

Recollection Bias and the Combat of Terrorism

W. Kip Viscusi and Richard J. Zeckhauser

ABSTRACT

Survey respondents assessed the risks of terrorist attacks and their consequences and were asked how their assessments changed from before September 11 to the present. This paper analyzes those current and recollected risk assessments. More than half of the respondents exhibited what we label “recollection bias”: looking backward from 2002, 2003, or 2004, they reported that their forward-looking risk assessments did not rise after September 11. However, government expenditures and policies and evidence from insurance markets suggest that there were major risk increases in estimated risks. Respondents were generally willing to support airplane passenger profiling when the time costs of alternative policies were great and were supportive of strengthened surveillance policies to address terrorism risks as well. However, individuals suffering from recollection bias are less supportive. We label as “recollection choice bias” a link between policy choices and recollection bias.

1. INTRODUCTION

The September 11, 2001, terrorist attack imposed devastating costs on the United States and dramatically increased policy makers’ assessments of the likelihood and magnitude of future terrorist attacks. A series of significant policy measures—such as the creation of the Department of Homeland Security, more intensive airport inspections, and a war that deposed the Taliban in Afghanistan—were taken to combat these risks.

The trade-offs involved in preventing terrorist attacks in the United

W. KIP VISCUSI is Cogan Professor of Law and Economics, Harvard Law School. RICHARD J. ZECKHAUSER is Ramsey Professor of Political Economy, Kennedy School of Government, Harvard University. Viscusi’s research is supported by the Harvard Olin Center for Law, Economics, and Business. We are indebted to Jeffrey Brown and Miriam Avins for helpful comments; to Anda Bordean, DeYett Law, and Paul Novosad for valuable research assistance; and to an insightful referee who led us to recast our paper. We are grateful to Kent Smetters for conducting the study for us at the Wharton School.

[*Journal of Legal Studies*, vol. 34 (January 2005)]

© 2005 by The University of Chicago. All rights reserved. 0047-2530/2005/3401-0002\$01.50

States are fundamental, yet often hard to grasp. While most policy trade-offs balance money and some policy goal, combating terrorism often creates conflicts between the fundamental attributes of safety and civil liberties. Advocates on each side claim that the goal they favor cannot be compromised, yet their claims frequently cannot be satisfied simultaneously. In addition, what policies are appropriate against terrorism will depend importantly on probabilities as well as preferences. Policy makers must implicitly make subjective assessments of risk to formulate policies. Moreover, their policies will also be influenced by the public's support, which in turn will depend on its perceptions of risks and of the effectiveness of measures to control them.¹

Terrorist risks—man-made events of low probability and high consequence—are hard to estimate, making decisions on how to respond particularly prone to biases that are not entirely rational. While sound policies should be based on carefully derived subjective estimates of the risks and the ways those risks are affected by various measures, we simply do not have a large data source telling us how the risk of a terrorist attack would change with some change in policy.² We lack answers to questions such as, If we seek to reduce terrorism risks through a variety of screening efforts with different levels of stringency, how will this change the expected number of fatalities?

This paper addresses two major questions. First, what probability judgments do people make about terrorism risk? Second, how do their judgments relate to their preferences for policies that balance the risk of terrorism against the cost of measures to reduce the risk?

We analyze the risk perceptions and policy preferences of groups of relatively sophisticated citizens, professional school students at Harvard University and undergraduate business students at the Wharton School of the University of Pennsylvania. We examine these respondents' patterns of response and search for biases and irrational judgments. The policy trade-offs involved in combating terrorism are more subtle than the choices in most policy choice contexts.

Our three primary findings are the following: (1) More than two-

1. Of course, it was naïve to assume such a low probability before September 11. For example, the simultaneous successful bombings of the U.S. embassies in Kenya and Tanzania, or indeed the unsuccessful 1993 World Trade Center bombing, should have led to much higher estimates.

2. In contrast, we have plenty of information about car accidents. For everyday risks, we get dozens or indeed many thousands of data points each year. Moreover, unlike terror risks, gauging them does not require estimates of human intentions.

thirds of respondents believe that their estimated risk of a future terrorist attack did not rise from before September 11 to after September 11. (2) Roughly three-quarters of respondents favor racial profiling in airports if the alternative is a 60-minute line. (3) There is significantly less support for profiling among those who state that their estimate of future terrorist risk did not rise after September 11.

How can we be confident that respondents' estimates of the risk of a terrorist attack should have increased dramatically since September 11, despite sustained efforts against Al Qaeda and sister organizations and generally tighter security in the United States? We rely on market information and the estimates implicit in the choice of vigorous and extremely expensive government antiterrorist policies, which possibly may include the war in Iraq. For example, terror risks were excluded from many new insurance policies after September 11. Difficulties in securing such coverage, and dramatic escalation in proposed rates, eventually led to a 3-year government reinsurance program.³ In 2004, the U.S. government is spending many tens of billions of dollars and implementing some unappealing policies to fight terrorism, which implies that the returns from such measures are far greater after September 11 than was thought prior to it. In addition, an attack that cost 3,000 lives and perhaps \$100 billion in damage was beyond the expectation of virtually all observers; this highly unlikely outcome greatly increases the expected damages from subsequent successful attacks. Evidence on attention or inattention to terrorist risks by newspapers and national survey institutions is also instructive. For example, there were 51 articles on terrorism from January 1, 2001, through August 31, 2001, in the *New York Times*. The equivalent figures for 2002 and 2003 were 393 and 234.⁴

The paper proceeds as follows. Section 2 discusses a statistical pro-

3. The dramatic effect this attack had on tightening insurance markets—one could buy less coverage at higher prices—is documented by Brown et al. (2003), Cummins and Lewis (2003), Doherty, Lamm-Tennant, and Starks (2003), and Smetters (2004). Prior to September 11, Chicago's airports had \$750 million of insurance for a premium of \$125,000. Coverage after September 11 was reduced to \$150 million and the premium raised to \$6.9 million (Jaffee and Russell 2002, pp. 2–3). In the reinsurance program, the government pays 90 percent above the insurance companies' deductible, which is 10 percent of a company's premiums in 2004. Brown et al. (2003) and Smetters (2004) provide excellent discussions of the legislation and of policy choices pertaining to insurance involving terrorism. Other possible policy responses are explored by Keohane and Zeckhauser (2003) and Kunreuther and Heal (2003).

4. Articles were researched through Lexis-Nexis. The pace through June 2004 is a bit over half that for 2003.

cedure for updating a risk after a catastrophe. It first describes the applicable methods of decision theory. It then presents a brief discussion of recollection bias and introduces a model that analyzes how such bias might reduce support for antiterrorism risk policies.

Section 3 details our assessment of individuals' beliefs about terror risk and examines the determinants of such beliefs. It pays particular attention to the possible influence of recollection bias.

Section 4 focuses on the key trade-offs in combating terrorism. Our main case study addresses respondents' willingness to support targeted passenger screening at airports, where the trade-off is between additional waiting time for all passengers and targeting passengers according to some demographic profile of potential terrorists. It focuses on recollection bias as an explanatory variable. Individuals who couple such bias with the belief that policies before September 11 were optimal for that time will feel that it is undesirable to compromise civil liberties now to promote safety. This constellation of perceptions and beliefs produces recollection choice bias.

