

Risk in a Free Society¹

William D. Ruckelshaus²

It is now a commonplace of political discourse that technological advances have had a profound effect on our democratic institutions. Mass communication is the familiar example. But I would like to draw your attention to another way in which technology may impinge upon a democratic society, one that is perhaps as serious, if more subtle; one that commands a huge proportion of my own attention. I refer to the chemical products and by-products of modern technology and the potential social disruption associated with the processes we have created to control them.

When I began my current, and second, tenure as Administrator of EPA, my first goal was the restoration of public confidence in the Agency, and it was impressed upon me that straightening out the way we handled health risk was central to achieving it. Needless to say, EPA's primary mission is the reduction of risk, whether to public health or the environment. Some in America were afraid. They were afraid that toxic chemicals in the environment were affecting their health, and more important, they suspected that the facts about the risks from such chemicals were not being accurately reported to them, that policy considerations were being inappropriately used in such reports, so as to make the risks seem less than they were and excuse the Agency from taking action. Even worse, some people thought that the process we had established to protect public health was being abused for crass political gain.

Whether this was true or not is almost beside the point; a substantial number of people believed it. Now in a society such as ours, where the people ultimately decide policy—what they want done about a particular situation—the fair exposition of policy choices is the job of public agencies. The public agency is the repository of the facts; you can't operate a democratic society, particularly a complex tech-

nological one, unless you have such a repository. Above all, the factual guardian must be trusted, a failure of trust courts chaos. Chaos, in turn, creates its own thirst for order, which craving in its more extreme forms threatens the very foundation of democratic freedom. So in a democracy a public agency that is not trusted, especially where the protection of public health is concerned, might as well close its doors.

I described a possible solution to this problem last June in a speech to the National Academy of Sciences. The academy had stated in a recent report that Federal agencies had often confused the assessment of risk with the management of risk. Risk assessment is the use of a base of scientific research to define the probability of some harm coming to an individual or a population as a result of exposure to a substance or situation. Risk management, in contrast, is the public process of deciding what to do where risk has been determined to exist. It includes integrating risk assessment with considerations of engineering feasibility and figuring out how to exercise our imperative to reduce risk in the light of social, economic, and political factors.

The report proposed that these two functions be formally separated within regulatory agencies. I said that this appeared to be a workable idea and that we would try to make it happen at EPA. This notion was attractive because the statutes administered by many Federal regulatory agencies typically force some action when scientific inquiry establishes the presence of a risk, as, for example, when a substance present in the environment, or the workplace, or the food chain, is found to cause cancer in animals. The statutes may require the agency to act according to some protective formula: to establish "margins of safety," or "prevent significant risk," or "eliminate the risk."

When the action so forced has dire economic or social consequences, the person who must make the decision may be sorely tempted to ask for a "reinterpretation" of the data. We should remember that risk

¹An address given at Princeton University on February 18, 1984.

²Administrator, U.S. Environmental Protection Agency, Washington, D.C. 20460.

assessment data can be like the captured spy: if you torture it long enough, it will tell you anything you want to know. So it is good public policy to so structure an agency that such temptation is avoided.

But we have found that separating the assessment of risk from its management is rather more difficult to accomplish in practice. In the first place, values, which are supposed to be safely sequestered in risk management, also appear as important influences on the outcomes of risk assessments. For example, let us suppose that a chemical in common use is tested on laboratory animals with the object of determining whether it can cause cancer. At the end of the test a proportion of the animals that have been exposed to the substance show evidence of tumor formation.

Now the problems begin. First, in tests like these, the doses given are extremely high, often close to the level the animal can tolerate for a lifetime without dying from toxic non-cancer effects. Environmental exposures are typically much lower, so in order to determine what the risk of cancer is at such lower exposures—that is, to determine the curve that relates a certain dose to a certain response—we must extrapolate down from the high-dose laboratory data. There are a number of statistical models for doing this, all of which fit the data, and all of which are open to debate. We simply do not *know* what the shape of the dose-response curve is at low doses, in the sense that we know, let's say, what the orbit of a satellite will be when we shoot it off.

Next, we must deal with the uncertainty of extrapolating cancer data from animals to man, for example, determining which of the many different kinds of lesions that may appear in animals is actually indicative of a probability that the substance in question may be a human carcinogen. Cancer is cancer to the public, but not to the pathologist.

Finally, we must deal with uncertainty about exposure. We have to determine, usually on the basis of very scant data, and very elaborate mathematical models, how much of the stuff is being produced, how it is being dispersed, changed or destroyed by natural processes, and how the actual dose that people get is changed by behavioral or population characteristics.

