



The Ecology of Terror Defense

NATHANIEL O. KEOHANE
Yale School of Management

RICHARD J. ZECKHAUSER* richard_zeckhauser@harvard.edu
John F. Kennedy School of Government, Harvard University, 79 John F. Kennedy Street,
Cambridge, MA 02138, USA

Abstract

We draw on an ecological metaphor to analyze terror defense by governments and individuals. Governments can combat terror in part by targeting what we call the “stock of terror capacity” accumulated by a terrorist organization. The optimal control of terror stocks relies on both periodic cleanup and ongoing abatement, a lesson derived from the optimal control of other stocks of public bads, such as pollution.

Citizens can protect themselves from terror by avoiding exposure to threats and by reducing the harm they suffer if an attack occurs—actions we call “avoidance” and “amelioration,” respectively. Such measures may exert positive or negative externalities on other individuals. A simple model shows how individual responses to collective threats may undermine the effectiveness of government policies to address such threats. The model uncovers a strong analogy between the problem of individual responses to terror and the familiar congestion externality.

Keywords: terrorism, stock control, economies of scale, avoidance, security externalities

JEL Classification: D62, H41, H56

1. Introduction

From the close of World War II to the collapse of the Soviet Union, the threat of massive retaliation protected the United States and its allies against external enemies. Though the concept of threats has been known for millennia, significant progress on models analyzing their use was made in the late 1950s and early 60s.¹ A threat is costly to the actor carrying it out, but is usually more costly to the other side. An actor has a variety of mechanisms to make his threat credible, including putting his reputation on the line to carry it out if the other side misbehaves.

This work on threats, and the use of threatened retaliation to encourage good behavior, grew out of the recognition that the world was confronting a previously unseen situation. The two opposing sides in the Cold War each had significant military, but that was nothing new. What was new was their weapons of mass destruction; they could impose vastly unacceptable tolls on each other.

Today’s world, where terrorism is the major threat to developed nations, offers dramatically new strategic realities.² Threatened retaliation is no longer feasible as the principal

*To whom correspondence should be addressed.

mechanism for containing enemies, if those enemies are terrorists. The threatened parties often cannot be found and may possess little in the way of conventional targets to be destroyed. To the extent that terrorists have hard assets (such as training camps, trained personnel, or weapons factories and stores), those assets are primarily a means to foment terror, rather than being valued in themselves. This implies that the threat of widespread asset destruction is unlikely to be a strong deterrent to terror. In the extreme, terrorists may even welcome the martyrdom that would come with death. Hence, we believe that to understand how to cope with terrorism, a new model is required.³ This paper takes a first step in that direction. It invokes the metaphor of an ecosystem. Two major elements comprise this heuristic metaphor.

First, from a government's perspective, terror can fruitfully be viewed as the product of a "stock" of terror capacity—much as many environmental ills are the product of a stock of pollution. Both stocks produce bads. Of course, the analogy is far from perfect. Most obviously, terrorists are strategic, while natural processes are not.⁴ In addition, though the output from either stock is uncertain, that from the terror stock is dramatically more so. Nonetheless, the analogy offers some insight into the policies governments can pursue in defending their citizens against terror. Drawing on previous work in the environmental realm, we argue that the optimal counterterrorism policy ought to include efforts aimed at both reducing the flow of new capacity by terrorist organizations and directly reducing the stock of capacity that already exists. Such reductions offer economies of scale over some range, e.g., launching a sufficient operation to wipe out or neutralize some terror organization. The optimal balance to be struck between controlling the flow and controlling the stock will be determined by the relative costs of those measures as well as by how fast the stock of terror is growing.

These efforts aimed at controlling the stock and flow of terror capacity are complemented by two other defensive measures available to governments: *averting actions*, which reduce the likelihood of a terror attack; and *amelioration*, which lessens the damages in the event of an attack. Like a prudent investor, the government should diversify its portfolio by incorporating all of these approaches. In this paper, we discuss all four tools but focus primarily on government efforts that seek to control the stock and flow of terror capacity.

The government takes actions on a public goods scale. In contrast, individual entities have the potential to protect themselves or those with whom they have contractual relationships, but their actions do not affect the terrorist or the terror capacity. Thus building owners can reinforce their structures or inspect visitors. Individuals can refuse to fly, stay away from crowds, or stock their own water supplies.

The second component of the ecological metaphor concerns the efforts of individual entities to defend themselves, much as organisms in an ecosystem struggle for survival but in doing so affect the survival of others. We develop a simple model of a population of individuals who are able (at some cost) to avoid a collective threat from terrorist activity, or to reduce the harm they will suffer given an attack. Such individual precautions have external effects: one person's decision to avoid a terrorist threat—by moving a corporate office to the suburbs, for example, or by avoiding a crowded public space that might be a target of terror—affects the safety of those who remain. This security externality turns out to have the same structure as a congestion externality, with similar implications for policy.

A striking conclusion from our model is that individual responses to government policies may undermine the social welfare purposes of those policies.

Our double-edged treatment of terror defense relates to disparate literatures. The model of the control of stocks and flows of terror capacity employs a model originally motivated by stock pollutants and more generally by the management of deteriorating capital stocks (Keohane, Van Roy, and Zeckhauser, 2002, henceforth KVZ). That paper, in turn, builds on a substantial body of work in environmental economics and in operations research, some of which we reference below.

The model of individual precautions against terror, meanwhile, focuses on security externalities generated by individual actions to avoid or diminish the harm from terrorist attacks. As noted in that section, our discussion has natural connections to the literature on crime prevention. Our treatment of individual responses to terror also complements concurrent work by Howard Kunreuther and Geoffrey Heal, who contemplate the externalities of protective measures when risks are interdependent (Heal and Kunreuther, 2002; Kunreuther and Heal, 2002). The security externality they identify results from the structure of relationships among potential victims. In contrast, the externalities we consider are driven by the relationship between the behavior of potential victims and the actions of terrorists. Hence our externalities flow from the strategic nature of the terrorist threat, although we do not need to model the terrorist's actions explicitly to derive the key results.

The next section develops our "capital stock" model of terror and considers the policy portfolio available to governments in defending their citizens against terrorism. Section 3 shifts gears to consider the efforts of individual agents to protect themselves from attack. Section 4 concludes. Throughout the paper, our discussion emphasizes the frameworks for considering terror defense; mathematical details are kept to a minimum or referenced from other works.

2. Controlling the "stock" of terror

In this section we propose a framework for considering policies by governments to defend their citizens against terror threats and particularly to reduce the effectiveness of terrorist organizations. We first develop this framework in conceptual terms and describe the control portfolio available to governments. We then focus on one aspect of that policy portfolio: namely, attempts to influence the stock of terror capacity, whether by reducing the stock directly or by curbing the flow. We sketch a model that seeks to capture major features of the problem facing governments, and draw out and discuss the key results.

2.1. The stock of terror capacity

The success of a terrorist organization in making threats and carrying out attacks depends critically on the resources it accumulates to support its cause: a network of supporters; financial capital; weapons, explosives, and materiel; destructive know-how; a communications network; the tacit approval or even active encouragement of a state or states; trained personnel; and a sufficient number of willing recruits willing to risk prison or death. We

refer to these accumulated resources as the “stock of capacity” available to a terrorist organization. The mix of resources may vary greatly from organization to organization, but some accumulated capacity is essential for terror activity.⁵ Note that this concept of capacity includes not just physical and monetary resources but also intangible assets such as the psychological support of the “man in the street.” The effectiveness of organizations such as the IRA or Hamas depends greatly on the popular support they enjoy; as that popular backing wanes, as the IRA’s appears to have done in the years leading up to the Northern Ireland peace accord, the ability of militant organizations to sustain terror campaigns falls dramatically.

In contemplating terror *defense*—as opposed to retaliation—what matters is some measure of the *expected* damage or social loss from terrorist activity, in advance of actual attacks. Even immediately after an attack takes place, defensive actions are driven by the desire to limit the threat of future damage.⁶ Moreover, the success and even the scale of an attack are often unknown *ex ante*—even by the terrorists themselves.⁷ In this light, a terrorist organization’s stock of capacity can usefully be viewed as an accumulation of the *potential* to plan and carry out terror attacks.

Given such uncertainties, the per-period loss from a particular stock of capital will have a probability distribution. For our purposes, a sufficient statistic for that distribution is its certainty equivalent (CE). We assume, as is conventional, that period utilities are separable, and additive on a discounted basis. Moreover, we assume that the expected scale and damage from such attacks increases with the organization’s resources, at an increasing rate. Thus, if we denote the stock of terror capital as S , and the certainty equivalent loss as $x = L(S)$, we have $L' > 0$ and $L'' > 0$. We take $L(S)$ as the natural “objective function” that governments seek to minimize.⁸

Terror capabilities are not neatly tallied in units such as dollars or trained operatives. Even if some absolute metric were found for terror stock, we believe that our convex functional form assumption for expected losses would apply: i.e., that no rescaling would be necessary.⁹ As a terrorist organization continues to accumulate resources and support, the threat is likely to grow even more quickly than the organization’s resources, at least until losses become monumental. Even if Al Qaeda’s resources and capabilities rose substantially between the 1998 embassy bombings in East Africa and the attack on the World Trade Center and Pentagon three and a half years later, that increase was dwarfed by the enormous increase in the resulting devastation. While the damages from a bioterror attack could conceivably be an order of magnitude greater than the September 11th attack, the required increase in terror *capacity* is probably much smaller, as suggested by the recent evidence that Al Qaeda had made some progress towards obtaining chemical, biological, and conceivably even nuclear weapons.

2.2. *The control portfolio*

A state can defend its citizens against terrorism in four fundamental ways.¹⁰ First, it can reduce the stock of terror capital directly. The most obvious form of stock control is military action directed against the terrorist organization or the states that support them. Prominent examples include invading Afghanistan to disrupt Al Qaeda and overthrow the Taliban,

rolling tanks into West Bank cities to destroy explosives laboratories, and assassinating terrorist leaders. Stock control need not involve military action. For example, construction of a wall or barrier along the West Bank, sharply limiting the capabilities of terrorists to carry out attacks within Israel, could be considered stock control.

