4-1-2012

The Interdependence Between Homeland Security Efforts of a State and a Terrorist's Choice of Attack

Timothy Mathews
Kennesaw State University, tmathew7@kennesaw.edu

Anton D. Lowenberg
California State University Northridge, anton.lowenberg@csun.edu

Recommended Citation

The Interdependence Between Homeland Security Efforts of a State and a Terrorist's Choice of Attack

April 8, 2011

Abstract: Consider a state that chooses security levels at two sites (Targets A and B), after which a terrorist chooses which site to attack (and potentially a scope of attack). The state values A more highly. If the state knows which target the terrorist values more highly, he will choose a higher level of security at this site. Under complete information, if the terrorist's only choice is which site to attack, the state will set security levels for which the terrorist prefers to attack A over B if and only if the ratio of the value of B to the value of A is greater for the state than for the terrorist. When the state has incomplete information on the terrorist's target values, the optimal security levels may be such that: a target is completely undefended (but attacked with positive probability); the probability of attack is greater at A than at B; and the expected damage from an attack is greater at A than at B. In total, the results reveal that the state's choice of security is heavily influenced by the terrorist's target valuations.

1. Introduction

There is strong evidence that terrorists are rational and respond to incentives. For example, if tougher measures are taken against skyjackings or attacks on airports, terrorists may instead target trains; if security at embassies is improved, terrorists may switch to kidnappings. Enders and Sandler (1993, 1995, 2004) and Sandler and Enders (2004) offer empirical support for terrorists' rationality in the form of observable responses to changes in constraints (e.g., substitution away from skyjackings to kidnappings after airports installed metal detectors).[1] Thus, expenditures on counterterrorism very often have the effect of displacing terrorist attacks from one target to another (*Economist*, March 8, 2008, p. 69).

Our purpose is to analyze a government's decision on how to allocate protective resources across different targets within a single country. These alternative targets can be interpreted either as geographical venues (e.g., cities versus borders) or as different types of attacks (e.g., nuclear versus biological).[2] In this respect, our approach differs from much of the existing literature on strategic counterterrorism, which focuses on interactions between different target countries or between targets within the same country that are treated as independent decisionmakers. Our interest in a centralized decisionmaker is motivated in part by a growing concern among scholars and the media with the costs and inefficiencies of counterterrorism policies. A high-profile critique of U.S. spending on the Iraq war[3] by Stiglitz and Bilmes (2008),[4] puts the cost of the war at an exorbitant $3 trillion and argues that the U.S. government adopted a shortsighted allocation of resources, economizing on the costs of equipment for the troops at

---

[1] Enders, Sandler, and Cauley (1990) use time-series techniques to evaluate short-, medium- and long-run effects of specific terrorist-thwarting policies, such as metal detectors, enhanced security for U.S. embassies and personnel, and the U.S. retaliatory strike against Libya in April 1986. Brandt and Sandler (2010) document empirically how the hardening of some targets has led to more attacks on other targets, for example, as government and military targets have been defended more successfully, terrorists have substituted toward increased attacks on businesses and private individuals.

[2] See Bier, Oliveros, and Samuelson (2007, p. 564).

[3] While the Iraq war may not be considered a counterterrorism expenditure, a case can be made that, at least at its outset, it was a component of the "war on terror." See Lowenberg and Mathews (2008).

[4] See also Bilmes and Stiglitz (2006).

the expense of higher casualties and medical care burdens in the future.[5] A study by Sandler, Arce, and Enders (2008) suggests that much of the spending on counterterrorism fails a simple cost-benefit test. They estimate that global spending on homeland security has risen by $65 billion to $200 billion per year since 2001. In their calculation, the aggressive increase in counterterrorism spending in 2002-03 following the 9/11 attacks produced only five to eight cents of benefits per dollar spent.[6] Thus, it is useful to closely examine how a target country should allocate counterterrorism resources, especially since terrorists' behavior can be influenced by the defensive actions of a target state.

Sandler and Arce (2007) provide an excellent survey of the vast literature on game theoretic analyses of terrorism and counterterrorism policies. More recent developments in the literature are discussed by Sandler and Siqueira (2009). As these authors note, the fundamental nature of the interaction between terrorists and their targets is strategic, often due to the presence of many uncompensated interdependencies in the form of public good and externality effects. An example, noted early in the history of such studies, is that, in a three-player game consisting of two potential target governments and a terrorist group, defensive policies initiated independently by one target government create negative spillover effects for other targets, by deflecting terrorist attacks to the relatively "soft" target (Sandler and Lapan, 1988; Sandler and Siqueira, 2006). That is, defensive policies are strategic complements: one target's defensive policies potentially deflect attacks to the other target, requiring the latter to respond with its own defensive measures. When defensive resources are chosen by two independent decisionmakers, the result is a costly and Pareto inferior deterrence race among target countries. Such a game is effectively an open-access commons problem, or congestion externality,[7] in which the socially efficient outcome is mutual inaction whereas the dominant strategy is to take action.[8]

However, this effect of "increasing defensive efforts at one target makes other targets more desirable"

---

[5]*Economist*, March 15, 2008, p. 98. Other analysts, however, have accused Stiglitz and Bilmes of overstating some costs.

[6]See *Economist*, March 8, 2008, p. 69.

[7]See Keohane and Zeckhauser (2003).

[8]Cadigan and Schmitt (2010) verify, using laboratory experiments, the finding that negative externalities produce overinvestment in deterrence expenditures.

is present even when levels of protective resources are chosen by a single decisionmaker. Consider a state setting security levels at two sites, Target A and Target B. When assessing the terrorists' desirability for attacking A, the state need not be worried about "external costs" stemming from the fact that A becomes a relatively more desirable target when B is more heavily defended, not because the costs are absent but rather because the costs are not external. That is, when choosing defensive efforts at each site, it is critical for the state to be cognizant of and carefully account for such cross target effects. In the present paper, we analyze the decision of a single government regarding levels of protection at alternative targets. We conclude this section by summarizing the recent work most closely related to the present study.

Powell (2007) analyzes a government's allocation of a fixed amount of defensive resources across multiple domestic target venues, under the assumption that the government is uninformed about the terrorists' preferences regarding choice of target. When making its choice, the government must account for: the loss incurred as a result of a successful attack on a given venue; the vulnerability of the venue (i.e., the probability an attack will succeed); and the probability of an attack occurring, which is determined by the terrorists' choice of target. Powell focuses on an environment in which a target becomes less attractive to the terrorists when it is more heavily defended.[9] Under this assumption, Powell clearly explains how increasing defensive resources at a venue has competing effects: the marginal expected loss to the government at the first venue declines, while that at the alternative venue increases. He extends his model to the N-target case, using a minimax algorithm to solve the game in order to illustrate how the government would optimally allocate defensive resources.

The two papers most closely related to the present study are Bier, Oliveros, and Samuelson (2007) and Zhuang and Bier (2007). In contrast to Powell (2007), in each of these studies the total amount of resources available to the defender is endogenously determined. Further, in Zhuang and Bier (2007) the

---

[9]This is a simplifying assumption, which is not always satisfied in practice. For example, terrorists could be motivated by the prestige of successfully attacking a heavily guarded site, in which case increasing defensive resources could make a target more attractive to terrorists. Such considerations are beyond the scope of both Powell (2007) and the present study.

terrorists get to choose both the location and scope of their attack (i.e., the amount of resources devoted to the attack), but these authors examine only an environment of complete information, in which both players know the true values of all relevant parameters. In contrast, Bier, Oliveros, and Samuelson (2007) consider an environment in which there is incomplete information regarding the value that the attacker places on each target, but they do not allow the attacker to have a choice regarding the scope of attack. In further contrast to Powell (2007), each of these studies provides some discussion of instances in which staging an attack is costly for the terrorists (so that an attack is not always staged with certainty).

Several interesting results are derived from these models. First, it is shown that the terrorists' attack efforts and the state's defensive investments may be strategic complements (i.e., increased defense expenditures by the government may lead to increased attack efforts by the terrorists and vice versa). In such cases, increased levels of activity by either party may ultimately be mutually welfare reducing. These authors also demonstrate that the state has a strategic interest in moving first, thereby inducing the terrorists to attack more vulnerable but less valuable targets. Due to this first-mover advantage, the government may prefer its defense allocations to be observable, as opposed to unobservable.[10] These models (as well as those of Farrow (2007) and Lee (2007)) also illustrate that a venue may be left undefended in equilibrium, even if it is subject to a positive probability of attack, if the expected losses to the government from an attack are low relative to the opportunity cost of enhanced defenses.

In the present study we focus on the choice of defensive resources across two targets by a centralized decisionmaker. As in Bier, Oliveros, and Samuelson (2007) and Zhuang and Bier (2007), the total amount of defensive resources is endogenously determined. Further, we generally allow for incomplete information regarding the terrorist's "type" and allow the terrorist to choose both a target and scope of

---

[10]However, Zhuang and Bier (2011) and Zhuang, Bier, and Alagoz (2010), allowing for two-sided uncertainty in which neither player knows the other's capabilities or efforts, show that secrecy by the government may be optimal as a consequence of the endogenous response of the terrorists. Bernhardt and Polborn (2010) argue that, due to non-convexities in defense strategies, a government that attaches similar values to alternative targets might benefit from concealing its defense allocations from the terrorists.

attack, assumptions which when made together differentiate this study from the existing literature.[11]

In Section 2 a theoretical model is specified and preliminary insights on the behavior of both the terrorist and state are obtained. It is shown that if the terrorist is certain to value a specific target more highly, then the state will allocate a greater amount of security resources to this target (even if this is not the target that the state values more highly). A closer examination of situations of complete information is provided in Section 3. Within this discussion, it is shown how the preference of the state regarding where an attack is ultimately staged in such a setting depends critically upon the ratio of target valuations for both the state and terrorist. Further, it is noted how it may be best for the state to leave a target undefended and may be best for the state to choose security levels sufficiently high so that no attack is staged. Additional insights on situations of incomplete information are provided in Section 4. It is noted how it may be best for the state to leave a target completely undefended, even if it will be attacked with positive probability (an insight obtained in previous studies). Further, we compare the equilibrium "probability of attack" and "expected damage" across targets (comparisons not made in previous studies) to illustrate that it may be best for the state to choose security levels: for which an attack is more likely to be staged on the target that is more highly valued by the state; and for which the expected damage from a terrorist attack is greater at the target which is more highly valued by the state. In total, our results reveal the strong degree to which the state's choice of security is driven by the target valuations of the terrorist. Section 5 concludes.

