BLINDSIDE

How to Anticipate Forcing Events and Wild Cards in Global Politics

Francis Fukuyama, *editor*

An *American Interest* Book

BROOKINGS INSTITUTION PRESS
Washington, D.C.

Contents

1

The Challenges of Uncertainty: An Introduction

Francis Fukuyama

The collapse of communism, the rapid emergence of China and India as major economic powers, the September 11 attacks, the appearance of relatively new diseases like HIV/AIDS and H5N1 bird flu, Hurricane Katrina—the past decade and a half has demonstrated that nothing is as certain as uncertainty in global politics. As the famous scatological bumper sticker suggests, bad things happen. But there are benign surprises as well, and these, no less than catastrophic events, challenge society's capacity to understand, to adapt, and to lock in good fortune.

Anticipating and dealing with what were thought to have been very low-probability events have clearly become central challenges for policymakers in public and private sectors alike all over the world. This book, sponsored by *The American Interest* magazine, addresses those challenges. The magazine's first annual event, held in Washington, D.C., in May 2006, brought together analysts, practitioners, policymakers, and unconventional thinkers from a variety of backgrounds and disciplines. The magazine assembled those who think about discrete uncertainties and who also have considered the very nature of uncertainty itself. The distinction may seem a minor one, but it is not. Many people, from intelligence analysts to investment bankers to corporate treasurers, are paid to think about discrete futures in their areas of professional responsibility. But far fewer people have ever tried to understand *why* the future is inherently difficult to anticipate, and how to mitigate our blindness

to its vicissitudes in a more systematic way than societies and governments have been able to do heretofore.

The task is a complex one. Those who deal professionally with global politics, foreign policy, and national security affairs have particular biases when it comes to thinking about the future. Those biases generate a perceptual incentive structure that throws off their general capacity for accurate prediction. Such analysts, after all, are seldom rewarded for predicting continuity or the sudden emergence of good news, but failure to predict bad news can be a career-ending mistake. No one wants to be in the position of Admiral Husband Kimmel, the commander of the U.S. Pacific Fleet, who was on duty the day of the Japanese attack on Pearl Harbor. Admiral Kimmel's code-breakers had deciphered the Japanese "winds" code, but he, the principal consumer of that intelligence, nonetheless failed to anticipate that the actual blow would land on the Pacific Fleet headquarters. The day that would live in ignominy well described the resting place of Admiral Kimmel's reputation for all time.

By contrast, no one in the U.S. intelligence community was cashiered for failing to predict that the Berlin Wall would come down in November 1989, though fail to predict it they did. This asymmetry in incentives leads the vast majority of those who work on national security issues to resort routinely to worst-case analysis as a means of covering themselves in case bad things happen on or just beyond their watch.

Investment bankers and corporations, in contrast, have much more balanced incentives to think about the future. For them, a failure to anticipate an opportunity can have very costly consequences—indeed, upside potential is often greater than downside, since one can only lose what one owns in an unleveraged situation. Businessmen may even be guilty of laboring under incentive structures that are biased toward excessive optimism.

The bias against optimism in the governmental intelligence world was nowhere more evident than in the estimates made of Iraqi weapons of mass destruction (WMD) before the 2003 Iraq invasion. After the entry of United Nations inspectors into Iraq in the wake of the 1991 Gulf War, the intelligence community found itself in a Husband Kimmel–type situation. Iraq was far closer to a nuclear capability than anyone outside of the Iraqi elite had realized before the 1991 Gulf War. Everyone involved was under enormous psychological pressure not to be duped again, and it was for that reason that everyone, from UN weapons inspectors to the U.S. intelligence community to their French and Russian counterparts, assumed that Saddam Hussein had more capabilities than the inspections were able to uncover. After the 2003 invasion

the world learned that the regime was incompetent, corrupt, and compartmentalized to the point where many senior Iraqi officials (including, at times, Saddam himself) believed their country possessed WMD capabilities that did not in fact exist. Before the invasion, though, it would have taken a brave (or foolhardy) intelligence analyst to aggressively downplay the danger represented by Iraq's WMD programs. One who did, Scott Ritter, had his motives severely impugned.

It is, of course, not possible to anticipate all the possible low-probability events that may litter world history in coming months and years. And even if one could anticipate many different futures, it would be impossible to hedge against all of them. Hedging is usually an expensive strategy in which high opportunity costs forsaken have to be weighed against other alternatives. So how does one deal practically with the problem of being blindsided?