Second, a state can curb the net flow of capacity into the terrorist organization. Identifying so-called “charitable organizations” with ties to terrorist groups and applying diplomatic pressure to state sponsors can cut off the flow of funds. Intercepting weapons shipments, such as Israel’s seizure of the *Karine A* in the Red Sea in January 2002, reduces the flow of materiel. Using propaganda to counter prevailing negative images of the United States—for example, broadcasting Arab-language pro-American cable television coverage of events in the Middle East—attempts to slow the accumulation of new recruits to the cause and perhaps even drain popular support for terror (i.e., not just slow the inflow of recruits but also hasten the outflow). Flow control might also consist of measures aimed directly at countering attempts by supporters of terrorist groups to attract new recruits. For example, tearing down the houses of relatives of terrorists might reduce the inflow of new volunteers—the mirror image of the practice (employed by Iraq and other terror supporters) of paying cash awards to families of “martyrs.” Most importantly, flow control encompasses projects aimed at education, economic development, and political openness, on the grounds that such aid will alleviate poverty and promote democracy, thus helping to eliminate support for radical groups such as terrorist organizations.

Both stock and flow control are aimed at reducing the potential or capacity of a terrorist organization to carry out attacks. Government defenses against terror can also more directly confront the likelihood of particular types of attack. To diminish the probability of a successful airplane hijacking, governments may require airplanes to X-ray baggage, impose more stringent identification requirements on passengers and subject them to more thorough searches, and so on. At a broader level, states can seek to reduce the likelihood of successful terror attack by tightening border controls. We call such actions “averting actions:” they aim to thwart attacks and thus reduce the terrorists’ probability of success without affecting the terrorist organization’s underlying capacity for terror. Averting actions may even include appeasement of terrorist groups or their fanatical supporters, the strategy followed by Saudi Arabia.

Finally, government policies can prepare for attacks when they come, to decrease the damages associated with a successful attack.¹¹ We call this category of approaches “amelioration.” Advance preparation for emergency response—clarifying chains of command, identifying evacuation routes, and so on—constitutes amelioration. Other examples include stockpiling Cipro as a safeguard against anthrax attacks and inoculating health care workers against smallpox. An excellent analysis of what we refer to as amelioration is provided by Kaplan, Craft, and Wein (2002) in the context of a smallpox attack. They demonstrate the importance of choosing the right amelioration measure: mass vaccination immediately following a smallpox attack would likely result in far fewer deaths than would a targeted “ring vaccination” program of the kind used to combat natural outbreaks in the past. Such a mass vaccination program would require considerable advance preparation and hence exemplifies amelioration efforts.

<i>Stock reduction</i>	<i>Arrest Al Qaeda operatives overseas Attack Taliban and Al Qaeda in Afghanistan</i>
<i>Flow control</i>	<i>Freeze assets of terror-supporting charities Build schools in Pakistan</i>
<i>Averting actions</i>	<i>Tighten airport security Improve border controls</i>
<i>Amelioration</i>	<i>Prepare mass smallpox vaccination program Stockpile Cipro</i>

Figure 1. Policy portfolio of government defenses against terror: illustrations from the case of combatting the threat from Al Qaeda.

These components of the government's "control portfolio" are illustrated in Figure 1, using the example of defending against attacks by Al Qaeda. Of course, these various components of the control portfolio are closely linked. Ameliorating the effects of a particular kind of attack diminishes the overall terror threat. For example, smallpox vaccination obviously reduces a population's vulnerability to a smallpox attack, and hence lowers somewhat the social loss or disutility associated with a terrorist organization's stock of capacity. Moreover, the potential scope of damage may enter into the terrorist's choice of target or decision of whether or not to attack. Terrorists are strategic opponents. If the target vaccinates for smallpox, terrorists may choose another means, perhaps raising the chances of an anthrax attack. Strong border controls might lead them to attack another nation.

Averting actions and amelioration affect the threat from terror, but not the underlying capacity. We now turn to a more detailed consideration of controls that aim directly at the capabilities of terrorists to plan and carry out attacks.

2.3. Controlling stocks and flows

To analyze the optimal policies towards terror stocks, we draw on the model developed in KVZ to study the control of environmental problems. That paper considers environmental quality as a stock subject to a flow of deterioration, and explores the optimal balance between controlling the flow (e.g., abating carbon dioxide emissions into the atmosphere, or curbing the rate of decline of an endangered species) and controlling the stock (e.g., removing accumulated carbon dioxide from the atmosphere, or restoring the population of the species). Framing terrorism as the product of a stock of terror capacity, we can employ the spirit of that model to explore how terror defense should trade off stock-control efforts versus flow-control efforts.

2.3.1. A model of the control of the terror stock. We start by assuming that over the long run the stock of terror capacity is growing, but that it fluctuates essentially randomly in the short term. For example, the long-term growth of Islamist terror groups is fueled by

underlying factors such as the political and economic conditions in Arab countries and the history of United States involvement in the Middle East. In the short run, however, a terror organization's ability to plan and carry out attacks may be influenced by shifts in public opinion, the capture or escape of a key operative, the ebb and flow of recruitment efforts, and so on.

In formal terms, we assume that absent defensive efforts to control stocks or flows, the stock of terror capacity S follows a Brownian motion with a positive drift rate μ . Here μ captures the rate of "natural" increase of terror capacity—the rate at which the terrorist organization's capital stock will grow (from new recruits, financial contributions, expertise, and so on) in the absence of any government intervention. We emphasize that μ is assumed to be positive: we are considering a situation in which the terrorist organization, if left unchallenged, will accumulate resources and capacity over time.

We shall consider two types of government policies: reducing the stock (alternatively, "restoring" the stock to its initial low level), and curbing the flow. The technologies for these two types of "cleaning up the terror stock" differ in important ways. Reductions of the stock are assumed to exhibit economies of scale, so that the costs of restoration depend little on the starting point. For simplicity, we take the cost of stock reduction or "restoration" to be constant, denoted by C ; but our results are robust to other forms of the cost function with sufficient economies of scale.

For example, launching an attack on Al Qaeda redoubts in Afghanistan involves enormous fixed costs of mobilizing troops and equipment, constructing bases in the field, and so on. The cost of a dozen bombing runs may be half as much as the cost of several hundred. Once an invasion is launched, the incremental costs of scouring another suspected terrorist camp or hideout are small relative to the overall costs of the operation. As an extreme example, the cost of building a wall along national borders to keep terrorists out is the same whether one terrorist or one hundred are prevented from entering.

Fixed costs may also arise from political or institutional factors. Any significant military action, for example, may require assembling a coalition of allies—whether it is an airstrike or a ground invasion that is proposed. Similarly, the diplomatic repercussions of invasion may be relatively insensitive to the extent of the action. The chorus of condemnations of an Israeli counterattack against Palestinian forces has been broadly similar, regardless of whether those actions have taken the form of a targeted assassination of a militant leader, a missile strike on apartment buildings, or full-blown tank movements leading to block-by-block fighting.

Curbing the net flow of new capacity, on the other hand, is assumed to be subject to increasing marginal costs. Discrete measures to curb will be taken sequentially, starting with those that promise the highest net benefits. For example, intercepting large weapons shipments and substantial asset transfers will be relatively inexpensive compared to the costs of ensuring that no weapons or finances reach a terrorist organization. Similarly, any attempt at reducing recruitment must deal with the distribution of potential recruits. Dissuading marginal recruits from joining Al Qaeda is likely to be much easier than discouraging diehards who are already radicalized.¹²

Increasing marginal costs are also likely to apply to attempts at undermining popular support for terrorism. Pro-American radio or cable television broadcasts are likely to be

a relatively cheap means of propaganda, but will have limited effectiveness. Persuading moderate Muslim clerics to speak out against terror would be more effective, but would also require a much more concerted (and expensive) effort. Winning over a large fraction of the population would likely require substantial, long-term investments in schools and infrastructure. In cases such as these, where the measures have variable scale, each measure will be pushed until the returns at the margin are equalized.

In formal terms, let q denote the efforts to abate the flow of new capacity. The costs of such efforts are given by $c(q)$, where $c'(q) > 0$ and $c''(q) > 0$. The net growth in terror capacity over time dt , given government action q , is then $(\mu - q)dt$. Hence the change in the stock S over a time interval dt is $dS = (\mu - q)dt + \sigma dw$, where dw follows a standard Brownian motion.

2.3.2. The optimal mix of stock and flow controls. This basic model corresponds closely to the model of stocks and flows examined in KVZ, and hence we can immediately apply many of the results in that paper to the current context. We shall outline the model and provide intuition, but refer readers to KVZ for elaborations and proofs.

First, the optimal policy will involve both stock and flow controls. Reductions in the stock follow a “threshold” policy: once the stock reaches or exceeds a certain “trigger level,” it is optimal to reduce it to a much lower level, denoted \underline{S} .¹³ This result is intuitive: the economies of scale in reducing the stock make it optimal to do so only periodically. It brings to mind the well-known (s, S) model of inventory management, in which the optimal policy in the face of stochastic demand for a product is to replenish inventory up to a level given by S every time it falls to (or below) s (Arrow, Harris, and Marschak, 1951; Scarf, 1960).¹⁴

The optimal level of flow control, meanwhile, varies with the stock. As the stock increases at first, the optimal abatement rate increases as well. At some point, however, optimal flow control slows down, and less and less abatement is done as the stock continues to grow. The intuition is that optimal abatement at any given point in time equates the marginal benefits (given the current stock level) with the marginal cost. While the marginal cost of abatement depends only on how much abatement is done, the marginal benefits vary with the stock. When the stock is small, marginal benefits of abatement rise as the stock increases, because the social loss function increases at a faster rate as the stock increases. As the stock nears the trigger level, however, restoration becomes imminent. Hence additions to the stock are likely to be very short-lived, and the gains from curbing the flow are reduced.

