

government should not react hysterically to people's fears — not least because a hysterical response would give terrorists a large additional victory. I will turn to this issue in due course. For now the important point to notice is that people's approach to the risks of catastrophe, including the risks associated with terrorism, can be better understood if we have a working knowledge of how people assess risks in general.

Reducing Risks Rationally

Often it is possible to resolve hard questions of law and policy without resolving deeply contested issues about justice, democracy, or the appropriate aims of the state. Often it is possible for people who disagree on many topics to agree on a social practice. In many areas of law and public policy, people can reach closure about what they should do despite their disagreement or uncertainty about why, exactly, they should do it. When closure is reached in this way, we might say that people find their way toward *incompletely theorized agreements*.¹

Thus people who sharply disagree about the purposes of the criminal law can agree that rape and murder should be punished, and punished more severely than theft and trespass. People can support an Endangered Species Act even if they disagree about whether the protection of endangered species is desirable for theological reasons; because of the rights of animals, plants, and species; or because of the value of animals, plants, and species for human beings. Incompletely theorized agreements are hardly foreign to government regulation. In the area of regulatory reform, deregulation of the trucking and airlines industry commanded agreement from many diverse people; in fact, it was spearheaded by Senator Edward M. Kennedy, who was enthusiastically joined by many conservatives. A great advantage of incompletely theorized agreements is that they allow people of diverse views to live together on mutually advantageous terms. An even greater advantage is that they allow people of diverse views to show one another a high degree of both humility and mutual respect.

In this chapter, I attempt to outline, in a preliminary and general way, government responses to the problems discussed thus far. The basic tasks are to find ways to ensure against cascade effects based on insufficient evidence; to promote governmental responses to significant risks even if the public has

¹ I discuss this point in detail in Cass R. Sunstein, *Legal Reasoning and Political Conflict* (New York: Oxford Univ. Press, 1996).

shown little interest in reducing them; and to place tradeoffs on-screen, including the risks that are created by risk regulation. Three principles, sketched in the introduction, provide the basic orientation. First, government should attempt to assess the magnitudes of risks; it should not concern itself with risks that are statistically small. Second, government should examine all the effects of risk reduction, including costs and additional hazards created by risk reduction itself. Third, government should explore alternatives to any proposed action, including alternatives that involve less intrusive tools for achieving the same basic ends. To adopt these principles is to move much of the way toward cost-benefit analysis, and to do so without saying anything especially controversial about the proper ends of government.

But let us begin with some ideas that, while extremely influential, seem to me quite inadequate: pollution prevention, the precautionary principle, and sustainable development. Each of these ideas is often proposed as an alternative to the analytical techniques of cost-benefit analysis – an alternative that is less “cold,” more attuned to potential harms to real people, and more friendly to the environment. Despite their popularity, and their extraordinary international prominence, I think that the three ideas are unhelpful, sometimes even ludicrous. An understanding of why this is so helps to prepare the way for the more affirmative discussion that follows.

POLLUTION PREVENTION, RISK PREVENTION

Should pollution be “prevented” rather than “cured”? With respect to many problems, including illness, prevention does indeed seem best – cheapest and most effective. It is usually better to have a flu shot than to treat flu after the fact. For most people, a good diet and exercise, alongside a refusal to smoke, are a lot better than heart surgery and chemotherapy. Perhaps prevention should be the preferred approach in the domain of social risks.

Barry Commoner, a well-known scientist and environmentalist, has urged that government regulation of risks should be fundamentally remodeled, with the idea of “prevention” at its core.² The idea has been influential in many circles. In fact Congress has enacted a statute – unimaginatively but informatively named the Pollution Prevention Act – that puts a premium on pollution prevention.³ The goal of pollution prevention is to ensure that regulators prevent pollution before it even enters the system. As real-world examples, consider the phaseout of lead in gasoline, the use of solar power, and the substitution of electric cars for cars powered by gasoline. According to Commoner and many others, this

approach is far more promising than “end-of-the-pipe” controls imposed on polluting technologies. Advocates of pollution prevention argue that pollution prevention promises larger and more dramatic pollution reductions and does not rely on imperfectly reliable, after-the-fact technological “fixes.”

Thus Commoner contends that the “real improvements have been achieved not by adding controls or concealing pollutants but by simply eliminating them. The reason there is so much less lead in the environment – and in children’s blood – is that lead has been almost entirely eliminated from the manufacture of gasoline. The reason why DDT and related pesticides are now much less prevalent in wildlife and our own bodies is that their use has been banned.”⁴ Prevention works “because it is directed at the origin of the pollutant in the production process itself.” By contrast, controls “yield little or no improvement” because they “are only one element in a larger system that can readily counteract their apparent efficiency.” For Commoner, the overall lesson is that if “you don’t put something in the environment, it isn’t there.” To take one illustration, Commoner urges that if we would really like to reduce nitrogen oxide emissions from cars, we will build “smogless engines that do not produce nitrogen oxides,” rather than relying on catalytic converters.⁵ It would be easy to build on the idea of pollution prevention to suggest that in many contexts, regulators should move toward “risk prevention.” Instead of managing risks that are already in the system, they should attempt to eliminate enterprises and activities that produce (unacceptably high?) risks.

Often pollution prevention makes sense. The EPA was right to eliminate lead from gasoline; the government was also right to stop the use of CFCs, which contribute to destruction of the ozone layer. If asbestos is a potent carcinogen, and if it can be eliminated without causing serious problems, there is a good argument for eliminating asbestos. But sometimes pollution prevention would be extremely unappealing, even ridiculous, simply because it is not worthwhile, all things considered. Consider some examples. The best way to prevent automobile pollution would be to eliminate the internal combustion engines that power most trucks and cars. The best way to prevent pollution from current power sources would be to stop relying on fossil fuels, now used by utility power plants. The best way to prevent the risks of genetically modified plants would be to ban the genetic modification of plants. Should the EPA be told to ban the internal combustion engine and coal combustion? Should national legislatures forbid the genetic modification of plants? If these would be ridiculous conclusions – as I think they would be – it is because the costs of the bans would dwarf the benefits. Pollution prevention is not worthwhile as such; it is worthwhile when it is better, all things considered, than the alternatives.

² See Barry Commoner, *Making Peace with the Planet* (New York: Pantheon Books, 1990).

³ 42 USC 13101.

⁴ Commoner, *supra* note 2 in this chapter, at 42.

⁵ *Id.* at 99.

There is a further point. In many contexts, the idea of pollution prevention, or risk prevention more generally, is literally paralyzing because no approach will actually "prevent" pollution or risk. If the internal combustion engine is banned, substitutes will have to be introduced, and electric cars cause pollution of their own, above all because they currently require considerable energy use. Perhaps electric cars are, on balance, better from the environmental point of view. But if they are better, it is because they produce less (not no) pollution; and that is a different claim from the claim that "prevention" is always or generally best. As we have seen, a ban on the genetic modification of plants might increase overall risks, if it is the case that genetically modified plants reduce risks compared to the alternative.

In fact the strongest argument for pollution prevention rests, at bottom, on cost-benefit balancing. Those who urge pollution prevention are thinking of instances where this approach has large benefits and small costs. In the context of eliminating lead from gasoline, cost-benefit analysis firmly supported pollution prevention, because the benefits dwarfed the costs. The same is true for banning CFCs. But where the balance does not support regulation, pollution prevention is a mistake. Properly understood, pollution prevention is what is recommended by cost-benefit analysis – but only sometimes.

None of this is to deny that sometimes projections of the future will involve a degree of guesswork and speculation. As we have seen, guesswork and speculation are the rule rather than the exception. But when this is so, good cost-benefit analysis calls for attention to the range of possibilities. What is not justified is to "prevent" pollution without an inquiry into the consequences, good and bad, of prevention. The most that can be said is that pollution prevention should always be considered as a possible option and that in some cases, it will turn out to be preferred. But as a general approach to environmental protection, or to the reduction of social risks, it is obtuse. In many contexts, pollution prevention would be literally dangerous, in fact a disaster.

not too strong an argument with cars.

THE PRECAUTIONARY PRINCIPLE

All over the world, there is increasing interest in a simple idea for the regulation of risk: In the case of doubt, follow *the precautionary principle*.⁶ Avoid steps that will create a risk of harm. Until safety is established, be cautious. In a catchphrase: Better safe than sorry. In ordinary life, pleas of this kind seem quite sensible. People buy smoke alarms and insurance. Shouldn't the same approach be followed by regulators as well?

⁶ Protecting Public Health & the Environment: Implementing the Precautionary Principle, Carolyn Raffensberger & Joel Tickner, eds. (Washington, D.C.: Island Press, 1999).

