy, *Why Edward Snowden Is a Hero* (New Yorker, June 10, 2013), archived at <http://perma.cc/4YGH-Q3X7>.

⁴¹ See *Review Group Report* at *75 (cited in note 34) (noting that “although recent disclosures and commentary have created the impression in some quarters that NSA surveillance is indiscriminate and pervasive across the globe, that is not the case”).

and in the context of national security, the effort at specification creates serious challenges. Call this the *epistemic difficulty*; it produces formidable problems for sensible risk management in this context.

IV. EXPECTED VALUES AND WORST-CASE SCENARIOS

A. The Epistemic Difficulty, Contextualized

Ideally, of course, we would be able to identify a range of possible outcomes, assign probabilities to each, and come up with some kind of common metric by which to make sensible comparisons. In regulatory policy in general, there is now a broad consensus in favor of cost-benefit analysis, understood as an effort to assess the costs and benefits of various options and to weigh the two against each other.⁴² Suppose, for example, that the monetized costs of an airline safety regulation are \$400 million and that the monetized benefits are \$70 million. If so, the regulation is unlikely to proceed, at least unless the law requires it or unless nonquantifiable benefits can be invoked to tip the balance.⁴³

One of the hardest challenges for cost-benefit analysis, of course, is that many of the variables at stake can be difficult or perhaps even impossible to monetize, thus producing one kind of epistemic difficulty.⁴⁴ In some of the most challenging cases, we might not be able to specify the relevant quantities even before we turn them into monetary equivalents. It might be unclear whether an air pollution regulation will save five hundred lives, or a thousand lives, or two thousand lives, or three thousand lives. If the value of a statistical life is \$9 million,⁴⁵ then the monetized mortality benefits range from \$4.5 billion to \$27 billion—a stunningly wide range. And in some regulatory settings, benefits cannot be quantified in any helpful way, simply because

⁴² See Executive Order 13563 § 1(a), 3 CFR 215, 215 (“Our regulatory system . . . must take into account benefits and costs, both quantitative and qualitative.”); Cass R. Sunstein, *Valuing Life: Humanizing the Regulatory State* 36–42 (Chicago 2014) (discussing the decisionmaking process within government and the impact of Executive Order 13563).

⁴³ See Sunstein, *Valuing Life* at 36–42 (cited in note 42).

⁴⁴ See Cass R. Sunstein, *The Limits of Quantification*, 102 Cal L Rev 1369, 1373–85 (2014) (examining the challenges that arise in attempts to quantify regulatory benefits).

⁴⁵ See, for example, Peter Rogoff and Kathryn Thomson, Memorandum to Secretarial Officers and Modal Administrators *1 (Department of Transportation, June 13, 2014), archived at <http://perma.cc/2XA8-GDXN> (revising the value of a statistical life for 2014 to \$9.2 million); Sunstein, *Valuing Life* at 94 (cited in note 42).

regulators lack the relevant knowledge. Here, the epistemic difficulty turns out to be formidable.⁴⁶

In the context of national security, the challenge of quantification can be even more daunting. Suppose that the relevant risk is a terrorist attack. In advance, it might be exceedingly difficult to quantify the costs of such an attack. How many lives are at risk? Ten? Two hundred? Three thousand? More? Even if the number is at the low end of the scale, we have seen that any terrorist attack has proliferating costs, some of them involving life itself.⁴⁷ Of course, assessment of the expected value of precautions must also engage with the question of probability: If an initiative is undertaken, what is the reduction in the probability of a successful terrorist attack? Officials might not be able to specify the answer to that question. In this respect, the domain of national security overlaps with that of financial regulation, in which the identification of the benefits of regulatory safeguards can also be daunting.⁴⁸

There are second-order effects as well as first-order effects: What kinds of social consequences follow from a successful terrorist attack? Do they include long-term economic costs? Do they include intrusions on privacy and liberty? If so, how should these be counted in the risk management calculation? Should a civil libertarian favor national-security safeguards that appear to threaten civil liberties, on the ground that if they are successful, those very safeguards will help to preserve civil liberties against further intrusions? These questions might prove difficult to answer when policymakers are assessing particular programs.

