2006

Efficient Responses to Catastrophic Risk

Richard A. Posner

Efficient Responses to Catastrophic Risk

Richard A. Posner*

The Indian Ocean tsunami of December 2004 has focused attention on a type of disaster to which policymakers pay too little attention—a disaster that has a very low or unknown probability of occurring but that if it does occur creates enormous losses. Great as the death toll, physical and emotional suffering of survivors, and property damage caused by the tsunami were, even greater losses could be inflicted by other disasters of low (but not negligible) or unknown probability. The asteroid that exploded above Siberia in 1908 with the force of a hydrogen bomb might have killed millions of people had it exploded above a major city. Yet that asteroid was only about two hundred feet in diameter, and a much larger one (among the thousands of dangerously large asteroids in orbits that intersect the earth's orbit) could strike the earth and cause the total extinction of the human race through a combination of shock waves, fire, tsunamis, and blockage of sunlight wherever it struck. Other catastrophic risks, besides earthquakes such as the one that caused the recent tsunami, include natural epidemics (the 1918–1919 Spanish influenza epidemic killed between twenty and forty million people), nuclear or biological attacks by terrorists, certain types of lab accidents (including one discussed later in this Article), and abrupt global warming. The probability of catastrophes resulting, whether or not intentionally, from human activity appears to be increasing because of the rapidity and direction of technological advances.

It is natural to suppose that the prediction, assessment, prevention, and mitigation of catastrophes is the province of science, but in this Article, which is based on my book *Catastrophe: Risk and Response*, I will argue that economic

---

* Richard A. Posner is a judge on the United States Seventh Circuit Court of Appeals, a senior lecturer at the University of Chicago law school, and the author of *Catastrophe: Risk and Response* (Oxford 2004). This Article is adapted from the text of a talk that was delivered at the AEI–Brookings Joint Center on Regulation on March 1, 2005. The author thanks Meghan Maloney for her excellent research assistance.

analysis has an important role to play, as well.[1] Able scientists can commit analytical errors when discussing policy that economists would easily avoid. Thus Barry Bloom, dean of the Harvard School of Public Health, has criticized the editors of leading scientific journals for having taken the position that "an editor may conclude that the potential harm of publication outweighs the potential societal benefits."[2] (The specific reference is to publications from which terrorists could learn how to create lethal biological weapons.) Bloom calls this "a chilling example of the impact of terrorism on the freedom of inquiry and dissemination of knowledge that today challenges every research university."[3] He appears to believe that freedom of scientific research should enjoy absolute priority over every other social value. Such a belief comes naturally to people who derive career advantages from being able to engage in a particular activity without hindrance, but this belief arbitrarily refuses to weigh costs and so consider the need to make tradeoffs.

Bloom is particularly incensed at limitations on allowing foreigners to study science in American universities. Under the rubric of "Advancing Openness," he advocates changes in existing regulations to enable any foreigner who obtains a visa for studying science in the United States to pursue any area of scientific research, however sensitive and whatever the student's likely motive.[4] Bloom's concern is understandable in terms of professional self-interest. American universities, and especially their graduate programs, are heavily dependent on foreigners. According to Bloom, 24 percent of the graduate students at Harvard are foreign.[5] But professional self-interest is not a sure guide to sound public policy.

I will begin my analysis of the catastrophic risk problem with the recent tsunami. Suppose that a tsunami as destructive as the one that struck the Indian Ocean occurs, on average, once a century and kills 250,000 people. That is an

---

1 For a general discussion, see Richard A. Posner, *Catastrophe: Risk and Response* 139–98 (Oxford 2004).

2 Barry R. Bloom, *Bioterrorism and the University: The Threats to Security—and to Openness*, Harvard Magazine 48, 51 (Nov–Dec 2003), quoting an uncited agreement among editors of thirty-two major biological science journals.

3 Id.

4 Id at 52 ("It is my view that anyone here for training, from any country, who is bright enough and promising enough to be admitted . . . and who is cleared by the government agency that has been set to screen applicants . . . ought to be able to pursue any area of study and research under faculty supervision in the university."). I take it that by "cleared by the government agency that has been set to screen applicants" Bloom simply means that the applicant has received a visa permitting him or her to study science, as distinct from having received a security clearance. Obviously the government cannot conduct a full background investigation of every foreign student who wants to study science in the United States.

