

5 Assessing and Communicating the Risks of Terrorism

Baruch Fischhoff

Assessing and communicating the risks of terrorism requires the collaboration of the full spectrum of social and behavioral sciences. This chapter discusses research that we could and, were we to be responsible, should apply to the problem of terrorism. I will also give a feeling for how we might do the kind of systems engineering that Dr. Lewis Branscomb calls for in Chapter 2 of this volume. This will require integrating the social, behavioral, natural, and engineering sciences.

This chapter examines, in turn, the psychology of risk (both for experts and the public); risk analysis and risk communication (and how the two must be integrated for either to be effective); special considerations in the domain of terrorism; how we might begin to apply these perspectives to bioterrorism; and, finally, some areas where we might immediately begin to develop applications and conduct the supporting basic science.

The Psychology of Risk—The Public

The public is important in our response to terror, both as actors and as audiences. We need to communicate with people well in advance of any terror-related crisis. They need to have some idea of

Baruch Fischhoff is university professor in the Department of Social and Decision Sciences and the Department of Engineering and Public Policy at Carnegie Mellon University. This chapter is based on remarks delivered at the 27th Annual AAAS Colloquium on Science and Technology Policy held April 11-12, 2002, in Washington, DC.

what is going on, in order to have a chance to make effective decisions, being as heroic as they choose—without feeling they were misled or incompletely informed.

Understanding the psychology of the public is also important for anticipating how people will respond to our plans. For example, Dr. Eugene Spafford's paper (Chapter 4 in this volume) discussed various plans for protecting our computer systems. Each plan assumes some behavior by the computers' human operators, such as respecting one another's privacy and protecting one's own. If we do not understand human behavior, then we have *behaviorally unrealistic* plans. Including the human sciences adds a level of complexity to already complicated planning processes. Yet without them, we are blinding, and perhaps deluding, ourselves.

One long-standing focus of research into the psychology of decision making is how people's current beliefs shape their future understanding. Knowing the details of these processes is essential for effective communication. If we do not know where people are coming from, it is very difficult for us to get them to another place. People's ability to process risk communications depends on their numeracy and literacy. Numeracy is required to understand how big risks are (and how much risk-reduction measures will cost). Language literacy is required to process written messages. Scientific literacy is needed to grasp the content of messages that, with terror, can involve a large number of domains. For example, we need to know something about anthrax, about the (foreign or domestic) people who may be behind an outbreak, about diffusion rates for small particles, about the management responsibility of various government agencies, and so on. How well one can understand the anthrax crisis depends on one's literacy in these different domains.

People's responses are also constrained by their limited cognitive capacity, which can shrink further under crisis conditions. Given these limits, people manage to function either by acquiring domain-specific knowledge or by relying on robust, but imperfect *heuristics*. These "rules of thumb" simplify problems and provide approximate answers. But they can also produce biases. For example, people seem to count, almost automatically, how frequently they see various events. Those estimates can be useful in estimating the frequency of such events—unless appearances are deceiving, such that some

events are disproportionately visible, leading to overestimating their actual frequency. People may not often think about the representativeness of the evidence that they see. When they do ask that question, they may not adjust adequately from what they have seen to what is actually there.

Researchers relying on psychological theories and methods have found it possible to increase people's understanding of many risks. Nonetheless, some concepts are inherently difficult to communicate. One challenge is giving a feeling for very low probabilities. We have 285 million people in the denominator, when thinking about the risk of terror for an individual in the United States. However, our perception of these risks may be unduly influenced by a relatively small number of very salient incidents in the numerator. A second challenge is conveying notions of cumulative risk, arising from repeated exposure. A particular event might be very unlikely on a given day (or trip), but over time, those tiny probabilities can mount up—and at a rate that people do not realize. A third problem arises with verbal quantifiers like “likely,” which can be very confusing. “Likely” means different things to different people, and different things to the same people in different situations. If you try to communicate the size of risks using words instead of numbers, you are setting a trap for your audience.