B. Break-Even Analysis (and Its Discontents)

Even if such questions do not have clear answers, officials may not be entirely at sea. Within the federal government, it is standard to speak of break-even analysis, by which officials ask: *What would the benefits have to be for it to be worthwhile to*

⁴⁶ See John C. Coates IV, *Cost-Benefit Analysis of Financial Regulation: Case Studies and Implications*, 124 Yale L J 882, 888 (2015) (“These features [of cost-benefit analysis] undermine the ability of science to precisely and reliably estimate the effects of financial regulations, even retrospectively.”).

⁴⁷ See generally, for example, Gerd Gigerenzer, *Dread Risk, September 11, and Fatal Traffic Accidents*, 15 Psychological Sci 286 (2004) (finding that in the three months following 9/11, a rise in traffic deaths resulted in a number of traffic deaths that exceeded the number of plane passengers killed in the attacks, because more individuals elected to drive rather than fly during those months).

⁴⁸ See Coates, 124 Yale L J at 893–95 (cited in note 46).

*impose the costs?*⁴⁹ Suppose, for example, that the costs of a rule that would protect against some environmental risk were \$200 million but that the benefits could not be specified. We might be able to say that at its upper bound, the cost of the environmental damage, if it were to occur, would be \$100 million. If so, the rule could not easily be justified.⁵⁰

Now suppose that at its upper bound, the cost of the environmental damage would be \$900 million. If so, it is not clear that the benefits would fail to justify the costs. An obvious question would be: What kind of contribution would the rule have to make to the prevention of that damage? If the rule can be taken to reduce the risk by 10 percent, the costs and the benefits would be fairly close. Of course, the agency might not be able to specify any such percentage. But perhaps it is able to identify lower and upper bounds. If the lower bound in terms of risk reduction is (say) 15 percent, then the benefits do seem to justify the costs.

With approaches of this kind, break-even analysis can make seemingly intractable problems far more tractable. Suppose, for example, that officials know the upper or lower bound of the costs associated with a risk if it comes to fruition, or suppose that they know the number of people or activities that might be affected (even if they do not know the costs of the per-person or per-activity impact). If so, break-even analysis might prove feasible, and it might suggest that a regulation is either clearly desirable or clearly a mistake. Even in standard settings, it is possible that regulators will know too little to make use of that form of analysis; but if they have even small pockets of knowledge, the approach can greatly clarify their judgments.

At least in theory, break-even analysis can play a role in the context of national security as well. There are many complexities here, so let us consider a highly stylized example. Suppose that the cost of a terrorist attack, if it were to occur, is at least \$200 billion, and suppose that the measure in question would reduce the probability of its occurrence by 10 percent. (Nothing turns on these particular numbers, which are introduced simply for purposes of analysis.) Suppose too that the measure in

⁴⁹ Sunstein, *Valuing Life* at 65–67 (cited in note 42).

⁵⁰ For present purposes, I am putting to one side questions about distribution and equity as well as questions about nonquantifiable benefits. For relevant discussions of these questions, see *id.* at 127–30 (discussing distribution and equity concerns in cost-benefit analysis); *id.* at 58–60, 66–67 (discussing nonquantifiable benefits).

question would consist of security precautions at airports or certain forms of surveillance. We might ask: Is the cost of the invasion of privacy in excess of \$20 billion? Of course, there is no purely arithmetic answer to that question⁵¹—but the question itself might turn out to be helpful, at least if we know something about the nature of the risk to privacy. Advocates of the measure will press a legitimate question: Is it even plausible to think that the risk to privacy is worth more than \$20 billion?

This is of course an artificial example, and in this context, break-even analysis runs into particular trouble, at least if we indulge the reasonable assumptions that a great deal of important information is missing and that moral valuations will play an inescapable role. If hard-to-quantify costs are on both sides of the ledger—as they are in the contexts under discussion—then break-even analysis becomes especially hard to undertake. If so, its chief advantage is that it may promote transparency about the issues involved.