The form of the optimal flow-control policy is illustrated by Figure 2, which shows the optimal abatement rate in the case of a quadratic cost function and exponential utility. Although the specific parameters used for that figure are arbitrary,¹⁵ they offer some insight into how the optimal policy might play out in the real world. In the simulation depicted in Figure 2, the destination for stock reduction \underline{S} was assumed to equal 250. The drift rate was set at 1.5. Thus if we interpret the length of the time period dt as one week, this figure corresponds to a model in which stock reduction lowers terror capacity to the level it would reach in a little over three years (in expectation) if left unabated. The optimal policy, given these parameters, is to curb the growth of the stock somewhat, but never fully offset it in expectation: the optimal rate of flow control always lies below the drift rate (represented by the dotted line in the figure). Stock reduction occurs optimally when the terror

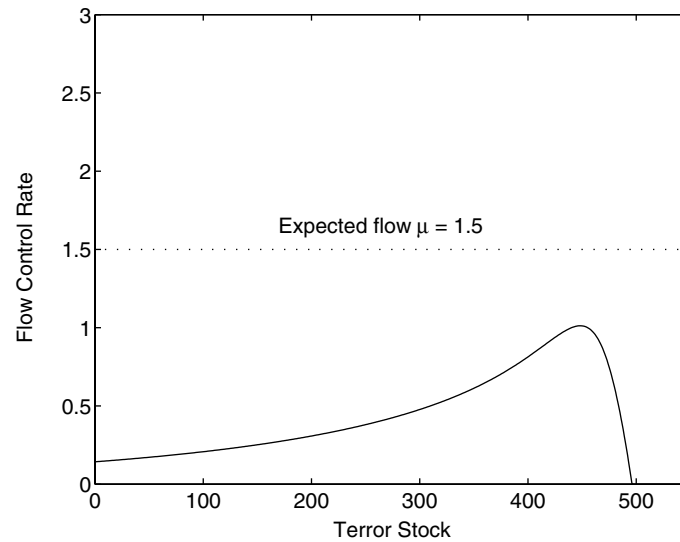


Figure 2. Optimal abatement policy as a function of the terror stock.

stock rises to approximately 500—that is, roughly twice the level to which it returns with restoration.

The exact abatement path and the optimal trigger for stock reduction depend in natural ways on the key parameters of the model. As the rate of “natural” stock growth μ increases, abatement becomes less important relative to stock reduction. This is illustrated by Figure 3, which shows the fraction of the flow abated as a function of the terror stock, for three flow rates. (The other parameters of the model remain the same.) Intuitively, the cost of periodically reducing the stock decreases relative to the cost of managing the flow.

Recall that optimal flow control initially increases with the stock. For low drift rates, optimal abatement eventually offsets the expected flow, so that the optimal policy maintains an equilibrium level of the stock in expectation. In Figure 3, this “expectation equilibrium” is located where the rising abatement path meets the horizontal line representing the drift rate. To the left of this point, abatement is less than expected flow, so the stock continues to grow in expectation. To the right of the point, in the range where the abatement path lies above the horizontal line, the expected net growth in the terror stock is *negative*, pushing the stock back towards the equilibrium. In such a situation, restoration only becomes part of the optimal policy when there is a large stochastic increase in the stock level.

For higher drift rates, on the other hand, fully offsetting the average drift is too costly. Instead, the optimal policy involves periodic reductions of the stock along with some control of the flow. As the drift rate increases further, the optimal policy involves abatement, and more and more frequent (in expectation, at least) restoration.

An important implication of these results is that the optimal policy does not in general involve a “target” level of the stock. Rather, if the average flow of new capacity is too high to be offset by government abatement policies, then the stock will tend to increase over time

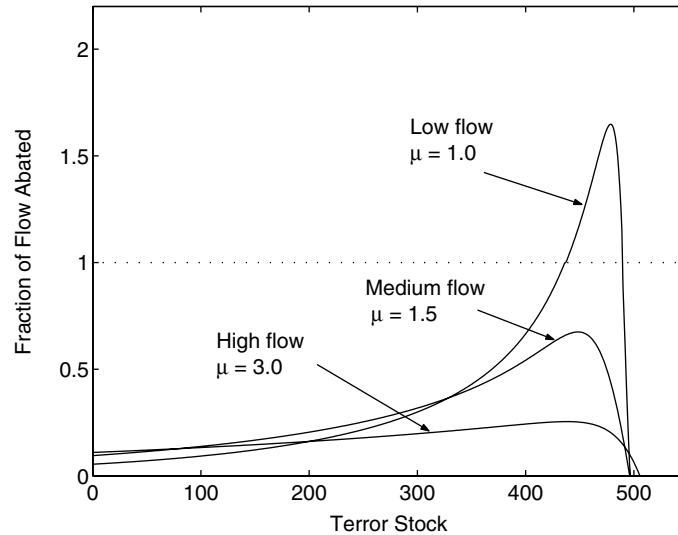


Figure 3. Fraction of the flow abated as a function of the terror stock, for three flow rates.

under the optimal policy, with a corresponding increase in the certainty equivalent losses associated with the terror stock. When the stock reaches the trigger level, a major effort is undertaken to curtail it. This jerky policy capitalizes on the economies of scale in cleaning up the stock of terror.

Moreover, as Figure 3 demonstrates, a relatively modest increase in the rate at which the terror stock grows can have substantial effects on the optimal policy. In terms of the simulations depicted here, a three-fold increase (from $\mu = 1.0$ to $\mu = 3.0$) shifts the optimal policy from one centered around an “expectation equilibrium” to one characterized almost solely by periodic restoration.

Figure 4 illustrates the steady-state distributions of the stock level under the same three flow rates. When flows are relatively high, the stock will increase steadily between periodic cleanups. Hence the distribution of states approaches a uniform distribution, with all states between the cleaned-up level and the trigger equally frequent in expectation. Only if flows are relatively low will the stock tend to stabilize around some level; and even then, a sudden increase in the stock may lead to restoration.

As the cost of restoration C increases, the trigger stock level increases as well. This result is intuitive, but it has an especially important implication in a model of terror stocks. As noted above, political costs (in the form of diplomatic sanctions or the costs of assembling a coalition) may be an important component of the cost of stock reductions. These political costs are likely to be highly sensitive to the incidence of terror attacks. Al Qaeda’s stock of terror capacity was presumably not much different in April 2001 than in October later that year; but the costs of launching an invasion against Afghanistan were much lower in the aftermath of September 11th.

If such feedback effects exist—i.e., terror attacks reduce the costs of controlling the stock—then the use of terror will undermine itself. An interesting question is where the

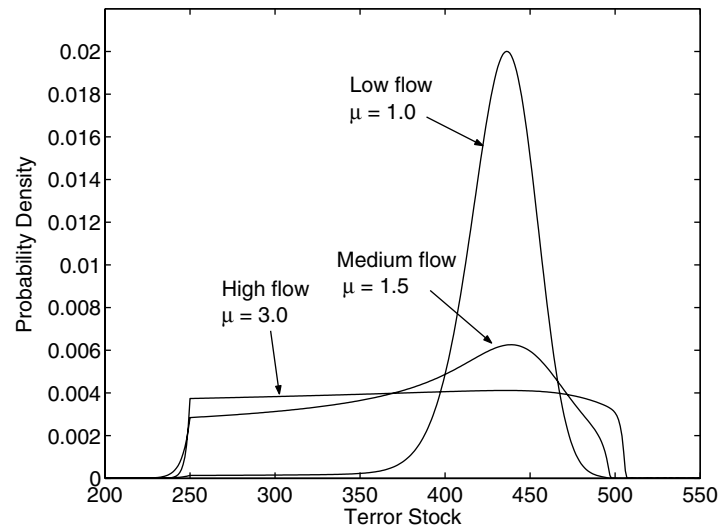


Figure 4. Frequency distributions of terror stocks (states) under optimal policies for three flow rates. (Note that the horizontal scale is truncated at 200.)

threshold lies for such feedback effects to kick in. For example, apparently neither the 1998 bombings of American embassies in Nairobi and Dar es Salaam, nor the subsequent attack on the U.S.S. Cole were enough to bring the costs of a response down sufficiently to support a sustained attempt at reducing the stock.

2.4. Key implications of stock-and-flow model

Our model of government responses to terror focuses on the underlying stock of capacity that gives rise to terror incidents, rather than on those incidents themselves. If economies of scale prevail in reducing the terror stock, as we argue is likely, then a government's best response to the threat posed by terrorists is not to seek to maintain some "manageable" level of terror activity. Nor is the best policy one that aims to destroy terror organizations with one blow. Rather, the optimal policy involves ongoing attempts to staunch the growth in the terror stock, punctuated by periodic reductions in the stock itself. The recurrent and even jerky nature of such a policy represents a fundamental break from the conventional approach to waging war, despite the frequent use of the war metaphor to describe the fight against terror.

3. Individual responses to terror

We have considered the responses to terror that could be undertaken by a centralized entity, presumably a state or alliance of states.¹⁶ We now shift our perspective to consider how individual agents, such as citizens or firms, respond to threats of terror. Clearly these agents

will not undertake stock or flow controls. Those measures are highly expensive, and are public goods that are inevitably undertaken by governments. However, individual entities can take measures to remove themselves from harm (what we term “avoidance”), or to reduce the harm they suffer (“amelioration”).¹⁷

We begin by addressing the security externalities associated with amelioration and avoidance. A simple model demonstrates that the positive externality shares a surprising but intuitive affinity with familiar congestion externalities—with the same bleak implications for the effectiveness of government policies. We conclude with several extensions that relax several assumptions of the basic model without altering this central insight.

Throughout this section, we use the term “target” to delineate the set of individuals exposed to a single threat. The physical connotation of this term is well-suited to many cases, as when the target is downtown Manhattan at risk from a dirty bomb detonation, or the residents of a large prominent building. However, we shall also apply the term more generally: a “target” could be a mode of transport (plane or train), or a population unprotected against a certain type of attack (individuals lacking vaccination against smallpox). It is important to stress that our usage of target is generic rather than specific.¹⁸

3.1. Security externalities from individual precautions

Individual entities—hereafter often called individuals even if possibly firms or families—may seek to lessen the damage to themselves in the event of a terrorist attack. For example, a family may create its own store of Cipro to protect itself against anthrax. The owner of a prominent skyscraper may reinforce the structure to reduce the harm from a car bomb.¹⁹ Individuals may also lower their chances of being harmed by avoiding likely terrorist targets. For example, travelers may choose to drive rather than fly, in response to the threat of airplane bombings or hijackings. Urban dwellers may leave the city for the perceived safety of suburban areas. Corporations may disperse their operations among several buildings, rather than concentrate them in a prominent skyscraper; or they may move to the suburbs. Israelis may steer clear of popular nightclubs and rush-hour buses.

3.1.1. Externalities from avoidance and amelioration. Like other protective measures taken by individual agents, amelioration and avoidance impose security externalities on others. These externalities may in principle be beneficial or harmful. If corporations or residents flee a city to avoid the threat of a terror attack, their departure makes an attack on the city less attractive to the terrorist and hence less likely. Their actions will make individuals in other cities worse off, however, if the terrorist’s efforts are simply shifted elsewhere.