People have difficulty making decisions about events that they have never experienced. It is hard to project oneself into unfamiliar situations. As a result, we have difficulty predicting our own responses to events. A growing literature shows how decision-making difficulty can reflect uncertainty about ourselves (or “value uncertainty”) as well as uncertainty about the world. In effect, we do not really know what we want in many novel situations. Difficult medical decisions often evoke such feelings.

Our responses to risks reflect our emotions, as well as our beliefs. Emotions can both confound and support our cognition. They can empower us to act, but also paralyze us. They both affect and reflect our beliefs. Terrorism evokes a wide range of emotions, which must be understood if we are to aid and predict citizens' choices. These emotions include fear of the effects of terrorism, frustration with ourselves and our authorities, mourning, and solidarity with our fellow citizens.

The Psychology of Risk—The Experts

Experts have uncertain beliefs and emotions, too. Novel situations may draw experts into areas that no one understands very well. They may need to interface with other experts, from unfamiliar disciplines. As citizens, we need to understand the psychological processes of our experts, in order to decide how much we can trust them. As experts, we need to have—and to convey—a realistic assessment of our own competence, if we are to merit the public's trust. We need to define our domain of expertise and be willing to coordinate with experts from other domains.

We also need to have a clear, consistent public role. We can try to inform people or to persuade them. That is, we can provide facts or we can provide spin. But we cannot mix these roles. If we do, we will confuse our audience, who will not know how to interpret our claims.

Because the problems of terror are so new, they force new groups of experts to communicate with the public and, in doing so, to demonstrate their competence and honesty. There is, however, a learning curve for talking to people about risk. Experts must make rapid progress on this curve if they are to earn trust that is hard to restore, once lost. Unfortunately, a natural first response for many experts is telling citizens to “go away while we do our work.” If people persist in wanting to hear about the risk, it is tempting to tell them what they ought to think—rather than leveling with them, and providing the facts that they need for independent choices.

It is tempting, sometimes, to magnify risks in order to motivate citizens. At other times, it is tempting to trivialize their worries, with comparisons like “why get so exercised about terror, when you're still smoking” or “only five people have died from anthrax [so far], compared with 40,000 annually from motor vehicle accidents.” Except for extreme situations (e.g., rapid evacuations), experts must seek a partnership with the public.

We saw the result of confused expert roles in the anthrax crisis. We also saw experts disparaging the public, reflecting the limits to their own psychology. Experts may be biased by their limited opportunities to observe the public—and their failure to recognize the un-

representativeness of what they see (i.e., how citizens appear in heated controversies, “person in the street” interviews, or responses to ambiguous survey questions). Experts may also have ego involvement in how the public is viewed. They may show professional arrogance and defensiveness. They may be concerned with defending their own expert status, hence benefit from deprecating the public. They are just people, also feeling the pressures of the times, whatever their professional training.

Risk Analysis and Communication

One way of disciplining expert judgment is to perform formal risk analyses. That means identifying valued outcomes, the processes affecting them, and the experts with the best understanding of each. These experts must be asked to pool their beliefs, uncertainties, controversies, and omissions—then subject their work to independent peer review. In principle, risk analysis is no different than any other modeling process. Yet, in any new area, analytical conventions need to evolve. Without them, relying on experts’ intuitive judgments may create misleading pictures.

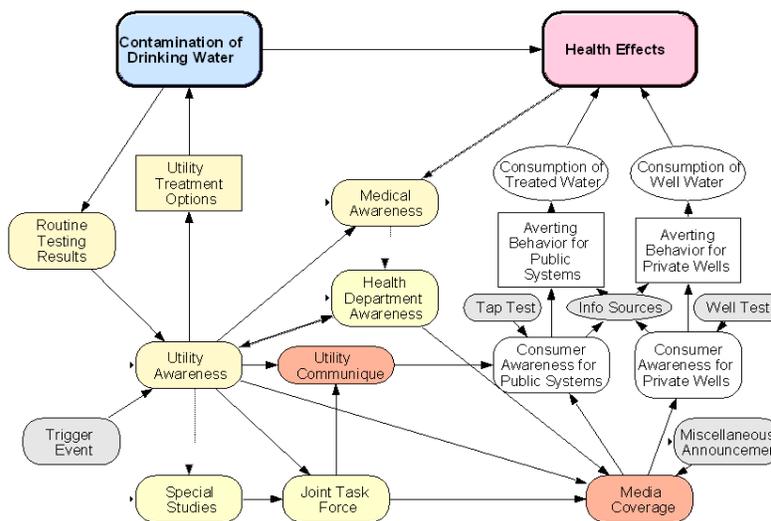
For example, we recently conducted a risk analysis for contamination of domestic water supplies by *cryptosporidium*, a water-borne parasite. Milwaukee had an outbreak about 10 years ago, where about 100 people died and 400,000 got sick. We were asked to develop the perfect “boil water” notice. Through interviews, we found that many people did not know how to boil water effectively. We also found that many people would want to know who produced a notice, before deciding how seriously to take it. Thus, a boil-water notice might need to explain the risk management system that produced it, in addition to instructions about what to do.