2. Theoretical Model

Consider a terrorist organization (denoted T) looking to stage an attack within the homeland of a state

---

[11]Additionally, several minor modeling assumptions differ from previous studies. For example, when modeling the choice of security, Bier, Oliveros, and Samuelson (2007) allow the state to have a direct choice of the "success probability for the terrorists' attack" while incurring security costs that are non-linear in this probability. In contrast, we assume security costs are a linear function while the terrorists' success probability is a non-linear function of the government's security level.

(denoted S).[12] Suppose T will attack at most one of two targets (denoted Target A and Target B). Let $V_{Si} > 0$ denote the value that S places on Target i, for $i \in \{A, B\}$. Likewise, let V_{Ti} denote the value that T places on Target i, for $i \in \{A, B\}$. We assume throughout that the values of V_{SA} and V_{SB} are known by both players, but we allow for incomplete information regarding the true values of V_{TA} and V_{TB}.[13] The true pair of (V_{TA}, V_{TB}) identifies the type of T. Let V_T denote the set of possible types of T, and suppose $(V_{TA}, V_{TB}) \in V_T$ is determined as the realization of a single draw from the joint distribution function $G_T(v_{TA}, v_{TB})$. Assume V_T is a convex set, with $G_T(v_{TA}, v_{TB})$ placing strictly positive probability on all points in V_T. By incomplete information, we mean that while T knows the true values of (V_{TA}, V_{TB}), S does not know these exact values but rather knows only the distribution function $G_T(v_{TA}, v_{TB})$ from which the type of T is determined.

Before T chooses which target to attack, S sets levels of homeland security at each site ($h_i \geq 0$ denotes the level at Target i). After observing the security levels, T chooses which target to attack (with no attack being an option) along with the scope of attack ($r_i \geq 0$ denotes the amount of resources devoted to attacking Target i). The expected damage from an attack on Target i, of scope r_i, defended at h_i is $D(h_i, r_i)$, meaning that such an attack leads to an expected loss of $V_{Si}D(h_i, r_i)$ for S and an expected gain of $V_{Ti}D(h_i, r_i)$ for T.

Assume that the function $D(h, r) : [0, \infty) \times [0, \infty) \to [0, 1]$ satisfies the following properties: $\frac{\partial D(h,r)}{\partial h} < 0$ and $\frac{\partial^2 D(h,r)}{\partial h^2} > 0$ for all $r > 0$ (increasing homeland security decreases the damage from an attack, but does so with diminishing returns); either $\frac{\partial D(h,r)}{\partial r} > 0$ and $\frac{\partial^2 D(h,r)}{\partial r^2} < 0$ for all $h > 0$ (increasing the scope of the attack increases the damage from the attack, but does so at a decreasing rate) or $\frac{\partial D(h,r)}{\partial r} = 0$ for all $h > 0$ (the damage from the attack does not depend at all upon the scope of the attack, so that in practice T is simply choosing where to attack); $D(0, r) = 1$ for all $r \geq 0$ (an attack on an undefended target imposes the full loss of V_{Si} on S and gives T the full gain of V_{Ti}, even an attack of

---

[12]A summary of all notations is provided in Table 1.

[13]As will be discussed below, because of the way in which the timing of the decisions is modeled, the analysis and results would not differ if we additionally allowed there to be incomplete information with respect to the values of V_{SA} and V_{SB}.

6

zero scope); and $D(h, r) \to 0$ as $h \to \infty$ for all $r \geq 0$.

Suppose S incurs a constant marginal cost of $C_S > 0$ from increasing homeland security at either target, so that choosing (h_A, h_B) costs $[h_A + h_B]C_S$. These costs are incurred regardless of T's subsequent choice. Thus, S has a payoff of: $\pi_{SA}(h_A, h_B, r_A) = -D(h_A, r_A)V_{SA} - [h_A + h_B]C_S$ if A is attacked with scope r_A; $\pi_{SB}(h_A, h_B, r_B) = -D(h_B, r_B)V_{SB} - [h_A + h_B]C_S$ if B is attacked with scope r_B; and $\pi_{S\emptyset}(h_A, h_B) = -[h_A + h_B]C_S$ if no attack is staged.

Suppose T incurs fixed costs of F_T from staging an attack, a value which depends upon neither which target is attacked nor the level of security at the targeted site.[14] Additionally, T incurs a marginal cost of $C_T > 0$ when increasing the scope of an attack at either target. Thus, T realizes a payoff of: $\pi_{TA}(h_A, r_A) = D(h_A, r_A)V_{TA} - r_A C_T - F_T$ from attacking A with r_A; $\pi_{TB}(h_B, r_B) = D(h_B, r_B)V_{TB} - r_B C_T - F_T$ from attacking B with r_B; and $\pi_{T\emptyset} = 0$ from staging no attack.

We focus on a single period, sequential decision making environment, in which S first chooses (h_A, h_B). After S makes this observable choice, T then chooses to attack A, B, or neither target (and, if staging an attack, chooses the scope of the attack). We generally consider an environment of incomplete information in which S chooses (h_A, h_B) without knowing the actual values of (V_{TA}, V_{TB}) but rather knowing only the distribution function, $G_T(v_{TA}, v_{TB})$, from which the type of T is drawn.[15]

The game is analyzed via backward induction.[16] When S chooses (h_A, h_B), he does so without knowing the actual (V_{TA}, V_{TB}) but with the correct recognition that the subsequent behavior of T depends critically on the chosen (h_A, h_B) and actual values of (V_{TA}, V_{TB}). Thus, S bases his choice on a computation of his expected payoff, supposing T will behave rationally (given his actual type) and with

---

[14]These fixed costs can be entirely avoided by staging no attack. Note, the assumption that T attacks at most one target can be justified by assuming that either: the fixed costs of staging a second attack are so large, that doing so is never desirable; or the marginal benefit to T of attacking and damaging a second target is so small, that doing so is never desirable.

[15]Complete information can be thought of as a special case, in which V_T consists of a single pair of values (V_{TA}, V_{TB}).

[16]We technically have a sequential game of incomplete information (with observable actions). The appropriate equilibrium concept is Perfect Bayesian Equilibrium. For brevity, we refer to simply the equilibrium throughout the discussion.

an appropriate weight assigned to each possible realization of type for T (according to the distribution $G_T(v_{TA}, v_{TB})$). Let h_A^* and h_B^* denote the equilibrium levels of homeland security for S. We aim to gain insight into how h_A^* and h_B^* relate to each other and depend upon the parameters of the model. When doing so, it is necessary to understand how this choice by S impacts the choice of target and scope by T.

2.1. Initial Insights on the Choice of Target/Scope by T

Focusing on the choice by T in the terminal node of the game, we can think of T first determining the ideal scope for attacking each possible target.[17] Let $r_i^*(h_i)$ denote the optimal scope of attack on Target i, and let $\pi_{Ti}^*(h_i) = \pi_{Ti}(h_i, r_i^*(h_i))$ denote the resulting payoff of T from staging such an attack. Recall, $\pi_{Ti}(h_i, r_i) = D(h_i, r_i)V_{Ti} - r_i C_T - F_T$. It follows that when $\frac{\partial D(h,r)}{\partial r} > 0$, $r_i^*(h_i)$ must (at an interior solution) satisfy the first order condition $\frac{\partial D(h_i, r_i)}{\partial r_i} V_{Ti} = C_T$. If instead $\frac{\partial D(h,r)}{\partial r} = 0$ for all $r \geq 0$, then the optimal choice of scope by T is simply the corner solution of $r_i^* = 0$.

After identifying $r_A^*(h_A)$ and $r_B^*(h_B)$ and determining $\pi_{TA}^*(h_A)$ and $\pi_{TB}^*(h_B)$, T will: attack A with scope $r_A^*(h_A)$ if $\pi_{TA}^*(h_A) \geq \max\{\pi_{TB}^*(h_B), 0\}$; attack B with scope $r_B^*(h_B)$ if $\pi_{TB}^*(h_B) \geq \max\{\pi_{TA}^*(h_A), 0\}$; and stage no attack if $\max\{\pi_{TA}^*(h_A), \pi_{TB}^*(h_B)\} \leq 0$.[18] It is instructive to see how $\pi_{Ti}^*(h_i) = D(h_i, r_i^*(h_i))V_{Ti} - r_i^*(h_i)C_T - F_T$ depends upon both h_i and V_{Ti}. Lemma 1 addresses the first of these issues.

Lemma 1. $\pi_{Ti}^*(h_i) = D(h_i, r_i^*(h_i))V_{Ti} - r_i^*(h_i)C_T - F_T$ *is decreasing in* h_i.

Proof. Consider an arbitrary $h_i = \tilde{h}$, for which $\tilde{r} = r_i^*(\tilde{h})$ and $\pi_{Ti}^*(\tilde{h}) = D(\tilde{h}, \tilde{r})V_{Ti} - \tilde{r}C_T - F_T$.

---

[17]Returning attention to the impact of the informational structure on the behavior of the players, recognize that: the payoff of T does not directly depend upon V_{SA} and V_{SB}; T is only called upon to act once, in the terminal node of the game; and T is able to observe (h_A, h_B) before deciding how to act. Thus, the behavior of T could never depend upon the beliefs that T holds regarding the values of V_{SA} and V_{SB}, implying that the analysis would proceed in an identical manner if we additionally allowed there to be incomplete information regarding the type of S.

[18]If any of these conditions hold with equality, then T has two or more options that yield his maximal payoff. In such instances, T is indifferent with respect to choosing among these best options.