We distinguish between concentrated and collective threats and between local and global effects.

Concentrated versus collective threats. A crucial distinction concerns the nature of the threat: in particular, whether it is *concentrated* or *collective*. By a “concentrated” threat we mean one that will ultimately affect only a particular individual or individuals at a time in a given target location. For example, most conventional crimes—car theft, mugging, break-ins, and so on—constitute concentrated threats. Note that a threat can be concentrated

even if the individual attacked is chosen randomly: the point is not how but whether an individual is singled out. For such threats, the likelihood of harm to any particular individual in the target population decreases as the number of individuals exposed to the threat increases.

To illustrate, consider a set of skyscrapers exposed to the possibility of terrorist attack, where the tallest skyscrapers (or perhaps the most distinctive) are the most likely targets. The owners of skyscrapers can protect the buildings by hardening them against attack. For example, Citigroup Center, a familiar part of the New York City skyline and home to one of the world's most prominent financial institutions, is undergoing structural reinforcements in an apparent effort to help it withstand a terrorist attack.²⁰ Because these reinforcements are publicly observable, they are likely to reduce the probability of an attack on Citigroup Center, as well as lessening the damage if an attack does occur. They constitute avoidance as well as amelioration.

The total social benefit from such protection, however, may be substantially less than the gains to the owners and occupants of the Citigroup building, if shielding that building from an attack simply makes an attack on other prominent skyscrapers (the Chrysler Building, say, or the Sears Tower in Chicago) more likely. More generally, when the effects of terror attacks are concentrated on particular individuals, avoidance carries a negative externality for the rest of the target population: in avoiding an attack, the agent diverts it onto others.²¹

In contrast, consider a *collective* threat, affecting all exposed individuals in a city or an area at once. In this case, we suppose that as the number of people exposed increases, the attractiveness of the target to the terrorist (and hence the likelihood of attack) increases as well. This clustering of people makes major sports events natural targets for terrorist threats, quite apart from their salience. Hence, the likelihood of an attack increases as the number of people exposed grows. As a result, individual precautions can discourage a collective attack, producing a positive externality to those at a particular target.

Anthrax attacks exemplify a collective terror threat. Individuals within a target population can protect themselves against the threat by securing private stores of Cipro. Such amelioration confers a positive externality on other members of the same target population by reducing the risk of attack. Similarly, in the example already given of city dwellers fleeing to the suburbs, avoidance of the collective threat to the city benefits those who remain.

Local versus global effects. The examples above considered *local* externalities; i.e., effects on "nearby" or similar entities. Precautions can have wider external effects, however. Hardening of skyscrapers may lead terrorists to shift their efforts towards other targets altogether. Individuals who stockpile Cipro in one community may simply redirect an anthrax attack to another, less well-protected one. Alternatively, they may lower the likelihood of an anthrax attack, but raise that of a smallpox attack. Individuals who forsake planes for trains may (in the aggregate) make trains more likely targets of attack.

Such negative externalities are *global* in scope: they affect individuals in other target populations. Note also that these global effects are harmful regardless of whether the threat is a concentrated or collective one.

3.1.2. A unifying framework. From the examples discussed thus far, we can draw out three distinct security externalities from avoidance and amelioration. Individual precautions

against terror can reduce the probability of an attack on a target, generating a positive local externality we refer to as a *discouragement effect*. Such precautions can also shift the likelihood of attack onto others—an effect we call *diversion* if the shift is local (i.e., onto other individuals at the same target), and *displacement* if the shift is global.²²

A fourth type of security externality, referred to here as a *containment effect*, is examined in a pair of papers by Howard Kunreuther and Geoffrey Heal (Heal and Kunreuther, 2002; Kunreuther and Heal, 2002). They contemplate the externality generated by the interdependence of one agent's security on the protective measures taken by others. In their motivating example, an airline will be less likely to use a baggage screening device if its rivals also have not installed such screening: the possibility that bags from another unprotected carrier will be transferred onto the airline's own plane reduces the gains from screening. Hence a decision by one individual to protect himself yields external benefits for others. Kunreuther and Heal (2002) show in a simple model that for some parameter values the problem has the structure of a prisoners' dilemma: no agent protects himself in Nash equilibrium, even though full protection would be optimal. Moreover, if agents vary in the probabilities of harm and in the costs of security measures, a "tipping phenomenon" may exist, where the decisions of a fraction of the population can determine whether all agents invest in security or none do (Heal and Kunreuther, 2002).

We can elucidate the connections among these various security externalities with a bit of formal notation. Index possible targets by t and the total population of individuals (across all targets) by i . Let h_{it} represent the harm to individual i given an attack on target t . Let the probability that individual i is harmed by an attack on t be given by $\pi_t \cdot p_{i|t}$, where π_t is the probability that the terrorist attacks target t and $p_{i|t}$ denotes the conditional probability that individual i is attacked, given an attack on t . The expected harm to individual i from being located at t is evidently $\pi_t \cdot p_{i|t} \cdot h_{it}$. The total expected harm from an attack on target t , meanwhile, is $\sum_i p_{i|t} h_{it}$.

Using this framework, concentrated threats are those for which $p_{i|t}$ varies among agents at location t , as a result of both individual characteristics and precautionary measures. For collective threats, on the other hand, $p_{i|t} = 1$ for all individuals at t . In both cases, amelioration reduces h_{it} , while avoidance lowers $p_{i|t}$. Note that avoidance of a collective threat entails removal from the targeted location, and hence lowers $p_{i|t}$ to zero; avoidance of concentrated threats can also result from observable precautions (as in the case of hardening buildings), and so need not eliminate risk or exposure completely.

Self-protection generates externalities to the extent that π and p depend on the expected damages from attacks at various locations. Suppose that terrorists seek to maximize the harm they inflict (an assumption we discuss in detail below) and face a binding resource constraint. Then a decrease in the total expected harm from an attack on target t will lower π_t and raise π_τ , where τ signifies any other target location. The former effect represents the local positive externality from discouragement; the latter effect represents the global negative externality from displacement.²³

For a concentrated threat, avoidance can generate two other externalities. If a reduction in $p_{i|t}$ raises $p_{j|t}$, avoidance entails a local negative externality from diversion. This effect will hold when the terrorist singles out one individual entity for attack, as in the skyscraper example above.

<i>Nature of threat</i>	<i>Concentrated</i>	<i>Collective</i>
<i>Scope of effect</i>		
<i>Local (among individuals at the same target)</i>	<i>discouragement (+)</i> <i>diversion (-)</i> <i>containment (+)</i>	<i>discouragement (+)</i>
<i>Global (among targets)</i>	<i>displacement (-)</i>	<i>displacement (-)</i>

Figure 5. Types of security externalities, classified by the nature of the threat and the scope of the effect. The direction of each externality is in parentheses.

On the other hand, the possibility of contamination between agents (as in Kunreuther and Heal's example of a bomb in luggage that may be transferred from one airplane to another) implies that a reduction in $p_{i|t}$ lowers $p_{j|t}$, so that avoidance yields a positive local externality from containment. Note that this effect, unlike the others considered so far, is nonstrategic: it results not from a change in the terrorist's behavior but from the structure of relationships among potential victims.

This discussion is summarized in Figure 5, which classifies these externalities based on the scope of the effect (global or local) and the nature of the threat (collective or concentrated).

3.2. *Collective threats: Discrete precautions and congestion effects*²⁴

Terror threats are often collective threats: they are targeted at, and result in damages to, populations (or groups of people) rather than individuals.²⁵ Because of the externalities associated with self-protection, individual responses to terror in the absence of government policy will be inefficient.

If individuals could vary their precautions continuously, then in Nash equilibrium each individual would spend just enough on avoidance or amelioration to equate the marginal cost of doing so with the marginal private benefit. Government policies that reduced the marginal private benefit of precaution (say by reducing per-capita harm from terror attacks in a way that lowered the marginal benefit at every level of precaution) would thereby reduce individual expenditures on protection (assuming diminishing marginal returns to such expenditure). Hence such policies would provide social benefits, to be weighed against the costs of government action.

Many measures to protect against terror, however, are inherently discrete. A corporation can choose to locate in a city or not;²⁶ a family can store antibiotics (or gas masks) or not; an individual may travel by airplane or drive. We consider such discrete measures in this section. A striking conclusion from our model is that government policies that reduce the probability of attack or the harm from an attack may have very little impact on social welfare (as measured by expected utility of a representative individual). Indeed, in the simplest case where individuals are identical, government policies that fall short of complete protection, so that some avoidance or individual amelioration is still undertaken, will improve welfare not

at all. We show that there is a strong analogy between the problem of individual responses to terror and the familiar congestion externality.

3.2.1. A simple model of binary location choice. We begin with a simple model in which individuals can choose between two locations. One location is the target for a possible terrorist attack; the other is a “safe haven” which is certain not to be attacked. In concrete terms, this choice might represent living in the city versus living in the country; taking the airplane versus driving a car; locating corporate headquarters in a downtown skyscraper versus a suburban office park; or going to a nightclub in Jerusalem versus staying home. To simplify exposition, we shall refer to the discrete precaution as “avoidance.” However, the results go through unchanged for a discrete amelioration measure, with the natural adjustments to the notation.²⁷

A number of identical agents independently choose the risky or safe location. The former take the “expose” strategy, the latter the “avoid” strategy. Because only the fraction of individuals choosing an action matters, not the absolute number, we normalize the size of the population to unity. One of the two locations, denoted location 1, is a target of a possible terrorist attack. The per capita harm from a successful attack equals $h_1 > 0$ for individuals at location 1. The alternative location—denoted location 0—is a safe haven: the probability of attack is zero, so that the expected harm from terror is zero to individuals at that location. However, there is a cost of avoidance, denoted by a_0 . This represents chiefly the opportunity cost associated with choosing a location (or mode of travel, etc.) that would otherwise be inferior. We will restrict our attention to the interesting case in which $a_0 \in (0, h_1)$. If $a_0 = 0$, the safe haven would be just as good as the target in the absence of terror, and hence strictly preferred under threat of attack; while if $a_0 \geq h_1$, avoidance would be so costly (or the expected loss from the attack so minimal) that people would choose exposure even if attack were certain.

