Risk Analysis, Vol. 32, No. 8, 2012 DOI: 10.1111/j.1539-6924.2012.01849.x

Optimal Security Investments and Extreme Risk

Hamid Mohtadi[1,*] and Swati Agiwal[2,†]

In the aftermath of 9/11, concern over security increased dramatically in both the public and the private sector. Yet, no clear algorithm exists to inform firms on the amount and the timing of security investments to mitigate the impact of catastrophic risks. The goal of this article is to devise an optimum investment strategy for firms to mitigate exposure to catastrophic risks, focusing on how much to invest and when to invest. The latter question addresses the issue of whether postponing a risk mitigating decision is an optimal strategy or not. Accordingly, we develop and estimate both a one-period model and a multiperiod model within the framework of extreme value theory (EVT). We calibrate these models using probability measures for catastrophic terrorism risks associated with attacks on the food sector. We then compare our findings with the purchase of catastrophic risk insurance.

KEY WORDS: Catastrophic risk; extreme value statistics; food sector; multiperiod; optimum investments

1. INTRODUCTION

In the aftermath of September 11, the greatest terrorist attack of all time, and Katrina, the greatest natural disaster in recent history, insurance and reinsurance companies have become increasingly reluctant to cover catastrophic risks. The numerous difficulties that have plagued catastrophic risk insurance, and particularly terrorism risk insurance,(1,2) lead us to the need to consider what firms and businesses can do by way of *ex ante* investments in risk mitigating strategies. However, little is known about how much investments for self-protection (self-insurance), or other hedging strategies, firms need when insurance markets fail to provide the necessary protection.

[1]University of Wisconsin, Milwaukee, WI, USA.
[2]Quantitative Research, US Bank Corp. Minneapolis, MN, USA.
[†]At the time of this research was a PhD Student at the Department of Applied Economics, University of Minnesota, Minneapolis, MN, USA.
[*]Address correspondence to Hamid Mohtadi, Professor of Economics, PO Box 413, University of Wisconsin, Milwaukee, WI 53201 USA and Visiting Professor of Applied Economics, University of Minnesota, Minneapolis, MN; mohtadi@umn.edu.

One of the principal reasons why catastrophic and especially terrorism insurance markets have not been adequately developed is the difficulty in assessing low-frequency high-impact risks.[3] Even within the world of catastrophic insurance, there are major differences between natural and terrorism risks. In 2003, Kunreuther(3) outlined some key differences between terrorism risk and natural disasters as a way of explaining why terrorism risk is so difficult to insure. Chief among them were three points that are relevant to our article. First was the absence of historical data for catastrophic terrorism risk, versus the availability of some historical data for natural disasters. Thus, the risk of occurrence could be reasonably well specified for natural disasters based on historical data and expert estimates, whereas for terrorism risk, there is considerable ambiguity of risk. Second, catastrophe modeling has been developed for natural risk since the late 1980s and early 1990s, whereas the first models to address terrorism risk were developed

[3] Of course, as elaborated in the concluding section, catastrophic insurance is an annual cost, whereas risk mitigation investments involve a one-time commitment of resources.

in 2002. Third, there is some geographic specificity of natural disasters (e.g., California for earthquakes or Florida for hurricanes), whereas terrorists may attack anywhere, any time.

Since 2003, a number of improvements have taken place in both the data and the risk modeling/measurement world that address some of the concerns that were raised in the Kunreuther *et al.*'s[(3)] study. First, historical terrorism data are now far more readily available than in 2003. The most prominent source for such data is the MIPT database as well as the data that have become available through the University of Maryland's START center.[3] Second, risk modeling and measurement has now been extended to the analysis of terrorism risk.[(4,5)]

Yet, obtaining terrorism risk insurance—at least at reasonable prices—remains an illusive goal, perhaps due other limitations spelled out by Kunreuther *et al*. Given the persistence of these difficulties, we must still rely on risk mitigation measures by the firms themselves—in the form of self-protection. However, even on this point, the Kunreuther *et al.* paper[(3)] implied limited success. Thus, they stated: "For natural disasters, the insureds can invest in well-known mitigation measures, whereas for terrorism risk, weapons and configurations are numerous, there are negative externalities of self-protection effort, and there is substitution effect in terrorist activity. Thus, the insureds may have difficulty in choosing measures to reduce consequences of an attack."

In this article, we try to address the self-protection challenge as well as the substitution in terrorist activity by reducing the ambiguity of terrorism risk, i.e., by narrowing the universe of risks. Specifically, we focus on the risk of intentional attacks on the food sector using chemical, biological, and radionuclear (CBRN) tools. The more focused the risk is, the more tangible it is to the firms in that sector. This facilitates firms' ability to *internalize* the risk and thus to invest in risk-mitigation for self-protection, as compared to a general form of risk such as, for example, a bomb threat. Because of this specificity, the negative externalities of self-protection effort are reduced. Also, the span of the substitution effect in terrorist activity is now limited to the universe of food-related risks. Finally, the "regional ambiguity" associated with terrorism risk at large (the third point that was raised in the Kunreuther *et al.*[(3)]) may also be lessened because food supply is typically associated with some geographically specifity, either in its production or in its distribution.

But this raises another challenge, namely, availability (i) of data on CBRN specific events, and (ii) of risk measurement/modeling of CBRN events. On data, Mohtadi and Murshid[(6)] complied and produced what were the most comprehensive, if not the first, available CBRN data and made them publicly available.[5] Subsequently, Mohtadi and Murshid[(7)] addressed the risk measurement question by developing quantitative probability measures that utilized the statistical properties of extreme heavy tail distributions from extreme value theory (EVT) to calculate catastrophic risks from intentional contamination of food by chemical or biological agents. This form of statistical analysis has gained in popularity in the recent past, owing to its ability to fit the data on rare events that would otherwise be considered outliers in the context of a normal bell-shaped curve. However, what complicated the analysis of man-made events, such as terrorism, is that these events are inherently different from naturally occurring risks. The critical point of separation is the element of intent, present in acts of terrorism, but absent in natural disasters. Unlike Mother Nature, terrorists adapt. They do so by reorganizing and restructuring, by redefining parameters of attack, by adopting new tactics, and by acquiring new weaponry. These innovations can be part of an ongoing effort to increase the potency of attacks, or they can be triggered by a change in counter terrorist practices that disturb the status quo.[(8)] At the same time an ever-changing geopolitical landscape seeds new political/terrorist movements with new ideologies and a higher propensity for violence. What all of this means is that the distribution of terrorism is unstable, more so than for any naturally occurring phenomenon. The critical issue is whether we can characterize this instability. Unfortunately, capabilities of terrorists as well as counter terrorist agencies are typically unobservable except in the terrorism data themselves. But to the extent that the probability law governing extreme events changes smoothly over time[6] it may be feasible to forecast current and/or

[4] A more detailed discussion of available terrorism data is presented in Mohtadi and Murshid.[(5)]

[5] Currently, START, at the University of Maryland, has data through 2007 which include new observations not included in Mohtadi and Murshid.[(6)] See http://www.start.umd.edu/gtd/. However, the sum total of CBRN events in that data set is 240. By contrast, the Mohtadi-Murshid data set has 447 CBRN observations.

[6] This statement may appear implausible at first given the dramatic nature of catastrophic events. However, what we mean by

near-term risks on the basis of established patterns within existing data. In this sense, EVT techniques provide a useful approach, especially as alternative methods for measuring extreme risk are not available.

This article takes the next step: it uses the previous advances to address the question of how much and when should firms invest to mitigate the impact of catastrophic risks on their operations. However, it also addresses the issue of the reliability of EVT modeling by providing an upper and a lower bound for the probability values that are derived from these models. We carry through these bounds into the calculation of the security investments. We first develop a simple analytical model. We then calibrate this model with the data on probability measures from the previous two studies. (For one of the parameters we also use a survey of food manufacturers based on the results from a working paper[9] that will be discussed later.) The decision of whether to invest at the present time to protect one's assets or to postpone such investments may appear to posses a "quasi" option value character. However, an option value formulation would provide a decision rule, given the amount of investments, whereas we are also interested in finding the optimum amount of investments and their optimum timing. To do this, we extend our model to incorporate the benefits of investing at some future time periods in a multiperiod framework. Our multiperiod model is based on a decision-tree approach. Interestingly, the multiperiod results suggest that self-protection needs to start from the very first period. In other words, we find that the benefits of waiting by not committing to large investments in security are outweighed by the rising risk of a major CBRN attack on the food supply. The fact that a multiperiod model yields a first period result as the optimum outcome is information that could not be inferred from the single-period model.

In relation to the literature, a study by Gordon and Loeb[10] considers optimal investment decisions in the area of information security breaches for IT systems within firms. Their study has a similar framework to our static framework. But two features distinguish our work. First, is the addition of the multiperiod framework, as mentioned, and second is the estimation of the actual values, given the probabilities of extreme risk mapped to the food sector.

