Defending Against Terrorist Attacks with Limited Resources*

Robert Powell

August 2005

Department of Political Science
210 Barrows Hall, 1950
University of California
Berkeley, CA 94720-1950
RPowell@Berkeley.edu

* For helpful comments, criticisms and discussion, I thank Manas Baveja, Mariano-Florentino Cuéllar, Ernesto Dal Bó, Lynn Eden, James Fearon, William Kastenberg, Victoria Kingsley, Lance Kim, Michael May, Gerard Padro-i-Miquel, Per Peterson, Jacob Shapiro, and Dean Wilkening. I also gratefully acknowledge the support of the National Science Foundation.

Defending Against Terrorist Attacks with Limited Resources

Abstract

This paper develops a framework for analyzing a defender's allocation of scarce resources against a strategic adversary like a terrorist group in four settings: (i) a baseline case in which the sites the defender tries to guard are "independent" in that resources dedicated to protecting one site have no effect on any other site; (ii) if the defender can allocate its resources to border defense, counter-terrorist operations, or intelligence activities in addition to defending specific sites; (iii) if threats have strategic and non-strategic components (e.g., the threat to public health from bio-terror attacks and the natural outbreak of new diseases); and (iv) if the defender is unsure of the terrorists' preferred targets. The analysis characterizes the defender's optimal (equilibrium) allocations in these settings as well as an algorithm or approach to finding them. These characterizations and accompanying approaches provide a general way of thinking about the resource-allocation problem in these settings.

Defending Against Terrorist Attacks with Limited Resources

How should a state allocate limited resources to defend against a strategic adversary like a terrorist group? This is part of the challenge facing the United States in the aftermath of the attacks of September 11, 2001. Since then, the federal government has spent over $150 billion on homeland security and is now spending about $15 billion a year on protecting the country's critical infrastructure and key assets alone (OMB 2003, 2004, 2005). But the scope of the task is vast. Prior to 9/11, the Environmental Protection Agency reported on the risks posed by "worst-case" chemical accidents, finding that an accidental release at over 2,300 facilities would threaten between 10 and 100 thousand people, 600 sites threatened between 100 thousand and a million people, and 123 sites threatened more than a million (Belke 2000; GAO 2004, 16; Kocieniewski 2005). By 2004, the national infrastructure database contained 33,000 sites. Of these, the Department of Homeland Security had identified 1,700 as the most critical and planned to conduct vulnerability assessments of them (Moteff 2004). As Department of Homeland Security Secretary Michael Chertoff summarized the situation when reporting the results of his review of the Department's operations, policies, and structures, "Although we have substantial resources to provide security, these resources are not unlimited. Therefore, as a nation, we must make tough choices about how to invest finite human and financial capital to attain the optimal state of preparedness" (Chertoff 2005b, 3).

The strategic nature of the adversary complicates these choices. "One fact dominates all homeland security threat assessments," states the *National Strategy for Homeland Security*, "terrorists are strategic actors" (White House 2002, 7). They will try to strike where the defense is weak and the expected gains are high. Protecting one site may therefore shift the risk of attack to another site. "Increasing the security of a particular type of target, such as aircraft or buildings, makes it more likely that terrorists will seek a different target. Increasing countermeasures to a particular terrorist tactic, such as hijacking, makes it more likely that terrorists will favor a different tactic" (White House

2002, 29).[1]

This paper develops a framework for analyzing the problem of allocating limited resources against a strategic adversary in four settings.[2] In the first baseline case, the sites the defender is trying to protect are "independent." Resources dedicated to guarding one site have no direct effect on any other site. A very simple algorithm leads to the optimal allocation in these circumstances. A strategic adversary will attack the weakest link, i.e., the site that offers the highest expected payoff. The defender, therefore, should invest in hardening or strengthening this link. But the more the defender spends on this link, the stronger it becomes. Eventually it will be just as strong as the second weakest link. At this point, it does no good to harden one of these links without hardening the other. So, the defender must now invest in strengthening both links. Eventually, these links will be as strong as the third weakest link at which point the defender must invest in hardening all three of these links. The defender continues to allocate the rest of its budget in this way by spending on and hardening more and more sites so that all of the weak links are equally strong.

