

Overreaction to Fearsome Risks

Cass R. Sunstein and Richard Zeckhauser
Harvard University

“...the only thing we have to fear is fear itself”
President Franklin D. Roosevelt, First Inaugural Address, March 4, 1933

1. The Problem and the Thesis

When risks threaten, some cognitive mechanisms push people toward action; others push them toward inaction. The availability heuristic can do either. When relevant events are cognitively available, people will be inclined to act; availability bias might lead them to value the risk excessively and thereby take excessive precautions. But if relevant events are not available, the unavailability bias will predominate. The risk will be slighted, and action will be inhibited.

Most of the literature on individual and social responses to risks stresses the inaction case, and focuses on the need to develop mechanisms to ensure that serious dangers receive sufficient response. President Roosevelt’s concerns about fear, expressed in his very next words, were: “nameless, unreasoning, unjustified terror which paralyzes needed efforts to convert retreat into advance.” Roosevelt was right to see that this “paralysis” phenomenon is important. But because cognitive biases also push in the opposite direction, producing overreactions or what we might think of as panic, mechanisms are also needed to dampen public demands and policy responses. Panic and passivity are the opposite ends of a spectrum. Each is unwelcome.

This essay focuses on fearsome risks -- those that stimulate strong emotional responses, such as fear and anxiety.¹ Such risks, which usually involve high consequences, tend to have extremely low probabilities, since life today is no longer nasty, brutish and short. We aim to show here that in the face of a fearsome risk, people often exaggerate the benefits of preventive, risk-reducing, or ameliorative measures. In both personal life and politics, the result is damaging overreactions to risks. In this essay, we focus on a form of decision bias that goes by the name of probability neglect.

Probabilities perplex people. Puny probabilities prove particularly perplexing. When risk probabilities are extremely low, as they are for most high-consequence negative events, those who fall victim to probability neglect will give up too much to avoid the risk. Frequently,

¹ The “risk as feelings” hypothesis highlights the “role of affect experienced at the moment of decision making,” (Loewenstein et al., 2001, p. 267), as opposed to a cognitive assessment of a risk.

they will take excessive preventive action.² This will be true of governments and corporations as well as individuals, in part because they respond to individuals, in part because of their own tendencies.

One salient manifestation of probability neglect is that in two situations involving the same fearsome risk, one much more likely than the other, individuals may value risk elimination little differently even though probabilities may differ by a factor of 20 or more. The reason is that they focus on the bad outcome itself, and they are inattentive to the question of how likely it is to occur – hence their overreaction when the risk is low. Such overreactions in general can be categorized as “action bias.” That bias is especially likely if the relevant actors will be able to obtain credit for responding to the risk. That credit may come from themselves or from the public more generally (Patt and Zeckhauser, 2000).

An understanding of action bias in the context of low-probability risks, and its common ingredient probability neglect, has important implications for both law and policy. It is predictable that in the aftermath of a terrorist attack, the public will both alter its behavior and demand a substantial governmental response. That will be true even if the magnitude of the risk does not warrant such a response, and even if the danger is far less than that presented by other hazards that do not greatly concern people, perhaps because they do not get much public attention. Consider, for example, the possibility that extensive security precautions at airports will lead people to drive rather than to fly; because flying is much safer than driving, such precautions might sacrifice many lives on balance. The monies spent in recent years on airplane security might be out of scale with the level of risk reduction produced, particularly since numerous tests have found that the screening routinely fails to find weapons.

Perhaps such screening, however low the risk or ineffective the preventive, does fulfill the function of reassuring the public. If so, it serves a positive function, not unlike the nighttime hoof claps of mounted police. Squad cars may be better at deterring or catching criminals, but do much less to reduce public anxiety. Hoof claps are splendid for fear reduction. The same points apply of course to many other purported forms of risk reduction, including measures to prevent financial crises, local steps to reduce greenhouse gas emissions, and regulation of abandoned hazardous waste dumps. Financial crises have a distinctive element: fear-related emotions themselves may stimulate the crisis, as in 2008, making reassurance that much more critical.

In the personal as opposed to social domain, we can find many analogues, as when people alter their travel plans to avoid slight risks of crime, restructure their portfolios to avoid small risks of big financial losses, or change their diet to avoid minimal health risks. In all of these cases, a form of action bias, fueled by probability neglect, may lead to overreactions. The costs of overreaction may be financial (restructuring the portfolio), increased risk (driving rather than flying), or sacrificed pleasure (children foregoing Halloween due to extremely rare razor-blade incidents).

² The converse is surely true. If emotions lead to the neglect of probabilities, there will be insufficient response to high probability risks.