Section 5 discusses qualifications to our study and suggests how some concerns could be resolved in the future. The concluding Section 6 highlights a major finding: individuals' preferences among policies to combat terrorism do not reflect current assessments of terror risks. Those preferences, however, are strongly influenced by whether they perceive that the risk level has increased. If so, more aggressive policies are supported; if not, they are not.

2. UPDATING RISK, RECOLLECTION BIAS, AND POLICY PREFERENCES

In estimating the risk of a catastrophe in a period, say, a month, one is essentially estimating an unknown probability. As a reference case, assume that the underlying risk of catastrophe does not change from period to period. Since the value of p is unknown, the subjectively perceived risk of catastrophe should change over time. In any period when the catastrophe does not occur, the estimated p declines; on the rare occasions when it does occur, the estimate increases dramatically.⁵

5. Such an updating process is often captured by a beta distribution with parameters a and b , where the initial p is $a/(a + b)$. After a nonoccurrence, it goes to $a/(a + b + 1)$, and after an occurrence, it goes to $(a + 1)/(a + b + 1)$, with the process repeating the next period. With the parameters $a = .1$ and $b = 20$, the initial likelihood of an attack in a month would be $1/200$. A month later, assuming no attack, the probability would be $1/210$. But if there were an attack, the probability would jump to $11/210$, a more than

2.1. Recollection Bias and Hindsight Bias

Our surveys were conducted after the tremendous shock of September 11, so respondents were unlikely to recall what they thought the risk of a major terrorist attack was before September 11. Indeed, we speculate that many of our respondents never thought about that issue, just as many heart attack victims do not think about that possibility until they have had a heart attack. Therefore, we needed to ask people to compare their current estimates of risk with what they think their estimates were prior to the catastrophe. While there is no conclusive way to learn what our respondents' *a priori* estimates really were, we can compare the patterns of their responses to those of their colleagues, to what Bayesian analysis might suggest, and to changing external estimates of risk, such as insurance rates for terror losses.

Our major conjecture is that many respondents will think that their *a priori* risk estimates were much higher than they really were. That is, their stated recollection of their pre-event probability and their assessed postevent probability are too close together. We label this cognitive error "recollection bias."

Recollection bias could be diagnosed easily if it was simply poor updating of an unchanging probability. However, the probability will likely shift. First, even if risk prevention measures are held fixed, actual risk levels may change between the before and after period. For example, whatever his risk perceptions may be, an individual who has had a heart attack is at much higher risk than a medical peer—same age, cholesterol, and so on—who has not. After an earthquake in a known zone, the immediate earthquake risk diminishes. Terrorists may learn from past attacks or gather new adherents and resources from successes. Second, risk prevention measures may be increased after an unlikely bad outcome. Our heart patient may begin exercising. The United States and some allies went after Al Qaeda after September 11. While we assert that assessments of future terrorist risk should have increased noticeably after September 11, there is no completely objective way to demonstrate this, particularly since there are no surveys of the risk from before September 11 of which we are aware.

Recollection bias is a first cousin of the widely studied hindsight bias

10-fold increase. The model, and its assumption of a constant probability, is too simple to capture the complexities of terror risks, but the qualitative result that estimated risk should jump following an unexpected attack is robust.

(Fischhoff 1975).⁶ Hindsight bias states that after an unlikely event, individuals will overestimate the extent to which the event was foreseeable. Like hindsight bias, recollection bias addresses individuals' ex post assessments of their prior estimates. The most important difference is that recollection bias compares ex post assessments to current estimates of the risk going forward, not to actual prior estimates.

Studies of hindsight bias have not generally tracked risk beliefs over time; rather, the approach has been experimental and typically has been conducted across subjects. One group of subjects considers a hypothetical scenario that describes a situation before a risk event has occurred and asks respondents to assess the risk. A second group of subjects considers the risk event after the lottery has been run, where the outcome such as an accident is known, and is asked what their preaccident risk assessment would have been.⁷ Thus, the comparison is between assessed risk when given an ex ante experimental scenario and assessed prior risks when given an ex post experimental scenario. We cannot duplicate this format for terrorism risks: we could hardly construct a credible hypothetical scenario in which people are ignorant of the World Trade Center attacks. As a result, we studied recollection bias: we compared whether respondents' current assessments of the terrorism risk are greater than they recollect that they were before the attack. The Appendix presents a schematic view of hindsight bias and of recollection bias.

In addition to recollection bias, there are at least two other reasons why people did not raise their estimates of the risk of a terrorist attack after September 11. First, individuals might have thought that policy measures, such as increased screening at airports or the wars in Afghanistan and Iraq, lowered terror risks sufficiently to bring them to or below September 11 levels.

Second, although the occurrence of such an extreme outcome of an event with such a low and hard-to-estimate probability should significantly shift rationally computed risk beliefs, some people will simply ignore the new information. Those whose risk beliefs did not change are "resistant learners" who ignore information.⁸

6. We thank a referee for encouraging us to look for a term that distinguishes our concept from traditional hindsight bias.

7. In one version of the experiment, the same individuals are asked their values before and after a hypothetical accident.

8. Resistant learning could result from respondents' anchoring their risk beliefs on their preaccident risk assessments. An early discussion of anchoring appears in Tversky




Figure 1. Terror risk versus civil liberties

2.2. Risk Assessments and Preferred Policies

Why might recollection bias arise? As with hindsight bias, individuals might wish to magnify their own ability to assess the world. For example, after the World Series, sports fans might believe that their prior estimates of the winning team's prospects were far higher than were the Las Vegas odds. Recollection bias might also be a product of motivated reasoning: an individual who did not want to change policies might assert—even to himself—that probabilities had not changed.

Figure 1 illustrates how risk perceptions might influence preferences on civil rights and antiterror policies. The original set of trade-offs between civil liberties and expected terrorism losses is characterized by the curve *xx*. As the level of civil liberties increases, one would expect higher expected terrorism losses because screening of passengers and other anti-terrorism efforts will be less vigorous in such a world. Individuals have

and Kahneman (1974). Resistant learning could have been expected here. In experimental tests, Kelman, Fallas, and Folger (1998) found little evidence of hindsight bias for gambles in which the risks were transparent. When the odds were more difficult to calculate, which is surely the case with terrorist calamities, respondents tended to exhibit much more hindsight bias.

indifference curves with respect to combinations of civil liberties and terrorism losses, which are indicated by I_1 , I_2 , I_3 , and I_4 in the figure. All individuals prefer lower levels of terrorism losses and higher levels of civil liberties, so preferences take on a higher value moving in the south-easterly direction of the figure. The indifference curves are vertical offsets of each other. This relationship implies, as would seem reasonable though not inevitable, that the trade-off between lives and liberties does not depend on the expected number of lives already lost. Given the preferences indicated, the optimal choice before the September 11 attack is at point A, offering expected losses of a .