5 Id.

average of twenty-five hundred deaths per 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 twenty-five hundred could be substantially reduced at a moderate 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. This is because the annual costs of protection must be matched with annual, not total, expected costs of tsunamis.

I can be a little more precise about how one might determine the costs of catastrophes. There is now 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 ($70/0.00001), which is in fact the median estimate of the value of life of an American.[6] The value of this transformation is simply that when we calculate a risk, we can immediately read off its cost by multiplying it by the value of life.

But there is significant non-linearity to be considered at both ends of the risk spectrum. At the high end, ask yourself what you would demand to play one round of Russian roulette; the answer is probably a good deal more than 1/6 of $7 million. At the other, low-probability end of the risk spectrum, there may be a tendency to write the cost of the risk down to or near zero. In other words, the studies from which the $7 million figure is derived may not be robust with respect to risks of death either much larger, or much smaller, than the 1 in 10,000 to 1 in 100,000 range of most of the studies—and we do not know what the risk of death from a tsunami was to the people who were killed, though it was probably toward the low end of the range.

Moreover, because value of life is positively correlated with income,[7] the $7 million figure cannot be used to estimate the value of life of the people killed by

---

6 W. Kip Viscusi and Joseph E. Aldy, *The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World*, 27 J Risk & Uncertainty 5, 63 (2003).

7 Id at 40.

the Indian Ocean tsunami—or at least most of them.[8] Additional complications arise from the fact that the deaths were only a part of the cost inflicted by the disaster. The injuries, the suffering, and the property damage that resulted from the tsunami have to be estimated in conjunction with the efficacy and expense of precautionary measures that would have been feasible. In addition, the risks of smaller but still destructive tsunamis that such measures might protect against must be factored in; nor am I confident about my "once a century" risk estimate. Nevertheless, it seems apparent that the total cost of the recent tsunami is high enough to indicate that precautionary measures would have been cost-justified.

Why, then, were such measures not taken in anticipation of a tsunami on the scale that occurred? Tsunamis are a common consequence of earthquakes, which are themselves common, and tsunamis can have causes besides earthquakes—a major asteroid strike in an ocean would create a tsunami that would dwarf the one in the Indian Ocean. Again, economics can yield some useful insights.

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 at some time in the first decade of the century than at some time in the last nine decades of the century. (The point is simply that the probability is greater the longer the interval being considered; one is more likely to catch a cold in the next year than in the next forty-eight hours.) Politicians with limited terms of office, and thus foreshortened political horizons, are likely to discount low risk disaster possibilities steeply since the risk of harm 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, officials are likely to have a high threshold of attention below which risks are simply ignored.

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 richer and larger countries. Knowing this, the

---

8 Viscusi and Aldy have estimated that the income elasticity of the value of a statistical life is between 0.5 and 0.6, meaning that a 1 percent increase in income is associated with a 0.5–0.6 percent increase in the value of life. Id. The per capita GDP of the United States is about 11.5 times the per capita GDP of Indonesia (one of the countries devastated by the tsunami). Thus, if the Viscusi-Aldy study is correct, the value of life in Indonesia is likely to be no more than about $1 million. GDP figures are from the CIA World Factbook, available online at <http:// www.cia.gov/cia/publications/factbook/index.html> (visited Nov 27, 2005).

latter countries may be reluctant to take precautionary measures that would reward and thus encourage free riding.

Fourth, countries are often poor because they are run by weak, inefficient, or corrupt governments, and such governments may disable poor nations from taking cost-justified precautions.

And fifth, the positive correlation of per capita income with value of life suggests that it is quite rational for even a well-governed poor country to devote proportionately less resources to averting calamities than a rich country.

An even more dramatic example of neglect of low-probability/high-cost risks concerns the asteroid menace, which is analytically similar to the menace of tsunamis. NASA, with an annual budget of more than $10 billion, spends only $4 million a year mapping dangerously close large asteroids.[9] At that rate, NASA may not complete the task for another decade, even though such mapping is the key to an asteroid defense that may give us years of warning. Deflecting an asteroid from its orbit when it is still hundreds of millions of miles away from hitting the Earth appears to be a feasible undertaking.[10] Although asteroid strikes are less frequent than tsunamis, there have been enough of them to enable the annual probabilities of various magnitudes of such strikes to be estimated, and from these estimates, an expected cost of asteroid damage can be calculated.[11] As in the case of tsunamis, if there are measures, beyond those being taken already, that can reduce the expected cost of asteroid damage at a lower cost, thus yielding a net benefit, the measures should be taken, or at least seriously considered. Later I will show that such an analysis indicates that NASA should be spending much more on asteroid mapping.