In general, economic modeling based on purely profit maximizing or utilitarian assumptions is only one reason why firms invest in risk mitigation strategies. In practice, firms are making investments for certain types of catastrophic risks because, without any economic analysis, they have determined that a particular event outcome is unacceptable. In this respect, one way to motivate our analysis is to view it as a normative (prescriptive) rather than positive (descriptive) analysis, in the sense of using economic criteria to provide an *economically* optimal "best practice" benchmark against which the level of risk mitigation investments carried out by firms can be judged.

Our resulting estimates suggest that smaller but more probable potential losses may require a larger outlay of security investments to protect against than larger but less probable potential losses. For example, a one-time capital investment aimed at mitigating extreme risk is about 42 cents per \$100.00 for a potential loss of \$1.5 billion in a five-year time horizon, whereas it is only 6 cents per \$100.00 for a potential loss of \$3.75 billion in a one-year time horizon. The key is the associated probabilities. The event with a higher risk mitigation cost has a probability of 3.4% whereas the event with the lower percentage cost has a probability of only 0.5%.

Another interesting result is that for any given time horizon, optimum investments in risk mitigation by food companies as a percent of potential losses initially decline for some range with potentially higher losses but start to rise with greater losses beyond that range, thus producing an inverted U pattern. As it turns out, this pattern results because of the nature of "tail" events, as they are captured in extreme value statistical analysis. This and other issues are detailed in the text.

Our results suggest that in comparison to a regulated rate for terrorism risk insurance, the optimum level of investments to mitigate risk is somewhat larger, but does depend on the *level* of risk. However, the limitations that exist on the availability of catastrophic insurance at these regulated rates, and their high level of deductibility, together with the one-time nature of the alternative risk mitigating investment

a "smoothly evolving probability law" is that the realization of, say, a new massive event, whereas it might result in changes in the estimated parameter of an EVT distribution (e.g., a heavier tail estimate or a rightward shift of the location—i.e., mean—parameter of the distribution) than would otherwise be the case, it does *not* inherently alter the nature of the distribution. For example, Mohtadi and Murshid[7] show that for CBRN events, the tail parameter of the associated EVT distribution has a statistically significant positive time dependence, meaning that the distribution contains more extreme events over time. But the integrity of the distribution itself remains intact.

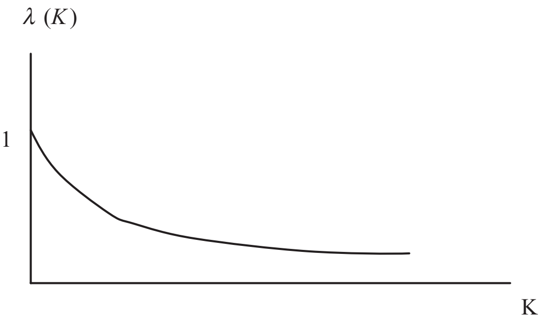

Fig. 1. Vulnerability and security investments.

strategies suggest that such investments should be undertaken whether or not catastrophic risk insurance is available, particularly as their impact on risk financing may not be negligible at all.

Following the development of the analytical one-period model in Section 2, the estimation method and the results are presented in Section 3. In Section 4, we extend the basic single-period model to account for multiple periods in a finite time horizon. Here, we address the question of whether postponing a risk mitigation investment is an optimal strategy or not. Section 5 makes some concluding observations based on the comparison of our results with those of catastrophic insurance.

2. A SIMPLE ANALYTICAL MODEL

Consider investing K in infrastructure/business practice/etc. *ex ante* to improve responsiveness to a catastrophic event, should such an event occur. Let L be the nominal loss in the case of an event. Then an initial investment of K could reduce the impact of the loss to a fraction $\lambda(K)$ such that $\lambda(K)L < L$ and $\lambda(K) < 1$. Higher values of initial investments reduce the impact of the loss L so that $\lambda'(K) < 0$. Thus, $\lambda(K)$ may be interpreted as the "*loss coefficient.*" However, this reduction in loss is bounded from below (by zero) as the law of diminishing marginal productivity would imply. Thus, without loss of generality one can assume that $\lambda''(K) > 0$ and that $\lambda(K)_{\lim K\to\infty} \longrightarrow 0$, as shown in Fig. 1.

Many functional forms satisfy the above requirement. We shall return to this issue shortly.

If π_L is the probability of an event of magnitude L, then the expected loss in the absence of investment is:

$$E(L|\pi_L)^{no\ invest} = \pi_L L. \quad (1)$$

In the case of investing K the expected loss is:

$$E(L|\pi_L)^{invest} = \pi_L\lambda(K)L + K. \quad (2)$$

Thus, though exposure to the event cannot be fully avoided, prior mitigating action can reduce the extent of loss from the event. Suppose the overall stream of profits from production is Π_o. Then, expected profits from investing and non-investing in security measures are, respectively, $E(\Pi^{invest}) = \Pi_o - E(L|\pi_L)^{invest}$ and $E(\Pi^{no\ invest}) = \Pi_o - E(L|\pi_L)^{no\ invest}$. Then expected *net* gain from investing, given the probability π_L, is denoted by $G(\pi_L)$ and is found from:

$$\begin{aligned} G(\pi_L) &= E(L|\pi_L)^{no\ invest} - E(L|\pi_L)^{invest} \\ &= \pi_L[1-\lambda(K)]L - K. \end{aligned} \quad (3)$$

Notice that $G(\pi_L)$ is $\gtreqless 0$ if $\pi_L \gtreqless \bar{\pi}$ where $\bar{\pi}$ is the threshold value of event probability and is given by $\bar{\pi} \equiv (K/L) \times (1/[1-\lambda(K)]$. This implies that for events with very low probability, a firm could lose by investing K!

Catastrophic events, which are the focus of this work, are by nature low-frequency high-impact events. Developing accurate metrics of catastrophic risk is not easy. The challenge arise because of the need to extrapolate from observed levels in data to unobserved levels. Classical statistical methods are not well-suited for this task.(11,12) Instead, the appropriate approach is to estimate catastrophic risk using extreme value (EV) analysis. By exploiting limiting arguments, EV models can provide an approximate description of the stochastic behavior of extremes.[7] This is both discussed and measured elsewhere.(5,7,13) Thus, $\pi(L) \sim f(L)$,where $f(L)$ is an EV probability distribution function. The expected gain from an event that leads to a loss in the range L_o to L_1 is obtained from the EV probability distribution function such that:

$$G(\pi_{L_0<L<L_1}) = [1-\lambda(K)]\int_{L_o}^{L_1} f(L).L.dL - K. \quad (4)$$

Owing to the complexity of the term inside the integral, we will approximate G by replacing L inside the integral with its linear mean, $\bar{L} = (L_o + L_1)/2$, to get:

$$\begin{aligned} G(\pi_{L_0<L<L_1}) \cong{}& [1-\lambda(K)]\bar{L}[F(L|_{L>L_o}) \\ & - F(L|_{L>L_1})] - K, \end{aligned} \quad (5)$$

[7] See Section 1 for a discussion of the EV modeling assumptions and limitations, and how we propose to address these by measuring an upper and lower bound for our probability values.

where $F(L)$ is the "anti-cumulative" distribution function of L, i.e., 1-cumulative density function of L.

2.1. Optimum Investments

Maximizing the net gains from security investments, $Max_{(K)}\, G(\pi_{L_0<L<L_1})(K)$, leads to a value K^* that satisfies the first-order condition below:[8]

$$\lambda'(K^*)\bar{L}[F(L|_{L>L_o}) - F(L|_{L>L_1})] = -1. \quad (6)$$

Deriving an explicit expression for optimum capital expenditures, K^*, depends on the functional forms used. But such functions must all satisfy the criteria discussed earlier. If we use the following form,

$$\lambda(K) = \frac{1}{1+\theta K^{\alpha}} \quad \theta > 0,\ 0 < \alpha < 1, \quad (7)$$

then the requirement that $\lambda' < 0$, $\lambda'' > 0$, $\lambda(0) = 1$, and $\lim_{K\to\infty} \lambda(K) \longrightarrow 0$, are all satisfied. Several functional specifications for such a loss coefficient with decreasing marginal rate of return as specified above exist in literature.(10,14) In particular, Gordon and Loeb,(10) in relating reductions in security breach as a function of information security investments, have implied a version of $\lambda(K)$ that is similar to the one here. For our specification in Equation (7), the first-order condition yields:

$$(1+\theta K^{*\alpha})^2 = [F(L|_{L>L_o}) - F(L|_{L>L_1})]\bar{L}\theta\alpha K^{*\alpha-1}. (8)$$

We can use Equation (8) along with Equation (7) to both simplify the expression for K^* and also to substitute for the value of θ in terms of λ. The reason for the latter is that, as we shall see below, the empirical results from the "Benchmarking Survey" are closely associated with the values of λ rather than θ. Substituting from Equation (7) into Equation (8) we find:

$$\frac{1}{\lambda^2} = [F(L|_{L>L_o}) - F(L|_{L>L_1})]\bar{L}\theta\alpha K^{*\alpha-1}. \quad (9)$$