These uncertainties inherent in risk assessment combine to produce an enormously wide range of risk estimates in most cases. For example, the National Academy of Sciences report on saccharin concluded that over the next 70 years the expected number of cases of human bladder cancer resulting from daily exposure to 120 mg of saccharin might

range from 0.22 to 1,144,000. This sort of range is of limited use to the policy maker, and risk assessment scientists are at some pains to make choices among possibilities so as to produce conclusions that are both scientifically supportable and usable.

Such choices are influenced by values, which may be affected by professional training, or by ideas about what constitutes "good science," and, of course, by the same complex of experience and individual traits that gives rise to personal values in all of us. An oncologist, for example, who values highly the ability to distinguish between different sorts of lesions, may discount certain test results as being irrelevant to decisions about human carcinogenicity. A public health epidemiologist may look at the same data and come to quite different conclusions.

Historically at EPA it has been thought prudent to make what have been called conservative assumptions; that is, *our* values lead us, in a situation of unavoidable uncertainty, to couch our conclusions in terms of a plausible upper bound. This means that when we generate a number that expresses the potency of some substance in causing disease, we can state that it is unlikely that the risk projected is any greater.

This is fine when the risks projected are vanishingly small; it's always nice to hear that some chemical is *not* a national crisis. But when the risks estimated through such assessments are substantial, so that some action may be in the offing, the stacking of conservative assumptions one on top of another, becomes a problem for the policymaker. If I am going to propose controls that may have serious economic and social effects, I need to have some idea how much confidence should be placed in the estimates of risk that prompted those controls. I need to know how likely *real* damage is to occur in the uncontrolled, partially controlled, and fully controlled cases. Only then can I apply the balancing judgments that are the essence of my job. This, of course, tends to insert the policymaker back into the guts of risk assessment, which we've agreed is less than wise.

This is a real quandary. I now believe that the main road out of it lies through a marked improvement in the way we communicate the realities of risk analysis to the public. The goal is public understanding. We will only retain the administrative flexibility we need to effectively protect the public health and welfare if the public believes we are trying to act in the public interest. There is an argument, in contradiction, that the best way to protection lies in

increased legislative specificity, in closely directing the Agency as to what to control and how much to control it. If we fail to command public confidence, this argument will prevail, and in my opinion it would be a bad thing if it did. You can't squeeze the complexity inherent in managing environmental risks between the pages of a statute book.

How then do we encourage confidence? Generally speaking there are two ways to do it. First, we could assign guardianship of the Agency's integrity—its risk assessment task—to a panel of disinterested experts who are above reproach in the public eye. This is the quasi-judicial, blue-ribbon approach, which has a strong tradition in our society. If we have a complex issue, we don't have to think about it very much, just give it to the experts, who deliberate and provide the answer, which most will accept because of the inherent prestige of the panel.

The discomfort associated with imagining, in 1984, a conclave of Big Brothers to watch over us only strengthens my conviction that such panels cannot serve the general purpose of restoring and maintaining confidence. It turns out that the experts don't agree, so instead of an unimpeachable and disinterested consensus you get dissenting advocacy. Once again, experts have values too.

Alternatively, we could all become a lot smarter about risk. The Agency could put much more effort into explaining what it is doing and what it does, and does not know. Here I do not mean "public involvement" in the usual and formal sense. This is embodied in administrative law and has always been part of our ordinary procedure in promulgating rules. Nor do I mean a mere public relations campaign to popularize Agency decisions. Public relations smoothes over—I think we need to dig up. We have to expose the assumptions that go into risk assessments. We have to admit our uncertainties and confront the public with the complex nature of decisions about risk.

Living in a technological society is like riding a bucking bronco. I don't believe we can afford to get off, and I doubt that someone will magically appear who can lead it about on a leash. The question is: how do we become better bronco busters? I think a great part of the answer is to bring about a major improvement in the quality of public debate on environmental risk.

This will not be easy. Risk assessment is a probabilistic calculation, but people don't respond to risks "as they should" if such calculations were the sole criterion of rationality. Most people are not comfort-

table with mathematical probability as a guide to living, and the risk assessment lingo we throw at them does not increase their comfort. Tell somebody that their risk of cancer from a 70-year exposure to a carcinogen at ambient levels ranges between 10^{-5} and 10^{-7} , and they are likely to come back at you with, "Yeah, but will I get cancer if I drink the water?" Also, attitudes toward risk are subjective and highly colored by personal experience and other factors not fully captured by risk assessments.

We have some research on this, which points out that people tend to overestimate the probability of unfamiliar, catastrophic, and well-publicized events and underestimate the probability of unspectacular or familiar events that claim one victim at a time. Many people are afraid to fly with commercial airlines, but practically nobody is afraid of driving in cars, a victory of subjectivity over actuarial statistics.