Often it is not possible to estimate the probability or magnitude of a possible catastrophe; how then can cost-benefit analysis, or other techniques of economic analysis, help us in devising responses to such a possibility? Well, it, and they, can. The probability of bioterrorism or nuclear terrorism, for example,

---

9 NASA Office of Space Science, *NASA Announcements Opportunity, Appendix A.2.8: Near-Earth Object Observations*, available online at <http://research.hq.nasa.gov/code_s/nra/current/nra-03-oss-01/appendA2.html> (visited Oct 28, 2005); Hearing on Near Earth Objects (NEO) before the Subcommittee on Science, Technology, and Space of the Senate Committee on Commerce, Science, and Transportation, 108th Cong, 2d Sess (Apr 7, 2004) (statement of Dr. Lindley Johnson, Program Manager, Near Earth Objects Observation Program, National Aeronautics and Space Administration), available online at <http://commerce.senate.gov/hearings/testimony.cfm?id=1147&wit_id=3241> (visited Oct 28, 2005).

10 See Edward T. Lu and Stanley G. Love, *Gravitational Tractor for Towing Asteroids*, 438 Nature 177, 177–78 (2005); Russell L. Schweickart, et al, *The Asteroid Tugboat*, Scientific American 54, 54–61 (Nov 2003).

11 See Harry Atkinson, Crispin Tickell, and David Williams, eds, *Report of the Task Force on Potentially Hazardous Near Earth Objects* 20 (British Natl Space Centre 2000), available online at <http://www.nearearthobjects.co.uk/report/resources_task_intro.cfm> (visited Oct 28, 2005).

cannot be quantified, but we have some sense of the range of possible losses that such terrorism would inflict (there really is no upper limit short of the extinction of the human race). We can infer from this that even if the probability of such a terrorist attack is small, the expected cost—the product of the probability of the attack and of the consequences if the attack occurs—probably is quite high. Then we can focus on trying to quantify the costs of remedial measures, which in the case of bioterrorism include stocking vaccines and, *pace* Barry Bloom, limiting access of foreign students to lethal pathogens in US university laboratories.

Let me give two examples of the use of economic analysis to frame public policy toward an unquantifiable risk. The first involves global warming. The Kyoto Protocol ("Protocol"),[12] which recently came into effect by its terms when Russia signed it, though the United States has not, requires signatory nations to reduce their carbon dioxide emissions to a level 7 to 10 percent below what they were at 6 years ago,[13] but exempts developing countries, such as China, a large and growing emitter, and Brazil, which is destroying large reaches of the Amazon rain forest, much of it by burning.[14] The effect of carbon dioxide emissions on the atmospheric concentration of the gas is cumulative because carbon dioxide leaves the atmosphere (by being absorbed into the oceans) at a much lower rate than it enters the atmosphere. Therefore, even if the annual rate of emission is cut substantially, carbon dioxide concentration will continue to grow.[15] Between this phenomenon and the exemptions, there is a widespread belief that the Protocol will only have a slight effect in arresting global warming; yet the tax or other regulatory measures required to reduce emissions below their 1990 level will be very costly.[16]

12 Kyoto Protocol to the United Nations Framework Convention on Climate Change (1997), 37 ILM 22 (1998) (hereinafter Kyoto Protocol).

13 William D. Nordhaus and Joseph G. Boyer, *Requiem for Kyoto: An Economic Analysis of the Kyoto Protocol*, Energy J 93 (Special Issue 1999).

14 Bruce Yandle, *The Precautionary Principle as a Force for Global Political Centralization: A Case-Study of the Kyoto Protocol*, in Julian Morris, ed, *Rethinking Risk and the Precautionary Principle* 167, 170 (Butterworth-Heinemann 2000).

