

An Analysis of Short-Term Responses to Threats of Terrorism

Edieal J. Pinker

Simon School of Business, University of Rochester, CS3-345 Carol Simon Hall,
Rochester, New York 14627, pinker@simon.rochester.edu

Two important defensive mechanisms available to governments combating terrorism are warnings and the deployment of physical resources. Warnings are relatively inexpensive to issue but their effectiveness suffers from false alarms. Physical deployments of trained security personnel can directly thwart attacks but are expensive and need to be targeted to specific locations. In this paper, we model the joint optimization of defenses against terrorist attacks based on warnings and physical deployments when there is uncertainty in the timing and location of attacks. We model both private warnings issued to security forces and public warnings broadcast to the general public. By structuring the trade-offs faced by decision makers in a formal way, we try to shed light on an important public policy problem. We show that the interaction between the use of warnings and physical defenses is complex and significant. For public warnings, we also model the possible response of terrorists and show how these responses influence the effectiveness of such warnings.

Key words: decision analysis; risk; military; defense systems; government; defense

History: Accepted by Linda V. Green, public sector applications; received August 8, 2005. This paper was with the author 2 months and 3 weeks for 1 revision.

1. Introduction

Keohane and Zeckhauser (2003) define four mechanisms by which a government may combat terrorism: reducing the stock of terrorists, limiting the flow of resources to terrorists, taking averting actions, and taking ameliorating steps. In this paper, we do not consider stocks and flows of terrorism but focus on two defensive measures—deployment of physical resources, an example of an averting action, and issuance of warnings, an example of amelioration. The purpose of this paper is to help create a better understanding about the purpose and effectiveness of terror warnings and resource deployments. We believe that this is the first paper to model joint decisions about warnings and resource deployments to minimize the costs of terrorism. By developing a theoretical model of the interplay between resource deployment, terror warnings, and the information available to decision makers we create a structured way to think about these defensive mechanisms and identify when and how they can be used most effectively to provide for public security.

Physically deploying security forces in the field can have several benefits. If the deployment is very visible it can serve as a sign that the government is “doing something” to protect the public and is a visible sign that perhaps the government is anticipating an attack. Similarly, these deployments may deter terrorists from striking. Clearly, the physical presence of

security forces can also prevent or limit the extent of a terrorist attack. Physical resources are also expensive, and their cost effectiveness has much to do with the intelligence available on likely threats. Placing a battalion of the New York State National Guard around the Statue of Liberty will not be very effective if terrorists are planning an attack on Los Angeles International Airport. Likewise, if the U.S. government believes that there is an imminent attack planned on a fast food restaurant in the continental United States, deploying a squad of armed troops at every McDonalds and Burger King in the country will be a very expensive exercise. Physical deployments may also reveal defensive procedures if terrorists are probing defenses with feints.

Since March 2002, the United States has adopted a system of color-coded terror warning levels. At any point in time, the Department of Homeland Security (DHS) states the current assessment of the likelihood of a terrorist attack on the United States. These assessments are based on information and analyses from the various intelligence agencies of the U.S. government. In Israel, the security services regularly report on the number of active terror threats they are tracking (see, for example, Dudkevitch 2004). In Israel it is also not uncommon to hear news reports that refer to recent security operations as related to some previous warnings (Myre 2004). Warnings can serve many of the same purposes as physical deployments.

Public warnings can create the perception that government authorities are active and informed in a struggle that is often conducted in a secretive way. Public warnings can also deter terror attacks. If a terrorist cell, in the midst of planning an attack, comes to believe (as a result of the warning) their plans have been partially revealed, they may abort them. Warnings can also serve as a mechanism to directly aid in the prevention of terrorist attacks and the reduction in the damage caused by terrorist attacks. The first goal is achieved by increasing the alertness of security personnel at possible targets. The second goal is clearly expressed by the DHS in its description of its alert system. In its literature, the DHS describes the steps that all government agencies should take to prepare their facilities and employees at different levels of threat, with similar guidelines for the general public. Raising preparedness means that if important infrastructure is damaged, government and civil society can continue to function thereby minimizing the damage of the attack. It can also mean canceling or altering activities so that fewer people are present at likely targets, again minimizing the damage from an attack.

Compared to physical deployments, warnings have an important cost advantage. Given the modern communications technologies available, disseminating a warning across a broad geographic area is relatively inexpensive. However, the effectiveness of warnings is limited by their credibility. Pate-Cornell (1986) and Pate-Cornell and Benito-Claudio (1984) model the optimization of warning systems. In particular, they provide a structure for understanding how warning systems work and can be evaluated. Pate-Cornell (1986, p. 223) characterize the costs and benefits of a warning system as being a function of “(1) the quality of the signals issued and (2) the response of the people to whom warnings are directed.” These two factors determine the trade-off between failures to issue alarms when something bad happens and the issuance of false alarms. If the public does not believe the warning it will not take the precautions dictated by the raised alert. This holds true for security personnel as well whose work is typically quite monotonous. Therefore, while warnings may be relatively cost effective they cannot be used indiscriminately or their credibility and utility will be diminished.

In this paper, we analyze two types of warnings—private and public. Private warnings are defined as warnings issued to security forces to boost their alertness and preparedness for possible attacks. Public warnings are defined as warnings issued to the general public to similarly boost their alertness and preparedness for possible attacks. Being public, these warnings also serve as warnings to security forces and are observed by terrorists, and as such, may affect

terrorist behavior. The economic analysis of terrorism, and responses to it, has found that substitution is an important phenomenon in terrorist behavior (Enders and Sandler 1993). Defensive actions that make one type of target more difficult to attack lead terrorists to different targets. Terrorists are also forced to make intertemporal substitutions shifting attacks from one time period to another in response to attacks on them by security forces. Enders and Sandler (2004) give an overview of the development of this literature. Therefore, in this paper, when we model public warnings we include the possibility that terrorists will shift their attacks to periods in which there is no public warning.

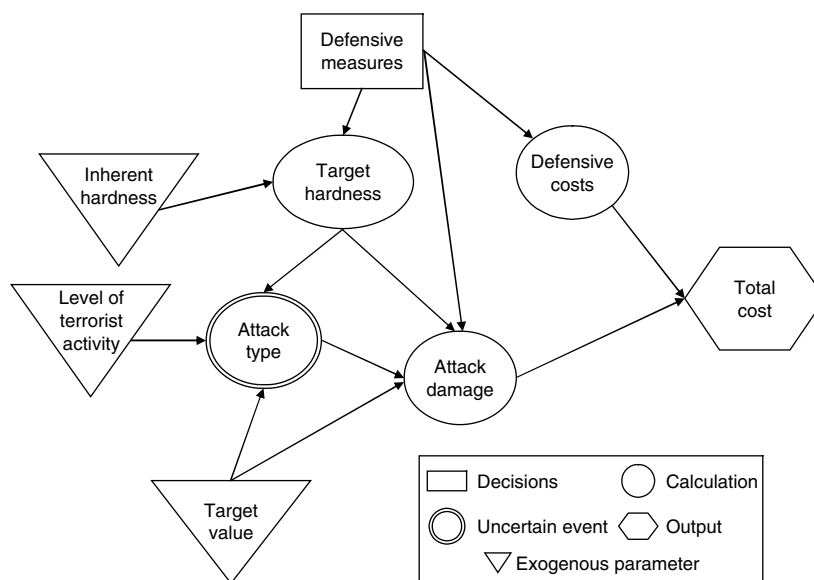
In §2, we develop a conceptual model of terrorism risks and defenses using influence diagrams to create a broad structure for the main issues facing a decision maker. In §§3 and 4, we develop and analyze mathematical models that capture the key elements of the conceptual model described in §2. While the mathematical models are further abstractions of the situation faced by society and the decision maker, their formalism enables us to investigate the issues more deeply and precisely than the conceptual model. Our goal is to improve our understanding of the interaction between warnings and physical deployments and to identify prescriptions for decision makers. Given the abstraction of the models, we should not expect decision makers to adopt the prescriptions without modification. However, we should expect the prescriptions to shape and spur discussion and further investigation of these important issues. In §3, we formulate a model of private warnings and physical defenses. We analyze this model to develop insight into how governments can best utilize warnings and physical defensive resources. We also compare the performance of our model to an alternative heuristic demonstrating the importance of integrating warning and physical deployment decisions and effective intelligence management. In §4, we extend our model to include public warnings and develop rules of thumb for the optimal use of public warnings. In particular, we show how expectations about terrorist responses to public warnings can influence the optimal use of public warnings. We conclude in §5.

2. Conceptual Model

To better structure our thinking about terrorism defenses, we now develop conceptual models of long-term and short-term risk analysis. These conceptual models will form the basis for our mathematical modeling and analysis in the following sections. We begin by considering a single potential target and then expand our discussion to a spectrum of multiple targets.