C. Avoid Gratuitous Costs

Diverse people should be willing to converge on a simple principle: *avoid gratuitous costs*. In the environmental context, that seemingly self-evident principle turns out to have real bite. Suppose, for example, that on reflection, certain environmental risks turn out to be *de minimis*, in the sense that they are trivial.⁵² It makes sense to say that the government should not regulate those risks, at least if regulation itself imposes costs. The principle is also important in the context of national security. Suppose that some forms of surveillance produce no benefits or only *de minimis* benefits. Suppose that their only function is to pick up information that cannot plausibly contribute to the prevention of terrorist attacks. If so, there would seem to be no reason that they should be continued.

The principle does not only inform the scope of surveillance activities. It should also inform the design of relevant institutions.

⁵¹ I am putting to one side questions about the use of “willingness to pay” to value privacy issues. For a relevant discussion of this question, see *id.* at 91–92 (explaining that one method of producing monetary amounts for statistical risks is to “ask people how much they are willing to pay to reduce statistical risks”).

⁵² In the easiest cases, the judgment of triviality comes from the fact that even if these risks come to fruition, they do not involve much harm. In harder cases, the judgment of triviality comes from a calculation of expected value. If such a calculation is possible, of course, then the epistemic difficulty is not so large.

For example, the Review Group recommended that metadata should be held not by the government itself but rather by the phone companies, with access by the government on the basis of the appropriate showing of need.⁵³ We can understand this recommendation as an outcome of the no-gratuitous-costs principle. On optimistic (but not unrealistic) assumptions, it would deprive the government of exactly nothing that it is important for the government to have, while also providing a layer of protection against risks to privacy and free speech. Even if the government does not hold the metadata, it can obtain it on a showing of need, and indeed if time requires (for example, under emergency conditions), it need not obtain judicial authorization in advance. Under these assumptions, the Review Group's recommendation flows directly from the no-gratuitous-harm principle.

That principle is a sensible way to provide a layer of privacy protection without threatening national security. A much more controversial question still exists: Can the principle be used to *scale back* some kinds of apparent privacy protections, on the ground that they do no real good in terms of privacy but also impose some costs (in the form of national-security risks)? However uncomfortable it may be, this question deserves attention.

D. Avoiding Illicit Grounds

If the purpose of surveillance is to protect national security, then some grounds for and uses of surveillance are automatically off-limits. They do not count in the balance at all.⁵⁴ This is an exceedingly important idea, because it captures and takes directly on board some of the most plausible judgments behind a privacy precautionary principle. More specifically, it addresses several concerns that motivate that principle.

The major categories are straightforward.⁵⁵ Surveillance cannot legitimately be used to punish people because of their political views or religious convictions. Under current conditions, surveillance that is designed to reduce risks to national security

⁵³ *Review Group Report* at *17 (cited in note 34) (“We recommend that Congress should end such storage [of metadata] and transition to a system in which such metadata is held privately for the government to query when necessary for national security purposes.”).

⁵⁴ For a discussion of “exclusionary reasons,” see Joseph Raz, *Practical Reason and Norms* § 1.2 at 39–40 (Princeton 1990) (“An *exclusionary reason* is a second-order reason to refrain from acting for some reason.”).

⁵⁵ See *Review Group Report* at *16 (cited in note 34) (“[S]ome safeguards are not subject to balancing at all.”).

should probably not be designed to protect against criminal activity that raises no national-security issue.⁵⁶ If the underlying activity involves unlawful gambling or tax evasion, there are established routes by which the government may obtain the relevant information. It is generally agreed that surveillance should not be designed to give a commercial advantage to American firms.⁵⁷ In these and other respects, the interest in national security—which is what motivates surveillance in this context—also limits and disciplines the permissible grounds for surveillance. No sensible form of Cheneyism should reject those limits.

