

33. Facts versus fears: Understanding perceived risk

Paul Slovic, Baruch Fischhoff, and Sarah Lichtenstein

People respond to the hazards they perceive. If their perceptions are faulty, efforts at personal, public, and environmental protection are likely to be misdirected. For some hazards, such as motor vehicle accidents, extensive statistical data are readily available. For other familiar activities, such as the consumption of alcohol and tobacco, assessment of risk requires complex epidemiological and experimental studies. However, even when statistical data are plentiful, the "hard" facts can only go so far toward developing policy. At some point, human judgment is needed to interpret the findings and determine their relevance.

Still other hazards, such as those associated with recombinant DNA research or nuclear power, are so new that risk assessment must be based on complex theoretical analyses such as fault trees (see Figure 1), rather than on direct experience. Despite an appearance of objectivity, these analyses, too, include a large component of judgment. Someone, relying on educated intuition, must determine the structure of the problem, the consequences to be considered, and the importance of the various branches of the fault tree. Once the analyses have been performed, they must be communicated to those who actually manage hazards, including industrialists, environmentalists, regulators, legislators, and voters. If these people do not understand or believe the data they are shown, then distrust, conflict, and ineffective hazard management are likely.

This chapter explores some psychological elements of the risk-assessment process. Its basic premises are that both the public and the experts

This is a revised version of a paper that originally appeared in R. Schwing and W. A. Albers Jr. (Eds.), *Societal Risk Assessment: How safe is safe enough?* New York: Plenum Press, 1980. Copyright © 1980 by Plenum Press. Reprinted by permission.

Support for this work was provided by the Technology Assessment and Risk Analysis Program of the National Science Foundation under Grant PRA79-11934 to Clark University under subcontract to Perception, Inc.

biases with serious implications for decision making in areas as diverse as financial analysis (Slovic, 1972c) and the management of natural hazards (Slovic, Kunreuther, & White, 1974).

Availability

One heuristic that has special relevance for risk perception is called availability (Tversky & Kahneman, 1973, 11). People using this heuristic judge an event as likely or frequent if instances of it are easy to imagine or recall. Because frequently occurring events are generally easier to imagine and recall than are rare events, availability is often an appropriate cue. However, availability is also affected by numerous factors unrelated to frequency of occurrence. For example, a recent disaster or a vivid film, such as *Jaws* or *The China Syndrome*, could seriously distort risk judgments.

Availability bias helps explain people's misperceptions and faulty decisions with regard to certain natural hazards. For example, in discussing flood plain residents, Kates (1962) wrote:

A major limitation to human ability to use improved flood hazard information is a basic reliance on experience. Men on flood plains appear to be very much prisoners of their experience. . . . Recently experienced floods appear to set an upward bound to the size of loss with which managers believe they ought to be concerned. (p. 140)

Kates attributed much of the difficulty in improving flood control to the "inability of individuals to conceptualize floods that have never occurred" (Kates, 1962, p. 92). He observed that individuals forecasting flood potential "are strongly conditioned by their immediate past and limit their extrapolation to simplified constructs, seeing the future as a mirror of that past" (p. 88). Similarly, the purchase of earthquake insurance increases sharply after a quake and then decreases steadily as memories fade (Steinbrugge, McClure, & Snow, 1969).

One particularly important implication of the availability heuristic is that discussion of a low-probability hazard may increase its memorability and imaginability and hence its perceived riskiness, regardless of what the evidence indicates. For example, leaders in the field of recombinant DNA research quickly regretted ever bringing to public attention the remote risks of contamination by newly created organisms. Rosenberg (1978) summarized the reaction that followed the revelation of such hypothetical risks:

Initially, the response was one of praise for the . . . social responsibility shown by the scientists involved. . . . Gradually and predictably, however, the debate became heated. Speculation abounded and the scarier the scenario, the wider the publicity it received. Many of the discussions of the issue completely lost sight of the fact that the dangers were hypothetical in the first place and assumed that recombinant

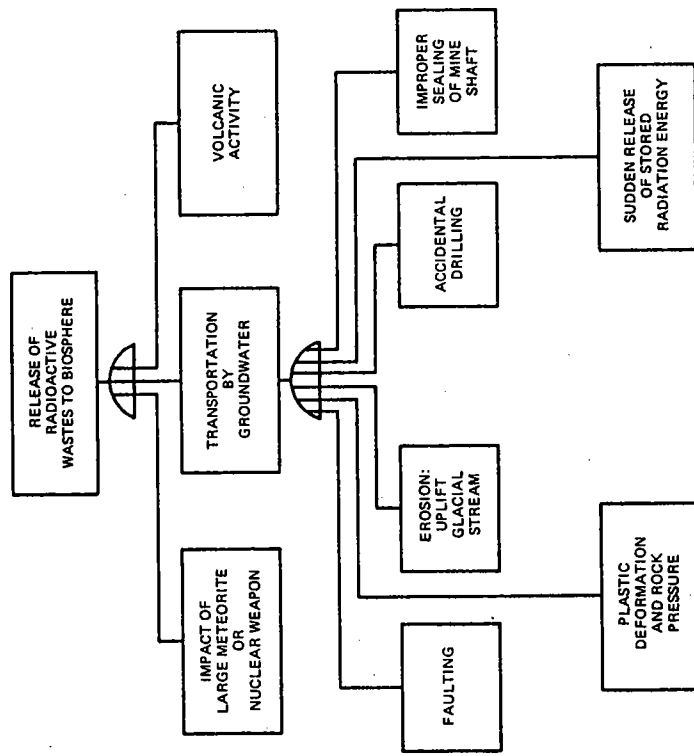


Figure 1. A fault tree indicating the various ways in which radioactive material might accidentally be released from nuclear wastes buried within a salt deposit. Each of the possible initiating events in the bottom two rows can lead to the transportation of radioactivity by groundwater. This transport can in turn release radioactivity to the biosphere. As indicated by the second level of boxes, release of radioactivity can also be produced directly (without the help of groundwater) through the impact of a large meteorite, a nuclear weapon, or a volcanic eruption. (Source: McGrath, 1974.)