This volume is organized into five sections. The first, of which this chapter is a part, introduces the book and examines the fact that surprise is, almost by definition, a psychological problem. Richard Posner, author of *Catastrophe*, observes, for example, that even though one can show that it would be cost-effective to hedge against a low-probability event like an asteroid strike, policymakers and the politicians who hire them are unwilling to pay the cost because they simply cannot imagine such a contingency becoming real.[1] It often takes a Hollywood movie or a similar event occurring in a different country to enable people to visualize a contingency and thus to act on it.

Important institutional constraints, moreover, make it difficult to act even when some people can and do accurately anticipate a low-probability, high-impact contingency. One might call this a form of "socio-surprise" characteristic of collective psychologies. Hurricane Katrina, for example, was one of the most fully predictable and scenario-tested natural disasters in American history, but that fact still did not lead to appropriate preparatory actions or adequate crisis responses on the part of responsible officials at the local, state, or federal levels.

The following section, "Cases: Looking Back," looks more closely at some historical examples of surprise—upside as well as downside—and asks why the social and economic impacts of emergent technologies and events like the collapse of the former Soviet Union and the 1997–98 Asian financial crisis were not anticipated. David Landes, Bruce Berkowitz, and David Hale draw on their knowledge of history and policy to pinpoint those institutional, and not just personal, failures that prevented policymakers and others from properly anticipating major events of the time.

The third section discusses potential future cases of surprise. William Bonvillian, director of MIT's Washington office, suggests ways to set up institutions so that they can deliberately create surprises—positive surprises—in this case for dealing with energy technology and policy issues. Based on the connected-science model that led to applied scientific advances during World War II and thereafter to the establishment and flourishing of DARPA (Defense Advanced Research Projects Agency), Dr. Bonvillian teases out the essence of effective innovation systems. He then applies this essence to what a DARPA-like innovation-generating agency would look like if applied to energy technology. The section also includes chapters by Gal Luft and Anne Korin on other aspects of energy policy, and by Scott Barrett on the uncertainties and dangers concerning new, potentially global-scale diseases.

For reasons rehearsed by Richard Posner, it is clear that psychological preparedness for low-probability events—even ones generated on purpose—is both extremely important, and extremely difficult to achieve. The next section, "Forecasting," tackles this problem.

It is not possible to anticipate all possible futures or to hedge against even a small proportion of them. The incentives to do so are not always present either. After all, politicians need to get reelected in the near term and therefore seldom have the incentive to worry about costs that will be incurred after they have departed the stage. Yet there are nonetheless systematic ways of looking at the future. One is through traditional cost-benefit analysis of the sort Posner outlines, but with proper discounting of future costs and opportunities. Another way of approaching the problem is through scenario methodology. This section thus begins with an essay by Peter Schwartz and Doug Randall of Global Business Network.

Peter Schwartz, whose career started in Royal Dutch/Shell's planning division, has built a business model and career around scenario planning. He observed many years ago that if one proceeded on the basis of a straightforward rational choice model in which one thought through different futures and assigned probability weights to them, senior decisionmakers would simply stop thinking about the low-probability ones. The chief problem is to overcome the psychological resistance to thinking about low-probability futures; the company he created, Global Business Network, has engaged in scenario planning that deliberately ignores probabilities to do precisely that. The focus, Schwartz and Randall argue, needs to be on the decisionmakers themselves and on the institutional constraints they face that allow them to

avoid thinking creatively about the future. In this volume, they reflect on the experience of scenario planning over the past two decades.

Another way of thinking systematically is to select hedges, as Robert Lempert recommends in his essay, that are robust over the largest number of possible futures. Like Schwartz and Randall, he is interested in how scenario methodologies can help policymakers find the proper balance between boldness and care in their planning functions.

Technological change has driven much economic and political change. No wonder, then, that technological forecasting has become a staple of our world. And yet even with skillful cost-benefit analysis and scenario techniques, technological forecasters continue to get most things wrong, failing to anticipate major trends and overestimating the importance of the "latest great thing." Could this be because political and economic change also drives technological change, that while technologists are reading the tea leaves from left to right, reality is proceeding from right to left? Why technology forecasting is so poor and why it will probably continue to be poor is the subject of the essay by Mitchell Waldrop, who uses innovation in information technology as his base example to explain why forecasting is so difficult to get right.