Figure 1 shows the top level of our risk analysis model (led by Liz Casman, a microbiologist.) It integrates the engineering science of managing water supplies, the biological science of parasites’ effects on our bodies, the communication science of how messages get organized and disseminated, and, finally, the psychology of what happens when people make such risky decisions.

Once we completed the model, we asked what contribution a perfect “boil water” notice would make, assuming that everyone e-

ceived and believed the notice, then boiled water the right way. How would this affect a typical epidemic? We ran the numbers and found that it had no effect whatsoever. It took us a couple of anxious weeks to understand what was happening. We hoped our sponsor would not call during this period. We finally realized that with current detection systems, *cryptosporidium* would have already done its damage before anyone knew that it was there.

Figure 1
A Risk Management System for *Cryptosporidium*



Thus, the tests are good for forensic value, determining what hit you after an epidemic, but not for protecting public health. We had a public health system built on unrealistic psychological and engineering assumptions. However, we did not realize this until we actually did the analytical work. Having done so, we realized that the system had misplaced its priorities. In this case, we should not be putting money into better communication, but into better testing technology or into land-use practices that reduce intrusions. Believing in the cur-

rent system means not routinely providing immunocompromised people with bottled water, needlessly exposing them because we had not analyzed the system.

At this time, I was reading some brochures published by the National Cancer Institute (NCI) on how to deal with chemotherapy. They were very nice brochures, but said nothing about risk from water-borne sources. I called some colleagues at NCI and at the Centers for Disease Control and Prevention. They said that they had thought about this risk, but somehow that concern had never found its way into the system. I talked to the person responsible for the brochures who said, "That's a good point, but we just printed several million of them. Call us back when the supply runs out." This example shows that we have to set priorities explicitly. If we are not systematic about this, we may invest our money in the wrong places and expose people to needless risk.

Once we have figured out which facts are important to know, developing risk communications is relatively straightforward. First, we need to determine what information is common knowledge, hence goes without saying. Knowing what people already know avoids wasting their time and losing their respect (by not giving them credit for that knowledge). Such common knowledge can be identified with open-ended interviews, allowing the full expression of intuitive beliefs, values, and formulations. Structured surveys allow estimating the frequency of different beliefs.

The next step in designing communications is characterizing the critical gaps in lay beliefs, representing what is worth knowing. Many risk situations require understanding both quantitative information (how big a risk is, how much it will cost to reduce it) and qualitative information (what determines the risk, where it comes from, how it is assessed, how it reveals itself in everyday life). One then needs a story line to communicate the information that matters. People need a coherent mental model, giving qualitative meaning to the quantitative statistics, building on the constructive processes of learning and memory. The success of any communication must be empirically evaluated. We have been studying these topics for 25-plus years and still are surprised by what we learn from the testing process.

The Special Challenges of Terrorism

What is special about terrorism within this general context of risk analysis and risk communication? One important feature is the diversity of people who must work together in order to address these very complicated problems. They need to create broadly shared mental models in order to coordinate their actions and beliefs. They also need to reconcile their mixed motives. In addition to fighting terror, their actions will affect their own status and our society. For example, airport security is about flying safety, but also about the respective roles of the public and private sectors in our society. As citizens, we, too, have mixed motives. We want facts, but also reassurance. We want to know whom to blame, but also to feel solidarity with our fellow citizens.