Let the probability of a successful attack on the target be given by $\pi(n_1, h_1)$, where n_1 is the fraction of people who locate at 1 and expose themselves to the threat, and h_1 is the per-person harm. (Recall that we are considering only collective threats, so that $p_{i|1}$ equals 1 for individuals at the target and 0 for those who avoid the attack.) For the time being, we shall assume that h_1 is constant across individuals, and therefore shall suppress it in the probability-of-attack function π .

Total expected harm from the terrorist attack is $\pi(n_1)h_1n_1$; the total expected welfare loss is given by $L(n_1) = \pi(n_1)h_1n_1 + a_0(1 - n_1)$. We assume that the greater is the potential harm, the more attractive is the target to the terrorist. Thus, $\pi(n_1)$ increases with n_1 .

Two assumptions deserve emphasis. The first is that individuals accurately gauge the probability of attack $\pi(n_1)$. This rationality assumption is standard in a game-theoretic context. Psychological research, however, has suggested that individuals tend to overestimate low-probability events (Kahneman and Tversky, 1979), and fail to distinguish sufficiently among small but positive probabilities. Taking such a bias into account would alter the implications of our model for social welfare, but not the mechanics of the model itself. That is, *given individuals' perceptions of risk*, our model implies that they will tend to choose too little avoidance as long as they ignore the impact their choice has on the probability of attack.²⁸ If those perceptions consistently overestimate actual risk, the avoidance

externality we identify may be alleviated. In this light, one implication of our model is that irrational responses to risk may improve social welfare. Where avoidance externalities pertain, ignorance may be preferable to accuracy in public perception of risk.

The second critical assumption is that the likelihood of attack rises as more people are exposed to the threat. This will clearly hold if inflicting harm is a primary objective of the terrorist. Many analysts (e.g., Falkenrath, 2001) argue on the contrary that terrorists do not generally seek to maximize the direct harm they inflict, on the grounds that harming too many people would undermine their cause. For example, a terrorist group may bomb a building in the middle of the night, to make a statement without imposing too many casualties; or it may alert authorities before a bomb explodes. Other possible objectives of terrorists include promoting fear and disruption in the target population, e.g., widespread refusal to fly because airplanes seem unsafe; and impressing their supporters, hence the pursuit of spectacular and photogenic attacks. Given the range of possible motives for terrorism, our model will not apply to all terror activity. Of course, an assumption that terrorists maximize expected damages does not necessarily imply that they seek to kill as many people as possible: damage to the functioning of political and economic institutions or even to physical structures may also be the goal.²⁹

More generally, while our model requires that the probability of attack rise with exposure, that correlation need not be causal. Even if terrorists choose targets for symbolic value, symbolism and damage may be closely linked. The size of the World Trade Center towers was inextricably linked with both the symbolic power of their destruction and the attendant human toll. Finally, our model may apply within a class of targets, even if the choice among classes is dictated by other criteria. Thus, although a terrorist might target sporting events for their symbolic value, the most densely packed arenas would be the likeliest targets.³⁰

The model we explore here has clear relevance to prominent present-day examples of terrorism. Palestinian suicide bombers, for example, target crowded buses, popular night spots, and densely packed markets in Israel. It appears that they seek to maximize the human toll of their efforts. Although largely unsuccessful, the 1995 poison gas attack by Aum Shinrikyo targeted crowded rush-hour subways in Tokyo. Most prominently, Al Qaeda operatives flew planes into the largest office towers in one of the most densely populated areas in the U.S. The threats made by Osama bin Laden and his associates emphasize the magnitude of the damage they will inflict, and their organization has made efforts to acquire nuclear and chemical weapons. The view that terrorists do not seek to inflict maximum harm may fit the conventional politically motivated terrorist organizations of previous decades. It seems much less appropriate for groups like Al Qaeda, whose expressed aims are not the achievement of political representation or self-governance, but rather the killing of large numbers of civilians.³¹

3.2.2. Results: Inefficient equilibrium and ineffective policies. At a Nash equilibrium, every individual must be indifferent between exposure and avoidance. Hence equilibrium exposure n_1^e is the level of exposure at which expected harm from the terrorist attack equals the cost of avoidance:

$$\pi(n_1^e)h_1 \equiv a_0.$$

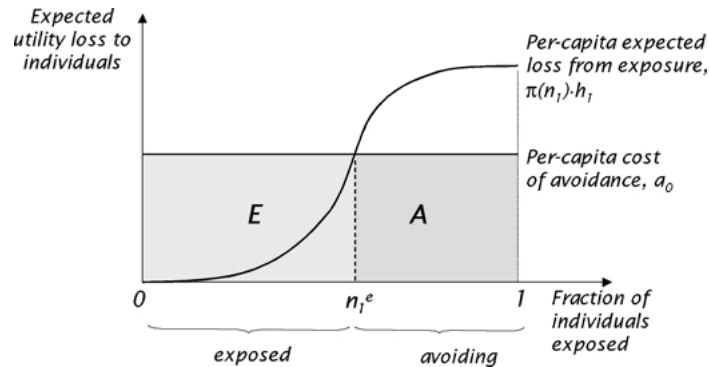


Figure 6. Exposure and avoidance equilibrium. The equilibrium level of exposure is n_1^e . Area E is the total expected loss to exposed individuals; area A is the corresponding loss to avoiders.

The per-capita welfare loss in equilibrium is $L^e = a_0$.³² Note that even though only a fraction $1 - n_1^e$ of the population locates at the safe haven, every individual incurs the same expected loss a_0 . Those at the safe haven suffer this loss as the opportunity cost of avoiding the threat; those at the target location suffer this loss as the expected harm from exposing themselves to the threat.

Figure 6 illustrates this result, employing a form familiar from Schelling's diagrams of binary choice games in *Micromotives and Macrobehavior* (1978). It displays the average (per capita) expected loss from exposure and from avoidance. The expected loss from exposure depends on the fraction of the population exposed, and is given by $\pi(n_1)h_1$. This is an increasing function of n_1 , given our assumption that the probability of attack increases with the numbers hit. The expected loss from avoidance, meanwhile, is constant and equal to a_0 . Equilibrium is given by the intersection of these functions at n_1^e , the point at which the individual is indifferent between the two locations. Area E , to the left of the equilibrium point, represents the expected welfare loss to exposed individuals; area A represents the expected loss to avoiders. Taken together, the shaded areas represent the total expected welfare loss from the threat of terror in equilibrium.

Several implications follow. First, the equilibrium is socially inefficient. By reducing the likelihood of an attack, one individual's avoidance of the target benefits those who remain. Because individuals do not capture this benefit, they choose too little avoidance in equilibrium. At the margin, no one is willing to incur the cost of avoidance for the sake of reducing the expected harm to others.³³

A reduction in the fraction of people exposed, n_1 , raises social welfare relative to the equilibrium outcome.³⁴ Following an increase in avoidance, the people who were formerly exposed are just as badly off as before (suffering a_0), while those who remain at the target enjoy greater safety than before. Similarly, any increase in exposure makes everyone worse off. Figure 7 illustrates the optimal level of exposure. The optimum is at n_1^* , where the sum of the two shaded rectangles (i.e., total expected losses) is minimized.

Second, avoidance complicates centralized attempts (by governments, say) to counter the terror threat. For example, consider averting actions by government (e.g., greater

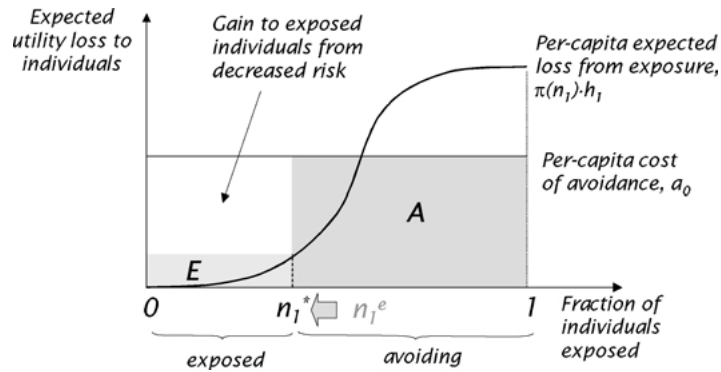


Figure 7. Optimal levels of individuals exposed and avoiding. Optimal exposure n_1^* is less than the equilibrium level, due to the positive avoidance externality. Total social loss ($E + A$) is much smaller than in equilibrium.

surveillance leading to terrorist arrests) that make a successful attack less likely. This shifts the probability-of-attack function π downward. As long as the expected harm from an attack remains greater than the avoidance cost when everyone is exposed to the threat, the total welfare loss will remain unchanged and equal to a_0 . Of course, more people will locate at the target; but in equilibrium, exposure will increase to make citizens indifferent between exposure and avoidance, so that everyone will be exactly as badly off, in expectation, as they were before the government’s action. Government averting actions, even though effective in reducing the likelihood of an attack, will only reduce citizens’ losses if they are so effective that no one chooses the avoidance strategy in equilibrium.³⁵

Other means the government might employ to reduce the expected harm from terror would produce similar results. For example, the government might seek to ameliorate the damages of an anthrax attack by stockpiling Cipro. In our model, amelioration by the government reduces h_1 and hence lowers the average expected harm curve $\pi(n_1)h_1$. This will have no effect on welfare unless harm h_1 is lowered far enough that exposure is less costly than avoidance even when everyone is exposed. Hence like averting actions, government amelioration must be sufficiently effective and employed on a sufficiently large scale, to reduce citizens’ total losses.

This simple model, therefore, produces a striking result over a wide range of government policies (where policies are interpreted as shifts in $\pi(\cdot)$ or h_1). Government measures that are successful in averting attacks or ameliorating their effects—i.e., making an attack less likely or less harmful—may nevertheless not improve expected welfare. This result is somewhat analogous to the well-known “tragedy of the commons.” Open access to an ungoverned fishery results in overfishing and the dissipation of surplus, since fishermen ignore the detrimental effect their fishing has on other fishermen. Here the scarce resource is the target location: individuals continue to expose themselves to the threat until any gains from being at the target rather than the safe haven are completely dissipated.

As the results above demonstrate, however, the problem of “overexposure” in response to a terror threat may be harder for governments to correct than the traditional open-access problem. In this respect the avoidance externality has the structural attributes of a congestion

externality. Commuters will continue to clog congested highways until their progress is slow enough that they would be just as well off on back roads or on mass transit. Adding more lanes may increase the number of commuters on the road, but often does not reduce the resulting congestion, or offer any net benefits. Similarly, in the model considered here, individuals will respond to government policies by increasing their exposure to the terror threat until the expected harm from exposure rises again to the cost of avoidance.

