The Impact of Extreme Events in Decisions Under Uncertainty: A Cognitive Perspective

Craig R. Fox

Fuqua School of Business
Duke University Box 90120
Durham, NC 27708-0120
919/660-7833
cfox@mail.duke.edu

In the month following the terrorist attacks of September 11, domestic air travel in the United States dropped more than 30% relative to the same period the previous year. More than six months later, revenue passenger miles remain more than 10% below what they were last Spring (Air Transport Association, March 6, 2002). Such reluctance to fly has obviously contributed to the recent recession in the U.S. economy. Ironically, it may also contribute to an increase in travel-related fatalities as people substitute flying with driving. As Myers (2001) observes:

> "If we now fly 20 percent less and instead drive half those unflown miles, we will spend 2 percent more time in motor vehicles. This translates into 800 more people dying as passengers and pedestrians. So, in the next year the terrorists may indirectly kill three times more people on our highways than died on those four fated planes."

What factors influence the fear of extreme events such as an airplane hijacking? The experience of risk is difficult to define and measure, though some fruitful axiomatic (e.g., Pollatsek & Tversky, 1970) and psychometric (e.g., Slovic, 1987) approaches have been advanced. Rather than dwell on phenomenological questions of how people perceive risk, I prefer to focus on a more positivistic question: what influences people's willingness to take protective action against extreme events? This latter question can be couched in terms of simple gambles that can be analyzed using tools of behavioral decision theory. For instance: "How much would you be willing to pay to eliminate the possibility that the next airplane you board will be destroyed by terrorists?" In this case, an extreme consequence (a terrifying death) is contingent on an uncertain event (terrorists destroy the plane on which you are flying). Though it may seem perverse to ask people if they would be willing to pay $\$x$ to reduce or eliminate the possibility of a terrorist attack, travelers have been indirectly making this choice for the past several months. Last November President Bush signed the Aviation and Transportation Security Act, authorizing a "September 11" surcharge of $2.50 per commercial flight segment to help pay for new federal aviation security measures—a surcharge passed on directly to consumers. Have you been willing to pay $2.50 per flight segment to reduce the threat of terrorism by an unspecified amount? For that matter, have you been willing to tolerate longer lines and more intrusive security checks?

Viewed from a behavioral decision making perspective, willingness to take protective action should be a function of both the magnitude of the perceived threat (e.g., injury or death) and the impact of the target event on which this threat depends (e.g., hijacking). Hence, rather than ask "what factors influence the perceived risk posed by extreme events?" one might ask "what factors influence the impact of an extreme event in the decision to take protective action?" In the discussion that follows I will sketch a cognitive perspective on the impact of extreme events (i.e., their weight in decision making). I will argue that the impact of an extreme event is governed by three factors: (1) attention to the event—that is, whether or not the decision maker recognizes the possibility that the target event could conceivably occur; (2) perceived likelihood that the event will occur—that is, the judged probability of the target event; (3) subjective response to the risk posed—that is, the weight that the decision maker places on the probability of the target event.

Recognizing the possibility of extreme events

All of us face myriad risks every day, regardless of how cautiously we live our lives. For instance, in the coming year if you are the "average American," the odds are approximately 1 in 400 that you will injure yourself in a bed or chair, 1 in 3,000 that you will be injured by your clothing, and 1 in 160,000 that

you will choke to death on food (Lauden, 1994). Of course, we don't typically attend to every minute risk that we face; to do so would be crippling. Most potential hazards simply never occur to people, whereas other potential hazards (e.g., a tire blowout) are recognized but edited by people to "essentially nil" (Stone, Yates, & Parker, 1994). It is true that some clinical populations obsess over statistically remote hazards such as deadly germs or falling onto subway tracks; however, most people afford extreme events little thought in the absence of some external prompting. Of course, cueing often arrives unsolicited. For instance, media coverage of news events (e.g., the war on terrorism) may automatically prime fears of semantically related risks (e.g., airplane hijacking). It is worth noting that the word "terrorism" appeared in the headline of news articles in the New York Times only 19 times in the six months preceding September 11, but 246 times in the six months following that date. It is virtually impossible for a literate member of our society to escape daily reminders of the threat of terrorism. Moreover, once awareness of a potential hazard is planted, it may be difficult for even non-clinical populations to succeed in ignoring the risk, much as the instruction not to think of a "white bear" makes such an image difficult to suppress (e.g., Wegner *et al.,* 1987).