15 Posner, *Catastrophe* at 51 (cited in note 1).

16 Murray Weidenbaum, *Scientific Uncertainties and Policy Controversies over Global Warming*, USA Today 12 (July 2005); Richard Black, *Giant Leaps Needed Post-Kyoto*, BBC News (Feb 16, 2005), available online at <http://news.bbc.co.uk/1/hi/sci/tech/4268707.stm> (visited Nov 22, 2005); Martin I. Hoffert, et al, *Advanced Technology Paths to Global Climate Stability: Energy for a Greenhouse Planet*, 298 Science 981 (2002); *Blair Issues Global Warming Challenge*, BBC News (Sept 1, 2002), available online at <http://news.bbc.co.uk/1/hi/uk_politics/2228741.stm> (visited Nov 22, 2005); Richard Black, *Climate Treaty's 'Minimal' Impact*, BBC News (Nov 8, 2001), available online at <http://news.bbc.co.uk/1/hi/sci/tech/1646029.stm> (visited Nov 22, 2005).

The Protocol could certainly be improved, but I consider it on balance a significant step in the right direction—at least if the United States ratifies it. (The United States has thus far refused to ratify the Protocol and shows no signs of doing so in the future.) But my reasoning differs from that of most of the Protocol's supporters. They are content to slow the rate of global warming through heavy taxes (for example, on gasoline or coal) or other measures (such as quotas) that will make fossil fuels more expensive to consumers, and thereby encourage conservation measures, such as driving less or driving more fuel-efficient cars.[17] This is either too much or too little. It is too much if, as most scientists believe, global warming will continue to be a gradual process producing really serious effects—including the destruction of tropical agriculture, the spread of tropical diseases such as malaria to currently temperate zones, dramatic increases in violent storm activity (increased atmospheric temperatures, by increasing the amount of water vapor in the atmosphere, increases precipitation), and a rise in sea levels (eventually to the point of inundating most coastal cities)—only toward the end of the century.[18] For by that time, science, without prodding by governments, is likely to have developed economical "clean" substitutes for fossil fuels (we already *have* a clean substitute—nuclear power) and economic technology that will either prevent carbon dioxide from being emitted into the atmosphere by the burning of fossil fuels or will remove carbon dioxide from the atmosphere. Furthermore, if the focus is changed from gradual to abrupt global warming, the Protocol, at least without the participation of the United States, is too limited a response to global warming.

At various times in the Earth's history, drastic temperature changes have occurred in the course of just a few years. The most recent of these periods, called the Younger Dryas (Dryas is a flower that flourished then), took place about eleven thousand years ago, shortly after the end of the last Ice Age. In a period of no more than a decade, global temperatures soared by about 14 degrees Fahrenheit.[19] Because the earth was still cool from the Ice Age, the effect of the increased warmth on the human population was positive. But a similar increase in a modern decade would have devastating effects on

17 Posner, *Catastrophe* at 156 (cited in note 1).

18 See Bjørn Lomborg, *The Skeptical Environmentalist: Measuring the Real State of the World* 288–90 (Cambridge 1998); Samuel Fankhauser, *Valuing Climate Change: The Economics of the Greenhouse* 27–28 (Earthscan 1995); John F.B. Mitchell, *General Circulation Modeling of the Atmosphere*, in M.E. Schlesinger, ed, *Climate-Ocean Interaction* 67, 80 (Kluwer 1990).

19 Steven Mithen, *After the Ice: A Global Human History, 20,000–5,000 BC* 12–13 (Weidenfeld & Nicolson 2003); National Research Council of the National Academies, *Abrupt Climate Change: Inevitable Surprises* 27 (Natl Acad 2002).

agriculture and on coastal cities, and might even cause a shift in the Gulf Stream that would result in giving all of Europe a Siberian climate.[20]

Because of the enormous complexity of the forces that determine climate, and the historically unprecedented magnitude of human effects on the concentration of greenhouse gases, one cannot exclude the possibility that continued growth in greenhouse gas concentration could precipitate a sudden warming similar to that of the Younger Dryas—and within the near rather than the distant future. Indeed, no probability, high or low, can be assigned to such a catastrophe. But it may be significant that while dissent continues, many climate scientists are now predicting dramatic effects from global warming within the next twenty to forty years, rather than just by the end of the century.[21] It may be prudent, therefore, to try to stimulate increases in development rates for economical substitutes for fossil fuels, technology that limits the emission of carbon dioxide by fossil fuels when they are burned in internal combustion engines or electrical generating plants, and technology for removing carbon dioxide from the atmosphere. This can be accomplished through stiff taxes on carbon dioxide emissions. Such taxes give the energy industries and their customers (such as airlines and motor vehicle manufacturers) a strong incentive to finance research and development ("R&D") designed to create economical clean substitutes for fossil fuels and devices to "trap" emissions at their source, before they enter the atmosphere. Given the technological predominance of the United States, it is important that these taxes be imposed on US firms, which would be the case if we ratified the Protocol.