2. Overreaction to Environmental Risks

Overreaction to risk is frequently found in the environmental realm. A dramatic example is provided by the Three Mile Island accident in 1979. Significant amounts of reactor coolant were leaked to the environment, including releases of cancer-inducing agents. The Kemeny Commission Report, created under Presidential order, concluded that in expectation less than one case of cancer would be created. Yet the accident affected public and political opinion sufficiently to terminate the construction of any new nuclear plants in the United States for 30 years. The coal- and oil-fueled plants built in their stead surely caused many more health problems, looking only at the air pollution they produced. (Today, nuclear power is poised to make a comeback given concerns about global warming due significantly to CO₂ emissions from conventional power plants.) The impact of Three Mile Island was reinforced due to the release of the movie The China Syndrome the same month as the accident. The movie made a catastrophic accident – narrowly avoided due to courageous action of the movie’s heroes -- “available” in the public mind. When an example of an event can readily be brought to mind, it is judged to be much more likely (Tversky and Kahneman, 1973 and 1974).

Probabilities are often hard to calculate in environmental contexts. Risks can emanate from thousands of sources and affect millions of people. Ideal policy, say in setting a standard, would look at the consequences from a slight change in a risk-producing activity. However, epidemiological evidence will rarely be conclusive in establishing such consequences, for multiple reasons: (1) most risks have their impacts over long time periods, (2) if large numbers of people are affected, all very slightly, changes in risk are hard to perceive, (3) other factors are likely changing at the same time when the risk under study changes, (4) controlled experiments imposing risks are unethical, and (5) the risks that we are concerned about have very low probabilities, requiring extremely large samples to detect changes in their levels.

Given the virtual impossibility of deciphering risks levels from mere observation, we often extrapolate risks to exposed human populations from quite different contexts. Thus, to judge the danger of a carcinogen, we are likely to expose an animal population – mice being a favorite subject pool – to extremely high concentrations. Such concentrations make it likely that effects, if there are such, will show up in a small pool of subjects, say 500 mice. The challenge to this approach is that we then must extrapolate from high to low doses and from animals to man. Such extrapolations require heroic assumptions, and are subject to fierce debate. A risk that is judged to be 1/1,000,000 by one researcher might be judged 100 times as high by another.

Overreaction to fearsome risks in the environment can come about for many reasons: (1) Environmental risks are usually imposed unwillingly on external parties. Indignation over such an imposition stirs excess reaction. This proclivity is reinforced in the policy realm, where the party demanding action (the externality recipient) is not the party paying for it. Hence, paying much more to avoid a risk than the benefits secured is not troubling. (2) We misperceive the risk, as happened after Three Mile Island. (3) Though we calculate the risk, as opposed to merely grasping it through some gut process, through a series of conservative assumptions we end up with a probability estimate that is far above what an expected value calculation would yield. Scientific practice, say when proving a theory, is to employ conservative assumptions. However, rational decision theory requires that one act on the expected value of a probability.

(4) Though we may have an accurate estimate of the risk, the emotions stirred by its fearsome aspect lead us to neglect probability values. This essay is devoted to the fourth reason. It focuses on low-probability risks. Therefore, not considering their magnitude leads to overreaction.

In general, we overreact to some risks and virtually ignore others. “Often too much weight is placed on risks of low-probability risks but high salience (such as trace carcinogens)... Too little effort is spent ameliorating voluntary risks, such as those involving automobiles and diet.” (Zeckhauser and Viscusi, 1990, p. 559) Many environmental risks would seem to fall in the overweighted category.

3. Demonstrating probability neglect

Prospect theory (Kahneman and Tversky, 1979) tells us that the perceived benefits of risk elimination will be much less than proportional to the risk avoided, since the probability weighting function takes a downward leap at 0. However, prospect theory alone gives no indication that the ratio of valuations would change dramatically with the nature of a risk, or with the way it was described.³

Experiments on probability neglect seek to assess whether attention to probability could be overshadowed by attention to the affective goodness or badness of the outcome, quite contrary to what leading theories of decision making posit. To make the same point in broader and more metaphorical terms, emotional activity dampens cognitive activity. Loewenstein and Lerner (2003) observe that: “As the intensity of immediate emotions intensifies, they progressively take control of decision making and override rational decision making.” We would expand this assertion to include overriding well-documented behavioral patterns in decision making, such as those described by prospect theory. If such overshadowing or “taking control” is found, then a dire risk could swamp or at least temper the importance of dramatic probability differences.