Judging the likelihood of extreme events

Several factors influence the perceived likelihood of extreme events. Most people are unrealistically optimistic that unfavorable events will not befall them (Weinstein, 1980; see also Taylor & Brown, 1988), and so they may underestimate the probability that they will be personally affected by generic extreme events (e.g., believing that a hurricane is less likely to harm me than my neighbors or people in other parts of the state). However, when a particular extreme event is experienced personally or vicariously, the perceived likelihood may increase to unreasonable levels due to heuristic evaluation of evidence, an inappropriate partition of the sample space, and a misperception that the underlying base rate has shifted.

Heuristic evaluation of probability

The heuristics and biases literature (Kahneman, Slovic, & Tversky, 1982) suggests that people reduce the complicated task of probability estimation to simpler tasks such as memory retrieval and similarity judgment. Typically, people assess the likelihood of an extreme event by the ease with which they can recall instances of that event. Prior to September 11, the notion of suicide hijacking, while theoretically plausible, was almost unthinkable. However, since September 11 vivid images of airplanes slamming into buildings have been seared indelibly into our collective consciousness so that future terrorist acts are much easier to imagine and therefore seem to us to be much more likely. Availability of instances may also be enhanced when the target event is unpacked into a disjunction of more specific scenarios (Tversky & Koehler, 1994; Rottenstreich & Tversky, 1997). For instance, Johnson, Hershey, Meszaros, and Kunreuther (1993) report that their respondents said they were willing to pay more to for an insurance policy that covers hospitalization for "any disease or accident" than for a policy that covers hospitalization "for any reason." Presumably, unpacking provides a more effective retrieval cue when recalling of instances of past events that required hospitalization.

The exaggerated subjective likelihood conferred by the hyper-availability of a particular extreme event may be reinforced by the natural tendency for people to ruminate concerning how a subsequent extreme event might occur. When assessing the probability of a specific threat (e.g., this plane will be hijacked), people typically adopt the *inside view,* sketching a "representative scenario that captures the essential elements of the history of the future," rather than the *outside view* that "essentially ignores the details of the case at hand… [and instead] focuses on the statistics of a class of cases chosen to be similar in relevant respects to the present one." (Kahneman & Lovallo, 1993) If the passenger sitting next to me on a plane physically resembles one of the hijackers of September 11, it is easy for me to elaborate a scenario in which he will seize controls of my plane. The more I elaborate my imagined scenario (e.g., he could have hidden a ceramic knife in his bag), the more closely it will resemble a real hijacking and therefore more likely it will subjectively appear. Of course, adding conjunctive features to a scenario can only render it more statistically remote (see Tversky & Kahneman, 1983).

Partition dependence

The perceived probability of an extreme event may be influenced not only by an assessment of the nature of the event but also by the number of interchangeable events that the judge considers. Some of my recent research suggests that judged probabilities are typically biased toward an "ignorance prior" probability that is defined by maximum entropy of belief across all elementary events into which the sample space is subjectively partitioned. Usually, people partition the sample space into a two-fold "case" partition (e.g., this plane will be hijacked; this plane will *not* be hijacked) which may bias probability estimates upward (i.e., toward 50-50). Extreme events may seem more remote, however, for people who invoke a "class" partition of many comparable events. For instance, noting that there are well over 20,000 domestic airline departures per day in the United States, the consumer might arrive at an "ignorance prior" probability of 1 in 20,000 that the next hijacking, should it occur today, will happen to be *my* flight. We have found that the way in which people subjectively partition the sample space can be manipulated by the way in which a probability query is couched (Fox & Rottenstreich, forthcoming). I suspect that people would be willing to pay more to cut in half the probability that "your flight will be destroyed by terrorists" than they would to cut in half the probability that "the next Domestic flight to be destroyed by terrorists will be yours."