The Blindside conference took to heart the very nature of the subject, which calls attention to the cognitive challenges of dealing with low-probability events and so put a premium on getting beyond the usual solipsistic habits of academe. As the program itself featured a debate and two discussions-in-the-round, so those events are represented in this book. The final section, "What Could Be," begins with an edited transcript of a debate between James Kurth of Swarthmore College and Gregg Easterbrook of the Brookings Institution on what may fairly be called the philosophy of forecasting. Kurth uses his wide-ranging intellect to creatively join developments in the geopolitical and spiritual realms and spin out a gloomy prognosis for a declining West. Easterbrook, drawing on his book *The Progress Paradox*, points out that those who look to the facts, rather than to their fears and instincts, will find that the world has been getting better by any number of measures and is likely to continue doing so.[2]

The second and third chapters of the section feature discussions by members of *The American Interest*'s editorial board, dealing in turn with international and American scenarios for low-probability, high-impact events to come. Both the debate and the two discussions evoked novel observations that no one participant would likely have hit upon alone.

There is no easy or obvious methodology that will prevent us from being blindsided in the future. It is important, however, to understand the specific obstacles, both psychological and institutional, that prevent us from first seeing the future clearly and then acting on our insights in a responsible way. The essays in this volume lay out the conceptual problem of anticipating unexpected events, provide glimpses of different possible futures across a range of regions and issues, and may even offer up some creatively practical advice about how to plan for those futures. It is to this kind of creative thinking that *The American Interest* is dedicated.

2

Thinking about Catastrophe

Richard A. Posner

A catastrophe, as I use the term, is an unexpected event that causes great harm. The two parts of the definition fit logically because most harmful events that are expected can be mitigated by preventive measures, often or at least sometimes rendering them less than catastrophic. As science advances, enabling greater predictive accuracy at least over the natural world, scientists may be able to predict catastrophes that cannot be prevented (as well as prevent some man-made catastrophes that cannot be predicted). So to be more precise, let me substitute for "unexpected event" the phrase "event of low or unknown probability"—that is, an event that is either low-risk or uncertain, in the statisticians' useful distinction between risk and uncertainty.

Catastrophes have been common enough in human history, and there is no paradox in describing as "common" a series of events of low or unknown probability. That is because the very low probabilities of a very large number of possibilities can aggregate to a probability close to 1. But—and here *is* a paradox—the problem of dealing with catastrophes has gotten bigger simply because the range of catastrophes that cannot be averted or mitigated has gotten smaller. Until recently, for example, there was no reason to worry about an asteroid strike because nothing could be done to prevent it and little if anything could be done to mitigate its effects. Now, by careful mapping of the orbits of asteroids whose orbits intercept Earth's orbit, and by skillful deployment of rockets to nudge dangerous asteroids out of their current orbits, catastrophic asteroid strikes can be prevented.[1] Early-warning systems for tsunamis are now

available, as well. So now the question is whether to take these defensive measures, given that the required investment of human and technical resources could also be used elsewhere. The larger the number of preventable catastrophes that are identified, the more difficult the issue of resource allocation becomes.

The march of science is doing something else: It is creating *more* catastrophic risks. This is obvious in the case of modern weapons of mass destruction, but it is also happening in more esoteric areas of scientific advance such as genetically modified crops, nanotechnology, and robotics. Sheer economic and population growth is doing its share, too, by contributing to global warming but also by making certain places more vulnerable to catastrophe and by making weapons of mass destruction more affordable. As to vulnerability, the August-September 2005 flooding of New Orleans was a consequence in part of economic development that had eliminated natural barriers to flooding. As to affordability, there is a particularly sinister conjunction with respect to biological weapons, because they are becoming cheaper at the same time that they are becoming more lethal and that people and nations are becoming wealthier. The Unabomber attacked with letter bombs; a biological Unabomber with bioengineering skills may soon be able to attack with aerosols of synthesized smallpox virus.

On top of all this, geopolitical changes, in conjunction with the increased availability of such weapons, are increasing the likelihood of catastrophic attacks. Here I refer specifically to the rise of global terrorism with apocalyptic aims and suicidal means that preclude effective deterrence, the growing instability and hostility (to the West and its allies) of the vast Muslim world, and the emergence of heavily armed "rogue states," notably (at the moment) Iran and North Korea.

In sum, there are many more catastrophic risks that we can productively worry about than ever before: more risks, and a greater proportion of preventable risks. This is an uncomfortable situation for three fundamental reasons, one psychological, a second political, and a third analytical.

Psychological Discomforts

The psychological problem is the difficulty of getting people, even most officials and many intellectually gifted people, to think seriously about catastrophes that have not yet happened. It takes no imagination to think seriously and productively about airplane crashes, forest fires, cardiac arrest, and

other common catastrophes (large or small), because these things happen frequently enough to make them part of ordinary experience. In other words, these events do not need to be *imagined*. But thinking that absolutely requires imagination is another matter altogether.