A second feature of terror-related events is that they challenge the validity of our experience. This is novel ground, even for the professional community. We must deal with unfamiliar topics, unfamiliar people and places, and unfamiliar pathogens. As a result, terrorism requires theoretical understanding to augment the historical statistics. Because the old statistics may not be valid for the estimating, say, aviation or anthrax risks, we need models integrating these theories. Unless we recognize and interpret these pieces, we are working at cross-purposes.

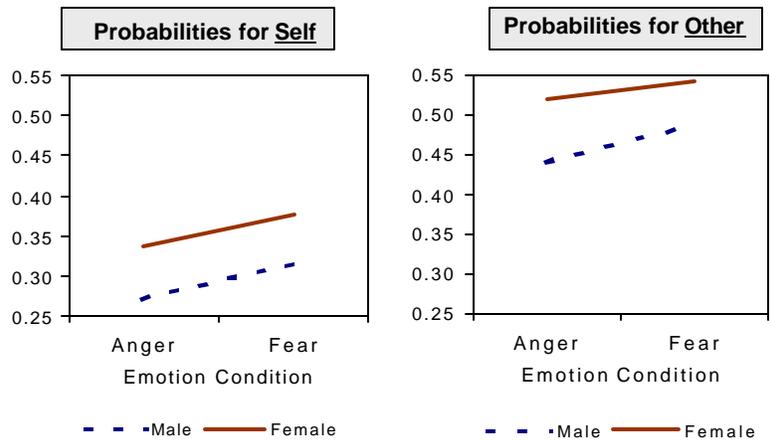
A third distinguishing feature of terror-related events is the intensity of the emotions involved.

We were able to conduct a study examining the effects of these emotions on risk judgments (with support from the National Science Foundation). We had access to a random sample of a thousand Americans, through WebTVs in their homes. We experimentally heightened three emotions that are naturally associated with terrorism—fear, anger, and sadness. For example, the anger group saw a picture of people supporting Osama bin Laden, with a voice-over of quotes from CNN and *The New York Times*. They also wrote about their anger regarding the events of September 11. A second group received a fear prime (showing gloved postal clerks sorting mail), while a third group received a sadness prime (a woman reading a letter from her husband who had died at the World Trade Center).

A precursor to our study was recent research showing that angry people tend to be more optimistic, and see lower probabilities of bad things happening to them. Before this finding, researchers generally believed that all bad emotions went together: People who are depressed also feel hopeless, and so on. But anger is a different kind of negative emotion.

Figure 2 shows these individuals mean judgments of the probabilities of eight risky events, summarized so that higher mean probabilities indicate higher perceived risk. The figure shows that people made angry about the attacks were more optimistic about their prospects.

Emotion & Gender Effects on Terror Risk Judgments (mid-November 2001; n=973)



The figure also shows that women see greater risks than men. This result is consistent with other research finding that women report lower degrees of anger than men—as was found here as well, both for the anger that they brought with them to the study and their responses to the emotion primes.

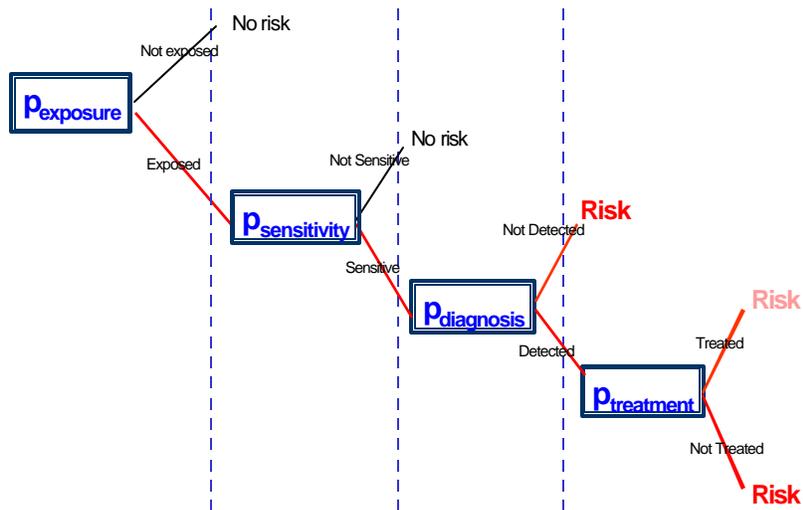
Many previous studies have found that people believe that others are at greater risk than they are. The comparison between the two graphs shows a replication of this result. The graph on the right shows the mean risk for the “average American.” The graph on the

left shows the average of these average Americans' judgments of their personal risks.