After the terrorist attack on September 11, the trade-off curve between expected terrorism losses and civil liberties is characterized by curve yy . For any given level of civil liberties, the expected terrorism losses have increased because we have raised our assessed risks from a terrorist attack. This curve is twisted upward: for any civil liberties level, the yy trade-off curve is steeper than the xx curve. This tilting reflects our assumption that a reduction in civil liberties, like any expensive safety measure, reduces losses more when risks are high. It also reflects the need for caution in a perilous situation. If we were to keep civil liberties at their optimal pre-September 11 level at point A, we would land at point C. Instead, the new optimal combination of civil liberties and expected terrorism losses is at D. It provides for a lower level of civil liberties than before, yet expected terrorism losses, at d , are greater.

What if individuals do not shift their risk estimates? Resistant learners, who fail to update their beliefs from September 11, will set risk estimates well below those of the rest of the population. They would wish to operate at A, as before September 11.

How would people who fall prey to recollection bias balance terror risk and civil liberties after September 11? In recollection bias, although a person's risk beliefs do not differ from those of the rest of the population, remembered *ex ante* beliefs are raised relative to current beliefs, possibly even exceeding them. Let risk beliefs R_A and R_P represent *ex ante* and *ex post* beliefs, with $R_A(i)$ being perceptions of *ex ante* views at time i , where $i = A, P$, and similarly for $R_P(i)$. For expositional simplicity, we deal with the case in which *ex post* and remembered *ex ante* beliefs are equal. That is, those suffering recollection bias have $R_A(P)$ equal to $R_P(P)$. They believe that both before and after the attack they were operating on curve zz and had chosen point B. Since they assume that society behaved optimally before the attack and there is no change

in risk after the attack, then there is no reason to move from point B.⁹ The associated civil liberties level remains at *d*.

2.3. Judgments of External and Internal Decisions

Individuals' judgments about the likelihood of a bad outcome may depend on who made the decision that might have affected the bad outcome. When individuals were external to decisions made before the low-probability event occurred, if they never approved of them or were in no way invested in them, then they will be termed "external." In making normative judgments, they will be judging others, and their judgments might suffer from hindsight bias.¹⁰ This phenomenon is akin to what Kelman, Fallas, and Folger (1998) refer to as "third-party hindsight bias" in the risk perception context. For example, jurors in negligence cases often indicate that they believe that risk levels were apparent to the defendant before a significant accident and that the defendant should have taken appropriate efforts to reduce these risks.

When individuals are judging their own past decisions or others' past decisions of which they approved, we say they are internal, or invested. Recollection bias is more likely to emerge with such decisions. Our respondents composed a well-informed audience on government-citizen relations. Virtually all of them were familiar with the broad sweep of American civil liberties policies prior to September 11, and we suspect that virtually all of them approved of the protections then in place. Thus, we think of our respondents' decisions as being internal.

Our primary interest in this paper is forward-looking decisions, not past decisions. We employ the term "recollection choice bias" to refer to the tendency of recollection-biased individuals to favor different policies than other individuals with the same forward-looking probabilities.

9. There is an important distinction here with how hindsight bias would operate in jurors deliberating about negligence. In that case, the juror never made a judgment about the level of precautions taken. Hence, if he now believes the risk was higher than others felt before, he can blame the responsible party for laxity. It is interesting to note that jurors in experimental contexts have quite different attitudes to risk-taking decisions they are asked to make *ex ante* and failed risky decisions others have made. See footnote 11. In the terrorism case, we assume that respondents approved, as did virtually all Americans, of the levels of civil liberties before September 11. Loss aversion could also be at play. Given loss aversion in a situation with multiple attributes, it is far more costly to sacrifice on one dimension than it is beneficial to gain on another. Hence, there is a tendency to stick with the perceived status quo policy choice.

10. The potentially powerful role of hindsight bias with respect to assessments of liability and punitive damages is documented by Rachlinski (1998), Hastie, Schkade, and Payne (1999), Viscusi (1999), and Sunstein et al. (2002).

People who are invested in past policies and who do not think that risk levels have changed presumably judge past policies to be still optimal.¹¹

For example, two sedentary, overweight 60-year-old males, A and B, each with high cholesterol and blood pressure despite medication, suffer heart attacks. Each is told his risk of recurrence is substantial. A does not want to change his ways. B wants to follow a stringent diet and exercise regime. Presumably, B has appropriately raised his assessment of risk given his old bad lifestyle. A, by contrast, may well exhibit recollection choice bias; his estimates before and after may be too close. He believes that he always knew he was at high risk, so the heart attack may not justify a change in policies.

3. THE DETERMINANTS OF RISK BELIEFS

The discussion below draws on three surveys from students at Harvard Law School, Harvard's Kennedy School of Government, and the Wharton School at the University of Pennsylvania. The first survey was run at Harvard Law School in April 2002 and had a sample size of 95 students, most of them first-year law students. The second survey was taken by 56 Harvard Law School students who were in their second or third year and 61 Kennedy School students in November 2003, for a total 2003 sample of 117. The third survey was completed in March 2004 by 122 Wharton School students; roughly 90 percent were undergraduates. Respondents considered a written survey instrument and were assured of the confidentiality of their responses.¹²

3.1. Recollection Bias Measures

Although actual risk beliefs are of interest and are reported, our principal concern is recollection bias. We assessed this bias with respect to ter-

11. Respondents' policy preferences may have influenced their stated risk beliefs. If they do not wish to allow new curbs on civil liberties, they convince themselves or state that current risk levels are no higher than those before the salient event. They adjust remembered ex ante levels upward to achieve this equality. However, the survey questions were not framed in a manner that would foster such an effect.

12. The law students were enrolled in Viscusi's classes "Analytical Methods for Lawyers" in 2002 and "Economics of Regulation and Antitrust" in 2003. The Kennedy School students were masters and doctoral students in Zeckhauser's elective course "Analytic Frameworks for Policy." The Wharton students were in Kent Smetters' course "Risk Management and Insurance."

Table 1. Percentage Change in Terrorist Risk Estimates after September 11

	2002 Sample ^a (N = 94)	2003 Sample ^b (N = 117)	2004 Sample ^c (N = 122)	Pooled Sample (N = 333)
Risk is now higher	43	54	33	43
Risk is same as before September 11	17	24	34	26
Risk is now lower	40	22	33	31

^a Harvard Law School, where 95 people were surveyed but there were only 94 responses to this question.

^b 56 at Harvard Law School, 61 at Harvard Kennedy School.

^c University of Pennsylvania Wharton School.

rorism by asking the following question:¹³

Take yourself back to the World Trade Center disaster. Do you believe that the risk of a terrorist attack over the next year on an airplane is higher or lower than you thought it was before the September 11th disaster?

Higher ____ The Same ____ Lower ____

The results, reported in Table 1, are startling. A substantial fraction of respondents—57 percent in the pooled sample—believed that the risk was lower or the same as what they remember they thought it was before the attack.¹⁴ In regressions reported below, the variable September 11 recollection bias will take on a value of one if the respondent believed the risks were the same or lower than they were before the attack and a value of zero otherwise.

