The Journal of Risk and Uncertainty, 26:2/3; 99–120, 2003

Sacrificing Civil Liberties to Reduce Terrorism Risks

W. KIP VISCUSI Kip@law.harvard.edu
Cogan Professor of Law & Economics, Harvard Law School, 1575 Massachusetts Avenue, Hauser 302, Cambridge, MA 02138, USA

RICHARD J. ZECKHAUSER
Ramsey Professor of Political Economy, Kennedy School of Government, Harvard University, Cambridge, MA 02138, USA; Visiting Professor, Harvard Business School, USA

Abstract

Our survey results demonstrate that targeted screening of airline passengers raises conflicting concerns of efficiency and equity. Support for profiling increases if there is a substantial reduction in avoided delays to other passengers. The time cost and benefit components of targeting affect support for targeted screening in an efficiency-oriented manner. Nonwhite respondents are more reluctant than whites to support targeting or to be targeted. Terrorism risk assessments are highly diffuse, reflecting considerable risk ambiguity. People fear highly severe worst case terrorism outcomes, but their best estimates of the risk are more closely related to their lower bound estimates than their upper bound estimates. Anomalies evident in other risk perception contexts, such as hindsight biases and embeddedness effects, are particularly evident for terrorism risk beliefs.

Keywords: terrorism, risk beliefs, ambiguity

JEL Classification: A10, D8, J28, K00

1. Introduction

The 9/11/2001 terrorist attack on the United States dramatically affected the nation and the rest of the world. The attack stimulated a concerted effort to adopt measures that would reduce the risks in the future. Many of these precautions involve costs that are not financial, and some involve a reduction in civil liberties of various kinds. This article examines people's willingness to sacrifice civil liberties in an effort to reduce terrorism risks, and also explores aspects of individuals' terrorism risk perceptions that govern the character of their responses.

The scale of terrorism risks might seem to be small relative to the attention they command. While 3,000 deaths is clearly a catastrophic outcome, it is considerably smaller than many other mortality risks. More Americans are killed every month in automobile accidents and, according to public health officials, more than 100 times more Americans are killed each year by cigarette smoking. What makes terrorism deaths different from the deaths either from automobile accidents or cigarette smoking is that these risks contain a voluntary element. Consumers of these products obtain some valued attribute such as increased mobility or satisfaction from smoking, which compensates for the risks associated with the activity. There is no voluntary market transaction involved with risks of terrorism and no form of

compensatory benefit. Concern with terrorism risks is also stimulated by their character and media coverage.[1]

The deaths associated with 9/11 commanded public attention. The deaths were dramatic and clustered—a large number of deaths occurred in one place at once. The deaths were also accompanied by the destruction of two of the tallest buildings in the world, and took place in one of the media centers of the world, ensuring substantial coverage. The extraordinary publicity given to the attack and the accompanying losses produced the kind of risk that people are likely to severely misestimate in the future.

These terrorist attacks also forced people to completely rethink their risk beliefs. Just prior to the attack, it is highly unlikely that terrorist attacks of this magnitude were even among the possibilities that ordinary citizens contemplated. Thus, it is not a question of whether people assigned a low probability to this event. Rather, it is likely that this event was not even in the set of possible outcomes that people took into account when assessing the likely risks they faced. Much of the uncertainty created by the event is the realization that the events set that we previously thought was possible was incomplete and that the future may contain many other severely adverse events that are currently unanticipated.

Because of the apparent ease with which the hijackings took place, there has been a dramatic change in the precautions for air travel. These include much more rigorous passenger screening, which creates intrusions such as searches of people and their belongings, and which are paid for with ticket surcharges. These searches have also raised the issue of whether it is appropriate to target classes of passengers considered at high risk, e.g., Middle-Eastern-looking men, as opposed to undertaking searches randomly or generally. Whether targeting and other measures that compromise civil liberties are desirable depends in large part on our assessment of and valuation of the risks, and on our willingness to sacrifice civil liberties to reduce these risks. This article examines some survey results that begin to shed some light on this matter. The subjects were Harvard Law School students who were surveyed in the spring 2002, or roughly seven months after the 9/11 attack.[2] One would expect this sample to be more sensitive to civil liberties than the population at large.

Section 2 of the article outlines the basics of the civil liberties/terrorism risk tradeoff. The optimal balance between these competing concerns depends in large part on individual preferences, the subject of our survey results in this section. We found that people's attitudes toward increased screening involve both efficiency concerns and issues of equity. The length of time involved in the screening, an efficiency concern, is consequential, as is whether particular groups are going to be targeted for screening, an equity issue. People's attitude toward such targeting also depends on whether they belong to a group that has been singled out in the past, and perhaps whether they are currently accorded equal treatment in other contexts.

Terrorism risks are highly imprecise and difficult to predict. The dramatic differences between the treatments of terrorism risks and mortality risks for which we have a well established statistical basis is reflected in the structure of terrorism risk beliefs, a subject we explore in Section 3. Since there is little hard evidence to rely upon, people's assessments of terrorism risks are highly variable. Estimates of terrorism risks, e.g., number of deaths expected in a year, clearly produce a situation of considerable risk ambiguity. However,

this ambiguity is not randomly and symmetrically distributed around some mean risk value. Rather, there is a long tail—tiny probabilities of extremely catastrophic outcomes.

One would expect that terrorism risks would tend to frustrate rational decision making under uncertainty. We have little statistical basis on which to make a judgment regarding such risks, as 9/11 makes evident. Even the insurance industry, which is thoroughly acquainted with estimating unusual risks, has a hard time gauging the risk of terrorism losses. Terrorism presents a situation of tremendous uncertainty, or perhaps a better phrase is "ignorance," since many states of the world are not defined. Given this, attempts to estimate terrorism risks will fall prey to some of the more salient biases and anomalies that have been identified in the risk and uncertainty literature. Section 4 examines some of these biases, particularly those relating to hindsight effects and the substantial premium paid for zero-risk outcomes.

2. Sacrificing civil liberties to curtail terrorism

Searches of individuals and their luggage at airports are a sensible precaution to reduce terrorism risks. Why search policies became more stringent in the wake of the 9/11 disaster is clear on logical grounds. The main civil liberties issue is not whether searches should be undertaken, but whether a particular population group should be targeted in these searches. How much of a tradeoff is warranted is an empirical issue that must be explored directly.

