

Cost-Benefit Analysis and the Environment*

Cass R. Sunstein

In the United States, cost-benefit analysis (CBA) is in the ascendancy. For over twenty years, American presidents have required agencies to perform CBA for major regulations; indeed, they have told agencies to regulate only if the benefits of regulation justify its costs.¹ Congress has also shown considerable interest in CBA, most prominently in the Safe Drinking Water Act, which asks agencies to produce quantitative assessments of both costs and benefits. For their part, federal courts have adopted a series of principles that promote CBA, saying that if Congress has not been clear, agencies may consider costs, take account of the substitute risks introduced by regulation, and exempt trivial risks from governmental control.

In its enthusiasm for cost-benefit analysis, the United States provides a sharp contrast to Europe, which has shown intense interest in a quite different organizing principle for environmental protection: the Precautionary Principle.² According to the Precautionary Principle, regulation is required even in the face of scientific uncertainty—even if it is not yet clear that environmental risks are serious. A central point of the Precautionary Principle is to recognize the limitations of existing knowledge and to protect against harm that cannot yet be established as such.

Cost-benefit analysis and the Precautionary Principle can lead in

* I am grateful to Elizabeth Emens, Charles Larmore, Martha Nussbaum, and Richard Posner for extremely valuable comments on an earlier draft of this article.

1. The developments discussed in this paragraph are traced in Cass R. Sunstein, *Risk and Reason* (Cambridge: Cambridge University Press, 2002).

2. See, e.g., Arie Trouwborst, *Evolution and Status of the Precautionary Principle in International Law* (London: Kluwer Law International, 2002); Poul Harremoës et al., eds., *The Precautionary Principle in the 20th Century: Late Lessons from Early Warnings* (London: Earthscan, 2002).

radically different directions. For example, many Europeans argue that the consequences of genetic modification are uncertain, that real harm is possible, and hence that stringent regulation is readily justified. By contrast, many Americans respond that the likely benefits of genetic modification are far greater than the likely harms and that stringent regulation is therefore unsupportable. Or consider global warming. Many European leaders have argued in favor of precautions, even extremely expensive ones, simply to reduce the risk of catastrophe. But under President George W. Bush, American officials have called for continuing research on the costs and benefits of controlling greenhouse gas emissions.

The tension between CBA and the Precautionary Principle raises serious questions about the theory and practice of environmental protection. To engage in cost-benefit analysis, regulators must make difficult and often speculative judgments about the likely effects of alternative regulatory strategies.³ The easiest task is often the identification of costs, but even here there are formidable empirical problems. It is difficult to project the expense of regulations of different levels of stringency, especially because environmental protection often spurs technological innovation, greatly reducing the cost of pollution reduction. The identification of benefits presents even harder empirical problems—and knotty normative and conceptual ones as well. At a minimum, agencies must estimate the savings that are likely to result from regulation, including reductions in mortality and morbidity, along with improvements in terms of visibility, recreation, aesthetics, animal welfare, property values, and more. When science leaves room for doubt, as it often does, agencies typically specify a range of possibilities, representing low-end estimates and high-end estimates in addition to the best point estimate. Agencies might, for example, project that a certain regulation will save as many as eighty lives each year and as few as zero, with a preferred estimate of twenty-five.⁴ These numbers inevitably involve a degree of guesswork.

After specifying the likely benefits, CBA requires agencies to engage in multiple acts of conversion, assigning economic values to human lives, human morbidity, and a range of harms to the environment. Typically American agencies assign monetary values on the basis of private “willingness to pay” (WTP).⁵ For example, the Environmental Protection

3. The Office of Management and Budget has issued guidelines to govern and to standardize the use of cost-benefit balancing. See Office of Management and Budget, “Regulatory Analysis,” Circular A-4 (September 17, 2003), available at <http://www.whitehouse.gov/omb/inforeg/regpol.html#rr>.

4. See Cass R. Sunstein, “The Arithmetic of Arsenic,” in *Risk and Reason*, pp. 153–90.

5. See W. Kip Viscusi, *Fatal Tradeoffs* (New York: Oxford University Press, 1993).

Agency (EPA) values a human life at about \$6.1 million, a figure that comes from real-world markets.⁶ In the workplace and for consumer goods, additional safety has a price; market evidence is investigated to identify that price. The \$6.1 million figure, known as the value of a statistical life (VSL), is mostly a product of studies of actual risks from the workplace and consumer products, attempting to determine how much workers and consumers are paid to assume mortality hazards. Suppose that people must be paid \$600, on average, to eliminate risks of 1/10,000; suppose, for example, that workers who face risks of that magnitude generally receive \$600 in additional wages each year. If so, the VSL would be said to be \$6 million. Where market evidence is unavailable, agencies often produce monetary valuations on the basis of contingent valuation surveys, which ask people how much they are willing to pay to save coral reefs or endangered species, to eliminate a risk of chronic bronchitis or curable lung cancer, and much more. Drawing on market evidence and contingent valuation studies, the EPA has recently valued a case of chronic bronchitis at \$260,000, an emergency hospital visit for asthma at \$9,000, hospital admission for pneumonia at \$13,400, a lost work-day at \$83, and a specified decrease in visibility at \$14.⁷

Once a CBA is produced, what should be done with it? The most ambitious answer is that agencies should adopt regulations only when the likely benefits exceed the likely costs—and that if several regulations meet this test, agencies should select the one that “maximizes net benefits.” On this approach, CBA provides a clear rule of decision, one by which regulators should be bound. A more cautious response would be that agencies should generally require benefits to exceed costs, and also seek to maximize net benefits, but that they need not do so; on this view, the outcome of the CBA provides a presumption but no more. The presumption could be rebutted by showing that the particular situation justifies a departure from the result indicated by CBA—as, plausibly, in cases in which poor people would stand to gain a great deal. A still more cautious approach would be that in deciding what to do, agencies should consider the outcome of CBA alongside such other variables as they deem relevant. There are important differences between those who would make CBA determinative and those who would merely make it relevant. But even on the most cautious understandings of the role of CBA, government’s choices would be significantly affected by the translation of environmental benefits into monetary equivalents.

6. See Frank Ackerman and Lisa Heinzerling, *Priceless: On Knowing the Price of Everything and the Value of Nothing* (New York: New Press, 2003).

7. See Cass R. Sunstein, *The Cost-Benefit State* (Washington, D.C.: American Bar Association, 2002), p. 145.

To say the least, it is highly controversial to say that people's protection against environmental risks is properly measured by their WTP to avoid those risks. It is at least equally controversial to use WTP as the basis for policies protecting endangered species, nature, and wildlife. But the Precautionary Principle raises serious problems of its own. How much precaution is the right level of precaution? Are costs relevant to the answer? In any case human beings face a number of risks, not simply one, and any effort to reduce one risk might well increase another risk. Is it possible, even in principle, to take precautions against all risks, rather than a subset? If all risks cannot be reduced at once, how should regulators set priorities?

In this essay, I approach these questions through a discussion of three illuminating books that offer radically different approaches to environmental protection. Frank Ackerman and Lisa Heinzerling believe that CBA is a hopelessly crude tool, one that buries indefensible judgments of morality and politics.⁸ Drawing on the war on terrorism, they argue for the Precautionary Principle instead. By contrast, Adam Burgess uses the controversy over cell phones to suggest that the Precautionary Principle capitulates to, and even promotes, baseless public fears.⁹ Objecting to what he sees as excessive fear of new technologies, Burgess argues for careful attention to scientific evidence and for regulation only when the risk is real. Richard Posner argues for CBA and economic analysis in a context in which it seems least promising: catastrophic risk.¹⁰ He contends that global warming, and other potentially catastrophic problems, cannot sensibly be approached without a disciplined effort to quantify and monetize both costs and benefits. But where Ackerman and Heinzerling see CBA as an excuse for regulatory inaction, Posner invokes CBA on behalf of aggressive controls on greenhouse gases and other sources of potentially serious danger. Indeed his central goal is to draw private and public attention to catastrophic risks that are exceedingly unlikely to come to fruition.

Building on the arguments made by Burgess and Posner, I shall mount a qualified defense of CBA here. Without some sense of both costs and benefits—both nonmonetized and monetized—regulators will be making a stab in the dark. Human beings have a great deal of difficulty in assessing risks, making them prone to both hysteria and neglect; CBA does not supply definite answers, but it can help to establish

8. See Ackerman and Heinzerling.

9. See Adam Burgess, *Cellular Phones, Public Fears, and a Culture of Precaution* (Cambridge: Cambridge University Press, 2004).

10. See Richard A. Posner, *Catastrophe: Risk and Response* (New York: Oxford University Press, 2004).

which risks are serious and which are not.¹¹ By contrast, the Precautionary Principle approaches incoherence. Because risks are on all sides of social situations, and because regulation itself increases risks of various sorts, the principle condemns the very steps that it seems to require.

But building on the arguments made by Ackerman and Heinzerling, I shall suggest that there are two serious problems with CBA. The first is that WTP is sometimes an inappropriate basis for environmental policy. Human beings are citizens, not merely consumers, and their consumption choices, as measured by WTP, might be trumped by their reflective judgments as citizens. In any case, WTP is dependent on ability to pay; when the poorest members of societies stand to gain from environmental protection, they should be protected even if their poverty ensures that their WTP is low. The second problem is that regulators cannot always assign probabilities to environmental outcomes. If probabilities cannot be assigned, regulators are unable to engage in CBA; they might do well to follow the maximin principle, taking steps to avoid the worst-case scenario. This point helps pave the way toward a narrower and more refined use of the Precautionary Principle, one that has important real-world applications and that provides a valuable complement to approaches based on CBA.