Some of the relevant experiments explore whether varying the probability of harm would matter less in settings that trigger strong emotions than in those that are relatively emotion-free. One such study explored people’s willingness to pay to avoid electric shocks, in an effort to test the relevance of variations in probability to “affect rich” decisions (Rottenstreich and Hsee, 2001). In the “strong emotion” setting, participants were asked to imagine that they would participate in an experiment involving some chance of a “short, painful, but not dangerous electric shock.” In the relatively emotion-free setting, they were told that the experiment entailed some chance of a \$20 penalty. Participants were asked to say how much they would be willing to pay to avoid participating in the relevant experiment. Some participants were told that there was

³ The Russian roulette example (p. 283) described to illustrate the probability weighting function that plays a central role in the theory does inject emotion into a decision. The subject is told she is compelled to play Russian roulette, but can pay to remove one bullet from the loaded gun. She is asked whether she would pay more to reduce the number of bullets from 4 to 3 than from 1 to 0. Most people pay substantially more for the reduction in the probability of death from 1/6 to 0 than for a reduction from 4/6 to 3/6. This reverses rationality if money is worth less if one is dead. – Surely the probability weighting function plays a significant role in explaining the payment disparity, but the emotion of the life-and-death decision may magnify it.

a 1% chance of receiving the bad outcome (either the \$20 loss or the electric shock); others were told that the chance was 99%.

The central result was that variations in probability affected those facing the relatively emotion-free injury, the \$20 penalty, far more than they affected people facing the more emotionally evocative outcome of an electric shock. For the cash penalty, the difference between the median payment for a 1% chance and the median payment for a 99% chance was predictably large. The median subject paid \$1 to avoid a 1% chance, and \$18 to avoid a 99% chance. For the electric shock, by contrast, the difference in probability made little difference to median willingness to pay: \$7 to avoid a 1% chance, and \$10 to avoid a 99% chance -- only 1½ as much in the affect-rich setting as opposed to 18 times as much when little emotion was involved. (Of course, a fully rational response would produce a ratio much greater than 18, since income effects are likely trivial for such small gambles.)

The conclusion is that when a hazard stirs strong emotions, most people will pay an amount to avoid it that varies little even with extreme differences in the starting probability. What we are stressing here is that when the probability of loss is very low, people will show action bias. They will favor precautionary steps even if those steps are not justified by any plausible analysis of expected utility.

For either social or personal risks, the implication is clear. When the potential loss is likely to trigger strong emotions, action bias threatens, as it does when the loss is an economic meltdown, environmental catastrophe, terrorist attack, contracting cancer, or getting killed in a plane crash. Even if the likelihood of a terrible outcome were extremely low, people would be willing to pay a great deal to avoid it, whether through public or private action. Once a risk is in people's minds, their willingness to pay to avoid it will often be relatively impervious to significant changes in probability. The significant and often expensive precautions taken against possible sniper attacks by citizens of the Washington, DC area in October 2002 provide a dramatic example; they attest to the phenomenon of probability neglect in the face of a vivid threat. Indeed, some of these precautions, such as driving great distances to a gas station in Virginia, almost certainly increased mortality risks on balance.⁴

Probability neglect and hence action bias can be found for willingness to reduce, and not merely to eliminate, a risk. To be sure, prospect theory shows that people will pay a special premium to eliminate a risk, but where the risk is emotionally gripping, risk reduction will not be sensitive to the question of probability. To investigate the possibility of value inflation in response to risks, we asked a large number of law students to state their maximum willingness to pay to reduce levels of arsenic in drinking water.

The appropriate standard for arsenic in drinking water drew substantial public attention a few years back. Three days before the end of the Clinton Administration, it proposed tightening the standard from 50 parts per billion to 10 parts per billion. The new Bush Administration immediately suspended the standard. However, in response to a public outcry it reversed its

⁴ When the risk is imposed by malicious people, there is often a negative externality from the precautions taken by any individual. Those who went to Virginia to fill made it more dangerous for D.C. fillers. When few citizens walk in an urban area at night, those who still walk find such activity more dangerous.

decision, and the 10 parts per billion standard is now the law. In the academic community, there is considerable controversy on two issues relating to the standard. First, what does this tighter standard cost per life saved? The amount could be high since the risk is low. Second, the costs of meeting the standard per capita vary dramatically from one jurisdiction to another. One estimate was that the cost per household would be \$0.86/year in districts with more than one million households, but \$163/year for districts with between 101 and 500 households.⁵ Cost-benefit analysis would suggest the standard should be much stricter in large as opposed to small districts.

Our goal is not to determine the appropriate standard, but to illustrate the real-life relevance of our questions. They were based on the actual choice confronting the Environmental Protection Agency, and employ risk numbers within the ballpark of actual figures used by the agency itself.