When analyzing the long-term risk of terrorism to a particular target, several factors are important: the

Figure 1a Influence Diagram of Long-Term Risk of Damage from Terrorism to a Single Target



probability of an attack, the difficulty of doing substantial damage to the target (its hardness), and the economic and political value of the target. In Figure 1a, we use an influence diagram to depict how these factors work together to determine the cost of terrorism from the perspective of a single target.

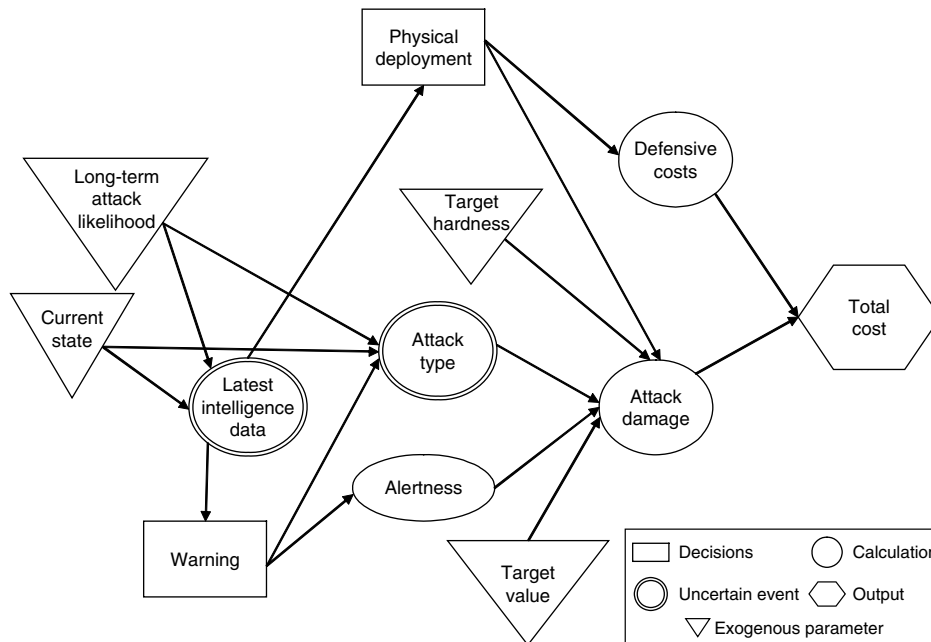
The total cost of terrorism is comprised of damage from attacks and the cost of defensive measures taken. Defensive measures are in addition to the inherent hardness of the target. For example, a shopping mall, designed for open access to the general public, is not as difficult to damage as a nuclear power plant that has restricted access and many safeguards to make it robust. A shopping mall can be hardened by installing concrete barriers that prevent vehicle bombs from getting too close or by installing metal detectors in its entrance ways. Such measures combined with the inherent hardness of the target will influence the damage an attack causes. The damage caused by an attack is also influenced by the value of the target. Damage caused to a high-value target is by definition more costly to society than damage caused to a low-value target. The value of a target is determined by its economic significance as well as its political or symbolic significance. Loss of life has both economic and political significance. Finally, attack damage is influenced by the type of attack, where “no attack” is a possible type of attack. The type of attack is of course uncertain but is itself influenced by the hardness of the target and its value because we assume that the terrorists themselves perform a cost-benefit analysis when selecting a target. Terrorists want to cause significant damage but want a high probability of success.

The analysis outlined above, when applied to all potential targets in a geographic area, will yield what we can term a *risk profile* for that space. Several insurance companies are currently developing such risk profiles to determine insurance coverage criterion and parameters (see Risk Management Solutions 2005). The purpose of the risk profile is to assess the expected damage costs of terrorism to each target over the lifetime of the insurance policy.

A short-term perspective assumes that the outcome of the long-term risk analysis outlined above has been conducted and operationalized and that its outcome is known by the decision maker and to a large extent the terrorist groups as well. In Figure 1b, we use an influence diagram to depict how intelligence information influences decisions about short-term defensive measures.

Given the long-term risk profile, there is a current state of the world that influences the short-term likelihood of attack and type of attack. Intelligence data gives the decision maker imperfect knowledge of this state that influences his decision making. For example, a long-term risk profile, as developed by insurance companies, may indicate that the set of high-rise office buildings in Manhattan are high-risk targets for a truck bomb. However, security surveillance tapes may indicate suspicious activity by small groups of men filming specific buildings in the financial district. Such information merged with the long-term risk profile would yield an updated probability of an attack on high-rise office buildings in the financial district. Looking at Figure 1b, this means that the probability of an attack on a specific target is influenced by the long-term risk for that target and

Figure 1b Influence Diagram of Short-Term Risk of Damage from Terrorism to a Single Target



the current intelligence data. Target value and target hardness have already influenced the probability of attack through the long-term risk analysis. In the short-term perspective of Figure 1b, these factors only influence the damage from an attack. In the short-term setting, the decision maker must decide whether and how to deploy additional physical defensive resources and whether or not to issue warnings. The physical deployment has the effect of further hardening the target and potentially reducing damage, while warnings influence alertness which also potentially reduces damage from an attack.

The structure provided by the influence diagrams allow us to illustrate several important observations about defending against terrorism. First, there is a distinction between long-term and short-term decision making that closely resembles planning for natural disasters. For example, there is a historical record of hurricanes that informs our expectations about where and when they will hit and the damage they will cause. Such expectations influence insurance assessments, disaster planning, and building codes. However, when real-time intelligence is available that a hurricane has, for example, formed in the Caribbean, the government updates its assessment of where it may strike, issues warnings to the public, deploys emergency services, and prepares shelters for storm refugees all in an effort to mitigate the potential damage from the hurricane. Second, the influence diagrams indicate a key difference between terrorism and natural disasters, namely, terrorists are agents that decide when and where to attack. Long-term measures to harden the defenses of some targets will

influence terrorist's choices of targets because they perform their own cost-benefit analysis of targets. Short-term intelligence gives the decision maker an indication of the outcome of the terrorists' target analysis. Third, we note that both warnings and physical deployments work together to influence the outcomes of terrorist attacks but through different mechanisms. As a result, decision-making processes for these two defense modalities must be integrated.

3. Private Warnings Model

In any time period, decision makers face the following situation. They have a risk profile for a wide range of potential targets that defines the expected damage to these targets from terrorism. They also have current intelligence on some of the activities of known terrorist groups. Analysis of this intelligence yields an assessment of the likelihood of an attack during that time period and a revised risk profile for the possible target locations. Given these data, the decision maker must decide where to deploy additional defensive resources, if at all, and whether or not to issue a warning of an impending attack. Notation is summarized in Table 1. We assume that, in any time period, there are discrete possible threat states of the world $s \in [0, k]$. For each state s , there is a probability p_s that a terrorist attack will be attempted in that time period. Without loss of generality, we assume that $p_s > p_{s'}$ for $s > s'$. The probability of being in state s is given by π_s . The vectors p_s and π_s in effect determine what the decision maker knows about the possible timing of an attack.

Table 1 Notation

q	Decision variable, vector of the range of deployment of defensive resources in different threat states.
w	Decision variable, vector of the warnings given in each threat state.
W	Set of threat states for which a private warning is issued.
s	Threat state of the world.
π_s	Probability of being in state s .
p_s	Probability of an attempted attack given in state s .
$f(x)$	Damage density function.
$f_i(x)$	Damage density function sorted highest to lowest.
$f_i^*(x)$	Normalized sorted damage density function.
C	Total expected damage cost if a terrorist attack occurs.
R	Cost, relative to C , of deploying physical resources per unit of territory covered.
g	Reduction in expected damage from attack for a target with extra physical defensive resources.
$\gamma(\phi)$	Fraction of damage caused by successful attack when warning is given.
p_d	Probability an attack will be deferred when a public warning is given.
β	Reduction in attack damage caused by disruption of attack deferral.

We define the set of targets as a set of locations that are uniform in terms of the resources required to defend them. For example, while some would designate both Los Angeles International Airport (LAX) and the Seattle Space Needle as potential terrorist targets, it makes little sense to view them as equivalent from a defensive standpoint. On the other hand, we can split these into subtargets like the northern, southern, eastern, and western borders of the LAX grounds, the airspace over LAX, the domestic terminal at LAX, the grounds of the Space Needle park, the airspace over the Space Needle, etc. These subtargets form a new set of targets that are less dissimilar in terms of the resources required to defend them. To simplify our analysis, we map this discrete, uniformized list of targets to the continuous unit segment $[0, 1]$. Underlying this simplification is the assumption that there are a sufficient number of targets to make a continuous model a reasonable approximation of what is inherently a discrete space.