To see if recollection bias was peculiar to the terrorism case or in fact generalized to less emotionally charged issues, we asked our 2003 and 2004 respondents about their updating of space shuttle risks after the 2003 *Columbia* disaster. The question posited that no new precautions were taken. An external indicator that the perceived risk had risen

13. As we observe in Viscusi and Zeckhauser (2003, p. 108), until September 11 there had been very few fatalities resulting from hijackings of U.S. aircraft since 1970: "In 1973, 30 people died as a result of a terrorist action while a Pan Am flight from Rome was loading at the gate. The 1985 hijacking of a TWA flight from Athens resulted in the death of one U.S. serviceman. In 1986 a hijacking of a Pan Am flight from Karachi, Pakistan led to 16 fatalities. The largest death toll from a pre-9/11 hijacking was the 43 people who were killed as a result of the hijacked Pacific Southwest flight by a former employee."

14. In the 2003 survey, fewer thought the risk was lower than in the 2002 survey. There was no terrorist attack in the United States between the two survey times, which should have lowered risk estimates. On the other hand, the war in Iraq probably raised them.

is that NASA suspended all shuttle flights until the safety issues were resolved. Here 55 percent thought their risk estimate would be unchanged; 6 percent thought it would be lower, and 39 percent higher. The 61 percent of the people who estimated risks to be unchanged or lower show recollection bias.

3.2. Determinants of Terrorism Risk Assessments

The respondents also considered a series of questions asking them to assess the likelihoods of various death tolls from terrorism attacks. Their answers were taken to be a measure of individuals' probabilistic risk beliefs. Below, we analyze how recollection bias regarding the World Trade Center attack influences risk beliefs and whether both recollection bias and risk beliefs affect attitudes toward protective measures against terrorism. The questions on risk beliefs are as follows:

Based on some estimates, the September 11, 2001 disaster led to 266 deaths in the planes and 2,717 deaths at the World Trade Center. The total number of deaths was about 3,000. Below is a series of questions about the number of deaths on the ground and to passengers in the U.S. in the next 12 months because of attacks by foreign terrorists on airplanes.

- a. Think of the best outcome in which the number of terrorism deaths could be low. Suppose there is only one chance in 20 that the number of terrorism deaths could be at this low level or below. What is your estimate of this low-end death toll in the U.S.? _____
- b. Now think of the worst outcome. Suppose there is only one chance in 20 that the number of terrorism deaths could be this high. What is your estimate of this high-end death toll in the U.S.? _____
- c. Your best estimate of the actual death toll will be somewhere between your estimate of the low-end death toll and your estimate of the worst death toll. What is your best estimate of the expected number of terrorism deaths over the next 12 months in the U.S.? _____
- d. If we were to average the best estimate of your classmates of the terrorism death toll in the U.S., that is, their answers to part c, what number do you think we will get? _____
- e. If a plane with 300 passengers were shot down by terrorists at the London Heathrow Airport, what would your best estimate be of the number of deaths on the ground and to passengers in the U.S. in the following 12 months because of attacks by non-U.S. citizens on airplanes? Do not count any passengers killed in this attack. _____

A second version of the survey, given to different respondents, was identical except that instead of the risks arising solely because of attacks

Table 2. Terrorism Risk Estimates of Fatalities in Next 12 Months: 2003 and 2004 Survey Results Combined

	Median	Mean	S.E.M.
Airplane risk:			
Lower bound (5th percentile)	1.5	95.95	33.71
Upper bound ^a (95th percentile)	2000	23,768.35	12,658.61
Best estimates (50th percentile)	100	451.59	99.98
Estimates of classmates	200	706.45	116.02
Post-London attack estimate	100	408.39	77.50
Multiple risks:			
Lower bound (5th percentile)	5	64.31	26.14
Upper bound ^a (95th percentile)	4000	33,419.97	12,657.38
Best estimates (50th percentile)	100	341.75	65.96
Estimates of classmates	200	1,306.87	819.23
Post-Los Angeles attack estimate	100	255.08	45.17
Pooled survey estimate:			
Lower bound (5th percentile)	5	79.73	21.18
Upper bound ^a (95th percentile)	3000	28,762.01	8,942.09
Best estimates (50th percentile)	100	395.52	59.39
Estimates of classmates	200	1,012.94	421.62
Postattack estimate	100	328.80	44.22

Note. S.E.M. = standard error of the mean.

^a Eight observations with estimates of over 1 million were excluded from these statistics, as was another observation with incomplete information.

by non-U.S. citizens on airplanes, it also included “violent terrorist attacks by non-U.S. citizens more generally, e.g., crashed airplanes, bombs, and bullets.” This expanded set of outcomes was designed to test for possible embedding effects. Embedding occurs when two or more risks are combined in a class, and the total estimated risk for the class falls far below the sum of the individual risks.¹⁵

Table 2 reports the responses to these questions in 2003 and 2004 for respondents considering airplane risks alone or multiple risks from terrorism (similar results for 2002 for several of these variables appear in Viscusi and Zeckhauser [2003]). The responses for airplane risks and multiple risks are very similar, as the subsequent regression results will also indicate. While some of the multiple-risk responses are greater than for airplane risks alone, others are not, and none of these observed differences are statistically significant. There is no evidence that people regard the hazards of multiple risks as being noticeably greater than the

15. As long as the risks are small, the possibility that two or more may occur in the same period does not reduce the total probability much. Kahneman and Knetsch (1992) discuss the embedding phenomenon.

risks posed by airplane risks alone. This failure to distinguish between one risk and a class of risks in which it is included is a classic embedding effect. Given its appearance, we pool the results for airplane risks and multiple risks and consider the pooled estimates at the bottom of the table.

The first set of results pertains to the lower-bound estimate, or what the respondents were told was the 5th percentile of their risk distribution for terrorism fatalities. The mean value for this is 80. The second set of values pertains to the 95th percentile of that distribution, which is the high-end risk estimate (the potential for alarmist responses to such a risk information situation is discussed in Viscusi [1997, 1998]). For the 2003 and 2004 sample, this mean value is 28,762.

The third estimate presented is the respondent's best estimate of the risk. While the survey did not indicate whether the "best estimate" should be the mean value or the median, the previous survey questions were in terms of percentiles of the distribution, which would be consistent with assessing the median. The best-estimate responses were much closer to the low end than were some average of high and low estimates, which suggests that respondents did in fact think in terms of the median, which we will assume for the following discussion.¹⁶ The mean of these median terrorism death tolls was 396. Even on a logarithmic basis, this value is closer to the lower-bound estimate than to the upper-bound worst case. Individuals' loss estimate distributions appear to have a long right tail. Respondents' best estimate of their classmates' estimates of terrorism risk are higher than their own estimates of the risk; respondents believe others will assess the risk as being greater than they do.