I. MONETIZATION AND ITS DISCONTENTS

Ackerman and Heinzerling do not object to efforts to specify the range of outcomes associated with alternative courses of action.¹² Their principal objection is to the WTP criterion. Insisting that human deaths are not mere “costs,” Ackerman and Heinzerling contend that CBA is morally obtuse. They argue that a well-functioning democracy should respect the informed judgments of citizens rather than aggregate private consumption choices. Ackerman and Heinzerling much prefer the Precautionary Principle, which, in their view, is “a more holistic analysis” that argues for regulation in the face of scientific uncertainty and that is “committed to fairness within and beyond this generation” (p. 234).

Ackerman and Heinzerling are aware that many people have turned to CBA because of widely publicized studies that purport to show a high

11. See Allan Gibbard, “Risk and Value,” in *Values at Risk*, ed. Douglas MacLean (Totowa, N.J.: Rowman & Allanheld, 1986), p. 94, for an exploration of “a principal rationale for wanting something like risk-cost-benefit analysis: for seeking a way to regiment our judgments about risk, and so to avoid the blatant irrationalities of unaided common sense.”

12. I discuss their book for a popular audience in Cass R. Sunstein, “Your Money or Your Life,” *New Republic* (March 15, 2004), p. 27; my treatment here borrows from that discussion.

level of arbitrariness in modern regulation.¹³ According to such studies, regulations in the United States are wildly inconsistent. Sometimes the United States spends \$100,000 (or less) to save a human life. Sometimes it spends tens of millions. Cost-benefit supporters ask: shouldn't nations be devoting their resources to serious health problems rather than trivial ones? If a nation can spend ten million dollars to save one thousand lives, shouldn't it do that rather than waste the money on a similarly priced program that saves only one or two people? In any case, human beings make many errors in assessing risks, using heuristics and demonstrating biases that make them exaggerate some dangers and underestimate others.¹⁴ These errors seem to be replicated in existing policies; CBA might be defended as a promising corrective to blunders in citizens' perception of risk. In these ways, interest in CBA has been fueled less by contentious claims of value than by the pragmatic suggestion that it can assist in more intelligent priority setting.¹⁵

Ackerman and Heinzerling believe that the attack on the current system is based on misleading studies, burying controversial and indeed implausible judgments of value. True, some regulations do not prevent many deaths, but they do prevent serious (nonfatal) harms to human health and also harms to ecosystems. The resulting benefits should not be disparaged. More fundamentally, Ackerman and Heinzerling argue that the key studies find low benefits partly because they greatly discount future gains to life and health. Everyone agrees that a dollar today is worth more than a dollar in twenty years; economists use a standard "discount rate" (often 7 percent annually) to convert future dollars into current equivalents. In calculating the benefits of regulation, they use the same discount rate for lives saved and illnesses averted. Ackerman and Heinzerling contend that this approach wrongly shrinks the value of regulations that will save people in the future. One of their central

13. The most well-known is John F. Morrall III, "A Review of the Record," *Regulation* (November/December 1986): 25-29, p. 30, table 4. For an updated treatment, see John F. Morrall III, "Saving Lives: A Review of the Record," *Journal of Risk and Uncertainty* 27 (2003): 221-37.

14. A good collection is Thomas Gilovich et al., eds., *Heuristics and Biases: The Psychology of Intuitive Judgment* (Cambridge: Cambridge University Press, 2002). For use of this idea in (qualified) defense of CBA, see Gibbard.

15. David M. Driesen, *The Economic Dynamics of Environmental Law* (Cambridge, Mass.: MIT Press, 2003), offers a powerful criticism of CBA insofar as it offers a static account of both costs and benefits and fails to see that regulation and other forces often produce innovation, thus reducing the expense of environmental protection. I believe that this argument is best taken as a reason for skepticism about existing figures about likely costs, rather than as an attack on CBA as such. See Matthew Adler, "Cost-Benefit Analysis, Static Efficiency, and the Goals of Environmental Law," *Environmental Affairs Law Review* 31 (2004): 591-600.

claims, then, is that the standard discount rate should not be applied to future savings in terms of life and health.

Suppose that their arguments are right—that once economic values are properly assigned to environmental gains, few existing regulations will be condemned as requiring huge investments for trivial benefits. Regulators still might want to use cost-benefit analysis to improve current decisions.¹⁶ Ackerman and Heinzerling complain that to do this, they will have to produce a dollar value for a human life—and any such effort will be arbitrary, offensive, or worse. They reject the view that WTP, based largely on workplace studies, produces information that agencies should use. In their view, workers often have little information about the risks that they face, and hence they cannot be charged with consciously trading hazards against dollars. Even when workers are informed, they may have few options and hence little choice. If they accept a job with significant hazards for a low premium, it is not because they are genuinely free to choose.

Some anomalies in the empirical literature are highly relevant here. Nonunionized workers have sometimes been found to receive little or nothing for the reduction of statistical risks, and African-Americans have been found to receive much less than white people do.¹⁷ Does it follow that regulators should treat the lives of nonunionized workers, or African-Americans, as worth especially little? Ackerman and Heinzerling add that the key studies ask only how much individuals care about risks to themselves. They ignore the fact that many of us value the lives of others too. I might be willing to pay only \$60 to eliminate a 100,000 risk that I face, but I might be willing to pay much more than that to eliminate the same risk from my child's life and substantial amounts to help reduce the risks of my friends. Altruism is ignored in the current calculations.

Ackerman and Heinzerling also contend that statistically equivalent risks should not be treated the same, because people's valuations of mortality risks depend not only on the probability of harm but also on their nature and their context. About 3,000 people died from the terrorist attacks of 9/11—a much smaller number than die each year from suicide (30,500), motor vehicle accidents (43,500), and emphysema (17,500). Ackerman and Heinzerling approve of the fact that the reaction of the United States to the 9/11 attacks was not based on simple numerical comparisons. Drawing on work by psychologist Paul Slovic, Ackerman and Heinzerling emphasize that the risk judgments of or-

16. See W. Kip Viscusi, *Rational Risk Regulation* (Cambridge: Cambridge University Press, 2001).

17. John D. Leeth and John Ruser, "Compensating Wage Differentials for Fatal and Nonfatal Injury Risk by Gender and Race," *Journal of Risk and Uncertainty* 27 (2003): 257–77.

dinary people diverge from the risk judgments of experts—not because ordinary people are stupid or confused, but because they have a different normative framework for evaluating risks.¹⁸ While experts focus on the number of deaths at stake, most people are especially averse to risks that are unfamiliar, uncontrollable, involuntary, irreversible, inequitably distributed, man-made, or catastrophic.¹⁹ Diverse valuations of diverse risks should play a role in regulatory policy.

For example, most of us are not greatly troubled by the cancer risks associated with x-rays, partly because they are voluntarily incurred. By contrast, the risks of terrorism and even pesticides and air pollution are more alarming because individuals cannot easily control them. And when a risk is faced by an identifiable community—as, for example, when landfills with toxic chemicals are located in largely poor areas—the public is especially likely to object to what it will perceive as unfairness.²⁰ Ackerman and Heinzeling thus complain that CBA disregards important qualitative differences among quantitatively identical risks. It also tends to ignore, and often to reinforce, patterns of social inequality, above all because it pays no attention to a key question, which is distributional: who receives the benefits and who incurs the costs? For both domestic and international environmental issues, Ackerman and Heinzeling emphasize the importance of fairness. If environmental threats mostly burden poor people, regulators should take that point into account, whatever the cost-benefit ratio.

Ackerman and Heinzeling are also concerned about how cost-benefit analysts value nature. How much will human beings pay to save an animal or a member of an endangered species? Economists have tried to answer the question by actually asking people. For example, one study finds that the average American family is willing to pay \$70 to protect the spotted owl, \$6 to protect the striped shiner (an endangered fish), and as much as \$115 per year to protect major parks against impairment of visibility from air pollution. Ackerman and Heinzeling ridicule these numbers, complaining that any precise monetary value fails to provide useful information. Bans on whaling, for example, are rooted in a widely shared ethical judgment, not on cost-benefit analysis. A democracy should make its decisions about the protection of nature

18. See Paul Slovic, *The Perception of Risk* (London: Earthscan, 2000).

19. But see Howard Margolis, *Dealing with Risk* (Chicago: University of Chicago Press, 1999), for a challenge to this account of the lay/expert division in risk perceptions.

20. Note, however, that the Not In My Backyard Syndrome—known in the trade as NIMBY—suggests that many people will make self-serving judgments about the proper location of environmentally risky activities. This point is related to the suggestion, developed below, that people tend to become intuitive cost-benefit analysts when both the benefits and the costs of environmental regulation are on-screen.

on the basis of such ethical judgments rather than by aggregating people's willingness to pay.

Ackerman and Heinzerling offer a final objection to CBA: the rights of future generations. I have noted that economists generally apply a discount rate to future gains and losses. With a 7 percent discount rate, \$1,000 in twenty years is worth only \$260 today. Cost-benefit analysts within the federal government have long applied a standard discount rate for money (7 percent) to the benefits of safety and health regulation, so that prevention of 1,000 fatal cancers in 2025 is equivalent to the prevention of 260 fatal cancers in 2005. Ackerman and Heinzerling respond that lives are not like money; they cannot be placed in a bank for the accumulation of interest. A discount rate of 7 percent radically shrinks the value of reductions in risk for those born, say, one hundred years from now. But current generations owe obligations to the future and should not discount measures that protect people not yet born.

Invoking the Precautionary Principle, Ackerman and Heinzerling argue that nations are obliged to take action against serious threats even before there is a scientific consensus. Above all, they want regulators to make regulatory decisions by attending to the worst-case scenario. If the worst case is extremely bad, aggressive regulation is desirable even if it might result in wasted money. When a nation spends too much on regulatory protection, it loses limited resources, which admittedly is undesirable; but waste is far better than catastrophe. Hence their "preference is to tilt toward overinvestment in protecting ourselves and our descendants" (p. 227). Ackerman and Heinzerling note that this approach was taken in the context of the military spending in the Cold War, arguing that the nation rightly prepared for the high-risk case. They see protection against terrorism in similar terms. Ackerman and Heinzerling want to treat health and environmental risks in the same way.