are necessary participants in that process, that assessment is inherently subjective, and that understanding judgmental limitations is crucial to effective decision making.

Judgmental biases in risk perception

When laypeople are asked to evaluate risks, they seldom have statistical evidence on hand. In most cases, they must make inferences based on what they remember hearing or observing about the risk in question. Psychological research, much of which has been described earlier in this book, has identified a number of very general inferential rules that people seem to use in such situations. These judgmental rules, known as heuristics, are employed to reduce difficult mental tasks to simpler ones. Although they are valid in some circumstances, in others they lead to large and persistent




Figure 2. Relationship between judged frequency and the actual number of deaths per year for 41 causes of death. If judged and actual frequencies were equal, the data would fall on the straight line. The points, and the curved line fitted to them, represent the averaged responses of a large number of laypeople. As an index of the variability across individuals, vertical bars are drawn to depict the 25th and 75th percentiles of the judgments for botulism, diabetes, and all accidents. The range of responses for the other 37 causes of death was similar.

DNA laboratories were full of raging beasts. Ultimately, the very scientists whose self-restraint had set the whole process in motion were vilified. (p. 29)

Judged frequency of lethal events. Availability bias is illustrated by several studies in which college students and members of the League of Women Voters judged the frequency of 41 causes of death (Lichtenstein et al., 1978). In one study, these people were first told the annual death toll from 1 cause (motor vehicle accidents) in the United States (50,000) and then were asked to estimate the frequency of the other 40. In another study, participants were asked to judge which of 2 causes of death was more frequent. In both studies, judgments were moderately accurate in a global sense: People usually knew which were the most and least frequently lethal events. Within this global picture, however, people made serious misjudgments, many of which seemed to reflect the influence of availability.

Figure 2 compares the judged number of deaths per year with the number reported in public health statistics. If the frequency judgments were accurate, they would equal the statistical rates, with all data points

Table 1. Bias in judged frequency of death

Most overestimated	Most underestimated
All accidents	Smallpox vaccination
Motor vehicle accidents	Diabetes
Pregnancy, childbirth, and abortion	Stomach cancer
Tornadoes	Lightning
Flood	Stroke
Botulism	Tuberculosis
All cancer	Asthma
Fire and flames	Emphysema
Venomous bite or sting	
Homicide	

Source: Slovic, Fischhoff, & Lichtenstein (1979).

falling on the identity line. Although more likely hazards generally evoked higher estimates, the points seem scattered about a curved line that lies sometimes above and sometimes below the line of accurate judgment. In general, rare causes of death were overestimated and common causes of death were underestimated.

In addition to this general bias, sizable specific biases were evident. For example, accidents were judged to cause as many deaths as diseases, whereas diseases actually take about 16 times as many lives. Homicides were incorrectly judged more frequent than diabetes and stomach cancer deaths. Homicides were also judged to be about as frequent as death by stroke, although the latter actually claims about 11 times as many lives. Frequencies of death from botulism, tornadoes, and pregnancy (including childbirth and abortion) were also greatly overestimated. Table 1 lists the lethal events whose frequencies were most poorly judged in our various studies. In keeping with availability considerations, overestimated causes of death were dramatic and sensational, whereas underestimated causes tended to be unspectacular events, which claim one victim at a time and are common in nonfatal form.

Biased newspaper coverage and biased judgments. The availability heuristic highlights the vital role of experience as a determinant of perceived risk. If one's experiences are biased, one's perceptions are likely to be inaccurate. Unfortunately, much of the information to which people are exposed provides a distorted picture of the world of hazards. Consider author Richard Bach's observation about the fear shown by a couple taking their first airplane ride:

In all that wind and engineblast and earth tilting and going small below us, I watched my Wisconsin lad and his girl, to see them change. Despite their laughter, they had been afraid of the airplane. Their only knowledge of flight came from

newspaper headlines, a knowledge of collisions and crashes and fatalities. They had never read a single report of a little airplane taking off, flying through the air and landing again safely. They could only believe that this must be possible, in spite of all the newspapers, and on that belief they staked their three dollars and their lives. (Bach, 1973, p. 37)

As a follow-up to the studies reported above, Combs and Slovic (1979) examined the reporting of causes of death in two newspapers on opposite coasts of the United States. Various indices of newspaper coverage were recorded for alternate months over a period of one year. The results indicated that both newspapers had similar biases in their coverage of life-threatening events. For example, examination of Table 2 reveals that many of the statistically frequent causes of death (e.g., diabetes, emphysema, various forms of cancer) were rarely reported by either paper during the period under study. In addition, violent, often catastrophic, events such as tornadoes, fires, drownings, homicides, motor vehicle accidents, and all accidents were reported much more frequently than less dramatic causes of death having similar (or even greater) statistical frequencies. For example, diseases take about 16 times as many lives as accidents, but there were more than 3 times as many articles about accidents, noting almost 7 times as many deaths. Among the more frequent events, homicides were the most heavily reported category in proportion to actual frequency. Although diseases claim almost 100 times as many lives as do homicides, there were about 3 times as many articles about homicides as about disease deaths. Furthermore, homicide articles tended to be more than twice as long as articles reporting disease and accident deaths.