Many people believe so. The Ministerial Declaration of the Second International Conference on the Protection of the North Sea, held in London in 1987, states: "Accepting that in order to protect the North Sea from possibly damaging effects of the most dangerous substances, a precautionary principle is necessary which may require action to control inputs of such substances even before a causal link has been established by absolutely clear scientific evidence."⁷ The closing Ministerial Declaration from the United Nations Economic Conference for Europe in 1990 asserted, "In order to achieve sustainable development, policies must be based on the precautionary principle. . . . Where there are threats of serious or irreversible damage, lack of full scientific certainty should not be used as a reason for postponing measures to prevent environmental degradation."⁸ The widely publicized Wingspread Declaration, from a meeting of environmentalists in 1998, went further still: "When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not established scientifically. In this context the proponent of the activity, rather than the public, should bear the burden of proof."⁹

There is some important truth in the precautionary principle. Sometimes it is much better to be safe than sorry. Certainly we should acknowledge that a small probability (say, 1 in 100,000) of a serious harm (say, 100,000 deaths) deserves extremely serious attention. It is worthwhile to spend a lot of money to eliminate that risk. The fact that a danger is unlikely to materialize is hardly a good objection to regulatory controls. But everything depends on the size of the investment and the speculativeness of the harm. Unless the harm would be truly catastrophic, a huge investment makes no sense for a harm that has a one in one billion chance of occurring. Taken literally, the precautionary principle would lead to indefensibly huge expenditures, exhausting our budget well before the menu of options could be thoroughly consulted. If we take costly steps to address all risks, however improbable they are, we will quickly impoverish ourselves. This is no less true for nations than for individuals.

But there is a larger problem. The precautionary principle can provide guidance only if we blinker ourselves and look at a subset of the harms involved. In real-world controversies, a failure to regulate will run afoul of the precautionary principle because potential risks are involved. But regulation itself will cause potential risks, and hence run afoul of the precautionary principle too; and

⁷ Quoted in Rethinking Risk and the Precautionary Principle 3, Julian Morris, ed. (Oxford: Butterworth-Heinemann, 2000).

⁸ Id. at 5.

⁹ Id.

the same is true for every step in between. Hence the precautionary principle, taken for all that it is worth, is literally paralyzing. It bans every imaginable step, including inaction itself.

A failure to allow genetic modification might well result in many deaths, and a small probability of many more. Hence the precautionary principle seems to argue both for and against banning genetic modification of food. The point is very general. An expensive regulation can have adverse effects on life and health, and hence a multimillion dollar expenditure for "precaution" has – as a worst case scenario – significant adverse health effects, with perhaps many tens of lives lost. If this is so, the precautionary principle seems to argue against the regulation. If the precautionary principle argues against any action that carries a small risk of significant harm, then we should be reluctant to spend a lot of money to reduce risks because those expenditures themselves carry risks. Like the idea of pollution prevention, the precautionary principle, understood in an extreme form, stands as an obstacle to regulation and nonregulation, and to everything in between. Indeed the most sensible understandings of the precautionary principle emphasize the need for an overall assessment and insist on exploring all of the risks at stake, including low-probability, potentially catastrophic risks.¹⁰

Perhaps this argument is missing something. In some contexts, regulation is indeed a form of insurance, or a way of placing special locks on a door. Consider the following choice. Would you rather have

- (a) a sure loss of \$20 or
- (b) a 1 percent chance of losing \$1,980?

In terms of expected value, (b), representing a statistical loss of \$19.80, is a bit less bad than (a); but most people would gladly choose the sure loss of \$20. People do not like to run a small risk of a large or catastrophic loss; this is why we buy insurance and take special precautions against serious harms. Where a risk would be very high, and when we can reduce or eliminate it with cash, it makes sense to endorse the precautionary principle.

If we speak in these terms, however, we are substantially qualifying the precautionary principle, treating it less like a dogma and more as the beginning of a serious analysis of how to approach risks. A competent cost-benefit analysis (CBA) takes good account of the precautionary principle by asking regulators to attend to low-probability risks of significant harms. CBA subsumes this risk, as it does all others, into the overall assessment. A special advantage of CBA is that it incorporates all risks, on all sides of the equation; it therefore prevents

¹⁰ See Indur Goklany, *The Precautionary Principle* (Washington, D.C.: The Cato Institute, 2001).

the kinds of myopia and tunnel vision that are threatened if we take precautions against some risks and thereby fail to take precautions against others. (Consider bans on the use of cell phones in cars, which may increase risks by, for example, preventing people from making or receiving calls in emergencies.) Because CBA is a source of information, and not something to be used mechanically as the basis for decisions, good regulators will often buy the equivalent of insurance. Nothing in CBA precludes a policymaker from concluding that we should suffer a sure loss of, say, \$200,000 to prevent a 1 percent risk of \$19,800,000 in monetary and other losses. Nor does CBA prevent regulators from deciding that a 1/10,000 risk of 100,000 deaths is worse, or less bad, than a 1/1,000 risk of 10,000 deaths. This is a political judgment, not a technical one to be decided by mechanical use of the numbers.

Aaron Wildavsky, a political scientist with a special interest in risk regulation, attacks the precautionary principle on grounds similar to those I have used here.¹¹ Wildavsky convincingly argues that often risks are on both sides of the equation, and hence the idea of "precaution" senselessly guards against one set of risks while ignoring the others. In Wildavsky's view, the notion of "precaution" should not merely be disregarded; it should also be replaced. It should be replaced with a principle of "resilience," based on an understanding that nature, and society, are quite able to incorporate even strong shocks, and that the ultimate dangers are therefore smaller than we are likely to fear. Unfortunately, the principle of "resilience" is no better than that of "precaution." Some systems are resilient, but many are not. Whether an ecosystem, or a society, is "resilient" cannot be decided in the abstract. In any case resilience is a matter of degree. Everything depends on the facts. As we will see, cost-benefit balancing is to be preferred, not because resilience is unlikely, but because it is necessary to inquire into the particular problem to know.

SUSTAINABLE DEVELOPMENT

The notion of "sustainable development" has had an extraordinary influence in international environmental debates, so much as that it now serves as a kind of symbol for any serious commitment to environmental protection. But the notion is highly ambiguous. What kind of development counts as sustainable? What counts as "unsustainable"? In a standard formulation, the idea of sustainable development is said to refer to "development that occurs on a scale that does not exceed the carrying capacity of the biosphere."¹²

¹¹ See Wildavsky, *supra* note 9 in Chapter I.

¹² See Robert Percival et al., *Environmental Regulation: Law, Science, and Policy* 1182 (Boston: Little, Brown & Co., 2000).

To the extent that endorsement of sustainable development is meant as a criticism of approaches that are literally "unsustainable" in the sense that future generations will lack environmental goods – clean air and water, for example – everyone should support sustainable development. Any minimally sensible policy will ensure decent lives and options for future generations. No sane person is opposed to that. Indeed, cost-benefit analysis itself calls for sustainable development, because a competent analyst incorporates the interests of members of future generations. But outside of the easy cases for environmental protection, the real question is not "sustainable" development or "unsustainable" development; it involves what level of resources to commit to environmental protection. Often there is no simple line to divide the sustainable from the unsustainable. If certain regulatory steps would increase "sustainability" but cause a great deal of suffering and misery, simply by virtue of their expense, a sensible administrator will take that point into account. And if the sensible administrator is thinking in this way, he is balancing the various considerations, very much in the way that cost-benefit analysis helps to do. In the international context, the administrator might even ask for financial help from wealthier countries, help that is required from the standpoint of justice.

Because cost-benefit balancing requires consideration of the interests of future generations, the goal of sustainable development is in no conflict with that form of balancing. Indeed, cost-benefit analysis strongly supports the idea that sustainability is a desirable goal. It also helps give content to the hard question, which is how much should be done to improve environmental quality in poor as well as wealthy nations.

To conclude: Pollution prevention can be a good idea, but it would sometimes lead to disaster, not least because it would increase serious risks. Sometimes it is better to be safe than sorry, and here the precautionary principle makes sense. People do, and should, buy insurance. But sometimes the precautionary principle leads to paralysis, and in any case some precautions are simply not worthwhile. Everyone should support sustainable development. But support for sustainable development does not answer any of the hard questions. It is important to ensure that policies are sustainable rather than the opposite. But in poor nations as well as rich ones, regulators need much better guidance than that.

WHAT IS COST-BENEFIT ANALYSIS? WHY FAVOR IT?

For the moment, let us understand cost-benefit analysis to entail a full accounting of the consequences of risk reduction, in both quantitative and qualitative terms. Officials should have this accounting before them when they make decisions. They should not

be bound by the "bottom-line" numbers. But if they are to proceed, they should be prepared to explain how the benefits justify the costs, or if not, why it is nonetheless worthwhile to go forward. I am therefore understanding cost-benefit analysis to require a certain procedure: a quantitative and qualitative accounting of the effects of regulation, together with a duty to explain the grounds for action unless the benefits justify the costs. On this view, the antonym to regulation guided by cost-benefit analysis is regulation undertaken without anything like a clear sense of the likely consequences – or regulation that amounts to a stab in the dark.

The most basic argument for cost-benefit balancing, thus understood, is cognitive. The goal is to overcome cognitive limitations by ensuring that people have a full, rather than limited, sense of what is at stake. People often miss the systemic effects of risk regulation; cost-benefit analysis is a way of putting those effects squarely on-screen. At the same time, cost-benefit analysis helps overcome the problems created by the availability heuristic, by allowing an accounting of the actual consequences of current risks, and of the effects of reducing them. If people are fearful of a risk that is actually quite small, cost-benefit analysis will supply a corrective. If people are indifferent to a risk that is actually quite large, cost-benefit analysis will help to stir them out of their torpor. One of the primary advantages of cost-benefit analysis, properly understood, is that it promotes an understanding of health-health tradeoffs, but it also does much more. To the extent that people's emotions are getting the better of them, by producing massive concern about small risks, cost-benefit analysis should help put things in perspective and at the same time might help to calm popular fears. To the extent that people's perspective is distorted by seeing problems in isolation, cost-benefit analysis can help to put isolated problems in the context of a full range of dangers, thus preventing the kinds of distortions that come from a narrow perspective. The result should be to help with cognitive distortions and to produce sensible priority-setting.