We formalize the concept of a short-term risk profile by defining $f(x)$, with $x \in [0, 1]$, to be the *damage density function* for the location of the attack, with $F(x) = \int_0^x f(l) dl$ defined as the associated *cumulative damage density function*. The damage density function (DDF) provides a measure of the expected damage to a target location given a terror attack occurs somewhere. As such, this function incorporates data on the value, hardness, and likelihood of attack for each target. We define $f_i(x)$ to be the DDF of target locations after they have been sorted from most likely to least likely and note that it is by definition decreasing. We define $F_i(x)$ as the corresponding cumulative DDF. The current state s and DDF $f_i(x)$ are outputs of the intelligence gathering process. The state information indicates what is known about the timing of an attack and the DDF captures what is known about its location. See Appendix A for descriptive diagrams.

Given that the likely modalities of terrorist attacks, e.g., bombings, hijackings, and CBRN (chemical, biological, radiological, and nuclear), are not changing rapidly and that the values and hardness of potential targets are not changing rapidly either, it is reasonable to assume that the long-term DDF is constant over time. We further assume that the ranked short-term DDF is time invariant. The implication of this assumption is that while intelligence data may boost the likelihood of attack at some locations in one period and different locations in a later period, the overall ranked DDF will not change significantly. If the short-term DDF changed over time, it would mean that the capabilities of terrorists and or security forces were changing over time as well. While this is entirely likely, and of interest, modeling such capability shifts is beyond the scope of this paper and involves analysis of longer time horizons than we consider here.

We define $C = \int_0^1 f_i(l) dl$ as the total expected damage from a terrorist attack given that it occurred. By scaling $f_i(x)$ by C , we can form a normalized damage density function $f_i^*(x)$ defined on $[0, 1]$, with $F_i^*(x)$ the associated cumulative DDF. We define R as the cost of defending all targets with physical resources to an acceptable standard. Here, R is not measured in dollars but relative to C . If $q \in [0, 1]$ is the range of targets defended by physical resources, then the defensive cost generated is Rq . The expected damage to any target that is defended by additional physical resources is assumed to be reduced by a constant factor g . That is, if the interval of targets defended by additional resources is $[0, q]$, the expected damage from an attempted terrorist attack is $gF(q) + (1 - F(q))$.

We first formulate the decision problem assuming that the only defensive mechanism available to the decision maker is deploying physical resources. In this case, the problem is

$$P1: \quad \text{Min}_q \sum_{s=0}^k \pi_s [p_s [gF_i^*(q_s) + (1 - F_i^*(q_s))] + q_s R] \quad (1)$$

$$\text{subject to } q_s \in [0, 1]. \quad (2)$$

In each state, the decision maker must select the deployment range that minimizes total costs. We will refer to the region receiving extra defensive resources (locations with positions $x < q_s$) as *the defended region* and the rest of the space as *the undefended region*.¹ We have assumed that the benefit of deploying physical resources is the same g reduction in expected damage regardless of location and that the level of resources expended per target defended R is an exogenously

¹ In this paper, the terms undefended and defended are defined with respect to the extra resources being deployed. All locations may have some basic level of security.

defined constant. The implication of these assumptions is that the decision maker cannot choose to defend one target more than another within the defended region and that the effect of defensive measures is proportional to the expected damage from the attack. Making g and R constant greatly simplifies the optimization, and relaxing this assumption is an avenue for extending the model in future research. Having a proportional effect is motivated by the observation that there is a difference between preventing an attack on a specific target and preventing damage from an attack by a terrorist cell that was targeting a specific location. For example, a terrorist cell may be planning a truck bomb attack on a landmark building. The planning involves preparing the explosives and the truck, mapping a route to the target building that best avoids security forces, identifying the optimal points to detonate the bomb, and conducting rehearsals. In theory, enough physical resources could be deployed in such a way that even the most carefully planned truck bombing attack would be thwarted from destroying the targeted landmark. However, there are many ways the terrorist could still cause damage in the attempted attack. First, the terrorists could cause less than catastrophic damage to the target by exploding the bomb farther away from the building than planned. Second, the terrorists could kill and injure security forces and bystanders when they were blocked from attacking the target directly. Third, the terrorists might at the last minute attack a nearby, less defended, less significant building when they observe very tight security at their planned target. Fourth, an attack, even if unsuccessful at its main objective, causes fear in society with significant economic and political impact. Because the destructive capability of a terrorist cell is probably proportional to the damage they are attempting to cause by attacking a particular target, it is reasonable to assume that the impact of the defensive resources is proportional to the expected damage to the target.

We have also assumed that the deployment does not affect the DDF by shifting the likelihood of attack to a different location. First, we note that as discussed above, part of the justification for a proportional defensive effect, g , is that terrorists may shift targets at the last minute in response to defensive deployment. Second, we note that even though deployment of physical resources tends to be observable, terrorists do not always have the ability to observe these resources in time to change their plans. For example, surprise roadblocks can catch unsuspecting terrorists on their way to an attack, and raids of suspected terrorist hideouts can thwart attacks in planning stages. Also, an increased presence of undercover security is very difficult to detect. For example, a hijacker will have difficulty detecting additional air marshals

placed on a flight that is believed to be at high risk by security forces.

Private warnings are another mechanism for defending against terror attacks. We model private warnings as having zero cost and reducing the probability of a successful terror attack in any location by a multiplicative factor $\gamma \in [0, 1]$. Define w_s as an indicator variable that takes value one if a private warning is issued when in state s and zero otherwise. The decision maker must decide in which threat states he will issue warnings. Call this set W , i.e., $W \equiv \{s \mid w_s = 1\}$. Given the vectors \mathbf{w} , $\boldsymbol{\pi}$, and \mathbf{p} there is a long-run rate ϕ at which false alarms are issued.² The greater the false alarm rate the less warnings will be heeded, and therefore the less of a benefit that will come from the warnings; thus $d\gamma(\phi)/d\phi > 0$. We reformulate the problem with warnings as follows.

$$\text{P2: } \underset{\mathbf{q}, \mathbf{w}}{\text{Min}} \sum_{s \in W} \pi_s [p_s \gamma(\phi) [gF_l^*(q_s) + (1 - F_l^*(q_s))] + q_s R] \\ + \sum_{s \notin W} \pi_s [p_s [gF_l^*(q_s) + (1 - F_l^*(q_s))] + q_s R] \quad (3)$$

$$\text{subject to } \phi = \sum_{s \in W} \pi_s (1 - p_s), \quad (4)$$

$$q_s \in [0, 1], \quad (5)$$

$$w_s \text{ is 0 or 1.} \quad (6)$$

In Problem P2, we are putting a formal structure on the situation faced by a decision makers facing potential terrorist attacks. While the formulation somewhat simplifies reality, we believe that structurally it does capture the inherent trade-offs in defending against terror attacks. In our model, deploying defensive resources has a direct cost that is directly related to the size of the space defended. The deployment of physical resources in one period has no effect on future periods. For example, the fact that a roadblock at an intersection is up on a Monday does not increase the ability to discover a car bomb at that intersection on Friday. Warnings are less costly than physical resources but do have an effect across periods. For example, a false alarm on Monday can reduce the impact of a warning on Friday.

In theory, warnings could have levels of severity, e.g., low, high, etc. Each threat state could receive a different level of warning instead of the warning/no warning dichotomy we model. We do not believe that differing warning levels are practical. It is difficult for

² We are interested in the rate at which false alarms are issued and not the fraction of warnings that are false alarms. We model false alarms this way because it is believed (Mileti and Peek 2000) that when false alarms occur more frequently, recipients will be more skeptical even if the fraction of warnings that are false alarms is the same.

individuals to calibrate themselves to different levels of alertness. In fact, this has been one of the criticisms of the U.S. DHS color-coded scheme (*New York Times* 2004). It is also worth noting that since its inception soon after 9/11/2001 the warning system has only taken on two settings—yellow and orange (*New York Times* 2004). These have formed a de facto two-level system—normal and heightened state of alert.

We do not capture the fact that threat probabilities may change over time. However, as long as the impact of false alarms on warning effectiveness operates on a time scale that is significantly shorter than the timescale in which there are significant shifts in the threat probabilities, our model should be applicable. We also do not model the effect that defensive actions have on the behavior of terrorists. It is possible that terrorists who are aware of a higher level of security will shift their planned attack to a later time. If we view warnings as being given to security forces only, then it is unlikely that terrorists will be aware of them.³ In §5, we expand our model to public warnings and explicitly consider terrorist reactions to warnings.

3.1. Analysis

In this section, we investigate two sets of questions about short-term defensive measures. First, we study how the use of physical resources and warnings interact with one another. Second, we study how the information available to the decision maker regarding the timing and location of attacks influences the optimal use of defensive measures. We then illustrate our observations with a set of numerical examples.