The second setting extends the baseline case to allow the defender to allocate resources to border defense, counter-terrorist operations, or intelligence in addition to spending on specific sites.[3] The key idea here is that investing in the first three activities reduces the probability of a successful attack on every site whereas hardening a specific site only affects that site. The third setting examines the allocation problem if the threat has a strategic and non-strategic component. For example, many of the steps that a state might want to take to protect public health against the outbreak or spread of a new

---

[1] For statistical evidence of these "substitution" effects see Enders and Sandler 1993, 2002, 2004; Enders, Sandler and Cauley, 1990; and Im, Cauley and Sandler, 1987.

[2] The present analysis focuses on understanding the first-best outcome, i.e., the optimal allocation for a single unitary defender. Other work examining the distorting effects of multiple actors with different interests is Bueno de Mesquita 2005, Lakdawalla and Talley 2005, Rosendorff and Sandler 2004, Sandler and Lapan 1988, and Sandler and Siqueira 2003. Sandler and Arce 2003 review game theoretic analyses of terrorism.

[3] In addition to "protecting critical infrastructure and key assets" and "defending against catastrophic threats," three of the other four critical missions specified in the *National Strategy for Homeland Security* are intelligence and warning, border and transportation security, and domestic counter-terrorism (White House 2002).

virus like SARS or a highly contagious avian flu are the same as it would want to take to defend against a bio-terror attack (Chyba 2001, 2002). But the former is non-strategic: viruses do not try to strike where a defender is weak and the expected gains are large. How should a defender allocate its resources if the threat has both a strategic and non-strategic component? Finally, the fourth setting studies the resource-allocation problem when the defender is unsure of the terrorists' priorities, i.e., about the way that it ranks potential targets. How, for example, should a defender allocate its resources if it is unsure whether its adversary is primarily interested in attacking political or economic targets or striking at nuclear or chemical facilities?

The analysis below characterizes the defender's optimal (equilibrium) allocations in each of these settings as well as an algorithm from finding them. These characterizations and accompanying algorithms provide a general way of thinking about the resource-allocation problem in these settings. As will be seen, we can often conceive of the defender in each of the three extensions of the baseline model as trying to equate the marginal gain from investing its resources in any one specific way, e.g., defending a particular site, with the marginal gain from investing them in any other way, e.g., in border defense. But the nature of these marginal gains can be quite subtle when facing a strategic adversary.

Strategic Versus Non-Strategic Terrorists

Although the *National Strategy for Homeland Security* emphasizes the dominating fact that terrorists are strategic, early spending decisions on homeland security and, especially, on critical infrastructure protection have been widely criticized for being dominated by pork-barrel politics. As the *9/11 Commission Report* delicately puts it, "In a free-for-all over money, it is understandable that representatives will work to protect the interests of their home states and districts. But this issue is too important for politics as usual to prevail. Resources must be allocated according to vulnerabilities" (2003, 396).[4] Secretary Chertoff underscored the same point shortly after taking office, "Risk management must

---

[4] More direct criticisms include Ervin 2005, Heyman and Carafano 2004, Kettl 2004, Khademain 2004, O'Beirne 2003, and editorials by the *Los Angeles Times* (2005), *New York Times* (2005a,b), and *Washington Post* (2005).

guide our decision making as we examine how we can best organize to prevent, respond and recover from an attack"(2005a, 2).

Risk management does take an adversary's intentions into account. But this approach typically does not treat adversaries as a fully strategic actors who try to counter a defender's decisions, who "modify their tactics and targets to exploit perceived vulnerabilities and observed strengths" (White House 2003, viii). And, not taking this into account can lead to a significant misallocation of defensive resources. This section illustrates the potential for misallocation with a simple model in which the defender only has two sites to protect. The next section extends this model to an arbitrary number of sites and develops a more general framework for analyzing the allocation problem.

In the two-site model, the defender has a total of R resources to allocate and suffers a loss of L_1 if site 1 is successfully attacked and L_2 if site 2 is successfully attacked. Take r_1 to be the resources dedicated to protecting the first site and r_2 to be the resources allocated to the second where $r_1 + r_2 = R$. Let $\delta_j(r_j)$ denote the conditional probability that an attack on site j succeeds if the defender allocates r_j to its defense. Finally, α is the probability of an attack on site 1 and $1 - \alpha$ is the probability of an attack on the site 2. All of this implies that the defender's expected loss to allocating r to site 1 and $R - r$ to site 2 is $L(r) = \alpha L_1 \delta_1(r) + (1 - \alpha) L_2 \delta_2(R - r)$. The first term on the right in this equality is the expected loss from an attack on the first site, $L_1 \delta_1(r)$, weighted by the probability of an attack on that site, and the second term is the expected loss from an attack on the second site weighted by the probability of that attack.[5]

The basic elements of this model are the same as those at the core of Secretary Chertoff's risk-management approach which centers on consequences, vulnerability, and threat.[6] Consequences are the outcomes of a successful attack and correspond in the model to L_1 and L_2. The likelihood of a successful attack, δ_1 and δ_2, reflect a site's

---

[5] It is convenient to work with expected losses which are simply the negative of expected gains or expected utilities. On utlity theory see Mas-Colell, Winston, and Green (1995).