In turn, we find from Equation (7) that $\theta K^{\alpha} = (1-\lambda)/\lambda$. Substituting this value into Equation (9) we find the optimum level of investments K^* to be:

$$K^* = \alpha\bar{L}[F(L|_{L>L_o}) - F(L|_{L>L_1})]\lambda(1-\lambda). \quad (10)$$

3. ESTIMATION AND RESULTS

There are two values we need to estimate K^* in Equation (10): the probability values of $F(L|_{L>L_o})$ and the parameter λ. First, we will discuss estimating $F(L|_{L>L_o})$. These values are derived from data compiled by Mohtadi and Murshid.(6) This work compiles 448 incidents of chemical, biological, and radionuclear (CBRN) attacks from the 1960s to 2005. Mohtadi and Murshid(5,7,13) then calculates the probability of any such attack worldwide of a given magnitude or larger (in terms of the number of injuries) for different time horizons. These probability values are presented in the third column of Table I. The fourth and the fifth column then convert this anti-cumulative probability to the probability values $F(L|_{L_o<L<L_1})$ compatible with Equation (10). Although any attack on the food sector will be of a CBRN nature, not every CBRN attack involves food. In fact, only 60 of the CBRN incidents revolved around food. Although there have been attacks on our food and water supply that have involved the use of conventional weapons, there is no reason in particular why terrorists should favor the food supply chain over other potential targets when using such conventional means of attack.[9] The real threat as far as the food chain is concerned is likely to come from chemical, biological, or radionuclear contaminants, which can exploit an already present distribution network to maximize the potential for disruption. Conceptually, a CBRN attack on the food supply chain will take advantage of the carrying capacity of a highly diffused system by entering it and spreading itself across a wide geographic area. In this way, the risk spreads from one location to many. By contrast, a conventional attack on the food system (e.g., bombing a milk production site) while disrupting the supply chain is not capable of spreading because it is not embodied in the system. The added uncertainties that arise in the detection of biological and chemical agents may also imply that such agents may create havoc by simply the suspicion of their entry into the food system. Of the 448

[8] The second-order condition is satisfied because $\lambda'' > 0$.

[9] This statement is valid only in the absence of specific prior information regarding the likelihood of a terrorist intent on the food system. To make this more clear, consider the Bayes's Rule: with $P(.)$ denoting probability, we have, $P(F|CBRN) = \frac{P(CBRN|F)P(F)}{P(CBRN)} = \frac{P(F)}{P(CBRN)}$, where the second equality follows from the reasonable assumption that nearly all terrorist food events would have to be of CBRN type and thus $P(CBRN|F) = 1$. From this equation, it follows that $P(F) = P(F|CBRN).P(CBRN)$. In the absence of any prior information about a particular terrorist event, $P(F|CBRN)$ is *uniformly* distributed (i.e., minimally informative). Thus, $P(F|CBRN)$ can be estimated as $P(F|CBRN) = \frac{\text{\# of Food events}}{\text{\# of CBRN events}} = \frac{60}{448}$. We adjust $P(CBRN)$ by this factor to find $P(F)$. This is given in the last column of Table I.

Table I Probabilities of a CBRN Attack of Various Magnitudes (Based on Extreme Value Analysis) and Extrapolated Probabilities of Attacks on Food Sources

Number of casualties	Time horizon	Anti-cumulative prob. of CBRN at various casualty levels $F(L \geq L_0)$	# of casualties L_0 to L_1	prob. of a CBRN event between L_0 to L_1 $F(L \geq L_0) - F(L \geq L_1)$	Adjusted prob. for attacks on food
1000	Current risk	0.310	1000-5000	0.143	0.019
5000		0.167	5000-10000	0.034	0.005
10000		0.133	10000-15000	0.055	0.007
15000		0.078			
1000	5-year forecast	0.546	1000-5000	0.251	0.034
5000		0.295	5000-10000	0.071	0.009
10000		0.225	10000-15000	0.095	0.013
15000		0.130			
1000	10-year forecast	0.732	1000-5000	0.211	0.028
5000		0.520	5000-10000	0.110	0.015
10000		0.410	10000-15000	0.119	0.016
15000		0.291			
1000	20-year forecast	0.863	1000-5000	0.095	0.013
5000		0.768	5000-10000	0.057	0.008
10000		0.712	10000-15000	0.077	0.010
15000		0.634			

Source for Column 3: Mohtadi and Murshid (2006b) and additional extrapolations.

biological, chemical, and radiological incidents that we recorded, 75 involved either a direct attack or a plan to attack the food or water supply chains. The majority of these attacks (50 altogether) involved the use of chemicals, eight attacks were carried out using biological agents and one suspected incident involved the release of plutonium into New York City's water reservoirs. In the remaining attacks the type of agent is unknown. Most experts concede that the food chain, from production to distribution and processing, is highly exposed. Indeed, there have been attacks before. In 1984, for instance, the Rajneeshees—an Oregon-based cult—contaminated food at 10 restaurants with Salmonella typhimurium, causing 751 people to get sick.[15,16] Court testimony suggests that members of this cult considered various other, and more deadly, pathogens, including Salmonella typhi (which causes typhoid) and the AIDs virus.[16] Since then there have been a number of serious attacks on the food chain. Some of these are outlined in Table I, which chronicles all CBRN attacks between 1950 and 2005 that led to at least 50 casualties. There are 23 such cases, of which 12 involve an attack on food. Many of these attacks have taken place in China, where too often disputes have been settled through the use of Dushuqiang, a strong rat poison. The most serious such attack occurred, as mentioned previously, in September 2002 on a fast food restaurant that led to 41 fatalities and left over 400 ill. In the United States, the most serious attack on the food chain since the Rajneeshee incident occurred at the Family Fare supermarket in Byron Center, just outside Grand Rapids. In that incident, Randy Bertram poisoned about 250 pounds of ground beef, which caused at least 111 people to become ill.

As mentioned above, the probabilities values in Table I are from Mohtadi and Murshid.[7] These values represent the *average* of seven different models with different specifications. Following a discussion of Table I and derivation of optimal investments K^*, we will return to these probability values once more in Section 3.1 and examine what happens if different assumptions were made in the underlying EV modeling. We will consider the *upper* and *lower* probability bounds, based on alternative EV modeling specifications. From there, we will rederive the corresponding upper and lower bounds of security investments. For now, let us focus on Table I.

Several features of Table I deserve attention. First, there are two key features of the anti-cumulative probability distribution of an event associated with x number of casualties or greater, in Column 3 of Table I, that deserve a mention. One, for a given time horizon, the anti-cumulative probability

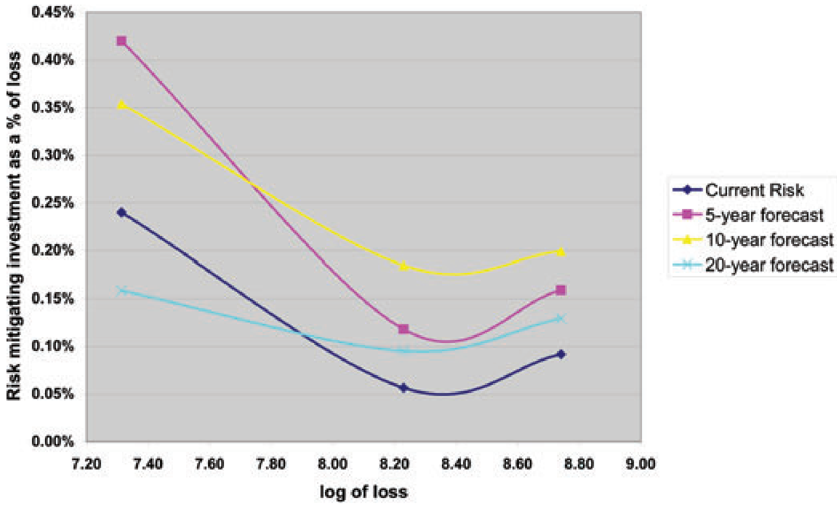

Fig. 2. Risk mitigating investments as a function of original loss.

falls, as one would expect: the more severe an event, the less likely it is to occur.[10] Two, moving across time horizons, the anti-cumulative probability for a given *x* rises with longer time horizons! This interesting but alarming result is a unique feature of the data and the fact that with the recent experience with terrorism and bioterrorism events, more severe events have become more likely in time (i.e., the tail parameter of the EV distribution shows a time trend). Thus, an event leading to casualty level of 15,000 or more has a chance of under 1% in the current time horizon, but over 63% in a 20-year time frame !