3.1.1. Interaction Between Warnings and Physical Deployments.

PROPOSITION 1. *Given a warning vector \mathbf{w} , the optimal defensive deployment in threat state s , $q_s^*(\mathbf{w})$ is 0, 1, or is given by*

$$q_s^*(\mathbf{w}) = \begin{cases} f_i^{*-1}\left(\frac{R}{p_s(1-g)}\right) & \text{for } w_s = 0, \\ f_i^{*-1}\left(\frac{R}{\gamma(\phi)p_s(1-g)}\right) & \text{for } w_s = 1. \end{cases} \quad (7)$$

³ To the degree that changes in physical deployments are broadly observable by security forces, terrorists, and the public, they can indirectly serve as warnings. If terrorists respond to physical deployments by deferring attacks and this occurs at the same time private warnings are issued, the false alarm rate may go up. On the other hand, the false alarm rate may go down if the terrorists defer their attack to periods in which a private warning has been issued. Because terrorists will not always be able to modify their plans in response to physical deployments and because such modifications will have a mixed effect on false alarms, ignoring the effect should not strongly affect our results.

PROOF. See the appendix.

The numerators in Equation (7) are the marginal costs of broader deployment of defensive resources, while the denominators are the difference between the expected damage in a defended region versus an undefended region with warnings and without. Because f_i^* is by definition, decreasing it means that greater deployment cost reduces the scope of the deployment, while greater expected damage increases the scope of deployment.

From Equation (7), we can also see that there is a complex relationship between the warning threshold and the deployment range. When a warning is given, the deployment of physical resources is smaller than when no warning is given because $\gamma(\phi) \in [0, 1]$. At the same time, the more frequently warnings are given, i.e., the larger the set W , the less effective the warning in any particular threat state, and therefore the greater the optimal deployment. We can also observe that a warning given in a particular state reduces costs for that state regardless of the false alarm rate. This means that it is always optimal to issue a warning in at least one of the states. We state this formally in Proposition 2.

PROPOSITION 2. *In the optimal warning policy \mathbf{w}^* ,*

$$\sum_s w_s^* \geq 1. \quad (8)$$

PROOF. Follows from discussion above.

Problem P2 integrates decision making about warnings and the deployment of physical resources. In practice, the organizational structure of a country's security services may cause these decisions to be made independently leading to inferior performance. We refer to an independent decision-making approach as a *segmented defense*. In the segmented defense, the optimal physical deployment is determined by solving Problem P1. This means that the deployment is insensitive to the shape of the function $\gamma(\phi)$. The warning policy in the segmented defense is determined by choosing the \mathbf{w} that solves Problem P2 for a fixed identical deployment in each state. This is equivalent to solving the following optimization problem:

$$\text{P3: } \text{Min}_w \sum_{s \in W}^{\pi_s} [p_s \gamma(\phi)] + \sum_{s \notin W}^{\pi_s} [p_s] \quad (9)$$

$$\text{subject to } \phi = \sum_{s \in W} \pi_s (1 - p_s), \quad (10)$$

$$w_s \text{ is 0 or 1.}$$

We expect that the states in which we have a high terror threat are typically the least likely or least frequent, that is, we expect that π_s is decreasing in s . Given this observation, we see that it makes sense

to deploy resources more widely if we are in a high threat state because they are most likely to be useful as shown in Proposition 1 and the amount of time we will have them deployed is relatively short. It also makes sense to issue a warning to boost alertness and effectiveness of the defenses because in high threat states there will be fewer false alarms. In fact, as per Proposition 1, the warnings make it possible to use fewer physical resources by deploying less widely. On the other hand, deploying physical resources during the more frequent low threat states is more costly and it is in these states that the low cost of warnings is attractive. This cost trade-off can lead to a somewhat counterintuitive result. Considering warnings independently of physical deployments, one would expect that as the probability of an attack increases, warnings would become more attractive. However, higher frequency states give warnings a larger cost advantage over physical resources. For example, let us say there are three states of the world—1, 2, and 3—with $\pi_s = (0.8, 0.15, 0.05)$ and $p_s = (0.05, 0.23, 0.5)$. The fraction of attacks that occur in each state i is given by

$$A_i = \frac{p_i \pi_i}{\sum_j p_j \pi_j}. \quad (11)$$

In the example, $\mathbf{A} = (0.4, 0.35, 0.25)$. Nearly half the attacks actually occur in the lowest threat state (state 1), so holding the false alarm rate constant, the most benefit from a warning comes in the lowest threat state. Therefore, there may be situations in which it is optimal to issue a warning in low-threat-level states, not to issue a warning in an intermediate threat level, and to issue a warning in a high-threat-level state, a warning vector $\mathbf{w} = (1, 0, 1)$. Such a strategy will lead to more false alarms than $\mathbf{w} = (0, 1, 1)$ but may, in total, reduce costs if warning effectiveness is not very sensitive to false alarms.

The interaction between warnings and physical deployment has another effect. If the overall likelihood of an attack increases, one would expect defensive deployments to increase as well and total costs to increase. However, more frequent attacks reduce the false alarm rate of warnings making them more effective and reducing the need to rely as much on physical resources. The net result is that it may be the case that societies that are at a greater state of alertness because of more frequent attacks may be more efficient at defending themselves against terrorism than societies that face less frequent attacks.

The complex interactions between physical deployments and warnings described here suggest that there is a potential for significant gaps between the performance of integrated and segmented defenses. Clearly, the size of this gap will depend on parameter values. If warnings are not effective and/or resource deployments are not costly, the gap will be small.

3.1.2. Effect of Intelligence on Defensive Measures. Our model of the information available to the decision maker posits that the long-term risk profile is updated with real-time intelligence data that, over the short term, determines the likelihood of an attack, i.e., identifies the threat state, and indicated the most likely and costly targets via the damage density function. Having more or less precise intelligence says something about the shape of the damage density function, $f_i^*(x)$, and the threat state structure. We expect that more precise intelligence leads to lower total costs from terrorism.

In an ideal world, the decision maker would know exactly when an attack would occur. This is equivalent to there being two threat states—one in which an attack always occurs in the current period, and one in which an attack never occurs. In such a situation, warnings would be given whenever there was a certainty of an attack and there would be no false alarms. The least informative situation would be one in which the threat states were indistinguishable, i.e., $p_s = p$ for all s . In such a situation, the decision maker cannot know if an attack is more likely in the current period versus later periods. According to Proposition 2, in this situation it is optimal to issue warnings sometimes just to get some benefit from the increased alertness. The rationale is that without specific attack timing information, randomly issuing warnings will occasionally thwart an attack and is better than never issuing a warning. The closer the decision maker is to the ideal world case, the more effective his use of warnings.

Information theory provides a mechanism for determining how close the threat state structure is to the ideal case. We can adapt the concept of *conditional information* (Brillouin 1962) to give a measure of how much uncertainty in the occurrence of an attack is resolved by knowledge of the threat state. For our purposes, we calculate this quantity, I_π , as follows:

$$I_\pi = 1 + \frac{1}{\ln(2)} \sum_s \pi_s [p_s \ln(p_s) + (1-p_s) \ln(1-p_s)]. \quad (12)$$

In the ideal world, $I_\pi = 1$. When threat states are indistinguishable, $I_\pi = 1 + (p \ln(p) + (1-p) \ln(1-p)) / \ln(2)$. If $p = 0.5$, $I_\pi = 0$, it is at its minimum because an attack is equally likely to no attack. In the numerical examples, we show that costs increase when I_π decreases.

In an ideal world, the decision maker would also know exactly where an attack was going to occur and its damage potential. In our model, this corresponds to having intelligence data that makes the probability of an attack, and therefore the damage density, zero everywhere except one point. In this case, $f_i^*(x)$ would be an impulse function at $x = 0$. Defensive resources would be deployed only at $x = 0$ and therefore deployment costs would be zero. The

least informative situation would be one in which the damage density function was uniform, i.e., $f_i^*(x) = 1$ for all x . In this case, no one target or set of targets stands out as a priority for defensive resources. The result is that either all targets are covered at great expense, or none are worth defending.

From Equation (7), we see that if $R/(p_s(1-g)) > f_i^*(0)$, then $q^*(w) = 0$; and if $R/(p_s(1-g)) < f_i^*(1)$, then $q^*(w) = 1$. The situation in which the decision maker is least informed about where an attack will occur is the case of $f_i^*(x) = 1$ for all $x \in [0, 1]$, i.e., a uniform distribution. In all other cases, $f_i^*(0) > 1$ and $f_i^*(1) < 1$ because $f_i^*(x)$ integrates to one and is decreasing by definition. As a result, when the DDF is uniform, the optimal deployment is an all or nothing strategy, $q^*(w) = 0$ or $q^*(w) = 1$, depending on the value of $R/(p_s(1-g))$. We generalize this observation to say that the less informed the decision maker is regarding where an attack will occur, the more likely the optimal deployment is all or nothing.