Moreover, the biases in newspaper coverage and people's judgments were quite similar. The correlation between judged frequency of death and the number of deaths reported in the newspapers was about .70. This high correlation was not due to a common association of both judged and reported deaths with statistical frequency. When the latter was held constant, the partial correlations between people's judgments and the number of deaths reported were .89 and .85 for the two newspapers. Although it is tempting to conclude from these correlations that media coverage biases perceptions of risk, it might also be the case that people's opinions about what is important influence the media. The journalism literature is replete with instances in which influence has occurred in each direction (Brucker, 1973).

It won't happen to me. People's judgments of causes of death may be about as good as could be expected, given that they are neither specialists in the hazards considered nor exposed to a representative sample of information. Accurate perception of misleading samples of information might also be seen to underlie another apparent judgmental bias, people's predilection to view themselves as personally immune to hazards. The great majority of individuals believe themselves to be better than average drivers

Table 2. Statistical frequency and newspaper coverage in the Eugene, Oregon, Register Guard and the New Bedford, Massachusetts, Standard Times for 41 causes of death

Cause of death	Rate per 2.05 x 10 ⁸ U.S. Res.	Subjects' estimates	Reported deaths				Occur- rences				Articles			
			R-G		S-T		R-G		S-T		R-G		S-T	
			0	57	0	0	0	0	0	0	0	0	0	0
1. Smallpox	0	57	0	0	0	0	0	0	0	0	0	0	0	
2. Poisoning by vitamins	1	102	0	0	0	0	0	0	0	0	0	0	0	
3. Botulism	2	183	0	0	0	0	0	0	0	0	0	0	0	
4. Measles	5	168	0	0	0	0	0	0	0	0	0	0	0	
5. Fireworks	6	160	0	0	0	0	0	0	0	0	0	0	0	
6. Smallpox vaccination	8	23	0	0	0	0	0	0	0	0	0	0	0	
7. Whooping cough	15	93	0	0	0	0	0	0	0	0	0	0	0	
8. Polio	17	97	0	0	0	0	0	0	0	0	0	0	0	
9. Venomous bite or sting	48	350	0	0	0	0	0	0	0	0	0	0	0	
10. Tornado	90	564	36	25	10	6	14	7	7	14	7	7	7	
11. Lightning	107	91	1	0	1	0	1	0	1	0	1	0	0	
12. Non-venomous animal	129	174	4	2	4	2	4	2	4	2	4	2	2	
13. Flood	205	736	4	10	2	2	2	2	2	2	2	2	2	
14. Excess cold	334	314	0	0	0	0	0	0	0	0	0	0	0	
15. Syphilis	410	492	0	0	0	0	0	0	0	0	0	0	0	
16. Pregnancy, birth & abortion	451	1,344	0	0	0	0	0	0	0	0	0	0	0	
17. Infectious hepatitis	677	545	0	0	0	0	0	0	0	0	0	0	0	
18. Appendicitis	902	605	0	0	0	0	0	0	0	0	0	0	0	
19. Electrocuton	1,025	766	5	0	5	0	6	0	6	0	6	0	0	
20. MV/train collision	1,517	689	0	1	0	1	0	1	0	1	0	1	0	
21. Asthma	1,886	506	1	0	1	0	1	0	1	0	1	0	0	
22. Firearm accident	2,255	1,345	8	1	8	1	9	1	9	1	9	1	1	
23. Poison by solid/liquid	2,563	1,013	3	3	1	1	1	1	1	1	1	1	1	
24. Tuberculosis	3,690	658	0	0	0	0	0	0	0	0	0	0	0	
25. Fire and flames	7,380	3,336	94	46	33	9	38	10	38	10	38	10	10	
26. Drowning	7,380	1,684	47	60	44	24	45	37	45	37	45	37	37	
27. Leukemia	14,555	2,496	1	0	1	0	1	0	1	0	1	0	0	
28. Accidental falls	17,425	2,675	15	7	15	6	16	9	16	9	16	9	9	
29. Homicide	18,860	5,582	278	208	167	122	329	199	329	199	329	199	199	
30. Emphysema	21,730	2,848	1	0	1	0	1	0	1	0	1	0	0	
31. Suicide	24,600	4,679	29	19	28	18	36	20	36	20	36	20	20	
32. Breast cancer	31,160	2,964	0	0	0	0	0	0	0	0	0	0	0	
33. Diabetes	38,950	1,476	0	1	0	1	0	1	0	1	0	1	1	
34. Motor vehicle accident	55,350	41,161	298	83	245	69	180	73	180	73	180	73	73	
35. Lung cancer	75,850	9,764	3	2	3	2	4	2	4	2	4	2	2	
36. Stomach cancer	95,120	3,283	0	1	0	1	0	1	0	1	0	1	1	
37. All accidents	112,750	88,879	715	596	421	152	374	177	374	177	374	177	177	
38. Stroke	209,100	7,109	12	4	12	4	13	4	13	4	13	4	4	
39. All cancer	328,000	45,609	25	12	25	12	26	15	26	15	26	15	15	
40. Heart disease	738,000	23,599	49	30	45	25	46	25	46	25	46	25	25	
41. All disease	1,740,450	88,838	111	87	100	76	104	78	104	76	104	78	78	
Total no. of reports (cases 10, 11, 13, 29, 31, 37 & 41)			1,174	945	729	376	860	483	860	483	860	483	483	
Correlations (R-G vs. S-T)			r = .97	r = .94	r = .94	r = .94	r = .94	r = .94	r = .94	r = .94	r = .94	r = .94	r = .98	

Note: R-G = Register Guard; S-T = Standard Times.
Source: Combs & Slovic (1979).