Our results also suggest measures governments could take to address avoidance externalities. For example, governments might levy a tax on individuals who choose to expose themselves to a threat, just as road tolls can be used to reduce traffic congestion. More attractive might be a subsidy to those who avoid risks that might attract an attack. For example, one could imagine a government program to subsidize telecommuters, in the hopes of reducing the attractiveness of downtown streets or rush-hour subways as terror targets.³⁶

3.2.3. Extensions. This section presents four extensions to the model of avoidance just developed. The first extension combines continuous amelioration with discrete measures. The second allows individuals to differ in their costs of avoidance. The third contemplates multiple targets, while the fourth considers multiple modes of attack. All four extensions are admittedly incomplete. We sketch the ways in which the results of our model will be affected, and defer more detailed work to future research. The fundamental insights from our model continue to hold for all four cases.

Continuous amelioration when discrete measures are possible. If both continuous amelioration and discrete measures are possible, the qualitative results remain much the same as when only discrete measures are available. The congestion effects persist. Thus partial amelioration measures may have very little effect on social welfare when avoiding the threat is also feasible. To see this, suppose that individuals are identical. Each individual chooses between avoidance and exposure, and exposed individuals can reduce their potential harm by exerting amelioration effort m . Measuring amelioration effort in the same units as avoidance cost, let the harm an individual suffers as a function of amelioration be given by a decreasing and convex function $h_1(m)$, with $h_1' < 0$, $h_1'' > 0$, $h_1(0) = h_1$, and $h(m) \geq 0 \forall m$. An exposed individual exerting amelioration effort m then has expected disutility of $g(m) = m + \pi_1(n_1, h_1(m)) \cdot h_1(m)$.

In equilibrium, amelioration m^e will be chosen to minimize expected disutility $g(m^e)$, so that $g'(m^e) = 0$. As in the case without amelioration, however, the number of people exposed will adjust to equate $g(m^e)$ with the avoidance cost a_0 . Hence in any interior solution with $n < 1$, per capita expected utility loss equals a_0 . Marginal shifts in the probability-of-attack function π_1 or in the unmitigated harm h_1 —e.g., due to government averting actions or amelioration—will again matter not at all for social welfare.³⁷ As before, government policies will be effective only when a corner solution is reached, which requires lowering π_1 or h_1 sufficiently that $g(m) < a_0$ when $n_1 = 1$.

Heterogeneous agents. One of the simplifying assumptions made in the models above is that individuals are identical. As a result, the avoidance cost does not vary as the fraction of the population exposed to the threat changes. Relaxing this assumption does not affect the inefficiency of equilibrium, but it does improve the outlook for government policies.




Figure 8. Exposure and avoidance equilibrium with heterogeneous agents. As before, areas E and A denote total expected losses to exposed and avoiding individuals, respectively.

The case of heterogeneous individuals (varying only in avoidance cost, i.e., in their relative preferences for the safe haven versus the target) is illustrated by Figure 8. Rather than being constant, the avoidance cost is now a function $a_0(n_1)$. Note that as a function of n_1 , the fraction of the population exposed to the threat, the avoidance cost is downward-sloping. (We assume that the allocation of individuals to locations is efficient, so that the agents with the lowest avoidance costs are the first ones to avoid the threat.) Hence at any given level of exposure n_1 , avoidance cost is greater for individuals who are still exposed (i.e., to the left of n_1 on a graph) and less for individuals who are avoiding the threat.

The equilibrium remains inefficient, since individuals continue to ignore the externality. Total expected losses at equilibrium are now given by $L(n_1^e) = n_1^e \pi(n_1^e) h_1 + \int_{n_1^e}^1 a_0(n) dn$. Further reduction in exposure will lower expected losses if the derivative of this expression is positive at the equilibrium outcome, i.e., if $L'(n_1^e) = n_1^e \pi'(n_1^e) h_1 + \pi(n_1^e) h_1 - a_0(n_1^e) > 0$. This condition does indeed hold. Because expected losses from exposure and avoidance are equal in equilibrium, $\pi(n_1^e) h_1 = a_0(n_1^e)$, implying that $L'(n_1^e) = n_1^e \pi'(n_1^e) h_1$. Just as in the homogeneous case, therefore, equilibrium exposure is inefficiently high.³⁸

The effectiveness of government policies, however, now depends on the slope of the avoidance cost function. Suppose a defensive policy shifts the probability of successful attack, increasing n_1^e . The resulting change in the equilibrium value of the loss function is given by $L'(n_1^e) = n_1^e a_0'(n_1^e) < 0$. It follows that the flatter is the avoidance cost function, the smaller is the improvement in expected welfare. Clearly, the slope of the avoidance cost function reflects the degree of heterogeneity in preferences. Hence the greater the heterogeneity in the neighborhood of the equilibrium, the more effective are government policies. Conversely, in the case of homogeneous agents considered above, $a_0' = 0$ by assumption and averting actions or ameliorations by government are ineffectual on the margin.

Multiple locations. Next, consider an extension to multiple target locations subject to a single threat. For example, a terrorist may threaten to attack one of several cities, or one of several modes of transport. As before, each individual must choose a single location. (For

simplicity, we revert to the assumption of identical individuals.) Denote target locations $1, \dots, T$ by the subscript t , and let location 0 be a “safe haven” as before. Associate with each target t a harm h_t and a probability-of-successful-attack function $\pi_t(n_t)$, where n_t is the fraction of the population at location t . Let location T be the location that would be preferred in the absence of terror, so that all other targets entail a positive avoidance cost. The disutility or expected loss at location t is therefore $\pi_t(n_t) + a_t$. Finally, to reflect the possibility of limited space, let each target have some maximum capacity n_t^{\max} (except for the safe haven, which is assumed to be able to absorb the whole population).

As long as individuals choose locations to maximize their utility, then in equilibrium the expected losses will equalize across all “active” locations.³⁹ Hence the results of the single-target model go through unchanged. Expected welfare loss in equilibrium is a_0 , the avoidance cost of the safe haven. Avoidance is less than socially optimal: too many people choose to locate themselves at the targets. Most importantly, small-scale government policies are again ineffective at reducing overall welfare: they result only in relocation of individuals from the safe haven to the targets, while leaving total expected welfare loss unaffected. In particular, averting and amelioration policies are ineffective as long as $a_0 - a_t \leq \pi_t(n_t^{\max})h_t$ for all t —that is, as long as at each target t , the expected harm from an attack if as many individuals as possible were to locate at t would outweigh the savings in avoidance costs. In other words, in the context of this model, averting and amelioration efforts that fail to make at least one target strictly preferable to the safe haven will not improve expected social welfare.

Hence this model suggests that when avoidance externalities are significant, government policy should focus on making one location at a time safer, rather than trying simultaneously to ameliorate conditions across a range of locations. Marginal improvements in several locations at once may merely reallocate individuals from the safe haven to those targets, without improving overall social welfare. To be effective, government policies must make at least one location strictly preferable to the “safe haven.” Indeed, the optimal government policy in this model involves a strict ranking of target locations, with each location ranked according to the gain in expected social welfare from making it preferable to the safe haven (net of the cost of such a policy).

What about situations in which no “safe haven” exists, and none can be created at affordable cost? In such cases, marginal reductions in expected harm at one location can lead to welfare gains—but these gains will be largely dissipated by individuals relocating in response to the policy. For example, a reduction in the probability of a successful attack at a particular target location will make that location safer and hence strictly more attractive to individuals than other sites. In response, individuals will migrate to that target, driving up the expected harm there until a new equilibrium is reached. Because in equilibrium those individuals will be drawn from every other site, the net gain in expected welfare from the policy may be slight. Hence a similar result holds as in the case of a safe haven: individual responses to government averting or amelioration policies will tend to undermine the effectiveness of those policies. The degree to which government protection will be attenuated at a given target depends on how readily people relocate among targets and on how sensitive the likelihoods of attacks on other locations are to the numbers of people at those locations.

Moreover, note that the existence of a “safe haven” can be treated as endogenous. Consider a model in which no safe haven exists. Now let government policies proceed until enough individuals have left one location that the probability of an attack there falls to zero. If the fraction of the population at this location is still positive, the location then serves the same role as a “safe haven.” Further amelioration or averting actions by government to protect other targets will prove ineffective until no more individuals are left in the safe location.

With multiple targets, the negative externality from avoidance becomes relevant. Reducing expected harm at one location increases the probability that other locations will be attacked: this is the global displacement effect. However, the significance of this externality is diminished in a manner very similar to the congestion effects already noted. When individuals can move among target sites, in equilibrium the expected threats will be equalized among locations. Hence the negative effects of making one location safer will be diminished, just as the positive effects are.

We showed above that a reduction in the likelihood of attack at a particular location (e.g., by government averting actions) may be ineffective in improving social welfare. In exactly the same way, the welfare effects of an *increase* in the probability of one type of attack due to displacement will be diminished when discrete precautions are feasible.

Multiple modes of attack. Finally, consider the case (of great practical relevance) of multiple simultaneous modes of attack. In contrast to the case just discussed of multiple locations subject to a similar threat, we have in mind here a terrorist choosing among several different *types* of attack, with no common amelioration or avoidance measure. For example, consider two very different biological terror attacks: anthrax and smallpox. Ameliorating the damages from anthrax (say by stockpiling Cipro) will be of little use against a smallpox attack.

The global displacement effect appears at first to be of major importance to this case. Indeed, if individuals protect themselves against one type of attack, that will raise the likelihood of another type of attack. However, the welfare effects of displacement will be diminished when discrete precautions are feasible against each type of attack. To the extent that each individual threat can be avoided at some cost (or ameliorated with discrete measures), *the analysis of the simple model above will hold for each threat independently.* Indeed, in the extreme case in which individuals are identical and some fraction of agents remains exposed to each threat, the global displacement effect will vanish. The negative externality from the increased likelihood of other types of attack will simply be absorbed by increases in avoidance of those threats, with no net reduction in expected welfare.

While this result is symmetric with the congestion effect from the simple model, it has much brighter implications for government policy. In particular, it suggests that policies that successfully encourage avoidance of prominent threats (e.g., through a telecommuter subsidy) may have lasting benefits that are not eroded by negative externalities elsewhere.