There are democratic advantages as well. As I have emphasized, interest groups often manipulate policy in their preferred directions, sometimes by exaggerating risks, sometimes by minimizing them, sometimes by using heuristics and biases strategically, so as to mobilize public sentiment in their preferred directions. An effort to produce a fair accounting of actual dangers should help to diminish the danger of interest-group manipulation. More generally, cost-benefit analysis should increase the likelihood that citizens generally, and officials in particular, will be informed of what is actually at stake. By itself this is a large democratic gain.

Of course interest groups will also try to manipulate the numbers in their preferred directions. Industry, which is typically well funded, will exaggerate the costs and minimize the risks. Public interest groups will do the opposite.

(I explore this point later in connection with the regulation of arsenic in drinking water.) A government that attempts to produce cost-benefit analysis will face a formidable task; it is possible that government will lack the information necessary to do this task well. But if there is a degree of accuracy, and if ranges are specified where there is uncertainty, cost-benefit analysis can be seen, not as some antidemocratic effort to tyrannize people with numbers, but instead as an indispensable tool of democratic self-government.

Of course, it is possible that in practice cost-benefit analysis will have excessive influence on government decisions, drowning out "soft variables."¹³ Since the numbers are not all that matters, any such effect would be a point against cost-benefit analysis. But to date, the actual record does not support this concern. To take just one example, the EPA's decision to go forward with new controls on arsenic in drinking water was supported partly on the ground that nonquantifiable variables tipped the balance.¹⁴ The risk that cost-benefit analysis will drown out relevant variables is not a reason to abandon the analysis, but to take steps to ensure against any such effect.

ARE RISKS SOCIALLY CONSTRUCTED? OF SCIENCE AND ITS DISCONTENTS

A large goal of cost-benefit analysis is to increase the role of science in risk regulation. This is a controversial goal. Many diverse people have been skeptical of the idea that risk regulation should be grounded, first and foremost, in an understanding of scientific fact. Some of these groups purport to be inspired by democratic, even populist ideals. They suggest that the underlying issues should be resolved by ordinary people, not by a technocratic elite.¹⁵ Some of these critics are skeptical of the whole idea of scientific objectivity. They doubt that risk is an objective concept, measuring something "out there" in the world.¹⁶ Others urge that scientists have values and agendas of their own, and that the effort to strengthen the role of science is actually an effort to fortify the place of those values and agendas. They insist that questions of risk regulation are questions of value, not of fact, and that the real issue is whose values will prevail.

¹³ Lisa Heinzerling, *Clean Air and the constitution*, 20 St. Louis U. Pub. L. Rev. 121, 151 (2001) ("cost-benefit analysis tends to underrate those things that cannot be so quantified and monetized"); Laurence H. Tribe, *Ways Not to Think About Plastic Trees: New Foundations for Environmental Law*, 83 Yale L. J. 1315, 1318-19 & n. 25 (1974).

¹⁴ See Chapte 7; see also *Corrosion Proof Fittings v. EPA*, 947 F. 2d 1201 (5th Cir., 1991).

¹⁵ Kristin Shrader-Frechette, *Risk and Rationality* (Berkeley, Calif.: Univ. California Press, 1991).

¹⁶ See Slovic, *supra* note 12 in Chapter 2, at p. 392.

Respect for "science" is said to be a smokescreen here. Some of these critics include powerful industries purporting to respect the facts but fearful that a careful look at the facts will compromise their economic self-interest; tobacco companies are only the most visible example. In a more theoretical vein, some people insist that risks are "socially constructed," and that any judgment about the magnitude or even the existence of a "risk" is a matter of social judgments, not of something external to what people think.¹⁷

Some of these objections are easy to handle. Of course there is a problem if some groups are manipulating science to promote their own selfish agendas. At the level of principle, little needs to be said, except perhaps that manipulation of this sort is deplorable. The real task is to ensure that any efforts at manipulation will fail. But some of the critiques of scientific objectivity, and some claims about "social construction," raise hard questions. For the moment let us simply insist that, with respect to risks, there really are facts, and it is crucial to know the truth as best we can. Some risks are genuinely serious, in the sense that many people are going to get hurt, get sick, or die as a result. Some risks are not serious, in the sense that few people, or no people, are in danger. Whether a risk is serious, in the sense that many people will be hurt or die, is a question of fact, not of values. Of course values affect any evaluative judgment about the seriousness of risks, but the factual questions are crucial and should be identified and evaluated as such.

Risks are "socially constructed" in the sense that the categories that we use to perceive them are our own. Moreover, both individuals and societies will be greatly concerned about some hazards but treat others as inevitable aspects of life, worth acknowledging but no more. There is a great deal to be said about the relationship between risk and culture, and I will deal very little with that important issue here.¹⁸ What is important for present purposes is that risk is not socially constructed, in any interesting sense, when a certain number of people will (actually) be injured or killed as a result. As best we can tell, hundreds of thousands of people will die each year, in the United States alone, as a result of cigarette smoking. To argue on behalf of cost-benefit analysis, it is necessary only to say that numbers of this sort have real meaning.

Now these points are not meant to deny that scientists have their own values, and that with different assumptions, risks can be estimated as large or small. It is easy to show that with small variations in assumptions, we will come up with dramatically different projections about the future, and that in the face of scientific uncertainty, reasonable people can come up with widely

¹⁷ See id.

¹⁸ For an influential discussion, see Mary Douglas & Aaron Wildavsky, *Risk and Culture* (New York: The Free Press, 1992).

varying judgments about the seriousness of risks. Indeed, Chapter 6 is devoted to an exploration of exactly this issue. Of course the assumptions that produce different judgments will sometimes be a product of values, not simply science. What is the discount rate for future benefits and costs? How should we proceed if we lack knowledge about the rate or nature of climate change? What is the level of risk from exposure to low levels of arsenic, when we have evidence only for high exposure levels? It is extremely important to ensure that assumptions are identified as such, not least to ensure public oversight of how to proceed in the face of uncertainty.

But none of this justifies skepticism about the role of science in risk regulation. There really are facts, whether or not we are able to identify them. Of course we should sometimes be skeptical of scientists' ability to get the facts right; scientists are themselves sometimes skeptical of their ability to get the facts right. When skepticism is justified, the best approach is not to do nothing, to let companies do whatever they wish, to adopt the precautionary principle, to regulate to the hilt, to flip a coin, or even to defer to public values, whatever they happen to be. The best approach is instead to specify the range of possible outcomes, with some sense of the likely probabilities. If there is a small chance of a catastrophic outcome, inaction is hardly the obvious course; people buy insurance, and risk regulation is a form of (preventative) insurance. In some cases, it might not be possible even to specify a range. Here science will leave us at sea. But many cases are not like that.

It bears emphasizing that science cannot by itself resolve normative questions. An understanding of likely consequences cannot resolve issues of value. But democratic judgments should be made with reference to the best understandings of the facts, rather than short-term panics and scare tactics. People should be allowed to see the range of legitimate uncertainty. For these purposes, cost-benefit analysis is an excellent place not to end but to start.

EIGHT PROPOSITIONS

Recall that by cost-benefit analysis, I mean a full accounting of the effects of regulation, both qualitative and quantitative, along with descriptions in both monetary and nonmonetary terms. My basic suggestion is that government should offer that accounting and also make it relevant for purposes of decision. To make these ideas more concrete, here are eight propositions, offered in the hope that they might attract support from diverse theoretical standpoints.

1. Agencies should not only identify the advantages and disadvantages of proposed efforts to reduce risks but also attempt to quantify the relevant effects to the extent that this is possible. When quantification is not possible,

agencies should discuss the relevant effects in qualitative terms. The statement should include the full range of beneficial effects.

2. The quantitative description should supplement rather than displace a qualitative description of relevant effects. Both qualitative and quantitative descriptions should be provided. It is important to know not only raw numbers but also the nature of the relevant effects (e.g., lost workdays, cancers averted, respiratory problems averted). To the extent possible, the qualitative description should give a concrete sense of who is helped and who is hurt (e.g., whether the beneficiaries are mostly or partly children, whether the regulation will lead to lost jobs, higher prices, more poverty, and so forth). Where the only possible information is speculative, this should be noted, along with the most reasonable speculations.

3. To improve the overall evaluation, agencies should attempt to convert nonmonetary values (involving, for example, lives saved, health gains, and aesthetic values) into dollar equivalents. There is nothing magical or rigid about the dollar equivalents. The conversion is simply a pragmatic tool to guide analysis and to allow informed comparisons.

4. Agencies entrusted with valuing life and health should be controlled, by statute or executive order, via presumptive floors and ceilings. For example, a statute might say that a statistical life will ordinarily be valued at no less than \$2 million and no more than \$10 million. Evidence of worker and consumer behavior, suggesting a valuation of between \$5 million and \$8 million per statistical life saved, is at least relevant here. The willingness to pay numbers are in this range, which is certainly relevant if not decisive. Similar numbers appear to represent the midpoint of agency practice. Thus both market and governmental measures point in the same basic direction. It is important in this regard that as the expenditures get very far above \$10 million, regulations threaten to become counterproductive. An expenditure of \$15 million per life saved threatens to take as many lives as it protects.