(Nääänen & Summala, 1975; Svenson, 1981), more likely than average to live past 80 (Weinstein, 1980), less likely than average to be harmed by products they use (Rethans, 1979), and so on. Although such perceptions are obviously unrealistic, the risks look very small from the perspective of each individual's experience. Consider automobile driving: Despite driving too fast, tailgating, etc., poor drivers make trip after trip without mishap. This personal experience demonstrates to them their exceptional skill and safety. Moreover, their indirect experience via the news media shows them that when accidents happen, they happen to others. Given such misleading experiences, people may feel quite justified in refusing to take protective actions such as wearing seat belts (Slovic, Fischhoff, & Lichtenstein, 1978).

Out of sight, out of mind. In some situations, failure to appreciate the limits of "available" data may lull people into complacency. In a study by Fischhoff, Slovic, and Lichtenstein (1978), three groups of college student subjects were asked to evaluate the completeness of a fault tree showing the risks associated with starting a car (see Figure 3). One group saw the full tree. Each of the other two groups received a different pruned tree. In one version, the starting, ignition, and mischief branches were missing; the other version lacked branches detailing battery, fuel, and other engine problems.

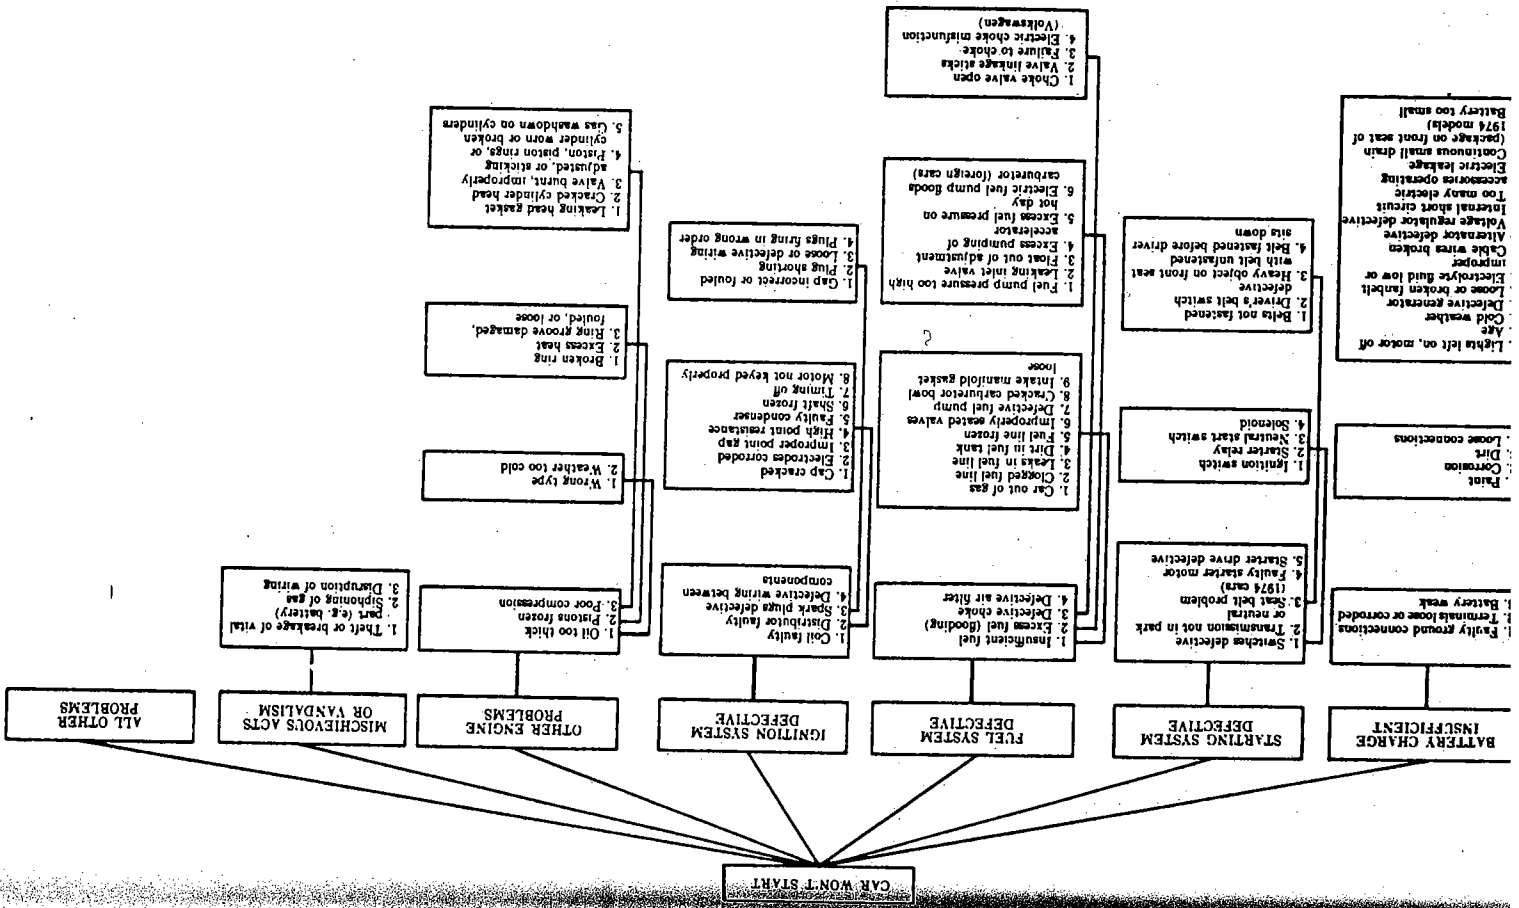
Instructions for the task read as follows (numbers in brackets were given to people who saw the pruned trees):