2.1. The theoretical basis tradeoff

Figure 1 illustrates the character of the civil liberties—terrorism risk tradeoff decision. The curve xx shows the pre-9/11 perceived initial relationship between civil liberties and terrorism risks. In this diagram, civil liberties can be increased over quite a range without dramatically affecting expected terrorism losses. However, even along this curve of moderate

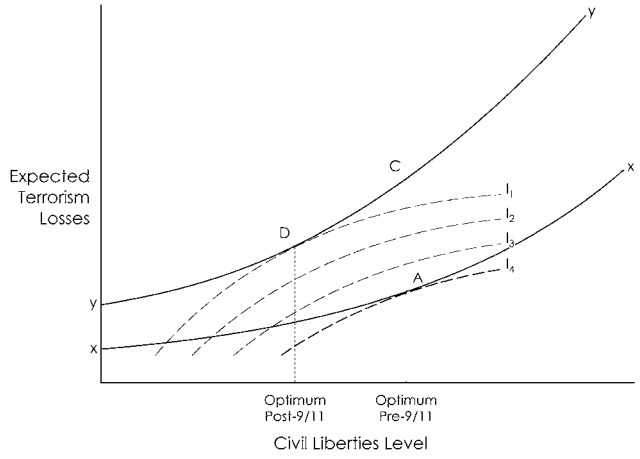

Figure 1. Terror risk versus civil liberties.

risk, had we increased civil liberties by not screening passengers for metal objects, expected terrorism losses would increase. In accord with normal economic assumptions, such losses increase at an increasing rate as civil liberties expand.

Individual preferences regarding airport searches incorporate two dimensions—civil liberties, a desirable attribute, and terrorism risks, an undesirable attribute. Therefore, utility increases as we move southeasterly. Moreover, indifference curves have the shape given by $I_1, I_2, \ldots$ where greater subscripts imply greater utility. In the situation before the 9/11 terrorism attack, when xx was the perceived opportunities frontier, the optimal choice was at point A. Civil liberties were high; the perceived terrorism risk was low. Indeed, prior to 9/11, there had not been a domestic plane hijacking in many years and never had a hijacked plane been crashed into a building. After the 9/11 terrorism attack, society's expectations of terrorism losses associated with any given level of civil liberties changed dramatically. Perceived risks rose for any level of civil liberties, and the marginal cost of civil liberties increased dramatically. The new perceived situation is indicated by curve yy, which lies above and is twisted counterclockwise from xx. Were we to maintain the pre-9/11 level of civil liberties, we would get to point C on curve yy, which involves a very high terrorism risk.

Given the character of individual preferences shown in Figure 1 the optimal decision is actually at point D. Reaching that outcome requires reducing civil liberties from its level at A. Yet, because our world has been recognized as much more dangerous since 9/11, expected terrorism losses at D are far greater than they were at A, though far less than they would be with no adjustment.

This figure shows why the optimal level of civil liberties is not at the highest possible value for this attribute. That would sacrifice too much on other valued attributes. The optimal level of civil liberties changes depending on the particular circumstances. For example, our society does not in general stop and inspect automobiles along roadways, but it would do so were there a serial killer on the loose. The willingness to sacrifice some civil liberties for other goals reflects the more general argument—articulated by Kaplow and Shavell (2002)—that many legal rights and liberties are not absolutes.

Civil liberties and the prevention of terrorism represent attributes for which society often makes extreme symbolic commitments toward the highest level. Many would argue that civil liberties are guaranteed rights, rights that cannot be compromised. In much the same way, advocates of risk control often claim that so long as any individual is at risk of being killed involuntarily, the risk must be reduced to ensure that we are in fact truly safe. Taken to the logical limit, this leads to the zero-risk mentality that pervades many legislative mandates of U.S. government risk and environmental regulation agencies, and is reflected in public risk attitudes as well.

These conflicting absolutes can not survive. Curve yy above can not be wished away; indeed, neither absolute is tenable. If we were to have a situation in which civil liberties were not compromised but were at the highest possible level, then the terrorism risks—as shown in the figure—would be enormous. In much the same way, completely eliminating the terrorism risk would require that we abandon most of our current civil liberties. The optimal outcome shown in this figure at D represents a trade-off between these two concerns. This tradeoff hinges on two aspects of the decision—first, the opportunities locus, and second the

shape of preferences. The tradeoff rate that is observed at point A before the 9/11 attack, and at point D after the attack simultaneously reflect the slope of both the indifference curves of individuals at that point as well as the slope of the opportunities locus there. Thus, these tradeoffs reflect influences of both supply and demand. Knowing the location of the curves xx or yy alone cannot tell us the optimal tradeoff rate.[3]

To our knowledge, market data does not exist that would enable us to derive either explicit or implicit estimates of these tradeoffs. Thus, we chose to survey people directly about their preferences.

2.2. Survey estimates of tradeoff rates

2.2.1. Airport screening. We examine civil liberties issues pertaining to the targeting of passengers for screening at airports based on their demographic characteristics, most often salient characteristics such as ethnic background and country of origin. Such targeting is a real policy concern. Assume that profiling of terrorists has some informational content. Then targeting passengers for screening selectively will reduce the expected terrorism risk, given any level of expenditure on screening. Such targeting based on ethnicity or race may, however, systematically impose differential costs on particular groups within the population, where these correlations with ethnicity and race create concerns with respect to civil liberties. Indeed, even if the inspection itself is relatively costless, being singled out for inspection is not.

The result is that the stated United States policy for screening is that it is largely random with respect to race and ethnicity rather than systematic. Such a policy would imply a horizontal coordinate toward the right in Figure 1. Whether current screening is in fact random is hard to determine. Official statements regarding the targeting procedures are not definitive. Attorney General John Ashcroft indicated that he opposed targeting "suspects solely on their race or ethnic origin." Similarly, FBI Director Robert Mueller stated: "We do not, have not, will not target people based solely on their ethnicity. Period." Neither official ruled out race and ethnicity as factors that might affect targeting, as they only ruled out such targeting *solely* on the basis of race or ethnicity. From a citizen's standpoint, whether it should be random depends both on how much that increases risk, and on that individual's preferences tradeoff rate between risk and civil liberties.

Our survey examines these tradeoffs. We divided the respondents into two groups, which received different variants of the screening question. For the first group there would be terrorism screening based on demographic characteristics, but the individual respondent would have a profile that would not lead that person to be singled out for selective screening. Rather, the costs—including the discomfort of being publicly identified as a risk—would be borne by others who would undergo the search. For this question, the respondent was asked whether this targeting of other passengers was desirable if the alternative were to undertake a random screening process that involved delays for all passengers. Each respondent considered situations in which the extra delay was 10 min., 30 min., and 1 h. In particular, the wording of the question was as 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 profiling. People who are singled out based on the racial profiles will have to undergo an extra 10 min. 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 min. so that all passengers could be screened?