If Congress does not act, OMB should establish presumptive floors and ceilings for various regulatory benefits. If an agency is going to spend (say) no more than \$500,000 per life saved, or more than \$20 million, it should have to explain itself.

5. Agencies should be permitted to make adjustments in the analysis on the basis of the various "qualitative" factors. For example, agencies might add a "pain and suffering premium" or increase the level of expenditure because children are disproportionately affected or because the victims are members of a disadvantaged group. It would be reasonable to conclude that because AIDS

has disproportionate adverse effects on homosexuals and poor people, special efforts should be made to ensure against AIDS-related deaths. To the extent possible, agencies should be precise about the nature of, and grounds for, the relevant adjustments, especially in light of the risk that interest-group pressures will convert allegedly qualitative adjustments in illegitimate directions. Agencies might incorporate these considerations in a "sensitivity analysis," showing how the numbers are sensitive to changes in assumptions.

one fund

6. Agencies should be required to show that the benefits justify the costs. If they do not, they should be required to show that the action is nonetheless reasonable, on the basis of a publicly articulated explanation. In most cases, the tabulation of costs and benefits should be accompanied by a showing that the benefits justify the costs. If the benefits do not justify the costs, the agency should explain itself. Perhaps the agency could claim that those who bear the costs can easily do so, whereas those who receive the benefits have a special claim to public help. The principal danger here is that well-organized groups will be able to use equitable arguments on behalf of their preferred adjustments. It is important to ensure a degree of discipline here. Perhaps the dangers of interest-group manipulation are serious enough to suggest that uniform numbers or ranges might be used.

7. Under ordinary circumstances, the appropriate response to social fear not based on evidence, and to related ripple effects, is education and reassurance rather than increased regulation. We have seen that public concern about certain risks can be widespread and very intense, even though the concern is not merited by the facts. The best response to unjustified fear is not to capitulate to it but to inform people that the fear is baseless. A government should not expend significant resources merely because an uninformed public thinks that it should. But if education and reassurance fail, increased regulation may be defensible as a way of providing a kind of reassurance in the face of intense fears, which can themselves impose high costs of various kinds. Recall, for example, the possibility that people who are afraid of risks of plane crashes will shift to driving, a more risky method of transportation; consider also the fact that the fear is itself both a problem and a cost. In the aftermath of an actual or apparent terrorist attack, involving airplanes or anthrax, the government might respond quite aggressively, even if it believes that the risk of another attack is very low. One reason for an aggressive response is that a low probability of a significant harm is itself worth addressing. But another reason is that widespread public fear has a wide range of ripple effects, far outrunning the material consequences of the risks themselves.

8. Unless the statute requires otherwise, judicial review of risk regulation should require a general showing that regulation has produced more good than harm, on a reasonable view about valuation of both benefits and costs. On this view, courts should generally require agencies to generate and to adhere to ceilings and floors. But they should also allow agencies to depart from conventional numbers (by, for example, valuing a life at less than \$3 million, or more than \$15 million) if and only if the agency has given a reasonable explanation of why it has done so. The ultimate task would be develop a kind of "common law" of cost-benefit analysis, authorizing agencies to be lawmaking institutions in the first instance.

If all this is correct, it should ultimately be possible to produce a convergence on a form of cost-benefit analysis that should be understood as a pragmatic instrument and that ought not to be terribly contentious—a form of cost-benefit analysis that does not take a stand on highly controversial questions about what government ought to do and that promises to attract support from people with diverse views about what is good and what is right. For those who reject proposals of the sort that I have outlined here, the question remains: On what other basis should government choose one or another approach to risk regulation?

Even if we endorse cost-benefit analysis, it is not enough simply to urge agencies to engage in it. Institutional reforms are necessary from each of the three branches of government. I offer a brief sketch here.

CONGRESS

As an institution, Congress is strikingly ill-equipped to place risks in comparative perspective. One reason is that legislators, like everyone else, are subject to cognitive biases. Another reason is that legislators are subject to the public demand for legislation, which is produced by those very biases as well as by well-organized groups that attempt to steer the legislature in their favor. Members of the national legislature are especially vulnerable to availability cascades, partly because they can produce an irresistible demand for legislation, and partly because isolated members often act as availability entrepreneurs, working in concert with private persons to further parochial agendas. Interest groups are in a good position to exploit the underlying mechanisms, heightening or dampening public fears, sometimes by manipulating people's affect about certain products and processes. At the same time, the committee system keeps Congress highly balkanized. It also increases its susceptibility to short-term pressures. Because the activities of committee members are carefully scrutinized by directly affected lobbies, they are often reluctant to publicize information, make statements, or take positions that can be used against them at election time.

An unfortunate consequence is that when an availability cascade, or some other pressure, fuels massive demand for acting on a matter under the committee's purview, committee members may yield, and legislative action may follow without any attempt to coordinate it with existing policies, to place it in a comparative context, or to assess the issue of tradeoffs. *Risk redistribution* becomes at least as likely as *risk reduction*. If the corn lobby has particular influence before the key committee, it might well be able to push for a requirement of ethanol in gasoline, whatever the health consequences. A balkanized Congress is ill-equipped in an excellent position to be giving out gifts to well-organized private groups, in the form of purportedly public-spirited regulation that really helps some groups while hurting others.

In its present form, Congress is entirely illequipped to consider the problem of health–health tradeoffs. Here too committee structure ensures a high degree of fragmentation and does not allow for deliberation on such tradeoffs. On the contrary, that structure makes substitute risks difficult to evaluate or, much worse, even to see; often those risks are thought to be subject to the jurisdiction of another committee, which means, in practice, that coordination is extremely difficult. What can be done? Here are several possibilities.

- *Risk oversight.* Congress should create a risk regulation committee that would be entrusted with compiling information about a wide range of risk levels and helping to produce sensible priorities. This committee would have authority over both substantive statutes and the appropriations process. It would thus operate as a check on short-term pressures by putting particular concerns in a broader context. Its basic goal would be to engage in risk ranking, to publicize misallocations, and to initiate legislative corrections. In its ideal form, the committee would rely heavily on prevailing scientific knowledge. Its essential function would be to ensure against myopic, unduly quick, and poorly reasoned responses, not to insulate risk regulation from social values. This committee should have the power to introduce corrective legislation when a statute, or agency action under a statute, has been shown to increase aggregate risks. Like any government body, the risk regulation committee would most certainly become the target of well-organized private groups seeking to mold policies to their own advantage. But the very acts of comparing risks and publicizing the comparisons would provide a measure of protection against well-organized private interests. If the social costs of accommodating a particular lobby's demands gain widespread recognition, counterlobbying may neutralize the lobby's political effectiveness.
- *Health–health comparisons.* Congress should also address the problem of health–health tradeoffs through a new directive in the Administrative

Procedure Act (APA), the basic law governing the legality of actions by administrative agencies. The new directive could build on a House bill introduced in 1995, which contains a subsection entitled “substitution risks.” This subsection says that “each significant risk assessment or risk characterization document shall include a statement of any substitution risks to human health, where information on such risks has been provided to the agency.” This would be a strikingly modest initiative. It does not require agencies to investigate substitute risks on their own. Nor does it say that agencies may not proceed unless the regulation yields net benefits. I suggest instead an amendment to the Administrative Procedure Act: “Agencies shall ensure, to the extent feasible, that regulations do not create countervailing risks that are greater than those of the regulated risk.” A forerunner of this idea can be found in the “clean fuels” provision of the Clean Air Act, which says that the administrator of EPA may not prohibit the use of a fuel or fuel additive “unless he finds . . . that in his judgment such prohibition will not cause the use of any other fuel or fuel additive which will produce emissions which will endanger the public health or welfare to the same or greater degree than the use of the fuel or fuel additive proposed to be prohibited.”¹⁹ This idea should be generalized. (For more details, see Chapter 6.)

- *Peer review.* Recent proposals have put considerable emphasis on requiring executive agencies to corroborate the evidence that underlies regulations through peer review.²⁰ In fact, many agencies have already experimented with peer review. The arguments thus far strongly support this initiative and suggest that the experiments should be expanded through a general congressional requirement. The most important point is that peer review can serve as a check on the mistakes that come from populist wildfires. Whether an availability cascade is in progress or has completed its course, peer review provides an important safeguard against policy responses that the facts do not justify. This is the basic case for peer review in the context of availability cascades; the Love Canal and Alar scares would have looked quite different if peer review had occurred at various stages. As we will see, peer review helped provide a great deal of information about the risks associated with arsenic. Indeed, peer review can easily be a prod rather than a brake, spurring agencies to deal with serious problem.
- *Cost–benefit mandates.* Congress has debated a number of bills designed to require agencies to engage in cost–benefit analysis. We have seen that

¹⁹ 42 USC 7545(c)(2)(C).

²⁰ See, for example, the bills discussed in Cass R. Sunstein, *Free Markets and Social Justice* chapter 10 (New York: Oxford Univ. Press, 1997).

the dynamics of availability cascades, alongside the risk of tradeoff neglect, provide a new and distinctive reason for some kind of cost-benefit mandate, not as a way of obtaining an uncontroversial assessment of policy options and not because economic efficiency is the only legitimate ground for regulation, but as a commonsensical brake on measures that would do little good and possibly considerable harm.