One interpretation for the scenario with a uniform DDF is that over a long period of time, a terror risk homeostasis is achieved. All high-value targets are hardened enough so that from the terrorist's perspective, there is little distinction between targets. For example, passenger planes have historically been popular targets for terrorists. Enormous investments in airline security have made it increasingly difficult for terrorists to commandeer or destroy a passenger jet. On the other hand, customers in check-in desk queues at most airports are vulnerable to attacks by suicide bombers. However, any place many civilians congregate—movie theaters, buses, train platforms, etc.—are equally vulnerable to suicide bombers. As a result, all these targets become equally likely choices for the terrorists with similar damage potential.

With risk homeostasis, the only reason the damage density functions could be nonuniform, or in other words, the only reason one location becomes a more likely target than another, is if short-term intelligence information indicates terrorist planning for a particular target and mode of attack. The implication of this observation is that while long-term defensive efforts may be effective in causing a general reduction in damage from terror attacks, the prevention of specific attacks will require short-term intelligence gathering and analysis capabilities. Such capabilities would apply to predicting both the location, mode of attack, and timing of attacks. In other words, a uniform damage density function is a symptom of intelligence failure. This failure could be any combination of failure to gather, analyze, or disseminate intelligence effectively.

3.2. Illustrative Numerical Experiments

In this section, we construct illustrative numerical examples to quantify and make more concrete the

qualitative observations we have made above.⁴ Quantifying the behavior of the model is very challenging because there is great uncertainty surrounding the key parameters. Some of these parameters, such as those related to costs, are frequently estimated by risk analysts, while others, such as the effect of warnings, are not well studied.

For example, Abt (2003) analyzes the threat of a nuclear attack on the United States via seaport freight transport. They estimate that a successful nuclear attack on a major seaport would cause from hundreds of billions to several trillion dollars of costs. These calculations include the value of statistical lives lost, direct property damage, the cost of trade disruption, and other indirect costs. The range of possible damage is quite large varying by a factor of at least 10 and would correspond to the parameter C . The study also estimates the costs and benefits of several recommended security measures. They estimate that by reducing the probability of a successful attack, the expected attack cost would be reduced by a factor of 10, corresponding to $g = 0.1$ in our model. They also estimate an annual cost of \$10 billion to implement the security measures. This corresponds to a defensive deployment cost R relative to C ranging from 0.005 to 0.1. This analysis, while perhaps accurate for seaports, greatly underestimates the costs of defending against a terrorist nuclear attack. A successful nuclear attack on a major U.S. city whether a seaport or not would probably lead to damage similar to those estimated in Abt (2003), however, sea freight is only one method of conducting a nuclear attack.⁵ Achieving a factor of 10, reduction in risk from all possible modes of nuclear attack would cost far more than \$10 billion per year.

Assessing the costs and benefits of defense against a potential attack is a meaningless exercise without an assessment of the likelihood or frequency of attack. An example of an attempt to address this question is a survey conducted by Senator Richard Lugar, the chairman of the senate foreign relations committee (Lugar 2005). Experts on security policy were surveyed about their opinions regarding the threats faced by the United States from weapons of mass destruction. The survey finds that on average these experts believe that the probability of a nuclear attack on the United States in the next decade is 29.2%, of a radiological attack 40%, of a biological attack 32.6%, and of a chemical attack 30.5%. 62% of the respondents thought the probability of a nuclear attack was between 10% and 50%. The high end of this range

⁴ Experiments with a large range of parameters have been done and the results presented here are representative.

⁵ For example, a nuclear weapon could be smuggled across a land border or by air.

would justify a five-fold increase in defensive spending over the low end of the range.

We can conclude from these examples that even in areas that risk analysts and policymakers are experienced in, there is considerable uncertainty about parameter values, and when we layer on the less well-understood aspects of the problem, the uncertainties increase. As a result, it is important to use great care in interpreting the results of the model and focus on identifying insights that are general and not very sensitive to particular parameter values. In the following, we use parameter values chosen for illustrative purposes but consistent with the studies reported above.

For the numerical examples, we consider an environment in which the overall probability of a terrorist attack in any period is $\sum \pi_s p_s = 0.1$, under two threat-state scenarios each with three states. In the “no information” scenario, the probability of an attack in a state s , $p_s = 0.1$ for all s . This means that the decision maker does not have the ability to detect activity leading to an attack and by default $\pi_s = 1/3$ for all s . In the “informative states” scenario, p_s varies with s , and we set $p_s = [0.05, 0.1, 0.8]$ and $\pi_s = [0.7, 0.25, 0.05]$. We also consider two normalized DDFs—one representing the case of no information about attack location and one representing the case in which the decision maker has information that makes some set of targets much more likely. The first case we model as a uniform distribution, and the second case as a beta distribution $\beta(1, 7.5)$ representing a situation in which 80% of the potential damage is concentrated in 20% of the possible targets.

For each pair of threat state and DDF, we determine the optimal defensive strategy for three values of the deployment cost, R : 0.01 C , 0.1 C , and 0.5 C . The effect of defensive resource deployment is assumed to be $g = 0.1$. We choose these values of R to capture a wide range of deployment costs relative to terror attack damage. As a reference point, we compare the cost of deploying over the entire target set in all states with the reduction in expected damage from an attack that the deployment provides, assuming no warnings. With an overall attack probability of 0.1, the benefit of a full deployment is $(0.1)(1 - g) C = 0.09 C$. $R = 0.01 C$ represents the case where deployment costs are significantly lower than the benefits they provide. In this case, terror damage costs dominate the objective function and a full deployment will be close to optimal. $R = 0.1 C$ represents the case where deployment costs are of a similar order of magnitude to the benefits they provide. In this case, the decision maker must trade off terror damage risk and deployment costs but the objective function is relatively flat so small deviations from the optimal deployment will not greatly increase expected costs. $R = 0.5 C$ represents the case

Table 2 Warning Response Functions $\gamma(\phi)$ Used in Numerical Examples

False alarm rate ϕ	Warning effect $\gamma(\phi)$	
	Case A	Case B
$\phi \leq 0.2$	0.3	0.1
$0.2 < \phi \leq 0.5$	0.5	0.35
$0.5 < \phi \leq 0.85$	0.8	0.45
$0.85 < \phi$	1.0	1.0

where deployment costs are high and thus optimization of deployments is most critical.

The warning response function can be characterized in two ways—the effectiveness of warnings at reducing damage from terror attacks and the sensitivity of the response to false alarms. Neither of these is a well-understood phenomenon, so for the purpose of the numerical examples presented here, we consider two cases—A and B defined in Table 2. In Case A, warnings tend to be less effective and their effectiveness is degraded significantly at a lower false alarm rate than Case B. We use a discrete representation of the warning response function because it is more realistic to assume that over ranges of false alarm rates there will be little change in people’s responses to warnings.

There has yet to be research on people’s responses to terror warnings and the resulting impact on the damage caused by terrorist attacks. Therefore, the functional form chosen here and any parameter values selected are purely speculative. The fact remains that many governments issue warnings of terror to security forces and to the general public to increase alertness and preparedness. If these warnings have little effect, then there is no reason to study them further. In this paper, we assume that they do have an effect and then investigate the best way to use them. We therefore select response functions such that warnings have a noticeable impact on the outcome of terrorist attacks.

In Tables 4 and 5, respectively, we present the optimal integrated and segmented defenses for four different information scenarios with warning response Case A. Within each information scenario (see Table 3), we display the optimal deployment in each state as a triplet on a $(0, 1)$ scale, and the total expected cost for three different deployment cost levels.