One advantage of the technology-forcing tax approach over public subsidies for R&D is that the government would not be in the business of selecting winners—the affected industries would decide what R&D to support. Another advantage is that the brunt of the taxes could be partly offset by reducing other taxes since emission taxes would raise revenue while inducing greater R&D expenditures.

It might seem that subsidies would be necessary for technologies that would have no market, such as technologies for removing carbon dioxide from the atmosphere. There would be no private demand for such technologies because, in contrast to ones that reduce emissions, technologies that remove already emitted carbon dioxide from the atmosphere would not reduce any

---

20 Alexander E. MacDonald, *The Wild Card in the Climate Change Debate*, Issues Sci & Tech 51, 52 (Summer 2001) (predicting a potential seven degree drop in average temperature across Europe if the Gulf Stream shifted).

21 Edward W. Lempinen, *Scientists on AAAS Panel Warn That Ocean Warming Is Having Dramatic Impact*, AAAS News Release (Feb 21, 2005), available online at <http://www.aaas.org/news/releases/2005/0217warmingwarning.shtml> (visited Oct 28, 2005).

emitter's tax burden. But this problem is easily solved by applying the tax to *net* emissions. An electricity generating plant or other emitter could then reduce its tax burden by removing carbon dioxide from the atmosphere as well as by reducing its own emissions of carbon dioxide into the atmosphere.

My second example of how economic analysis can produce insights even when catastrophic risks are non-quantifiable involves the Relativistic Heavy Ion Collider ("RHIC") that went into operation at Brookhaven National Laboratory in 2000. As explained by the distinguished English physicist Sir Martin Rees, collisions in RHIC might conceivably produce a shower of quarks that would:

> reassemble themselves into a very compressed object called a strangelet. . . . [A] strangelet could, by contagion, convert anything else it encountered into a strange new form of matter. . . . [A] hypothetical strangelet disaster could transform the entire planet Earth into an inert hyperdense sphere about one hundred metres across.[22]

Rees considers this "hypothetical scenario" exceedingly unlikely, yet points out that even a probability of one in fifty million is not wholly negligible when the result, should the improbable materialize, would be so total a disaster.[23]

Concern with such a possibility led John Marburger, the director of the Brookhaven National Laboratory and now the President's science advisor, to commission a risk assessment by a committee of physicists chaired by Robert Jaffe before authorizing RHIC to begin operating.[24] In a synopsis of the assessment, Marburger offered the following lucid summary of the strangelet doomsday scenario:

> All particles ever observed to contain "strange" quarks have been found to be unstable, but it is conceivable that under some conditions stable strangelets could exist. If such a particle were also negatively charged, it would be captured by an ordinary nucleus as if it were a heavy electron. Being heavier, it would move closer to the nucleus than an electron and eventually fuse with the nucleus, converting some of the "up" and "down" quarks in its protons and neutrons, releasing energy, and ending up as a

22 Martin Rees, *Our Final Hour: A Scientist's Warning: How Terror, Error, and Environmental Disaster Threaten Humankind's Future in this Century—On Earth and Beyond* 120–21 (Basic Books 2003). See also Ivan Carvalho, *Dr. Strangelet or: How I Learned to Stop Worrying and Love the Big Bang*, Wired 254 (May 2000).

23 Rees, *One Final Hour* at 12 (cited in note 22).

24 Bruce Lambert, *Lab's Chicken Littles Will Be Disappointed*, NY Times 14LI.2 (Oct 17, 1999); Brookhaven National Laboratory, *Statement on Committee Review of Speculative "Disaster Scenarios" at Brookhaven Lab's RHIC*, available online at <http://www.bnl.gov/rhic/disaster.htm> (visited Nov 22, 2005).

> larger strangelet. If the new strangelet were negatively charged, the process could go on forever.[25]

That is, the strangelet would keep growing until all matter was converted to strange matter.