In general, response to risks is most negative when the degree of risk is unknown and the consequences are particularly dreaded. Expert assessment does not seem to help here. People will fight like fury to keep a hazardous waste facility out of their neighborhood, despite expert assurances that it is safe, while people living under high dams located on earthquake faults pay scant attention to expert warnings.

Other hazard characteristics influence public perceptions of risk. For example, the voluntary or involuntary nature of the risk is important. People will accept far greater risks from driving an automobile than they will from breathing the emissions that come out of its tailpipe; the former is voluntary, the latter, involuntary. People also take into consideration whether the risk is distributed generally throughout the population or affects only a small identifiable group. Public response to the discovery of a toxicant that may result in 200 additional cancers nationwide is liable to be quite different from public response to the same number of cases in one county with a population of say, 3000.

The way risks and options are presented also influences perceptions. You might be worried if you heard that occupational exposure at your job doubled your risk of some serious disease; you might be less worried if you heard that it had increased from one in a million to two in a million. Surveys using physicians as subjects found that their preferences for treatment options changed markedly when the risks of these options were expressed in terms of lives saved rather than in terms of deaths occurring, even though the two forms of expression that were compared were mathematically identical. Finally, re-

search has shown that beliefs about risk are slow to change, and show extraordinary persistence in the face of contrary evidence.

Many people interested in environmental protection, having observed this mess, conclude that considerations of risk lead to nothing useful. After all, if the numbers are no good and the whole issue is so confusing, why not just eliminate all exposure to toxics to the extent that technology allows? The problem with such thinking is that, even setting aside what I have just said about the necessity for improving the national debate on the subject, risk estimates are the only way we have of directing the attention of risk management agencies toward significant problems.

There are thousands of substances in the environment that show toxicity in animals; we can't work on all of them at once, even with an EPA ten times its current size. More important, technology doesn't make the bad stuff "go away;" in most cases it just changes its form and location. We have to start keeping track of the flow of toxics through the environment, to what happens *after* they are "controlled." Risk management is the only way I know to do this.

In confused situations one must try to be guided by basic principles. One of my basic principles is reflected in a quotation from Thomas Jefferson: "If we think (the people) not enlightened enough to exercise their control with a wholesome discretion, the remedy is not to take it from them, but to inform their discretion." Easy for *him* to say. As we have seen, informing discretion about risk has itself a high risk of failure.

However, we do have some recent experience that supports the belief that better information inclines people to act more sensibly. In Tacoma, Washington, we have a situation where a copper smelter employing around 600 people is emitting substantial amounts of arsenic, which is a human carcinogen. We found that the best available technology did not reduce the risk of cancer to levels the public might find acceptable. In fact, it looked as if reducing to acceptable levels of risk might only be possible if the plant closed. I felt very strongly that the people in Tacoma whose lives were to be affected by my decision ought to have a deeper understanding of the case than they could get from the usual public hearing process.

Accordingly, we organized an extraordinary campaign of public education in Tacoma. Besides the required public hearing, we provided immense quantities

of information to all communications media, arranged meetings between community leaders and senior EPA officials, including myself, and held three workshops at which we laid out our view of the facts. I think most people appreciated this opportunity, and we certainly raised the level of discussion about risk. So unusual was this kind of event that some inferred that I was abdicating my responsibility for this decision, or that somehow the Tacoma people were going to vote on whether they wanted jobs or health. After some initial confusion on this score we made it clear that it was entirely my decision, and that while I wanted to hear, I was not committed to heed.

Although I suppose some would have been happier continuing in their fond belief that we could provide absolute safety with absolute certainty, and were disturbed by these proceedings, in all I would call it a qualified success. Those who participated came away with a better understanding of the anatomy of environmental decisions, and local groups were able to come up with options that increased protection while allowing the plant to remain open, options that are well worth considering as we put together our final decision.

What are the lessons of Tacoma? Shortly after we began the workshops, people started sporting buttons that said, "BOTH," meaning they were for both jobs and health. I took this as a good sign, that people were attending to the balance between economic realities and environmental protection. "Both" is a good idea, and in most cases we can have it, if we're smart. Another lesson is that we must improve the way we present risk calculations to the public. There was too much tendency to translate risks of cancer into cases, with no regard to qualifying assumptions and uncertainties. Cancer threats make great headlines and the inclination to infer certainty where none exists is very powerful. We must take seriously our obligation to generate lucid and unambiguous statements about risk. Finally, Tacoma shows that we have to prepare ourselves for the other Tacomas. Environmental stress falls unevenly across the land and we have a special responsibility to people in communities that suffer more than their share. We are prepared to make the extra effort in such communities, as we did in Tacoma.