Second, there are also two key features of the incremental probability distribution associated with area segments between two loss levels in Column 5 that also deserve attention: one, for any given time horizon, the probability segments first fall with severity and then rise, exhibiting a U-like behavior. Closer scrutiny reveals that this pattern results because in Column 3 the decline in the event probability, from a 5,000+ to a 10,000+ casualty event, is much smaller than the decline from a 10,000+ to a 15,000+ casualty event. This pattern turns out to have significant implications for the amount of investment food companies have to make as a fraction of potential losses they would like to avoid, and produces a U-shaped curve (Fig. 2) that we will come to later. Two, moving across time horizons, with one exception, the probability segments for any given segment (say, L_1 to L_2) *rise* for the first three time horizons (current, five-year, 10-year), but *fall* for the last time horizon.[11] The reason for the fall seems to be the very large probability of the fat tail event in the last time horizon (15,000 or more casualties). This large probability of 63% seems to take away areas from the other segments in this time horizon, again consistent with the time trend we have found for more extreme events.

Next, the "# of casualties" (Table I, Column 4) must be translated into dollar loss. The overall data consist of 254 fatalities and 9,733 injuries or illnesses. Thus, on average, fatalities and injuries/illnesses constitute 0.025 and 0.975 of the total number of casualties, respectively. In estimating the dollar loss, there are two challenges that we must address. First, how to value loss of life; second, how to value an injury or illness. On the first question, an influential study by Kenkel available on the ERS website discusses three methods: the "willingness to pay" approach used predominantly by EPA, a combined "human capital/willingness to pay approach" used predominantly by USDA, and a "quality adjusted life year" approach used predominantly by FDA.[(17)] These three approaches yield three different estimates, the first estimate of which is dramatically larger than the second and third. They are, respectively, $6.3 million (in 1999 dollars), $997,000 (1995 dollars), and $840 (1999 dollars), on average. The average of these three values is $2.21 million, which we will use for our fatality values. The second challenge is how to estimate injuries or illnesses that result from an intentional CBRN attack on food. This component constitutes a

[10] However, we must add these tail events still turn out to be *fatter* than a normal density tail, a feature of these EV distributions.

[11] The minor exception is the first segment (1,000–5,000 casualties), which rises only for the first two time horizons (current and five-year) but not the third (or fourth).

very large fraction of our total observations. For this, we rely on the quality adjusted life estimate in the above study, which also provides a value of \$100,000 for each year that is lost to illness.(17) Because many of the injuries or illnesses that are associated with food events are rather dramatic and have effects that may last four to five years, we use a value of \$450,000 for a injury on average associated with 4.5 year loss due to an injury or illness. Thus, the overall weighted average associated with each casualty figure in Table I is about \$494,000, which we round to \$500,000.

The underlying assumption here has been that an injury or death will generate a liability for which the firm is ultimately responsible. We recognize that this is an overly simplified approach. For example, the economic impact on the firm might include such factors as reputation effects, loss of future markets, etc., and this impact might be vastly greater than the liability loss associated with a severe illness or death. The result will be that the size of the expected gain (expected loss avoidance) from risk mitigating investments by firms may be significantly larger than the values obtained here. Thus, our estimates will be only a conservative lower bound for what is actually needed.

Values of λ are estimated from the median response value of a survey of 131 food manufacturers.[12] The primary focus of this survey was to collect firm-level responses on their various security practices. This is discussed in detail elsewhere.(9) Two other recent studies on food security practices, one on food service operations in Kansas,(18) and one on resilience in the U.K. food and drink industry,(19) have also used survey data on security practices. However, their scope of focus and the methodology is different in that they have either a narrower focus (food service operations in schools and hospital in Kansas) or a qualitative case-study approach (detailed interviews with only 25–30 organizations in the U.K. food and drink industry). Although there were several sections to the questionnaire, the section that relates to our value of λ here captures the impact of changes in security investments of the firms on *outcome* measures such as whether security investments by the firms increased resilience, reduced risk profile, or reduced number of security incidents experienced by them.[13] This section entails a large number of questions for which responses are categorical and range from 1 to 5. One of the most significant series of questions involves questions of the following type that would allow for the estimate of a "response function" from risk mitigating investments to actual results. An example of these types of questions is the following question: "Our firm's security investment has resulted in . . . loss/shrink."

Responses to questions (to be inserted in the blank space above), ranged from the value of 1 (significantly reduced) to 5 (significantly increased). Overall, the responses indicated a significant reduction in loss incidents by firms that have made security investments (see below).

Some examples of such security practices are:

1. having processes in place to prevent (detect) (respond to) a contamination/security event in the supply chain;
2. the ability to track and trace products one supplier up and one customer down the supply chain;
3. using closed circuit television (CCTV) to monitor activities on loading docks;
4. using technology (e.g., X-ray, RFID, etc. . . .) to verify trailer or container contents;
5. existence of decision trigger points and/or automated response actions in the event of a contamination/security incident;
6. existence of a senior management position focusing on security (e.g., Director of Security, Chief Security Officer);
7. having information systems that allow rapid sharing of information to all employees in case of security incidents or contamination;
8. incorporating information on preventing a contamination/security incident into employee food protection training; etc.

It can be appreciated by the reader that no security investment could be expected to focus exclusively on reducing the chances of a CBRN attack or, for that matter, of purely intentional attacks because protection against intentional attacks is at the same time protection against unintentional events. Although strictly speaking, such multifunctionality of risk mitigating investments would complicate

[12] The original survey included food retailers in addition to food manufaturers. However, the greater variance in the retailer response, together with the "agency" problems that cast doubt on the retailer's share of responsibility when a catastrophic event impacts a line of product they carry, led us to focus on food manufacturers only.

[13] The firms ranged from a minimum of 0–100 U.S. employees and \$20 million in annual revenues to over 50,000 U.S. employees and over \$1 billion in annual revenues, with a media range of 1,000–5,000 employees and \$100–\$500 million annual revenues.

analytical modeling, we are forced to abstract from this complication. Furthermore, there are already examples of mitigation strategies implemented solely based on food defense concerns (high pasteurization time/temperature for fluid milk).[14]

The median value of response was 2 from the survey. With a value of 3 associated with "neither reduced nor increased loss," and a value of 1 associated with maximum reduction in loss, we map these, respectively, to a value of $\lambda = 0$ to $\lambda = 100\%$ or $\lambda = 1$.[15] In this mapping framework, a median value of 2 in the survey corresponds to $\lambda = 0.5$. Thus, we use this value for our analysis. We now have sufficient information to estimate K^*. For $\lambda = 0.5$, and for the various levels of injury and time horizons given and the associated probabilities in Table I, the optimum level of risk mitigating investments K^* can be calculated from Equation (10); using values of $\alpha = 0.5$). The results are reported in the Table II.

To conduct a sensitivity analysis of the results with respect to the value of λ, we note from Equation (10) that $\partial K^*/\partial\lambda \propto \lambda(1-\lambda)$ and is, therefore, at its zero variation with respect to λ when $\lambda = 1/2$. In fact, this *is* the value of λ that we have found from the above exercise. It follows that as long as any alternative mapping assumption produces a value of λ that is close to 0.5, the values of K^* will not change by much. Further, it must be added that even for values of λ far from 0.5, whereas K^* values will be different from the above, the actual behavior of K^* over time and across probability space remains the same.[16]

Table II is constructed as follows. The first two columns are from Table I. The third column (mean # of casualties) simply takes the mid point in the range of casualties from L_o to L_1. Column 4 translates mean # of casualties into dollars, based on the extensive discussion earlier on loss value of injuries/illnesses (97.5% of our observations) and fatalities (2.5% of our observations). Column 5 repeats the last column of Table I. Column 6 (optimum investments in \$ million) is based on the value of K^* calculated from Equation (10). Column 7 is obtained by dividing Column 6 into Column 4.

As can be seen from the last column of Table II, a one-time capital expenditure that is needed to mitigate the impact of an extreme event ranges from a maximum of 0.042% (or 42 cents per \$100.00), corresponding to a \$1.5 billion potential loss (in a five-year horizon), to a minimum of 0.06% (or 6 cents per \$100.00), corresponding to a \$3.75 billion potential loss (in the current horizon). The larger percentage of investments that is needed to address smaller losses seems puzzling at first. The key to the puzzle is the corresponding probabilities: the less extreme event has a higher risk mitigation percentage cost because it has a higher probability of occurrence (3.4%), whereas the more extreme event has a lower risk mitigation percentage cost because it has a lower probability of occurrence (0.5%; Column 5).

Another interesting feature of Table II, which is also reflected in Fig. 2, is the fact that for any given time horizon, investments in risk mitigation as a percent of the potential loss decline for a range of losses only to rise beyond that range, producing an inverted U pattern. As mentioned in conjunction with Table I, this pattern has to do with the probability segments involved in both tables: for any given time horizon, the probability segments first fall with severity and then rise, producing the U-shaped curve. Closer scrutiny revealed that this pattern results because the decline in the event probability, from a 5,000+ to a 10,000+ casualty event, is much *smaller* than the decline from a 10,000+ to a 15,000+ casualty event (See Table I, Column 3). This, in turn, has to do with the larger than normal likelihoods of low-frequency, high-impact "tail" events, captured in EV statistical analysis.[17] A second feature of Fig. 2 is the nature of the shifts in at least three of the four curves. Specifically, notice that at the high end of the risk scale, extreme events associated with high monetary loss or high casualty (right end of X-axis) become more costly with longer time horizons because they become more likely (recall the time trend associated with fat tails discussed earlier). A third feature of the

[14] We thank an anonymous referee for pointing this out.