Yes ___ No ___

b. Would you favor terrorist risk profiling if the alternative was for you to wait in line an extra 30 min. so that all passengers could be screened?

Yes ___ No ___

c. Would you favor terrorist risk profiling if the alternative was for you to wait in line an extra 60 min. so that all passengers could be screened?

Yes ___ No ___

Table 1 summarizes the responses to this screening question. Consider the column in Panel A pertaining to the scenario in which the screening only affects others. With 10 min of delay from a random screening, 44.7% of the respondents would favor the targeted risk profiling. This percentage rises to 55.3% if the delay is 30 min., and reaches 73.9% if the alternative to the risk profiling was for all passengers to wait an extra hour in line so that they can be screened on a random basis. Thus, individuals did not hold absolute attitudes towards this civil liberty; 29.2% of the people would accept profiling if it saved an hour for all, but would not accept it if it saved merely 10 min.

The bottom Panel B of Table 1 indicates how these responses differ based upon whether the respondent is white or nonwhite. The level of support for targeting others is lower among nonwhites than it is for whites, which is not surprising given that nonwhites are more likely to have been targets of racial profiling in other contexts.

These results are different than the findings of Gallup public opinion polls in which 71 percent of black respondents favor more intensive security checks for Arabs, including Arab Americans, as compared to 57 percent of white respondents who favor such targeting.[4] Subsequent polls suggested that this level of support for screening of Arabs may be due in part to black respondents' greater lack of familiarity with Arabs and Arab Americans, leading them to envision an inaccurate stereotype. More important is that our survey focuses on targeting by race that will include black respondents and not be restricted to people of Arab descent.

The second version of the survey differed from the first only in the last sentence, which indicated that the respondent would be selected for the searches: "*You would* be singled out for such searches based on terrorist risk profiling." That is the respondent would bear

Table 1. Attitudes toward use of terrorism risk profiles.

Delay in line due to screening time	Percentage favoring risk profiling[a]			
	Screening only affects others	*N*	Screening affects respondent	*N*
	Panel A: General results			
10 min.	44.7	47	50.0	48
30 min.	55.3	47	52.1	48
60 min.	73.9	46	56.3	48

Delay in line due to screening time	Percentage favoring risk profiling			
	Screening only affects others[b]		Screening affects respondent[c]	
	White	Nonwhite	White	Nonwhite
	Panel B: Results by race			
10 min.	43.6	33.3	62.5	25.0
30 min.	56.4	33.3	62.5	31.3
60 min.	73.7	66.7	65.6	37.5

[a]Differences in percentages for 60 min. are significantly different at the 10% level, two-sided test.
[b]Means are not significantly different at the 5% level. There were 6 nonwhite respondents for this survey version.
[c]Means for 10 min. and 30 min. are significantly different at the 5% level, two-tailed test. Means for 60 min. are significantly different at the 10% level, two-tailed test. There were 16 nonwhite respondents in this survey version.

the differential cost, and the alternative, to avoid being targeted oneself, involves imposing costs on all other passengers ranging from 10 min. to an extra hour of delay.

Respondents' attitudes toward targeting when they would be subject to search are quite different from those expressed when the targeting affects others. In particular, as Panel A in Table 1 shows, respondents who are targets are little influenced by the length of time that other passengers must be subjected to searches. Rather, they appear to be driven by the concern whether time delays for all, of whatever length, are preferable to being targeted oneself. Support for targeting ranges from 50 to 56% of the respondents.[5]

However, the results in Panel B of Table 1 show that the attitudes of the white and nonwhite respondents differ substantially. White respondents are generally supportive of targeting that affects themselves. This level of support ranging from 62.5 to 65.6%, for the three different time delays involved is unrelated to the length of time. About half as many nonwhites support targeting; this is not surprising given that they have historically been singled out for discrimination. When targeted, neither group responds strongly to the waits that must be endured by others. Interestingly, the level of white support for profiling affecting themselves when the passenger delays are 10 min. is greater than their support of profiling affecting others when the delay is 10 min., as is reflected in a comparison of Panel B in Table 1. The level of nonwhite support for targeting affecting themselves is roughly half that of whites. It is relatively unresponsive to the delay time for screening.

To explore the efficiency and equity issues underlying support for passenger screening, Table 2 reports probit estimates of the probability that a respondent favors screening. Because the subjects' level of risk estimates did not affect these probabilities, these risk

Table 2. Determinants of the probability of favoring targeting of passengers for airport screening.

Independent variable	Coefficient (std. error)	
Waiting time (in min.)	0.0038**	0.0065***
	(0.0015)	(0.0022)
Respondents targeted for screening	0.0190	0.1870
	(0.0627)	(0.1140)
Nonwhite	−0.2653***	−0.2655***
	(0.0696)	(0.0697)
Waiting time × Respondent targeted for screening		−0.0052*
		(0.0030)
$N = 278$		

[a]All regressions are probit estimates, which also include a constant term. The coefficients reported are the transformed probit coefficients that correspond to the marginal effect of the variable on the probability of favoring targeting.
***Indicates coefficient is significantly different from zero at 1% level, two-sided test.
**Indicates coefficient is significantly different from zero at 5% level, two-sided test.
*Indicates coefficient is significantly different from zero at 10% level, two-sided test.

variables do not appear in the equation. Consider the first set of results in Table 2. The first explanatory variable is waiting time. Increasing waiting time boosts the support for screening; each additional 10 min. of waiting time increases the probability that the respondent supports targeting by 0.038. Nonwhite respondents have a 0.27 lower probability of supporting profiling, presumably because they view the personal costs to them of racial profiling as being high. Surprisingly, among all respondents aversion to screening was not significantly related to whether it was the respondent being selected for screening or whether screening would affect others.

The second equation in Table 2 adds an interactive term in which waiting time is interacted with whether screening will lead to targeting of the respondent. Support for screening affecting the respondent decreases as the waiting time that will be avoided for passengers as a group increases. The implication of the negative interaction term is that waiting time is much more influential when respondents would be part of a general screening policy. When the respondent is targeted for screening, there is little sensitivity to waiting time. The equity issue of targeting appears to be a substantial concern for both racial groups, but more so for nonwhites.