Participants were randomly sorted into four groups, representing the four conditions in a 2x2 experiment, where both the probability and the description of the risk varied. In the first condition, people were asked to state their maximum willingness to pay to eliminate a cancer risk of one in 1,000,000. In the second condition, people were asked to state their maximum willingness to pay to eliminate a cancer risk of one in 100,000. In the third condition, people were asked the first question, but the cancer was described in vivid terms, as “very gruesome and intensely painful, as the cancer eats away at the internal organs of the body.” In the fourth condition, people were asked the second question, but the cancer was described in the same vivid terms as in the third condition. In each condition, participants were asked to check off their willingness to pay among the following options: \$0, \$25, \$50, \$100, \$200, \$400, and \$800 or more. Notice that the description of the cancer in the “highly emotional” conditions added little information, simply describing many cancer deaths, though admittedly some participants might well have learned that these were especially horrific deaths.⁶

The first hypothesis, consistent with the probability weighting function of prospect theory, was that the ten-fold difference in probabilities – between 1/100,000 and 1/1,000,000 – would generate a much less than a ten-fold difference in willingness to pay. The second hypothesis was that the probability variations would matter less in the highly emotional conditions than in the less emotional conditions. More specifically, it was predicted that the highly emotional conditions would overshadow differences in probability, whereas such differences would have greater importance in the less emotional condition.

Here are the results in tabular form:

⁵ Oates (2002), p. 17.

⁶ Paul Slovic stressed to us that individuals cannot experience fear over a sustained period of time. However, we believe that our questions spotlight a risk, and do have the potential to stir severe emotions as individuals respond.

Table 1. Willingness to Pay in Dollars for Elimination of Arsenic Risks

Harvard Law School Results, 2008
Mean (Median)
[Number of Subjects]

Probability	Unemotional description	Emotional description
1/100,000	241.25 (100) [20]	250 (100) [13]
1/1,000,000	59.21 (25) [19]	211.67 (200) [15]

The study was conducted in two law school venues, University of Chicago (Sunstein, 2002) and at Harvard Law School. At Chicago, the medians were 25 and 100 for the unemotional description, and 100 and 100 for the emotional description. While the sample size was too small to permit firm conclusions, the qualitative results pointed in the hypothesized direction. The emotional description drove out responses to the quantitative difference in the risk.

At Harvard, as shown and hypothesized, the valuations for the emotional description hardly differed even though risks differed by a factor of 10. There was substantial difference in willingness to pay (WTP) for the unemotional description. A Mann-Whitney Wilcoxon Rank Sum test showed that the WTP was actually higher for the 1/1,000,000 risk given the emotional description, though far from significant. By contrast, and as expected, the difference for the emotional description was highly significant, with the 1/100,000 payment higher ($z = 3.398$, $p < 0.001$). Comparing the two results showed that the unemotional description gave a greater differential that was highly statistically significant.

It is important to note that the difference in WTP, even for the unemotional description, was far below the 10 to 1 odds ratio; for means it was roughly 4 to 1. Both hypotheses were therefore supported. First, varying the probability had an effect on WTP that was much less than rational decision theory would predict. (Future research should assess whether even mentioning the word “cancer” induced sufficient emotion to reduce a 10 to 1 ratio to 4 to 1.) Second, the effect of increasing the probability by a factor of ten had an effect that was highly significant in the unemotional condition -- but was completely insignificant in the emotional condition. When the cancer was described in affectively gripping terms, people were insensitive to probability variations.

These findings have two implications for overreactions. They suggest, first, that when extremely low probability risks give rise to intense fear, they are likely to trigger a larger behavioral response than do statistically identical comparisons involving less fearsome risks. Here, as in the experiment, there will be a kind of “emotion premium.” The findings suggest,

second, that probability neglect will play a role in the private and public reaction to emotionally gripping risks, and that many people will focus, much of the time, on the emotionally perceived severity of the outcome, rather than on its likelihood. In this light, it should not be surprising that our public figures and our cause advocates often describe tragic outcomes. Rarely do we here them quote probabilities. The latter, even if reasonably large, would have little salience in the public debate.

Think of an anti-drunk driving campaign that said: “When you drink don’t drive; you have one chance in 1,000,000 of getting into a fatal car crash, a much higher probability than most people believe.”⁷ The alternative campaign, catering to probability neglect, would skip probabilities altogether. It might show a car wrapped around a tree, and a grieving family standing nearby, with the simple statement: “When you drink don’t drive.” The tree-wrap, we predict, would be more effective.

Emotions not connected to fear may also drive probability neglect. Consider outrage, an emotion sometimes stirred when low probability risks are created from the outside, as they are with nuclear waste radiation. A similar risk from radon exposure comes from one’s own basement, hence no outrage. Outrage can overshadow probabilities in much the same way as a vivid risk can, reinforcing our metaphor about emotional activity dampening cognitive activity. A central finding of relevant empirical work is consistent with that stressed here: a large difference in probability had no effect on people’s judgments in a “high outrage” condition, involving nuclear waste, but a significant effect in a “low outrage” condition, involving radon. For nuclear waste, people responded the same way to a risk of 1 in 100,000 as to a risk of 1 in 1,000,000 (Sandman et al., 1998). Even when both the statistical risk and ultimate consequences were identical in the high outrage (nuclear waste) and low outrage (radon) cases, people in the nuclear waste case reported a much greater perceived threat and a much higher intention to act to reduce that threat (id.). Indeed, “the effect of outrage was practically as large as the effect of a 4000-fold difference in risk between the high-risk and low-risk conditions” (id.).⁸