[6] Chertoff (2005a,b). Moteff (2004) provides an overview of risk-based managment in relation to the Department of Homeland Security's critical infrastructure program. Guikema and Paté-Cornell's (2002) analysis of the allocation problem combines elements of risk analysis and game theory.

vulnerability. Indeed, the U.S. Coast Guard defines vulnerability precisely as "the conditional probability of success given that a threat scenario occurs," and vulnerability assessments can be seen in part as efforts to describe the conditional probabilities δ_1 and δ_2 (USCG 2003). Finally the probability of an attack α represents the threat.

Because resources are limited, investing more in one site means investing less in another. Figure 1(a) plots the defender's expected loss and the trade off it faces when the value of the first site relative to the second is 3/2 (i.e., $L_1 = 3$ and $L_2 = 2$), the probability of a successful attack on a site is $\delta_j(r_j) = 1 - \sqrt{r_j/R}$, and the total resources is $R = 10$. The relative odds of an attack on a site are also taken to be the same as that of the sites' relative value (i.e., $\alpha/(1-\alpha) = 3/2$ which means $\alpha = 3/5$ and $1 - \alpha = 2/5$).

The defender's optimal allocation $r^* \approx 8.3$ minimizes its expected loss. At this allocation, the marginal reduction in the defender's expected loss from spending more on protecting the first site is just offset by the marginal increase in the defender's expected loss from having to spend less on the second site. More formally, the optimal risk-management allocation r^* given the threat assessment α satisfies the first-order condition $dL(r^*)/dr = 0$ or, equivalently, r^* solves $\alpha L_1 \delta_1'(r^*) = (1-\alpha) L_2 \delta_2'(R - r^*)$.

But, equating the marginal costs and benefits in this way gives the wrong allocation against a strategic adversary. Figure 1(b) illustrates this misallocation. At r^*, a terrorist group can impose a higher expected loss on the defender by attacking the second site rather than the first. Therefore a strategic adversary seeking to impose the highest expected cost on the defender will attack the second site. The odds of an attack on the first site are not 3/2 and a spending decision based on these odds leads to a misallocation.

Put another way, the allocation r^* and attack probability $\alpha = 3/5$ are not a Nash equilibrium of the underlying game. In a Nash equilibrium each actor plays optimally against the other, i.e., each chooses a strategy which maximizes its payoff given the other actor's strategy. In this sense, r^* and α are only "half" an equilibrium: The defender is playing optimally against the attacker's strategy α, but the attacker is not playing optimally against the defender. The allocation r^* fails to reflect the strategic nature of the threat.

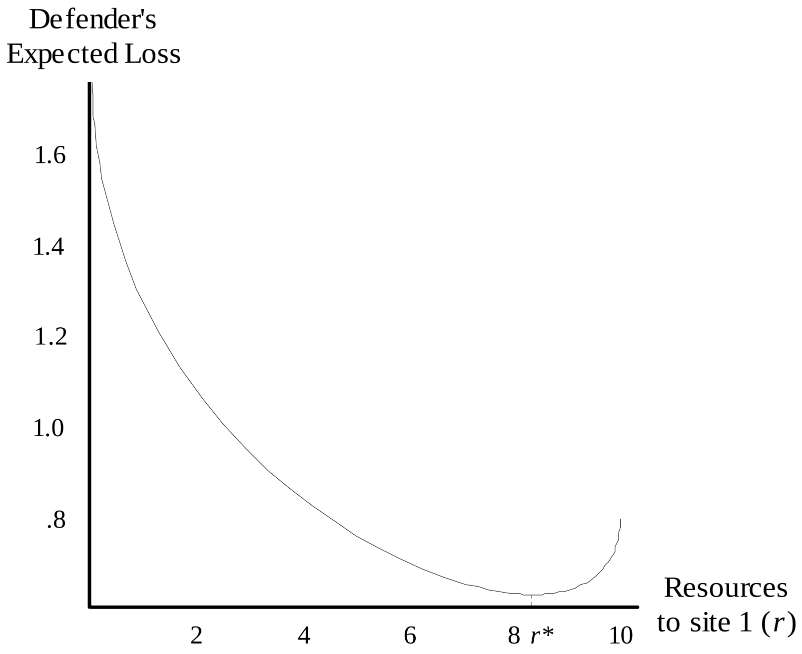

(a) Non-strategic terrorists.