In theory, the greater the risk one thought terrorism to be, and presumably the greater the risk reduction from targeted searches, the more individuals should be willing to sacrifice civil liberties for rigorous airplane searches. Surprisingly, we found no effect of individuals' risk estimates on their willingness to engage in targeted searches. That is why we delay our discussion of risk estimates until Section 3.[6]

2.2.2. Surveillance of communications. Many civil liberties concerns involve intrusions unrelated to air travel or physical personal searches. A second question included in the

survey for all respondents ascertained their attitude toward surveillance of their mail, e-mail, and phone communications. In particular, the respondents considered the following question:

> Would you support policies that make it easier for legal authorities to read mail, email, or tap phones without a person's knowledge so long as it was related to preventing terrorism?
>
> Yes __ No __

Table 3 summarizes the responses to this question. Overall, 36.2% of the sample support this surveillance policy. The level of support by white respondents is roughly double that for nonwhite respondents. As one would expect, respondents' attitudes toward profiling are correlated with their attitudes toward surveillance of mail, e-mail, and phones. People who answered yes to any of the profiling questions had 42.9% support overall for the surveillance efforts compared to 16.7% for those who did not support any of the profiling possibilities. There consequently appeared to be consistency in terms of support for different sacrifices of civil liberties, even though the components of the tradeoff differ across scenarios. In particular, civil liberties tradeoffs for airline passengers involve screening and public identification. By contrast, the surveillance of mail, e-mail, and phone calls would be less intrusive; they would not even be known to the person, would not impose any delays or other such costs, and would not entail being publicly identified as a target. However such secret snooping would impose unforeseen costs. For the plane-related terrorism risk, any risk avoided would benefit the passengers, including the respondent. By contrast, in the case of surveillance of mail, e-mail, and phones, any risk avoided would broadly affect the population, and the innocent respondent would have little stake in it.

Table 3. Attitudes toward surveillance of mail, e-mail, and phones.

	Percentage of respondents supporting policy
Full sample	36.2
White respondents	40.8*
Nonwhite respondents	18.2
Respondents answering yes to profiling[a]	42.9**
Respondents answering no to profiling	16.7

**(*) Indicates means are significantly different at the 5% (10%) level, two-sided tests.
[a]Respondents who answered yes to profiling to avoid any of the following: a 10-min. wait, a 30-min. wait, or a 60-min. wait.

The final question with respect to precautionary responses to terrorism did not involve civil liberties tradeoffs. Instead they involved costs that the respondent himself or herself was willing to incur to reduce risk. We asked to what extent would respondents change their behavior with respect to opening mail as a result of the anthrax incident in which anthrax spores enclosed in letters led to anthrax outbreaks in Florida and the Northeastern United States. The particular question was:

How much caution do you currently use in opening and handling your mail as compared to what you used before the Anthrax scare?

More Caution ____ Same Caution ____ Less Caution ____

As shown by the results in Table 4, the overwhelming percentage (91.5%) of the respondents exercised the same caution as before. No respondents exercised less caution than before, which is presumably rational given the increased risk; 8.5% exercised more caution.

Table 4. Degree of caution exercised with respect to opening mail.

Degree of caution compared to before the anthrax incident	Percentage[a]
More caution	8.51
Same caution	91.49
Less caution	0.00

[a]94 observations.

While respondents may appear lax with respect to these precautions, this low level of precautions seems appropriate to us. It is more likely that a terrorist would target a government official or some other visible figure rather than a student. Moreover, if there is another series of anthrax contaminations in the mail, all but the few early recipients would likely get some advance notice of generalized risk, as other people, either local or elsewhere, received contaminated letters.

3. Characteristics of terrorism risk beliefs

Assessing the risk of a terrorism attack properly requires that one makes judgments about events for which there is a very sparse informational base due to the rare nature of such events. In this section, we report on respondents' assessment of the risks of future attacks given that they have experienced the attacks in 2001. In the subsequent section we ask respondents to take themselves back to the situation before the 9/11 attack and assess the risks given the previous informational base.

For each instance, the historical number of terrorist attacks on airplanes that might inform such judgments is quite modest. Including the four hijackings on 9/11, since 1970 there have only been eight incidents of fatalities to U.S. citizens resulting from airplane hijackings. 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 the result of a hijacked Pacific Southwest flight by a former employee.

The two American Airlines flights from Boston and Washington, and the two United Airlines Flights from Boston and Newark that were hijacked on 9/11/2001 complete the set

of all hijackings since 1970. Thus, half of the fatal hijackings involving a U.S. carrier on a flight arriving to or departing from the United States took place on a single day. The total death toll from all such hijackings before 2001 was 90. To turn this figure from 1970-2000 into a risk per flight one also needs the denominator of the number of passengers enplaned on foreign and domestic U.S carrier flights over that period, which is 11.95 billion.[7] The hijacking fatality risk per flight was 7.5 per billion over the 1970-2000 period or just over 1 per 100 million flights. The hijackings on 9/11 have led to an upward reassessment of that risk by several orders of magnitude.

The number of lives that will be lost to terrorism in the future is extremely uncertain. There may be few or no terrorism deaths in the United States at all in the coming year, as there were from 9/12/01 to 9/11/02. Or there could be another wave of terrorism attacks that leaves thousands or conceivably millions dead. To gauge our respondents' estimates, we developed two sets of survey questions, that were given to two different groups of respondents. The first set of questions focused on terrorism attacks on airplanes, and the second group included terrorism attacks on airplanes, as well as terrorism attacks involving bombs and bullets. This structure was designed to test for an embeddedness effect whereby risk beliefs are not sensitive to the range of terrorism risk events included in the listing.

To analyze the range of people's expectations on terrorism losses, we asked respondents not only for their best estimate of the risks but also for their estimates at the 5th percentile of their distribution and at the 95th percentile of their distribution. The terrorism risk assessment question for the airplane risk scenario was 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 consequently 2,983, or about 3,000. Below is a series of questions about the number of people whom you believe will be killed in the next 12 months because of attacks by foreign terrorists on airplanes.
>
> (a) Think of the best-case 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? _____
> (b) Now think of the worst-case 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? _____
> (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 high-end death toll. What is your best estimate of the expected number of terrorism deaths over the next 12 months? _____

The top Panel of Table 5 presents the estimates for the airplane terrorism risk scenario responses. Consider the median responses. The lower bound estimate is zero deaths and the upper bound estimate is 4,000 deaths, with the best estimate being 75.5 deaths. Thus, people's best estimates of the terrorism risk on airplanes is well below the midpoint of their lower bound and upper bound estimates and is very close to the lower bound. Much the same is true of the mean values for this assessment, with the main difference being that

Table 5. Terrorism fatality estimates for the next 12 months.[a]

	Median	Mean	Std. error of mean	*N*
Estimates of total fatalities due to airplane terrorism				
Lower bound	0.00	16.22	6.89	46
Upper bound	4,000.00	16,596.02	5,285.07	46
Best estimate	75.50	188.83	41.03	46
Estimates of total fatalities due to all terrorism				
Lower bound	0.00	33.45	12.92	44
Upper bound	2,000.00	35,199.55	18,277.90	44
Best estimate	100.00	403.59	228.85	44
Estimates of total terrorism fatalities for pooled sample				
Lower bound	0.00	24.64	7.25	90
Upper bound	2,500.00	25,691.08	9,332.83	90
Best estimate	88.00	293.82	113.73	90

[a]The lower bound is the response provided for a one chance in 20 that the number of terrorism deaths could be this low. The upper bound is the response provided for a one chance in 20 that the number of terrorism deaths could be this high. Values exclude four outliers with values greater than 1 million. See appendix Table A1 for values including these outliers.

the upper bound estimate is much higher than the median value because of the influence of outliers.