[15] In this mapping framework the survey values of 4 and 5 would imply a negative λ! By focusing on the representative firm with a median value of 2 we avoid this complication. Moreover, a very small fraction of firms in the sample had answers in categories of 4 or 5. We can only assume that in such extreme cases other factors may have been at work to lead to such adverse results.

[16] We also conducted another exercise with the value of $\lambda = 0.67$, based on a mapping assumption that allowed only positive values of λ. The trends all remained the same, but the values of K^* were slightly smaller than the ones found in Table II, as would be expected, because a somewhat larger λ would mean a more efficient loss reduction coefficient.

[17] This result is best understood by comparing it against results that one would have obtained from a normal distribution. For a normal distribution, the tail is not only thin but also importantly it is *flat*. This means that for high impact low probability events, little incremental gain would be made by investing in security investments, (Equation (10)). This would have reduced the upward slope of the curves in Fig. 2 compared to what is seen in the present case.

Table II. Estimated Values of Secuity Investment for Different Loss Levels and Time Horizons

	Range of casualties L_o to L_1	Mean # of casualties	Monetary value of loss from casualties (in $ million)	Probability of loss between L_o and L_1	Optimum investments K* in $ million	Optimum investments K* as % of loss	Log of loss
current risk	1000–5000	3000	1,500.00	0.019	3.60	0.24%	7.31
	5000–10000	7500	3,750.00	0.005	2.12	0.06%	8.23
	10000–15000	12500	6,250.00	0.007	5.74	0.09%	8.74
5-year forecast	1000–5000	3000	1,500.00	0.034	6.30	0.42%	7.31
	5000–10000	7500	3,750.00	0.009	4.43	0.12%	8.23
	10000–15000	12500	6,250.00	0.013	9.94	0.16%	8.74
10-year forecast	1000–5000	3000	1,500.00	0.028	5.31	0.35%	7.31
	5000–10000	7500	3,750.00	0.015	6.93	0.18%	8.23
	10000–15000	12500	6,250.00	0.016	12.45	0.20%	8.74
20-year forecast	1000–5000	3000	1,500.00	0.013	2.38	0.16%	7.31
	5000–10000	7500	3,750.00	0.008	3.57	0.10%	8.23
	10000–15000	12500	6,250.00	0.010	8.08	0.13%	8.74

Source: Based on Table 1 and analysis in the text.

figure, namely, the crossing of the curves, is not as critical as it simply reflects the various tradeoffs between time horizons and event likelihood that occurs at the *intermediate* segments of risk.

3.1. Upper and Lower Bounds

As stated, the main probability values in Column 3 of Table I are selected by averaging seven different models originally derived in Mohtadi and Murshid.[7] These models make different assumptions about the underlying EV parameter estimates. Here, we focus on two of the models that produce the two extremes, i.e., the lowest and the highest, of the probability forecasts. The main difference between these two models arises from assumptions about the stationarity of the underlying EV distribution: the lower probability bound belongs to a model that assumes that only the "location" parameter in the underlying distribution changes over time. The upper probability bound belongs to a family of models that assumes a shift in both the location and the scale parameters of the underlying distribution over time. These estimates are presented in Table III.

Although the likelihood ratios and the QQ plots tend to favor the family of distributions that produce the upper-bound likelihoods, one cannot entirely rule out the lower-bound likelihoods. One striking feature of the upper-bound values in Table III (last column) is that in a 20-year horizon the likelihood of a CBRN-based bioterrorism attack is virtually certain. Even after these values are adjusted to reflect the chance of an attack on food, the constancy of the values in the case of a 20-year horizon will have an important implication for firms' incentives to invest in security. This will be discussed below in connection with Table IV. In this table, the resulting optimal security expenditures for the lower and upper

Table III. Lower & Upper Bounds for Probability of CBRN Attacks (from Table I)

Number of casualties	Time horizon	Anti-cumulative prob. of CBRN at various casualty levels $F(L \geq L_0)$	Upper & lower probability bounds based on different EVT models: Lower bound	Upper bound
1000	Current risk	0.310	0.210	0.400
5000		0.167	0.080	0.310
10000		0.133	0.050	0.280
15000		0.078	0.021	0.224
1000	5-year forecast	0.546	0.350	0.770
5000		0.295	0.140	0.380
10000		0.225	0.090	0.340
15000		0.130	0.039	0.198
1000	10-year forecast	0.732	0.570	1.000
5000		0.520	0.250	0.790
10000		0.410	0.170	0.620
15000		0.291	0.081	0.473
1000	20-year forecast	0.863	0.670	1.000
5000		0.768	0.660	1.000
10000		0.712	0.490	1.000
15000		0.634	0.425	1.000

Source: Mohtadi and Murshid (2006b) and additional extrapolations

Table IV. Optimal Security Expenditures for "Base Values" (Table II) and Lower and Upper Bound Probability Values

	Range of casualties L_o to L_1	Mean # of casualties	Loss value from casualties ($ million)	Prob. of loss L_0 to L_1 base value	K* ($ million)	K* as % of loss	Prob. of loss L_0 to L_1 lower bound	K* ($ million)	K* as % of loss	Prob. of loss L_0 to L_1 upper bound	K* ($ million)	K* as % of loss	Log of loss
Current	1000–5000	3000	1,500.00	0.019	3.60	0.24%	0.017	3.26	0.22%	0.012	2.26	0.15%	7.31
Risk	5000–10000	7500	3,750.00	0.005	2.12	0.06%	0.004	1.88	0.05%	0.004	1.88	0.05%	8.23
	10000–15000	12500	6,250.00	0.007	5.74	0.09%	0.004	3.06	0.05%	0.007	5.85	0.09%	8.74
5-year	1000–5000	3000	1,500.00	0.034	6.30	0.42%	0.028	5.27	0.35%	0.052	9.79	0.65%	7.31
Forecast	5000–10000	7500	3,750.00	0.009	4.43	0.12%	0.007	3.14	0.08%	0.005	2.51	0.07%	8.23
	10000–15000	12500	6,250.00	0.013	9.94	0.16%	0.007	5.32	0.09%	0.019	14.86	0.24%	8.74
10-year	1000–5000	3000	1,500.00	0.028	5.31	0.35%	0.043	8.03	0.54%	0.028	5.27	0.35%	7.31
Forecast	5000–10000	7500	3,750.00	0.015	6.93	0.18%	0.011	5.02	0.13%	0.023	10.67	0.28%	8.23
	10000–15000	12500	6,250.00	0.016	12.45	0.20%	0.012	9.34	0.15%	0.020	15.42	0.25%	8.74
20-year	1000–5000	3000	1,500.00	0.013	2.38	0.16%	0.001	0.25	0.02%	0.000	0.00	0.00%	7.31
Forecast	5000–10000	7500	3,750.00	0.008	3.57	0.10%	0.023	10.67	0.28%	0.000	0.00	0.00%	8.23
	10000–15000	12500	6,250.00	0.010	8.08	0.13%	0.009	6.80	0.11%	0.000	0.00	0.00%	8.74

probability bounds are obtained by repeating the previous procedure (see Table II). For comparison, the base values are also included.

First, we note that even though the lower and upper bound CBRN probability values in Table III "encompassed" the base values, as they should, this is not the case for Table IV. That is, there are some probability values for which the base case falls outside the range defined by the lower and upper bound probability values. (Compare Column 5 with Columns 8 and 11 in Table IV.) This is due to the fact that the probability values in Table IV are not probability *levels*, but probability *increments* associated with losses between two points (L_0 and L_1).

The above observation has an important implication for optimal security investments, K^*. Although the results for K^* for the case of lower bound probabilities are not too different from those for the "base" probabilities, the K^* values for the upper bound probabilities, for the 20-year risk horizon, are dramatically different and in fact equal zero for all the three risk categories in that time horizon (the bottom 3 rows of Column 11). Why is this so? The answer is that because at this upper end, the chance a CBRN event in any of the three magnitudes is virtually certain in a 20-year horizon (probability = 1 from Table III), even after mapping this probability to the probability of food attacks, the resulting probabilities remain constant, even if less than 1. As such, the incremental probabilities become zero due to the constancy of the probability levels. Thus, no security investment by firms can reduce the chance of a loss for this time horizon and, therefore, there is no incentive to undertake it (in Equation (10) $K^* \cong 0$ as $F(L|_{L>L_o}) \cong F(L|_{L>L_1})$). In effect, when the loss is extremely large there is very little reduction in the loss that can be achieved by security investments (i.e., the increments in loss reduction are extremely small). To elaborate, suppose that an equivalent of a 9/11-size attack or larger were to occur in the food industry by, say, an intentional act of introducing ricin or a deadly form of *Escherichia coli* into milk at many nodes across the food supply chain. If such an event were to occur, no single firm would be able to affect the risk of such a massive loss. As such, though public authorities may be able to affect aggregate risk collectively (a point we do not address in this article), no risk mitigating investment would be effective by any single firm in isolation (there is zero net marginal gain of such an investment in terms of loss reduction.) Naturally, what the optimal level of investments should be—i.e., based on lower-bound, average, or upper-bound probabilities—will depend on what modeling assumptions we believe in. Although for a distant time horizon of 20 years, the results are highly sensitive to these modeling assumptions, for shorter time horizons, there is somewhat smaller variance in the value of optimal investments, at least for some of the risk categories, suggesting a degree of consistency in the results. Nonetheless, one must interpret these results with some caution and understand that there are inherent underlying limitations in the data and in the modeling assumptions.