To be sure, we could imagine more-difficult cases. Suppose, not implausibly, that a certain set of political views or identifiable religious convictions are closely associated with a desire to do harm to the United States and its allies. If people are members of the Islamic State of Iraq and the Levant (ISIL), for example, the United States is entitled to focus on them by virtue of that fact. But the reason involves national security, not politics or religion as such. We can imagine cases that might test the clarity of that line, but the basic principle should not be obscured.

E. Avoid the Worst of the Worst Cases

Decision theorists sometimes distinguish between situations of *risk*, in which probabilities can be assigned to various outcomes, and situations of *uncertainty*, in which no such probabilities can be assigned.⁵⁸ In the domain of national security, we can imagine instances in which analysts cannot specify a usefully narrow range of probabilities and in which the extent of the harm from bad outcomes is also not susceptible to anything like precise prediction. Here again the analogy to financial regulation is plausible: analysts might be able to identify only an unhelpfully wide range of bad outcomes, and they might not be

⁵⁶ To be sure, there are imaginable complexities here, such as when surveillance that is meant to protect against national-security risks uncovers a plan to commit acts of violence that do not involve national security.

⁵⁷ See generally Samuel J. Rascoff, *The Norm against Economic Espionage for the Benefit of Private Firms: Some Theoretical Reflections*, 83 U Chi L Rev 251 (2016) (discussing US intelligence policy's apparent aversion to the collection of intelligence for the benefit of American firms).

⁵⁸ See Paul Davidson, *Is Probability Theory Relevant for Uncertainty? A Post Keynesian Perspective*, 5 J Econ Persp 129, 130 (Winter 1991) (differentiating "true uncertainty" from situations of probabilistic risk); Frank H. Knight, *Risk, Uncertainty and Profit* 233 (Houghton Mifflin 1921).

able to say a great deal about the contribution of a regulation to the prevention of such outcomes.⁵⁹

In situations of uncertainty, when existing knowledge does not permit regulators to assign probabilities to outcomes, it is sometimes suggested that rationality calls for invocation of the maximin principle: *choose the policy with the best worst-case outcome*.⁶⁰ Suppose that the worst case associated with one policy involves a successful terrorist attack on the United States, with a significant loss of lives. Suppose that the worst case associated with another policy involves a serious threat to privacy, in the form of (say) widespread official reading of private metadata (or more), leading to the official invasion of the private sphere. Suppose finally that we cannot say much about the probability that one or another worst case will occur. In a case of that kind, there is a good argument for Cheneyism and a much weaker one for Snowdenism. The reason is that the worst case associated with a successful terrorist attack is so much worse than the worst case associated with a breach of personal privacy.

To make this argument work, of course, we need to have thresholds of plausibility to discipline the universe of worst-case scenarios. We could construct a chain of causal links by which any bad thing leads to truly catastrophic bad things—as, for example, when practices of the NSA are taken to produce, as a worst-case scenario, the end of democracy in the United States. But it seems reasonable to say that the threshold of plausibility rules that case out of bounds.

Of course, the problem I have presented is artificial along multiple dimensions. We might be speaking of *bounded uncertainty*:⁶¹ In a particular period, the probability of a successful terrorist attack might not be between 0 percent and 100 percent but rather between 0 percent and 30 percent (though we might not be able to say much about where it falls within that range). We might be able to say that if a terrorist attack occurs, very bad outcomes would have a cost between \$X and \$Y, in which \$X

⁵⁹ See Coates, 124 Yale L J at 894–95 (cited in note 46) (discussing some of the difficulties involved in predicting the effects of financial regulations).

⁶⁰ For a helpful discussion, see Jon Elster, *Explaining Technical Change: A Case Study in the Philosophy of Science* 185–207 (Cambridge 1983). Within decision theory, this is not an uncontested view. I am bracketing the complexities here.

⁶¹ For a discussion of “bounded uncertainty,” see Mie Augier and Kristian Kreiner, *Rationality, Imagination and Intelligence: Some Boundaries in Human Decision-Making*, 9 *Indust & Corp Change* 659, 670–72 (2000) (describing how human decisionmaking is limited to individual knowledge and to the boundaries of what is considered possible).

is (say) \$100 billion, and in which \$Y is (say) \$950 billion (these numbers are merely illustrative). And while the contribution of a particular policy might not be susceptible to precise specification, policymakers might have an idea of a sensible range.