The second variant of the survey focused on responses in which all three risks of terrorism are included, not simply those due to airplanes. The one change was in the third sentence of the introductory discussion, which now read: "Below is a series of questions about the number of people you believe will be killed in the next 12 months because of violent terrorist attacks by non-US citizens, e.g., crashed airplanes, bombs, and bullets." By logic, since this question includes a much broader set of possible terrorism risks, the risk assessment should be much higher. As the middle set of results in Table 5 indicate, the assessed total fatalities due to all terrorism tend to be fairly similar to those for airplane terrorism only. Moreover, the pattern of responses is the same: the best estimates of the terrorism risk lie very close to the lower bound estimate.

The bottom section of Table 5 pools these results, where we see that the median lower bound for both scenarios is zero, and the median upper bound is 2,500. The mean estimated number of terrorism fatalities is 88.

To eliminate the influence of outliers, Table 5 excludes four respondents who assessed terrorism risks greater than one million. Such responses are not necessarily errors, as a nuclear attack or a major biological or chemical attack could lead to such substantial deaths. However, including these responses potentially may distort some of the patterns in Table 5. As is indicated by the results in Table A1, including these outliers greatly affects the mean values for the upper bound estimates of the terrorism risk due to all terrorism and for the pooled sample, but does not otherwise alter the general pattern of the results.

The distribution of the responses is illustrated in Figure 2(a) for the lower bound, Figure 2(b) for the upper bound, and Figure 2(c) for the best estimate by respondents. The lower bound mass is densely concentrated at zero deaths, with 10 fatalities being the next most frequent response. The upper bound risk in Figure 2b is also highly skewed, with the most frequent responses being 1,000 fatalities, 5,000 fatalities, and 10,000 fatalities. Some respondents also present extreme high values for their upper bound estimates of the terrorism fatality risk. The best estimate of the terrorism risks by respondents appears in Figure 2c. The principal spike in that distribution occurs at 100 fatalities, which more than half of the sample has selected as their expected terrorism risk. Some respondents believe the risk could be as low as zero, but the best estimate of the risk often extends to the hundreds of fatalities and even beyond 1,000 in some cases.

What these results suggest is that any particular individual's estimate of terrorism risk over the coming year is highly uncertain.[8] While most of the weight of the distribution is toward relatively few fatalities, there is also the expectation with a much smaller probability that the risk could in fact be quite high, even much higher than was experienced on September 11th, 2001.

How do people form their best estimates of the terrorism risk? To explore this question, we calculated regression estimates of the best estimate as a function of the respondent's lower bound estimate, the respondent's upper bound estimate, and an indicator variable for whether the respondents had version B of the survey in which the terrorism risk arose from crashed airplanes, bombs, and bullets, and not simply from airplanes alone. Because zero values are included in the dependent variable, we report tobit regression estimates of the equation in Table 6. In each instance we estimate a linear equation in which the dependent variable is the number of terrorism deaths and a log equation in which the dependent variable is the natural logarithm of one plus the number of terrorism deaths.

If respondents simply selected the midpoint between the lower bound and the upper bound estimates as their best estimate, then we would observe a coefficient of 0.5 in the linear equation for both the lower bound and the upper bound values. Such a result was not found. The upper bound risk estimate has no significant effect at all on the best estimate of the fatality risk. In contrast, the lower bound estimate is statistically significant, as respondents appear to simply increase the lower bound estimate by a factor of 5 in forming their best estimate of the risk.

In the log equation, both the lower bound and the upper bound values are statistically significant determinants of the best estimate of the fatality risk. The logarithmic transformation depresses the otherwise distorting influence of the upper bound outliers, so these upper bound values now play a significant role. But, the lower bound risk values are still more consequential. Respondents place a coefficient weight of 0.593 on the log value of the lower bound risk estimate as compared to a coefficient of 0.262 on the log value of the upper bound risk estimate.

4. Biases and anomalies in risk beliefs

An important question from the standpoint of analyzing whether people form risk beliefs sensibly is whether respondents take into account the range of risks included in the question.

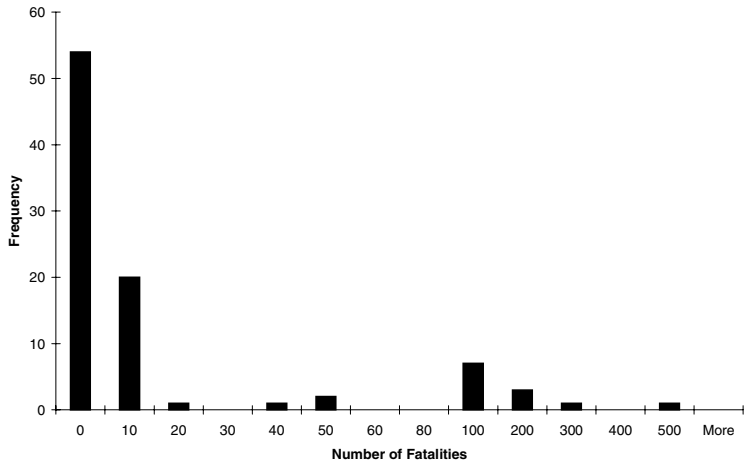

(a)

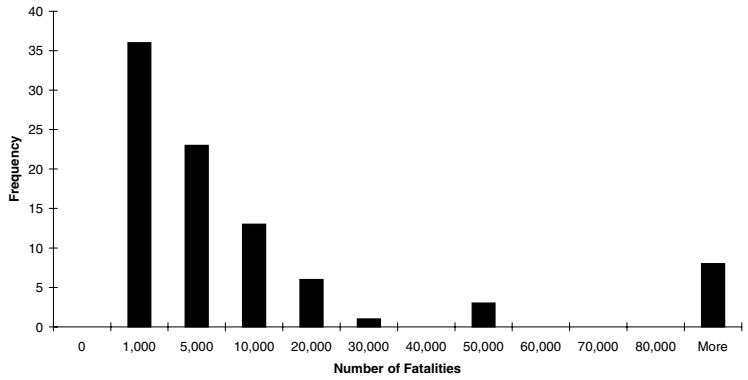

(b)