Every day, across the United States, millions of drivers perform the act of getting into an automobile, inserting a key in the ignition switch, and attempting to start the engine. Sometimes the engine fails to start, and the trip is delayed. We'd like you to think about the various problems that might be serious enough to cause a car to fail to start so that the driver's trip is delayed for at least 1 minute.

The chart on the next page is intended to help you think about this problem. It shows six [three] major deficiencies that cause a car's engine to fail to start. These major categories probably don't cover all possibilities, so we've included a seventh [fourth] category, All Other Problems.

Please examine this diagram carefully and answer the following question: For every 100 times that a trip is delayed due to "starting failure," estimate, on the average, how many of those delays are caused by each of the seven [four] factors. Make your estimates on the blank lines next to the factors named below. Your estimates should sum to 100.

If people who saw the pruned trees were properly sensitive to what had been omitted, the proportion of problems that they attributed to "other" would have equaled the sum of the proportions of problems attributed to the pruned branches and to "other" by those who saw the full tree. The results in Table 3 indicate that what was out of sight was effectively out of mind. For example, in pruned tree Group 1, "other" should have increased by a factor of six (from .078 to .468) to reflect the proportion of failures due



3. Fault tree indicating the ways in which a car might fail to start. It was used by the authors to study whether people are sensitive to the completeness of presentation. Omission of large sections of the diagram was found to have little influence on the judged degree of completeness. In effect, what was

Table 3. Attribution of starting failures for pruned and unpruned trees

Group	n	M proportion of starting failures by type				
		Battery system	Fuel system	Ignition system	Engine	Other
Unpruned tree	93	.264	.195	.144	.076	.051
Pruned tree 1	29	.432	.309	—	.116	—
Pruned tree 2	26	—	.357	.343	—	.073

Note: A dash indicates that the branch was deleted.

^aShould be .468.

^bShould be .611.

Source: Fischhoff, Slovic, and Lichtenstein (1978).

to starting and ignition problems and mischief, which had been omitted from the diagram. Instead, "other" was only doubled, whereas the importance of the three systems that were mentioned was substantially increased. A second study not only replicated these findings but showed that persons who observed pruned trees judged starting failure (due to all causes) to be less likely than did those who observed the unpruned tree.

Overconfidence

Knowing with certainty. A particularly pernicious aspect of heuristics is that people typically have great confidence in judgments based upon them. In another follow-up to the study on causes of death, people were asked to indicate the odds that they were correct in choosing the more frequent of two lethal events (Fischhoff, Slovic, & Lichtenstein, 1977). Table 4 shows the percentages of correct answers for each of the most frequently used odds categories. In Experiment 1, subjects were reasonably well calibrated when they gave odds of 1:1, 1.5:1, 2:1, and 3:1. That is, their percentage of correct answers was close to the appropriate percentage correct, given those odds. However, as odds increased from 3:1 to 100:1, there was little or no increase in accuracy. Only 73% of the answers assigned odds of 100:1 were correct (instead of 99.1%). Accuracy "jumped" to 81% at 1000:1 and to 87% at 10,000:1. For answers assigned odds of 1,000,000:1 or greater, accuracy was 90%; the appropriate degree of confidence would have been odds of 9:1. The 12% of responses that are not listed in Table 3 because they fell between the most common odds categories showed a similar pattern of overconfidence. In summary, subjects were frequently wrong at even the highest odds levels. Moreover, they gave many extreme odds responses. More than half of their judgments were greater than 50:1. Almost one-fourth were greater than 100:1.