Table 3 Information Scenarios

	Information scenario I	Information scenario II	Information scenario III	Information scenario IV
Damage density function	Uninformative Uniform(0, 1)	Uninformative Uniform(0, 1)	Informative Beta(1, 7.5)	Informative Beta(1, 7.5)
Threat states	Uninformative ($I_x = 0.53$)	Informative ($I_x = 0.65$)	Uninformative ($I_x = 0.53$)	Informative ($I_x = 0.65$)

Table 4 Integrated Defenses Under Different Information Scenarios (Warning Response Case A)

<i>R</i>	Information scenario I	Information scenario II	Information scenario III	Information scenario IV
Deploy				
0.01 C	(1.0, 1.0, 1.0)	(1.0, 1.0, 1.0)	(0.42, 0.48, 0.48)	(0.42, 0.42, 0.58)
0.1 C	(0, 0, 0)	(0, 0, 1.0)	(0.17, 0.26, 0.26)	(0.07, 0.17, 0.40)
0.5 C	(0, 0, 0)	(0, 0, 0)	(0, 0.05, 0.05)	(0, 0, 0.23)
Warn				
0.01 C	(1, 0, 0)	(0, 1, 1)	(1, 0, 0)	(0, 1, 1)
0.1 C	(1, 0, 0)	(0, 1, 1)	(1, 0, 0)	(0, 1, 1)
0.5 C	(1, 0, 0)	(0, 1, 1)	(1, 0, 0)	(0, 1, 1)
Cost				
0.01 C	0.0183 C	0.0168 C	0.0136 C	0.0118 C
0.1 C	0.0833 C	0.0545 C	0.0412 C	0.0354 C
0.5 C	0.0833 C	0.0675 C	0.0808 C	0.0578 C

Information scenario I represents the case in which the least information is available, the DDF is uniform, and attack probabilities are constant across all threat states. In scenario II, threat states are informative. In scenario III, the DDF is more concentrated, i.e., more informative but threat states are not. In scenario IV, both the threat states and DDF are informative. The optimal warning strategy for each state is displayed as a triplet of zeroes and ones representing, respectively, no warning and private warnings. Tables 6 and 7 display the same information but with more effective warnings, warning response Case B.

In Table 4, we see that less informative scenarios have higher costs than more informative ones. More precision regarding the location and/or timing of an attack makes more efficient resource allocation possible. When the DDF is uniform, we see that the optimal deployment is full or nothing. The results for the uninformative threat states indicate that even if the decision maker cannot distinguish one threat state

Table 5 Segmented Defenses Under Different Information Scenarios (Warning Response Case A)

<i>R</i>	Information scenario I	Information scenario II	Information scenario III	Information scenario IV
Deploy				
0.01 C	(1.0, 1.0, 1.0)	(1.0, 1.0, 1.0)	(0.48, 0.48, 0.48)	(0.42, 0.48, 0.62)
0.1 C	(0, 0, 0)	(0, 0, 1.0)	(0.26, 0.26, 0.26)	(0.17, 0.26, 0.46)
0.5 C	(0, 0, 0)	(0, 0, 1.0)	(0.05, 0.05, 0.05)	(0, 0.05, 0.31)
Warn				
0.01 C	(1, 0, 0)	(0, 1, 1)	(1, 0, 0)	(0, 1, 1)
0.1 C	(1, 0, 0)	(0, 1, 1)	(1, 0, 0)	(0, 1, 1)
0.5 C	(1, 0, 0)	(0, 1, 1)	(1, 0, 0)	(0, 1, 1)
Cost				
0.01 C	0.0183 C	0.0168 C	0.0137 C	0.0118 C
0.1 C	0.0833 C	0.0545 C	0.0421 C	0.0365 C
0.5 C	0.0833 C	0.0745 C	0.0839 C	0.0607 C

Table 6 Integrated Defenses Under Different Information Scenarios (Warning Response Case B)

<i>R</i>	Information scenario I	Information scenario II	Information scenario III	Information scenario IV
Deploy				
0.01 C	(1.0, 1.0, 1.0)	(1.0, 1.0, 1.0)	(0.48, 0.41, 0.41)	(0.34, 0.48, 0.57)
0.1 C	(0, 0, 0)	(0, 0, 1.0)	(0.26, 0.6, 0.16)	(0.06, 0.26, 0.39)
0.5 C	(0, 0, 0)	(0, 0, 0)	(0.05, 0, 0)	(0, 0.05, 0.22)
Warn				
0.01 C	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)	(1, 0, 1)
0.1 C	(0, 1, 1)	(1, 0, 1)	(0, 1, 1)	(1, 0, 1)
0.5 C	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)	(1, 0, 1)
Cost				
0.01 C	0.0163 C	0.0158 C	0.0114 C	0.0106 C
0.1 C	0.0633 C	0.0476 C	0.0361 C	0.0302 C
0.5 C	0.0633 C	0.0577 C	0.0621 C	0.0496 C

from another, it still makes sense to issue warnings some of the time, but deployments are lower when warnings are issued than when warnings are not issued. In Table 5, we see that the segmented defense costs are always greater than or equal to the integrated defense costs. Comparing with the results in Tables 6 and 7, for warning response Case B, we see that the gap between the integrated and segmented defenses increases as the effectiveness of warnings increases.

Scenario I in Tables 4 and 6 can be interpreted as a situation in which an intelligence management failure has occurred. Comparing the results of scenario I with scenario II indicates the value of conveying threat state information to those decision makers that control physical deployments. The comparison also indicates how better threat state intelligence can lead to more effective and selective use of warnings. Comparing the results of scenarios I and II with scenarios III and IV, we can see how an inability to

Table 7 Segmented Defenses Under Different Information Scenarios (Warning Response Case B)

<i>R</i>	Information scenario I	Information scenario II	Information scenario III	Information scenario IV
Deploy				
0.01 C	(1.0, 1.0, 1.0)	(1.0, 1.0, 1.0)	(0.48, 0.48, 0.48)	(0.42, 0.48, 0.62)
0.1 C	(0, 0, 0)	(0, 0, 1.0)	(0.26, 0.26, 0.26)	(0.17, 0.26, 0.46)
0.5 C	(0, 0, 0)	(0, 0, 1.0)	(0.05, 0.05, 0.05)	(0, 0.05, 0.31)
Warn				
0.01 C	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)
0.1 C	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)
0.5 C	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)	(0, 1, 1)
Cost				
0.01 C	0.0163 C	0.0158 C	0.0115 C	0.0108 C
0.1 C	0.0633 C	0.0502 C	0.0381 C	0.0351 C
0.5 C	0.0633 C	0.0702 C	0.0692 C	0.0570 C

turn intelligence into operational assessments of types and locations of attacks severely limits the use of physical deployments. For example, in the case of $R = 0.1$, it is optimal to use more physical resources when the DDF is concentrated than when it is uniform. The result is that when the DDF is uniform, a higher proportion of the costs incurred are from terror than when the DDF is more informative. Comparing scenarios II and IV in Table 6, we see that as the DDF changes the physical deployments can change dramatically but so does the optimal warning policy. These observations highlight the fact that warnings cannot be managed independently of physical deployments.

3.3. Summary

We have found that the intelligence available to decision makers strongly influences the optimal defensive measures taken. We also find that when intelligence is more precise, it can best be exploited by integrating the decision-making process for using warnings and physical deployments. This integration requires the ability to manage the complex interaction between warnings and deployments. These observations suggest that some organizational structures and decision-making processes for defense against terrorism may be more appropriate than others.

Our analysis also highlights some important political challenges in defending against terrorism. Issuing a warning because it is better than not issuing a warning, even if the government does not know the threat is any greater, seems dishonest. If deployment when threat states are high are not significantly different than when threat states are low because warnings have boosted alertness, the public may feel that the government is not doing enough to defend them. Not deploying any additional physical resources because the attack location uncertainty is large is also politically risky. In the next section, where we consider public warnings, the political challenges become even greater.

4. Public Warnings

We now broaden our analysis to include the decision of when to issue public warnings in addition to private warnings of terror attacks. Public warnings may reduce the expected damage from a terrorist attack further than a private warning because the entire population is potentially taking actions that could help thwart attacks and/or diminish their damage. However, public warnings differ from private warnings because terrorists themselves are aware of the warnings. A public warning may therefore have a behavioral effect on terrorists. Terrorists that are planning an attack in one period may defer their planned attack to a later period if a public warning is issued. When

terrorists defer their planned attack, we assume that they are self-imposing a delay between their planning and preparations for an attack and their actual conduct of the attack. During this delay, they are more difficult to detect because they are inactive. Deferring the attack also has potential to be disruptive to the terrorists. During the delay, the terrorists could be captured or their planning could become outdated and less effective. At the same time, if an attack is deferred when a public warning is issued, it makes that warning a false alarm, thus reducing the effectiveness of future warnings. In this section, we extend Problem P2 to include the issues related to public warnings described above.

4.1. Model Formulation

With public warnings, in each threat state the decision maker can issue no warning, a private warning, or a public warning corresponding to values 0, 1, or 2 for decision variable w_s . We assume that when a public warning is given, a private warning is given as well because the security forces are also part of the general public. We define separate false alarm rates for the security forces and the public ϕ_1 and ϕ_2 . When a public warning is issued and no attack occurs, it counts as a false alarm for both the public and the security forces. When a private warning is issued and no attack occurs, it only counts as a false alarm for the security forces. In the case when an attack was going to happen but a public warning was issued, we assume that with probability p_d the terrorists will defer the attack to a period in which no public warning is issued. We assume further that the decision maker will never issue public warnings in all states because if the government always issued public warnings, the terrorists would have no reason to defer their attacks and the analysis of terrorist response to public warnings would become trivial. Finally, we define $\beta \in [0, 1]$ as the factor by which the damage from an attack is reduced due to the disruption to the terrorists of deferring the attack.