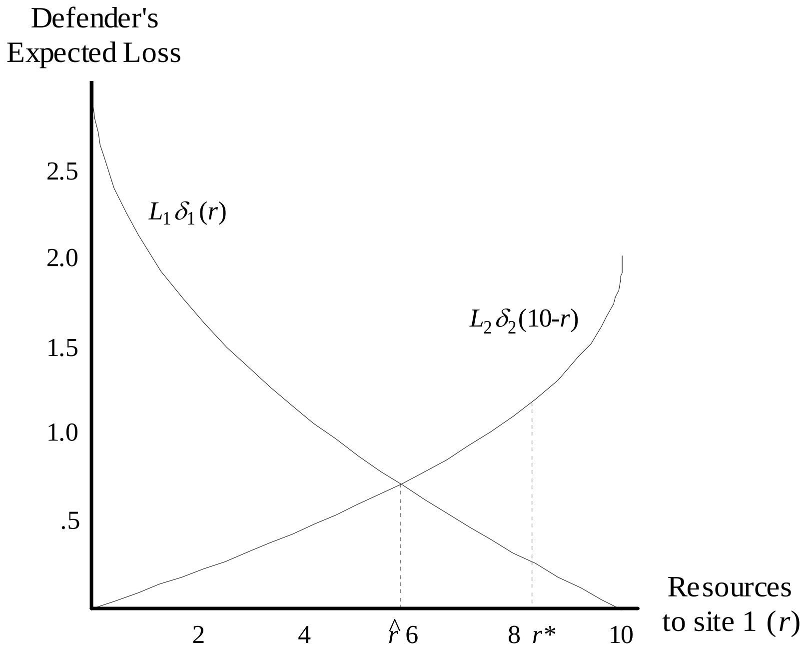

(b) Strategic terrorists.

Figure 1: Strategic versus non-strategic terrorists.

What is the optimal strategic allocation in this example? If the attacker is trying to impose the highest expected loss on the defender, the defender wants to minimize this maximum expected loss. Allocation $\widehat{r} \approx 5.8$ in Figure 1(b) does this. If the defender allocates less than this to protecting site 1, then that site will be the weakest link because that site offers the attacker its highest expected payoff. If the defender spends more than $\widehat{r}$ on site 1, say r^*, site 2 becomes the weakest link as striking that site now gives the attacker its highest expected payoff.

In brief, the risk-management allocation r^* does take the adversary's intentions into account (through α). But this approach does not treat the adversary as fully strategic: If the terrorists anticipated the allocation r^*, they would not want to behave in a way consistent with the threat assessment. Failing to take the strategic nature of the threat into account in this example leads to a significant misallocation of resources as the defender over spends on defending site 1 by more than 40% of the optimal allocation ($r^*/\widehat{r} \approx 1.42$).

A Basic Framework

This section presents a basic game-theoretic framework for thinking about and analyzing the resource-allocation problem. The solution to this baseline model also provides a simple algorithm for finding the unique optimal allocation of scarce resources against a strategic adversary. As will be seen, many sites may not receive any resources in the optimal allocation.

A defender must decide how to allocate R resources across N sites, and a terrorist group must decide which target to attack. These decisions are assumed to be made secretly although it turns that the results do not depend on this assumption and that the same results would obtain even if the defender had to move first and the attacker could observe the defender's allocation before deciding where to strike. A strategy for the defender is simply a resource allocation $(r_1, ..., r_N)$ such that $r_1 + \cdots + r_N \leq R$, and a strategy for the terrorist group is a set of probabilities $(\alpha_1, .., \alpha_N)$ where α_j is the probability that the terrorist group attacks site j and $\alpha_1 + \cdots + \alpha_N = 1$.