A second experiment attempted to improve performance by giving subjects more instruction. The experimental session began with a 20-

Table 4. Percentage of correct answers for major odds categories

Odds	Lethal events						General-knowledge questions					
	Experiment 1 ^a			Experiment 2 ^b			Experiment 1 ^a			Experiment 2 ^b		
	N	%N	% correct	N	%N	% correct	N	%N	% correct	N	%N	% correct
1:1	644	9	53	339	8	54	861	19	53	210	5	56
1.5:1	68	1	57	108	2.5	59	455	1	63	455	1	63
2:1	575	8	64	434	10	65	157	3.5	76	194	4	76
3:1	189	2	71	252	6	65	376	8	74	66	1.5	85
5:1	250	4	70	322	8	71	227	5	74	69	1.5	83
10:1	1,167	17	66	390	9	76	319	8	87	376	8	80
20:1	126	2	72	163	4	81	219	5	84	334	7	88
50:1	258	4	68	227	5	74	263	6	89	134	3	92
100:1	1,180	17	73	319	8	87	47	1	96	360	8	94
1,000:1	862	13	81	219	5	84	6,098	88	71.0	2,981	70	72.5
10,000:1	459	7	87	138	3	92	3,855	75	73.1	73.1		
100,000:1	163	2	85	23	.5	96						
1,000,000:1	157	2	90	47	1	96						
Total												
Overall % correct												

Note: % N refers to the percentage of odds judgments that fell in each of the major categories. There were 66 subjects in Experiment 1, 40 in Experiment 2, and 42 in Experiment 3.

^aFor well-calibrated subjects.

^bExperiments 1, 2, and 3 were labeled Experiments 2, 3, and 4 in the original report.

Source: Fischhoff, Slovic, and Lichtenstein (1977).

minute lecture in which the concepts of probability and odds were carefully explained. The subtleties of expressing one's feelings of uncertainty as judgments of numerical odds were discussed, with special emphasis on how to use small odds (between 1:1 and 2:1) when one is quite uncertain about the correct answer. A chart was provided showing the relationship between various odds and the corresponding probabilities. Finally, subjects were taught the concept of calibration (Chap. 22) and were urged to make odds judgments in a way that would lead them to be well calibrated. Although performance improved somewhat, subjects again exhibited unwarranted certainty (see Table 4). They assigned odds greater than or equal to 50:1 to approximately one-third of the items. Only 83% of the answers associated with these odds were correct.

In a third experiment, people proved to be just as overconfident when answering questions of general knowledge (e.g., Which magazine had the largest circulation in 1970? (a) *Playboy* or (b) *Time*) as when they answered questions about the frequency of lethal events (see Table 4). Additional studies by Fischhoff et al. tested people's faith in their odds assessments

by asking if they would stake money on them by playing the bet described below.

Instructions for "Trivia Question Hustling"

The experiment is over. You have just earned \$2.50, which you will be able to collect soon. But before you take the money and leave, I'd like you to consider whether you would be willing to play a certain game in order to possibly increase your earnings. The rules of the game are as follows:

1. Look at your answer sheet. Find the questions where you estimated the odds of your being correct as 50:1 or greater than 50:1. How many such questions were there? ——— (write number).
2. I'll give you the correct answers to these "50:1 or greater" questions. We'll count how many times your answers to these questions were wrong. Since a wrong answer in the face of such high certainty would be surprising, we'll call these wrong answers "your surprises."
3. I have a bag of poker chips in front of me. There are 100 white chips and 2 red chips in the bag. If I reach in and randomly select a chip, the odds that I will select a white chip are 100:2 or 50:1, just like the odds that your "50:1" answers are correct.
4. For every "50:1 or greater" answer you gave, I'll draw a chip out of the bag. (If you wish, you can draw the chips for me.) I'll put the chip back in the bag before I draw again, so the odds won't change. The probability of my drawing a red chip is 1/51. Since drawing a red chip is unlikely, every red chip I draw can be considered "my surprise."
5. Every time you are surprised by a wrong answer to a "50:1 or greater" question, you pay me \$1 (raised to \$2.50 in some conditions). Every time I am surprised by drawing a red chip, I'll pay you \$1.
6. If you are well calibrated, this game is advantageous to you. This is because I expect to lose \$1 about once out of every 51 times I draw a chip, on the average. But since your odds are sometimes higher than 50:1, you expect to lose less often than that.
7. Would you play this game?

This bet is advantageous for perfectly calibrated and underconfident participants and disadvantageous to overconfident ones. Most participants in our study were eager to play the game. Because their confidence was unjustified, they suffered sizable monetary losses (which we returned to them after the experiment was over).

Although the psychological basis for unwarranted certainty is complex, a key element seems to be people's lack of awareness that their knowledge is based on assumptions that are often quite tenuous. For example, 30% of the respondents in Experiment 1 gave odds greater than 50:1 to the incorrect assertion that homicides are more frequent than suicides. These individuals may have been misled by the greater ease of recalling

Table 5. Experts' insensitivity to omissions from the car-won't-start fault tree

Group	M proportion of starting failures by type							
	n	Battery	Starting system	Fuel system	Ignition system	Engine	Mischief	Other
Unpruned tree, ordinary subjects	93	.264	.195	.193	.144	.076	.051	.078
Unpruned tree, experts	13	.410	.108	.096	.248	.051	.025	.060
Pruned tree 1, experts	16	.483	—	.229	—	.073	—	.215 ^a

^aShould be .441.

Source: Fischhoff, Slovic, and Lichtenstein (1978).

instances of homicide, failing to appreciate that memorability is an imperfect basis for such an inference.

Hyperprecision. Overconfidence manifests itself in other ways as well. A typical task in estimating uncertain quantities such as failure rates is to set upper and lower bounds so that there is a 98% chance that the true value lies between them. Experiments with diverse groups of people making many different kinds of judgments have shown that, rather than 2% of true values falling outside the 98% confidence bounds, 20–50% do so (Chaps. 21 and 22). Thus people think that they can estimate such values with much greater precision than is actually the case. Tversky and Kahneman (1974, 1) have attributed such hyperprecision to reliance on the anchoring and adjustment heuristic.