Suppose that S had instead chosen $h_i = \hat{h}$ with $\hat{h} < \tilde{h}$. Since $\frac{\partial D(h,r)}{\partial h} < 0$, the fact that $\hat{h} < \tilde{h}$ implies $D(\hat{h}, \tilde{r}) > D(\tilde{h}, \tilde{r})$. Thus, even if T were to still choose $\tilde{r}$ when S chooses $\hat{h}$, T would realize a larger payoff: $\pi_{Ti}(\hat{h}, \tilde{r}) = D(\hat{h}, \tilde{r})V_{Ti} - \tilde{r}C_T - F_T > \pi_{Ti}^*(\tilde{h})$. Since $\tilde{r}$ need not be optimal when $h_i = \hat{h}$, T's payoff must be even greater for $\hat{r} = r_i^*(\hat{h})$. Thus, $\pi_{Ti}^*(\hat{h}) \geq \pi_{Ti}(\hat{h}, \tilde{r}) > \pi_{Ti}^*(\tilde{h})$. ∎

Lemma 1 implies that when S increases the level of security at a target, attacking the target becomes less desirable to T (even after accounting for the fact that T could alter his scope of attack). Inspection of $\pi_{Ti}(h_i, r_i) = D(h_i, r_i)V_{Ti} - r_iC_T - F_T$ provides the intuition for Lemma 1. Since $\frac{\partial D(h,r)}{\partial h} < 0$, increasing h_i decreases the value of $\pi_{Ti}(h_i, r_i)$ at each r_i. Thus, the maximum value of $\pi_{Ti}(h_i, r_i)$ (when maximized with respect to r_i) must be smaller when a larger value of h_i is chosen. Further, increasing h_i has no direct impact on the payoff of T from attacking the alternative target.

Let $\bar{h}_i$ denote the smallest value of h_i for which $\pi_{Ti}^*(h_i) \leq 0$.[19] If S chooses $h_i \geq \bar{h}_i$, then T prefers to stage no attack over attacking Target i. Thus: if $h_A \geq \bar{h}_A$ and $h_B \geq \bar{h}_B$, then T will not stage an attack; if $h_A < \bar{h}_A$ and $h_B \geq \bar{h}_B$, then T will attack A; and if $h_A \geq \bar{h}_A$ and $h_B < \bar{h}_B$, then T will attack B. Further, if $h_A < \bar{h}_A$ and $h_B < \bar{h}_B$, then attacking either target yields a positive payoff for T, in which case the choice of target is based upon a comparison of $\pi_{TA}^*(h_A)$ to $\pi_{TB}^*(h_B)$. That is, T prefers to attack A over B if and only if $\pi_{TA}^*(h_A) \geq \pi_{TB}^*(h_B)$, or equivalently

$$D(h_A, r_A^*(h_A))V_{TA} - r_A^*(h_A)C_T \geq D(h_B, r_B^*(h_B))V_{TB} - r_B^*(h_B)C_T. \qquad (1)$$

The combinations of (h_A, h_B) for which the condition in (1) holds with equality define a "locus of indifference in (h_A, h_B)-space" for T with respect to attacking A versus attacking B.[20] For $h_B \in [0, \bar{h}_B)$, let $h_A^L(h_B)$ denote the h_A for which (1) holds with equality. For each $h_B \in [0, \bar{h}_B)$, T will attack A if $h_A < h_A^L(h_B)$ and will attack B if $h_A > h_A^L(h_B)$. This target choice (dependent upon (h_A, h_B),

---

[19] Assuming $D(h, r) \to 0$ as $h \to \infty$ for all $r \geq 0$ guarantees that $\pi_{Ti}^*(h_i)$ becomes negative for arbitrarily large h_i.

[20] Since the expression on the left side of (1) is decreasing in h_A and does not depend upon h_B while the expression on the right side of (1) is decreasing in h_B and does not depend upon h_A, it follows that this locus must be upward sloping. Additionally, this locus must pass through $(h_A, h_B) = (\bar{h}_A, \bar{h}_B)$, since $\pi_{TA}^*(\bar{h}_A) = \pi_{TB}^*(\bar{h}_B) = 0$.

accounting for the actual values of (V_{TA}, V_{TB})) is illustrated in Figure 1. Inspection of Figure 1 reveals the following intuitive points – other factors fixed, T will: attack A when the h_A is relatively small; attack B when h_B is relatively small; and attack neither target when h_A and h_B are both relatively large.

We now address the preference of target for T when $h_A = h_B = h$, which allows us to determine whether this locus in (h_A, h_B)-space lies above or below the $45°$-line. Lemma 2 provides such a characterization, dependent upon the relation of $\alpha_T = \frac{V_{TB}}{V_{TA}}$ to a value of one.

Lemma 2. *If $\alpha_T = \frac{V_{TB}}{V_{TA}} = 1$, then $\pi^*_{TA}(h) = \pi^*_{TB}(h)$ for all $h > 0$. If $\alpha_T = \frac{V_{TB}}{V_{TA}} < 1$, then $\pi^*_{TA}(h) > \pi^*_{TB}(h)$ for all $h > 0$. If $\alpha_T = \frac{V_{TB}}{V_{TA}} > 1$, then $\pi^*_{TA}(h) < \pi^*_{TB}(h)$ for all $h > 0$.*

Proof. Regardless of the magnitude of $\alpha_T = \frac{V_{TB}}{V_{TA}}$, the desired insight depends upon a comparison of $\pi^*_{TA}(h)$ to $\pi^*_{TB}(h)$ (i.e., essentially the comparison made by condition (1) evaluated at $h_A = h_B = h$).

First consider $\alpha_T = 1$ (i.e., $V_{TA} = V_{TB}$). When $V_{TA} = V_{TB}$, $r^*_A(h) = r^*_B(h)$ (recall, when $\frac{\partial D(h,r)}{\partial r} > 0$ the optimal choice of scope must satisfy $\frac{\partial D(h,r_i)}{\partial r_i} = \frac{C_T}{V_{Ti}}$) and $D(h, r^*_A(h)) = D(h, r^*_B(h))$. Thus, $\pi^*_{TA}(h) = D(h, r^*_A(h))V_{TA} - r^*_A(h)C_T - F_T$ is equal to $\pi^*_{TB}(h) = D(h, r^*_B(h))V_{TB} - r^*_B(h)C_T - F_T$.

Next consider $\alpha_T < 1$ (i.e., $V_{TA} > V_{TB}$). Compare $\pi^*_{TB}(h) = \pi_{TB}(h, r^*_B(h))$ to $\pi_{TA}(h, r^*_B(h))$. This latter expression is the payoff to T from attacking A with scope $r^*_B(h)$ (the optimal scope for attacking B). Note, $\pi_{TA}(h, r^*_B(h)) > \pi_{TB}(h, r^*_B(h))$ if and only if $D(h, r^*_B(h))V_{TA} - r^*_B(h)C_T > D(h, r^*_B(h))V_{TB} - r^*_B(h)C_T$ or equivalently if and only if $V_{TA} > V_{TB}$ (which is the condition defining this case). Since $\pi^*_{TA}(h) = \pi_{TA}(h, r^*_A(h))$ must be greater than $\pi_{TA}(h, r^*_B(h))$ (by the optimality of $r^*_A(h)$), it follows that $\pi^*_{TA}(h) > \pi^*_{TB}(h)$. The proof of the third portion of the lemma proceeds along identical lines and is omitted for brevity. ∎

By Lemma 2, if S defends the targets equally, T prefers to attack the target which T values more. For example, when $V_{TA} > V_{TB}$ (i.e., $\alpha_T < 1$), the locus lies strictly below the $45°$-line (as illustrated in Figure 1). Thus, for $h_A = h_B = h$, T realizes a greater payoff from attacking A than from attacking B.

Since the choice of (h_A, h_B) is generally made by S in an environment of incomplete information, it

is instructive to recognize how (for chosen levels of h_A and h_B) the choice of target by T depends upon the realized V_{TA} and V_{TB}. Toward this end, recognize that (as stated by Lemma 3) for arbitrary levels of security, $\pi_{Ti}^*(h_i) = D(h_i, r_i^*(h_i))V_{Ti} - r_i^*(h_i)C_T - F_T$ is increasing in V_{Ti}.

Lemma 3. $\pi_{Ti}^*(h_i) = D(h_i, r_i^*(h_i))V_{Ti} - r_i^*(h_i)C_T - F_T$ *is increasing in* V_{Ti}.

Proof. To recognize why this lemma is not self evident, recall that generally $r_i^*(h_i)$ (and thus $D(h_i, r_i^*)$) depends upon V_{Ti}. Fixing h_i, let $\tilde{V}_{Ti}$ denote an arbitrary realized value of V_{Ti}, for which $\tilde{r}_i^*$ is the corresponding optimal scope when attacking Target i. This realization of $\tilde{V}_{Ti}$ and choice of $\tilde{r}_i^*$ lead to $\tilde{\pi}_{Ti}^*(h_i) = \pi_{Ti}(h_i, \tilde{r}_i^*(h_i))$. Suppose that $\hat{V}_{Ti}$ had been realized, with $\hat{V}_{Ti} > \tilde{V}_{Ti}$. Let $\hat{r}_i^*$ denote the optimal scope of attack when $\hat{V}_{Ti}$ is realized, for which $\hat{\pi}_{Ti}^*(h_i) = \pi_{Ti}(h_i, \hat{r}_i^*(h_i))$ results.

By comparison, $D(h_i, \tilde{r}_i^*(h_i))\hat{V}_{Ti} - \tilde{r}_i^*(h_i)C_T - F_T$ is greater than $D(h_i, \tilde{r}_i^*(h_i))\tilde{V}_{Ti} - \tilde{r}_i^*(h_i)C_T - F_T = \tilde{\pi}_{Ti}^*(h_i)$ (i.e., the payoff of T is larger for the larger V_{Ti}, without altering the scope of attack). Since $\hat{r}_i^*$ is the optimal scope when $\hat{V}_{Ti}$ is realized, it follows that $\hat{\pi}_{Ti}^*(h_i) = D(h_i, \hat{r}_i^*(h_i))\hat{V}_{Ti} - \hat{r}_i^*(h_i)C_T - F_T$ must be even greater than $D(h_i, \tilde{r}_i^*(h_i))\hat{V}_{Ti} - \tilde{r}_i^*(h_i)C_T - F_T$, implying $\hat{\pi}_{Ti}^*(h_i) > \tilde{\pi}_{Ti}^*(h_i)$. ∎

Lemma 3 states that for any arbitrarily fixed value of h_i, the payoff for T from staging an attack of optimal scope on Target i is increasing in the realized value of V_{Ti}. Further note that a larger realization of V_{Ti} does not impact the payoff to T from attacking the alternate site whatsoever. Finally, recall that staging no attack and realizing a payoff of $\pi_{T\emptyset} = 0$ is always an option.