To specify the player's payoffs, suppose that the defender suffers a loss L_j^D if site j

is successfully attacked. The more resources the defender allocates to protecting a site, the lower the probability that an attack, if there is one, will succeed. In symbols, the conditional probability that an attack on site j succeeds is $\delta_j(r_j)$ with $\delta'_j(r_j) < 0$ as long as this site is imperfectly defended (i.e., as long as $\delta_j(r_j) > 0$).[7] Thus, the defender's expected loss if the terrorists attack j is $L_j^D \delta_j(r_j)$.[8] So, the defender's expected loss if it plays strategy r and the terrorist group plays α is $L^D(r, \alpha) \equiv \sum_{j=1}^N \alpha_j L_j^D \delta_j(r_j)$. As for the attacker's payoffs, let A_j be the attacker's gain from successfully striking site j. Then the terrorist group's expected gain is $A(r, \alpha) \equiv \sum_{j=1}^N \alpha_j A_j \delta_j(r_j)$.[9]

Finally resources are scarce in the sense that the defender cannot perfectly defend every target. That is, any allocation of R leaves one or more sites imperfectly defended. Formally, for any allocation of R, $\delta_j(r_j) > 0$ for some j.

A Nash equilibrium is a pair of strategies (r, α) such that each player's strategy maximizes its payoff against the other player's strategy. If, therefore, the defender's allocation is $(r_1, ..., r_N)$, the terrorist group will attack the a site offering the highest expected gain. That is, the attacker strikes a site j such that $A_j \delta_j(r_j) = \max\{A_1\delta_1(r_1),\ ...,\ A_N\delta_N(r_N)\}$.

A simple algorithm yields the defender's optimal allocation. Assume that the defender has not allocated any of its resources and that the sites are indexed so that if the defender

---

[7] In principle, it might be possible to defend some targets perfectly in which case $\delta_j(\overline{r}) = 0$ for some level of resources $\overline{r}$. If so, then devoting additional resources to this site has no effect and $\delta_j(r) = 0$ for all $r > \overline{r}$.

[8] This formulation assumes that an attack either succeeds or fails. We might assume more generally that an attack on j could lead to any loss in the range of $[0, L_j^D]$. Let $F_j(x|r_j)$ be the cummulative distribution of this loss if the defender allocates r_j to this site. Then the expected loss if j is attacked is $\lambda_j^D(r_j) = \int x dF_j(x|r_j)$. As long as $d\lambda_j^D(r)/dr < 0$ if $\lambda_j(r) > 0$, the analysis in this section carries through.

[9] In a zero-sum game, the players' payoffs are dimetrically opposed: anything that lowers the defender's expected loss (and therefore increases its payoff) must decrease the attacker's payoff. A zero-sum assumption may at first seem quite plausible in the present context of defending against terrorist attacks. But this assumption turns out to impose some strong restrictions on the parameters which seem less plausible. Specifically, the attacker and defender must not only share the same overall ranking of the sites (i.e., $L_i^D > L_j^D$ implies $A_i > A_j$ and vice-versa), but they also must put precisely the same relative value on the sites. In symbols the zero-sum assumption requires the much stronger assumption that $L_i^D / L_j^D = A_i / A_j$ for any two sites i and j.

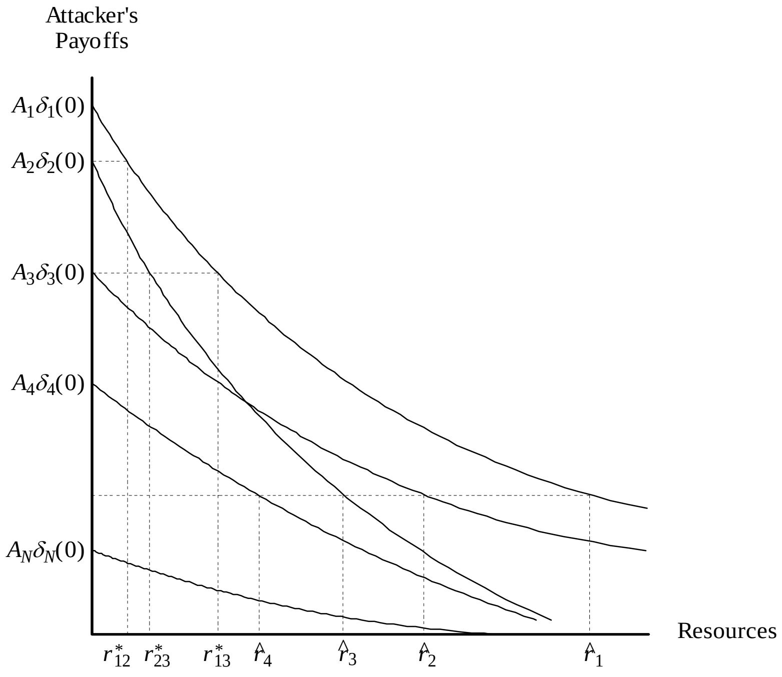

Figure 2: The optimal allocation.