Overconfident experts. Unfortunately, experts, once they are forced to go beyond their data and rely on judgment, may be as prone to overconfidence as laypeople. Fischhoff, Slovic, and Lichtenstein (1978) repeated their fault-tree study (Figure 3) with professional automobile mechanics (averaging about 15 years of experience) and found these experts to be almost as insensitive as laypersons to deletions from the tree (see Table 5). Hynes and Vanmarcke (1976) asked seven "internationally known" geotechnical engineers to predict the height of an embankment that would cause a clay foundation to fail and to specify confidence bounds around this estimate that were wide enough to have a 50% chance of enclosing the true failure height. None of the bounds specified by these individuals actually did enclose the true failure height. Figure 4 shows these results.

The multimillion dollar Reactor Safety Study (U.S. Nuclear Regulatory Commission, 1975), in assessing the probability of a core melt in a nuclear reactor, used the very procedure for setting confidence bounds that was




Figure 4. An example of overconfidence in expert judgment, as represented by the failure of error bars to contain the true value. The data represent estimates by seven "internationally known" geotechnical engineers of the height at which an embankment would fail. (Source: Hynes & Vanmarcke, 1976.)

shown in Chapters 21 and 22 to produce a high degree of overconfidence. In fact, the "Lewis Committee" concluded its review of the Reactor Safety Study by noting that despite the great advances made in that study "we are certain that the error bands are understated. We cannot say by how much. Reasons for this include an inadequate data base, a poor statistical

treatment, [and] an inconsistent propagation of uncertainties throughout the calculation" (U.S. Nuclear Regulatory Commission, 1978, p. vi).

Further anecdotal evidence of overconfidence may be found in many other technical risk assessments (Fischhoff, 1977a). Some common ways in which experts may overlook or misjudge pathways to disaster are shown in the list below.

Failure to consider the ways in which human errors can affect technological systems. Example: Because of inadequate training and control room design, operators at Three Mile Island repeatedly misdiagnosed the problems of the reactor and took inappropriate actions (Sheridan, 1980; President's Commission, 1979).

Overconfidence in current scientific knowledge. Example: Use of DDT came into widespread and uncontrolled use before scientists had even considered the possibility of the side effects that today make it look like a mixed and irreversible blessing (Dunlap, 1978).

Failure to appreciate how technological systems function as a whole. Example: The DC-10 failed in several early flights because its designers had not realized that decompression of the cargo compartment would destroy vital control systems (Hohenemser, 1975).

Slowness in detecting chronic, cumulative effects. Example: Although accidents to coal miners have long been recognized as one cost of operating fossil-fueled plants, the effects of acid rains on ecosystems were slow to be discovered.

Failure to anticipate human response to safety measures. Example: The partial protection afforded by dams and levees gives people a false sense of security and promotes development of the flood plain. Thus, although floods are rarer, damage per flood is so much greater that the average yearly dollar loss is larger than before the dams were built (Burton, Kates, & White, 1978).

Failure to anticipate "common-mode failures," which simultaneously afflict systems that are designed to be independent. Example: Because electrical cables controlling the multiple safety systems of the reactor at Browns Ferry, Alabama, were not spatially separated, all five emergency core cooling systems were damaged by a single fire (U.S. House of Representatives, 1975; Jennergren & Keeney, in press).

The 1976 collapse of the Teton Dam provides another tragic example of expert overconfidence. The Committee on Government Operations attributed this disaster to the unwarranted confidence of engineers who were absolutely certain they had solved the many serious problems that arose during construction (Committee on Government Operations, 1976). Fail-

ure probabilities are typically not even calculated for new dams even though about 1 in 300 fails when the reservoir is first filled.

Informing people about risks

Thinking clearly about risk is difficult. Unfortunately, it is also necessary. Radiation hazards, medical side effects, occupational diseases, food contaminants, toxic chemicals, and mechanical malfunctions increasingly fill our newspapers and our thoughts. Since the management of these hazards is vital to the well-being of individuals and society, people are presently asserting their right to play an active role in the decision-making process. As a result, the promoters and regulators of hazardous enterprises face growing pressure to inform people about the risks they face (see Figure 5). For example, in recent years:

The Food and Drug Administration mandated patient information inserts for an increased number of prescription drugs.

The Department of Housing and Urban Development began to require the sellers of homes built before 1950 to inform buyers about the presence of lead-based paints.

The proposed federal products liability law placed increased weight on adequately informing consumers and workers about risks they are likely to encounter.

The White House directed the Secretary of Health, Education, and Welfare to develop a public information program on the health effects of radiation exposure.

Despite these good intentions, creating effective informational programs may be quite difficult. Doing an adequate job means finding cogent ways of presenting complex, technical material that is often clouded by uncertainty. Not only is the allotted time sometimes very limited, but messages must confront the listeners' preconceptions (and perhaps misconceptions) about the hazard in question and its consequences. For example, in some situations, misleading personal experiences may promote a false sense of security, whereas in other circumstances, mere discussion of possible adverse consequences may enhance their apparent threat. Moreover, as Ross and Anderson (Chap. 9) have demonstrated, people's beliefs often change slowly and show extraordinary persistence in the face of contrary evidence. What follows is a brief overview of some additional challenges that information programs must confront.