These observations imply that for any chosen (h_A, h_B) there exists a cutoff $\bar{V}_{TA}(h_A)$ such that $\pi_{TA}^*(h_A) \geq \pi_{T\emptyset} = 0$ if and only if $V_{TA} \geq \bar{V}_{TA}(h_A)$ and a cutoff $\bar{V}_{TB}(h_B)$ such that $\pi_{TB}^*(h_B) \geq \pi_{T\emptyset} = 0$ if and only if $V_{TB} \geq \bar{V}_{TB}(h_B)$. Thus, T will only ever stage no attack if his true valuations are $V_{TA} < \bar{V}_{TA}(h_A)$ and $V_{TB} < \bar{V}_{TB}(h_B)$. If either $V_{TA} \geq \bar{V}_{TA}(h_A)$ or $V_{TB} \geq \bar{V}_{TB}(h_B)$, then T will stage an attack. If both $V_{TA} \geq \bar{V}_{TA}(h_A)$ and $V_{TB} \geq \bar{V}_{TB}(h_B)$, then attacking either target gives T a positive payoff, in which case T bases the choice of target upon a comparison of $\pi_{TA}^*(h_A) = D(h_A, r_A^*(h_A))V_{TA} - r_A^*(h_A)C_T - F_T$ to $\pi_{TB}^*(h_B) = D(h_B, r_B^*(h_B))V_{TB} - r_B^*(h_B)C_T - F_T$.

Recognizing that this comparison is identical to that described by (1), we can construct a "locus of indifference in (V_{TA}, V_{TB})-space" consisting of $(V_{TA}, V_{TB}) \in [\bar{V}_{TA}(h_A), \infty) \times [\bar{V}_{TB}(h_B), \infty)$ for which T is indifferent between attacking A versus B.[21] Let $V_{TB}^L(V_{TA}, h_A, h_B)$ specify the value of V_{TB} along this locus as a function of the chosen security levels and realized V_{TA}. For an arbitrary $V_{TA} \geq \bar{V}_{TA}(h_A)$, T will attack A if $V_{TB} < V_{TB}^L(V_{TA}, h_A, h_B)$ and will attack B if $V_{TB} > V_{TB}^L(V_{TA}, h_A, h_B)$. Figure 2 illustrates how (for chosen values of (h_A, h_B)) the choice of target depends upon the realized (V_{TA}, V_{TB}). For (V_{TA}, V_{TB}): in Area $A2$, T will attack A; in Area $B2$, T will attack B; and in Area $N2$, T will not stage an attack. Inspection of Figure 2 reveals the following intuitive points – other factors fixed, T will: attack A when V_{TA} is relatively large; attack B when V_{TB} is relatively large; and attack neither target when V_{TA} and V_{TB} are both relatively small.

2.2. Initial Insights on the Choice of Homeland Security Levels by S

To see how changes in h_A or h_B alter the choice of target by T, we need to determine how such changes impact the boundaries in Figure 2. Consider an increase in h_A from $\hat{h}_A$ to $\tilde{h}_A$ with h_B fixed.[22] Increasing h_A decreases $\pi_{TA}^*(h_A)$ for any value of V_{TA} (Lemma 1). Since $\bar{V}_{TA}(\hat{h}_A)$ denotes the value of V_{TA} for which $\pi_{TA}^*(\hat{h}_A) = 0$, it follows that attacking A would yield a negative payoff for T if $V_{TA} = \bar{V}_{TA}(\hat{h}_A)$ and $h_A = \tilde{h}_A$. Since $\pi_{TA}^*(h_A)$ is increasing in V_{TA} (Lemma 3), in the face of $\tilde{h}_A$ a larger value of V_{TA} is required to realize $\pi_{TA}^*(\tilde{h}_A) = 0$. Thus, $\bar{V}_{TA}(h_A)$ is increasing in h_A.

Continue to consider an increase in h_A from $\hat{h}_A$ to $\tilde{h}_A$ with h_B fixed, but now focus on the locus of indifference between attacking A versus B (defined by $\pi_{TA}^*(h_A) = \pi_{TB}^*(h_B)$). An increase in h_A decreases $\pi_{TA}^*(h_A)$ (Lemma 1) but has no impact on $\pi_{TB}^*(h_B)$. For combinations of (V_{TA}, V_{TB}) along the initial locus, when facing $h_A = \tilde{h}_A$ instead of $h_A = \hat{h}_A$ the payoff of T is now greater from attacking B as opposed to A. Since $\pi_{TA}^*(h_A)$ is increasing in V_{TA} (Lemma 3), in the face of $\tilde{h}_A$ a larger value of V_{TA} is needed to realize $\pi_{TA}^*(\tilde{h}_A) = \pi_{TB}^*(h_B)$. Thus, $V_{TB}^L(V_{TA}, h_A, h_B)$ must be decreasing in h_A.

---

[21] Lemma 3 implies that this locus must be upward sloping in (V_{TA}, V_{TB})-space.

[22] Because of the symmetric nature of the problem, similar insights would follow for an increase in h_B with h_A fixed.

Linking this discussion to Figure 2 reveals the two ways that increasing h_A alters the target choice of T. As h_A is increased, $\bar{V}_{TA}(h_A)$ increases (i.e., the boundary between Area $A2$ and Area $N2$ moves to the right) and $V_{TB}^L(V_{TA}, \hat{h}_A, h_B)$ decreases (i.e., the boundary between Area $A2$ and Area $B2$ moves down or equivalently to the right). For chosen (h_A, h_B), T will attack A with probability $P_A(h_A, h_B)$, attack B with probability $P_B(h_A, h_B)$, and stage no attack with probability $P_\emptyset(h_A, h_B)$ (with each probability determined by integrating over all (V_{TA}, V_{TB}) in the relevant area of Figure 2, according to $G_T(v_{TA}, v_{TB})$).

Starting at (h_A, h_B) for which $(\bar{V}_{TA}(h_A), \bar{V}_{TB}(h_B)) \in V_T$, each area in Figure 2 would have a non-empty intersection with V_T, so each probability would be strictly positive. As h_A is increased, there are some $(V_{TA}, V_{TB}) \in V_T$ for which T would have initially attacked A but for which he now attacks B and there are some $(V_{TA}, V_{TB}) \in V_T$ for which T would have initially attacked A but for which he now stages no attack.[23] Thus, increasing h_A has not only the two desired effects of decreasing $P_A(h_A, h_B)$ and increasing $P_\emptyset(h_A, h_B)$, but also the undesired effect of increasing $P_B(h_A, h_B)$.

If we instead started at (h_A, h_B) for which $(\bar{V}_{TA}(h_A), \bar{V}_{TB}(h_B)) \notin V_T$, then one or more of the probabilities may initially equal zero, and further, an increase in h_A might not alter the probabilities. For example, suppose that for an initial (h_A, h_B) the set V_T lies strictly within the interior of Area $A2$ in Figure 2, so that $P_A(h_A, h_B) = 1$ and $P_\emptyset(h_A, h_B) = P_B(h_A, h_B) = 0$. For a sufficiently small increase in h_A the entire set V_T will still lie strictly within Area $A2$, so that we still have $P_A(h_A, h_B) = 1$ and $P_\emptyset(h_A, h_B) = P_B(h_A, h_B) = 0$. Over this range, increasing h_A does not alter the choice of target by T at all, and only has the impact of decreasing the expected damage when A is ultimately attacked.

Recall, if both sites are equally defended, then T prefers to attack the target which he values more highly (Lemma 2). In terms of Figure 2, this implies that for $h_A = h_B = h$, the locus of indifference in (V_{TA}, V_{TB})-space is simply a 45°-line. Further: if $h_A > h_B$, then $\bar{V}_{TA}(h_A) > \bar{V}_{TB}(h_B)$ and $V_{TB}^L(V_{TA}, h_A, h_B) < V_{TA}$ (i.e., the locus lies below the 45°-line); and if $h_A < h_B$, then $\bar{V}_{TA}(h_A) <$

---

[23]This relies on the assumption that V_T is a convex set, with $G_T(v_{TA}, v_{TB})$ placing strictly positive probability on all points in V_T.

$\bar{V}_{TB}(h_B)$ and $V_{TB}^L(V_{TA}, h_A, h_B) > V_{TA}$ (i.e., the locus lies above the 45°-line). Finally, the situation in which S leaves both targets undefended (i.e., $h_A = h_B = 0$) is simply a special case of $h_A = h_B = h$. Under the assumption that $D(0, r) = 1$ for all $r \geq 0$, following a choice of $h_i = 0$ by S, if T were to attack Target i he would do so by choosing $r_i = 0$, resulting in $D(0, 0) = 1$ and $\pi_{Ti}^*(0) = V_{Ti} - F_T$. This implies $\bar{V}_{Ti}(0) = F_T$. That is, attacking Target i defended at $h_i = 0$ is better than staging no attack, so long as T's value for Target i is simply above the fixed cost of staging an attack.

Proposition 1 provides insight on (h_A^*, h_B^*) for two special cases with respect to V_T.

Proposition 1. *If $V_{TA} \geq V_{TB}$ (i.e., $\alpha_T \leq 1$) for all $(V_{TA}, V_{TB}) \in V_T$, then $h_A^* \geq h_B^*$. If $V_{TA} \leq V_{TB}$ (i.e., $\alpha_T \geq 1$) for all $(V_{TA}, V_{TB}) \in V_T$, then $h_A^* \leq h_B^*$.*

Proof. Suppose $V_{TA} \geq V_{TB}$ for all $(V_{TA}, V_{TB}) \in V_T$ (i.e., V_T consists of points on or below the 45°-line in (V_{TA}, V_{TB})-space). If S chooses $h_B > h_A$, then the boundary between Areas $A2$ and $B2$ in Figure 2 lies above the 45°-line. Thus, $P_B(h_A, h_B) = 0$. From here, S could decrease h_B to $h_A + \varepsilon$ without altering the choice of target or scope by T for any $(V_{TA}, V_{TB}) \in V_T$. However, this decrease in h_B directly decreases security costs for S (recall, costs of $[h_A + h_B]C_S$ are incurred by S regardless of the subsequent choice by T) and thereby directly increases the payoff of S. Since this is true starting at any $h_B > h_A$, it follows that $h_B^* \leq h_A^*$. A similar argument proves the second part of the Proposition. ∎

By Proposition 1, whenever S knows which target is more highly valued by T, S will choose a (weakly) higher level of security at the target which T values more highly. This begins to reveal the strong degree to which S's optimal choice depends upon the target valuations of T. This should not be entirely surprising, since in any one-shot, sequential move game, the initial choice of Player 1 depends greatly on the subsequent behavior of Player 2 (behavior which is based on the payoffs of Player 2).