Detecting a regime shift

Of course, it could be that the advent of an extreme event reflects a change in the underlying base rate of the relevant class of extreme events. One might argue that the terrorist attacks of September 11 have augured in a new era of suicide hijacking in which air travel really is less safe than it had been previously. Indeed, many media sources pronounced that the "world has changed" since that fateful date. Or has it? Perhaps Americans did nurture a false sense of security when boarding planes prior to September 11. On the other hand, counter-terrorism measures have been stepped up in ways that have probably made air travel much less susceptible to the suicide hijackings that we have come to fear, whereas we remain blissfully unaware of other, unimagined threats to airline security. As Gladwell (2001) observes, "The better we are at preventing and solving crimes before us, the more audacious criminals become…The history of attacks on aviation is the chronicle of a cat-and-mouse game, where the cat is busy blocking old holes and the mouse always succeeds in finding new ones." So it is possible, in fact, that the probability of airline attacks has remained relatively stable in the past year. This said, recent research has determined that in relatively stable regimes, people tend to overreact to proximate signals (e.g., a hijacking) when updating their beliefs that the underlying probability distribution has shifted (Massey & Wu, 2002).

Moreover, when people do infer a regime shift in dynamic systems they may make an attribution that is overly narrow. For instance, one might argue that following the 1993 bombing of the World Trade Center people overreacted in updating their beliefs that the buildings would be successfully "bombed," underappreciating the effectiveness of security measures were subsequently put into place. However, it could be that people also underreacted in updating their beliefs that the buildings would successfully "attacked" by other means, such as hijacking, for which heightened security was not put into place.

Weighting perceived probability of extreme events

According to prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) probabilities are weighted by an inverse-S shaped weighting function that overweights small probabilities and underweights moderate to large probabilities. More recent work (Tversky & Fox, 1995; Fox & Tversky, 1998) suggests that when objective probabilities are not available, people act as if they weight their own *judged* probabilities by the same function.

The weighting function is characterized by (1) a degree of curvature that reflects the decision maker's sensitivity changes in probability and (2) a degree of elevation that reflects the decision maker's overall willingness to act in the presence of a particular source of uncertainty. Although the qualitative shape of the weighting function is consistent across individuals, there are considerable individual differences in the curvature and elevation of decision weights (Gonzalez & Wu, 1999), and the elevation may vary depending on a person's perception of his or her relative knowledge judging the likelihood of the target event (Heath & Tversky, 1991; Fox & Tversky, 1995; Fox & Weber, 1998).

Curvature of the weighting function.

Curvature of the weighting function models diminishing sensitivity to increasing probabilities around the natural boundaries of 0 and 1. The impact of a particular probability is more pronounced when it turns an impossibility into a possibility (or a possibility into a certainty) than when it merely makes a possibility more likely. Hence, an increase in probability of an extreme event by 1% per year will induce more fear when added to a previous probability of zero than when added to a previous probability of 1%. Diminishing sensitivity implies that small probabilities will be overweighted by decision makers. Overweighting of low-probability gains and losses, respectively, explains why people are risk seeking when purchasing lottery tickets and risk averse when purchasing insurance.

Diminishing sensitivity seems to be more pronounced for more "affect-rich" consequences. For instance, Rottenstreich & Hsee (2001) report that their participants found an electric shock about as unattractive as a penalty of $20. However, participants were willing to pay seven times as much to avoid a 1% chance of an electric shock as they were to avoid a 1% chance of paying $20. If the prospect of a mild electric shock is sufficiently "affect rich" to distort the weighting function, the prospect of a terrifying death will surely exacerbate curvature of the weighting function and amplify overweighting of extreme events.

Diminishing sensitivity causes people to find the elimination of a particular risk more attractive when considered in isolation from the broader family of risks in which it is situated, a phenomenon known as the *pseudocertainty effect*. Kahneman and Tversky (1984) found that most people favored a 20% chance to win $45 over a 25% chance to win $30. However, when the gamble was reframed as a two-stage game so that one had a 25% chance of proceeding to a choice between a certain $30 or an 80% chance to win $45, preferences reversed. In a similar vein, suppose that you are considering taking off for the weekend for your annual ski trip. It may surprise you to learn that the average alpine skier will face approximately a 1 in 23,000 chance of dying in 15 hours of skiing. Such numbers may tempt you to give up skiing altogether so that you can reduce your risk from 1 in 23,000 to zero. However, viewed in a broader context, foreswearing skiing only reduces the aggregate annual probability of accidental death from 345 in a million to 301 in a million, and such a small reduction in risk may not seem worth the cost of relinquishing the pleasure of skiing.

Elevation of the weighting function.