II. POINTLESS PRECAUTIONS?

Ackerman and Heinzerling do not focus in detail on any particular regulatory issue. By contrast, Burgess explores the idea of "precaution" with close reference to a single controversy: the health risks associated with cellular phones. Burgess does not explicitly discuss CBA, but he is highly skeptical of the Precautionary Principle, which, in his view, leads regulators to capitulate to baseless public fear. One of Burgess's central claims is that public fears are often manufactured rather than found; for this reason, Burgess describes himself as a "social constructivist" (p. 11). But with respect to risk, Burgess is no constructivist at all. He believes that some risks are serious and that others are not, and that science is the best way to tell the difference.

Burgess contends that notwithstanding countless efforts, no repu-

table study has demonstrated significant health risks as a result of emissions from cell phones and cell phone towers.²¹ To date, much of the so-called evidence comes from anecdotes of the sort provided by anti-cell phone activist Debbie Collins, who contended that her daughter's health had significantly improved after she was removed from a school near a cell phone tower. Rejecting expert opinion, Collins stated: "She's a different child now—it's all the proof I need to convince me there is a link between those wretched masts and the health of children" (p. 1). Another mother said, "I needed no more proof than that. This term he started at a new school and I can already see the change in him. His memory has improved and his headaches have gone" (p. 2). Burgess is concerned that a precautionary approach, founded on statements of this kind, will both aggravate fear and impose costs for no good reason.

Burgess's tale begins with a media campaign. In the early 1990s, a number of newspaper stories in the United Kingdom contended, on the basis of little evidence, that mobile phones and base stations were producing harmful health effects. Apparently influenced by these stories, the European Commission in Brussels began an official inquiry in 1995, ultimately funding future research and concluding that adverse effects could not be ruled out. Public fears intensified in 1996 after the issue received attention in a consumer health program on the BBC and a widely read news story in the *Sunday Times* featuring the headline "Mobile phones cook your brain." In 1997, alarmist reports grew in the media, making the claim that cell phones could produce illness and premature mortality (and also reduce sex drives). These reports helped to spur citizen action. By 1999, local political campaigns against cell phone towers became prominent, and they received favorable coverage in local and regional newspapers, which further energized public concern.

These campaigns significantly affected both private and public institutions. The London Metropolitan Police service told its officers to limit cell phones as a "purely precautionary measure" (p. 87). Harrods banned cell phones from its premises. Speaking in explicitly precautionary terms, entrepreneur Richard Branson recommended the use of safety devices for his employees. Local governing councils across the United Kingdom attempted to ban or restrict mobile towers, particularly those near schools. At the national level, the minister for public health legitimated public fears, insisting that in such a context, "it is very im-

21. A different set of issues is raised by the risks associated with use of cell phones while driving. Here there is much stronger evidence of serious hazards. For an overview, see Robert Hahn and James Prieger, "The Impact of Driver Cell Phone Use on Accidents" (July 2004), available at <http://www.aei-brookings.org/publications/abstract.php?pid=806>.

portant that” officials “work very hard to keep ahead of public anxiety” (pp. 88–89). In Burgess’s account, precautionary responses by official institutions helped to fuel that very anxiety. Thus it “is only through being taken seriously by state bodies that the allegations about hypothetical risks have been able to command authority and acquire momentum beyond the immediate reactions of some individuals” (p. 222).

Burgess also makes some interesting and somewhat puzzling remarks about cross-cultural comparisons. In the United Kingdom, there was intense public focus on cell phone risks; similar concerns have been found in Australia, Italy, and South Africa. In Italy, the environment minister established a “green hotline” asking people to state their complaints about “abusive” siting of cell phone towers. The Australian government funded a large-scale research project on potentially adverse health effects. But in the United States, the brief burst of concern in the early 1990s rapidly dissipated, to the point where it is hard to find any serious private or public concern about health risks. And in Finland, no discernible public fear has arisen at all, even though Finland has the highest percentage of cell phone users in the world. (The fact that Nokia is Finland’s biggest company is highly relevant here—a point, bearing on both precautions and CBA, to which I will return.)

Burgess thinks that the cell phone controversy is merely one example of the misuse of precautionary thinking in domains in which scientific evidence fails to support people’s fears. For example, he challenges European skepticism about genetically modified food, describing it as “alarm” (p. 259); and he mounts a broader attack on what he sees as the unhelpful belief that it is wrong to interfere with nature. He is therefore troubled by a wide climate of sensitivity to small risks, especially those that are novel and associated with technological innovation. Precautionary thinking, he believes, helps to create a culture of fear.

III. CATASTROPHE, COSTS, AND BENEFITS

Richard Posner is one of the founders of the economic analysis of law, and he should be expected to be enthusiastic about CBA. In *Catastrophe: Risk and Response*, he does not disappoint that expectation. What makes the book noteworthy is its focus on the application of CBA to truly catastrophic risks—those that might threaten the survival of the human race. Posner covers an extraordinarily wide range of hazards, including genetically modified crops, robotics, and nanotechnology, but he focuses in particular on four: asteroid collisions, particle-accelerator disasters, global warming, and bioterrorism. Posner believes that none of these risks can be dismissed, and he thinks that cost-benefit analysis should be applied to each of them.

Consider, for example, the dangers associated with very powerful particle accelerators. It is extremely unlikely, but not impossible, that

such accelerators will produce a highly compressed object called a “strangelet,” which has the ability to convert whatever it encounters into a new form of matter. Posner quotes Martin Rees, professor of physics at the University of Cambridge, who writes, “A hypothetical strangelet disaster could transform the entire planet Earth into an inert hyperdense sphere about one hundred meters across.”²² Posner accepts the widely held view that a strangelet disaster is exceedingly improbable, but he insists that it cannot be ruled out. As a result, he thinks that nations should at least be willing to ask whether the benefits of very powerful particle accelerators justify incurring the risk. On that question, he is quite doubtful.

Posner’s discussion extends over a wide range. Because my topic is environmental protection, I shall focus on his treatment of global warming.²³ Posner believes that the associated risks should be taken seriously, above all because of the possibility of truly catastrophic harm. He acknowledges that the leading economic expert on global warming, William Nordhaus, estimates its total cost at \$4 trillion—a high figure, to be sure, but hardly astronomical, and one that allows cost-benefit analysis to get off the ground.²⁴ (The United States has an annual GDP of \$10 trillion, and as Posner points out, \$4 trillion represents present value, which might be compared with the present economic value of the United States, which is roughly \$100 trillion.) Nordhaus produces his \$4 trillion figure essentially through the methods that Ackerman and Heinzerling deplore, namely, using WTP and discounting the future.

Posner is concerned not with the objections made by Ackerman and Heinzerling, which he implicitly rejects, but with the possibility that Nordhaus’s estimate greatly understates the problem, above all because of the dangers of abrupt warming, which would be especially destructive. Thus Posner thinks that existing models do not rule out the possibility of (for example) very rapid changes in both temperature and sea levels, the evolution and migration of deadly pests, and even a runaway greenhouse effect, produced by melting of tundras, thus releasing large quantities of additional greenhouse gases. One worst-case scenario is “snowball earth,” covering the world with a layer of ice several kilometers thick, a result of massive increases in cloud cover, preventing sunlight from reaching the earth. Sounding very much like Ackerman and Heinzerling, Posner seeks to draw attention to the worst that might happen.

Many scientists and economists, including Nordhaus, believe that

22. See Martin Rees, *Our Final Hour* (New York: Basic Books, 2003), p. 120.

23. For a general discussion, see Stephen Gardiner, “Ethics and Global Climate Change,” *Ethics* 114 (2004): 555–600.

24. See William Nordhaus and Joseph Boyer, *Warming the World* (Cambridge, Mass.: MIT Press, 2000).

global warming is not likely to create catastrophic harm and that the real costs, human and economic, will be high but not intolerable. In their view, the worst-case scenarios can be responsibly described as improbable.²⁵ Posner disagrees. He believes that “no probabilities can be attached to the catastrophic global-warming scenarios, and without an estimate of probabilities an expected cost cannot be calculated.” In the terms of decision theory, Posner contends that global warming presents a situation of uncertainty, where probabilities cannot be assigned to outcomes, rather than risk, where such probabilities can be assigned.²⁶ In this way, global warming differs from other potentially catastrophic risks that Posner explores, such as the strangelet disaster, which everyone characterizes as exceedingly unlikely.

Posner does not claim that responses to catastrophic risks should be chosen solely by reference to CBA. But he proposes that CBA “is an indispensable step in rational decision making in this as in other areas of government regulation. Effective responses to most catastrophic risks are likely to be extremely costly, and it would be mad to adopt such responses without an effort to estimate the costs and benefits” (p. 139). While favoring CBA, Posner rejects the Precautionary Principle because of its “sponginess” (p. 140). He contends that once that principle becomes sensibly tempered, it turns into a form of CBA with risk aversion, that is, a form of CBA that creates a margin of safety to protect against those dangers that produce special concern. This understanding of the Precautionary Principle, he believes, is perfectly reasonable, but it turns the principle into a version of CBA, not a rival (as Ackerman and Heinzerling claim).

Posner emphasizes that any effort to apply CBA to catastrophic risks requires a great deal of guesswork. Consider the proposal to build a new and very powerful particle accelerator, Brookhaven’s Relativistic Heavy Ion Collider. Posner is concerned about the remote possibility that the Brookhaven collider will destroy the earth; he wants to evaluate the proposal by reference to CBA. He notes that no effort has been made to monetize its benefits, but he ventures a “wild guess” (p. 140) that they amount to \$250 million per year. (It is extremely hard to

25. *Ibid.*, p. 88. Nordhaus suggests a 1.2 percent probability of a catastrophic impact with 2.5° C warming and a 6.8 percent probability with 6° C warming.