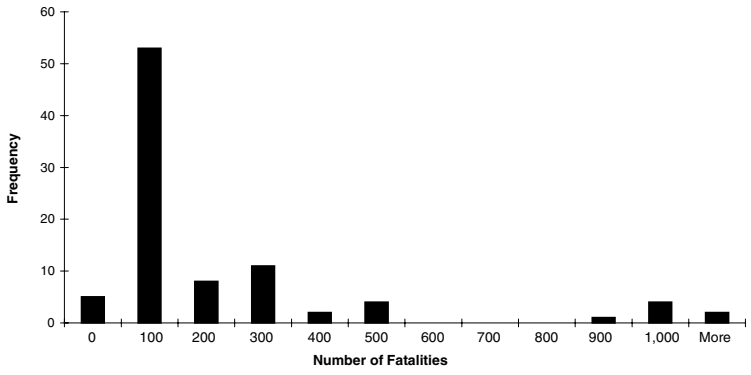

(c)

Figure 2. (a.) Total terrorism fatalities for pooled sample: Lower bound. (b) Total terrorism fatalities for pooled sample: Upper bound. (c) Total terrorism fatalities for pooled sample: Best estimate.

Table 6. Tobit regression estimates of best estimates of fatality risk from terrorism.

	Coefficient (std. error)	
Variable	Linear equation	Log equation[a]
Lower bound	5.213**	0.539**
	(1.628)	(0.093)
Upper bound	0.001	0.262**
	(0.001)	(0.079)
Version B indicator variable	120.622	−0.277
	(225.303)	(0.333)
Constant	31.374	1.486*
	(159.597)	(0.685)
Log likelihood	−715.001	−165.837
N	90	90

[a]The dependent variable, the lower bound value, and the upper bound value are all in natural logarithms.
**(*) Indicates coefficient is significantly different from zero at 1% (5%) level, two-sided test.

The first version of the survey included risks from airplanes only, whereas version B of the survey included all terrorism risks, including among other things, crashed airplanes, bombs, and bullets. The overall risk posed by this inclusive question consequently should be greater. However, as the results in Table 5 indicate, the responses to the two survey variants are quite similar. The inclusion of the version B indicator variable in the regression estimates reported in Table 6 provides a formal test of this hypothesis. In both the linear equation and the log equation there is no statistically significant effect of the version B indicator variable, which by logic should have a positive influence on the assessed terrorism fatality risk because more types of terrorism risks are included.

What these results imply is that respondents are prey to an embedding effect. Whether you ask about one risk or a larger category in which it is embedded, you get the same result. This type of phenomenon has been documented with respect to contingent valuation studies in which respondents are often not sensitive to the scope of the commodity that is being purchased in the contingent valuation survey.[9] Here we observe a similar phenomenon in which the risk assessment for terrorism risks from a narrowly defined category, in particular airplanes, is not statistically significant from an inclusive measure of the terrorism risk from all sources. Individuals consequently underestimate the additional terrorism risk that will arise once the scope of the risks spans a broader set of outcomes.

These embedding problems associated with people's terrorism risk beliefs are not the only logical failing in people's risk beliefs. There are surely other difficulties as well that arise, not surprisingly, because terrorism risks are poorly understood. Moreover, being highly improbable events as well as events for which the underlying risk probabilities appear to have changed over time, there is little reliable basis for estimating these risks.

Hindsight bias is a potential problem in risk assessment. People often tend to believe that they knew the risk all along even though the risk was completely unanticipated. This

phenomenon often arises with respect to accidents, and also with jury behavior.[10] In both contexts, people believe they could have anticipated a catastrophe before it had happened. Second guessing managerial decisions and the general phrase "Monday morning quarterbacking" capture this influence as well. A main reason why the 9/11 catastrophe had such a dramatic impact was because it was unanticipated. People consequently should believe that such an event is more likely to occur now than before, since previously the estimated likelihood of an occurrance should have been quite low, if not zero.

The formulation we used in our survey to test for hindsight effects was the following:

> Take yourself back before the World Trade Center disaster. Do you believe that the risk of a terrorist attack on an airplane is higher or lower than you thought it was before the September 11 disaster?
>
> Higher ____ The Same ____ Lower ____

Table 7 summarizes the responses to this question. While actual terrorism risk may now be lower than it was before the attack because of increased precautions, it is still higher than people probably believed it to be before 9/11. Thus the perceived terrorism risk today should still be greater than what people believed the risk to be before the attack. The responses in Table 7 fail to indicate such changes in risk beliefs. Indeed, the number of respondents who believe the risk is higher is almost the same as the number who believe it was lower than they thought it was before the 9/11 disaster. Overall, 42.6% of respondents believed the risk is higher and 40.4% believe the risk is lower—differences that are not statistically significant. The remaining 17% of respondents believe that the risk is the same as it was before the terrorism attack.

Appendix A summarizes the reasons why the respondents did or did not change their risk beliefs after the September 11 disaster. The respondents indicating a higher risk cite reasons such as not being aware of such terrorism risks or the weakness of security efforts. The potential for copycat terrorism attacks is also a major concern. The respondents who indicated that their risk beliefs are the same are more plausible. Many indicate that the risk is surprisingly large but that security measures should have some effect. While these explanations may be consistent with the actual risk, which may be the same, they are less

Table 7. Current assessed risk of terrorist attack on an airplane.

Assessed terrorism risk value compared to what respondent believed the risk was before 9/11	Percentage[a]
Higher	42.55
The same	17.02
Lower	40.43

[a]94 observations.

persuasive as reasons for why the perceived risk is the same. Confusing actual and perceived risks is only appropriate if people have perfect hindsight. Some respondents who estimate risks as being lower suggest that terrorists will switch to targets other than airplanes, which would not be a form of hindsight bias.

Another set of questions explored whether people's risk-money tradeoffs displayed a consistent pattern that would reflect a rational response to terrorism risks. One potential bias is the influence of certainty premiums whereby people are willing to pay much more for reducing the risk to zero than is warranted by the extent of the risk reduction. This zero risk mentality is usually viewed as a form of irrationality though there could be possible rational explanations for it, such as the elimination of anxiety with respect to a non-zero level of risk.