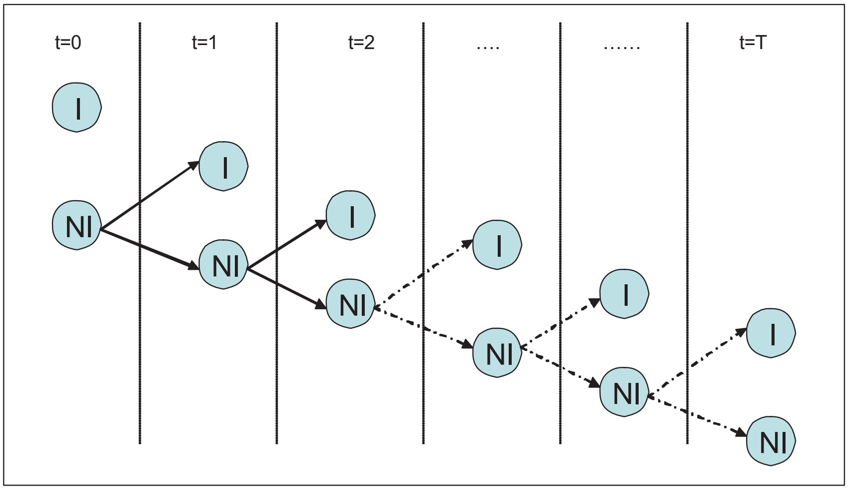

Fig. 3. Decision tree for optimum investments.

4. A MULTIPERIOD EXTENSION

Firms may choose to postpone investing in risk mitigation. The question that arises is whether such a strategy is an optimal strategy or not. To answer this question we consider a multiperiod extension of the previous model. Consider a firm with a finite time-horizon, say T, with the option to invest at the beginning of each time period, t, or to postpone the decision to future time periods. Because investing is an absorbing state, the decision tree of the firm can be represented by the chart in Fig. 3. In turn, this can be expressed, for any time t, as the number of periods, $t-1$, the firm has not invested, followed by the number of period $T-t$ starting from the time that the firm committed itself to security investments. This is presented in the Fig. 4.

Letting the probability of a finite loss, L, occurring in period t be given by π_t, the expected losses/gains arising from making the security investments in t must include the cost of investments, K, assumed to be made at the beginning of the time period, t, and the reduced losses, $\lambda(K)L$. Based on this and on Fig. 6, the present value of expected losses upon investing are given by:

$$PVE(L|\pi_t)^{invest} = \sum_{x=1}^{x=t-1} \delta^x \pi_x L + \sum_{x=t}^{x=T} \delta^x \pi_x \lambda(K) L + \delta^{t-1} K. \tag{11}$$

The present value of expected losses when no investments are made at all, is given by:

$$PVE(L|\pi)^{no\ invest} = \sum_{x=1}^{x=T} \delta^x \pi_x L. \tag{12}$$

Thus, the present value of expected net gain from investing, as arrived at in Equation (3) earlier, is given by:

$$\begin{aligned} PVE(G) &= PVE(L|\pi_t)^{no\ invest} - PVE(L|\pi)^{invest} \\ &= \sum_{x=1}^{x=T} \delta^x \pi_x L - \sum_{x=1}^{x=t-1} \delta^x \pi_x L - \sum_{x=t}^{x=T} \delta^x \pi_x \lambda(K) L - \delta^{t-1} K \\ &= \bar{L} \sum_{x=1}^{x=t-1} \delta^x \pi_x - \lambda(K) \bar{L} \sum_{x=t}^{x=T} \delta^x \pi_x - \delta^{t-1} K \\ &= (1-\lambda(K)) \sum_{x=t}^{x=T} \delta^x \pi_x L - \delta^{t-1} K. \end{aligned} \tag{13}$$

The firm must maximize PVE(G) by optimally choosing the investment level K^*, and the critical time period t^* at which to invest:

$$\underset{\{K,t\}}{Max} PVE(G) \Rightarrow t^*, K^*. \tag{14}$$

Consider a 10-year time horizon from the previous section. The probabilities of attacks on food supply for this horizon were given in Table II. From that table, the probability values for current risk, five-year risk forecast, and 10-year risk forecast are used to interpolate and obtain the values of π_t for $t = 1 \ldots 10$. This assumes that on average there are no further anticipated changes in the probability of catastrophic events occurring. There could certainly be changes in the probabilities contingent upon new information, but interpolating values of the forecasts implies that all available information is incorporated into the forecasts at the time the decision is made.

Parallel to Section 3, for catastrophic losses of sizes in the range L_0 to L_1, the EV probability

Fig. 4. Investment decision under alternative dynamic scenarios.

Period						
t=1	Not Invest	Not Invest	Not Invest	....	Not Invest	Invest
t=2	Not Invest	Not Invest	Not Invest		Not Invest	Invest
t=3	Not Invest	Not Invest	Not Invest		Invest	Invest
.	.	.	.		.	.
.	.	.	.		.	.
.	.	.	.		.	.
t=T-1	Not Invest	Not Invest	Invest		Invest	Invest
t=T	Not Invest	Invest	Invest	....	Invest	Invest

distribution function is obtained from Table II representing π_x. As before, we replace L with its linear mean, $\bar{L}$ yielding the present value of expected gains from investing in period t as:

$$PVE(G) = (1 - \lambda(K))\,\bar{L}\sum_{x=t}^{x=T}\delta^x\pi_x - \delta^{t-1}K. \quad (15)$$

First, let us focus on the optimum investment level $K^*(t)$. For this, we have the analytical solution as follows:

$$-\lambda'(K^*)\bar{L}\sum_{x=t}^{x=T}\delta^x\pi_x = \delta^{t-1}. \quad (16)$$

Using the form of $\lambda(K)$ as given by Equation (7), λ' can be expressed in terms of λ yielding:

$$K^*(t) = \lambda(1-\lambda)\bar{L}\sum_{x=t}^{x=T}\delta^{x-t+1}\pi_x. \quad (17)$$

Next, let us consider the optimum timing of such investments. In principle, maximizing for t per 14 should yield an equation $t^*(K)$, which together with Equation (17) could be solved to yield t^* and K^*. However, in practice, $t^*(K)$ may only be obtained numerically. We conducted this exercise for a 10-year horizon and for the loss values and their corresponding probability values given by Table II. Surprisingly, we found that unless the values of K are exceedingly larger than in the allowable range for which $K^*(t)$ with $t = 1\ldots10$ (see table below for this range), *the optimum timing t^* is always the very first period*. The estimation results are shown in Table V.

To elaborate on this result, notice also that the corresponding optimum $K^*(t)$ values are the maximum for these *first* periods. This is also shown in Fig. 5. Despite these large expenditures of capital, it is always better to invest in protection against catastrophic risk at the *beginning* of the period, rather than postpone such investments, as shown by the fact that net profits are largest (least negative) when $t = 1$ (last column).

Although larger $K^*(t)$ for smaller t values can be explained by lower present values of the cost of a fixed one-time investment compared to the gains (in terms of reduced risk) over repeated periods, the

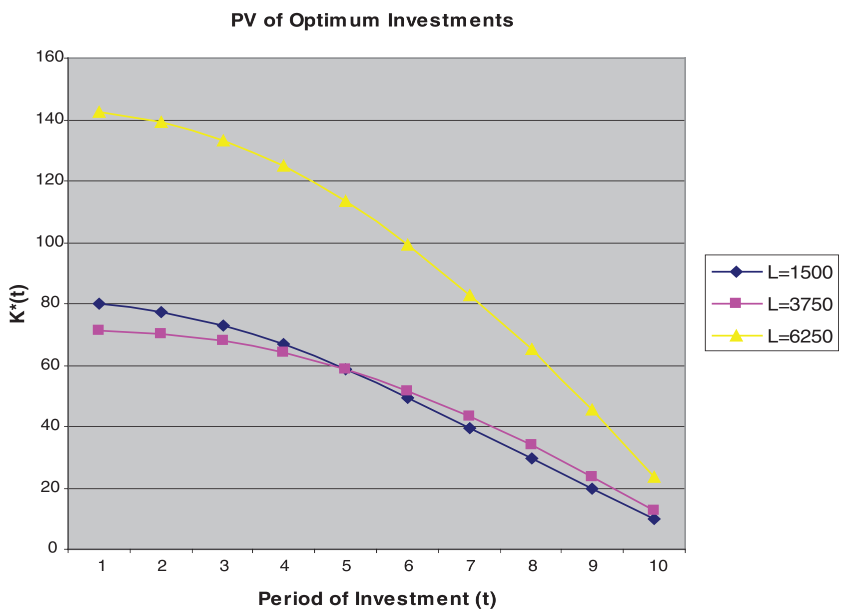

Fig. 5. PV of expected gains from investments.