Because of the risk that a salient event will cause a cascade, there is a good argument for attempting to understand and quantify the magnitude of the risk, and also for putting on-screen the various disadvantages of attempting to counteract it. In particular, cost-benefit analysis might well serve as a check on ill-advised availability campaigns; consider, for example, the very different findings of cost-benefit analysis for a lead phasedown (amply justified) and for eliminating asbestos (a far more mixed picture). An understanding of the relevant social mechanisms certainly does not prove that cost-benefit analysis is a good idea; everything depends on implementation. But such an understanding offers a new and largely institutional basis for a cost-benefit requirement, operating to widen the viewscreen of political actors and to contain availability errors.

THE EXECUTIVE BRANCH

The most effective response to the problems discussed here would probably involve the executive branch, which is in the best position to analyze risks comprehensively. Unlike courts, the executive branch can have a systematic overview of risk regulation; unlike Congress, the executive branch has institutions in place with which to undertake the process of rationalizing regulation. Of course whatever the executive branch does must fit within the law as enacted by Congress; in that sense the legislature has a kind of priority. But here are three complementary proposals.

OFFICE OF INFORMATION AND REGULATORY AFFAIRS

In Chaptre I, we saw that President Reagan created the Office of Information and Regulatory Affairs, assigning it the responsibility of overseeing risk regulation to ensure both coordination and rationality. Since that time, OIRA has helped to coordinate regulatory policies. But its functions have varied. Under President Reagan, OIRA operated essentially as a cost-benefit monitor that intervened in an ad hoc way to force the reconsideration of grossly inefficient regulations. President Clinton undertook a number of impressive steps to "reinvent government" to ensure greater attention to results than to processes; but this

particular role was deemphasized. In the Clinton Administration, OIRA was a weak institution, doing little of substantive value.

In view of the problems discussed here, OIRA should be reinvigorated, and its powers should be extended and strengthened, so as both to deter unreasonable regulation and to ensure that reasonable regulation is forthcoming. Specifically, OIRA should have, and be known to have, a degree of authority over priority-setting and cost-benefit balancing. It should work to mitigate the most unfortunate effects of cognitive biases, both by diverting resources and attention from small risks and by ensuring that serious but neglected risks receive attention. OIRA has used the idea of the "return letter," to require reconsideration of inadequately justified regulations, alongside the idea of the "prompt letter," to encourage agencies to respond where regulation would do more good than harm. Where experts working under OIRA lack confidence in risk judgments that are spreading and becoming embellished through a cascade, they should conduct fact finding exercises and, where necessary, publicize the inaccuracies of the popular beliefs. Thus OIRA's mission should also include the dissemination of systematic information concerning risks, including changes in what scientists know about the risks and the methods for lessening them. Finally, it should conduct systematic comparisons with other societies with an eye toward finding cross-country differences that might provide clues to misperceptions or policy flaws at home.

OIRA should also see, as one of its central assignments, the task of overcoming governmental tunnel vision, by ensuring that aggregate risks are reduced and that agency focus on particular risks does not mean that ancillary risks are ignored or increased. This is a more modest and particularized version of Justice Stephen Breyer's larger suggestion that government technocrats should have a power to set priorities by diverting resources from smaller problems to larger ones.²¹ No body in government is now entrusted with the authority of ensuring that risk regulation is managed so as to ensure global rationality and coherence. OIRA is well situated to take on that role, at least by attending to the possibility that regulation of some risks may make risk levels higher on balance. It is important to underline the complementary features of this proposal: OIRA should block regulations that cannot be justified, but also seek to energize agencies to produce regulations that would do more good than harm. The problem involves paranoia and neglect. Where there are targets of opportunity – areas in which significant risks can be reduced, without imposing unacceptable burdens and costs – OIRA should see, as one of its key missions, the encouragement of agency action. One possibility would be to designate one or more people as having a special responsibility for identifying areas where aggressive regulatory action could do some good.

²¹ Stephen G. Breyer, *Regulation and Its Reform* (Cambridge, Mass.: Harvard Univ. Press, 1992).

A NEW RISK AGENCY

If it proves infeasible to redesign OIRA along the proposed lines, and OIRA's role remains very limited, one might create a new institution that would publicize the inconsistencies of the prevailing regulatory system and focus attention on the most serious risks. Going well beyond the OIRA model, Justice Breyer has suggested the creation of a group of well-trained risk managers, versed in various disciplines and authorized to divert resources from small problems to large ones.²² Armed with federal authority, some kind of "Breyer group" would undertake the sort of analyses and educational activities proposed here. The group would have the authority to publicize its findings about the relative seriousness of risks, require agencies to engage in similar priority-setting, and recommend changes in statutes, regulations, and even appropriations. It might even be empowered to engage in some relocating on its own.

THE INTERNET

Why do people turn to nonexperts for information on various risks? One reason is that they have no easy way of finding statistically accurate and scientifically up-to-date judgments. How do the risks of driving compare with the risks of flying? How do the risks of pesticides compare with the risks of not eating a lot of fruits and vegetables? How are the risks of eating peanut butter compared with the risks of smoking cigarettes? Are organic foods healthier than genetically modified foods?

Without information on such questions, people cannot develop the capacity to compare risks and obtain sound understandings of the relative gravity of the risks associated with, for example, air travel, cell phone use, automobile driving, poor diet, and lack of exercise. To be sure, some such information may be found in the books shelved in a good library or bookstore. But people looking for it might have to read dozens of books to make the necessary comparisons; not only that, they would first have to identify the appropriate sources. Information should be made easily accessible in the quickest possible and most up-to-date form. In the modern era, the Internet would serve as the ideal vehicle.

The executive branch might well create a new website, dedicated to the listing of various risks and the identification of the probabilities, or range of probabilities, associated with each of them. The technology of the web allows the nesting of multiple levels of detail.²³ The most elementary level ought to be

²² See Breyer, *supra* note 8 in Chapter 2.

²³ The Harvard Center on Risk Analysis provides a great deal of information in this vein. See www.hcra.harvard.edu/.

extremely simple to follow — simple enough, perhaps, for high school graduates to check the latest scientific knowledge on, say, the risks associated with Alar or the damage caused by the spill at Love Canal. Where the scientific community is divided on a particular risk matter, as it is on global warming, the website should make the substance of the controversy as clear as possible. If the government fails to act, or if the government's information seems incorrect, such a website could be set up and operated by a nongovernmental organization, along the lines of *Consumer Reports*. It could even be a profit-oriented enterprise; after all, many profit-seeking credit bureaus enjoy a good reputation because they have a stake in its maintenance. Whether or not a nongovernmental entity is willing or able to undertake the task, the federal government could take the lead, though here there are obvious issues of trust.

The website's construction would constitute a substantial educational service. Potentially counteracting the irrational attitudes and decisions caused by cognitive biases and heuristics, it would allow individuals to form their risk judgments more rationally than they currently can. An even more ambitious enterprise would provide information also about the groups that have a stake in the risk judgments that drive regulatory policies. One could publicize analyses of the organized political activity concerning various risk issues. Finally, the website could contain systematic information on any discrepancies between private and public opinion on current risk issues. Such information would identify hidden currents of opinion, thus strengthening individual resistance to the biases of public discourse and pressures rooted in public opinion.

In a world in which all perceived risks are quickly evaluated by trusted scientists, who then post their knowledge in comprehensible form on a website known to everyone, availability errors will be less likely to develop. If a news program claims that apples carry a deadly poison, people can check the website to learn what is known about the identified risk. If scientific tests show the risk is nonexistent, few people will believe the claim, and no cascade will follow. If the claim is groundless, a possibly costly cascade will have been prevented without involving the courts. Of course some people might distrust what a website, private or governmental, has to say. But everything is to be gained by increasing information on risk-related questions.

COURTS

Courts, too, have a role to play in preventing excessive reactions to availability cascades and in ensuring attention to neglected tradeoffs. I will deal with this issue in some detail in Chapter 8; for now consider some basic points.

The most natural route involves the judicial review of administrative actions alleged to be "arbitrary" or "capricious" within the meaning of the

Administrative Procedure Act.²⁴ The ban on arbitrary or capricious action is increasingly being understood to require agencies to show that their action produces "more good than harm."²⁵ This notion might well be taken to embody a presumptive requirement that costs not be grossly disproportionate to benefits. Simple though it sounds, existing doctrine authorizes courts to invalidate the most extreme and the most poorly conceived regulatory proposals, at least when statutes do not require them. Courts should hardly be expected to identify cognitive errors and to invalidate an outcome for having been produced by them. But if regulations must be shown to make things better rather than worse, taking account of all relevant variables, a safeguard will be in place.

AN IMPORTANT PRAGMATIC QUALIFICATION

232
The arguments I have made for cost-benefit balancing are highly pragmatic. I have urged that if this form of balancing is done well, regulation is likely to be far more sensible than it now is. But this is not inevitable. Suppose, for example, that the principal effect of cost-benefit analysis is to stall desirable (as well as undesirable) regulation, ensuring that regulations would be more difficult to implement. Suppose too that if regulations would be more difficult to implement, the public would be worse off, simply because serious dangers would not be addressed. Under these assumptions, cost-benefit analysis would not be justified. It would fail cost-benefit analysis. And there can be no doubt that many of those who reject cost-benefit analysis do so on the ground that this "tool" would – in their judgment – simply provide yet another obstacle to desirable protection of safety, health, and the environment. Perhaps the skeptics would favor cost-benefit analysis if it could be done quickly and accurately, without distorting influences. But the skeptics might fear that the likelihood of rapid, accurate cost-benefit analysis is too low to justify the real risk, which is to stall and even block lifesaving interventions.