26. See Frank H. Knight, *Risk, Uncertainty, and Profit* (Boston: Houghton Mifflin, 1933); Paul Davidson, “Is Probability Theory Relevant for Uncertainty? A Post-Keynesian Perspective,” *Journal of Economic Perspectives* 5 (1991): 129–45. Some people object that uncertainty does not exist, because it is always possible for decision makers to produce probability assignments by proposing a series of lotteries over possible outcomes; but such assignments have no epistemic credentials if unrooted in either theory or repeated experiences, and some environmental problems, plausibly including global warming, are that sort of case.

produce a figure, monetized or nonmonetized, to capture the benefits of basic research; for this reason, Posner's guess is indeed wild.) With that amount, the collider would have a net present value of \$400 million: \$21.1 billion in benefits, assuming a 3 percent discount rate, over a projected ten-year span, minus the accelerator's construction and operating costs, which are \$1.1 billion. But what is the monetized value of the extinction risk? To answer that question, Posner needs to estimate both the probability of extinction and its monetized cost if it comes to fruition. For probability, he ventures a figure of 1 in 10 million, a figure that he also deems "arbitrary," though it is in line with several estimates by expert risk assessors. For monetized cost, based on WTP to reduce statistical risks and a 3 percent discount rate, he values the loss of the human race at \$600 trillion.²⁷ Doing the arithmetic, Posner believes that the net benefits of the Brookhaven collider are negative: -\$100 million. Thus he concludes that the collider should not be built.

Posner acknowledges that "global warming is the poster child for the limitations of cost-benefit analysis" (p. 155). But even here, he thinks that it is possible to make progress by attempting to be as quantitative as possible. Most economists, armed with CBA, oppose the Kyoto Protocol, arguing that its monetized costs probably would exceed its monetized benefits. Recall that the monetized costs of global warming are estimated at around \$4 trillion. For the world as a whole, the monetized benefits of the Kyoto Protocol are estimated at far less than those costs: only \$108 billion.²⁸ The reason is that the protocol would do relatively little about the problem of global warming. Greenhouse gases stay in the atmosphere for a long time, and the Kyoto Protocol would not, of course, affect those emissions that have already occurred. In addition, its provisions do not limit developing nations, which will soon be large sources of greenhouse gases, at all (a primary complaint of the Bush administration); and for the industrialized world, it would merely stabilize emissions modestly below 1990 levels. Hence the benefits of the Kyoto Protocol would be modest, consisting as they would of a mere reduction in the increase of global warming emissions. At the same time, the Kyoto Protocol would impose significant costs on those subject to it, producing a total global cost ranging from \$59 billion to \$884 billion.²⁹ A standard view is that the Kyoto Protocol fails CBA, because

27. To produce this number, Posner values an individual life at only \$50,000, based on an assumption of a very low WTP for tiny risks. He emphasizes that this is a quite conservative assumption and that it would be reasonable to choose higher values.

28. See Nordhaus and Boyer, p. 167.

29. *Ibid.*, p. 156. The low end of the range represents the cost with fully global emissions trading; the high end represents the cost without trading. If trading occurs within the nations listed in an annex to the protocol, the costs are estimated at \$217 billion.

it is likely, in its implementation, to inflict costs in excess of the \$108 billion gains.

Posner thinks this analysis is badly incomplete, because it ignores the possibility that government regulation will force technological innovation, thus producing dramatic decreases in greenhouse gas emissions; and dramatic decreases are necessary to reduce the risk of catastrophe. Posner is particularly interested in the potentially desirable effects of significant taxes on carbon emissions. Such taxes would create economic incentives to develop clean fuels and better methods of carbon sequestration. Posner acknowledges that in view of existing uncertainty and the high costs of emissions controls, it is tempting simply to wait for more scientific information (as the Bush administration has argued). One problem with this approach is that of irreversibility: once greenhouse gases are in the atmosphere, they stay there for a long time. In a key passage, he argues that making “shallower cuts now can be thought of as purchasing an option to enable global warming to be stopped or slowed at some future time at a lower cost” (p. 112).

Posner does not offer a formal CBA for various approaches to the global warming problem. The reason is that his fundamental concern is abrupt warming, to which he believes that no probability can be assigned. In contrast to his quantitative analysis of particle accelerators, his analysis of global warming does not offer many numbers. Indeed, his own form of balancing does not have a transparent structure; his major argument involves the option analysis just described, with the suggestion that current cuts give us the flexibility to reduce warming in the future if that is what we choose to do. He thus argues in favor of aggressive technology-forcing emissions taxes on greenhouse gases, above all to reduce the possibility of catastrophic risk.

IV. PROBLEMS WITH PRECAUTIONS

These three books cover three quite different issues: the idea of precaution, the translation of environmental harms into monetary equivalents, and the appropriate approach to environmental regulation in the face of scientific uncertainty. Let us explore these issues in turn.

Ackerman and Heinzerling argue in favor of the Precautionary Principle. Burgess rejects it as leading to nonsensical outcomes. Posner believes that it must be converted into a form of CBA, one that embodies an aversion to those risks that deserve particular concern. At first glance, it is tempting to say, with Burgess, that the idea of precaution will lead to excessive controls on small or nonexistent risks. It is equally tempting to say, with Posner, that the idea is simply too vague to provide guidance: how much precaution is enough? But the most serious problem lies

elsewhere. In many contexts, the Precautionary Principle is incoherent.³⁰ Risks are often on all sides of a social situation, and risk reduction itself produces risks. Hence the Precautionary Principle, taken for all that it is worth, forbids the very measures that it requires. Ackerman and Heinzerling neglect the fact that regulation can create dangers of its own, in a way that suggests that along some dimensions, many precautions are not precautionary at all. Advocates of precaution often emphasize the costs associated with a product or process, without seeing that it may have benefits as well; and sometimes those benefits involve the environment itself. Why should regulators examine only one side of the ledger?

For example, regulation often gives rise to substitute risks, in the form of hazards that materialize, or are increased, as a result of regulation. Consider the case of DDT, often banned or regulated in the interest of reducing risks to birds and human beings. From the standpoint of the Precautionary Principle, the problem with such bans is that, in poor nations, they eliminate what appears to be the most effective way of combating malaria—and thus significantly undermine public health.³¹ Or consider the United States EPA's effort to ban asbestos, a ban that, on health grounds, might well seem justified or even compelled by the Precautionary Principle. The difficulty, from the standpoint of that very principle, is that substitutes for asbestos also carry risks. The problem is pervasive. The Precautionary Principle is often invoked as a reason for banning genetic modification of food, on the ground that genetic modification creates risks to human health and to the environment. The problem is that genetic modification of food also promises benefits to human health and the environment—and, by eliminating those benefits, regulation itself threatens to run afoul of the Precautionary Principle. When the principle seems to give guidance, it is often because those who use it are focusing on one aspect of risk-related situations and neglecting others.

It is possible to go much further. A great deal of evidence suggests the possibility that an expensive regulation can have adverse effects on life and health.³² It has been urged that a statistical life can be lost for every expenditure of \$7 million;³³ one study suggests that an expenditure

30. I develop this claim in some detail in *Laws of Fear: Beyond the Precautionary Principle* (Cambridge: Cambridge University Press, forthcoming).

31. See Indur Goklany, *The Precautionary Principle* (Washington, D.C.: Cato, 2002).

32. Ralph Keeney, "Mortality Risks Induced by Economic Expenditures," *Risk Analysis* 10 (1990): 147–60; Randall Lutter and John F. Morrall III, "Health-Health Analysis: A New Way to Evaluate Health and Safety Regulation," *Journal of Risk and Uncertainty* 8 (1994): 43–69, p. 49, table 1.

33. See Keeney.

of \$15 million produces a loss of life.³⁴ Another suggests that poor people are especially vulnerable to this effect—that a regulation that reduces wealth for the poorest 20 percent of the population will have twice as large a mortality effect as a regulation that reduces wealth for the wealthiest 20 percent.³⁵ If poor people are paying a significant amount for modest environmental benefits, their health might be made worse rather than better. To be sure, both the phenomenon and the underlying mechanisms are reasonably disputed (and Ackerman and Heinzerling reasonably dispute it).³⁶ For purposes of applying the Precautionary Principle, the only point is that sensible people believe in that association. It follows that a multimillion dollar expenditure for “precaution” has—as a worst case scenario—significant adverse health effects, with an expenditure of \$200 million possibly leading to perhaps as many as twenty lives lost.

This point makes the Precautionary Principle hard to implement not merely where regulation introduces or increases substitute risks but in any case in which the regulation costs a significant amount. If this is so, the Precautionary Principle, for that very reason, raises doubts about many regulations. If the principle argues against any action that carries a small risk of imposing significant harm, then we should be reluctant to spend a great deal of money to reduce risks, simply because those expenditures themselves carry risks. Here is the sense in which the Precautionary Principle, taken for all that it is worth, is paralyzing: it stands as an obstacle to regulation and nonregulation, and to everything in between.³⁷

Ackerman and Heinzerling do not sufficiently appreciate this point. They neglect the possibility that expensive regulation will actually hurt real people. Consider their seemingly offhand remark about protection against workplace hazards: the “costs of the regulation would probably be borne by the employers who would be required to maintain safer workplaces” (p. 193). But the costs of regulation are often borne not only by employers but also by consumers, whose prices increase, and by workers, who might find fewer and less remunerative jobs. When government imposes large costs on “polluters,” consumers and workers are likely to pay part of the bill. And if prices increase, some risks will

34. See Robert W. Hahn et al., *Do Federal Regulations Reduce Mortality?* (Washington, D.C.: American Enterprise Institute, 2000).

35. See Kenneth S. Chapman and Govind Hariharan, “Do Poor People Have a Stronger Relationship between Income and Mortality than the Rich? Implications of Panel Data for Health-Health Analysis,” *Journal of Risk and Uncertainty* 12 (1996): 51, 58–63.

36. See Lutter and Morrall.

37. Posner’s distinctive concern involves catastrophic harms with a threat of extinction; here, of course, no equally serious risk is likely to be on the other side. I return to this point in the discussion of maximin below.

increase as well. To be sure, some environmental regulations do increase employment and decrease prices. But as a general rule, there is no reason to believe that regulatory imposition of high costs will benefit workers and consumers; the opposite is more likely to be true.