allocates no resources to any site, the attacker's expected gain from striking site 1 is larger than its expected gain to striking site 2 which is larger than its expected gain to striking site 3 and so on. That is, $A_1\delta_1(0) > A_2\delta_2(0) > \cdots > A_N\delta_N(0)$.[10] Then, to minimize the maximum loss the terrorists can inflict, the defender should begin by allocating resources to the site or sites that the terrorists would attack. And, the site that brings the highest expected gain at the outset of this allocation process is site 1 because $A_1\delta_1(0) = \max\{A_1\delta_1(0), ..., A_N\delta_N(0)\}$. As the defender devotes more resources to defending this weakest link, the attacker's expected gain from striking this target declines until the payoff to hitting site 1 just equals the attacker's payoff to going after site 2. This occurs at r^*_{12} in Figure 2 where $A_1\delta_1(r^*_{12}) = A_2\delta_2(0) = \max\{A_1\delta_1(r^*_{12}), A_2\delta_2(0), ..., A_N\delta_N(0)\}$. At this point, the defender must begin to allocate resources to both sites 1 and 2 so that $A_1\delta_1(r_1) = A_2\delta_2(r_2)$. Otherwise the attacker would strike either site 1 or 2 depending on which offered the highest expected gain, and the resources

---

[10] The inequalities are assumed to be strict in order to simplify the exposition. Generalizing the analysis if some of the inequalities are weak is straighforward.

allocated to the other site would effectively have been wasted.

The defender continues allocating resources to sites 1 and 2 keeping the attacker's payoffs to striking these sites equal until the expected gain of striking site 3 just equals that of attacking 1 or 2, i.e., until $A_1\delta_1(r_{13}^*) = A_2\delta_2(r_{23}^*) = A_3\delta_3(0) = \max\{A_1\delta_1(r_{13}^*), A_2\delta_2(r_{23}^*), A_3\delta_3(0), ..., A_N\delta_N(0)\}$. Now the defender allocates resources to sites 1, 2, and 3 so as to keep $A_1\delta_1(r_1) = A_2\delta_2(r_2) = A_3\delta_3(r_3)$. The defender continues dividing its resources in this way reducing the attacker's maximum expected gain and having to spread its resources across more and more sites until it fully allocates its resources.

This algorithm provides a simple way of allocating R so as to stengthen the weakest links as much as possible by minimize the ceiling on the attacker's maximum expected gain $\max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}$. The resulting division of resources, $\widehat{r}$, is also the unique equilibrium allocation. That is, there exists an attack strategy $\widehat{\alpha}$ such that $(\widehat{r}, \widehat{\alpha})$ is an equilibrium and there are no other equilibria (r', α') with $r' \neq \widehat{r}$. The following proposition summarizes the results. (Proofs of the propositions are in the appendix):

Proposition 1: *The unique equilibrium allocation strengthens the weakest link as much as possible by minimizing the ceiling* $\max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}$, *and this optimal allocation is given by the algorithm above.*

To specify the attacker's equilibrium strategy $\widehat{\alpha}$ and thereby establish that $\widehat{r}$ is an equilibrium allocation, let the optimal allocation $\widehat{r}$ be denoted by $(\widehat{r}_1, ..., \widehat{r}_T, 0, , ..0)$ where T is the number of sites allocated resources, $\widehat{r}_1 + \cdots + \widehat{r}_T = R$, and $A_1\delta_1(\widehat{r}_1) = \cdots = A_T\delta_T(\widehat{r}_T) > A_{T+1}\delta_{T+1}(0)$.[11] Given this allocation, the attacker's payoff to striking any of the first T sites is the same and higher than its payoff to going after any other site. This implies that the probability that the attacker will strike any site $j > T$ is zero, i.e., $\widehat{\alpha}_j = 0$.