Finally, the survey included questions about how the respondent's best estimate of the risk would change were an airplane to be shot down by terrorists, killing 300. The scenario reported in the top section of Table 2 was a terrorist attack at London Heathrow Airport, and the scenario in the middle section was an attack on Los Angeles Airport. Our hypothesis was that an attack in the United States would have greater impact; however, the risk estimates after attacks at the two airports are quite similar, with identical median values and somewhat higher means for a London attack. More surprising, the numbers hardly rise above the individuals' previous best estimates of terrorism risk. This

16. Indeed, since the upper-bound estimate is only the 95th percentile and the estimates were substantially skewed to the right, the mean estimate should be well above the average of the 5th and 95th percentiles.

Table 3. Tobit Results of Best Estimate of Terrorism Risk

	Without Demographics		With Demographics	
	Linear Equation	Log Equation ^a	Linear Equation	Log Equation ^a
Lower bound	1.8055* (.1707)	.4873* (.0456)	1.7980* (.1706)	.4858* (.0474)
Upper bound	.0004 (.0004)	.2689* (.0438)	.0004 (.0004)	.2569* (.0446)
Multiple risks	26.2666 (97.4742)	-.0495 (.1765)	32.3010 (97.4542)	-.0420 (.1761)
September 11 recollection bias	-1.8486 (100.1616)	-.2000 (.1837)	-15.8305 (101.5873)	-.2266 (.1855)
2003	-70.0136 (123.7554)	-.1982 (.2239)	-187.7072 (155.5459)	-.5231 ⁺ (.2787)
2004 (Wharton)	71.9063 (121.6734)	.2361 (.2251)	98.9954 (126.4293)	.2959 (.2318)
Female	. . .	. . .	260.0589 ⁺ (105.6914)	.2694 (.1915)
Nonwhite	. . .	. . .	-27.7425 (111.8196)	-.2937 (.2087)
U.S. citizen	. . .	. . .	-101.5829 (133.9680)	-.3030 (.2423)
Kennedy School	. . .	. . .	190.2480 (170.4961)	.5055 (.3073)

Note. Each equation also contains a constant term. Standard errors are in parentheses.

^a The natural logarithm is taken of the best estimate, lower bound, and upper bound, where all are +1.

⁺ Statistically significant at the 10% level, two-tailed test.

* Statistically significant at the 5% level, two-tailed test.

hypothetical attack barely budges risk estimates. Such inertia could reflect resistant learning, or perhaps people do not change risk perceptions much in response to hypothetical events.¹⁷

To analyze the determinants of the best estimates of terrorism risks, we prepared Tobit estimates regressing those estimates on other aspects of risk beliefs, as shown in Table 3. These Tobit results jointly estimate the influence of whether the respondent assesses a nonzero level of terrorism deaths and an estimate of the scale of these deaths. The first set of results includes only risk variables, survey year dummy variables, and survey group variables; the second set also includes important demo-

17. A weakness of many hindsight bias studies is that they ask questions about hypothetical happenings.

graphic variables. For each set, we estimate both a linear equation in which the dependent variable within the Tobit framework is the subject's best estimate of the terrorism risk and an equation in which the dependent variable is the logarithm of the subject's best estimate.

The lower- and upper-bound estimates of the terrorism risk compose the first pair of explanatory variables. The lower-bound variable is always a significant predictor of the best estimate of the terrorism death toll. The upper-bound estimate is significant only after taking the log transformation of it and the terrorism death toll, thus muting the influence of the high-end outliers. The magnitudes of the lower-bound coefficients are always greater than the upper-bound coefficients as well. These results are consistent with the sample statistics in Table 2, which indicates a much closer relationship of the best estimate to the lower-bound estimate than to the upper-bound estimate.

The third variable in the equation addressed the potential embeddedness effect. It is a dummy variable for the version of the survey in which respondents were asked to assess the terrorism risk deaths from multiple sources of risk, not just risks to passengers on airplanes. The multiple-risk coefficients fail the usual tests of statistical significance in every instance. Quite simply, respondents do not take into account the additional risk when multiple risks are considered instead of a single source of risk. This result reinforces findings about embeddedness.

The September 11 recollection variable takes the value of one when the *ex ante* risk estimate at least equals the *ex post* estimate. This variable is never statistically significant. Respondents who indicate that they believe the risk is lower or the same as they thought it was before the September 11 attack do not have lower values for their best estimate of the terrorism risk than do those who indicate that they have raised their estimates.¹⁸ This provides strong evidence that these people are not resistant learners, as their postattack risk beliefs are no lower than for respondents who recognize that they have altered their risk assessments.

Table 3 also includes dummy variables for the different survey years, 2003 and 2004; the 2004 dummy variable also serves as a Wharton School dummy variable, as all respondents in that year were at Wharton. There is also a dummy variable for the Kennedy School in 2003. The

18. This result also suggests that the people exhibiting September 11 recollection bias did not misinterpret the question of whether they thought the actual risks were the same or lower than their perceived risks before September 11.

omitted dummy variable categories are 2002 and Harvard Law School. There are no significant year or locale effects.¹⁹

4. SUPPORT FOR ANTITERRORISM RISK POLICIES

4.1. Use of Profiling for Passenger Screening

The principal antiterrorism policy that we analyze is the profiling of airline passengers on the basis of demographic characteristics to identify potential terrorist threats on airplanes. In particular, we explore the extent to which people are willing to trade off additional time waiting in line for the infringement on civil liberties that such profiling might represent. The basic survey question mentioned three different periods of waiting time—10 minutes, 30 minutes, and 60 minutes. In each instance, the respondent could avoid this additional wait in line if a policy were adopted to screen passengers on the basis of demographic profiling.

In one version of the survey the respondent is told: “*You would not* be singled out for such a search based on terrorist risk profiling.” In the second version of the survey, the respondent is told: “*You would be* singled out for such searches based on terrorist risk profiling.” The exact wording of the full question for the version in which the subject would not be singled out for searches follows:

One way of reducing terrorism risks to plane flights is better screening of passengers. The FBI has developed a profile of the chances that a passenger is a terrorist, taking into account the person’s age, race, gender, national origin, appearance, and baggage. Airlines either could screen all passengers, leading to additional delays in line, or they could screen passengers based on the terrorist risk profiling. Targeted screening would reduce the terrorist risk by as much as random searches, but would involve time delays for passengers. People who are singled out based on the terrorist risk profiles will have to undergo an extra 10 minutes of searches. *You would not* be singled out for such racial profiling.

- a. Would you favor terrorist risk profiling if the alternative was for you to wait in line an extra 10 minutes so that all passengers could be screened randomly?
Yes ___ No ___
- b. Would you favor terrorist risk profiling if the alternative was for you to wait in line an extra 30 minutes so that all passengers could be screened randomly?
Yes ___ No ___

19. Since there were no terrorist attacks in the United States between the survey dates—a pleasant surprise—we might have expected a decrease in risk estimates. The change was in the right direction but far from significant.