The results of Proposition 1 are driven by the fact that S will only ever devote security resources to a target if there is a benefit from doing so. Increasing h_i can have the potential benefits of decreasing the expected damage if Target i is attacked and decreasing the probability that Target i is attacked. For

example, consider increasing h_A, starting at a situation in which initially the entire set V_T is in Area $B2$ in Figure 2. Since $P_A = 0$ to start, there is no benefit from increasing h_A: since A is never attacked, decreasing the expected damage if A is attacked has no benefit; and since the probability of A being attacked is zero, there is no benefit of decreasing the probability with which A is attacked.

3. Complete Information

Suppose S knows the exact values of V_{TA} and V_{TB} before choosing (h_A, h_B). Recall that such an instance of complete information corresponds to a situation in which the entire set V_T consists of a single point (V_{TA}, V_{TB}), which is clearly covered by Proposition 1. Thus, for the case of complete information: if $V_{TA} > V_{TB}$, then $h_A^* \geq h_B^*$; and if $V_{TB} > V_{TA}$, then $h_B^* \geq h_A^*$. When attempting to further analyze S's choice in this setting, it is best to examine Figure 1. When S knows the values of (V_{TA}, V_{TB}), S is in the position of choosing which target is ultimately attacked (since S can infer with certainty where an attack will be staged for any (h_A, h_B)).

If S chooses (h_A, h_B) along the locus of indifference, then (by construction) T is indifferent between attacking A versus B. However, following such a choice of security levels, S may have a strict preference regarding which target is attacked. In order to ease the analysis, assume that following a choice of (h_A, h_B) along this locus, T chooses to attack the target which S prefers to be attacked.[24]

From Figure 1, we see that in the case of complete information (h_A^*, h_B^*) must be at one of the boundaries of the identified regions. For example, Area $N1$ depicts the (h_A, h_B) for which T attacks neither target. But since increased homeland security is costly for S, a choice of $(\bar{h}_A, \bar{h}_B)$ gives S a greater payoff than any other (h_A, h_B) in Area $N1$. Similarly, any (h_A, h_B) strictly within the interior of Area $B1$ cannot be best, since S could increase his payoff by decreasing h_A; any (h_A, h_B) strictly within

---

[24]If this were not assumed, then S would have to choose a slightly different level of security at one of the sites to give T a strict preference for attacking the site that gives S the larger payoff. By the continuity of the payoff functions the payoff of each player would be essentially equal to what results by assuming that when indifferent T stages the attack where S desires.

the interior of Area $A1$ cannot be best, since S could increase his payoff by decreasing h_B. The search for (h_A^*, h_B^*) is significantly narrowed by these insights. When $V_{TA} > V_{TB}$ (as illustrated in Figure 1), the optimal (h_A, h_B) must be either: along the horizontal axis between the origin and the point $(h_A^L(0), 0)$ (resulting in an attack on A); on the "locus of indifference in (h_A, h_B)-space" (leading to an attack on the target which is preferred by S); or $(\bar{h}_A, \bar{h}_B)$ (resulting in no attack).[25] From here the choice by S can be analyzed as a choice of $h_A \in [0, \bar{h}_A]$, under the constraint of h_B being determined as just described.

Consider gradual increases in h_A from $h_A = 0$ up to $h_A = \bar{h}_A$. As h_A is increased between $h_A = 0$ and $h_A = h_A^L(0)$ (with $h_B = 0$) an attack is ultimately staged on A and S realizes $\pi_S = -D(h_A, r_A^*(h_A))V_{SA} - h_A C_S$. The only benefit of increasing h_A over this range is decreasing the expected damage from an attack on A, while the marginal cost of increasing h_A is simply C_S. Let h_A^{low} denote the level of $h_A \in [0, h_A^L(0))$ which maximizes π_S, and let π_S^{low} denote the resulting payoff of S.

For $h_A \in [h_A^L(0), \bar{h}_A)$, $\pi_S = -\min\left\{D\left(h_A, r_A^*(h_A)\right)V_{SA}, D\left(h_B^L(h_A), r_B^*(h_B^L(h_A))\right)V_{SB}\right\} - [h_A + h_B^L(h_A)]C_S$ (where the first term reflects the assumption that the attack is staged on the target yielding the greater payoff for S). As h_A is increased from below $h_A^L(0)$ to above $h_A^L(0)$ the costs (and potentially the benefits) to S change. First, the effective marginal cost of increasing h_A becomes larger when $h_A > h_A^L(0)$, because now S must choose $h_B = h_B^L(h_A)$ (i.e., a positive and increasingly larger level of h_B) to remain on the locus as h_A is increased. Further, increasing h_A from below $h_A^L(0)$ to above $h_A^L(0)$ induces T to attack B as opposed to A if and only if an attack on B leads to a larger payoff for S than does an attack on A. As a result, if S prefers that the attack be staged on B for combinations of (h_A, h_B) along the locus, then π_S is characterized by a discrete increase in value at $h_A = h_A^L(0)$ (if instead S prefers that the attack be staged on A for combinations of (h_A, h_B) along the locus, then π_S does not have such a discontinuity at $h_A = h_A^L(0)$, but rather has only a change in the value of its slope,

---

[25]For $V_{TA} < V_{TB}$, the locus in Figure 1 lies above the $45°$-line (and has a positive vertical intercept of $h_B^L(0)$). In such cases, (h_A^*, h_B^*) would be either: along the vertical axis between the origin and $(0, h_B^L(0))$; on the locus of indifference; or $(\bar{h}_A, \bar{h}_B)$.

reflecting the effective increase in marginal cost above this level of h_A). Let h_A^{mid} denote the level of $h_A \in [h_A^L(0), \bar{h}_A)$ which maximizes π_S, and let π_S^{mid} denote the resulting payoff of S.

When analyzing S's choice under complete information, it is critical to determine the preference of S over whether an attack is staged on A versus B for points along the locus in Figure 1. Condition (1), which defines this locus, can be expressed as:

$$\frac{D(h_A, r_A^*(h_A))}{D(h_B, r_B^*(h_B))} = \frac{V_{TB}}{V_{TA}} + \frac{C_T[r_A^*(h_A) - r_B^*(h_B)]}{V_{TA}D(h_B, r_B^*(h_B))}. \tag{2}$$

For chosen (h_A, h_B), we have $\pi_{SA} > \pi_{SB}$ if and only if $D(h_B, r_B^*(h_B))V_{SB} > D(h_A, r_A^*(h_A))V_{SA}$ or equivalently $\frac{V_{SB}}{V_{SA}} > \frac{D(h_A, r_A^*(h_A))}{D(h_B, r_B^*(h_B))}$. Imposing (2) (and expressing $\frac{V_{SB}}{V_{SA}} = \alpha_S$ and $\frac{V_{TB}}{V_{TA}} = \alpha_T$), it follows that along this locus $\pi_{SA} > \pi_{SB}$ if and only if

$$\alpha_S > \alpha_T + \frac{C_T[r_A^*(h_A) - r_B^*(h_B)]}{V_{TA}D(h_B, r_B^*(h_B))}. \tag{3}$$

Consider $\frac{\partial D(h,r)}{\partial r} = 0$, in which case $r_A^*(h) = r_B^*(h) = 0$ for all $h \geq 0$. Thus, (3) reduces to $\alpha_S > \alpha_T$. This implies that on this locus upon which T is indifferent regarding his choice of target, S prefers to have an attack staged on A over B if and only if $\alpha_S > \alpha_T$. That is, for $(h_A, h_B) \in (0, \bar{h}_A) \times (0, \bar{h}_B)$, S would: choose (h_A, h_B) for which A will be attacked, if and only if the state's relative valuation for B is greater than the terrorist's relative valuation for B (i.e., $\alpha_S > \alpha_T$); and choose (h_A, h_B) for which B will be attacked, if and only if the state's relative valuation for B is less than the terrorist's relative valuation for B (i.e., $\alpha_S < \alpha_T$). Since $\alpha_S \leq 1$ by assumption, it follows that whenever $\alpha_T > 1$ (i.e., $V_{TB} > V_{TA}$), S would never choose security levels for which A is ultimately attacked. If instead $\alpha_T < 1$ (i.e., $V_{TB} < V_{TA}$), then choosing (h_A, h_B) along the locus, S may prefer to have the attack staged on A (if $\alpha_S > \alpha_T$) or may prefer to have the attack staged on B (if $\alpha_S < \alpha_T$).[26]

---

[26]If instead $\frac{\partial D(h,r)}{\partial r} > 0$, then the final term in (3) will generally not equal zero, in which case the identification of the preferred target of S along the locus is not determined by a simple comparison of α_S and α_T. For example, it is straightforward to show that for $D(h, r) = \frac{r}{r+h}$ the condition corresponding to (3) is $\alpha_S > \sqrt{\alpha_T}$, so that in this case for points along the locus S prefers an attack to be staged on A over B if and only if $\alpha_S > \sqrt{\alpha_T}$.

To see why the preference of S regarding the target choice by T along the locus reduces to a simple comparison of α_S to α_T when $\frac{\partial D(h,r)}{\partial r} = 0$, recall that what we are trying to determine is whether S prefers an attack to be staged on A or B, focusing on combinations of (h_A, h_B) for which T is indifferent regarding his choice of target. Since S must always incur the security costs at each target, the comparison of π_{SA} to π_{SB} at any (h_A, h_B) reduces to a comparison of the expected loss from an attack staged at each site (i.e., a comparison of $V_{SA}D(h_A, r_A)$ to $V_{SB}D(h_B, r_B)$). When $\frac{\partial D(h,r)}{\partial r} = 0$, $r_i^* = 0$ (for which $C_T r_i^* = 0$) at each target. Since T incurs identical costs of simply F_T from attacking A or B, the comparison of π_{TA} to π_{TB} reduces to a comparison of the expected gain from an attack staged at each site (i.e., a comparison of $V_{TA}D(h_A, r_A)$ to $V_{TB}D(h_B, r_B)$). It follows that the relevant comparison reduces to a comparison of α_S to α_T. If instead $\frac{\partial D(h,r)}{\partial r} \neq 0$, then $C_T r_A^*$ need not equal $C_T r_B^*$, so that the comparison of π_{TA} to π_{TB} no longer reduces to a comparison of $V_{TA}D(h_A, r_A)$ to $V_{TB}D(h_B, r_B)$ (but rather reduces to a comparison based upon the condition specified by (2)).