In this light, it is not surprising that visualization or imagery matters a great deal to people’s reactions to risks. Vivid images can produce palpable overreactions (Slovic et al., 2000). When an image of a bad outcome is easily accessible, people will become greatly concerned about a risk, holding probability constant (Loewenstein et al., 2001). An interesting anomaly is that when people are asked how much they will pay for flight insurance for losses resulting from “terrorism,” they will pay more than if they are asked how much they will pay for flight insurance from all causes (Johnson et al., 1993).⁹ The likely explanation for this peculiar

⁷ This number was roughly calculated assuming that there are 100 million drivers, each driving after drinking 15 days a year, and accounting for half the annual fatalities in the United States.

⁸ An alternative explanation is that individuals demand substantial compensation for their outrage, and that such compensation is both fairly independent of the probability and large relative to the compensation for risk. Note that efforts to communicate the meaning of differences in risk levels, by showing comparisons to normal risk levels, reduced the effect of outrage; but only modestly so. Outrage had nearly the same effect as a 2000-fold increase in risk (id.). Did this information provision improve cognitive uptake directly, or indirectly because it dampened the outrage? Further experiments will be required to tell.

⁹ This should bring to mind a component of the embeddedness phenomenon known for contingent valuation. If seals are appealing and easily visualized, it is not surprising that we might pay more to save them in an oil spill than to save all wildlife.

result is that the word “terrorism” evokes vivid images of disaster, outrage, or both, thus inhibiting judgments about probability differences. Note also that when people discuss a low-probability risk, their concern rises even if the discussion consists mostly of apparently trustworthy assurances that the likelihood of harm really is infinitesimal (Alkhami and Slovic, 1994). The discussion helps people to visualize the risk, thus making it more frightening. The most sensible conclusion is that with respect to risks of injury or harm, vivid images and concrete pictures of disaster can “crowd out” the cognitive activity required to conclude and consider the fact that the probability of disaster is really small.

3. Probability neglect and anxiety

We also sought to test the relationship between probability neglect and the emotion of anxiety. In order to do so, we followed Rottenstreich and Hsee (2001), who used a painful but not dangerous electric shock to produce emotion and presumably anxiety. In a subsequent class, the authors asked the same class of Harvard Law School students how much they would demand to accept a shock. If their demand price was below the experimenter’s payment price, the experiment would go ahead at the experimenter’s price, thus assuring incentive compatibility. This too was a 2x2 setup, with the probability and the timing of the shock each taking two values. The shock was either received for certain (100% probability) or with 1% probability. The shock was to be delivered immediately after class, or immediately after a class one year from the experiment. Unlike the arsenic study, the categories of payment were not predefined. Not surprisingly, the mean values substantially exceeded the medians due to a few severe outliers.

It was conjectured that contemplating receiving the shock in a year would raise anxiety beyond that of waiting to receive a potential shock at the end of class. Such anxiety, it was thought, might enhance probability neglect. It could also raise the demand price.

There were three hypotheses to be tested. They would be tested using rank order, non-parametric methods.

1. Subjects would show probability neglect. The differences in demand prices in the certainty condition and 1% chance condition would be far less than 100 to 1.
2. The contemplation period for receiving a shock in a year would increase the total costs of anxiety. Because people would want to reduce their own anxiety, the demand price would be higher for the shock to be delivered a year from today.

If hypothesis 2 was confirmed, a third hypothesis would be relevant.

3. Given that probability neglect is greater in emotionally gripping cases, the disparity between the certainty price and the 1% price would be more compressed for the shock to be received in a year (thus multiplying anxiety) than from the shock to be received today.

The results were as follows:

Table 2. Demand Price to Accept a Painful But Non Dangerous Electric Shock

Harvard Law School Results, 2008
Mean (Median)
[Number of Subjects]

Probability	Shock Today	Shock in a Year (Anxiety)
100%	1283.33 (50) [12]	1966.43 (100) [14]
1%	661.41 (50) [23]	824.05 (50) [21]

The first hypothesis, most important to our central claim here, was confirmed. There was nothing even close to a statistically significant difference between the certainty price and 1% conditions in both the Shock Today and Shock in a Year cases. Despite a 100 times greater chance of getting a shock, people demanded to be paid no more to receive it.¹⁰

The second hypothesis was rejected. Moreover, there was not close to statistical significance in the difference between the price in a year and the price today for either probability. This rendered the third hypothesis moot; moreover, no statistically significant difference was found in either direction.