- c. Would you favor terrorist risk profiling if the alternative was for you to wait in line an extra 60 minutes so that all passengers could be screened randomly?
Yes ___ No ___

Note that respondents are being asked to trade off convenience and profiling, not safety versus profiling. If freedom from profiling is regarded as a civil liberty, and if civil liberties are regarded as a fundamental value, not to be compromised unless some equivalently fundamental value is at stake, individuals should prefer even the 60-minute wait to airport profiling.²⁰ The profiling here reduces inconvenience and does not increase safety, as the alternative is to screen everyone. However, the profiling may be more acceptable because it takes place in a safety-related context. Still, we found it surprising that three-fourths of our respondents would allow profiling to avoid imposing a 60-minute wait on air travelers and that the numbers favoring profiling doubled when the wait increased from 10 to 60 minutes.

Table 4 summarizes the willingness of respondents to support terrorist screening policies on the basis of the length of time waiting in line and on whether the respondent would be affected by this screening policy. In every instance in this table, the support for targeting increases as the time waiting in line rises. Indeed, nearly two-fifths of the respondents change from opposing to favoring profiling as the wait increases. The civil liberties issue proves important but is not a trumping consideration.²¹

Somewhat surprisingly, respondents were not generally more willing to support screening that targets other people rather than themselves. In the 2002–2004 pooled survey results, for example, when the incremental cost of avoiding targeting was a 10-minute wait, 40 percent of the passengers supported targeting if the screening affected only others,

20. A more extreme version of this example would highlight the drawbacks to strict lexicographic preferences. We expect that even staunch civil libertarians might accept profiling rather than a 20-hour wait in line that yielded equivalent safety. By offering three scenarios—10-, 30-, and 60-minute waits—we may have inadvertently created a frame that made it acceptable to profile to avoid longer waits.

21. The authors recognize that by giving respondents three different times, they were to some extent framing a situation in which different answers were reasonable; neither civil liberties nor safety would be trumping concerns. To explore this possibility, in June 2004 the survey question was run on a sample of 17 judges in a cognitive-interview setting. That setting allowed respondents to discuss the question and its interpretation. The judges indicated that their preferences were not affected by the offered ranges. Indeed, 15 of the 17 judges favored profiling even if it saved no time in line, and two of the judges opposed profiling even if the alternative was to wait an hour in line.

Table 4. Attitudes toward Use of Terrorism Risk Profiles (Percentage Favoring Risk Profiling)

Delay in Line Due to Screening Time	Screening Affects Only Others	Screening Affects Respondents
2002 survey:		
10 minutes	40.91	41.46
30 minutes	59.09	58.54
60 minutes	77.27	65.85
2003 survey:		
10 minutes	35.19	48.89
30 minutes	59.26	68.89
60 minutes	75.93	80.00
2004 survey:		
10 minutes	45.24	47.83
30 minutes	61.90	69.57
60 minutes	66.67	76.09
Pooled results:		
10 minutes	40.00	46.21
30 minutes	60.00	65.91
60 minutes	73.57	74.24

while 46 percent supported screening if it affected the respondents themselves. By the time the additional wait in line reaches an hour, about three-fourths of all respondents supported screening in each instance. This pattern indicates that concern for civil liberties is far broader than mere self-interest. Indeed, some respondents may feel more comfortable with profiling—which brings efficiency advantages—if they would be a target, since that defuses the issue of unjust imposition.

The sample used for this survey is far from a random sample of the U.S. population. If anything, given their fields of study, our respondents are likely to be less supportive of the practice of profiling than the general population. Harvard Law School students and Kennedy School of Government students, whose fields are government and public policy, are likely to be particularly sensitive to civil liberties issues and the importance of preserving civil liberties from unnecessary infringement. One might suspect that Wharton business school students would be less sensitive to civil liberties, but that in fact was not the case. We expect that all these groups of students take airplanes much more frequently than the average American.

Table 5 presents the probit estimates of the probability that an individual with particular characteristics will favor targeting passengers for screening on the basis of profiling. The first set of estimates in Table

Table 5. Probit Estimates of the Probability of Favoring Targeting Passengers for Screening

	Pooled Estimates	Waiting Time (in Minutes)		
		10	30	60
Waiting time	.0064* (.0009)	. . .	. . .	. . .
Respondents targeted for screening	.0607+ (.0361)	.0687 (.0619)	.0812 (.0606)	.0230 (.0550)
September 11 recollection bias	-.0892* (.0368)	-.0547 (.0636)	-.1008 (.0616)	-.0993+ (.0551)
Best risk estimate	.00004+ (.00002)	.00004 (.00004)	.00005 (.00004)	.00003 (.00004)
2003	.0400 (.0544)	-.0076 (.0942)	.0847 (.0900)	.0368 (.0825)
2004 (Wharton)	.0611 (.0457)	.0952 (.0805)	.0925 (.0756)	-.0098 (.0702)
Female	-.0548 (.0396)	-.0187 (.0676)	-.0732 (.0669)	-.0664 (.0613)
Nonwhite	-.0918* (.0419)	-.1168+ (.0694)	-.0941 (.0709)	-.0492 (.0641)
U.S. citizen	.0023 (.0510)	.0065 (.0872)	.0245 (.0863)	-.0234 (.0762)

Note. All regressions include a variable for whether the respondent is a Kennedy School student, which is not statistically significant, as well as a constant term. Coefficients are transformed to reflect marginal probabilities. $N = 799$ for the pooled estimates. Standard errors are in parentheses.

+ Statistically significant at the 10% level, two-tailed test.

* Statistically significant at the 5% level, two-tailed test.

5 pools the various waiting time observations and includes a variable for the length of waiting time. The length of the wait is statistically significant; the coefficient implies that for every additional 10 minutes of waiting in line, a respondent is 6 percent more likely to favor profiling passengers for screening purposes.

Respondents who exhibit September 11 recollection bias are 9 percent less likely to favor targeting of passengers in the pooled estimates and 10 percent less likely to favor targeting in the 60-minute wait situation. Thus, individuals who do not recognize an increase in terrorist risk since September 11 are less likely to favor profiling. This is precisely the pattern predicted by recollection choice bias.

This negative linkage between September 11 recollection bias and willingness to use passenger profiling could be the result of motivated reasoning. The evidence is consistent with that hypothesis (see Figure

1).²² However, two aspects of the results suggest caution in concluding that respondents consciously or subconsciously justified their policy positions with faulty risk assessments. First, a more straightforward way to reject curbs on civil liberties would be to produce low risk estimates. However, people who exhibit recollection bias did not have lower assessments of current terrorism risk than did people who did not exhibit recollection bias. Second, even in the context of space shuttle launches, where the issues are much less charged, many people exhibit recollection bias. If recollection bias is unaffected by policy preferences, it will still be correlated with a preference for sticking with the policies preferred in the past, the pattern we observed.

It is interesting to note that the best risk estimate for an individual has a statistically significant effect on support for passenger screening, and the effect is small. A small or zero effect on risk estimates would be expected if people suffer from probability neglect. As hypothesized by Sunstein (2003), people facing catastrophic terrorism risks may focus more on the severity of the outcome than on the probability that it will occur. The expectation of an additional 1,000 deaths from future terrorism attacks makes respondents only 4 percent more likely to favor targeting of passengers for screening. It is not the level of risk estimates that drives support for antiterrorism policies but rather the perception that we are on a steeper portion of the risk/civil liberties trade-off curve since September 11. Once again, this result is consistent with the recollection choice bias framework, as outlined in Section 2. People who do not believe that there has been an upward shift in the risk level are less supportive of more intrusive passenger screening regardless of the value of their best risk estimate.