Moreover, maximin is most useful in cases in which the outcomes can easily be rendered commensurable. Suppose that a policymaker has two options, which would lead to different worst-case scenarios: (1) a loss of \$500 million or (2) a loss of \$900 million. The option that leads to the lower worst-case loss is better (and it is clear which is lower). Or suppose that with (1), the worst-case scenario involves a loss of one thousand lives, whereas with (2), the worst-case scenario would lose six thousand lives—no ambiguity there. But the issue is more difficult when the outcomes are not easily commensurable. Suppose that a policymaker has two options, with different worst-case scenarios: (1) a loss of \$600 million and also 200 lives and (2) a loss of \$900 million and also 150 lives. Are the 50 lives saved from (2) worth the \$300 million cost? The answer depends on the value of a statistical life. The government now values a statistical life at about \$9 million, so the answer is yes.

But far harder cases are imaginable. In the context at hand, suppose that with one approach, the worst-case scenario is a loss of significant numbers of lives whereas with another, the worst-case scenario is a massive intrusion into personal privacy. For progress to be made, both of these consequences would have to be specified. How many lives? One thousand, or five thousand, or forty thousand? More? And what kind of intrusion counts as massive? Issues of valuation cannot be avoided here. Official reading of (say) private metadata is far more alarming to some people than to others.

On one view, which I find reasonable, a certain degree of vulnerability with respect to private metadata does not involve anything like the worst-case scenarios associated with successful terrorist attacks. That view might be accompanied by a judgment that the risk of vulnerability with respect to private metadata can be sufficiently contained. But on a competing view, a cavalier approach to personal privacy threatens both liberty and self-government themselves, and so the worst-case scenario is very bad indeed (and cannot be ruled out).

Disagreements of this kind cannot be resolved by arithmetic. A reference to maximin will not do the trick. Perhaps the best that can be done is to attempt to identify safeguards with

respect to privacy that plausibly reduce the risks associated with worst-case scenarios while also allowing officials to do what must be done with respect to national security. In the abstract, it might well seem more difficult to achieve that goal than it is in practice.

CONCLUSION

In ordinary life, people take precautions, and sensibly so; insurance policies are often an excellent idea. The precautionary principle is animated by the reasonable idea that it is prudent to act even when it is far from certain that the underlying danger will come to fruition. The problem is that action can create dangers of its own. In the environmental context, the precautionary principle runs into self-evident trouble when efforts to reduce some environmental risks give rise to other environmental risks. But it is also problematic when those efforts create risks that have nothing to do with the environment. A wide viewscreen, rather than a narrow one, is indispensable in the regulatory domain.

In the area of national security, it may be especially tempting for public officials to adopt some kind of precautionary principle, not least because they are confronted with a dazzling array of low-probability risks. It would be both irresponsible and dangerous to ignore those risks. At the same time, some precautions create risks of their own, and they must be considered in an overall balance. Cheneyism, as I have described it here, runs afoul of the need for wide viewscreens. The same point certainly holds for those who embrace Snowdenism, which I have described as an insistence on a precautionary principle for privacy and civil liberties.

To the extent feasible, the best approach to risk management involves cost-benefit balancing. The challenge is that in some domains, both costs and benefits are exceedingly hard to quantify, much less to monetize. The epistemic difficulty is severe.

I have argued for four ideas that can help. First, despite its difficulties, officials should nevertheless consider the use of break-even analysis. Second, officials should not impose essentially gratuitous costs (including risks). Third, officials should ensure that illicit grounds are not being invoked to intrude on privacy, liberty, or anything else. Fourth, officials should take steps to prevent the worst of the worst-case scenarios. There are

no algorithms here, but a form of risk management, embodying these ideas, can help to avoid some of the pathologies of both Cheneyism and Snowdenism: precautionary principles of the most blinkered or myopic sorts.