What explains the rejection of the second hypothesis? It is conceivable that the mere thought of a shock is sufficient to induce an emotional state, and that in light of that thought, subjects were not influenced by either time lapses or by probability. Note that standard theory might predict that people would discount the future and hence be more willing to pay to avoid an imminent shock than a future one; though we hypothesized the opposite, it is noteworthy that there was no discounting of an adverse event in the future. It is also possible that although anxiety induces emotion directly, anxiety about anxiety is much less powerful. That is, individuals may have a difficult time thinking about their loss from being anxious. Future experiments might examine the role of anxiety as an emotion-inducing experience in the period before an unfavorable lottery is resolved.

4. The Demand For and Supply of Law

If probability neglect characterizes individual judgment under certain circumstances, government and law are likely to be neglecting probability under those same circumstances. If

¹⁰ The contrast with the Rottenstreich and Hsee (2001) results is instructive. Their study differed in two significant ways: (1) They asked willingness-to-pay rather than willingness-to-accept. They got much lower median values, namely \$7 and \$19.86. (They also tried a median of \$10 for a 99% chance of a shock.) (2) They did not use an incentive-compatible procedure. Hence, their values may be low for strategic reasons.

people show unusually strong reactions to low-probability catastrophes, a democratic government is likely to act accordingly, either because it is responding to the public, or because its officials suffer the same proclivities. Our discussions of Three Mile Island and the arsenic standard above are salient examples of such responsiveness.

Recall that if government actors are able to claim credit for acting, or if they would be blamed for not acting (as was the Bush Administration on arsenic), the likelihood of action bias increases (Patt and Zeckhauser, 2000). We suspect that in many domains, government responses to emotionally gripping problems with low probability of occurrence can be explained in this way. (See Hamilton and Viscusi, 1999, for examples.) We also suspect that ill-considered, future, or counterproductive reactions to past, present, or imminent risks derive in part from this phenomenon.

In the environmental area, there has been an intense debate about whether the National Environmental Policy Act requires agencies to discuss the worst-case scenario in environmental impact statements. Environmental groups sought to ensure discussion of that scenario. They did so in part to stimulate public concern, with the knowledge that the worst case might well have a great deal of salience, however unlikely it might be, that it would be “available” in the public mind. For its part, the government originally required discussion of the worst case, but changed in its mind, with the apparent understanding that people are too likely to overreact. Hence the current approach, upheld by the Supreme Court,¹¹ requires consideration of low-probability events, but only if they are not entirely remote and speculative. Only time, and future cases will tell how these terms are to be interpreted.

At least at first glance, the current approach, and the Supreme Court’s decision, seem entirely reasonable.¹² (On some of the complexities here, see Sunstein, 2007.) If the chance that the worst case will come to fruition is truly miniscule, it is plausible to say that it need not be discussed in environmental impact statements, for the principal effect of the discussion would be to activate fear, which is by hypothesis unwarranted by the facts. Worst-case analysis should not apply when accumulated evidence shows that risks are quite low, as say with a major asteroid hitting the Earth in the next year. Yet in the context of terrorism and other emotionally laden hazards, people neglect the role of probability even when the evidence suggests that the probability is quite small (Rothchild, 2001).

A good deal of legislation and regulation can be explained partly by reference to the neglect of low probabilities when emotions are running high. Consider a few additional examples¹³:

- In the aftermath of news report about emotionally gripping adverse health effects allegedly caused by abandoned hazardous waste in Love Canal, the government

¹¹ *Robertson v. Mathow Valley Citizens Council*, 490 U.S. 332, 354-56 (1989).

¹² Jon Elster has explored the possibility that the worst-case analysis still should be employed as a maximin approach when the probabilities of the various outcomes cannot be assessed (Elster, 1983). Dyed-in-the-wool Bayesians, by contrast, would say that subjective probabilities should still be assessed and used.

¹³ In pointing to the role of probability neglect in these cases, we do not mean to reach any final conclusion on what the government ought to have done in any of them.

responded with an aggressive program for cleaning up abandoned hazardous waste sites, without closely examining the probability that illness and other harm would actually occur. In fact little was accomplished by early efforts to assure people of the low probability of harm (Kuran and Sunstein, 1999). When the local health department publicized controlled studies showing little evidence of adverse effects, the publicity did not dampen concern, because the numbers “had no meaning” (Gibbs, 1998). In fact the numbers seemed to aggravate fear: “One woman, divorced and with three sick children, looked at the piece of paper with numbers and started crying hysterically: ‘No wonder my children are sick. Am I going to die? What’s going to happen to my children?’” (id.). Questions of this sort contributed to the enactment of new legislation to control abandoned hazardous waste sites, legislation that did not embody careful consideration of the probability of significant health or environmental benefits (Kuran and Sunstein, 1999). Even now, law and policy are affected by interest group pressures and public alarm; the government has sometimes neglected the probability of significant harm in making clean-up decisions (Hamilton and Viscusi, 1998).