This observation has implications that are often underestimated. Imagination is a very scarce resource, and also a highly imperfect one, because thinking about things that have not happened is inherently more difficult than thinking about things that have. For one thing, probabilities—things that may or may not happen—are far more difficult to estimate than frequencies—things that will happen sooner or later. With probabilities, too, the human imagination is quickly overwhelmed because, while only a finite number of things *has* happened, the range of things that *may* happen is literally infinite. It is impossible for an individual, a government, or even a supercomputer to think about an infinite number of things. So the larger the array of possible preventable catastrophes, the more there is to think about, implying heavy demand and potential exhaustion of the imagination capacity of the society.

This problem of cognitive overload is exacerbated—and here I verge into the political impediments to responding to catastrophic risks—by the fact that the human mind has great difficulty thinking in probabilistic terms, especially when the probabilities are low. This is a problem not of having to think about too many things at once, but of having to think about one thing that is of low or unknown probability. Human brains did not evolve to deal easily with such events because in the ancestral environment, as evolutionary biologists call it, when human brains assumed their approximate current structure, there was no payoff to being quick-witted about probabilistic events about which one could do very little or nothing. A large and convincing literature in cognitive psychology shows that nonexperts handle probabilistic dangers very badly, sometimes exaggerating them unreasonably but more often writing them down to zero, that is, ignoring them.[2]

In a democratic society, the reactions of nonexperts, that is, of ordinary people, have a significant impact on public policy. President George W. Bush's science adviser once told me that while he appreciated that asteroid strikes were a menace that might justify a greater investment of national resources in detection and prevention, the investment would not be made because the American people simply do not worry about asteroid strikes, even though such a strike could, depending on the size of the asteroid, do incalculable damage up to and including the extinction of the human race.

Political Obstacles

The political obstacles to responding intelligently to catastrophic risks are magnified by the short horizons of politicians and many of their constituents. The probability of an event is a function of the time interval under consideration. A biological attack on the United States is much more likely within the next ten years than within the next week, and is less likely (though this is little more than a guess) within the next six months than an unfavorable outcome of the war in Iraq. So the natural tendency is to focus more on the war than on the threat of a biological attack. Likewise, if global warming is not likely to cause serious harm for another century or so, as many scientists believe (though there is some unknown danger of earlier, abrupt climate change), politicians are unlikely to take costly measures to combat it even if deferring the measures would greatly increase the harm. Civil servants often have longer horizons than politicians, but not much longer.

Politicians and civil servants are not the only ones with truncated horizons; ordinary individuals have them, as well. Their horizons are longer (as politicians' horizons would be if political office were hereditary) because of altruism toward one's children and grandchildren and perhaps even remoter descendants. But they are not infinite, which means that neither the current generation nor its political representatives internalize the welfare of remote unborn generations.

Some of the political obstacles to responding effectively to catastrophe can be more fully appreciated by examining the tsunami that ravaged the coastline of the Indian Ocean in December 2004. Suppose that a tsunami of that destructiveness occurs on average once a century and kills 250,000 people. That is an average of 2,500 deaths a year. Even without attempting a sophisticated estimate of the value of life to the people exposed to the risk, one can say with some confidence that if an annual death toll of 2,500 could be substantially reduced at a moderate annual cost, the investment would be worthwhile. A combination of educating the residents of low-lying coastal areas about the warning signs of a tsunami (tremors and a sudden recession in the ocean); establishing a warning system involving emergency broadcasts, telephoned warnings, and air-raid-type sirens; and improving emergency response systems would have saved many of the people killed by the Indian Ocean tsunami, probably at a total cost below any reasonable estimate of the average losses that can be expected from tsunamis. Relocating people away from coasts would be even more efficacious, but except in the most vulnerable areas or in areas in which residential or commercial uses have only marginal value, the costs would

probably exceed the benefits because annual costs of protection must be matched with annual, not total, expected costs of tsunamis.

So why were such systems not in place when the 2004 tsunami struck? First, although a once-in-a-century event is *as* likely to occur at the beginning of the century as at any other time, it is much *less* likely to occur during the first decade of the century than at some time in the last nine. Politicians with limited terms of office and thus foreshortened political horizons tend to discount low-risk disaster possibilities steeply because the risk of damage to their careers from failing to take precautionary measures is truncated.

Second, to the extent that effective precautions require governmental action, the fact that government is a centralized system of control makes it difficult for officials to respond to the full spectrum of possible risks against which cost-justified measures might be taken. Given the variety of matters to which they must attend, senior officials are likely to have a high threshold of attention below which risks are simply ignored (and the more senior they are, the higher the threshold). The upper levels of the U.S. government, preoccupied with terrorist threats, paid insufficient attention to the risk of a disastrous flood in New Orleans, even though the risk was understood to be significant.