If we define S_i as the set of states s for which $w_s = i$, we have

$$\phi_1 = \sum_{s \in S_1} \pi_s (1 - p_s) + \sum_{s \in S_2} \pi_s \left(1 - p_s + p_s p_d \frac{\sum_{s \in S_0} \pi_s}{1 - \sum_{s \in S_2} \pi_s} \right), \quad (13)$$

$$\phi_2 = \sum_{s \in S_2} \pi_s (1 - p_s + p_s p_d). \quad (14)$$

We assume that there are distinct warning effect functions for private and public warnings, $\gamma_1(\phi_1)$ and $\gamma_2(\phi_2)$, respectively, that have a multiplicative effect. To simplify notation, we define the following expres-

sions for the expected cost in a state s for each warning level:⁶

$$\begin{aligned} H_0(s) &= [gF_l^*(q_s) + (1 - F_l^*(q_s))], \\ H_1(s) &= \gamma_1(\phi_1)[gF_l^*(q_s) + (1 - F_l^*(q_s))], \\ H_2(s) &= \gamma_1(\phi_1)\gamma_2(\phi_2)[gF_l^*(q_s) + (1 - F_l^*(q_s))]. \end{aligned} \quad (15)$$

The decision problem with private and public warnings is therefore

$$\begin{aligned} \text{P4: Min}_{q,w} \sum_{s \in S_0} \pi_s(p_s H_0(s) + q_s R) &+ \sum_{s \in S_1} \pi_s(p_s H_1(s) + q_s R) \\ &+ \sum_{s \in S_2} \pi_s \left[q_s R + (1 - p_d)p_s H_2(s) + (1 - \beta)p_d p_s \right. \\ &\quad \left. \cdot \frac{(\sum_{s' \in S_0} \pi_{s'} H_0(s') + \sum_{s' \in S_1} \pi_{s'} H_1(s'))}{\sum_{s' \notin S_2} \pi_{s'}} \right]. \end{aligned} \quad (16)$$

The first term of Equation (16) is simply the expected cost in states with no warnings, while the second term is the expected cost in states with a private warning. The third term is the expected cost if a public warning is given which itself reflects three cases. Case (1): with probability $(1 - p_s)$, no attack occurs and the only cost comes from deployment of physical resources. Case (2): with probability $(1 - p_d)p_s$, an attack is intended, is not deferred, and occurs leading to an expected cost $H_2(s)$ that includes the joint effects of public and private warnings. Case (3): with probability $p_d p_s$, an intended attack is deferred to a later period with no public warning and is carried out with damage reduced by β . The cost in Case (3) depends on whether the attack is deferred to a state in which there is no warning in effect or a state in which there is a private warning in effect. We can see from the formulation of Problem P4 that the effect of public warnings is similar to private warnings. A public warning reduces the expected damage from an attack made in that period but this reduction depends on the false alarm rate. However, public warnings lead to some attacks being deferred to other periods. If the period in which the deferred attack eventually takes place is lightly defended, then the public warning could lead to more damage than if the warning had never been given. If β is large, a public warning may lead to a reduction in damage because it disrupts the terrorist's plans, giving security forces more time to apprehend them.

To simplify the model, the parameters β and p_d are exogenous. This does not prevent us from considering how they may be related in reality. For example, if β is high and p_d is low, it suggests that terrorists do not defer attacks when there is a public warning

because they believe that they may soon be apprehended or the delay would be too disruptive in other ways. If β is low and p_d is low, it reflects a situation in which the terrorist believes that the public warning is an indication that the security forces are closing in on them while in reality they are not. In such a setting, the public warning serves mainly to raise the public's alertness. If β is high and p_d is high, it reflects a situation in which the terrorists do not realize that security forces are closing in and so are more worried about the heightened alertness in the current period indicated by the public warning. When β is low and p_d is high, it reflects a situation in which terrorists believe that they are hard to detect or in which they go deeper underground making themselves harder to catch until a more opportune time arrives for an attack. The probability of deferring an attack, p_d , can also be low if terrorists believe that public warnings are not viewed as credible by the public. In the following, we conduct numerical experiments to determine how the values of β and p_d affect the decision making related to public warnings.

4.2. Analysis

In §3, we have shown how increased uncertainty about the timing and location of attacks increases the cost of terrorism. Because public warnings may lead to the temporal shifting of terror attacks, they indirectly alter the threat state structure. Therefore, they can in effect increase or decrease the uncertainty in the timing of attacks. The optimal use of public warning requires an understanding of this phenomenon.

If the probability that terrorists will defer their attack, p_d , is low, then public warnings essentially function like private warnings. When p_d is high and the disruption caused by attack deferral is high, β low in our model, using public warnings to cause disruption to terrorists becomes the dominant strategy. The situation is more complex when p_d is moderately high or high and the disruption effect is small, β high. In the following, we focus on this case.

If the threat states are, for example, uninformative, issuing a public warning in one state can lead to a greater concentration of attacks in the states without private warnings, making it easier to defend. There is a limit to using public warnings in this way because too many false alarms can make both the use of private and public warnings less effective. In situations in which the threat states are informative, using public warnings in a high-threat state can have the opposite effect. The timing of attacks will become less certain because a fraction of the attacks in the most likely time periods, high-threat states, are shifted to another time.

It is possible to glean some useful rules of thumb regarding public warnings from these observations.

⁶ For a public warning, $H_2(s)$ is only a component of the expected cost.

First, if the government is not close to catching the terrorists and is uncertain how the terrorists will respond to the public warning, then it is probably best not to issue public warnings. Second, when the government has good intelligence about the timing of an attack, i.e., the society is in a significantly higher threat state, it is better to issue a private warning, deploy physical resources, and try to stop any attacks rather than take the risk that a public warning will cause the terrorists to defer their attacks and introduce more uncertainty into when an attack will take actually place. Third, if there is great uncertainty about the timing of an attack, public warnings can be used to influence terrorists so that the timing will be made more predictable. In essence, these rules of thumb are stating that the role public warnings play in influencing terrorist behavior is at least as important as its role in preparing the public for possible attack.

4.3. Numerical Experiments

Table 8 displays the optimal deployment ranges, warning levels, and expected cost for different values of p_d for the four information scenarios examined in Tables 3 and 4. We set $R = 0.1$, $C = 1$, $g = 0.1$, and $\beta = 1$ and use warning response Case B for the effect of both the public and private warnings. The deployment ranges are displayed on $[0, 1]$ and warning levels (0, 1, 2) denote no warning, private warning only, and both public and private warning, respectively.

In Table 8, we can see examples of several properties of public warnings. We see that for all the information scenarios except the least informative one (upper left quadrant), certain patterns appear consistently. We have that the least costly outcome is when the probability of deferral is at its highest and the most costly outcome is when $p_d = 0.5$. When $p_d = 0.95$, the decision maker is able to influence the timing of the attacks to his advantage, i.e., to periods

in which there is the greatest deployment of physical resources. When $p_d = 0.5$, the decision maker is least certain what the effect of the public warning will be and thus may even prefer not to use public warnings. In the scenario with the least information (upper left quadrant), it is optimal to rely only on warnings as a defensive measure. This means that false alarms are especially bad and so when p_d is high, it is optimal to use only private warnings. The optimal warning and deployment policies also tend to place the lowest deployments when public warnings are issued. This approach would be politically difficult to accept by the public. We also see in Table 8 that the optimal public warning policy is more sensitive to the information scenarios and problem parameters than in the private warning examples.

4.4. Summary

Mathematical models of public warnings are problematic because their outcomes are determined by the responses of the public and terrorists. Neither of these sets of agents are well understood, and so our predictive ability is limited. For example, we have modeled the terrorist response to a public warning, p_d , as an exogenous parameter. A game-theoretic approach to this problem would endogenize this parameter. In such a framework, the terrorist would be trying to strategically influence the false alarm rate when deciding to defer an attack. While such an analysis is beyond the scope of this paper, it is not clear that such a sophisticated behavior on the part of a terrorist cell is realistic. A major obstacle to such behavior by terrorists is that it is difficult for them to assess if a public warning was caused by their feints or not, while at the same time they are exposing themselves to additional risk of capture. Green and Armstrong (2004) and Green (2002) demonstrate that role playing simulations may be the best methodology to apply to analyzing conflict situations such as modeled in this paper.