In my book, I attempt a cost-benefit analysis of RHIC with the strangelet disaster possibility factored in on the cost side. RHIC cost $600 million to build[26] and its annual operating costs were expected to be $130 million.[27] I did not previously try to monetize the expected benefits of experiments conducted at RHIC, but to begin the analysis, I make a wild guess that the benefits can be valued at $250 million per year. Second, I make the extremely conservative estimate, which biases the analysis in favor of RHIC's passing a cost-benefit test, that the cost of the extinction of the human race would be $600 trillion and that the annual probability of a strangelet disaster at RHIC is 1 in 10 million.

I grant that this probability estimate is arbitrary, which is why I consider the RHIC case as one of non-quantifiable risk. The physicist Arnon Dar and his colleagues have estimated the probability of a strangelet disaster during RHIC's planned ten-year life as no more than one in fifty million, which on an annual basis would be roughly one in five hundred million.[28] Robert Jaffe and his colleagues, the official risk-assessment team for RHIC, have offered a series of upper-bound estimates, including a one in five hundred thousand probability of a strangelet disaster over the ten-year period, which translates into an annual probability of approximately one in five million.[29] A 1 in 10 million estimate yields an annual expected extinction cost of $60 million for 10 years to add to the $130 million in annual operating costs and the initial investment of $600 million—and with the addition of that expected cost, the total costs of the project exceed its benefits.

Of course, this conclusion could easily be reversed by increasing the estimate of the project's benefits from my "wild guess" figure of $250 million. But here is the really interesting economic point: it is unclear whether RHIC will yield any social benefits and whether, if it will, the federal government should

---

25 John Marburger, *Synopsis of the Committee Report on Speculative "Disaster Scenarios" at RHIC*, available online at <http://www.phys.utk.edu/rhip/Articles/RHICNews/BNL_rhicreport.html> (visited Oct 28, 2005).

26 David Voss, *Making the Stuff of the Big Bang*, 285 Science 1194, 1194 (1999).

27 Brookhaven National Laboratory, *The Relativistic Heavy Ion Collider (RHIC): A Premier Facility for Nuclear Physics Research*, available online at <http://www.bnl.gov/bnlweb/PDF/Factsheet/FS-RHIC.pdf> (visited Oct 28, 2005).

28 Arnon Dar, A. De Rújula, and Ulrich Heinz, *Will Relativistic Heavy-Ion Colliders Destroy Our Planet?*, 470 Physics Letters B 142, 146 (1999).

29 R.L. Jaffe, et al, *Review of Speculative "Disaster Scenarios" at RHIC*, 72 Rev Modern Physics 1125, 1138 (2000).

subsidize particle accelerator research. As I understand it, the purpose of RHIC is not to produce useful products, as earlier particle accelerator research undoubtedly did, but to yield insights into the earliest history of the universe. In other words, the purpose is to quench scientific curiosity. Obviously, this research benefits scientists, or at least high-energy physicists. But how does such research benefit society as a whole? And why, in any event, should it be paid for by the taxpayer instead of being financed by the universities that employ the physicists who are interested in conducting theoretical research? I have not seen or heard any satisfactory answers to these questions. If there are no good answers, the fact that such research poses even a slight risk of global catastrophe becomes a compelling argument against its continued subsidization.

Returning to global warming for a moment, notice how by focusing on the risk of abrupt global warming we can largely elide the vexing problem of discount rate. But I do not want to ignore this problem, which is particularly acute when concern focuses on gradual global warming. Suppose that a $10 billion expenditure on capping emissions today would have no effect on human welfare during this century but, by slowing global warming, would produce a savings in social costs of $100 billion in 2100. At a discount rate of 3 percent, the present value of $100 billion a century from now is only $5 billion. That figure would make the expenditure of $10 billion today seem like a very poor investment. (For the sake of simplicity I ignore benefits that are expected to accrue after 2100.) The same amount of money invested in financial instruments could be expected to grow to $192 billion by 2100, assuming a 3 percent real interest rate for the next 100 years (though in fact interest rates cannot be forecast over such a lengthy period). If the fund were then disbursed to the victims of global warming, they would be better off than if the $100 billion cost of global warming assumed to be incurred in that year had been averted. Less conservative investments, moreover, would yield larger expected returns—10 percent or more rather than 3 percent.

But it is not actually feasible to invest $10 billion in a fund for future victims of global warming in lieu of spending the money now. No such fund will be created and the victims will not be compensated. In circumstances such as this one, discounting future to present values is not a method of helping people to decide how to manage their affairs in the way most conducive to maximizing their welfare. Rather, it is a method of maximizing global wealth without regard to its distribution among persons. In the case of gradual global warming, the victims are likely to be concentrated in poor countries. Thus, basing policy on the discounted costs of global warming would further immiserate the future inhabitants of those countries by increasing the authorized level of emissions harmful to them.