Returning to the choice of h_A, recognize that at the point at which h_A is increased to $\bar{h}_A$ (along with h_B increased to $\bar{h}_B$), π_S is characterized by a discrete increase in value to $\pi_S^{high} = -[\bar{h}_A + \bar{h}_B]C_S$. This change in π_S is certain to be an increase, since as h_A and h_B are increased from $h_A = \bar{h}_A - \varepsilon$ and $h_B = h_B^L(\bar{h}_A - \varepsilon)$ to $h_A = \bar{h}_A$ and $h_B = h_B^L(\bar{h}_A) = \bar{h}_B$ security costs increase by only a small, continuous amount, while the expected damage from an attack decreases from a strictly positive level down to zero.

In practice, (h_A^*, h_B^*) can be identified by comparing π_S^{low}, π_S^{mid}, and π_S^{high}. It is possible for the optimal choice of S to be characterized by $h_A^* < h_A^L(0)$, $h_A^* \in [h_A^L(0), \bar{h}_A)$, or $h_A^* = \bar{h}_A$. Further, for $h_A^* \in [h_A^L(0), \bar{h}_A)$ the equilibrium could involve an attack on A or an attack on B. The potential for these different outcomes to arise can be illustrated by considering some examples with $D(h, r) = \frac{1}{1+h}$ (for which the locus in Figure 1 is defined by $h_A^L(h_B) = \frac{1}{\alpha_T}(1 + h_B) - 1$, implying $h_A^L(0) = \frac{1-\alpha_T}{\alpha_T}$ and $(\bar{h}_A, \bar{h}_B) = \left(\frac{V_{TA}-F_T}{F_T}, \frac{V_{TB}-F_T}{F_T}\right)$). In each example, assume $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, and $V_{TA} = 1$.

Example 1: $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, $V_{TA} = 1$, $V_{TB} = \frac{1}{4}$, and $F_T = \frac{1}{5}$. It can be shown that

$h_A^* = h_A^{low} = \frac{3}{2}$ and $h_B^* = 0$, resulting in an attack on A and $\pi_S^* = \pi_S^{low} = -2\sqrt{C_S V_{SA}} + C_S = \frac{-16}{25}$. As this example illustrates, S may want to leave a target completely undefended.

Example 2: $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, $V_{TA} = 1$, $V_{TB} = \frac{9}{16}$, and $F_T = \frac{1}{5}$. Example 2 differs from Example 1 in that V_{TB} is larger. In this case, $\alpha_T = \frac{9}{16} < \frac{4}{5} = \alpha_S$, so that along the locus in Figure 1, S prefers A to be attacked. It can be shown that $h_A^* = \sqrt{\frac{V_{SA}}{(1+\alpha_T)C_S}} - 1 = 1$ and $h_B^* = h_B^L(h_A^*) = (1 + h_A^*)\alpha_T - 1 = \frac{1}{8}$, resulting in an attack on A and $\pi_S^* = -2\sqrt{(1 + \alpha_T)C_S V_{SA}} + 2C_S = \frac{-17}{25}$.

Example 3: $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, $V_{TA} = 1$, $V_{TB} = \frac{9}{11}$, and $F_T = \frac{1}{5}$. Example 3 differs from Example 2 in that V_{TB} is even larger. In this case, $\alpha_T = \frac{9}{11} > \frac{4}{5} = \alpha_S$, so that along the locus in Figure 1, S prefers B to be attacked. It can be shown that $h_A^* = \sqrt{\frac{V_{SB}}{(1+\alpha_T)\alpha_T C_S}} - 1 = \frac{5}{6}$ and $h_B^* = h_B^L(h_A^*) = (1 + h_A^*)\alpha_T - 1 = \frac{1}{2}$, resulting in an attack on B and $\pi_S^* = -2\sqrt{(1 + \alpha_T)\frac{1}{\alpha_T}C_S V_{SB}} + 2C_S = \frac{-56}{75}$. This example illustrates how S may prefer to induce an attack on B as opposed to A.

Example 4: $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, $V_{TA} = 1$, $V_{TB} = \frac{9}{11}$, and $F_T = \frac{4}{11}$. Example 4 differs from Example 3 in that F_T is larger. When the fixed cost for staging an attack is equal to this larger value, $h_A^* = \bar{h}_A = \frac{V_{TA} - F_T}{F_T} = \frac{7}{4}$ and $h_B^* = \bar{h}_B = \frac{V_{TB} - F_T}{F_T} = \frac{7}{4}$, for which no attack is staged and $\pi_S^* = \frac{-12}{25}$. Thus, it clearly may be best for S to choose security levels sufficiently high so that no attack is staged.

4. Further Insights for Situations of Incomplete Information

Recall, under incomplete information S chooses (h_A, h_B) to maximize his expected payoff by integrating over all possible types of T (according to $G_T(v_{TA}, v_{TB})$), anticipating the subsequent behavior of T (for the chosen (h_A, h_B) and realized (V_{TA}, V_{TB})).

When $\frac{\partial D(h,r)}{\partial r} > 0$ the optimal scope of T must satisfy $\frac{\partial D(h_i,r_i)}{\partial r_i}V_{Ti} = C_T$. Assuming $\frac{\partial^2 D(h,r)}{\partial r^2} < 0$, T will choose a larger r_i when V_{Ti} is larger, leading to a larger $D(h_i, r_i)$. Thus, when $\frac{\partial D(h,r)}{\partial r} > 0$, the type of T impacts the payoff of S by not only influencing the target choice but by also altering $D(h_i, r_i)$. If instead $\frac{\partial D(h,r)}{\partial r} = 0$, then T essentially chooses only where to attack so that the realization of (V_{TA}, V_{TB}) impacts the payoff of S only by altering where the attack is staged and not by altering $D(h_i, r_i)$. Thus, for

$\frac{\partial D(h,r)}{\partial r} = 0$ the expected payoff of S is $\pi_S = -V_{SA}D(h_A, 0)P_A(h_A, h_B) - V_{SB}D(h_B, 0)P_B(h_A, h_B) -$

$[h_A + h_B] C_S$, with the values of $P_A(h_A, h_B)$ and $P_B(h_A, h_B)$ determined by integrating over the relevant regions illustrated in Figure 2 according to $G_T(v_{TA}, v_{TB})$. Therefore, for analytical convenience we assume $D(h, r) = \frac{1}{1+h}$, for which we further have (with respect to the boundaries in Figure 2):

$\bar{V}_{TA}(h_A) = (1 + h_A)F_T$, $\bar{V}_{TB}(h_B) = (1 + h_B)F_T$, and $V_{TB}^L(V_{TA}, h_A, h_B) = \frac{1+h_B}{1+h_A}V_{TA}$.

Further assume: V_{TA} is determined by a draw from $G_{TA}(v) = \left(\frac{v}{U_{TA}}\right)^2$ (with $0 < U_{TA} < \infty$); and α_T (and thus $V_{TB} = \alpha_T V_{TA}$) is determined independently of V_{TA} by a draw from $G_{T\alpha}(x) = \frac{x}{\hat{\alpha}}$ (with $\hat{\alpha} > 1$). The parameter $\hat{\alpha}$ is the upper bound on the possible realizations of α_T. By focusing on $\hat{\alpha} > 1$, we are allowing both $V_{TA} > V_{TB}$ and $V_{TB} \geq V_{TA}$ to occur with positive probability (since, in Proposition 1, we already obtained preliminary insights when S knew for certain which target was more highly valued by T). It follows that the set V_T is a triangle with corners in (V_{TA}, V_{TB})-space at $(0, 0)$, $(U_{TA}, 0)$, and $(U_{TA}, \hat{\alpha}U_{TA})$, and each $(V_{TA}, V_{TB}) \in V_T$ is equally likely.[27] These assumptions greatly ease the determination of $P_A(h_A, h_B)$ and $P_B(h_A, h_B)$, since each probability is simply equal to the ratio of the intersection of the relevant area in Figure 2 with V_T to the entire area of the triangle representing V_T.

A choice of $h_A = h_B = 0$ leads to $\left(\bar{V}_{TA}(0), \bar{V}_{TB}(0)\right) = (F_T, F_T)$, a point which lies within V_T. Further, S would never choose $h_A > \frac{U_{TA} - F_T}{F_T}$, since $h_A = \frac{U_{TA} - F_T}{F_T}$ makes the payoff of T from attacking A negative for all possible V_{TA}. Similarly, for any chosen $h_A \in \left[0, \frac{U_{TA} - F_T}{F_T}\right]$, S would never choose $h_B > \hat{\alpha}h_A + \hat{\alpha} - 1$ (since $h_B = \hat{\alpha}h_A + \hat{\alpha} - 1$ guarantees that B is never attacked). Thus, the chosen (h_A, h_B) must lead to $\left(\bar{V}_{TA}(h_A), \bar{V}_{TB}(h_B)\right) \in V_T$. As a result, $P_A(h_A, h_B) = \frac{1+h_B}{(1+h_A)\hat{\alpha}} - \frac{(1+h_A)(1+h_B)}{\hat{\alpha}}\left(\frac{F_T}{U_{TA}}\right)^2$ and

$P_B(h_A, h_B) = 1 - \frac{1+h_B}{(1+h_A)\hat{\alpha}} - \frac{(1+h_A)(1+h_B)}{\hat{\alpha}}\left(\frac{F_T}{U_{TA}}\right)^2 + \left(\frac{1+h_B}{\hat{\alpha}}\right)^2\left(\frac{F_T}{U_{TA}}\right)^2$.

Even when we restrict attention to $\frac{\partial D(h,r)}{\partial r} = 0$ and make convenient assumptions on $G_T(v_{TA}, v_{TB})$,

---

[27]Note, we are assuming V_{TA} is the larger of two independent draws of a random variable distributed $U[0, U_{TA}]$ and α is a single draw of a random variable distributed $U[0, \hat{\alpha}]$. Thus, if $\hat{\alpha} = 1$ we have the special case in which V_{TA} is the larger and V_{TB} is the smaller of two independent draws from a $U[0, U_{TA}]$ distribution.

the expression for π_S is still quite complex and a general analysis of the choice of (h_A, h_B) by S is intractable. However, for fixed parameter values, h_A^* and h_B^* can be determined numerically. Fixing $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, $U_{TA} = 1$, and $F_T = \frac{1}{5}$, such numerical results were obtained for various $\hat{\alpha}$ (the results are reported in Table 2). Parameter values were intentionally chosen with the aim of illustrating the wide range of qualitatively different outcomes that can arise.