In the pooled and 10-minute regressions, nonwhites are less supportive of profiling. The other demographic variables are never statistically significant. Other variables that are never statistically significant are the 2003 and 2004 dummy variables and the variable indicating whether the respondent will or will not be targeted for screening. As in Table 4, there is no evidence that self-interest influences individuals' trade-off between civil liberties and waiting time.

An important finding that emerges from these results is consistent with our hypothesis that an individual displaying recollection bias after

22. Values and probability assessments are often linked. For example, those who think capital punishment is immoral also tend to think it is ineffective. Eldar Shafir (personal communication, May 25, 2004).

Table 6. Attitude toward Surveillance of Mail, E-Mail, and Phones (Percentage of Respondents Supporting Policy)

	2002 Survey	2003 Survey	2004 Survey	Pooled Results
Full sample	36.1	27.6	38.4	33.7
Respondents answering yes to profiling	43.3	32.5	41.0	38.4
Respondents answering no to profiling	17.4	9.5	32.0	20.3
Support both profiling and surveillance	31.3	25.5	29.1	28.5
Oppose both profiling and surveillance	22.9	19.4	19.8	20.6
Support profiling but not surveillance	41.0	53.1	41.9	45.7
Support surveillance but not profiling	4.8	2.0	9.3	5.2
September 11 recollection bias	31.9	21.3	36.4	30.2
No September 11 recollection bias	41.7	33.3	41.9	38.4

Note. The profiling question pertains to respondents who answered yes to profiling to avoid a 60-minute wait. To be included in any calculations appearing in this table, respondents had to be consistent; for example, if they favored profiling for a 10-minute wait, they should favor profiling for a 30- or 60-minute wait.

a catastrophic event is less likely to favor an expensive policy shift that makes that catastrophe less likely in the future.

4.2. Attitudes toward Surveillance Measures

Although the main focus of the survey was on targeting passengers for airplane screening, the survey also asked about the respondents' support for various kinds of surveillance measures. In particular, each respondent in all surveys considered the following question:

Would you support policies that made it easier for legal authorities to read mail, e-mail, or tap phones without a person's knowledge so long as it was intended to prevent terrorism? Yes ___ No ___

Table 6 reports the responses to these questions for the 2002 survey, the 2003 surveys, the 2004 survey, and the pooled results. Every analysis of the data showed somewhat greater support for these forms of sur-

veillance in 2002 than in 2003.²³ As the September 11 attack receded into memory and no new attacks took place in the United States, support for antiterrorism surveillance may have diminished.²⁴ However, responses for the 2004 full sample and for those who say yes to profiling at airports are very similar to the answers in 2002.

The pooled results from the surveys for the full sample indicate that just over one-third of all respondents support surveillance of mail, e-mail, and telephones. Of those who supported passenger profiling to avoid a 60-minute wait, 38 percent support surveillance, while among respondents who opposed such profiling, only 20 percent support surveillance measures. This difference is significant at the 1 percent level ($t = 2.88$). As expected, respondents' attitudes toward civil liberties are consistent in the surveillance and profiling scenarios. Those who showed recollection bias were less likely to support surveillance (30.2 versus 38.4 percent for those not showing bias). On a one-tailed test, this difference is significant at the .052 level.

Respondents were much less likely to support surveillance efforts than profiling. Nevertheless, there is a positive correlation between support for surveillance policies and support for passenger profiling. Of the sample, 29 percent support both surveillance and passenger screening (based on the 60-minute wait), and 21 percent of the sample oppose both. However, among respondents whose attitudes differ for the two forms of civil liberties infringement, far more support passenger profiling—46 percent versus 5 percent.

5. QUALIFICATIONS AND FUTURE ENHANCEMENTS²⁵

Clearly, it would have been desirable to have started our research asking individuals their terror risk assessments before September 11. Then we could tell how much their memories shaded their prior beliefs using a within-subjects format or how their beliefs related to the actual prior beliefs of matched others in an across-subjects format. Alas, it would require a supernatural degree of foresight to undertake such before-the-

23. Conceivably this difference could arise because of a difference in attitudes between law and Kennedy School of Government students, since the latter were included only in 2003. In fact, a Kennedy School of Government dummy variable proved insignificant.

24. Views may have also changed about government policies, such as the Patriot Act, which was passed in the interim. The war in Iraq may have also altered attitudes.

25. We thank the editor and the referee for suggesting that we include the material in this section.

fact surveys; the study of unique rare bad events will generally remain retrospective. Recollection bias, with all its challenges, is likely to be a component of such studies.

Our study had no controls on a variety of variables that would have provided firmer insights into behavior. For example, a quantitatively meaningful attitude scale for general preferences for civil liberties would be useful as a control on respondents' attitudes.²⁶ However, the best measure of civil liberties preferences is that revealed by attitudes toward actual policy choices, and the results in this paper present a first step at developing such trade-off values.

Our sample consisted of 333 individuals. Though this sample size adequately identifies a number of important relationships, it may be too small to identify others. Say that the sample consists of one-third As and two-thirds Bs, where As avoid recollection bias and Bs exhibit it. Say that we know that As are .5 likely to favor a policy, and we want to find out whether Bs are different. If the true p for Bs is .6, the likelihood that we reject the null of no difference, that is, the power of the test, is 84.9 percent.

Most policy discussion in this arena addresses the trade-off between civil liberties and safety. However, the airport profiling scenario in our study trades off civil liberties against convenience, not safety. Using actual trade-offs with safety would have provided additional insight, particularly if our interest is fundamental civil liberties.

Fortunately, other inferences are available on the civil liberties/safety trade-off. If the expected trade-off between convenience and safety is relatively constant across individuals, then trade-offs between civil liberties and convenience translate to those between civil liberties and safety. Table 6, which reports acceptance of surveillance, is reassuring in this regard. Respondents who were willing to permit profiling for convenience were dramatically more willing to accept surveillance. Thus, the trade-off between civil liberties and convenience predicted the trade-off between civil liberties and safety, at least in the context of surveillance. Moreover, in 2002 and 2003, those showing September 11 recollection bias were less likely to support surveillance, which implies a direct trade-off between safety and civil liberties.

We have sketched a broad set of relationships and drawn inferences

26. We could also directly measure personality differences of our subjects, say, to assess a relationship with recollection bias. Neither past studies of hindsight bias nor the extremely weak correlation between September 11 and space shuttle recollection bias suggests that personality traits would be strong predictors.

on the basis of survey data. In some contexts, more direct measures of attitudes and preferences would provide firmer insights. In addition, since we merely employed a survey, we learn only what people say they would do, not what they actually would do. However, since the questions involve public goods and public bads over which individuals do not have control, for example, risk profiling at airports, we cannot delegate real-world decisions to ordinary citizens and draw inferences from their actual behaviors.