A discount rate based on market interest rates tends to obliterate the interests of remote future generations. The implications are drastic. At a discount rate of 5 percent, one death next year counts for more than one billion deaths in 500 years. On this view, catastrophes in the distant future can now be regarded as morally trivial. (What right would the Romans have had to regard our lives as worthless in deciding whether to conduct dangerous experiments?) The tradeoff is only slightly less extreme if one substitutes 100 years for 500. At a 5 percent discount rate, the present value of $1 to be received in 100 years is only three-quarters of a cent—and if for money we substitute lives, then to save one life this year we should be willing to sacrifice almost 150 lives a century hence.

But to refuse to discount future costs at all would be absurd, certainly as a practical political matter. For then the present value of benefits conferred on our remote descendants would approach infinity. Measures taken today to arrest global warming would confer benefits not only in 2100, but in every subsequent year, perhaps for millions of years. The present value of $100 billion received every year for a million years at a discount rate of 0 percent is $100 quadrillion, which is more than even Greenpeace wants spent on limiting emissions of greenhouse gases.

But maybe the vexing problem of how much weight to accord the welfare of remote future generations can be finessed, at least to some extent, if not solved. A discounted present value can be equated to an undiscounted present value simply by shortening the time horizon for the consideration of costs and benefits. For example, the present value of an infinite stream of costs discounted at 4 percent a year is equal to the undiscounted sum of those costs for 25 years while the present value of an infinite stream of costs discounted at 1 percent a year is equal to the undiscounted sum of those costs for 100 years. The formula for the present value of $1 per year forever is $1/*r*, where *r* is the discount rate. So if *r* is 4 percent, the present value is $25, and this is equal to an undiscounted stream of $1 per year for 25 years. If *r* is 1 percent, the undiscounted equivalent is 100 years.

One way to argue for the 4 percent rate (that is, for truncating our concern for future welfare at 25 years) is to say that we are willing to weight the welfare of the next generation as heavily as our own welfare, but that is the extent of our regard for the future. One way to argue for the 1 percent rate is to say that we are willing to give equal weight to the welfare of everyone living in this century, which will include us, our children, and our grandchildren, but that we are indifferent beyond that point. Looking at future welfare in this way, we may be inclined toward the lower rate—which would have dramatic implications for our willingness to invest today in limiting global warming. The lower rate could even be regarded as a ceiling. Most of us have some regard for human welfare, or at least the survival of some human civilization, in future centuries. We are grateful that the Romans did not exterminate the human race in chagrin at the impending collapse of their empire.

Another way to bring future consequences into focus without conventional discounting is by aggregating risks over time rather than expressing them in annualized terms. If we are concerned about what may happen over the next century, then instead of asking what the annual probability of a collision with a ten kilometer-wide asteroid is, we might ask what the probability is that such a collision will occur within the next one hundred years. An annual probability of one in seventy-five million translates into a century probability of roughly one in seven hundred and fifty thousand. That may be high enough—considering the consequences if the risk materializes—to justify spending several hundred million, perhaps even several billion, dollars to avert it.

A helpful approach to cost-benefit analysis under conditions of extreme uncertainty is what I shall call "inverse cost-benefit analysis." Analogous to extracting probability estimates from insurance premiums, this approach involves dividing what the government is spending to prevent a particular catastrophic risk from materializing by what the social cost of the catastrophe would be if it did materialize. The result is an approximation of the implied probability of the catastrophe. Remember that expected cost is the product of probability and consequence (loss): $C = PL$. If P and L are known, C can be calculated. If instead C and L are known, P can be calculated: if \$1 billion (C) is being spent to avert a disaster that if it occurs will impose a loss (L) of \$100 billion, then $P = C/L = 0.01$.