I have not established that these skeptics are wrong. And if they are right, we should abandon cost-benefit analysis. Whether they are right depends on empirical issues on which we lack clear evidence. In response, I have pointed to cases in which cost-benefit analysis has done considerable good, not least by spurring regulation. But the real burden lies on public officials who seek to use cost-benefit analysis as a tool for improving decisions. If they use the tool properly, and if both lives and dollars are saved as a

²⁴ 5 USC 706.

²⁵ See Edward Warren & Gary Merchant, "More Good Than Harm": A First Principle for Environmental Agencies and Reviewing Courts, 20 *Ecol. L. Q.* 379 (1993); Margolis, *supra* note 3 in Chapter 3.

result, the skeptics' concerns will be met on their own, sensibly pragmatic grounds.

FOUR OBJECTIONS

The argument made thus far, cautious though it may seem, runs into four immediate objections. The first involves democratic considerations. The second points to the limitations of quantification. The third involves the question of who bears the relevant costs and who gets the relevant benefits. The fourth involves the possibility that some tradeoffs, including some trades of dollars for lives, should be treated as "taboo." I take up these objections in sequence.

POPULISM

The first objection, populist in character, would be that in a democracy, government properly responds to the social "demand" for law. Government does not legitimately reject that demand on the ground that cost-benefit analysis suggests that it should not act. On this view, a democratic government should be accountable. Any approach that uses efficiency, or technocratically driven judgments, as a brake on accountability is fatally undemocratic.

The problem with this objection is that it rests on a controversial and even unacceptable conception of democracy, one that sees responsiveness to citizens' demands, whatever their factual basis, as the foundation of political legitimacy. If those demands are uninformed, or based on unreliable intuitions about risks, it is perfectly appropriate for government to resist them. Indeed, it is far from clear that reasonable citizens want, or would want, their government to respond to their uninformed demands. After the analysis and information have been generated, and public officials have taken them into account, democratic safeguards continue to be available, and electoral sanctions can be brought to bear against those who have violated the public will. At the very least, cost-benefit analysis should be an ingredient in the analysis, showing people that the consequences of various approaches might be different from what they seem.

QUANTIFICATION AND EXPRESSIVE RATIONALITY

In an extensive discussion, Lisa Heinzerling has raised a number of objections to cost-benefit balancing.²⁶ Heinzerling argues that many of the values depend on controversial judgments of value. Making reference to Table 2.1 in particular, she urges that cost-benefit analysis masks those judgments. Heinzerling contends

²⁶ See Heinzerling, *supra* note 5 in Chapter 2, at 198I.

Table 5.I Corrected (?) table on cost-effectiveness of regulations

Regulation	Adjusted Cost Estimate (thousands of 1995 dollars)
Asbestos (OSHA 1972)	700
Benzene (OSHA 1985)	2,570
Arsenic/glass plant (EPA 1986)	6,610
Ethylene oxide (OSHA 1984)	3,020-5,780
Uranium mill tailings/inactive (EPA 1983)	2,410
Acrylonitrile (OSHA 1978)	8,570
Uranium mill tailings/active (EPA 1983)	3,840
Coke ovens (OSHA 1976)	12,420
Asbestos (OSHA 1986)	3,860
Arsenic (OSHA 1978)	24,490
Arsenic/los-arsenic copper (EPA 1986)	5,740
Land disposal (EPA 1986)	3,280
Formaldehyde (OSHA 1985)	31,100

that many of the key numbers depend on controversial judgments about how to "discount" future benefits, that is, how to treat benefits that will occur in the future. This is indeed a highly disputed issue. The numbers in Table 2.I depend on a 10% discount rate, whereas the agencies tended to use a lower discount rate, or not to discount at all. Heinzerling also suggests that Table 2.I depends on downward adjustment of the agency's estimates of risk. Her own estimates result in Table 5.I, which is a corrected risk table, adjusted for inflation.

Table 5.I may be more accurate than Table 2.I; certainly there are problems with any approach that assumes a 10 percent discount rate, which seems far too high (see Chapter 8 for details). But even if Heinzerling's table is better, it offers an ironic lesson, serving largely to confirm the point that current regulatory policy suffers from poor priority-setting. The disparities here are not as dramatic as in Table 2.I, and they certainly do not establish pervasive overregulation; but they do support the view that resources are being misallocated and that there is a serious problem to be solved.

Heinzerling does not, however, conclude that this revised table is the appropriate basis for evaluating regulatory policy. Her aim is not to come up with a better table from which to reassess government behavior. On the contrary, she takes her argument to be a basis for rejecting cost-benefit analysis altogether. This, then, is a lesson about "the perils of precision."²⁷ Heinzerling suggests that it "would be better if we left the picture blurry, and declined to connect the dots between all the confusing and sometimes conflicting intuitions

and evidence."²⁸ She is concerned that "some, probably many, people will be fooled into believing that numerical estimates of risks, costs, and benefits are impartial reflections of factual reality, in which case the likely result of increased reliance on quantification in setting regulatory policy will be that the side that best obscures the value choices implicit in its numbers will prevail."²⁹

There is considerable truth here and an important pragmatic warning; but I think that Heinzerling's lesson is greatly overdrawn. Truth first: If an agency says that the cost of regulation is \$100 million and the benefit \$70 million, we still know much less than we should. It is important to know who bears these costs, and if possible with what consequences. Will wages be lower? Whose wages? Will prices be higher? Of what products? A disaggregated picture of the benefits would also be important; what does the seventy million dollar figure represent? Consider, for example, a recent table explaining that the costs of skin cancer, from health effects of reducing tropospheric ozone, are between \$290 million and \$1.1 billion, with dollar subtotals for skin cancers and cataracts.³⁰ By itself, this table is insufficiently informative to tell people what they need to know.

Heinzerling is therefore on firm ground if she means to suggest that the dollar numbers cannot substitute for a fuller inquiry into what is at stake. Any cost-benefit analysis should include more than the monetary values by, for example, showing what the values are about, such as life-years saved and accidents averted. But her own table suggests that the general conclusion — that cost-benefit analysis can illuminate inquiry — remains unassailable. If regulation ranges from tens of thousands to tens of millions per life saved, at least there is an issue to be addressed. One of the functions of cost-benefit balancing is to help show where limited resources should go. In fact, a regulation of particulates is hard to evaluate without knowing, for example, the number of deaths averted and the range of consequences for morbidity: How many work-days will be saved that would otherwise be lost? How many hospitalizations will be avoided? How many asthma attacks will be prevented? It could even be useful to attempt to describe these effects in terms of "quality-adjusted life-years" (see Chapter 9), knowing that here, too, a good analyst will go back and forth between bottom lines and the judgments that go into their creation.

I suspect that theoretical claims lie behind Heinzerling's skepticism about quantification. She may believe that many of the goods at stake in regulation (e.g., human and animal life and health, recreational and aesthetic opportunities)

²⁸ Id. at 2069.

²⁹ Id. at 2068.

³⁰ See Lutter & Wolz, *supra* note 20 in Chapter 2, at 142A, 145. In fairness to the authors, it should be noted that a previous table in their essay describes adverse health effects in quantitative terms by listing the numbers of cases averted.

²⁷ Id. at 2042.

are not merely commodities, that people do not value these goods in the same way that they value cash, and that cost-benefit analysis, by its reductionism, is inconsistent with people's reflective judgments about the issues at stake. Arguments of this sort have been developed in some philosophical challenges to cost-benefit analysis.³¹ Such arguments are convincing if cost-benefit analysis is taken to suggest a controversial position in favor of the commensurability of all goods. Part of what people express, in their daily lives, is a resistance to this form of commensurability, and some goods are believed to have intrinsic as well as instrumental value. The existence of qualitative differences among goods fortifies the claim that any bottom line about costs and benefits should be supplemented with a more qualitative description of the variables involved. But cost-benefit analysis should not be seen as embodying a reductionist account of the good, and much less as a suggestion that everything is simply a "commodity" for human use. It is best taken as pragmatic instrument, agnostic on the deep issues and designed to assist people in making complex judgments where multiple goods are involved.

We should conclude that the final number may provide less information than the ingredients that went into it, and that officials should have and present cost-benefit analysis in sufficiently full terms to enable people to have a concrete sense of the effects of regulation. This is an argument against some overambitious understandings of what cost-benefit balancing entails. But it is not an argument against cost-benefit balancing.

WHOSE COSTS? WHOSE BENEFITS?

The third objection, hinted at earlier, is perhaps the most obvious. Suppose that the costs of additional reductions in sulfur dioxide emissions would be \$500 million. Suppose that the health benefits of these reductions, once monetized, would be \$250 million. Suppose too that we do not have any trouble with the underlying calculations – that we have accurately identified the health benefits and that we have properly "translated" them into monetary equivalents. If the costs are double the benefits, is it clear that government should not go forward? A critic of cost-benefit analysis might think that this is not clear at all. Everything depends on *who* would bear the relevant costs and *who* would enjoy the relevant benefits. Suppose that if the regulation is imposed, companies will lose hundreds of millions of dollars in profits, but a nontrivial number of people will live who would otherwise die. If the costs mean lower profits, and the benefits mean longer lives, it might seem clear that the government should act, even if the costs are higher than the benefits.

³¹ See Anderson, *supra* note 4 in Chapter 3.