The first column in Table 2 lists each $\hat{\alpha}$ considered, from smallest to largest. Reading down each remaining column reveals how the value of the corresponding endogenous variable depends upon $\hat{\alpha}$. Since a larger $\hat{\alpha}$ implies that T is more likely to place a relatively higher value on B, the results provide insight on how the equilibrium depends upon the distribution from which the type of T is drawn.

The columns labeled h_A^* and h_B^* report equilibrium levels of security. As $\hat{\alpha}$ increases, h_A^* decreases and h_B^* increases. Further, for relatively small $\hat{\alpha}$, S chooses a higher level of security at the target which he values more highly (i.e., $h_A^* > h_B^*$), while for relatively large $\hat{\alpha}$, S instead chooses a lower level of security at the target which he values more highly (i.e., $h_A^* < h_B^*$).[28] This observation reinforces the degree to which the choice by S is driven by the target valuations of T.

The next two columns report the corresponding values of $D_A^* = \frac{1}{1+h_A^*}$ and $D_B^* = \frac{1}{1+h_B^*}$, the expected damage at each target if the target is attacked. Since $D_i^* = \frac{1}{1+h_i^*}$ is decreasing in h_i^*, it directly follows from the reported h_A^* and h_B^* that: D_A^* is increasing in $\hat{\alpha}$; D_B^* is decreasing in $\hat{\alpha}$; $D_A^* < D_B^*$ for relatively small values of $\hat{\alpha}$; and $D_A^* > D_B^*$ for relatively large values of $\hat{\alpha}$.

The next three columns report the equilibrium probabilities that T stages an attack on A, stages an attack on B, and attacks neither target. As $\hat{\alpha}$ increases, we have the intuitive results that: P_A^* decreases; P_B^* increases; and P^* decreases. Further, for relatively small $\hat{\alpha}$, we have $P_A^* > P_B^*$.[29] That is, it may be best for S to choose security levels for which an attack is more likely to be staged on the target that S values more highly. This observation is anticipated, since it was already noted that in an environment of

---

[28]Additional calculations suggest that for the chosen parameter values $h_A^* > h_B^*$ for $\hat{\alpha} \leq 2.7630$.

[29]Additional calculations suggest that for the chosen parameter values $P_A^* > P_B^*$ for $\hat{\alpha} \leq 1.5690$.

complete information it may be best for S to choose security levels for which A is certain to be attacked.

Returning to the columns labeled h_A^* and h_B^*, it may be best for S to leave a target completely undefended. However, when S does so under incomplete information, he is leaving a target undefended even though it is attacked with positive probability. For example, for $\hat{\alpha} = 20$, $h_A^* = 0$ even though $P_A^* = .0941 > 0.$[30] For such large $\hat{\alpha}$, T is almost certain to place such a larger value on B than on A, that even a completely undefended Target A will rarely be attacked, to the point where it is best for S to allocate no defensive resources to this target that he values more highly.[31]

The final two columns in Table 2 report $P_i^* D_i^*$, equilibrium values of the expected damage at each target (i.e., the probability that a target is attacked multiplied by the expected damage to the target when the target is attacked). From these results, $P_A^* D_A^*$ appears to decrease and $P_B^* D_B^*$ appears to typically increase as $\hat{\alpha}$ increases.[32] Finally, $P_A^* D_A^* > P_B^* D_B^*$ may arise (see the first two rows of results). That is, it may be best for S to choose security levels for which the equilibrium expected damage is greater at the target which S values more highly.[33] Again, a parallel can be drawn to the case of complete information where the expected damage is greater at A whenever S chooses security levels for which A is attacked.

As illustrated above, a general analysis of the choice of (h_A, h_B) in an environment of incomplete information is intractable. However, insights on S's choice and the resulting equilibrium were obtained by a numerical analysis. The equilibrium could qualitatively differ in several dimensions, in that: either $h_A^* > h_B^*$ or $h_A^* < h_B^*$ is possible (which, with $\frac{\partial D(h,r)}{\partial r} = 0$, directly implies either $D_A^* < D_B^*$ or $D_A^* > D_B^*$ is possible); either $P_A^* > P_B^*$ or $P_A^* < P_B^*$ is possible; $h_A^* = 0$ even though $P_A^* > 0$ may be best; $h_B^* = 0$ even though $P_B^* > 0$ may be best; and either $P_A^* D_A^* > P_B^* D_B^*$ or $P_A^* D_A^* < P_B^* D_B^*$ is possible.

---

[30]Additional calculations suggest that for the chosen parameter values $h_A^* = 0$ for $\hat{\alpha} \geq 19.6611$.

[31]While not reported in Table 2, $h_B^* = 0$ could also be best. If instead $V_{SB} = \frac{1}{3}$ (along with $V_{SA} = 1$, $C_S = \frac{4}{25}$, $U_{TA} = 1$, and $F_T = \frac{1}{5}$), then $h_B^* = 0$ (even though $P_B^* > 0$) for $\hat{\alpha} \leq 1.6269$.

[32]More precisely, $P_B^* D_B^*$ appears to increase in $\hat{\alpha}$ until the point at which $h_A^* = 0$, beyond which further increases in $\hat{\alpha}$ appear to result in a decrease in $P_B^* D_B^*$.

[33]Additional calculations suggest that for the chosen parameter values $P_A^* D_A^* > P_B^* D_B^*$ for $\hat{\alpha} \leq 1.2968$.

5. Conclusion

A single period interaction between a state (S) and a terrorist organization (T) was analyzed, in which S initially chose security levels at two sites and T subsequently chose the target (and potentially scope) of attack. Increasing security at a site is costly for S, but has the potential benefits of decreasing the expected damage if the site is attacked and decreasing the probability with which the site is attacked. However, there is an undesired side effect of this second benefit: increasing security at one target may increase the probability that an alternative target is attacked. Thus, a single decisionmaker allocating resources across multiple targets must carefully account for such cross target effects.

If S knows which site is more highly valued by T, then S allocates more security to this site (even if this is not the site that S values more highly). Under complete information, S can correctly infer which target T will attack. A closer examination under complete information revealed the strong degree to which the preference by S over which site is attacked depends on the relative valuation of each player for each target: when the expected damage from an attack does not depend upon the scope of the attack (so that T is effectively choosing only where to attack), then S will set security levels for which T prefers to attack Target A (the target that S values more highly) over Target B if and only if the state's relative valuation for Target B is greater than the terrorist's relative valuation for Target B (i.e., if and only if $\alpha_S > \alpha_T$). Further, it was shown how S may want to leave one of the targets completely undefended and how S may want to choose security levels sufficiently high so that no attack is staged.

A setting of incomplete information was subsequently analyzed in greater detail. A numerical analysis revealed that in equilibrium it may be best: for S to leave a target completely undefended (even if it is attacked with positive probability); for S to choose security levels for which an attack is more likely to be staged on the target that he values more highly; or for S to choose security levels for which the expected damage from a terrorist attack is greater at the target that he values more highly.

As an avenue for further research, our model does not allow for the terrorists' valuation of alternative

targets to be affected by the government's allocation of defensive resources (on this point see also Bier, Oliveros, and Samuelson, 2007, p. 585). Terrorists may possibly value not only damaging a target but also destroying defensive resources at a target. One possible extension of our model would involve intertemporal optimization on the part of both the terrorists and the state, in which the choices in the current period influence target values in the future.

References:

Bernhardt, Dan and Polborn, Mattias K., "Non-Convexities and the Gains from Concealing Defenses from Committed Terrorists," *Economics Letters*, April 2010, *107*, 52-54.

Bier, Vicki, Oliveros, Santiago and Samuelson, Larry, "Choosing What to Protect: Strategic Defensive Allocation against an Unknown Attacker," *Journal of Public Economic Theory*, August 2007, *9*, 563-87.

Bilmes, Linda J. and Stiglitz, Joseph E., "The Economic Costs of the Iraq War: An Appraisal Three Years After the Beginning of the Conflict," NBER Working Paper No. 12054, National Bureau of Economic Research, Cambridge, MA, February 2006.

Brandt, Patrick T. and Sandler, Todd, "What Do Transnational Terrorists Target? Has It Changed? Are We Safer?" *Journal of Conflict Resolution*, April 2010, *54*, 214-36.

Cadigan, John and Schmitt, Pamela M., "Strategic Entry Deterrence and Terrorism: Theory and Experimental Evidence," *Public Choice*, April 2010, *143*, 3-22.

The Economist, March 8, 2008, p. 69; March 15, 2008, p. 98.

Enders, Walter and Sandler, Todd, "The Effectiveness of Antiterrorism Policies: A Vector-Autoregression-Intervention Analysis," *American Political Science Review*, December 1993, *87*, 829-44.

Enders, Walter and Sandler, Todd, "Terrorism: Theory and Applications," in Keith Hartley and Todd Sandler, eds., *Handbook of Defense Economics, Vol. 1*, Amsterdam: Elsevier, 1995, pp. 213-49.

Enders, Walter and Sandler, Todd, "What Do We Know About the Substitution Effect in Transnational Terrorism?" in Andrew Silke, ed., *Research on Terrorism: Trends, Achievements and Failures*, London: Frank Cass, 2004, pp. 119-37.

Enders, Walter, Sandler, Todd and Cauley, Jon, "Assessing the Impact of Terrorist-Thwarting Policies: An Intervention Time Series Approach," *Defence Economics*, December 1990, *2*, 1-18.

Farrow, Scott, "The Economics of Homeland Security Expenditures: Foundational Expected Cost-Effectiveness Approaches," *Contemporary Economic Policy*, January 2007, *25*, 14-26.

Keohane, Nathaniel O. and Zeckhauser, Richard J., "The Ecology of Terror Defense," *Journal of Risk and Uncertainty*, March 2003, *26*, 201-29.

Lee, Sanghoon, "Counterterrorism Alliance," Working Paper, School of Business Administration, Holy Family University, Philadelphia, PA, 2007.

Lowenberg, Anton D. and Mathews, Timothy, "Why Iraq?" *Defence and Peace Economics*, February 2008, *19*, 1-20.