Third, where risks are regional or global rather than local, many national governments, especially in poorer and smaller countries, may drag their heels in the hope of taking a free ride on larger and richer countries. Knowing this, and not wishing to reward and thus encourage free riding, the richer countries may be reluctant to take precautionary measures. Again, there is a U.S. parallel: State and local governments may stint on devoting resources to emergency response, expecting aid from other state and local governments and the federal government.

Fourth, countries are poor often because of weak, inefficient, or corrupt government, characteristics that may disable poor nations from taking cost-justified precautions. Again there is a U.S. parallel: Louisiana is a poor state and New Orleans, which has a large poor population, has a reputation for having an inefficient and corrupt government.

And fifth, the positive correlation of per capita income with value of life (see below) suggests that it is rational (though not always easy to explain as such) for even a well-governed poor country to devote proportionately fewer resources to averting calamities than rich countries do. This would also be true of a poor state or city in the United States.

Unfortunately, some of these political problems afflict not only poor countries, states, and cities; they afflict the mighty U. S. federal government as well, as the Hurricane Katrina debacle illustrates.[3] Set aside for a moment (I return

to the point later) the question whether the levees should have been strengthened or other measures taken to reduce the risk of a major flood; that is the analytical question. For current purposes the important point is that such measures were not taken. Hence the risk of such a flood was not eliminated; this was known (a 2002 series in the New Orleans *Times-Picayune* had explained the risk of a disastrous flood in the city in great detail[4]); and it followed that it might be necessary to respond to such a flood. Yet four years after 9/11, and two and a half years after the creation of the Department of Homeland Security, the federal government had yet to devise an executable plan for responding to a catastrophic event in New Orleans, or, I imagine, in any other threatened city in the United States. (I had thought Washington, D.C., was an exception, but it turns out it was not and is not.)[5]

Now, this failure *seems* incomprehensible. Planning an evacuation would not have been costly. It would not even have stepped on any big political toes. The need for emergency planning was not only apparent; it was explicitly acknowledged at every level of officialdom. So why did nothing happen? One reason I have already discussed is that a democratic (perhaps any) government is incapable of taking effective measures against novel threats. They do not have to be *really* novel; it is enough that no major American city had recently been inundated. The human mind, as already noted, has trouble thinking in terms of probabilities as distinct from frequencies, and politicians have foreshortened horizons (the probability of a disastrous flood in New Orleans was less than 10 percent over a period of thirty years, far longer than a politician's horizon). Policy myopia is thus built into democratic politics and is aggravated by the rapid turnover of appointed officials as they cycle between public sector and private sector jobs. An official who spends only two years in a job is unlikely to worry about what may happen decades hence. He will receive no current benefit from planning to deal with contingencies, however ominous, that seem to lie in the remote future.

Another obstacle to responding effectively to catastrophic risk is the pressure of the immediate. Officials are continuously harassed by members of Congress, the media, and White House staff to deal with the crisis du jour. They are not given leisure to address future contingencies, even to the extent of planning for them. Most of the offices in the federal government that are formally charged with conducting long-range planning do very little of it, in part because senior policymakers, with their truncated horizons and urgent distractions, do not pay serious attention to such planning efforts. The problem is aggravated by the sheer number of possible catastrophes, which makes it difficult to think systematically about responding.

Another obstacle, though this one is limited to after-the-fact as distinct from preventive responses, is cultural. Americans are not fatalists. They accept, for example, that national defense requires reserve forces and standby resources such as manned missile silos. But they would find it difficult to understand the use of government funds to establish a standby disaster command whose members sit around waiting for Seattle to be engulfed in a volcanic eruption by Mount Rainier or New Orleans to be inundated by a storm surge. Americans accept the inevitability of evil, but not of disaster.

And then there is the deficient *political* culture (in part a consequence of the social culture) that has produced the Department of Homeland Security in its current form. The creation of the department in 2003 was among a number of responses to the need to "do something" in the wake of the September 11, 2001, terrorist attacks. The need for better coordination of the numerous agencies responsible for protecting the nation's borders, and of the agencies that have responsibilities for responding to catastrophes, whether natural or man-made, was real enough. After all, America's three main border agencies—the Border Patrol, the Customs Service, and the Coast Guard—were located in three different executive departments (Justice, Treasury, and Transportation, respectively) and were incapable of effective coordination on any level. But that did not justify placing twenty-two agencies, including the Federal Emergency Management Agency (FEMA), a heretofore independent agency reporting directly to the president, in a gigantic new department. The department is centrally managed, hierarchical, with information flowing upward from the brontosaurus's tail to its tiny head (a handful of people, albeit some very able, trying to control more than 180,000 civil servants), and a response groping its way back down.