If P so calculated diverges sharply from independent estimates of it, this is a clue that society may be spending too much or too little on avoiding L. It is just a clue because of the distinction between marginal and total costs and benefits. The optimal expenditure on a measure is the expenditure that equates marginal cost to marginal benefit. Suppose we happen to know that P is not 0.01 but 0.1, so that the expected cost of the catastrophe is not \$1 billion but \$10 billion. It does not follow that we should be spending \$10 billion, or indeed anything more than \$1 billion, to avert the catastrophe. Maybe spending just \$1 billion would reduce the expected cost of the catastrophe from \$10 billion all the way down to \$500 million and no further expenditure would bring about a further reduction, or at least a cost-justified reduction. For example, if spending another \$1 billion would reduce the expected cost from \$500 million to zero, that would be a bad investment, at least if risk aversion is ignored.

The federal government is spending about \$2 billion a year to prevent a bioterrorist attack. (However, the president has requested another \$2.5 billion for 2005 under the rubric of "Project BioShield."[30]) I say "about \$2 billion" because while the \$2.6 billion that the president sought from Congress for 2004

30 Philip Shenon, *The President's Budget Proposal: Domestic Security; Plan Seeks Big Rise in Antiterror Spending*, NY Times A15 (Feb 3, 2004).

for combating "catastrophic threats" covers chemical, nuclear, and radiological threats as well as bioterrorism, the emphasis is on the last.[31] The goal is to protect Americans, so, in assessing the benefits of this expenditure, I shall ignore casualties in other countries. Suppose the most destructive biological attack that seems reasonably possible—on the basis of what little we now know about terrorist intentions and capabilities—would kill one hundred million Americans. We know that value-of-life estimates may have to be radically discounted when the probability of death is exceedingly slight. But there is no convincing reason for supposing the probability of such an attack to be less than, say, 1 in 100,000; and we know (well, think) that the value of life that is derived by dividing the cost that Americans will incur to avoid a risk of death of that magnitude by the risk is about $7 million. Then if the attack occurred, the total costs would be $700 trillion—and that is actually too low an estimate because the death of a third of the population would have all sorts of collateral consequences, mainly negative. Let us, still conservatively however, refigure the total costs as $1 quadrillion. The result of dividing the money being spent to prevent such an attack, $2 billion, by $1 quadrillion is 1/500,000. Is there only a 1 in 500,000 probability of a bioterrorist attack of that magnitude in the next year? One does not know for certain, but the figure seems too low.

It does not follow that $2 billion a year is too little to be spending to prevent a bioterrorist attack; we must not overlook the distinction between total and marginal costs. Suppose that the $2 billion expenditure reduces the probability of such an attack from 0.01 to 0.0001. The expected cost of the attack would still be very high—$1 quadrillion multiplied by 0.0001 is $100 billion—but spending more than $2 billion might not reduce the residual probability of 0.0001 at all. For there might be no feasible further measures to take to combat bioterrorism, especially when we remember that increasing the number of people involved in defending against bioterrorism—including not only scientific and technical personnel, but also security guards in laboratories where lethal pathogens are stored—also increases the number of people capable, alone or in conjunction with others, of mounting biological attacks. But there *are* other response measures that should be considered seriously. And one must also bear in mind that expenditures used to combat bioterrorism do more than prevent mega-attacks; the lesser attacks, which would still be very costly, both singly and cumulatively, would also be prevented.

Costs, moreover, tend to be inverse to time. It would cost a great deal more to build an asteroid defense in one year than in ten years because of the

31 Office of Management and Budget, *2003 Report to Congress on Combating Terrorism* 37 (Sept 2003), available online at <http://www.whitehouse.gov/omb/inforeg/2003_combat_terr.pdf> (visited Oct 28, 2005).

extra costs that would be required for a hasty reallocation of the necessary labor and capital from the current projects in which they are employed. And this inverse relationship between costs and time would also apply to other crash efforts to prevent catastrophes. Placing a lid on current expenditures would have the incidental benefit of enabling additional expenditures to be deferred to a time when, because more will be known about both the catastrophic risks and the optimal responses to them, considerable cost savings may be possible. The case for such a ceiling derives from comparing marginal benefits to marginal costs; the latter may be sharply increasing in the short run.

To conclude, catastrophic risks—in the sense of low-probability events that if they occur will inflict catastrophic harm—are, despite their low probability, well worth the careful attention of policymakers. There are, however, a variety of psychological and political obstacles to such attention. In addition, there is a sense that the uncertainties surrounding catastrophic risks are so great as to make such risks analytically intractable. My purpose in this Article has been to contest that sense. There are a variety of useful analytical techniques for dealing with catastrophic risks; greater use of those techniques would enable a rational response to those risks.

CJIL