The basic point is right; we do need to know who would bear the costs and enjoy the benefits. That has been and will continue to be one of my principal themes. Recall that I have not urged that the monetized numbers should be decisive. But for those who seek to reject cost-benefit analysis, things are a bit more complicated. The reason is not that redistribution through regulation is undesirable, or that we should care a lot about the profits of "companies." The reason is that when expensive regulation is imposed, "companies" are not going to be the only ones who pay. Sometimes workers will lose their jobs – and increases in unemployment are associated with a wide range of social ills, including crime and suicide. Sometimes workers will have lower wages – and wage decreases are especially hard on people who are living at the margin. Sometimes prices will go up, and price increases are especially hard on poor people. If regulation is very expensive, some people might even die. When companies are asked to spend hundreds of millions of dollars on risk reduction, ordinary people will inevitably be affected, usually for the worse. The "cost" side of the cost-benefit calculus means much more than a decrease in profits.

Of course, we cannot know, in the abstract, whether the cost of regulation will be borne by consumers, current workers, the unemployed, or companies themselves. Much depends on what the market will allow companies to do. If prices can be increased without significantly reducing sales, companies will increase prices; if wages can be cut without losing good workers, wages will be cut; if companies can do what they have done before with a smaller workforce, employment will be reduced. Inevitably companies will select what is, for them, the cheapest way to respond to the costs of regulation.

This is no mere conjecture. Consider, for example, the fact that workers' compensation programs are an excellent way to increase safety for workers because they increase employers' incentive to maintain a safe workplace – indeed, such programs probably cut workers' fatalities by 30 percent or more.³² In this way, workers' compensation programs should be counted as a huge success. But there is also evidence that such programs contributed to a substantial decrease in workers' wages – in some sectors, a dollar-for-dollar decrease, meaning that workers lost in salary what they gained in benefits.³³ To take a more recent example outside of the domain of health and safety, there is evidence that the parental leave program has produced a similar wage cut for affected workers.³⁴ None of this establishes that a workers' compensation program or a parental

³² See W. Kip Viscusi, *Reforming Products Liability*, p. 178, (Cambridge, Mass.: Harvard Univ. Press, 1995).

³³ See Fishback & Kantor, *supra* note 27 in Chapter 3.

³⁴ See Jonathan Gruber, *The Incidence of Mandated Maternity Benefits*, 84 *Am. Econ. Rev.* 622 (1994).

Propose a change for jury

leave program or a costly safety regulation is a bad idea. Even if the result is to produce wage reductions or increased unemployment, the effort to reduce risks might be justified on balance. My claim is only that the dollar costs associated with risk reduction are extremely important, and that they are important not because we should care about some abstraction called "dollars" or because company profits are important as such, but because costly regulations are translated, all too often, into harmful consequences for real people. The point is that those harmful consequences should be placed on-screen so that we do produce harmful unintended consequences and, on occasion, real calamities.

TABOO TRADEOFFS AND TRADEOFF AVERSION

For many regulations, there appears to be a relatively simple choice: dollars or risks? When the question is put in this form, many people find it tempting to reject cost-benefit analysis. In fact, rejecting that form of analysis seems, to many, to be the appropriate way to put priorities exactly where they belong — on the protection of life and health. In recent years some intriguing evidence suggests that this is how most people think. Some tradeoffs, between risks and dollars, actually seem to be "taboo," in the sense that some people, some of the time, will absolutely refuse to allow some goods to be traded for money.³⁵ How much would you have to be paid to allow your daughter, or mother, to suffer for a month, or to lose five years from her life? When a tradeoff is taboo, any proposed deal is not even considered. It would violate a firmly held moral principle.

In any case, people frequently display not merely tradeoff neglect but also *tradeoff aversion*. When a government commission suggests that a certain amount is "too high" to spend to protect human lives, almost three-quarters of subjects, in one experiment, rejected the suggestion — not because the commission was wrong on the numbers, but simply because it was weighing lives against dollars.³⁶ Jonathan Baron and his collaborators have done a great deal of work on "protected values" — values that people seek to insulate from ordinary tradeoffs.³⁷ When values are protected, people tend to believe that they are absolute; they deny the need for tradeoffs and show considerable anger at violations of those values. Would you agree with a company's decision to refuse to spend money

on improving the safety of the workplace on the ground that a cost-benefit analysis suggests that the improvement simply isn't worthwhile? Studies of jury behavior have shown that most people will not. In fact, jurors will punish, often severely, companies that have engaged in cost-benefit analysis before marketing products that carry a risk of harm.³⁸

Suppose, for example, that an automobile company is aware that, in extremely unusual circumstances, its brake system will fail, thus causing injury and death, and that the company could correct the problem with a costly expenditure, one that would, say, cost \$15 million per life saved. Suppose that the company refuses to make the expenditure, on the ground that the benefits do not justify it. A jury will likely be aghast at the company's behavior. Indeed, the jury might well award punitive damages. Note that this is not because the company has placed too low a value on human life. Even if the company's willingness to pay is quite high, juries will be unmoved. The problem is the fact that the company has made an explicit tradeoff between risk and cost. This finding is duplicated by others showing that people are extremely unhappy with government decisions to make such explicit tradeoffs, even though they would be willing to approve of the relevant decisions if those decisions are made on less transparent grounds.

I think that this kind of tradeoff aversion is a pernicious form of tradeoff neglect, with harmful effects on risk reduction efforts. But we have to be careful here. It would be far too simple to say that those who oppose explicit tradeoffs of risks and dollars are making some inexplicable "mistake." In ordinary life, anyone who talks explicitly in cost-benefit terms is likely to seem cold and calculating, perhaps even worse. A parent who says, "I will not buy a Volvo because the additional safety for my child just isn't worth \$600," would seem a bit strange. It would be much stranger for the parent to claim, "If I am paid enough, I will be willing to subject my child to a small danger." In surveys, a significant percentage of people actually say that they would not accept any amount of money to subject themselves to an small increase in risk, or to allow the environment to be harmed.

But why is this? People, including parents, trade risks for dollars all the time. We choose how much to spend on cars, knowing that safety is expensive; we decide how much to spend on security systems in the home; we choose where to live, knowing that some areas are safer than others; we go out at night, even though we know that by doing so, we increase our risks; when the cost of risk reduction is too high, we will not incur that cost even to protect our own children. What seems forbidden is not behavior that embodies tradeoffs, but rather unduly explicit talk to the effect. The taboo on such talk may well serve salutary social functions by helping to establish and maintain certain attitudes, in which life

³⁵ See Philip Tetlock, *Coping With Tradeoffs*, in *Elements of Reason: Cognition, Choice, and the Bounds of Rationality* 239, Arthur Lupia et al., eds. (Cambridge: Cambridge Univ. Press, 2000).

³⁶ *Id.* at 255.

³⁷ Jonathan Baron & Mark Spranca, *Protected Values*, 70 *Org. Behavior Human Decision Processes* 1 (1997).

³⁸ See W. Kip Viscusi, *Corporate Risk Analysis: A Reckless Act?*, 52 *Stan. L. Rev.* 547 (2000).

and health are not seen as simple commodities, qualitatively indistinguishable from money and other things that are simply for use. But we should not be fooled by the fact that people are nervous about explicit talk of reducing safety for money. Tradeoffs of money and risks are exceedingly common.

A more careful look at taboo tradeoffs suggests that when people disapprove of trading money for risks, they are generalizing from a set of moral principles that are generally sound, and even useful, but that work poorly in some cases.³⁹ Consider the following moral principle: *Do not knowingly cause a human death.* In the key studies, people disapprove of companies that fail to improve safety when they are fully aware that deaths will result — whereas people do not disapprove of those who fail to improve safety while appearing not to know, for certain, that deaths are going to ensue. When people object to risky action taken after cost-benefit analysis, it seems to be because that very analysis puts the number of expected deaths squarely on-screen. Companies that fail to do such analysis, but that are aware that a “risk” exists, do not make clear, to themselves or to jurors, that they caused deaths with full knowledge that this was what they were going to do. People disapprove, above all, of companies that cause death knowingly.

The problem here is that it is not always unacceptable to cause death knowingly, at least if the deaths are relatively few and an unintended byproduct of generally desirable activity. If government allows the marketing of SUVs, or small, fuel-efficient cars, it knows that deaths will occur; indeed, a government that requires small, fuel-efficient cars will produce some deaths as a result. If government allows new highways to be built, it will know that people will die on those highways; if government allows new power plants to be built, it will know that some people will die from the resulting pollution; if companies produce tobacco products, and if government does not ban those products, hundreds of thousands of people will die; alcohol kills people too. Much of what is done, by both industry and government, is likely to result in one or more deaths. Of course it would make sense, in most or all of these domains, to take extra steps to reduce risks. But that proposition does not support the implausible claim that we should disapprove, from the moral point of view, of any action taken when deaths are foreseeable.

The question is whether the statistical risk is worthwhile. Of course, the raw numbers cannot resolve that question. But without the raw numbers, any attempt to provide an answer is a stab in the dark. I believe that taboo tradeoffs and tradeoff aversion, far from being an objection to cost-benefit analysis, are a large part of the reason that it is so important.

³⁹ See Jonathan Baron, *Nonconsequentialist Decisions*, 17 *Behavioral and Brain Sciences* 1 (1994).