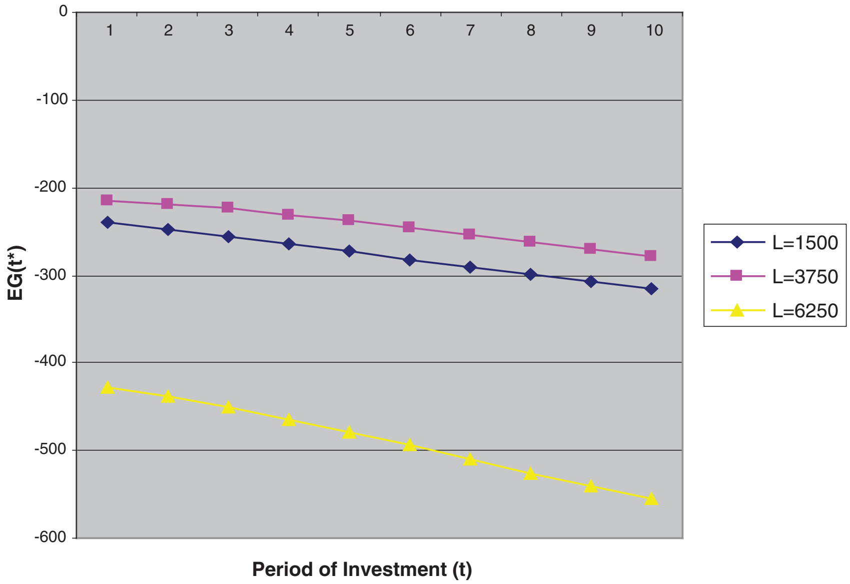

Fig. 6. PV of expected gains from investments.

case of $t^* = 1$ requires some further elaboration. Focusing on Equation (15), evidently, as the period of investment is postponed (moves closer to the time horizon T—in this case $T = 10$), the ongoing nature of the risk is sufficiently overwhelming as to favor an early investment decision rather than a later one. Interestingly, we tried many different parameters and the $t^* = 1$ results remained quite robust. This result is depicted in Fig. 6.

Finally, it is instructive to compare the one-period and multiperiod models with respect to the size of risk mitigating investments as a *percentage* of the initial loss. Comparing the corresponding columns in both Tables II and V (next to the last column in both tables), we see that, as a fraction of the original loss, risk mitigating investment levels in the multiperiod model are of the same order of magnitude as those in the one-period model. Notice, however, that the actual level of investments in this case is substantially larger (of the order of 10 times). This is, of course, not surprising because $K^*(t)$ in the multiperiod model mitigates the risk of a repeated exposure to risk over the entire 10-year horizon.

4.1. Upper and Lower Bounds

As in the case of the static model, it is important to examine the implication of considering lower and upper probability bounds for the dynamic decision horizon. As in the base model, the behavior of dynamic decision horizon is similar and suggests that greatest protection is afforded by early investments (Figs. 7(c) and (d)), even though this would be costlier (Figs. 7(a) and (b)).

5. COMPARISON WITH CATASTROPHIC INSURANCE: CONCLUDING REMARKS

How do expenditures aimed at risk mitigation compare with the purchase of catastrophic risk insurance? To gain some insights into the answer to this question consider the fact that in 2002 the Insurance Service Office assigned insurance costs of approximately 10 cents per $100 loss for the highest risk cities, but after discussions with the regulators, these rates were later adjusted downwards to less than 3 cents per $100 of loss (see Auerswald *et al.*,[2] p. 283). This suggests that loss mitigation costs are as high or higher than the cost of catastrophic insurance. Although risk mitigation and insurance are not entirely comparable strategies due to continued risk exposure under risk mitigation strategies, still the high deductibles of catastrophic insurance make the two approaches more comparable than might appear at the first glance. In this respect, several points need to be stressed.

First, risk mitigation investments are one-time actions. This is true both in the single-period and the multiperiod model. For example, to mitigate the risk of an event that could cause $3.75 billion loss in the current time horizon, a one-time capital expenditure of 6 cents per $100 under the base

Table V. Estimated Present Values (PV) of Secuity Investment for Different Loss Levels and Time

Range of casualties L_o to L_1	Mean # of casualties	Monetary value of loss from casualties (in \$ million)	PV of loss (assume r = 5%)	Probability of loss between L_o and L_1	Investments made at the beginning of period t	PV of optimum investments $K^*(t)$ in \$ million	PV of optimum investments $K^*(t)$ as % of loss	PV of expected gains from optimum investments $E(G(K^*(t)))$
1000-5000	3000	1,500.00	11,436.00	0.0190	1	80.16	0.15%	−240.48
				0.0228	2	77.26	0.15%	−247.25
				0.0265	3	72.79	0.14%	−254.95
				0.0303	4	66.68	0.13%	−263.47
				0.0340	5	58.85	0.12%	−272.71
				0.0328	6	49.20	0.10%	−282.58
				0.0316	7	39.49	0.09%	−291.62
				0.0304	8	29.71	0.07%	−299.90
				0.0292	9	19.88	0.05%	−307.46
				0.0280	10	9.98	0.02%	−314.36
5000-10000	7500	3,750.00	28,589.99	0.0050	1	71.34	0.15%	−214.01
				0.0065	2	70.40	0.15%	−218.46
				0.0080	3	68.01	0.14%	−223.96
				0.0095	4	64.09	0.13%	−230.39
				0.0110	5	58.56	0.12%	−237.64
				0.0116	6	51.33	0.10%	−245.62
				0.0122	7	43.16	0.09%	−253.62
				0.0128	8	33.99	0.07%	−261.60
				0.0134	9	23.78	0.05%	−269.56
				0.0140	10	12.47	0.02%	−277.48
10000-15000	12500	6,250.00	47649.98847	0.0070	1	142.72	0.15%	−428.15
				0.0085	2	139.29	0.15%	−438.54
				0.0100	3	133.34	0.14%	−450.53
				0.0115	4	124.73	0.13%	−463.92
				0.0130	5	113.33	0.12%	−478.56
				0.0136	6	98.98	0.10%	−494.28
				0.0142	7	82.94	0.09%	−509.90
				0.0148	8	65.12	0.07%	−525.39
				0.0154	9	45.42	0.05%	−540.73
				0.0160	10	23.75	0.02%	−555.90

model, or 5 cents per \$100 under the upper and lower bound models, from Table IV, pay off when compared with the purchase of catastrophic insurance after the fifth year. The key of course is that risk mitigation is not foolproof so that risk is not eliminated, but only reduced. As mentioned, however, this fact needs to be balanced against the fact that catastrophic risk insurance often entails large deductibles and, therefore, at least up to the deductible level, risk mitigation expenditures remain relevant whether or not catastrophic insurance is purchased.

Second, the availability of catastrophic insurance has become more limited after the 9/11 attacks, due to a variety of complicating factors, including (i) the evidence of an increase in risk trend[5] as discussed by Bogen and Jones,[20] (ii) the large size of the losses, requiring far greater pooling of resources than has been available to insurers and reinsurers, (iii) the uncertainties associated with calculating the probabilities of catastrophic insurance,[18] and (d) the asymmetric information between insurance firms and the insured on the one hand, and the reinsurance firms and insurance firms on the other, resulting in a double moral hazard problem.[2]

Third, as stated previously, issues such as reputation effects cannot be overlooked. These affects cannot be easily addressed by the purchase of insurance, especially if they are the outcome of inadequate risk mitigation by firms in the first place.

[18] On this point, we hope that this contribution will be of value. We argue that terrorism risk is quantifiable in the same manner as, for example, weather risk may be.