In the context of cell phones, this point helps illuminate a quite remarkable fact, one to which Burgess gives too little attention. Notwithstanding the popularity of precautionary thinking, and the apparent intensity of public fears, those fears did not, in fact, produce large-scale controls on either phones or towers. Burgess offers no explanation of why such controls did not materialize, but his brief discussion of Finland provides a useful clue. Is it really a paradox, or an irony, that fears of cell phones were especially weak in a nation that has the largest percentage of cell phone users in the world? I do not believe so. The Finnish economy is heavily dependent on Nokia and thus the cell phone industry; people in Finland do not want the Finnish economy to collapse. And if most citizens depend on cell phones, they are far less likely to accept sensationalistic claims of risk, simply because they have so much to lose from regulation. (Imagine, by way of comparison, the likely public reaction to a current suggestion by an American politician that cell phones should be banned because they pose a cancer risk.) If the benefits of cell phone use are evident to all or most, then people will demand a great deal of evidence that the harm is real. In short, the very idea of precaution loses some of its appeal when people are aware that precaution imposes costs and even risks of its own.³⁸ When people are aware of that fact, some kind of balancing, involving both costs and benefits, is likely to emerge.

In a brief but illuminating discussion of another environmental issue, Burgess strongly supports the general point. He refers to a mining town in Colorado whose citizens were deemed, by the EPA, to be at risk from toxic contamination. The town's citizens, already suffering from serious economic decline, responded not with fear—and much less with enthusiasm for a precautionary approach—but by demonizing the EPA, which it regarded as “the devil incarnate. Grimly they recounted how government bureaucrats had invaded their town uninvited, threatening residents with the prospect of condemned property, involuntary relocation, and unwelcome new legal requirements. . . . And all, they claimed, over a hazard ‘that doesn’t exist’” (p. 272). Far from succumbing to panic, citizens of the mining town were well aware of how much they had to lose from aggressive regulation; hence they sought to dismiss real evidence of harm. Precautions and precautionary thinking seem far less attractive when people believe that precautions would produce significant costs and risks.

38. See the treatment of fungibility in Margolis.

Burgess does not draw attention to one of the remarkable lessons of his story, which is that the cell phone scare did not produce aggressive regulation not only because the evidence of harm was weak but also and still more fundamentally because a growing number of people use cell phones and would be inconvenienced, or far worse, by such measures. (Consider the fact that cell phones are often used to obtain help in emergency situations.) In the context of genetically modified food, by contrast, the costs of regulation are not highly visible, at least not to Europeans. Precautionary thinking, in short, is most appealing when the costs and burdens of precautions are not visible. When both costs and benefits are on the public viewscreen, people become intuitive cost-benefit analysts, and they tend to be cautious about precautions—unless the evidence in their favor is strong. This point brings us directly to the questions raised by CBA.

V. COSTS AND BENEFITS

As Ackerman and Heinzerling stress, some of the most difficult questions for CBA involve the translation of risks into monetary equivalents. Recall that under current practice, the monetary values come mostly from real-world markets, producing evidence of compensation levels for actual risks. It is important to see that in basing CBA on calculations of this kind, regulators are not, in fact, producing a “value of a statistical life.” In fact they are not “valuing life” at all. They are not saying that the average American would pay \$6 million to avoid death, or that a human life is, in some metaphysical sense, worth that amount. Instead they are generating numbers that reflect the market value of statistical risks. Typically agencies are dealing with low-level risks, on the order of 1/100,000, and when they “value a life” at \$6 million, they are really saying that the evidence suggests that people must be paid \$60 to be subject to a risk of that magnitude—and that government will build on that evidence in making regulatory decisions.

A. *The Argument for WTP*

Ackerman and Heinzerling think that this practice is a form of madness, and hence they do not pause to ask why regulators in a democratic society might care about market valuations of statistical risks. But there are two possible answers, both connected with individual choice, and both growing directly out of prominent strands in liberal theory. The first involves welfare; the second involves autonomy.

To clarify the point, assume a society in which people face multiple risks of 1/100,000 and in which every person is both adequately informed and willing to pay no more and no less than \$60 to eliminate each of those risks. Assume too that the cost of eliminating these 1/100,000 risks is widely variable, ranging from close to zero to hundreds

of millions of dollars. Assume finally that the cost of eliminating any risk is borne entirely by those who benefit from risk elimination. Under that assumption, regulation imposes the equivalent of a user's fee; for example, people's water bills will entirely reflect the costs of a policy that eliminates a 1/100,000 of getting cancer from arsenic in drinking water. If the per-person cost is \$100, each water bill will be increased by exactly that amount.

At first glance, use of WTP, under the stated assumptions, is easy to defend. Why should people be forced to pay an amount for regulation that exceeds their WTP? Of course we might believe that a measure of redistribution is appropriate—that private sources, or government, should provide people with regulatory protection for free. But regulation need not, and often does not, amount to a subsidy to those who benefit from it. After the enactment of workers' compensation regulation, nonunionized workers faced a dollar-for-dollar wage reduction, corresponding almost perfectly to the expected value of the benefits they received.³⁹ For drinking water regulation, something similar is involved. The cost of regulations is passed onto consumers in the form of higher water bills.⁴⁰

More particularly, those who are interested in welfare will insist on the relevance of WTP under the stated assumptions.⁴¹ If people are willing to pay \$60, but no more, to eliminate a risk of 1/100,000, then we have good reason to think that their welfare is increased by asking them to pay that amount—and that their welfare is decreased by asking them to pay more. There are many demands on people's budgets, and if they refuse to spend more than \$60 on a 1/100,000 risk, it may be because they would like to use their money for food, shelter, recreation, education, or any number of other goods. Regulation can operate as a forced exchange and, by hypothesis, a forced exchange on terms that people dislike will make them worse off.⁴²

For purposes of evaluating regulation, it does not matter if the existing distribution of income is unjust or if poor people are, in an intelligible sense, coerced to run certain risks. The remedy for unjust

39. Price Fishback and Shawn Everett Kantor, *A Prelude to the Welfare State* (Chicago: University of Chicago Press, 1998).

40. See Sunstein, "The Arithmetic of Arsenic."

41. See Gibbard, p. 97, for a valuable discussion of how CBA might be taken as "a rough surrogate for expected total intrinsic-reward maximization." To the same general effect, for defenses of CBA on welfare grounds, see Matthew Adler and Eric A. Posner, "Implementing Cost-Benefit Analysis When Preferences Are Distorted," *Journal of Legal Studies* 29 (2000): 1105-45; Matthew Adler and Eric A. Posner, "Rethinking Cost-Benefit Analysis," *Yale Law Journal* 109 (1999): 167-246.

42. As we shall see, it matters a great deal whether it actually so operates; I explore this issue in detail below.

distributions, and for that form of coercion, is not to require people to buy regulatory benefits on terms that they find unacceptable. Suppose that people are willing to pay only \$60 to eliminate a 1/100,000 risk because they are not rich and that, if they had twice their current wealth, they would be willing to pay \$120. Even if this is so, government does people no favors by forcing them to pay the amount that they would pay if they had more money. In ridiculing WTP, Ackerman and Heinzerling devote too little attention to this point.

If we reject the argument from welfare, we might nonetheless rely on WTP on grounds of personal autonomy.⁴³ On this view, people should be sovereign over their own lives, and this principle means that government should respect personal choices about how to use limited resources (again so long as those choices are adequately informed). When people decline to devote more than \$60 to eliminate a 1/100,000 risk, it is because they would prefer to spend the money in a way that seems to them more desirable. If regulators reject people's actual judgments, then they are insulting their dignity. The use of WTP therefore can claim a simultaneous defense from both utilitarian and deontological considerations.

When the assumptions just outlined are met, we have what might be described as easy cases for the use of the WTP criterion. Some people contend that money and health are incommensurable—that our reflective judgments do not permit us to line up dollars and health risks along a single metric.⁴⁴ Suppose that this is so. To see the easy cases as such, it is not necessary to make controversial arguments about commensurability or to venture into controversial philosophical territory. The underlying claim is a simple pragmatic one, to the effect that people are willing to trade money against decreases in statistical risks. If people actually make those trades, then government might well build on their practices in designing policies.

B. Objections

There are several possible objections. Perhaps the most obvious would point to people's rights. On one view, people have a right not to be subjected to risks of a certain magnitude, and the use of WTP will violate that right. In fact it is fully reasonable to say that, whatever their WTP, human beings have a right not to be subject to risks above a particular level. Imagine, for example, that poor people live in a place where they face a 1/20 annual risk of dying from water pollution; it makes sense

43. See Ronald Dworkin, *Sovereign Virtue* (Cambridge, Mass.: Harvard University Press, 2002).

44. See Elizabeth Anderson, *Value in Ethics and in Economics* (Cambridge, Mass.: Harvard University Press, 1994).

to say that the government is required to reduce that risk even if people are willing to pay only \$1 to eliminate it and the per-person cost is \$5. The only qualification is that, in practice, rights are resource dependent. What rights people are able to claim, against their government, is a product of the amount of available money, and hence people's legitimate arguments for protection are inevitably affected by the level of resources in the society. But let us simply assume here that risks above a certain level should count as violative of rights.

It might be added that people have a right not to be subjected to the intentional or reckless imposition of harm, whatever their WTP. If a company subjects the citizens of a town to a high danger, and it does so maliciously or without the slightest concern for their welfare, the rights of those citizens have been violated, even if their WTP is low. Indeed, some such systems impose strict liability for harms.

As abstract claims about people's rights, these objections are entirely plausible. Something has gone badly wrong if people are exposed to serious risks and if their low WTP prevents them from doing anything in response. Things are even worse if government uses their low WTP to justify inaction in the face of those risks. It would be ludicrous to suggest that WTP is determinative of the appropriate use of government subsidies; a redistributive policy hardly tracks people's WTP. (Would it make sense to say that government will give poor people a check for \$100 only if they are willing to pay \$100 for that check?)