Despite the limitations of the methodology applied here, our analysis does point to some inherent problems with the use of public warnings. Uncertainty about the quality of intelligence on attack location and timing and terrorist behavior causes significant uncertainty about when and how to use public warnings. At the same time, the optimal (from the perspective of economic cost minimization) use of warnings leads to policies that may be politically untenable. Similarly, the politically palatable policies of combining public warnings with larger physical deployments lead to poor use of defensive resources and revelation of information to terrorists. We conclude that governments may be better off avoiding public warnings altogether, relying primarily on private warnings instead.

Table 8 Use of Public Warnings Under Different Information Scenarios (Warning Response Case B)

p_d	Information scenario I	Information scenario II	Information scenario III	Information scenario IV
Deploy				
0.05	(0, 0, 0)	(0, 0, 1.0)	(0.03, 0.03, 0.26)	(0, 0.26, 0.30)
0.50	(0, 0, 0)	(0, 0, 1.0)	(0, 0, 0.29)	(0, 0.27, 0.39)
0.95	(0, 0, 0)	(0, 0, 1.0)	(0, 0, 0.39)	(0, 0, 0.46)
Warn				
0.05	(2, 2, 0)	(2, 0, 1)	(2, 2, 0)	(2, 0, 2)
0.50	(2, 1, 1)	(2, 2, 0)	(2, 2, 0)	(2, 0, 1)
0.95	(1, 1, 0)	(2, 2, 0)	(2, 2, 0)	(2, 2, 0)
Cost				
0.05	0.0495 C	0.0400 C	0.0283 C	0.0217 C
0.50	0.0630 C	0.0450 C	0.0287 C	0.0224 C
0.95	0.0633 C	0.0234 C	0.0280 C	0.0185 C

5. Conclusion

In this paper, we have developed a structured framework for assessing options for defending against terrorist attacks. We have modeled three key decisions governments must make in defending against terrorist attacks: how widely to deploy physical defensive resources, when to issue warnings to the security forces in the field, and when to issue warnings to the general population. Physical resources can be selectively deployed but are costly. Warnings are applied broadly and are inexpensive but are limited in effectiveness by false alarms. The model takes the three major forms of intelligence data a government may have about a terrorist attack—its timing, its location, and its potential severity as inputs. The model uses these inputs to choose defensive actions and considers how security forces, the public, and terrorists may respond to these actions. By developing and analyzing such a structured quantitative model, we have derived a number of results.

We have argued that optimal defensive measures and outcomes are sensitive to the intelligence data available. We have also argued that the interaction between warnings and physical deployments are complex and that there is value to integrating decision making about these defensive measures. We have identified limitations of public warnings and the conditions under which they are more or less useful. We have also argued that determining the optimal defensive measures involves analyzing subtle trade-offs that do not always yield outcomes that are easily accepted politically.

The gathering and synthesizing of intelligence in real time is a process that will at best yield incomplete assessments with great uncertainty about the timing and location of terrorist attacks. This uncertainty forces government decision makers to make difficult cost-benefit trade-offs in the defensive measures they take. From a political standpoint, failure to deploy physical resources because the threat is not deemed great enough to justify the cost may be unacceptable if an attack indeed occurs in the undefended region. The political challenges are further exacerbated in the case of warnings. Warnings are made known to a broad audience who, if warnings are used optimally as dictated by our model, to some degree are being manipulated by the government. Warnings are the government's way of making the recipients or warnings co-producers of the defense against terrorism. To the degree that warnings serve as substitutes for physical deployments, the warning recipients are in effect being asked to accept being more alert so that physical deployments can be reduced. As discussed in §3, it is also possible that it is optimal to issue a warning in states with lower attack probabilities. Such a strategy takes advantage of the fact that warning

recipients believe that a warning indicates an attack is more likely.

Finding a way to effectively combat terrorism while maintaining the quality and standard of life is one of the greatest challenges facing democratic societies today. Because of their openness, democratic societies are the most vulnerable to terrorist attack. A society that becomes significantly less democratic and/or spends a large proportion of its wealth on anti-terror efforts cannot be viewed as having successfully combated terrorism. Therefore, it is important for a government to make the best use of its available defensive resources. An important element in this effort is a rational decision-making process. In an open society, public opinion plays a strong role in the political system and as a result common perceptions about how best to defend against terrorism have influence regardless of their correctness. An informed decision-making process that is robust to uncertainty in the quality of intelligence is valuable in helping decision makers avoid being guided entirely by the demands of public opinion. The modeling and analysis presented in this paper provide a step forward in defining the information needed to effectively defend against terrorism, illuminating the relationships between the major factors, and in quantifying the value of better intelligence management.

Acknowledgments

The author thanks the department editor, the area editor, and two anonymous reviewers for thoughtful and constructive critiques and suggestions that greatly improved this paper.

Appendix

PROOF OF PROPOSITION 1. When no warning is given in state s and defensive deployment is q_s , the expected cost is given by

$$V(q_s) = p_s[gF_i^*(q_s) + (1 - F_i^*(q_s))] + q_s R,$$

$$\frac{dV}{dq_s} = p_s(g - 1)f_i^*(q_s) + R,$$

$$\frac{d^2V}{dq_s^2} = p_s(g - 1)f_i^{*'}(q_s).$$

By definition, $f_i^{*'}(q_s) \geq 0$ because we rank location in decreasing order of likelihood of attack and $g < 1$. We can then apply the first-order condition and have that if the optimum deployment q_s^* is an interior solution, then it satisfies

$$\frac{R}{p_s(1 - g)} = f_i^*(q_s^*).$$

The same logic applies to the case when a warning is given in state s .

Figures A1 through A4 show the logical progression from a discrete target list with associated expected damage from terror attacks to the sorted DDF $f_i(\cdot)$. Figure A1 shows a set

Figure A1 Discrete Expected Damage Distribution




Figure A2 Uniformized Discrete Expected Damage Distribution






Figure A3 Uniformized Discrete Expected Damage Distribution (Ordered)



of target types and represents the expected damage combining the likelihood that a target is actually attacked and the potential for damage on some scale. Figure A2 is the same after target types have been broken down into uniformized targets that are similar in terms of the defensive resources required. This clearly leads to a substantially greater number of targets. Figure A3 is the same as Figure A2 but is

Figure A4 Damage Density Function (Ordered)



sorted in descending order. Figure A4 is the continuous approximation of Figure A3 mapped to the unit segment and represents a DDF prior to normalization.

References

Abt, C. C. 2003. *The Economic Impact of Nuclear Terrorist Attacks on Freight Transport Systems in an Age of Seaport Vulnerability—Executive Summary*. Abt Associates, Cambridge, MA.

Brillouin, L. 1962. *Science and Information Theory*, 2nd ed. Academic Press, New York.

Dudkevitch, M. 2004. Official: Israel facing 57 daily terror attack warnings. *The Jerusalem Post* (July 22).

Enders, W., T. Sandler. 1993. The effectiveness of antiterrorism policies: A vector-autoregression intervention analysis. *Amer. Political Sci. Rev.* 87 829–844.

Enders, W., T. Sandler. 2004. What do we know about the substitution effect in transnational terrorism? A. Silke, ed. *Research on Terrorism: Trends, Achievements, and Failures*. Frank Cass Publishers, London, UK, 119–137.

Green, K. C. 2002. Forecasting decisions in conflict situations: A comparison of game theory, role-playing, and unaided judgement. *Internat. J. Forecasting* 18 321–344.

Green, K. C., J. S. Armstrong. 2004. Value of expertise for forecasting decisions in conflicts. Working Paper 27/04, Department of Econometrics and Business Statistics, Monash University, Clayton, Australia.

Keohane, N. O., R. J. Zeckhauser. 2003. The ecology of terror defense. *J. Risk Uncertainty* 26(2) 201–229.

Lugar, R. 2005. *The Lugar Survey on Proliferation Threats and Responses*. U.S. Senate Foreign Relations Committee, Washington, D.C.

Mileti, D. S., L. Peek. 2000. The social psychology of public response to warnings of a nuclear power plant accident. *J. Hazardous Materials* 75 181–194.

Myre, G. 2004. With about 50 warnings daily, Israel handles most quietly. *The New York Times* (August 6) A9.

New York Times. 2004. Editorial: The Terror Alerts. *The New York Times* (August 5) A22.

Pate-Cornell, M. E. 1986. Warning systems in risk management. *Risk Management* 6(2) 223–234.

Pate-Cornell, M. E., C. P. Benito-Claudio. 1984. Warning systems: Response models and optimization. *Proc. Soc. for Risk Anal. Internat. Workshop on Uncertainty in Risk Assessment, Risk Management, and Decision Making*, Knoxville, TN, 457–468.

Risk Management Solutions. 2005. Managing terrorism risk, http://www.rms.com/publications/terrorism_risk_modeling.pdf.