To examine these risk-money tradeoffs we formulated the following question:

> Security on airplanes can be increased through investments in technologies that are better at scanning luggage and passengers. Suppose that such screening was financed by a surcharge that increases the ticket price by the same amount for all tickets.
>
> (a) What price increase in the ticket price would you be willing to pay for screening measures that would decrease the risk of a terrorism attack on an airplane by 50%? _____
> (b) What price increase in ticket price would you be willing to pay for screening measures that would decrease the risk of a terrorism attack on an airplane to 1 chance in 10,000,000 for each plane flight? _____
> (c) What price increase in ticket price would you be willing to pay for screening measures that would decrease the risk of a terrorism attack on an airplane to 0? _____

In an alternate version of the survey, question B was replaced by having the chance be 1 chance in 1 million rather than 1 chance in 10 million per plane flight.

Table 8 summarizes the responses to these questions. Before assessing whether respondents are rational, it is worthwhile to indicate what the level of the risk reduction is in each instance and what kinds of patterns should be expected. Suppose the current level of risk is R, which is not observable or known. Then the risk reduction achieved by reducing the risk by 50% of its current level is 0.5R. Similarly, the risk reduction from the fourth possibility in Table 8 reducing the risk to zero is R. Reducing the risk to zero as opposed to 50% of its current level consequently achieves twice the value of the risk reduction. For small probabilities such as this, the willingness to pay for the risk reduction should be a relatively invariant amount per unit risk reduction so that the willingness to pay to reach zero should be roughly double the value for achieving a 50% reduction. Both the mean and median values of reducing the risk to zero are more than double that for the 50% reduction. The standard errors are, however, quite large given the broad range of responses. The point estimates for the mean and median values are consistent with there being a substantial premium for reaching a zero risk level.

These effects are highly consistent with results in the literature pertaining to premiums for reaching a zero risk level. Claims that risk will be completely eliminated receive much greater public support than claims that the risk is simply being reduced. Not surprisingly,

Table 8. Absolute price increase willing to incur for reduced risk by screening.

Post-screening risk of terrorist attack	Median	Mean	Std. error of mean	*N*
50 Percent of current level	20.0	24.95	3.47	44
1 in a million per flight	25.0	38.42	13.15	12
1 in 10 million per flight	27.5	53.15	14.15	34
Zero	50.0	69.43	14.68	45

after the terrorism attacks, public officials ventured forth with policy measures that were designed to eliminate terrorism risks. There were no claims, for example, that terrorism risks would be simply cut in half or would be restricted to one crashed jet liner per year. Public support is much greater when there are promises, however unrealistic, that the risk will be completely eliminated.

For the other two risk scenarios in Table 8 the risk reduction is R-10 $^{-6}$ in the case where the risk is reduced to 1 in a million per flight, and the risk reduction is R-10^{-7} when the risk is reduced to 1 in 10 million per flight. The willingness-to-pay responses are fairly similar in terms of the median values, though there are somewhat greater differences in terms of the mean. Each of the patterns is plausible, as the price increases that people are willing to incur are somewhat greater when the risk is reduced to 1 in 10 million than when it is reduced to 1 in 1 million. Whether the relative values of the responses are rational depends on the base risk R, so one cannot say whether these willingness-to-pay values are rational.

However, if one inspects the mean values for all four different post-screening risks, thus avoiding the influence of outliers, one finds a pattern that is more disturbing. Whether the terrorism risk is reduced to 50% of its current level, 1 in 1 million per flight, or 1 in 10 million per flight or zero, is not significantly different. Indeed, a regression analysis of the price increase the respondent would accept does not differ significantly across these four different post-screening risk values shown in Table 8. Doing something about terrorism risks that is incomplete but beneficial consequently has a fairly similar attractiveness across these three options. However, a policy that would completely eliminate the terrorism risk is much more attractive and commands a much higher willingness-to-pay value.

5. Conclusion

The 9/11 attack was unprecedented in terms of the nature of the attack and the number of U.S lives lost. However, the structural component of the risk management problems arising from the attack reveals many elements common to risk policy more generally. Most important, any risk reduction policy surely will involve tradeoffs, as a zero risk level will be infeasible or astronomically expensive.

The principal tradeoff considered here was with respect to civil liberties, which may also be viewed as a concern that cannot be compromised. Respondents indicated a willingness to trade off civil liberties concerns, especially when there were significant efficiency gains in terms of reduced waiting time. Whether the respondent would be profiled rather than

others was not a salient concern. Nonwhites, who may have had more past experiences with such targeting, are less supportive of targeted screening policies.

Many of the other aspects of terrorism risk beliefs illustrate phenomena widely studied in the existing literature on risk and uncertainty. Terrorism risks, which are highly unpredictable and hardly subject to conventional statistical assessment, must be gauged as subjective probabilities, whether by experts or citizens. In making such assessments, our respondents display the familiar biases of embeddedness, hindsight bias, and certainty premia.

Our anti-terrorism policies, whatever their form, should be guided by best estimates of the terrorism risk, and should recognize that the optimal policy must involve tradeoffs of some concerns that their advocates claim can never be compromised.

Appendix A

Question: Take yourself back before 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 11 disaster? Please provide 2 or 3 sentences for why your estimates changed or stayed the same.

Higher:

- "Retaliation for U.S. response to Sept. 11. Copycat attacks more likely."
- "While the actual risk might be lower due to increased security, my perception of the risk before Sept. 11 of the chance of a terrorist attack on an American plane was around zero."
- "Though there is more security now, that security is often failing. Also, the situation in the country just seems less politically stable."
- "I didn't realize that there were terrorists targeting the U.S. prior to Sept. 11."
- "I was previously unaware of the laxity of security procedures and the ease of evading them."
- "The success of 9–11 may lead other attackers to go through with a plan they otherwise wouldn't have because security is seen as more vulnerable."
- "The attack exposed a vast network of terrorists and an amount of anti-American sentiment that I had not known existed. It exposed our airport security and INS procedures as grossly inadequate and insecure."
- "I thought plane attacks had become passé—much easier to bomb embassies or public places or start shooting up things."
- "Current terrorists are inspiring more people to commit terrorism."
- "Before 9–11, I didn't think that this sort of airplane-terrorism was plausible. Now I know that it can happen—regardless of how many precautions are taken."

The Same:

- "I think the risk was certainly high before and although terrorists don't seem as likely to use planes again in the near future, I think the risk is almost the same. Terrorists are smart and well-financed."