- During a highly publicized campaign designed to show a connection between Alar, a pesticide, and cancer in children, the public demand for action was not much affected by the EPA’s cautionary notes about the low probability of getting that disease (Wildavsky, 1995). The mere idea that children might die, as a result of apple consumption, had a significant effect on behavior, with probabilistic information seeming not to reduce people’s fears.
- In the fall of 2001, vivid images of summer shark attacks created a public outcry about new risks for ocean swimmers. This was so notwithstanding the exceedingly low probability of a shark attack, and the absence of any reliable evidence of an increase in shark attacks in the summer of 2001. Predictably, there was considerable discussion of new legislation to control the problem, and eventually such legislation was enacted in Florida. Public fears and anxieties were not impervious to the fact that the underlying risk was miniscule; but the emotional response greatly exceeded the statistical risk.

With respect to terrorism, the anthrax scare of October, 2001, which grew out of exceedingly few incidents, provides dramatic evidence. Only four people died of the infection; only about a dozen others fell ill. The probability of being infected was exceedingly low. Nonetheless, anxiety proliferated; people focused their attention on the outcome rather than the extremely low probability of the harm. The government responded accordingly, investing significant resources in ensuring against anthrax infections. Private institutions reacted the same way, asking people to take extraordinary care in opening the mail even though the statistical risks were tiny.

In any particular case, such as anthrax, it is hard to say that precautions were excessive. This could be hindsight bias. Maybe we just got lucky that there was not some large anthrax conspiracy. But if we look across dozens of cases, we can observe a pattern in which salient but extremely low probability risks are sometimes met with excessive responses. We should reiterate our horse-clop point made earlier. If the preventive measures significantly assuage public fears, perhaps they are worthwhile even if they reduce risk little if at all. But in some

cases, we suspect that to the contrary, some such measures stir rather than reduce fears by making the threat salient. Think of posting a police officer at the doors of college buildings to protect against campus shootings.

Arguably, the most severe recent example of overreaction to a risk threat is the follow-on to the terrorist attacks of September 11, 2001. Public fears and anxieties helped to produce the Iraq War, and to private and public costs that were orders of magnitude higher than the costs of the attacks themselves. A full explanation of the Iraq War would of course have to include a number of factors, but any such explanation would point, in part, to action bias and probability neglect.

What might be done, recognizing the widespread tendency toward overreaction to emotional risks? We do not have the space to answer fully, but with respect to regulatory policy, institutional safeguards are the best way of ensuring against the harmful consequences of probability neglect. The Office of Information and Regulatory Affairs, within the Office of Management and Budget, monitors agency action to ensure that it is directed against genuinely significant problems.¹⁴ A general requirement of cost-benefit balancing, with careful attention to the best estimates of relevant probabilities, should provide a check on regulations that deviate substantially from objective evidence, providing far too little risk reduction for the resources required. (Such requirements should also provide an impetus to preventive measures that the public might not seek.) Reduction in public fears should count, but they are not a trump card. If government wants to protect against hysterical precautions, analytic requirements and institutional checks will provide a start.

5. Conclusion

Our central goal here has been to understand overreactions to fearsome risks. We have suggested that when risks are vivid, people are likely to be insensitive to the probability of harm, particularly when their emotions are activated.¹⁵ If terrible outcomes are easy to visualize, large-scale changes in thought and behavior are to be expected, even if the statistical risks are dramatically lower than those associated with many activities where the stakes are equivalent but do not raise public concern. This claim about action bias helps explain public overreaction to certain highly publicized, low-probability risks, including those posed by sniper attacks, nuclear power, abandoned hazardous waste dumps, anthrax, and perhaps terrorism more generally. With financial crises, as late 2008 made tragically clear, fears and anxieties, and the action bias they induce, may dramatically magnify both the likelihood and size of a severe adverse outcome.

It follows that government regulation, affected as it is by the public demand for law, is likely to stumble on the challenge of low probability harms as well. The government should not swiftly capitulate if the public is demonstrating action bias and showing an excessive response to

¹⁴ As this paper was being completed, Cass Sunstein was appointed to head the Office of Information and Regulatory Affairs in the Obama Administration.

¹⁵ In future work, we expect to examine the complementary concept of payoff neglect: when emotions run high the size of potential losses will tend to be slighted. The emotion may be stimulated by anger due to the source of the risk, or merely a vivid description of the risk itself, apart from its magnitude.

a risk whose expected value is quite modest. A critical component of government response should be information and education. But if public fear remains high, the government should determine which measures can reduce it most cost effectively, almost in the spirit of looking for a placebo that may do little for risk but do a lot to reduce fear. Valued attributes for such measures will be high visibility, low cost, and perceived effectiveness. Reducing fear offers two major benefits: (1) Fear itself imposes significant costs. (2) Both private and public responses in the face of fearsome risks are likely to be far from rational. These observations lead to the difficult questions of how to monetize and reduce public fear. The answers lie well beyond the current topic.