A Worked Example: Bioterrorism

Last fall, I was on a panel with Dave Piposar from the Allegheny County Department of Health. He described how his department was dealing with the flood of anthrax-related calls. He said that one thing that people did not realize is that they have to be exposed in order to have a risk. In thinking about how to organize a risk analysis that would serve the communication needs of these callers, we produced the very simple model of Figure 3.

Figure 3
A Common Structure for Bioterrorism Risk

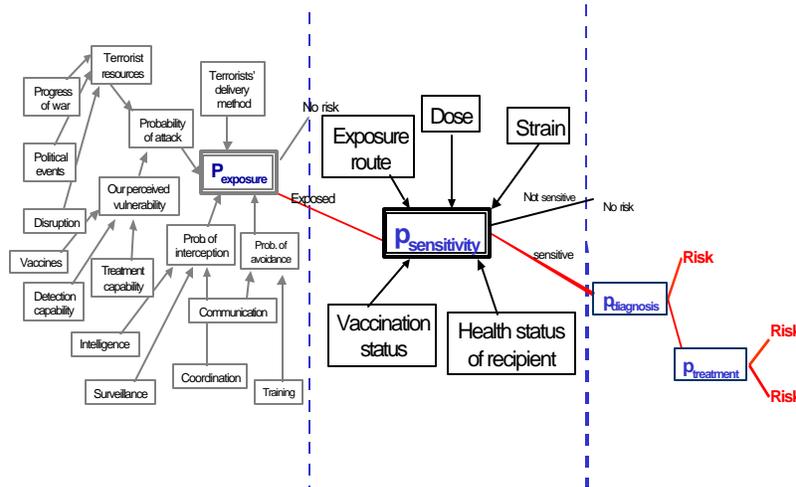


In order to have a risk from bioterrorism, you need to be exposed and you need to be sensitive at that dose. You need to miss the diagnosis or get the diagnosis right and have treatment fail. Thus, these risks have a multiplicative structure: everything needs to go

wrong. An effective communication would recall the full picture, even if focused on a specific issue in it.

Figure 4 shows a set of things that one would like to know in order to predict the probability of exposure. Conducting such analyses requires a team with expertise in each node. If we have not assembled that team, then we are not doing our job. Missing issues entirely may be more problematic than inaccurately measuring an issue that is considered.

Figure 4
What Determines the Probability of Exposure?



Producing such an integrated model is important, even if we do not “run the numbers” to produce quantitative estimates. It forces rigor in specifying variables and relationships, and in identifying relevant expertise. Given such a model, one could rapidly update it, if the situation changed. One could use it as a basis for estimating the risks of an outbreak of smallpox. Smallpox is different than anthrax because it is contagious. But many of the processes in the anthrax model recur with smallpox and many other disease agents, even when the details are different. The science that went into building the first

model could be used again in others. One could also look for processes that recur in multiple models, such as the enemy's probability of attack, our response capability (monitoring, coordination, etc.), and other valued consequences. We should invest in studying those crosscutting issues.

Like other risk analyses, such a model provides a guide for risk communications. It shows how to structure the long-term education needed to give citizens the big picture represented by the model as a whole, prepare them for the varying faces of the long struggle to come. It shows how to get ahead of the game, preparing and evaluating communications for crisis use. To the extent possible, such real-time risk communications should have a common (and pre-tested) structure for topics like exposure and detection, and a common format for expressing risk levels and uncertainties. Everything should reflect the science of risk communication.

Another topic requiring advance research is dealing with false alarms. We need to maintain consistent alarm standards (as proposed by Governor Ridge's color-coded alert levels). Explaining the philosophy underlying them will help citizens to understand why false alarms are inevitable and minimize their costs—including cumulative apathy. When false alarms do occur, we need to minimize both disruption and the perception of cover-ups.