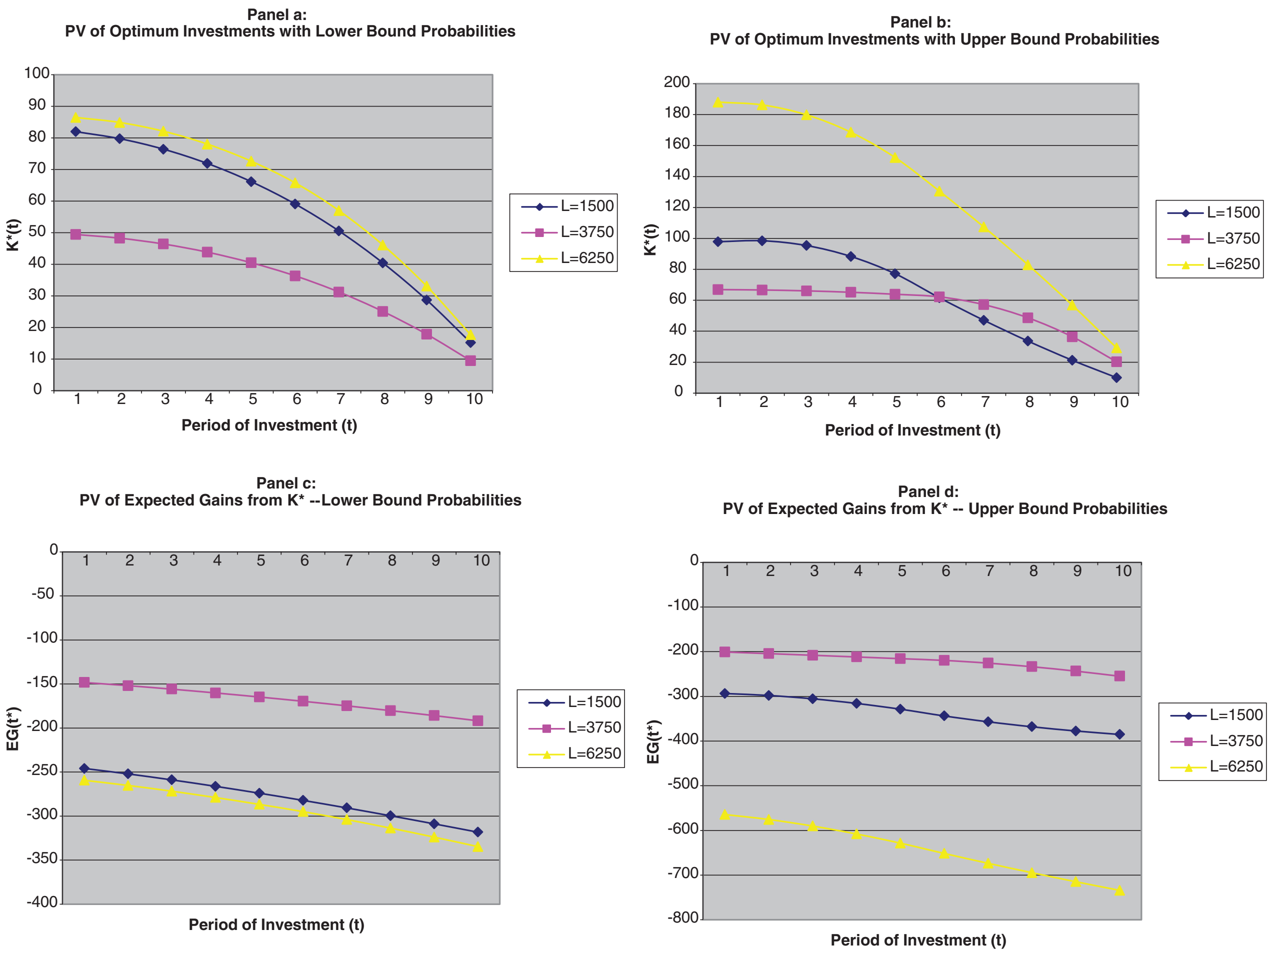

Fig. 7. Upper and lower bounds for security investments and optimal gains in dynamic model.

Fourth, the adequate coverage of catastrophic insurance together with the large potential losses and the associated "risk externalities" call for the involvement of the federal, state, and local governments (see Auerswald *et al.*,(2) p. 284), as this becomes a clear case of market failure. In fact, the Terrorism Risk Insurance Act (TRIA) and its 11th-hour extension in December 2005 was aimed at addressing the gaps in this market. However, many obstacles remain and catastrophic insurance remains a highly imperfect tool with limited or no availability in many cases. Thus, until and unless these issues are resolved, risk mitigation becomes an important and essential strategy.

Fifth, though risk mitigation is not well correlated with insurance costs, due to moral hazard problems associated with asymmetric information between the insured and the insurer (again see Auerswald *et al.*,(2) p. 283), it is likely that catastrophic risk mitigation will have a positive effect on the cost of risk financing, if not risk insurance. For example, Moody's risk ratings are highly affected by the ability of firms to prepare for and respond to risk. For all these reasons, the importance of catastrophic risk mitigation strategies cannot be overstated and such strategies must be an essential part of firms' overall strategy.

ACKNOWLEDGMENTS

This article was presented at the American Economic Association's annual meetings, January 2008, New Orleans, LA. A static version of this article was presented at the annual meeting of the Sloan Industry Group, April 25–27, 2007, Cambridge, MA. We thank Walter Enders, Howard Kunreuther, Erwann Michel-Kerjan, and Larry Samueslson for their comments and feedback. We also thank two anonymous referees for their constructive comments. This research is supported by the U.S. Department of Homeland Security (Grant number N-00014-04-1-0659), through a grant awarded to the National Center for Food Protection and Defense at the University of Minnesota. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not represent the policy or position of the Department of Homeland Security.

REFERENCES

1. Doherty E, Goldsmith S, Harrington P, Kleindorfer E, Kunreuther H, Michel-Kerjan M, Pauly MV, Rosenthal I, Schmeidler P. TRIA and Beyond: Terrorism Risk Financing in the US. Philadelphia, PA: Report issued by Wharton Risk Management and Decision Process Center, University of Pennsylvania, August 1, 2005.
2. Auerswald P, Lewis B, La Porte T, Michel-Karjan E. Seeds of Disaster, Roots of Response: How Private Action Can Reduce Vulnerability. Cambridge, UK: Cambridge University Press, 2006.
3. Kunreuther H, Michel-Kerjan E, Porter B. Assessing, Managing, and Financing Extreme Events: Dealing with Terrorism. NBER Working Paper No. 10179, December 2003.
4. Johnson N, Spagat M, Restrepo JA, Bohórquez J, Suárez N, Restrepo E, Zarama R. From Old Wars to New Wars and Global Terrorism. Working Paper, Universidad Javeriana–Bogotá, Documentos de Economía 002339, 2005.
5. Mohtadi H, Murshid AP. The risk of catastrophic terrorism: An extreme value approach. Journal of Applied Econometrics, 2009; 24(4): 537–559.
6. Mohtadi H, Murshid AP. A Global Chronology of Incidents of Chemical, Biological, Radioactive and Nuclear Attacks: 1961–2005. Minneapolis, MN: University of Minnesota National Center for Food Protection and Defense (NCFPD), 2006.
7. Mohtadi H, Murshid AP. A risk analysis of chemical, biological, or radionuclear threats: Implications for food security. Risk Analysis, 2009; 29(9): 1317–1335.
8. Michel-Kerjan E. Large scale terrorism: Risk sharing and public policy. Revue d'Economic Politique, 2003; 113(5): 625–648.
9. Agiwal S, Mohtadi H, Kinsey J. Resilience and Catastrophe: How Do Firms and Supply Chains Cope? Working Paper, University of Minnesota, St. Paul, MN, 2007.
10. Gordon L, Loeb MP. The economics of information security investment. ACM Transactions on Information and System Security, 2002; 5(2): 438–457.
11. Coles SG. An Introduction to Statistical Modeling of Extreme Values. London, UK: Springer Verlag, 2001.
12. Embrechts P, Klüppelberg C, Mikosch T. Modelling Extremal Events for Insurance and Finance. Berlin, Germany: Spring Verlag, 1997.
13. Mohtadi H, Murshid AP. Analyzing catastrophic terrorist events with applications to the food industry. Pp. 182–207 in Richardson HW, Gordon, Peter, Moore, JE II (eds). The Economic Costs and Consequences of Terrorism. Cheltenham, UK: Edward Elgar, 2007.
14. Hausken K. Returns to information security investment: The effect of alternative security breach functions on optimal investment and sensitivity to vulnerability. Information Systems Frontier, 2006; 8(5): 338–349.
15. Carus W. The Rajneeshees (1984). Pp. 115–137 in Tucker JB (ed). Toxic Terror: Assessing Terrorist Use of Chemical and Biological Weapons. Cambridge, MA: MIT Press, 2000.
16. Carus W. Bioterrorism and Biocrimes: The Illicit Use of Biological Agents Since 1900. Amsterdam, The Netherlands: Fredonia Books, 2002.
17. Kennel D. Using Estimates of the Value of a Statistical Life in Evaluating Regulatory Effects Economic Research Service. In Kachler F (ed). Valuing the Health Benefits of Food Safety. Washington, DC: USDA (ERS), April 2001. Available at http://www.ers.usda.gov/publications/mp1570/mp1570d.pdf.
18. Yoon E, Shanklin C. Food security practice in Kansas schools and health care facilities. Journal of American Dietetic Association, 2007; 107(2): 325–329.
19. Peck H. Resilience in the UK Food and Drink Industry. Shrivenham, UK: Canfield University. Final Report to the Department for Environment, Food and Rural Affairs, Department of Defence Management and Security Analysis, 2006.
20. Bogen K, Jones E. Risks of mortality and morbidity from worldwide terrorism: 1968–2004. Risk Analysis, 2006; 26(1): 45–59.