4. Conclusion

Governments can protect their citizens against terror. Individuals can protect themselves. Each set of actions can be illuminated by an ecological metaphor. An organization’s ability to wage terror rests on an accumulation of resources, a stock that includes physical

assets such as weapons and explosives, but also includes intangibles such as networks of supporters and the approval of at least some segments of the population. A government's efforts to combat terror share important features with efforts to control any stock of public bads, such as pollution. In particular, the optimal control of terror stocks will rely on both ongoing abatement and periodic cleanup, a lesson derived from a new model of optimal pollution control. The government's optimal policy portfolio also includes averting actions (reducing the probability of successful attacks) and amelioration (reducing the harm from an attack).

Curtailling a terrorist threat is a public good to the targeted population. As such, individual citizens will not provide it. However, individuals can avoid exposure to terror where threats are high, or reduce the harm they will suffer from an attack. Such precautionary measures, which we call "avoidance" and "amelioration," may exert harmful or beneficial externalities. The effect on individuals subject to different threats will be a negative one: by making one target safer, individuals displace risk onto other targets. The more relevant effects, however, operate at the local level. Whether one person's precautions harm or benefit others nearby hinges on how the likelihood of harm to an individual varies with the number of individuals similarly exposed. When terrorists single out an entity for attack, one agent's precautions divert the risk to others. Thus, when premier skyscrapers are strengthened, other skyscrapers become priority targets. By contrast, when the threat is a collective one directed at a total population, individual avoidance or amelioration helps similar others by making the group a less attractive target.

The latter externality is reminiscent of congestion externalities. No individual can significantly affect another, but the decisions of many can generate significant welfare effects. Moreover, individual responses to terror tend to undermine the effectiveness of government policies. In the extreme case where agents have identical costs, government policies will yield no benefits, unless they are so effective that all citizens choose exposure. This insight continues to hold, though perhaps attenuated, when we extend the model to heterogeneous individuals or multiple threats.

The two aspects of the ecological metaphor complement each other. For example, one implication of avoidance behavior is that when individuals determine their exposure among potential targets of a threat, the government should seek to eliminate the threat at the targets where it is most serious, rather than dispersing effort across a range of targets. As in many realms where the government protects, effective policies to defend against terrorism must account for the reactions of the individuals they are meant to protect, not merely the responses of the perpetrators.

Appendix: Computations of Section 2.3

Computations that generated the plots appearing in Section 2.3 were conducted using specific functional forms and fixed parameter values. The exception was the drift rate μ , which varied as indicated in the figures. The functional forms were a quadratic function for the abatement cost and a natural exponential function for utility. The fixed parameter values and exact functional forms are provided in Table 1. In accordance with the formal model in Keohane, Van Roy, and Zeckhauser (2000), abatement was assumed to be capped at some

Table 1. Parameter values.

Variance	σ^2	9
Discount rate	α	0.005
Utility	$u(x)$	$-e^{0.05(s-400)}$
Abatement cost	$c(q)$	$40q^2$
Abatement ceiling	\bar{a}	20
Restoration cost	C	13000
Restoration destination	\underline{s}	250

level (the “abatement ceiling”), although as explained in that paper the assumption of such a ceiling is done purely to simplify the analysis.

Value functions were computed via policy iteration on a “locally consistent” approximating Markov chain (see, e.g., Kushner and Dupuis, 1992). The program used to generate the plots was based on one written by Benjamin Van Roy.

Acknowledgments

We thank Miriam Avins, Geoffrey Heal, Ed Kaplan, Robert Keohane, Thomas Schelling, and participants at the Faculty Seminar at the Kennedy School of Government and at the American Economic Association annual meetings for helpful comments. Some portions of our analysis, as noted below, draw heavily on models developed in conjunction with Benjamin Van Roy (environmental quality) and Luke Stein (terror avoidance).

Notes

1. The most lasting work on this subject is Thomas Schelling’s *The Strategy of Conflict* (1963), in part because of the breadth of its vision.
2. The definition of the term “terrorism” has itself been the subject of much dispute. We define terrorism as violence conducted by organized nonstate actors and directed against incidental civilian targets. By “incidental,” we mean that victims of terror activity are anonymous from the terrorist’s point of view. Of course, the physical objects of attack may be intrinsic to the terrorist’s intent—e.g., the World Trade Center towers. Nonetheless, the particular identities of the human victims are incidental to the terrorist. This distinguishes terrorism from, say, gang warfare.
For an alternative definition, see Schmid and Jongman (1988, p. 28), who put greater emphasis on motive, distinguishing intimidation, coercion, and propaganda as the primary goals of terrorists.
3. Deterrence may remain effective against state *sponsors* of terror organizations, such as the Taliban regime in Afghanistan. The practical significance of such deterrence, however, is limited by the difficulty of marshalling enough evidence of a state’s involvement in terrorism to justify attacking it.
4. Although it is tangential to the thrust of our argument, one can note that while “Nature” is nonstrategic, the firms that emit pollution are not—although they may well be uncoordinated in their responses to regulation.
5. The contrast between the perpetrators of the two major successful acts of terrorism on United States soil illustrates this point. The Oklahoma City bombing was the product of a simple organization with limited resources, while the September 11th attacks were the work of a deep-rooted and well-funded network of terrorists with worldwide reach. At the same time, these disparities in resources are reflected not only in the magnitude of the attacks but more importantly in their aftermath. Al Qaeda’s vastly greater resources make

it a continuing threat, while the homegrown terrorism of Timothy McVeigh appears to represent minimal continuing expected danger.

6. Of course, motives may be mixed. Part of the impetus behind the U.S. invasion of Afghanistan may have been raw retaliation against aggression and its supporters. But the systematic attempts to bomb empty caves and raze deserted training camps (not to mention the ongoing efforts to uncover and disrupt the Qaeda network) are defensive actions to make the next attack less likely or less damaging in expectation.
7. Indeed, videotape footage of Osama bin Laden indicates that even Al Qaeda were unsure of the scale of their attacks on the World Trade Center—i.e., whether the towers themselves would actually fall. If the September 11th attacks were the “worst case scenario,” the 1993 bombing of the trade center fizzled by comparison; although there were casualties, the effects were presumably less than intended and, one imagines, less than they might have been under slightly different circumstances.
8. A terror stock also induces anxiety in the target population about potential attacks. At least for modest levels of the terror stock, anxiety may be significant relative to the CE loss. Moreover, the “anxiety cost” may well exhibit *decreasing* returns to scale as a function of the stock of terror capacity. A 5% probability of a “dirty bomb” being detonated in Manhattan may cause almost as much worry as a 25% probability. Stein (2002) contemplates the implications of concavity in the cost function for anxiety.
9. Our results require a less stringent condition. Even if the expected loss function loses convexity in some range when some absolute measure of the stock is used, a rescaling may be available that restores the required convex shape. The potential difficulty lies in simultaneously preserving the assumed concavity of the cost function for reducing the terror stock (i.e., economies of scale). As long as the underlying loss function is not too far from convex, and the underlying cost of terror reduction function is sufficiently concave, such a rescaling will be possible. To illustrate, it is straightforward to show that a logarithmic rescaling of the stock will restore the conditions we impose as long as $-\frac{L''}{L'} < \frac{1}{S} < -\frac{C''}{C'}$ over the relevant range of S , where C' and C'' are the first and second partial derivatives of the cost of stock reduction with respect to the initial stock. (Note that in the text we assume for simplicity that C is constant, but we allow for the case of a concave cost function.) Here $-\frac{L''}{L'}$, for example, is a conventional measure of the curvature of the loss function.
10. Our discussion here draws on many motivating examples from recent campaigns against terror—in particular, those conducted by the United States and its allies against Al Qaeda, and by Israel against Palestinian terrorists. It is important to stress that these are meant as *examples* to illustrate our framework, rather than as particular models for action. We are interested here in the optimal tradeoff between controlling stocks and flows of terror capacity. We make no claim that the particular examples cited represent “optimal” actions.
11. Amelioration efforts by governments raise similar issues as the individual precautions discussed in Section 3. Public efforts to reduce the harm associated with a particular type of attack may diminish the likelihood of such an attack taking place. But such efforts may also make other types of attack more probable—an effect we term “displacement” below.
12. In this respect, slowing recruitment is like pulling out weeds: the first 70% are much easier to get than the last 30%.
13. Note that \underline{S} might be well above zero: as the stock becomes smaller and smaller, at some point further reductions may become prohibitively expensive (that is, although economies of scale in restoration prevail over most of the range, marginal costs may start to increase as the stock approaches zero).
14. There are also instructive parallels between our model and models of optimal savings when investments exhibit economies of scale; see Dixit, Mirrlees, and Stern (1975).
15. This and other figures in this section are based on simulation programs written by Benjamin Van Roy. Full parameter values and other details are described in the appendix.
16. We set aside the question of whether action will be taken multilaterally or by individual states. While that issue is a crucial one, we now wish to consider an even more fundamental distinction—between measures taken by a centralized authority that affect the terror threat for an entire population, and measures taken by individuals within that population.
17. Like averting actions by government, avoidance by individuals reduces the likelihood of an attack. For clarity, and because the concepts are slightly distinct even beyond their scale, we use different terms for the individual and government actions.
18. Thus, we will consider threats against skyscrapers as a *class*, rather than against any individual building. The virtues of the generic approach will be apparent when we classify the externalities. From the perspective of

the owner of a skyscraper, a building is one of several individual entities that are all subject to the same type of threat—just as a car parked on a city street may be one of several in the same neighborhood that is at risk of being stolen.

19. Atomic bomb shelters provide another example of amelioration, albeit from an era of very different threats.
20. See the *New York Times* article by Glanz and Lipton (2002). Means of hardening buildings against terrorist attack have been explored systematically by the National Research Council (1995).
21. Note that the negative avoidance externality here hinges on the building reinforcements being *known* to the terrorist. Otherwise, such measures could not divert an attack from a particular building, although they might discourage an attack on skyscrapers as a whole. Indeed, the owner of the Citigroup Building was presumably not upset to have his project displayed prominently in newspaper reports, since an aura of invulnerability will enhance the building's protection. On the other hand, it might be in the public interest to keep secret the identities of buildings that have been hardened, in order to reduce the diversion effect relative to discouragement. Hence avoidance, and the diversion effect that may result, require that the terrorist or criminal discern and act upon the identity of the individual(s) taking a precaution. Unidentifiable measures can only be amelioration.