Acknowledgements

Our discussion of probability neglect draws on Sunstein (2002b). Thanks to Chris Robert, Paul Slovic, Ngoc Anh Tran, Adrian Vermeule, participants in the December 2008 conference “The Irrational Economist” in honor of Howard Kunreuther, and others for valuable discussions. A brief distillation of some of the points made in this paper will appear as Sunstein and Zeckhauser (2009).

References

- Alkhami, A.S. and P. Slovic (1994). “A Psychological Study of the Inverse Relationship Between Perceived Risk and Perceived Benefit,” *Risk Analysis* 14, 1085-1096.
- Elster, J. (1983). *Explaining Technical Change*. Cambridge: Cambridge University Press.
- Foster, K., D. Bernstein, and P. Huber (eds.) (1993). *Phantom Risk: Scientific Inference and the Law*. Cambridge, Mass.: MIT Press.
- Gibbs, L.M. (1998). *Love Canal: The Story Continues*. New York: New Society Publishers.
- Hamilton, J. and W. K. Viscusi (1999). *Calculating Risks: The Spatial and Political Dimensions of Hazardous Waste Policy*. Cambridge, Mass: MIT Press.
- Johnson, E., J. Hershey, J. Meszaros, and H. Kunreuther (1993). “Framing, Probability Distortions, and Insurance Decisions,” *Journal of Risk and Uncertainty* 7, 35-51.
- Kahneman D. and A. Tversky (1979). “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica* 47, 263-291.
- Kuran, T. and C. Sunstein (1999). “Availability Cascades and Risk Regulation,” *Stanford Law Review* 51, 683-768.
- Loewenstein, G. and J.S. Lerner (2003). “The Role of Affect in Decision Making.” In R. Davidson, H. Goldsmith, and K. Scherer (eds.), *Handbook of Affective Science* (pp. 619-642). Oxford: Oxford University Press.
- Loewenstein, G. F., E.U. Weber, C.K. Hsee, and E.S. Welch (2001). “Risk as Feelings,” *Psychological Bulletin* 127, 267-286
- Margolis, H. (1993). *Dealing With Risk*. Chicago: University of Chicago Press.
- Oates, W. (2002). “The Arsenic Rule: A Case for Decentralized Standard Setting,” *Resources* 147, 16-18.

Patt, A. and R. Zeckhauser (2000). "Action Bias and Environmental Decisions," *Journal of Risk and Uncertainty* 21(1), 45-72.

Rothschild, M. (2001). "Terrorism and You – The Real Odds," *Policy Matters*, AEI-Brookings Joint Center for Regulatory Studies, available at <http://www.aei-brookings.org/policy/page.php?id=19#top>.

Rottenstreich, Y. and C. Hsee (2001). "Money, Kisses, and Electric Shocks: On the Affective Psychology of Risk," *Psychological Science* 12, 185-190.

Sandman, P., N.D. Weinstein, and W.K. Hallman (1998). "Communications to Reduce Risk Underestimation and Overestimation," *Risk Decision and Policy* 3, 93–108.

Sandman, P., P.M. Miller, B.B. Johnson, and N.D. Weinstein (1994). "Agency Communication, Community Outrage, and Perception of Risk: Three Simulation experiments," *Risk Analysis*, 13, 589–602.

Slovic, P. (2000). *The Perception of Risk*. London: Earthscan Publications.

Sunstein, C. (2002a). *Risk and Reason: Safety, Law, and the Environment*. Cambridge: Cambridge University Press.

Sunstein, C. (2002b). "Probability Neglect: Emotions, Worst Cases, and Law," *Yale Law Journal* 112, 61-107.

Sunstein, C. 2007. *Worst-Case Scenarios*. Cambridge: Harvard University Press.

Sunstein, C. and R. Zeckhauser (2009). "Dreadful Possibilities, Neglected Probabilities," in *The Irrational Economist*, E. Kerjan and P. Slovic (eds.), Public Affairs Press.

Tversky, A. and D. Kahneman (1973). "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology* 5, 207-232.

Tversky, A., and D. Kahneman (1974). "Judgment Under Uncertainty: Heuristics and Biases," *Science* 185, 1124-1131.

Viscusi, W.K. (2000). "Corporate Risk Analysis: A Reckless Act?" *Stanford Law Review* 52, 547-97.

Wildavsky, A. (1995). *But Is It True? A Citizen's Guide to Environmental Health and Safety Issues*. Cambridge, Mass.: Harvard University Press.

Zeckhauser, R. and W.K. Viscusi (1990). "Risk Within Reason," *Science* 248, 559-64.