A third research topic is how to deal with the second-guessing that follows terror-related events, neither succumbing to hindsight bias nor hiding behind it. Learning from the experience requires considering our leaders' entire decision-making process. Could they have acted on their knowledge? Was leadership possible? What did they know? What could they reasonably have known? How clear was the signal? We must judge the quality of their choice, not its outcome. These inquiries, too, should be guided by the science of decision making.

As mentioned, the public may be second-guessed by experts who doubt citizens' competence. Here, we must distinguish between ignorance and stupidity. How good was the communication to citizens? How defensible are their misunderstandings? How good is our evidence about the public? Evaluating citizens' actions requires knowing what problems they are solving. This includes understanding their options, values, and beliefs.

A fourth topic for advance research is setting priorities, linking public opinion with public policy. We need to present citizens with a full range of options and predicted outcomes, if we hope to know how they want us to respond to terror. We need a structured public discussion for how to pursue the long task ahead of us.

Our national survey asked about the relative importance of four priorities. We found strong support for two. One is that Americans want honest and accurate information about terror-related situations, even if that information worries them. This was true whether people were in the anger or the fear condition of our experiment. Second, we found very strong support for investing in general capabilities (like stronger public health) rather than in specific solutions (like smallpox vaccination). This policy was supported slightly less by people in the anger condition. Two priorities given somewhat less importance, were deporting foreigners who lack visas from the United States, and strengthening ties with Muslim countries.

Conclusion

Effective risk analysis and communication require quantitative estimates of risk (including the attendant uncertainties) and explicit representation of the processes shaping those risks. Producing them requires suitably diverse expertise. It also requires integrating risk analysis and communication, so we can solve people's problems and secure their trust.

Acknowledgements

I would like to thank Liz Casman, Matt Dombroski, Sara Eggers, Dalia Patina Echerverri, Paul Fischbeck, Roxana Gonzalez, Umit Guvenc, Jennifer Lerner, Deborah Small, Conrad Steenkamp, and the National Science Foundation, for their roles in the work reported here. The views expressed are my own.

Additional Resources

- Casman, E., Fischhoff, B., Palmgren, C., Small, M., & Wu, F. (2000). Integrated risk model of a drinking waterborne *Cryptosporidiosis* outbreak. *Risk Analysis*, 20, 493-509.
- Fischhoff, B. (1992). Giving advice: Decision theory perspectives on sexual assault. *American Psychologist*, 47, 577-588.
- Fischhoff, B. (1995). Risk perception and communication unplugged: Twenty years of process. *Risk Analysis*, 15, 137-145.
- Fischhoff, B. (1998). Communicate unto others... *Reliability Engineering and System Safety*, 59, 63-72.
- Fischhoff, B. (1999). What do patients want? Help in making effective choices. *Effective Clinical Practice*, 2(3), 198-200.
- Fischhoff, B. (2000). Scientific management of science? *Policy Sciences*, 33, 73-87.
- Fischhoff, B., Bostrom, A., & Quadrel, M.J. (2002). Risk perception and communication. In R. Detels, J. McEwen, R. Beaglehole & H. Tanaka (Eds.), *Oxford Textbook of Public Health* London: Oxford University Press
- Henrion, M. & Fischhoff, B. (1986). Assessing uncertainty in physical constants. *American Journal of Physics*, 54, 791-798.
- Lerner, J. S., & Keltner, D. (2001). Fear, anger, and risk. *Journal of Personality & Social Psychology*, 81(1), 146-159.
- Morgan, M.G., Fischhoff, B., Bostrom, A., & Atman, C. (2001). *Risk communication: The Mental Models Approach*. New York: Cambridge University Press.
- National Research Council. (1996), *Understanding Risk*. Washington, DC: Author.
- Organisation for Economic Co-operation and Development. (2002). *Guidance Document on Risk Communication for Chemical Risk Management*. Draft, 6 March.
- Performance and Innovation Unit. (2002). *Risk and Uncertainty*. London: Parliament.
- Slovic, P. (Ed.). (2001). *The Perception of Risk*. London: Earthscan.