6. CONCLUSION

The September 11 attack dramatically transformed the recognized scale of terrorism risks facing U.S. citizens. As a result, the policy options with respect to terrorism risk and civil liberties deteriorated substantially. Assuming preferences between safety and civil liberties of the usual shape and that the productivity of antiterror measures increases with the level of risk, optimal antiterrorism measures after September 11 will be stricter than those before. Hence, civil liberties will be reduced, but terror risks are likely to be higher as well.

Among people with common preferences, assuming rational choice, levels of support for antiterrorism efforts should be positively related to perceived risks. However, our results indicate that the perceived risk level is not a predominant factor affecting support. Indeed, it is never significant in our empirical analyses.

Quite remarkably given the magnitude and surprise of the September 11 attack, 57 percent of respondents exhibited recollection bias: they stated that the risks were no greater post- than preattack. The current risk estimates of those exhibiting recollection bias were no lower than the estimates of those who did not exhibit such bias; the bias was not due to a failure to perceive the risks after September 11 as others do. Rather, it was a failure to recognize that they had updated their risk beliefs, just as others had.

A striking finding is that individuals who suffer recollection bias are more reluctant than their peers to accept racial profiling or surveillance policies. We posit that these individuals were internal to or invested in past choices; they approved of the civil liberties level relative to terrorism in effect prior to September 11. Since they do not perceive changed risk levels, they see no need to change their favored policy.




Figure A1. Hindsight bias

APPENDIX

We illustrate the differences between hindsight and recollection bias with Figures A1 and A2, which show the two phenomena in schematic form. In the figures, A represents the actual forward-looking estimate of the risk, H represents a hypothetical pre-event estimate judged from a postevent position, and R represents the recollection of a pre-event estimate. Judgments by individuals I and J are indicated by subscripts. The darkened arrows represent the magnitudes of the two biases.

Figure A1 compares I's postevent conjecture of what the risk was pre-event, which is a hypothetical assessment, and the actual pre-event estimate by person J. Hindsight bias is the tendency for the postevent estimate to be far higher. The magnitude of the bias is indicated by the darkened, double-headed arrow. A famed experimental test of hindsight (see the studies reported in Sunstein et al. [2002]) had respondents A consider a fictional railroad situation. They were then asked to assess the accident risk. Respondents B then consider the same situation, but after an accident has occurred. They are then asked to assess what they thought the risk was given the information available before the accident. The actual preaccident risk assessments by respondents A in such experiments are




Figure A2. Recollection bias

well below the hypothetical pre-event risk assessments by respondents B, which indicates hindsight bias.

The test for recollection bias compares different components. It involves neither hypothetical judgments nor those by individual J. The test first examines whether I's forward-looking risk beliefs are in line with external estimates; if so, I does not merely have low estimates. In Figure A2, the shaded boxes indicate the estimates by the external environment, including market and government estimates. The postevent box surrounds the postevent circle, showing that the individual's current estimates are in line with the external benchmark. If that test is passed, I's recollection of his pre-event risk estimate is subtracted from I's current risk estimate. If this difference is well below the external measure of the increase in risk between the preaccident and postaccident situation, then I is said to exhibit recollection bias. While railway safety experts will assess the future risk as being greater than they thought it was before, perhaps because they learned that the railroad's safety practices were less effective than they thought, respondent I, who exhibits recollection bias, believes that the prospective risk is that same as he thought it was before the accident.

Figure A2 shows the case where the perceived risk is the same before and after the low-probability event occurs. Recollection bias is indicated as the dif-

ference between the increase in risk according to external evidence and the respondent's report of perceived increase in risk. Here again, the darkened, double-headed arrow indicates the magnitude of the bias.

REFERENCES

- Brown, Jeffrey R., J. David Cummins, Christopher M. Lewis, and Ran Wei. 2003. An Empirical Analysis of the Economic Impact of Federal Terrorism Reinsurance. Working paper. University of Pennsylvania, Wharton Financial Institutions Center, Philadelphia. November 4.
- Cummins, J. David, and Christopher M. Lewis. 2003. Catastrophic Events, Parameter Uncertainty and the Breakdown of Implicit Long-Term Contracting: The Case of Terrorism Insurance. *Journal of Risk and Uncertainty* 26: 153–78.
- Doherty, Neil A., Joan Lamm-Tennant, and Laura T. Starks. 2003. Insuring September 11th: Market Recovery and Transparency. *Journal of Risk and Uncertainty* 26:179–99.
- Fischhoff, Baruch. 1975. Recollection $\neq$ Foresight: The Effect of Outcome Knowledge on Judgment under Uncertainty. *Journal of Experimental Psychology, Human Perceptions and Performance* 1:288–99.
- Hastie, Reid, David A. Schkade, and John W. Payne. 1999. Juror Judgments in Civil Cases: Hindsight Effects on Judgments of Liability for Punitive Damages. *Law and Human Behavior* 23:597–614.
- Jaffee, Dwight, and Thomas Russell. 2002. Extreme Events and the Market for Terrorist Insurance. Paper presented at National Bureau of Economic Research Insurance Conference. February.
- Kahneman, Daniel, and Jack L. Knetsch. 1992. Valuing Public Goods: The Purchase of Moral Satisfaction. *Journal of Environmental Economics and Management* 22:57–70.
- Kelman, Mark, David E. Fallas, and Hilary Folger. 1998. Decomposing Hindsight Bias. *Journal of Risk and Uncertainty* 16:251–69.
- Keohane, Nathaniel D., and Richard J. Zeckhauser. 2003. The Ecology of Terror Defense. *Journal of Risk and Uncertainty* 26:201–29.
- Kunreuther, Howard, and Geoffrey Heal. 2003. Interdependent Security. *Journal of Risk and Uncertainty* 26:231–49.
- Rachlinski, Jeffrey J. 1998. A Positive Psychological Theory of Judging in Hindsight. *University of Chicago Law Review* 65:571–625.
- Smetters, Kent. 2004. Insuring against Terrorism: The Policy Challenge. Working paper. University of Pennsylvania, Wharton School, Philadelphia, and National Bureau of Economic Research, Cambridge, Mass. February 2.
- Sunstein, Cass R. 2003. Terrorism and Probability Neglect. *Journal of Risk and Uncertainty* 26:121–36.

- Sunstein, Cass R., Reid Hastie, John Payne, David Schkade, and W. Kip Viscusi. 2002. *Punitive Damages: How Juries Decide*. Chicago: University of Chicago Press.
- Tversky, A., and D. Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185:1124–31.
- Viscusi, W. K. 1997. Alarmist Decisions with Divergent Risk Information. *Economic Journal* 107:1657–70.
- . 1998. *Rational Risk Policy: The 1996 Anne Ryde Memorial Lectures*. Oxford: Clarendon Press.
- . 1999. How Do Judges Think about Risk? *American Law and Economics Review* 1:26–62.
- Viscusi, W. Kip, and Richard J. Zeckhauser. 2003. Sacrificing Civil Liberties to Reduce Terrorism Risks. *Journal of Risk and Uncertainty* 26:99–120.