- "I don't think any new measures that have been taken will effectively prevent terrorism. Until airports are totally privatized, I don't think anything is safer. The government is too inefficient to really make any substantial changes."
- "If risk level is constant, then occurrence of event doesn't change results."
- "Not sure any security measures will prevent determined terrorists. Not sure if a successful mission (9/11) contents terrorists to rest on their laurels or makes them hungrier for more blood."
- "The event is unlikely to reoccur. Instead, terrorists will spend 2–3 years planning the next attack which will be in a different form, such as an Ebola-injected suicide bomber entering the Fleet Center during a sold out basketball game and infecting all those present."
- "Previously, I thought security was lacking, but motivation was too. Now, it seems that security is better, but motivation of terrorists is higher too."
- "More people are likely to try to blow stuff up—9/11 gave people ideas. But security measures have improved. So, my guess is that it's roughly the same. The two balance each other out."
- "Much of the safety changes are meaningless—cosmetic changes only—ex. INS still gave out visas to the dead terrorists, same level of incompetence in government agencies."
- "Even though security is screening people, it is still not that tight and there are still incidences of people running through security. Also, terrorists are likely to take more efforts in concealing weapons due to tighter security. Finally, not all terrorists fit the profile."
- "Although I think people are more scared of and aware of the risk of attacks, studies show that it is still quite possible for weapon-bearing passengers to get through airport security."

Lower:

- "Before the attack, I never considered a terrorist attack from planes flying into buildings. I thought the chance of terrorist bombing on plane was extremely remote. However, I now think that a similar terrorist attack is even more unlikely since passengers and crew would not let plane be taken over."
- "There are better security measures in place and more vigilance/awareness on the part of consumers that make me feel safer."
- "Everyone is more sensitive to terrorism. There is increased security. Higher likelihood that passengers will take security into their own hands."
- "With at least the appearance of heightened security, plus the flying public's (the other passengers) unwillingness to let a hijacker take over a plane, plus the ample opportunity to do harm in ways we haven't anticipated, rational terrorists will resort to other tactics."
- "Because now at least the lame half-hearted random, but somewhat improved measures may intimidate some potential terrorists (on Sept. 11, I was surprised something similar hadn't happened sooner.)"
- "Greater security. Disruption of terrorist networks."
- "Lightning doesn't strike twice."
- "Terrorists will find new ways to terrorize—there are easier ways than hijacking planes."
- "Terrorists will conduct future attacks via unexpected methods. Attacks with airplanes are old but there are plenty of alternative ways to cause destruction (e.g., poisoning water supplies)."

- "Extra security measures have probably decreased the risk. There may be more attempts, but a greater percentage of them will be foiled."

Table A1. Terrorism fatality estimates for the next 12 months.[a]

	Median	Mean	Std. error of mean	*N*
Estimates of total fatalities due to airplane terrorism				
Lower bound	0.00	16.22	6.89	46
Upper bound	4,000.00	16,596.02	5,285.07	46
Best estimate	75.50	188.83	41.03	46
Estimates of total fatalities due to all terrorism				
Lower bound	0.00	31.75	11.89	48
Upper bound	2,500.00	21,344,766.25	20,824,334.22	48
Best estimate	100.00	397.15	210.25	48
Estimates of total terrorism fatalities for pooled sample				
Lower bound	0.00	24.15	6.96	94
Upper bound	3,000.00	10,907,576.56	10,636,385.78	94
Best estimate	100.00	295.20	109.19	94

[a]The lower bound is the response provided for a one chance in 20 that the number of terrorism deaths could be this low. The upper bound is the response provided for a one chance in 20 that the number of terrorism deaths could be this high. Values include four outliers with values greater than 1 million.

Acknowledgments

Professor Viscusi's research is supported by the Harvard Olin Center for Law, Business, and Economics and the Sheldon Seevak Research Fund. DeYett Law, Baruch Fischhoff, Russell Archer, and participants in the Harvard Law and Economics workshop provided excellent comments.

Notes

1. The difficulty of estimating low probability risks is stressed by Kunreuther et al. (1978). The role of media attention is discussed in Fischhoff et al. (1981).
2. Most of the students were largely members of the analytic methods for lawyers class, which consists primarily of first year law school students who lack formal economic training. There were 15 additional respondents from the treatment of scientific evidence seminar.
3. If the opportunities locus is linear, then the rate is determined. But even for this case we need to know preferences to know where to operate on the locus.
4. Ann Sales, "Profiling of Arabs: Polls Say Blacks Tend to Favor Checks." *Boston Globe*, Sept 30, 2001.
5. It would be inappropriate to conclude that individuals who were really natural targets of profiling would have responded in this way.
6. In particular, we included in separate regressions respondents' best estimates of the expected number of terrorism risks in the coming year, the lower bound of this estimate, and the upper bound of this estimate, where these were both in linear and in log form. In addition, a 0–1 indicator for whether the respondent assessed a risk level above the median best estimate level was also included. None of the variables was significant.

7. This information was calculated using data posted by the Air Transport Association at www.air-transport.org.
8. Risk ambiguity has been a continuing concern in the literature on biases and anomalies in risk beliefs. See, among others, Camerer and Weber (1992) and Viscusi (1997, 1998).
9. Kahneman and Knetsch (1992) are among the many authors who have documented this phenomenon.
10. For the jury studies, see Chapters 6 and 11 of Sunstein et al. (2002). More generally, also see Fischhoff (1975), Kelman, Fallas, and Folger (1998), and Rachlinski (1998).

References

Camerer, Colin F. and Weber, Martin. (1992). "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty* 5, 325–370.

Fischhoff, Baruch. (1975). "Hindsight $\neq$ Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty," *Journal of Experimental Psychology, Human Perceptions and Performance* 1, 288–299.

Fischhoff, Baruch et al. (1981). *Acceptable Risk*. Cambridge: Cambridge University Press.

Kahneman, Daniel and Jack L. Knetsch. (1992). "Valuing Public Goods: The Purchase of Moral Satisfaction," *Journal of Environmental Economics and Management* 22, 57–70.

Kaplow, Louis and Steven, Shavell. (2002). *Fairness and Welfare*. Cambridge: Harvard University Press.

Kelman, Mark, David Fallas, and Hilary Folger. (1998). "Decomposing Hindsight Bias," *Journal of Risk and Uncertainty* 16, 251–269.

Kunreuther, Howard et al. (1978). *Disaster Insurance Protection: Public Policy Lessons*. New York: John Wiley.

Rachlinski, Jeffrey. (1998). "A Positive Psychological Theory of Judging in Hindsight," *University of Chicago Law Review* 65(2), 571–625.

Sunstein, Cass, Reid Hastie, John Payne, David Schkade, and W. Kip Viscusi. (2002). *Punitive Damages: How Juries Decide*. Chicago: University of Chicago Press.

Viscusi, W. Kip. (1997). "Alarmist Decisions with Divergent Risk Information," *The Economic Journal* 445, 1657–1670.

Viscusi, W. Kip. (1998). *Rational Risk Policy.* Oxford: Oxford University Press.