In many cases of environmental regulation, however, rights violations are not involved; we are speaking here of statistically small risks. Even if rights are involved when people are subject to small risks, people should be permitted to waive those rights at an agreeable price (at least on the assumptions that I am making). The proper response to an apparent rights violation is not to force people to buy protection that they do not want but to provide a subsidy that will give them the benefit for free or enable them to receive the benefit at what is, for them, an acceptable price. But regulation—and this is the key point—often does no such thing; and for the easy cases, the question is one of regulation under the stated assumptions. So long as that is the question, use of WTP does not violate anyone's rights.

What about environmental wrongdoers? If a company has intentionally, recklessly, or negligently exposed people to harm, it should be held accountable through the payment of damages, even if the WTP of the affected population is low. It is possible for the tort system to go much further. A sensible legal system might well choose to force companies to internalize the costs of their activities by requiring them to pay for the harms they have caused, even if there has been neither intentional nor negligent wrongdoing. Within tort theory, there is an active debate on this question, and it is possible to support strict liability

by reference to a range of theoretical positions.⁴⁵ But the subject here is regulation, not compensation via the tort system. It would be odd to say that people have a right to be required to pay more for risk reduction than they are willing to pay, at least if they are adequately informed. If people are willing to pay only \$25 to eliminate a risk of 1/100,000, a reference to their rights cannot plausibly justify the conclusion that government should impose a regulation that costs them \$75.

An independent objection would stress, as Ackerman and Heinzerling do, that people are citizens, not merely consumers. On this view, regulatory choices should be made, not after aggregating WTP, but after citizens have deliberated with one another about their preferences and values. The argument against forced exchanges treats people as consumers; it sees their decisions about safety as the same as their decisions about all other commodities. For some decisions, this approach is badly misconceived. The American constitutional system is a deliberative democracy, not a maximization machine, and many social judgments should be made by citizens engaged in deliberative discussion with one another rather than by aggregating the individual choices of consumers.⁴⁶

In the context of racial and sex discrimination, for example, sensible societies do not aggregate people's WTP. The level of permissible discrimination is not set by using market evidence to see how much people would be willing to pay to discriminate (or to be free from discrimination). Even if discriminatory employers would be willing to pay a great deal to avoid associating with members of unpopular groups, such discrimination is banned. Nor is the protection of endangered species chosen on the basis of aggregated WTP. Whether and when to protect members of endangered species is a moral question to be resolved through democratic discussion, not through exercises in consumer sovereignty. In many environmental contexts, use of WTP would wrongly see people as consumers, purchasing products, rather than as citizens, deliberating about values. Speaking in this vein, Amartya Sen emphasizes that "discussions and exchange, and even political arguments, contribute to the formation and revision of values."⁴⁷ He urges that, in the particular context of environmental protection, solutions require us "to go beyond looking only for the best reflection of existing

45. See Richard A. Epstein, "A Theory of Strict Liability," *Journal of Legal Studies* 2 (1973): 151–82; Richard Posner, *Economic Analysis of Law*, 4th ed. (New York: Aspen, 1992), pp. 175–82.

46. See William Bessette, *The Mild Voice of Reason* (Chicago: University of Chicago Press, 1992).

47. Amartya Sen, *Rationality and Freedom* (Cambridge, Mass.: Harvard University Press, 2001), p. 287.

individual preferences, or the most acceptable procedures for choices based on those preferences.”⁴⁸

Sen’s claims identify some serious limitations on the use of WTP. But such objections should not be read for more than they are worth. In trading off safety and health in our own private lives, we do not have static values and preferences. Much of the time, our choices are a product of reflection, even if we are simply acting as consumers. To be sure, moral questions are not to be resolved by aggregating private WTP. Some preferences, even though backed by WTP, are morally off-limits, and policy should not take account of them. In addition, many people are unwilling to pay a great deal for goods that have strong moral justifications; animal welfare is an example. In these circumstances, the market model is inapplicable, and WTP tells us very little.

But what about the easy cases? Do these arguments suggest that government should override individual choices about how much to spend to eliminate low-level risks, even when those choices are adequately informed? For environmental protection generally, it is indeed important to go beyond “the best reflection of existing individual preferences.” But this point does not establish that people should be required to pay (for example) \$100 to eliminate mortality risks of 1/100,000 when they are willing to pay only \$75. If people’s WTP reflects impulsiveness, recklessness, an absence of information, or insufficient deliberation, then it is important for other people, in government as elsewhere, to draw their attention to that fact. And in some cases, a low WTP might be overridden on the ground that it is rooted in errors, factual or otherwise. But these points should be taken neither as a general objection to my conclusion about the easy cases nor to suggest that government should force people to reduce statistical risks at an expense that they deem excessive.

A final objection would emphasize the possibility that people’s preferences have adapted to limitations in existing opportunities, including deprivation.⁴⁹ Perhaps people show a low WTP for environmental goods, including health improvements, simply because they have adjusted to environmental bads, including health risks. Perhaps people’s WTP reflects an effort to reduce cognitive dissonance through the conclusion that risks are lower than they actually are.⁵⁰ To generalize the objection, perhaps people suffer from a problem of “miswanting”; they want things that do not promote their welfare, and they do not want things that

48. *Ibid.*, p. 289.

49. See Jon Elster, *Sour Grapes* (Cambridge: Cambridge University Press, 1983); Martha Nussbaum, *Women and Human Development* (Cambridge: Cambridge University Press, 2001).

50. See George A. Akerlof, *An Economic Theorist’s Book of Tales* (Cambridge: Cambridge University Press, 1984), pp. 123–37.

would promote their welfare.⁵¹ If this is so, then WTP loses much of its underlying justification; people's decisions do not actually promote their welfare, and their autonomy, properly understood, may not require respect for their decisions, which may be nonautonomous.⁵² In other words, the idea of autonomy requires not merely respect for whatever preferences people happen to have, but also social conditions that allow preferences to be developed in a way that does not reflect coercion or injustice. With respect to some risks, the relevant preferences are nonautonomous; consider the fact that many women face a risk of male violence under circumstances in which they believe that little can be done and hence adapt. If government can be confident that people are not willing to pay for goods from which they would greatly benefit, government should probably abandon WTP.

In the context of ordinary regulatory policy, however, this objection has more theoretical than practical interest. Typically we are speaking of steps that would reduce low-level mortality risks (say, 1/50,000). Much of the time, there is no reason to believe that the use of informed WTP (say, \$100) is a product of adaptive preferences. When there is such a reason, the judgment about the easy cases must be revised.

C. Harder Cases

There is an obvious artificiality in the assumptions behind the easy cases. Most important, people do not always bear the full social costs of the regulatory benefits they receive. Sometimes they pay only a fraction of those costs—or possibly even nothing at all. When this is so, the analysis is much more complicated. In the context of air pollution regulation, for example, there is a complex set of distributional effects, and, on balance, poor people and members of minority communities appear to be net gainers.⁵³ A CBA, based on WTP, might not produce an adequate account of the welfare effects of air pollution regulation.⁵⁴ And even if it does, an account of welfare effects would not end the normative question, because the distributional gains are important to consider.

Suppose, for example, that beneficiaries of a proposed drinking water regulation are willing to pay only \$80 to eliminate a risk of 1/

51. Daniel T. Gilbert and T. D. Wilson, "Miswanting," in *Thinking and Feeling: The Role of Affect in Social Cognition*, ed. Joseph P. Forgas (Cambridge: Cambridge University Press, 2000), pp. 178-92; Timothy D. Wilson and Daniel T. Gilbert, "Affective Forecasting, Advances in Experimental Social Psychology," June 2003, at p. 345.

52. For general discussion, see Daniel Kahneman, Peter P. Wakker, and Rakesh Sarin, "Back to Bentham? Explorations of Experienced Utility," *Quarterly Journal of Economics* 112 (1997): 375-405, pp. 375, 379-80.

53. See Matthew E. Kahn, "The Beneficiaries of Clean Air Act Regulation," *Regulation* 24 (2001): 34-37.

54. On the relationship between welfare and CBA, see Gibbard; Adler and Posner.

50,000 in drinking water, that the per-person cost of eliminating a 1/50,000 risk is \$100, but that, for every dollar of that cost, the beneficiaries pay only 70 cents. The remaining 30 cents might be paid by water companies themselves, in the form of reduced profits, or by employees of the water companies, in the form of reduced wages and fewer jobs. In this example, the costs of the regulation exceed the benefits; it is inefficient. But, by hypothesis, the regulation makes its beneficiaries of the regulation better off. If CBA provides the rule of decision, and if the WTP criterion is used, the fact that the monetized costs exceed the monetized benefits is decisive. But as a normative matter, the analysis here is far harder than in the easy cases. On what assumption should the WTP numbers be decisive?

The assumption must be that economic efficiency is the goal of government, at least in the context of environmental regulation—that in order to know what to do, we should aggregate the benefits and costs of regulation, and act if and only if the benefits exceed the costs. When using the WTP numbers, government is acting as a maximization machine, aggregating all benefits and costs as measured by the WTP criterion. But this is an implausible understanding of what government should be doing. In fact, it represents a shift from the relatively uncontroversial Pareto criterion, exemplified in the easy cases, to a version of the far more controversial Kaldor-Hicks criterion, which assesses policy by asking this question: are the gainers winning more than the losers are losing?⁵⁵ The Kaldor-Hicks criterion is sometimes described as potential Pareto superiority, because it asks whether, in principle, the winners could compensate the losers and a surplus could be left over. The difficulty of course is that Pareto superiority is merely potential. Some people really are losing and others are gaining.