We must also improve debate on the national level. This may prove more difficult, as Washington is a most contentious place. Also, at the national level things tend to polarize perhaps more than they should, given how much we know about environmental health questions. Typically, where we obtain evidence of an

environmental threat, opinion divides between those who want to eliminate the risk as quickly as possible, with little concern about cost, and those who deny the threat exists. Fights between these groups can go on for a long time, time during which the object of the battle, the pollutant, remains in the environment. Acid rain threatens to become this kind of dispute.

And so too was the case of ethylene dibromide. As you may know, we recently banned the major uses of EDB, a grain and fruit fumigant that has been identified as a carcinogen, and which enters the human diet through residues in food and via ground-water contamination. By means of that ban, which applied to grain fumigation, we insured that EDB would immediately begin to diminish in the human food supply. Since there is still EDB in the grain products already in storage or on grocers' shelves, we set maximum acceptable residue levels for different products, the levels getting lower in products closer to the point of consumption. We will act soon on the use of EDB as a citrus fruit fumigant, its only remaining use in connection with the human food chain.

Needless to say, we were criticized both for going too far and for not going far enough. But in cases such as this, my personal predilection is to avoid the extremes and act to reduce, as quickly as possible, environmental exposure to substances that appear unacceptably risky, and to do so with as little social or economic disruption as possible. This generally satisfies no one, but I am convinced it is in the long term public interest.

What *was* dissatisfying about the EDB case was the substantial confusion surrounding the risk issues involved. Some say that we stir up cans of worms when we expose the risk judgments we make. I think we must do better than we have done, and let the worms crawl where they may. Let me now propose some principles for more reasonable discussions about risk.

First, we must insist on risk calculations being expressed as distributions of estimates and not as magic numbers that can be manipulated without regard to what they really mean. We must try to display more realistic estimates of risk to show a range of probabilities. To help do this we need new tools for quantifying and ordering sources of uncertainty and for putting them in perspective.

Second, we must expose to public scrutiny the assumptions that underly our analysis and management of risk. If we have made a series of conservative assumptions within the risk assessment, so that it

represents an upper bound estimate of risk, we should try to communicate this and explain why we did it. Although public health protection is our primary value, any particular action to control a pollutant may have effects on other values, such as community stability, employment, natural resources, or the integrity of the ecosystem. We have to get away from the idea that we do quantitative analysis to find the "right" decision, which we will then be obliged to make if we want to call ourselves rational beings. But we are not clockwork mandarins. The point of such analysis is, in fact, the orderly exposition of the values we hold, and the reasoning that travels from some set of values and measurements to a decision.

Third, we must demonstrate that reduction of risk is our main concern and that we are not driven by narrow cost-benefit considerations. Of course cost is a factor, because we are obliged to be efficient with our resources and those of society in general. Where we decline to control some risk at present, we should do so only because there are better targets; we are really balancing risk against risk, aiming to get at the greatest first.

Finally, we should understand the limits of quantification; there are some cherished values that will resist being squeezed into a benefits column, but are no less real because of it. Walter Lippman once pointed out that in a democracy "the people" as in "We the People," refers not only to the working majority that actually makes current decisions, and not only to the whole living populations, but to those who came before us, who provided our traditions and our physical patrimony as a nation, and to those who will come after us, and inherit. Many of the major decisions we make on environmental affairs touch on this broader sense of public responsibility.

I suppose that the ultimate goal of this effort is to get the American people to understand the difference between a safe world and a zero-risk world with respect to environmental pollutants. We have to define what safe means in light of our increasing ability to detect minute quantities of substances in the environment and to associate carcinogenesis with an enormous variety of substances in common use. According to Bruce Ames, the biochemist and cancer expert, the human diet is loaded with toxics of all kinds, including many carcinogens, mutagens, and teratogens. Among them are such foodstuffs as black pepper, mushrooms, celery, parsnips, peanut butter, figs, parsley, potatoes, rhubarb, coffee, tea, fats, browned meat, and alfalfa sprouts. The list goes on: my point is that it would be hard to find a diet that

would support life and at the same time impose no risk on the consumer.

So what is safe? Are we all safe at this instant? Most of us would agree that we are, although we are subjected to calculable risks of various sorts of catastrophes that can happen to people listening to lectures in buildings. We might be able to reduce some of them by additional effort, but in general we consider that we have (to coin a phrase) an “adequate margin of safety” sitting in a structure that is, for example, protected against lightning bolts but exposed to meteorites.

I think we can get people to start making those judgments of safety about the arcane products of

modern technology. I don't think we are even going to get agreement about values; a continuing debate about values is the essence of a democratic policy. But I think we must do better in showing how different values lead rationally to different policy outcomes. And we can only do that if we are able to build up a reservoir of trust, if people believe that we have presented what facts we have fairly, that we have exposed our values to their view, and that we have respected their values, whether or not such values can be incorporated finally in our decisions. We have, I hope, begun to build that sort of trust at EPA.