Presentation format is important

The precise manner in which risks are expressed can have a major impact on perceptions and behavior. For example, an action increasing one's



Figure 5. Drawing by S. Harris, © 1979 *The New Yorker Magazine*.

annual chances of death from 1 in 10,000 to 1.3 in 10,000 would probably be seen as much more risky if it were described, instead, as producing a 30% increase in annual mortality risk. A sampling of format effects from the literature is presented below.

Fault trees. The designers of a fault tree like that in Figure 3 must make numerous discretionary decisions regarding how to organize and present the various sources of trouble. One such decision that apparently makes little difference is how much detail to offer; Fischhoff, Slovic, and Lichtenstein (1978) found similar perceptions with varying levels of detail. Merely mentioning a branch allowed people to estimate accurately how troublesome that branch would look when fully detailed. However, fusing branches (e.g., combining starting system and ignition system into one broader category) or splitting branches (e.g., separating ignition system into ignition system [coil faulty and spark plugs defective, see Figure 3]

and distribution system [distributor faulty and wiring defective] did make a difference. A given set of problems was judged to account for about 30% more failures when it was presented as two branches than when it was presented as one.

Seat belts. A second demonstration of the importance of presentation format comes from a study of attitudes toward the use of automobile seat belts (Slovic, Fischhoff, & Lichtenstein, 1978). Drawing upon previous research demonstrating the critical importance of probability of harm in triggering protective action (Slovic, Fischhoff, Lichtenstein, Corrigan, & Combs, 1977), Slovic, Fischhoff, and Lichtenstein argued that people's reluctance to wear seat belts voluntarily might be due to the extremely small probability of incurring a fatal accident on a single automobile trip. Since a fatal accident occurs only about once in every 3.5 million person trips and a disabling injury only once in every 100,000 person trips, refusing to buckle one's seat belt may seem quite reasonable. It looks less reasonable, however, if one adopts a multiple-trip perspective and considers the substantial probability of an accident on some trip. Over 50 years of driving (about 40,000 trips), the probability of being killed rises to .01 and the probability of experiencing at least one disabling injury is .33. In a pilot study, Slovic, Fischhoff, and Lichtenstein showed that people asked to consider this lifetime perspective responded more favorably toward seat belts (and air bags) than did people asked to consider a trip-by-trip perspective. Whether the favorable attitudes toward seat belts induced by a lengthened time perspective would be maintained and translated into behavior remains to be seen.

Pseudocertainty. According to "prospect theory" (Kahneman & Tversky, 1979b), outcomes that are merely probable are underweighted in comparison with outcomes that are obtained with certainty. As a result, any protective action that reduces the probability of harm from, say, .01 to zero, will be valued more highly than an action reducing the probability of the same harm from .02 to .01.

Tversky and Kahneman (1981) note that mental representations of protective actions may be easily manipulated so as to vary the apparent certainty with which they prevent harm. For example, an insurance policy that covers fire but not flood could be presented either as full protection against the specific risk of fire or as a reduction in the overall probability of property loss. Prospect theory predicts that the policy will appear more attractive in the former perspective (labeled "pseudocertainty"), in which it offers unconditional protection against a restricted set of problems.

We have tested this conjecture in the context of one particular kind of protection, vaccination. Two forms of a "vaccination questionnaire" were created. Form I (probabilistic protection) described a disease expected to

afflict 20% of the population and asked people whether they would volunteer to receive a vaccine that protects half of the people receiving it. According to Form II (pseudocertainty), there were two mutually exclusive and equiprobable strains of the disease, each likely to afflict 10% of the population; the vaccination was said to give complete protection against one strain and no protection against the other. The participants in this study were recruited by an advertisement in the University of Oregon student newspaper. Half received Form I; the other half received Form II. After reading the description, they rated the likelihood that they would get vaccinated in such a situation, using a scale ranging from 1 ("almost certainly would not get vaccinated") to 7 ("almost certainly would get vaccinated").

Although both forms indicated that vaccination reduced one's overall risk from 20% to 10%, we expected that vaccination would appear more attractive to those who received Form II (pseudocertainty) than to those who received Form I (probabilistic protection). The results confirmed this prediction: 57% of those who received Form II indicated they would get vaccinated compared with 40% of those who received Form I.

The pseudocertainty effect highlights the contrast between the reduction and the elimination of risk. As Tversky and Kahneman have indicated, this distinction is difficult to justify on any normative grounds. Moreover, manipulations of certainty would seem to have important implications for the design and description of other forms of protection (e.g., medical treatments, insurance, flood- and earthquake-proofing activities).

Anchoring. One of the most general of presentation artifacts is the tendency of judgments to be anchored on initially presented values (Poulton, 1968; Tversky & Kahneman, 1974, 1). In another condition of the experiment presented in Figure 2, Lichtenstein et al. (1978) asked a second group of people to estimate the frequency of death in the United States from each of the 40 different causes. However, instead of being told that about 50,000 people die annually in motor vehicle accidents, these individuals were told about the 1,000 annual deaths from electrocution. Although both reports were accurate, provision of a smaller number reduced respondents' estimates of most frequencies. Such anchoring on the original number led the estimates of the two groups to differ by as much as a factor of 5 in some cases.

Fischhoff and MacGregor (1980) asked people to judge the lethality of various potential causes of death using one of four formally equivalent formats (e.g., For each afflicted person who dies, how many survive? For each 100,000 people afflicted, how many will die?). Table 6 expresses their judgments in a common format and reveals even more dramatic effects of question phrasing on expressed risk perceptions. For example, when