Powell, Robert, "Defending Against Terrorist Attacks with Limited Resources," *American Political Science Review*, August 2007, *101*, 527-41.

Sandler, Todd and Arce M., Daniel G., "Terrorism: A Game-Theoretic Approach," in Todd Sandler and Keith Hartley, eds., *Handbook of Defense Economics, Vol. 2: Defense in a Globalized World*, Amsterdam: Elsevier North-Holland, 2007, pp. 775-813.

Sandler, Todd, Arce M., Daniel G. and Enders, Walter, "Transnational Terrorism," Copenhagen Consensus Challenge Paper, Copenhagen Business School, Frederiksberg, Denmark, March 6, 2008, http://www.copenhagenconsensus.com/Default.aspx?ID=968.

Sandler, Todd and Enders, Walter, "An Economic Perspective on Transnational Terrorism," *European Journal of Political Economy*, June 2004, *20*, 301-16.

Sandler, Todd and Lapan, Harvey E., "The Calculus of Dissent: An Analysis of Terrorists' Choice of Targets," *Synthèse*, 1988, *76*, 245-61.

Sandler, Todd and Siqueira, Kevin, "Global Terrorism: Deterrence Versus Pre-emption," *Canadian Journal of Economics*, November 2006, *39*, 1370-87.

Sandler, Todd and Siqueira, Kevin, "Games and Terrorism: Recent Developments," *Simulation and Gaming*, April 2009, *40*, 164-92.

Stiglitz, Joseph E. and Bilmes, Linda J., *The Three Trillion Dollar War: The True Cost of the Iraq Conflict*, New York: W. W. Norton, 2008.

Zhuang, Jun and Bier, Vicki M., "Balancing Terrorism and Natural Disasters - Defensive Strategy with Endogenous Attacker Effort," *Operations Research*, September-October 2007, *55*, 976-91.

Zhuang, Jun and Bier, Vicki M., "Secrecy and Deception at Equilibrium, with Applications to Anti-Terrorism Resource Allocation," *Defence and Peace Economics*, February 2011, *22*, 43-61.

Zhuang, Jun, Bier, Vicki M. and Alagoz, Oguzhan, "Modeling Secrecy and Deception in a Multiple-Period Attacker-Defender Signaling Game," *European Journal of Operational Research*, June 2010, *203*, 409-418.

Table 1: Summary of Notations.

Variable/Expression	Meaning
T	Terrorist Organization
S	Target State
A	Target of greater value to S
B	Target of lesser value to S
V_{Si}	Value that S places on Target i, for $i \in A, B$
V_{Ti}	Value that T places on Target i, for $i \in A, B$
V_T	Set of possible values for (V_{TA}, V_{TB})
$G_T(v_{TA}, v_{TB})$	Joint distribution function from which (V_{TA}, V_{TB}) is determined
h_i	Homeland security resources allocated by S to Target i, for $i \in A, B$
r_i	Resources allocated by T to attacking Target i (i.e., the scope of attack on Target i), for $i \in A, B$
$D(h, r)$	Expected damage from an attack of scope r on a target defended with security resources of h
C_S	Marginal cost to S of increasing homeland security
F_T	Fixed costs to T of staging an attack
C_T	Marginal cost to T of increasing scope of attack
$\pi_{Si}(h_A, h_B, r_i)$	Payoff to S if Target i is attacked, for $i \in A, B$
$\pi_{S\emptyset}(h_A, h_B)$	Payoff to S if neither target is attacked
$\pi_{Ti}(h_i, r_i)$	Payoff to T from attacking Target i, for $i \in A, B$
$\pi_{T\emptyset}$	Payoff to T from staging no attack
$r_i^*(h_i)$	Optimal scope of attack on Target i defended with h_i, for $i \in A, B$
$\pi_{Ti}^*(h_i)$	Payoff to T from staging an attack on Target i of the optimal scope $r_i^*(h_i)$, for $i \in A, B$
α_T	Relative valuation of T for Target B (defined as $\frac{V_{TB}}{V_{TA}}$)
α_S	Relative valuation of S for Target B (defined as $\frac{V_{SB}}{V_{SA}}$)
$\bar{h}_i$	Level of homeland security at Target i above which an attack is never staged on the target, for $i \in A, B$
$\bar{V}_{Ti}(h_i)$	Valuation of T for Target i below which an attack is not staged on the target, for $i \in A, B$
$P_i(h_A, h_B)$	Probability with which Target i is attacked for chosen (h_A, h_B), for $i \in A, B$
$P_\emptyset(h_A, h_B)$	Probability that no attack is staged for chosen (h_A, h_B)
π_S	Expected payoff to S for chosen (h_A, h_B), accounting for $G_T(v_{TA}, v_{TB})$ and subsequent behavior of T
h_i^*	Optimal level of homeland security at Target i (i.e., level which maximizes π_S), for $i \in A, B$
r_i^*	Equilibrium scope of attack if Target i is attacked, for $i \in A, B$
P_i^*	Equilibrium probability with which Target i is attacked, for $i \in A, B$
$P_\emptyset$	Equilibrium probability that no attack is staged
D_i^*	Equilibrium expected damage from an attack on Target i (if the target is attacked), for $i \in A, B$
U_{TA}	Upper bound on V_{TA} (within Section 4)
$\hat{\alpha}$	Upper bound on α_T (within Section 4)
$G_{TA}(v)$	Distribution function from which V_{TA} is determined (within Section 4)
$G_{T\alpha}(x)$	Distribution function from which α_T is determined (within Section 4)

Table 2: Numerical Results, with $V_{SA} = 1$, $V_{SB} = \frac{4}{5}$, $C_S = \frac{4}{25}$, $U_{TA} = 1$, and $F_T = \frac{1}{5}$.

$\hat{\alpha}$	h_A^*	h_B^*	D_A^*	D_B^*	P_A^*	P_B^*	$P_\emptyset^*$	$(P_A^*)(D_A^*)$	$(P_B^*)(D_B^*)$
1	1.1244	.4634	.4707	.6834	.5645	.2725	.1630	.2657	.1862
1.25	1.0431	.5336	.4895	.6520	.5002	.3594	.1403	.2448	.2344
1.5	.9697	.5827	.5077	.6318	.4525	.4257	.1217	.2297	.2690
1.75	.9051	.6195	.5249	.6175	.4152	.4780	.1068	.2180	.2952
2	.8483	.6486	.5410	.6066	.3850	.5203	.0947	.2083	.3156
2.5	.7530	.6925	.5704	.5908	.3387	.5847	.0766	.1932	.3455
3	.6760	.7249	.5967	.5798	.3045	.6316	.0639	.1817	.3662
4	.5572	.7709	.6422	.5647	.2567	.6960	.0473	.1649	.3930
5	.4684	.8032	.6810	.5546	.2244	.7384	.0372	.1528	.4095
10	.2142	.8892	.8236	.5293	.1464	.8367	.0169	.1206	.4429
15	.0817	.9315	.9244	.5177	.1135	.8761	.0105	.1049	.4536
20	0	.9608	1	.5100	.0941	.8984	.0075	.0941	.4582
25	0	1.0078	1	.4981	.0771	.9167	.0062	.0771	.4566

FIGURE 1. Choice of target by T, dependent upon h_A and h_B

h_B

Area N1
(neither
target attacked)

$\bar{h}_B$

Area A1
(attack staged
on Target A)

$h_A = h_A^L(h_B)$
or equivalently
$h_B = h_B^L(h_A)$

Area B1
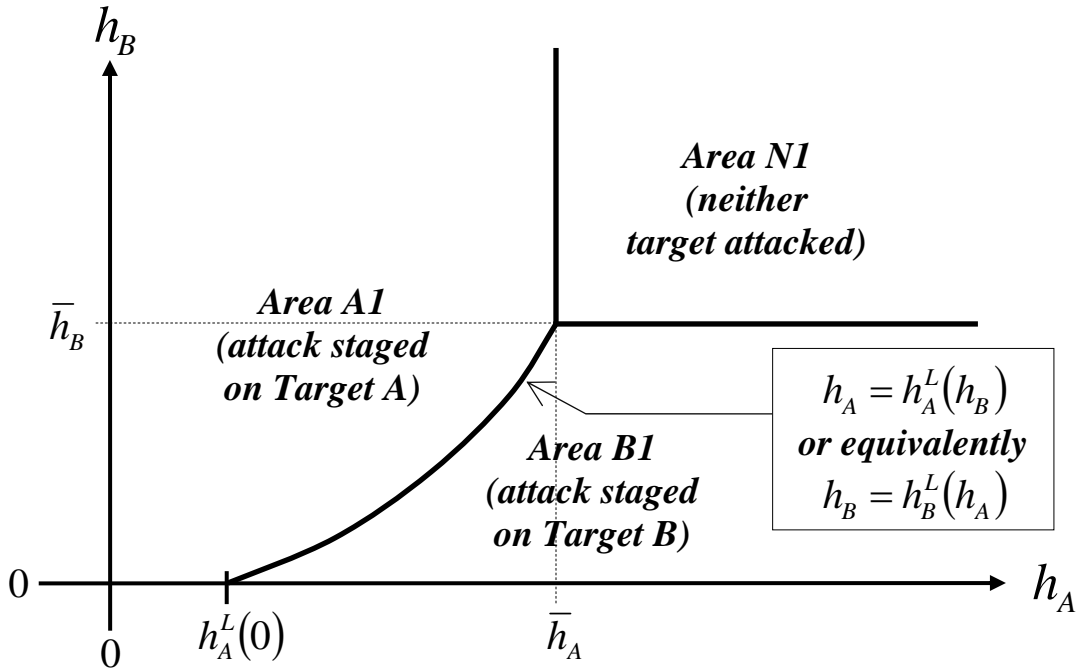(attack staged
on Target B)

0

h_A

$h_A^L(0)$

$\bar{h}_A$

0

FIGURE 2. Choice of target by T, dependent upon V_{TA} and V_{TB}

V_{TB}

Area B2
(attack staged
on Target B)

$V_{TB} = V_{TB}^L(V_{TA}, h_A, h_B)$
or equivalently
$V_{TA} = V_{TA}^L(V_{TB}, h_A, h_B)$

$\bar{V}_{TB}(h_B)$

Area A2
(attack staged
on Target A)

Area N2
(neither
target attacked)

0

V_{TA}

$\bar{V}_{TA}(h_A)$

0