In the harder cases, the gainers are gaining less (in monetary terms) than the losers are losing—and hence CBA suggests that regulation is unjustified. Under the assumptions I have given, the regulation is indeed inefficient: its social cost is higher than its social benefit. But is the regulation undesirable? This is not at all clear. The first problem is that WTP is measuring gains and losses in monetary terms, rather than in welfare terms. It is possible that those who gain, in the harder cases, are gaining more welfare than the losers lose; WTP is not dispositive on that question.⁵⁶ The second problem is distributional. Suppose that in terms of overall welfare, the regulation is not desirable; it makes

55. It is only a version of that criterion, because it is measuring welfare in monetary equivalents. A direct assessment of welfare, if it were possible, might show that the regulation in question is justified under the Kaldor-Hicks criterion.

56. On CBA and welfare, and for a welfarist defense of CBA as a decision procedure, see Adler and Posner.

aggregate welfare lower rather than higher. But suppose too that those who benefit are poorer and more disadvantaged than those who lose. If, for example, those who are willing to pay \$80 are disproportionately poor, and those who pay the remainder are disproportionately wealthy, the regulation might be plausibly justified despite the welfare loss.

It is natural to respond here that if redistribution is what is sought, then it should be produced not through regulation but through the tax system, which is a more efficient way of transferring resources to those who need help.⁵⁷ But suppose that redistribution is not going to happen through the tax system. If so, then the regulation in the harder cases cannot be ruled off-limits despite its inefficiency. The fact that a regulation is helpful to the most disadvantaged is not decisive in its favor. If it is trivially helpful, and if it inflicts huge costs on everyone else, little can be said for it. But everything depends on the magnitude of the relevant effects. A program that produces large gains for the least well-off would seem to be justified even if it imposes, in terms of WTP, slightly higher costs than benefits on balance.

The simple conclusion is that the argument for using WTP is most plausible in cases in which the beneficiaries of regulation pay all or most of its cost. In such cases, WTP is reasonably used so long as people are adequately informed and so long as the question is how much they should be forced to spend to avoid existing risks. The analysis must be different when the beneficiaries of regulation are paying only a small fraction of its costs. In such cases, it is possible that the regulation can be justified as a redistributive measure or on welfare grounds. To know whether it can be so justified, it is necessary to go beyond CBA and to identify the winners and losers.⁵⁸ Ackerman and Heinzerling neglect to distinguish between the easy cases and the harder ones and thus fail to appreciate the arguments that lie behind use of WTP.

D. On Individuation and Discount Rates

Many of the most interesting arguments offered by Ackerman and Heinzerling are best seen as attacks on CBA as currently practiced, not on CBA as such. Consider, for example, their suggestion that it is foolish to extrapolate, from workplace studies, a single figure for the value of

57. See, e.g., Louis Kaplow and Steven Shavell, "Why the Legal System Is Less Efficient Than the Income Tax in Redistributing Income," *Journal of Legal Studies* 23 (1994): 667–81, p. 667.

58. A possible defense of CBA would be that individualized identification of winners and losers would be extremely difficult, that agencies would not make defensible distributional judgments, and that CBA works better, on balance, than any approach that attempts to make finer distinctions. A great deal of empirical work would be necessary to make this defense convincing. My goal here is to sort out the theoretical issues; CBA might look especially good, or especially bad, when we investigate empirical issues as well.

statistical risks. Ackerman and Heinzerling note that people care not only about the magnitude of the risk (is it 1/10,000 or 1/100,000?) but also about its nature and context. A risk of death from cancer might well be worse than a statistically equivalent risk of a sudden, unanticipated death. A risk of death from air pollution, or drinking water, might well produce a higher WTP than a statistically equivalent risk of death from a workplace accident.

This claim is plausible, but it is most sensibly taken as an argument for a more refined version of CBA, one that insists on variations among statistically equivalent risks.⁵⁹ A single number is genuinely obtuse; in fact it is inconsistent with the very theory that gives rise to the use of WTP in the first place. Recall that if WTP is relevant, it is because its use promotes welfare, autonomy, or both. If this is so, regulators should consult actual WTP, which varies across risks, rather than a single or unitary WTP, which grows only out of one set of risks, and which (as Ackerman and Heinzerling say) cannot plausibly be applied to every risk of a given statistical magnitude. The real question is not whether to have more differentiated monetary values for qualitatively identical risks but where to find reliable evidence on which to base those values. Economists are starting to fill the relevant gaps, in a way that supports the suggestion that a single WTP is far too crude.⁶⁰ It would not be difficult to continue to use WTP and to take account of the fact that it varies across risks, even if they are statistically identical.

Ackerman and Heinzerling also object, plausibly, to the application of the standard discount rate for money to the valuation of future gains in terms of both mortality and morbidity.⁶¹ But suppose that no discount rate is appropriate—that deaths in 2050 should be valued the same way as deaths in 2010. If so, the analysis of costs and benefits would not be the same; but it would remain possible to calculate both costs and benefits. In any case, the analysis of discounting must make a distinction to which Ackerman and Heinzerling devote insufficient attention.⁶² Sometimes environmental regulation protects living people from latent harms—risks that will come to fruition not now, but ten, twenty, or thirty years from now. It seems clear that some discount rate should be applied to latent harms. Most people, intuitively and on reflection, would much

59. I develop this point in *Laws of Fear*.

60. *Ibid.*

61. See Tyler Cowen and Derek Parfit, "Against the Social Discount Rate," in *Justice between Age Groups and Generations*, ed. Peter Laslett and James S. Fishkin (New Haven, Conn.: Yale University Press, 1992), pp. 144–45.

62. See the illuminating discussion in Richard L. Revesz, "Environmental Regulation, Cost-Benefit Analysis, and the Discounting of Human Lives," *Columbia Law Review* 99 (1999): 941, 962–74.

prefer to face a harm in the future rather than immediately.⁶³ Hence some kind of discount rate makes a great deal of sense for harms that will not come to fruition for a long time (even if the discount rate for such harms does not turn out to be the same as the discount rate for money). But sometimes environmental regulation protects members of future generations; and this is a quite different problem. If a program would save one hundred people born in 2020, it is not clear that it deserves less enthusiasm than a program that would save one hundred people born in 2002.

At the same time, a refusal to use a discount rate creates a number of logical and practical conundrums, especially if it would impose high costs on current generations—a particular problem in light of the fact that if current generations face high costs, posterity is likely to be hurt too. Hence aggressive regulation, at least if it is extremely costly, may not help future generations at all. In any case, future generations are almost inevitably going to be wealthier than our own; should the relatively poor present redistribute resources to the relatively rich future? It has also been argued that if regulators are indifferent as to lives saved now or in the future, but discount costs at some positive rate, then it makes sense for them to delay life-saving expenditures indefinitely.⁶⁴ In any case, it has been suggested that instead of discounting lives, regulators might simply use the future discounted (monetary) cost of saving lives at the time when lives are saved—an approach that is mathematically identical and hence produces the same analysis.⁶⁵

I cannot resolve here the complex questions raised by individuation of WTP and by discount rates.⁶⁶ Ackerman and Heinzerling are right

63. It is difficult to separate, in practice, the fact that people would prefer more life-years rather than fewer from the fact that people would prefer a distant death rather than an immediate one. If most people would much prefer a fatal illness in 2050 to a fatal illness in 2010, there may be no discounting; the choice might be based on the fact that the later date ensures more life-years. The question of morbidity presents a simpler test of discounting: most people would rather have a month of illness in 2010 than a month of illness in 2005.

64. Emmett B. Keeler and Shan Cretin, "Discounting of Life-Saving and Other Non-monetary Effects," *Management Science* 29 (1983): 300–318. Ackerman and Heinzerling discuss this claim and reject it (pp. 193–94), in part on the ground that allowing numerous current deaths would be politically unacceptable; but the claim is one of the logical implication of refusing to discount, and the fact that it entails a politically unacceptable outcome does not mean that it is wrong.

65. See Morrall, "Saving Lives." Note that many people believe that because of technological advances, future risks are unlikely to come to fruition, simply because new technologies will permit us to prevent them. This is not, however, a point about discounting itself.

66. An especially good discussion is found in Revesz.

to raise questions about existing practice.⁶⁷ It seems clear that greater individuation is justified for statistically equivalent risks that people consider to be qualitatively different; it also seems clear that a discount rate should be used for latent harms. The appropriate approach to risks faced by future generations remains unsettled. The simplest point is that if these questions are properly answered, CBA will be mended, not ended.

VI. UNCERTAINTY, CATASTROPHE, AND MAXIMIN

Thus far I have emphasized issues of monetization, as highlighted by Ackerman and Heinzerling. But as Posner's discussion demonstrates, some of the most interesting problems raised by CBA have nothing to do with the translation of risks into monetary equivalents. Recall that Posner's assessment of the risks of particle accelerators contains what he himself describes as arbitrariness. Nor is the problem limited to unconventional problems of this kind. Consider a fairly mundane issue under the Safe Drinking Water Act. If government reduces permissible levels of arsenic in drinking water from 50 parts per billion (ppb) to 10 ppb, what, exactly, are the expected benefits? On the basis of existing evidence, many answers are scientifically respectable.⁶⁸ For a regulation mandating that reduction, the EPA estimated that it would prevent about twenty-five premature deaths and roughly an equivalent number of non-fatal cancer cases. But on the basis of the same inconclusive evidence that was before the EPA, it would have been reasonable to project that the regulation would prevent as few as six deaths or as many as 110. The evidence suggested a range, not a specific estimate.

As a result, critics of CBA might contend the method gives only the illusion of precision. Even before deaths are translated into monetary equivalents, regulators might well be required to make judgments of value, not merely fact, in projecting the likely effects of regulatory protection. But this point should not be read for more than it is worth. When specific estimates are not feasible, the evidence often permits agencies to specify a range. For the arsenic rule, they could say, for example, that a 10 ppb standard is likely to prevent a minimum of six and a maximum of 110 deaths, and they might undertake CBA with reference to the range. Such an analysis would not resolve the question of what to do, but it would greatly discipline the inquiry.

In some circumstances, however, existing information puts regu-

67. They are also right to contend that regulators should consider not only the WTP of those who face mortality risks, but also that of others who care about them. See Eric A. Posner and Cass R. Sunstein, "Dollars and Death," *University of Chicago Law Review* (forthcoming).