The importance of precautions against attack being "observable" has been noted in the literature on crime prevention (Shavell, 1991; Ayres and Levitt, 1998). For example, a car alarm must be visible to divert thieves. The term "identifiable" seems more apt than "observable," however. The use of such hidden measures as Lojack homing devices must be "observable" in the aggregate for their use to have any preventive effect at all. Similarly, in our model of precautions against terror, individual avoidance (by leaving a city, for example, or switching modes of transportation) may often be observable in principle, but the identities of the individuals are irrelevant to the terrorist: only the aggregate level of avoidance matters.

22. Externalities similar to those we discuss arise from crime prevention. Individual precautions against crime may discourage crime in general, divert crime onto neighboring victims, or displace criminals to other areas. Indeed, our terms "diversion" and "displacement" are taken from the literature on the economics of crime and its prevention. What we call "discouragement" is referred to in that literature as "deterrence"; we prefer the former term because deterrence connotes a threat of punishment. On the deterrence and diversion effects, see the articles by Shavell (1991), and Hui-wen and Png (1994). On displacement, see Ayres and Levitt (1998) and the articles cited therein. Clotfelter (1978) is generally credited with initiating this line of literature.
23. Amelioration exerts both discouragement and displacement effects, since a reduction in h_{it} unequivocally reduces $\sum_i p_{it}h_{it}$. In general, however, avoidance has these effects only for collective threats, where one individual's reduction in p_{it} from one to zero leaves the values of p_{jt} unaffected for all other individuals j .
24. This section has benefited greatly from the work of Luke C.D. Stein in his Harvard College A.B. Thesis titled "The Ecology of Terror Avoidance: Individual Responses and Group Consequences" (2002), for which both authors were advisors.
25. Whether a threat is collective or concentrated may depend on whose perspective is taken. The attack on the World Trade Center was a collective threat from the point of view of individual workers and the firms that employed them; similarly, the continuing danger of airborne attacks on major cities constitutes a collective threat. From the perspective of individual building owners, on the other hand, an airplane flying into a skyscraper appears as a concentrated threat, as already noted. Because our analysis focuses on responses by individuals, there is no contradiction in characterizing the same event as a collective threat to some of its potential victims and a concentrated threat to others.
26. Although one might argue that a corporation could choose a location along a continuous spectrum of urbanization, in practical terms the terrorist threat facing large cities such as New York is qualitatively different than that facing smaller cities or suburbs; hence the relevant choice is the discrete one of whether or not to locate in a city.
27. In particular, the model below could be applied to amelioration by first defining population 1 as those individuals lacking the protection offered by the amelioration measure, and population 0 as protected individuals. Then let a_0 denote the cost of amelioration, and normalize the harm suffered by protected individuals to be zero (equivalently, let h_1 denote the difference in harm between the two locations). The model then goes through exactly as in the case of avoidance.
28. In formal terms, one could introduce the bias identified by Kahneman and Tversky into our model by distinguishing the subjective perception of the probability of an attack, say π^s , from the actual probability π .

Equilibrium would then go through as before with π^s substituted for π . The implications of the avoidance externality for social welfare would evidently depend on the divergence between π and π^s .

29. The assumption that terrorists maximize expected damages may apply over some but not all of the range of possible harm. For example, a terrorist may be unwilling to cause damages beyond some ceiling, but prefers to inflict as much harm as possible up to that level. In this case, our model would apply to the set of threats that fall below the terrorist's self-imposed ceiling.

Indeed, this may fit the case of Al Qaeda. Top Qaeda leaders said in interviews with Al Jazeera that they considered but rejected attacking nuclear facilities for fear that they would "get out of hand." It is possible that they feared that such an attack would violate an unspoken "nuclear taboo." On the other hand, they have not ruled out targeting nuclear plants in the future. Moreover, Al Qaeda clearly sought to inflict enormous harm on the United States on September 11th.
30. One might still claim that the value of an attack on the Super Bowl lies in the value of the "spectacle" witnessed by billions of television viewers, rather than in the tens of thousands that would be killed in the stadium. But the value of the spectacle would be much reduced if the stands were empty (as would be the case if the authorities, anticipating a possible attack, closed the event to live spectators).
31. In his call to arms in February 1998, entitled "Jihad Against Jews and Crusaders," Osama bin Laden declared that "The ruling to kill the Americans and their allies—civilians and military—is an individual duty for every Muslim who can do it in any country in which it is possible to do it... We—with Allah's help—call on every Muslim who believes in Allah and wishes to be rewarded to comply with Allah's order to kill the Americans and plunder their money wherever and whenever they find it." English translation posted by the Federation of American Scientists at <http://www.fas.org/irp/world/para/docs/980223-fatwa.htm> (accessed September, 2002).
32. By our normalization of the population to unity, this is also the *total* welfare loss.
33. In terms of our prior typology, this simple model explores only the local discouragement effect of avoidance. With only one threat, the displacement effect cannot arise. We return to the global externality when we discuss the cases of multiple locations and modes of attack.
34. The derivative of the expected loss function is $L'(n_1) = \pi(n_1)h_1 + n_1\pi'(n_1)h_1 - a_0$, which equals $n_1^*\pi'(n_1^*)h_1$ at the equilibrium level of exposure. This is positive, by the assumption that π is increasing in the number of people exposed. Increases in exposure, meanwhile, are clearly welfare-reducing, since $\pi(n_1)h_1 > a_0$ for all $n_1 > n_1^*$.
35. That is, government actions must shift the probability-of-attack function $\pi(n_1)$ far enough that the expected loss from exposure is less than the cost of avoidance when everyone is exposed: i.e., $\pi(1) < a_0$.
36. Of course, there are myriad other considerations involved. Strict attention to the implications of avoidance externalities might suggest subsidizing the movement of large firms out of cities; but such a policy would exacerbate other externalities associated with actual congestion and sprawl, and would ignore the considerable positive effects of conglomeration.
37. Note that in this case individual amelioration, like avoidance, creates a positive externality by reducing the likelihood of attack. But these external benefits (along with the individual benefits from amelioration) are dissipated in equilibrium, as individuals who were formerly avoiding the threat now respond to the lower likelihood of attack by exposing themselves to it. Hence although amelioration is inefficiently low for any *fixed* level of exposure, a marginal increase in amelioration (on the intensive margin) will be offset by an increase in exposure (on the extensive margin).
38. Of course, this does not imply that the optimum is the same in the heterogeneous-agent case as it is when agents are identical. On the contrary, for a given probability-of-attack function $\pi(n_1)$, it is straightforward to show that the optimal level of exposure will be higher with heterogeneous agents (and hence closer to the equilibrium level), since avoidance is more expensive on the margin.
39. There may be some locations that are so easily attacked that no one will locate there.

References

- Arrow, Kenneth J., Theodore Harris, and Jacob Marschak. (1951). "Optimal Inventory Policy," *Econometrica* 19, 250–272.

- Ayres, Ian and Steven D. Levitt. (1998). "Measuring Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of Lojack," *Quarterly Journal of Economics* 113, 43–78.
- Clotfelter, Charles T. (1978). "Private Security and the Public Safety," *Journal of Urban Economics* 5, 388–402.
- Dixit, Avinash, James Mirrlees, and Nicholas Stern. (1978). "Optimum Saving with Economies of Scale," *Review of Economic Studies* 42, 303–325.
- Falkenrath, Richard. (2001). "Analytic Models and Policy Prescription: Understanding Recent Innovation in U.S. Counterterrorism," *Studies in Conflict & Terrorism* 24, 159–181.
- Glanz, James and Eric Lipton. (2002). "A Midtown Skyscraper Quietly Adds Armor," *The New York Times*, August 15, page A1.
- Heal, Geoffrey and Howard Kunreuther. (2002). "You Only Die Once: Managing Discrete Interdependent Risks," Mimeo, Columbia Graduate School of Business, April.
- Hui-wen, Koo and I.P.L. Png. (1994). "Private Security: Deterrent or Diversion?" *International Review of Law and Economics* 14, 87–101.
- Kahneman, Daniel and Amos Tversky. (1979). "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47, 263–291.
- Kaplan, Edward H., David L. Craft, and Lawrence M. Wein. (2002). "Emergency Response to a Smallpox Attack: The Case for Mass Vaccination," *Proceedings of the National Academy of Sciences* 99, 10935–10940.
- Keohane, Nathaniel O., Benjamin Van Roy, and Richard J. Zeckhauser. (2000). "Controlling Stocks and Flows to Promote Quality: The Environment, with Applications to Physical and Human Capital," NBER Working Paper No. 7727 (June).
- Kunreuther, Howard and Geoffrey Heal. (2002). "Interdependent Security: The Case of Identical Agents." NBER Working Paper No. 8871 (April).
- Kushner, Harold J. and Paul G. Dupuis. (1992). *Numerical Methods for Stochastic Control Problems in Continuous Time*. New York: Springer-Verlag.
- National Research Council. (1995). "Protecting Buildings from Bomb Damage: Transfer of Blast-Effect Mitigation Technologies from Military to Civilian Applications," Report of the Committee on Feasibility of Applying Blast-Mitigating Technologies and Design Methodologies from Military Facilities to Civilian Buildings. Washington, DC: National Academy Press.
- Scarf, Herbert. (1960). "The Optimality of (S,s) Policies for the Dynamic Inventory Problem." In *Proceedings of the First Stanford Symposium on Mathematical Methods in Social Sciences*. Stanford, CA: Stanford University Press.
- Schelling, Thomas C. (1963). *The Strategy of Conflict*. Oxford: Oxford University Press.
- Schelling, Thomas C. (1978). *Micromotives and Macrobehavior*. New York: W.W. Norton & Co.
- Schmid, Alex P. and Albert J. Jongman. (1988). *Political Terrorism: A New Guide to Actors, Authors, Concepts, Data Bases, Theories and Literature*. Amsterdam: North-Holland.
- Shavell, Steven. (1991). "Individual Precautions to Prevent Theft: Private versus Socially Optimal Behavior," *International Review of Law and Economics* 11, 123–132.
- Stein, Luke C.D. (2002). "The Ecology of Terror Avoidance: Individual Responses and Group Consequences." A.B. Thesis, Harvard College (April).