With the department's formation, FEMA was effectively demoted in the governmental hierarchy, losing much of its perceived importance in the process. Appointments to its senior managerial jobs could now be used to pay small political debts—and for the further reason that emergency response, though a challenging specialty, is not yet a formal profession like medicine or law, so there is a less definite sense of the proper credentials for the officials.[6] And now FEMA had to stand in line, waiting its turn for the attention of the beleaguered secretary of Homeland Security, who was struggling to assert control over his new far-flung domain. A plan formulated by FEMA for responding to large-scale catastrophe would have to be approved not only by the White House (which has its own Homeland Security Council, whose role in the response to Hurricane Katrina remains obscure), but also by the Secretary of Homeland Security. And the secretary was unlikely to be an expert in

emergency response, given the breadth of his responsibilities. But this meant that when disaster struck, the head of FEMA, an amateur in emergency response because the job was no longer considered very important, the agency having dropped a rung in the hierarchy of government agencies, had to consult a higher official, also an amateur in emergency response.

There was another problem with sticking FEMA in an immense new department. The department's emphasis was on fighting terrorism, and so preparing to deal with natural disasters got sidetracked, even though both natural and terrorist disasters can require similar responses.[7] This illustrates how placing a bureaucratic layer over heretofore independent agencies can undermine efforts to prevent catastrophes. The people at the top have a limited span of attention and control and so may be inclined to focus on a single mission, thereby curtailing the spectrum of risks that are dealt with.

What I am calling a deficient political culture is in part a result of the "do something" attitude of a nation of nonfatalists but probably in greater part a result of the interaction between a decentralized government structure designed in the eighteenth century and the enormous challenges to government thrust upon it by the complexity and diversity of modern America and its position in the world. The separation of the legislative and executive branches (which are effectively fused in a parliamentary system, such as that of the United Kingdom), aggressive judiciaries, and the distribution of government power among federal agencies, the states, and local governments make timely and coherent government action difficult at best, and perhaps impossible in dealing with subtle novel challenges.

Analytical Problems

The analytical problem of dealing effectively with catastrophic risks lies in the limitations of cost-benefit analysis. A cost-benefit analysis is the rational way to determine what if anything to do about the risk of something bad happening. In a simple analysis, first multiply the cost of the event if it materializes (say, a flood) by the probability that it will materialize if no (or no additional) preventive measures are taken. That will give the expected cost of the bad event. Next, calculate the cost of the measures necessary to prevent the event from occurring (that is, to eliminate the risk) and compare the two figures. If the expected cost of the event exceeds the cost of prevention (a cost measured by the value that the resources used for prevention would earn in their best alternative employment), adopt the measures. So, for example, say there is a 1 percent chance of a flood that would cause $1 billion in damages, making the expected cost of the

flood $10 million, and also assume the cost of averting the flood would be $9 million. As a first approximation, the preventive measure should be adopted. (It is only a first approximation because one must consider whether that $9 million might be invested even more productively elsewhere, of which more below.) A slight complication is that the measure may reduce rather than eliminate the risk; in that case, the reduction in expected cost is what is to be compared with the cost of (partial) prevention, that is, risk reduction.

Cost-benefit analysis of responding to catastrophic risk is often feasible. The flood caused by Hurricane Katrina is an example. In 1998 it was estimated that preventing such a flood would cost $14 billion; the estimated "economic" cost (that is, ignoring loss of life and physical and emotional suffering) of the flood was estimated at $100 billion–$200 billion;[8] and the Army Corps of Engineers estimated the annual probability of such a flood at 1 in 300.[9] Taking the lower cost and assuming that the $14 billion investment would eliminate the probability of a flood within thirty years—a period in which the probability of a flood (if the measures were not taken) would be a shade under 10 percent—yields an expected benefit from the flood-control measures of $10 billion. The proposed measures therefore flunked a cost-benefit test, since $10 billion is less than $14 billion.

They should not have flunked. The calculation of future benefits ignored the fact that the benefits are likely to grow—a flood that occurred thirty years hence would be likely to do more damage because property values would have increased—although these enhanced future benefits would have to be discounted to present value. Worse, the analysis ignored the expected loss of life, and other human suffering, that a massive flood would cause. There is a substantial economic literature inferring the value of life from the costs people are willing to incur to avoid small risks of death; if from behavior toward risk one infers that a person would pay $70 to avoid a 1-in-100,000 risk of death, his value of life would be estimated at $7 million ($70/.00001), which is in fact the median estimate of the value of life of an American. (These estimates are sensitive to incomes, as I noted earlier; the less money people have, the less they will allocate toward minimizing risks of death, which will automatically depress the measured value of life.)