68. See Sunstein, "The Arithmetic of Arsenic."

lators in a far more difficult situation. These are cases of genuine uncertainty, in which probabilities cannot be assigned to the expected outcomes. Posner is much concerned with these situations. While he does not spell out the argument, his treatment of catastrophic risks points to a promising possibility for a narrower, and more appealing, version of the Precautionary Principle, a kind of Anti-Catastrophe Principle. Suppose that citizens face catastrophic risks to which probabilities cannot be assigned; suppose, that is, that they are operating under conditions of uncertainty rather than risk. If regulators are operating under such conditions, they might well do best to follow maximin, identifying the worst-case scenarios and choosing the approach that eliminates the worst of these. It follows that, if aggressive measures are justified to reduce the risks associated with global warming, one reason is that those risks are potentially catastrophic and existing science does not enable us to assign probabilities to the worst-case scenarios. Maximin is an appealing decision rule whenever uncertainty is present, but in the regulatory context, it is particularly important for extremely bad outcomes. When Ackerman and Heinzerling suggest the value of focusing on the worst case, they are offering unfortunate advice under circumstances of risk; but if they are understood to be speaking of uncertainty, they are on much firmer ground (see pp. 225–26).

In an extremely illuminating effort to recast the Precautionary Principle,⁶⁹ Stephen Gardiner adapts John Rawls's argument that when "grave risks" are involved, and when probabilities cannot be assigned to the occurrence of those risks, maximin is the appropriate decision rule, at least if the chooser "cares very little, if anything, for what he might gain among the minimum stipend that he can, in fact, be sure of by following the maximin rule."⁷⁰ Applying Rawls's claims about the original position and distributive justice to the environmental setting, Gardiner urges that maximin, and hence a "core" Precautionary Principle, is justified (1) in the face of potentially catastrophic outcomes, (2) where probabilities cannot be assigned, and (3) when the loss, from following maximin, is a matter of relative indifference. Gardiner adds, sensibly, that to justify maximin, the threats that are potentially catastrophic must satisfy some minimal threshold of plausibility. If they can be dismissed as unrealistic, then maximin should not be followed. Gardiner believes that the problem of global warming can be usefully an-

69. See Stephen Gardiner, "A Core Precautionary Principle" (unpublished manuscript, University of Washington, 2004).

70. See John Rawls, *A Theory of Justice* (Cambridge: Harvard University Press, revised edition 1999), p. 134. Rawls draws in turn on William Fellner, *Probability and Profit* (Homewood, Ill.: Irwin, 1965).

alyzed in these terms and that it presents a good case for the application of maximin.

This argument seems to me on the right track, but its conclusion, as stated, risks triviality, above all because of condition 3. If individuals and societies can eliminate an uncertain danger of catastrophe for essentially no cost, then of course they should eliminate that risk. But the real world rarely presents problems of this form. In real disputes, the elimination of uncertain dangers of catastrophe imposes both costs and risks. In the context of global warming, for example, it is implausible to say that regulatory choosers can or should care “very little, if anything,” for what might be lost by following maximin. If we followed maximin for global warming, we would spend a great deal to reduce greenhouse gas emissions, and the result would almost certainly be higher prices for gasoline and energy, probably producing increases in unemployment and poverty.

For environmental problems, does Gardiner’s argument for maximin provide help beyond the trivial cases? I believe that, if properly reformulated, it does, for one simple reason: condition 3 is too stringent and should be abandoned. Even if the costs of following maximin are significant, and even if choosers care a great deal about incurring those costs, it makes sense to follow maximin when they face uncertain dangers of catastrophe. The hardest question here is: under circumstances of uncertainty, how much cost does it make sense to incur in the service of maximin? Consider a straightforward case: the catastrophic dangers associated with global warming could be eliminated if every nation contributed \$2 million to a fund to combat that risk. Surely that cost would be acceptable. Consider a very different case: the catastrophic dangers associated with global warming could be eliminated only if every nation contributed enough resources to reduce standards of living by 50 percent world-wide, with a corresponding increase in global poverty. If global warming really does pose an uncertain danger of total catastrophe, the logic of maximin argues in favor of this extraordinary reduction in world-wide standards of living; but it is not clear that following that logic would be reasonable. To incur costs of this magnitude, we might want to insist that the danger of catastrophe rises above a minimal threshold—that there be demonstrable probability, and a not-so-low one, that the catastrophic risk will occur. It would seem far more sensible to take less costly steps now and to engage in further research, attempting to learn enough to know more about the probability that the catastrophic outcomes will occur.⁷¹

71. Sometimes the Precautionary Principle is urged in situations in which one or another course seems irreversible; and Posner attends to issues of reversibility in discussing global warming. But the issue of irreversibility raises many problems. Any death, of any

For global warming, Posner's interesting recommendations seem vulnerable for this reason. Though a firm defender of quantification, he offers too little in the way of numbers here. To evaluate his proposal for new taxes on greenhouse gas emissions, designed to produce technological innovation, it would be valuable to know both the costs of that initiative and the likely benefits. Assessment of costs would not be easy, because we cannot project the rate of technological innovation; but if the taxes are significant, large increases should be expected in the price of energy, including gasoline, with particularly serious effects on poor people. Because of the range of uncertainties in the science of global warming, assessment of benefits is even harder. But at the very least, it should be possible to measure the likely effects of such taxes on greenhouse gas emissions. If the relevant taxes can be projected to spur significant reductions, then the argument for them is certainly strengthened.

Unlike Posner, I suspect that the likelihood of real catastrophe from global warming is low, and hence that he is wrong to say that no probability can be assigned to it. But I am far from an expert on the underlying science, and in any case Posner convincingly argues that some kind of positive tax on carbon emissions would be cost justified. The larger point is that an Anti-Catastrophe Principle has a legitimate place in environmental regulation, applying to uncertain dangers of catastrophe, at least when the costs of reducing those dangers are not huge and when incurring those costs does not divert resources from more pressing problems. The Anti-Catastrophe Principle is not the Precautionary Principle; it is far narrower than that, and it covers only a small set of environmental problems. But it nonetheless deserves to play a

living creature, is irreversible, and those who invoke irreversibility do not intend the notion of irreversible harm to apply to each and every mortality risk. And because time is linear, every decision is, in an intelligible sense, irreversible. If I play tennis at 11 A.M. today, that decision cannot be reversed, and what might have been done at that time will have been permanently lost. If government builds a new highway in upstate New York in May, that particular decision will be irreversible, even though the highway can be replaced or eliminated. Those who are concerned about irreversibility have something far more particular in mind. They mean something like a large-scale alteration in environmental conditions, one that imposes permanent, or nearly permanent, changes on those subject to them. But irreversibility in this sense is not a sufficient reason for a highly precautionary approach. At a minimum, the irreversible change has to be for the worse, and it must also rise to a certain level of magnitude. A truly minuscule change in the global temperature, even if permanent, would not justify expensive precautions if it is benign or if it imposes little in the way of harm. The idea of irreversibility is really important for two reasons. The first, referred to by Posner, draws on the analogy to stock options, and suggests that it is worthwhile to spend resources on (bounded) precautions to wait for more information to emerge before incurring a substantial and irreversible loss. The second reason involves the relationship between irreversibility and catastrophic harm; a harm is unlikely to be catastrophic if it can be reversed. For discussion, see Sunstein, *Laws of Fear*.

role in environmental protection, plausibly including global warming—calling for significant steps now, accompanied by further research to obtain a better understanding of the likelihood of real disaster.

Four qualifications are important. First, the Anti-Catastrophe Principle must be attentive to the full range of social risks; it makes no sense to take steps to avert catastrophe if those very steps would create catastrophic risks of their own. Second, use of the principle should be closely attentive to the idea of cost-effectiveness, which requires regulators to choose the least costly means of achieving their ends. In the context of global warming, there are many methods by which to reduce the relevant risks. Both nations and international institutions should choose those methods that minimize costs. Third, distributional considerations matter. The principle should be applied in a way that reduces extreme burdens on those least able to bear them. For global warming, there is a particular need to ensure that citizens of poor nations are not required to pay a great deal to contribute to the solution of a problem for which wealthy nations are most responsible—partly because the latter caused the problem in the first place, but also because poor people, faced with a global risk, need and deserve help from those who can provide it. Fourth, costs matter. The extent of precautions cannot reasonably be divorced from their expense. When the worst-case scenario is truly catastrophic and when probabilities cannot be assigned, a large margin of safety makes a great deal of sense.

VII. CONCLUSION

Because regulation itself often introduces new hazards, the Precautionary Principle risks incoherence; it forbids the very steps that it requires. For its part, CBA runs into two serious difficulties. The first involves the specification of both probabilities and outcomes. The second involves the translation of environmental risks into monetary equivalents.

For many of the problems involved in environmental regulation, it is possible to identify a range of outcomes in a way that allows CBA to get off the ground. But when catastrophic outcomes are possible and when regulators are operating under circumstances of uncertainty, it may well make sense to follow maximin. Even in such circumstances, however, an inquiry into costs cannot sensibly be avoided, not least because nations that impose high costs might increase mortality and morbidity risks as a result.

With respect to monetization, I have suggested that it is important to distinguish between the easy cases for using WTP and the harder ones. When the beneficiaries of environmental regulation pay all or most of its cost, the argument for using WTP is especially strong. In such cases, beneficiaries are unlikely to be helped if they are required to pay more than they are willing to pay; and requiring them to do so

is an insult to their autonomy.⁷² But when the beneficiaries of environmental regulation pay little or none of the cost, the regulation might be justified even if it fails CBA. To decide whether it is, it is necessary to identify the likely winners and losers. The most general conclusion is that CBA does not tell regulators all that they need to know; but without it, they will know far too little.

72. As I have emphasized, a low WTP does not mean that government should not subsidize the good; but regulation is not a subsidy.