CATASTROPHES, TERRORISM, AND UNCERTAINTY

I have been assuming throughout that the numbers can actually be assigned — that government is able to specify a range of possible outcomes, and that this range can help discipline the decision about that to do. Often this is true. But sometimes the range is extremely wide, and government is not able to assign probabilities or to come up with ranges with any confidence. In the aftermath of the terrorist attacks of September 11, 2001, what was the probability of a further attack within a week? A month? Six months? A year? Suppose that government is attempting to control the risks from biochemical attacks, or from the worst-case scenario involving global warming. Is cost-benefit analysis helpful? If not, what kind of analysis might be helpful instead?

Regulators are sometimes acting in a situation of “risk,” where probabilities can be assigned to various outcomes, but sometimes they are acting in a situation of “uncertainty,” where no such probabilities can be assigned. In a situation of uncertainty, it is exceedingly hard to do cost-benefit analysis.⁴⁰ And indeed it would be possible to think that the argument I have made applies only in cases when it is feasible to come up with numbers. But careful analysis will be helpful even when full cost-benefit balancing is impossible. If a catastrophe has a very low probability of occurring, it does not deserve the level of attention that it should command if its probability is higher. And even if cost-benefit analysis is not feasible, it is important to pursue cost-effectiveness, by ensuring that the steps being taken are both the cheapest and the most effective means of producing the goal at hand. This is no less true in the context of risks from terrorism than elsewhere. For global warming, a “cap-and-trade” system (see Chapter 10) seems to be the most promising, in part because it is so much less expensive than the alternatives. Indeed, the administration of George W. Bush is attempting to reduce the risk of global warming not through mandates but through information, moral suasion, and tax incentives. The Bush approach is, in my view, inadequate; but at least it is possible that these steps will produce some good relatively cheaply.

There is a related problem. Prospective estimates of both costs and benefits often turn out to be wrong. This is not merely because of interest-group pressures. One reason is that officials lack the extensive information that would permit them to make accurate predictions; indeed, the informational demand on agencies is overwhelming, especially because technologies change over time.

⁴⁰ It is standard, in such circumstances, to follow the maximin principle (choose the policy with the best worst-case outcome). See Jon Elster, *Explaining Technical Change*, 185–207 (Cambridge: Cambridge Univ. Press, 1983), for a helpful discussion.

the basic point

An enduring problem for regulatory policy is the absence of precise information on the cost or benefit sides. Industry, which typically has the most information about costs, has a strong incentive to overstate its burden. This point should be taken not as a criticism of cost-benefit analysis as such, but as a reason for skepticism about industry's numbers and for continuous monitoring and updating. Hence any system for cost-benefit balancing should create a mechanism to ensure that the prospective analysis is not way off – and should allow for correction if, as often happens, things turn out to be very different from what was anticipated.

Of course it is possible that, in practice, cost-benefit analysis will be used to block sensible regulation, and will therefore do more harm than good. If cost-benefit analysis simply makes it harder for agencies to protect the public, and mostly increases the power of regulated groups to block desirable regulation, it is senseless to celebrate cost-benefit analysis. But there are reasons to believe that cost-benefit analysis is not simply an antiregulatory tool. I have mentioned that cost-benefit analysis helped to spur the removal of lead from gasoline and dramatic steps, pushed by the United States, to eliminate CFCs, which contribute to depletion of the ozone layer. More recently, the Office of Information and Regulatory Affairs has pioneered the idea of a "prompt letter" – letters designed to prompt agencies to act in cases in which the benefits of action seem to outweigh the costs. Inspired by tentative cost-benefit analysis, OIRA has asked OSHA to consider requiring automatic defibrillators to be placed in workplaces; has urged the FDA to issue a final rule requiring disclosure of the level of trans fatty acids in foods; and has asked the Department of Transportation to take steps to improve automobile safety by establishing a high-speed, frontal offset crash test. In any case, cost-benefit analysis has often driven policy in more sensible directions by showing the best means of achieving regulatory goals. For those who have enthusiasm for cost-benefit analysis, there is a continuing duty to show that this is a tool for better outcomes and more sensible priority-setting, rather than a recipe for inaction.

BEYOND COST-BENEFIT ANALYSIS: A HAYEKIAN TURN?

I have suggested that the choice of "smarter tools" remains one of the largest imperatives for risk reduction. Cost-benefit analysis can help to identify those tools, and for that reason governmental acceptance of cost-benefit balancing might well spur the selection of more effective, cheaper means of accomplishing regulatory goals. Within EPA, this has happened on many occasions (see Chapter 10). But there are some deeper issues here, and they suggest the need for some fundamental redirection of the cost-benefit state. The basic problem

is that an adequate cost-benefit analysis imposes large demands on government, which may well lack the necessary information. What might be done instead?

To answer that question, imagine a small community consisting of farmers, each of whom owns a dozen cows, and all of whom own, jointly, the land on which the cows graze. It is easy to see that the farmers will soon find themselves facing a problem of overgrazing. Because the land is owned in common, each farmer will receive all the benefits of grazing by his cows but will pay only a fraction of the cost of that grazing. If farmers are acting in their rational self-interest, all of them will think that way, and before long the grass will be depleted. This is the famous "tragedy of the commons," and it accounts for many environmental problems.⁴¹

How might the tragedy be averted? Perhaps the small community could develop certain social norms, which would ensure limitations on grazing activity. Good norms can prevent tragedy and greatly reduce risks. Perhaps the small community could charge people for grazing, so as to ensure that the available stock is not depleted. In the spirit of this chapter, the community might even conduct a cost-benefit analysis of different levels of grazing and mandate levels of grazing, from each farmer, that would be consistent with the outcome of the analysis. It is easily imaginable that the latter approach would be better than unrestricted grazing. But it would also be full of pitfalls. Can the community really make accurate calculations? Won't things change over time? Doesn't the analysis, and then the set of mandates, seem a bit too reminiscent of Soviet-style planning?

Friedrich Hayek, winner of the Nobel Prize and critic of "planning" in all its guises, emphasized government's pervasive lack of information, certainly as compared with the information held by the numerous people who participate in a marketplace.⁴² In the context of our farming community, it would be possible to reject planning, even of the sort suggested by cost-benefit analysis, and to urge instead a kind of Hayekian turn: Create private property rights, and let the farmers operate as they wish, within the constraints of those rights. More specifically, each farmer might be given a right to a certain plot of land and be allowed to use the land of others only with permission or as a result of a voluntary agreement. No trespassing would be allowed. The beauty of this solution is that it solves the tragedy of the commons, by ensuring that there is no commons at all. If each farmer is responsible for his own land and his own cattle, he has the right incentive, to use the land in a way that does not deplete it.

⁴¹ Garrett Hardin, *The Tragedy of the Commons*, 162 *Science* 1243 (1968).

⁴² F. A. Hayek, *The Use of Knowledge in Society*, 35 *Am. Econ. Rev.* 519 (1945).

The movement toward "free market environmentalism" is designed to generalize from this tale.⁴³ The basic claim is that the risks of environmental degradation might best be handled through creating property rights, so as to ensure that people have responsibility for environmental amenities that might otherwise be subject to the tragedy of the commons. If, for example, environmental groups are allowed to purchase a right not to cut timber or to pay ranchers for destruction of property at the hands of wolves, it might be possible to provide far more effective protection than could ever come from interest group struggles over government mandates. I believe that free market environmentalism, although not a panacea, holds out a great deal of promise, and that along with other emerging tools, it suggests the possibility of a general movement in Hayekian directions for risk reduction policies. I will turn to these issues in detail in Chapter 10.

⁴³ Terry Anderson & Donald R. Leal, *Free Market Environmentalism* (New York: Palgrave, 2001).

Health-Health Tradeoffs

I have referred to a pervasive problem in risk regulation, one that has only started to receive public attention. The problem arises when the diminution of one health risk simultaneously increases another health risk. The phenomenon is common in daily life, where risk reduction strategies may produce risks of their own. If you engage in an aggressive program of exercise to lose weight, you should be careful to make sure that the exercise is not itself unduly hazardous. If you change your diet so as to avoid pesticides, you might find yourself eating foods that carry higher risks on balance. The problem can be found in many domains of policy and law.

Thus, for example, fuel economy standards, designed to reduce environmental risks, may make automobiles less safe, and in that way increase risks to life and health.¹ If government bans the manufacture and use of asbestos, it may lead companies to use more dangerous substitutes,² and efforts to remove asbestos from schools may cause serious risks to workers. Regulations designed to control the spread of AIDS and hepatitis among health-care providers may increase the costs of health care, and thus make health care less widely available, and thus cost lives.³ Regulation of nuclear power may make nuclear power safer; but by increasing the cost of nuclear power, such regulation will ensure reliance on other energy sources, such as coal-fired power plants, which carry risks of

¹ See Robert W. Crandall, *Policy Watch: Corporate Average Fuel Economy Standards*, 6 J. Econ. Persp. 171, 178 (1992); Robert W. Crandall & John D. Graham, *The Effect of Fuel Economy Standards on Automobile Safety*, 32 J. L. & Econ. 97 (1989).

² See *Corrosion Proof Fittings v. EPA*, 947 F.2d 1201 (5th Cir. 1991).

³ See *ADA v. Martin*, 984 F.2d 823, 826 (7th Cir. 1993): "OSHA also exaggerated the number of lives likely to be saved by the rule by ignoring lives likely to be lost by it. since the increased cost of medical care, to the extent passed on to consumers, will reduce the demand for medical care, and some people may lose their lives as a result."