The utility of this transformation is simply that, once a risk is calculated, its expected cost is instantly derived simply by multiplying the risk by the value of life. But a more intuitive way to understand the "value of life" estimates is as a summation of the value that people place on avoiding slight risks. A 1-in-100,000 risk of death implies that if 100,000 people are exposed to the risk, 1 will die. If each of the 100,000 would demand $70 to bear the 1-in-100,000 risk

of death, then the total demanded would be $7 million, and we can simply call this, for simplicity's sake, the value of the life of the 1 person who did die.

Now, the method just outlined—call it classic cost-benefit analysis—for dealing rationally with potential adversity unfortunately tends to break down in catastrophic situations. The stumbling blocks include difficulties in estimating probabilities (less commonly in estimating costs and benefits), in prioritizing risks, and in discounting for futurity; these are apart from the cognitive (and other psychological) and political factors discussed above.

Two common confusions need first to be dispelled. The first is that all catastrophic risks are low-probability events; the second is that public policy should focus on high-probability events, such as heart attacks.

Statisticians distinguish between risk and uncertainty. A risk is a contingent event to which a numerical probability can be assigned; a contingent event is uncertain if no probability can be assigned. When someone says that he is much more likely to die of a heart attack than be killed by a terrorist, he is implicitly assigning a low probability to a terrorist attack. That is a mistake; the probability of such an attack cannot be determined. Terrorist attacks have been infrequent causes of death, relative to heart attacks, in the past, but there is no basis for thinking that the future of terrorism will be similar to its past.

Even if it were known that terrorist attacks would continue to be infrequent and to inflict relatively limited damage (relative to the damage that would be inflicted, say, by a terrorist attack with weapons of mass destruction), it would not follow that resources should be reallocated from the struggle against terrorism to the struggle against heart disease. Probability is not the only factor in cost-benefit analysis. Considering the average age of heart-attack victims and the ability to reduce the likelihood of such an attack by modifying one's behavior, society may well already be spending the cost-justified amount of money (or more) on the prevention of heart attacks.

Some catastrophic risks can be quantified: the flooding of New Orleans is one example, and another is the range of possible asteroid collisions, about which a fair amount is known because of the long geological history of the earth and the moon. But for most catastrophes, risks cannot be estimated; some cannot even be bounded (except between 0 and 1!). Those are the cases of genuine uncertainty in the statistician's sense.

If one cannot estimate risk, one cannot do cost-benefit analysis. But that does not leave us completely stymied. Two remedial techniques merit more widespread use. One I call "inverse cost-benefit analysis." It involves calculating the *implied* risk from data on the cost of the catastrophe if it occurs and on the amount of money currently being spent to avert the risk of its occurring.

Catastrophe (deaths)	*Costs being spent to reduce risk*	*Loss if risk materalizes*	*Probability that it will materialize (implied)*
Bioterrorist attack (100 million)	$2 billion	$1 quadrillion (U.S. loss only)	.000002 (1 in 500,000)
Asteroid collision (1.5 billion)	$3.9 million	$3 quadrillion	.0000000013 (1 in 769 million)
Strangelet disaster in particle accelerator	$0	$600 trillion	0
Catastrophic global warming	$1.7 billion	$66.6 trillion (U.S. loss only)	.00000255 (1 in 388,000)

That is, if the government is spending $1 million on trying to prevent a catastrophe that, if it occurs, will cause damage of $100 million, the implication is that the government reckons the risk of the catastrophe as 1 percent or less, since the expected cost of a 1 percent risk of a $100 million loss is $1 million. If the real risk is higher and could be averted by an additional expenditure less than that added risk, the government is underspending. This explanation is a little oversimplified, but it will do to illustrate the potential utility of the technique. The table, taken from my book *Catastrophe*, shows how application of the technique strongly suggests government underspending.

Another technique for approximating cost-benefit analysis under conditions of uncertainty goes by the name "tolerable windows" and is illustrated in the figure. The marginal benefits (*mb*) and marginal costs (*mc*) of measures to reduce or eliminate some catastrophic risk are shown as functions of the quantity of precautions taken, with the optimal level of precautions (q^*) given by the intersection of the two functions. Suppose the optimum cannot be determined because of uncertainty about costs, benefits, the discount rate, or probabilities. Nevertheless enough may be known about the benefits and costs to be able to create the "window" formed by the two vertical lines.[10] Notice that