Elevation of the weighting function models overall attitudes toward uncertainty throughout the probability range. Individuals for whom uncertainty is particularly unpleasant would be characterized by a highly elevated weighting function for losses (i.e., a greater overall tendency to overweight possible losses). Recent studies have demonstrated that willingness to act under uncertainty is influenced by the decision maker's perception of his or her relative ignorance concerning the uncertain event. People do not fear the unknown *per se*; they fear what they *know* they do not know (Heath & Tversky, 1991; Fox & Tversky, 1995; cf. Ellsberg, 1961). Last November a friend of mine who studies legal decision making declined to attend a conference on the opposite coast, citing a new fear of flying. When I attempted to persuade him that statistically this fear may be less than rational, he responded, "I know the probability of disaster is very low. What bothers me is that we don't *know* what that probability is." Furthermore, the perception of comparative ignorance is highly context-dependent. In one study (Fox & Weber, in press) we found that people would much rather bet on their prediction of the winner of the upcoming Russian election (a moderately familiar event) if they had previously been reminded of an upcoming election in the Dominican Republic (a less familiar domain) than if they had been previously reminded of the upcoming American election (a very familiar domain). One suspects that reminding people of more familiar hazards (e.g., auto accidents) can only serve to make novel hazards (e.g., suicide hijacking) more daunting.

People can buffer fear of many extreme events with the illusion of control. Apparently, people would rather play a game of chance against a shy and awkward opponent than a confident and outgoing one, and would rather bet on a lottery ticket that they picked themselves than one that was assigned to them at random (Langer, 1975). Subjects would rather bet on a dart-throwing task than a chance gamble, holding the probability of winning constant (Howell, 1971). Similarly, managers typically treat risk as a "challenge to be overcome" and reject the gambling metaphor (March & Shapira, 1987). Hence, even

though I know that the chances of my dying in an auto accident this year are approximately 1 in 6,000, I find this risk less frightening when I am behind the wheel (although most people think that they are better-than-average drivers, they should be more concerned about other drivers and the possibility of mechanical failure than their own driving). In contrast, I am never behind the wheel of a commercial airplane and cannot readily delude myself into believing that I can control hurricanes in North Carolina, so these extreme events seem more frightening.

Summing up: Isolation errors and rational decision making

Willingness to take protective action against an extreme event is a function of the event's impact on the decision maker's assessment of possible consequences. For better or worse, it is almost impossible to fear that which we cannot imagine. But when we do acknowledge the possibility of a potential hazard, we often overestimate its likelihood and overweight its probability. Residents of New Madrid, Missouri in the late 1980's probably gave the possibility of a major earthquake little thought. The last major earthquake to hit that region had been in 1812. However, that quake was the largest ever in the continental United States, registering 7.9 on the Richter scale. Perhaps residents were living with a false sense of security. In November 1989 a Memphis newspaper ran an article saying that "a scientist who correctly predicted October's San Francisco earthquake" was predicting an earthquake on the New Madrid fault on or around December 3 1990. The scientist was Iben Browning, a New Mexico climatologist who sold long-range weather forecasts to businesses. The story went on to interview real earthquake experts who discredited Browning's forecast, but somehow these caveats were forgotten. The Earthquake became the story of the year in the region, leading to a rash of emergency planning drills, plans by many to leave the area, and a precipitous increase in the purchase of earthquake insurance. In this case, one could argue that these were rational responses to an irrational fear.

I have argued that willingness to take protective action against extreme events requires: (1) an awareness of a particular hazard; (2) assessment of its likelihood; and (3) weighting of the perceived likelihood. Any rational policy or corrective measure must address potential inconsistencies at all three levels. First, one might eliminate biases due to weighting of probabilities by consciously weighting outcomes by their respective probabilities (i.e., apply expected value maximization to cost-benefit analysis). This strategy alone will be insufficient to purge bias from decisions because it fails to address inconsistencies of belief. To wit, a few years ago we (Fox, Rogers & Tversky, 1996) observed a sample of professional options traders who applied expected value maximization when assessing selling prices for chance gambles. However, when pricing prospects contingent on the future close of familiar stocks, these options traders apparently applied expected value maximization to subadditive judgements of probability, thereby offering higher selling prices for more specific events. Second, one might attempt to expunge biases of belief by researching base rate statistics for all relevant risks, where available. Here again, this approach could fail if the decision maker neglects to attend to other, unspecified hazards.