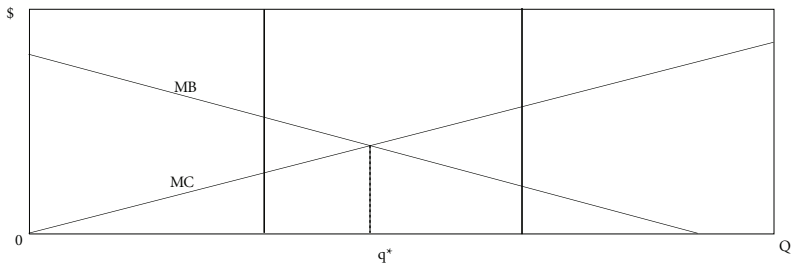

at the left side of the window frame, the benefits of a further effort to eliminate or prevent the catastrophe in question comfortably exceed the costs, while at the right side costs exceed benefits. Staying within the window gives some basis for confidence that the measures, while they may not be optimal, are neither grossly inadequate nor grossly excessive. This technique might be plausibly applied to the current funding of asteroid defense—which is probably well to the left of the left side of the window.

Using cost-benefit analysis to design optimal responses to catastrophic risks presents a priorities problem that arises from the infinite number of such risks. By discussing just a few, I have made the analytical task seem simpler than it is. Suppose the number of catastrophic risks worth worrying about, considering their magnitude and the feasibility of reducing them, is set at 100; even if preventive measures for each of them pass a cost-benefit test, it will be unclear how many of those measures should actually be adopted, since their benefits have to be compared with the benefits of alternative use of resources, uses that may have nothing to do with catastrophes. Policymakers could impoverish the country and the world by focusing excessively on eliminating catastrophic risks.

A partial answer is that the catastrophic risks tend to come in clusters from a prevention or response standpoint. Many of the measures for preventing or mitigating a bioterror attack will also prevent or mitigate a natural epidemic. An evacuation plan for a city such as Seattle, San Francisco, or New Orleans that is vulnerable to a natural disaster will also serve if the city is the victim of a WMD attack. Measures to reduce demand for oil will both alleviate (though I fear only slightly) global warming and reduce the likelihood of a catastrophic energy shortage resulting from our dependence on unstable and potentially hostile oil-exporting nations.

Finally as to the problem of discounting: If there is a temporal mismatch between costs and benefits, discounting will be required to enable a cost-benefit analysis to be conducted. The benefits of eliminating a catastrophic risk consist of the benefits of eliminating, say, a 1 percent chance of a flood this year, a 1 percent chance of a flood next year, and so on indefinitely. The probability of an event tends to rise with the interval over which the probability is assessed. The probability that an event with a 1 percent annual probability of occurring will in fact occur over the next century is not 100 percent but is close to it. This means that most of the expected cost of the event will lie in the future, perhaps the very distant future, even if the *annual* probability stays the same. If society values the future less than the present, the sum of the benefits of eliminating the risk will thus be smaller than simply the multiple of the first-

year benefits and the number of years the risk is expected to persist if nothing is done about it.

So the question becomes how heavily to weight expected costs that will be borne in the distant future. In noncommercial settings, as where one is trying to decide what weight to give to the expected cost of a tsunami in the third millennium, there is no intellectually satisfying answer. The only remotely satisfactory approach goes by the name "limited horizons."

To understand what is meant by "limited horizons," consider that the present value of an infinite stream of costs discounted at 4 percent is equal to the undiscounted sum of those costs for twenty-five years, while the present value of an infinite stream of costs discounted at 1 percent is equal to the undiscounted sum of those costs for one hundred years. So one might argue for the 4 percent rate (that is, for truncating concern for future welfare at twenty-five years) by saying that the current generation is willing to weight the welfare of the next generation as heavily as its own welfare, but that is the extent of its regard for the future. Or one might argue for the 1 percent rate by saying that the current generation is willing to give equal weight to the welfare of everyone living in this century—including themselves, their children, and their grandchildren—but beyond that they don't care. Looking at future welfare in this way, one may incline toward the lower rate—which would have dramatic implications for the willingness to invest today in limiting global warming. The lower rate could even be regarded as a ceiling.

After all, most people have some regard for human welfare, or at least the survival of some human civilization, in future centuries. We ought to be grateful that the Romans did not try to exterminate the human race in chagrin at the impending collapse of their empire. Since we owe future generations something like that consideration, we simply cannot afford in this day and age not to think about catastrophic risks. But thinking about them is very difficult, doing something practical about them even more so. Society therefore faces a great challenge, to which it had better rise. There cannot be any assurance that policymakers will heed cost-benefit analyses that reveal clear and substantial benefits from measures of prevention or mitigation of catastrophic risks. Such analyses are not a sufficient condition of wise policy, but they are probably a necessary condition.