The research reviewed here suggests that the impact of extreme events is enhanced when attention is directed toward a specific threat, so that people typically commit what has been called "isolation errors" (e.g., Kahneman & Lovallo, 1993). First, people selectively attend to some extreme events and neglect others. This may lead to an irrational fear of particular, salient events (e.g., a suicide hijacking) and a false sense of security concerning other less salient events (e.g., poisoning of the water supply). Second, when assessing likelihood, people typically construct scenarios to support a particular extreme event and they underappreciate base rates. As decision makers elaborate imagined scenarios extreme events may seem more likely. Third, people typically weight low probability events in isolation from other relevant events. Hence, the public may call on officials to marshal all available resources to reduce the probability of an extreme event (e.g., airplane crash) due to a particular cause (e.g., hijacking) to near zero, but neglect to consider that the probability of the extreme event due to other, less available causes (e.g., mechanical failure, midair collision, sabotage, surface-to-air missile) remains significant. Although the tremendous resources that have been allocated since September 11 to airport security may assuage public fears of a particular salient threat, one wonders whether the measures taken are the most efficient means of protecting the public against the more general threat of dying in an airplane crash or dying in equally terrifying ways.

References

Air Transport Association (2002). State of the industry: Recent trends for U.S. Carriers. http://www.air-transport.org/

Ellsberg D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics, 75,* 643-669.

Fox, C.R., & Rottenstreich, Y. (forthcoming). Partition priming in judgment under uncertainty. Accepted pending minor revision, *Psychological Science.*

Fox C.R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science, 44,* 879-895.

Fox C. R., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics, 110,* 585-603.

Fox, C.R., Rogers, B.A., & Tversky, A. (1996). Options traders exhibit subadditive decision weights. *Journal of Risk and Uncertainty, 13,* 5-17.

Fox, C.R., & Weber, M. (in press). Ambiguity aversion, comparative ignorance, and decision context. *Organizational Behavior and Human Decision Processes.*

Gladwell, M. (2001). Safety in the skies: How far can airline security go? *The New Yorker,* October 1, 50-53.

Gonzalez, R. & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology, 38,* 129-166.

Heath, C., & Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty, 4*, 5-28.

Howell, W. (1971). Uncertainty from internal and external sources: A clear case of overconfidence." *Journal of Experimental Psychology, 89, i240-243.*

Johnson, E.J., Hershey, J.. Meszaros, J., & Kunreuther, H. (1993). Framing probability distortions, and insurance decisions. *Journal of Risk and Uncertainty, 7,* 35-51.

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science, 39,* 17-31.

Kahneman, D., Slovic, P. & Tversky, A. (1982) (Eds.). *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist, 39,* 341-350.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 4,* 263-291.

Langer, E. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32,* 311-328.

Lauden, L. (1994). *The Book of Risk.* New York: John Wiley & Sons.

March, J., & Shapira, Z. (1987). Managerial perspectives on risk and risk taking. *Management Science, 33,* 1404-1418.

Massey, C. & Wu, G. (2002). Detecting regime shifts: The psychology of over- and under-reaction. Unpublished manuscript, Fuqua School of Business.

Myers, D.G. (2001). Do we fear the right things? *APS Observer,* November 2001.

Pollatsek, A., & Tversky, A (1970). A theory of risk. *Journal of Mathematical Psychology, 7,* 540-553.

Rottenstreich, Y. & Hsee, C.K. (2001). Money, kisses, and electric shocks: On the affective psychology of risk. *Psychological Science, 12,* 185-190.

Slovic, P. (1987). Perception of risk. *Science, 236,* 280-285.

Rottenstreich, Y., & Tversky, A. (1997). Unpacking, repacking, and anchoring: Advances in support theory. *Psychological Review, 2,* 406-415.

Stix, G., 1991. Future Shock: Fear of the "Big One" fuels sales of insurance and velcro, in Scientific American. January 1991. p. 111.

Stone, E.R., Yates, J.F., & Parker, A.M. (1994). Risk communication: Absolute versus relative expressions of low-probability risks. *Organizational Behavior and Human Decision Processes, 60,* 387-408.

Taylor, S.E. & Brown, J.D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103,* 193-210.

Tversky, A., & Fox, C.R. (1995). Weighing risk and uncertainty. *Psychological Review, 102,* 269-283.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty, 5, 297-323.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90,* 293-315.

Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review, 101*, 547-567

Wegner, D.M., Schneider, D.J., Carter, S.R., & White, T.L. (1987). Paradoxical effects of thought suppression. *Journal of Personality and Social Psychology, 53,* 5-13.

Weinstein, N.D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology, 39,* 806-820.