As for the probabilities of striking the first T sites, if $\widehat{r}$ is to be an equilibrium strategy, it must minimize the defender's expected loss $L^D(r, \widehat{\alpha})$ against the attacker's strategy $\widehat{\alpha}$. Thus, $\partial L^D/\partial r_j = \partial L^D/\partial r_i$ for all $i, j \leq T$ where $r_i > 0$ and $r_j > 0$. Consequently,

---

[11] In principle, $A_T\delta_T(\widehat{r}_T)$ might equal $A_{T+1}\delta_{T+1}(0)$ in which case the attacker's payoffs to stiking sites 1 through $T + 1$ would be the same. To ease the exposition, we assume the inequality to be strict which will be the case generically.

$\widehat{\alpha}_1 L_1^D \delta_1' = \widehat{\alpha}_j L_j^D \delta_j'$ for $j \leq T$. Solving for $\widehat{\alpha}_j$ in terms of $\widehat{\alpha}_1$, adding the $\widehat{\alpha}_j$'s for $j \leq T$, and using the fact that this sum equals one give $\widehat{\alpha}_j = \left(\sum_{n=1}^{T} L_j^D \delta_j' / L_n^D \delta_n' \right)^{-1}$.

In sum, a simple algorithm leads to the unique optimal allocation of resources against a strategic adversary in this basic framework where the sites are "independent" in that spending on one site has no effect on any other site. The optimal allocation in this baseline case by hardening the weakest links as much as possible by minimizing the ceiling $\max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}$.

Border Defense, Counter-Terrorist Operations, and Intelligence

In addition to defending individual sites, the United States is attempting to "harden" its borders by making it more difficult for terrorists to enter and operate in the country. The United States also engages in significant counter-terrorist activities throughout the world. And, good intelligence is seen to be critically important to disrupting terrorist attacks. Indeed, "intelligence and warning," "border and transportation security," and "domestic counter-terrorism" are three of the six critical mission areas specified in the *National Strategy for Homeland Security.*[12] This section examines the trade off between defending individual sites as in the basic framework above and allocating resources to border defense, counter terrorist operations, or intelligence activities.

The key to introducing this trade off into the baseline model is that investing in border defense, counter-terror operations, or intelligence lowers the probability of a successful attack on all of the sites whereas investing in defending some sites only helps protect those sites. These activities "tend to reduce the overall level of risk without having to know in advance what the targets are, while also complementing site defenses" (O'Hanlan *et. al.* 2003, 2). The central result is the that optimal allocation equates the marginal gain from investing in border defense with the marginal gain from site defense.[13] But as the analysis of the baseline case shows, spending on site defense must be spread across

---

[12] White House 2002. The United States is also devoting substantial resources to combatting terrorists abroad (OMB 2003).

[13] To ease the exposition, "border defense" will sometimes be used broadly to include counter-terrorist and intelligence activities as well.

multiple sites in order to keep the attacker's expected payoff to attacking any one of these sites equal to its expected payoff to striking any other site. Consequently, the marginal gain to site defense depends on the number of sites being defended and decreases as that number increases. Indeed, if the number of defended sites is quite large, the marginal gain to additional spending on site defense will be very small and is likely to be less than the marginal gain to border defense even if the latter is also quite small.

The equilibrium conditions also lead to two useful approximations expressed in terms of the easiest and hardest sites to defend. Because resources devoted to site defense must be spread across multiple threatened sites, say T, the marginal gain from site defense decreases like $1/T$. It follows that the defender prefers investing more in border defense, counter-terrorist operations, or intelligence activities than site defense if the marginal gain from the former is larger than $1/T$-th of the marginal gain from trying to harden the easiest site to protect.

If, alternatively, one site is much harder to defend than the others, then most additional spending on site defense will go toward defending this site (in order to keep the attacker's expected gain constant across the defended sites). Hence, the marginal gain from site defense approximates the marginal gain from defending this site. If this gain is less than the marginal gain from border defense, the defender prefers to allocate more to border defense.

To incorporate border defense, counter-terrorist activities, or intelligence in the basic framework, suppose that the defender now has to decide how to allocate resources R to N individual sites and to border defense. The probability that an attack on a site succeeds now depends on the resources allocated to that individual site and to border defense. More specifically, the conditional probability of a successful attack on j is the probability that the terrorist group penetrates the border (or evades the defender's counter-terrorist operations and intelligence efforts) times the probability of a successful attack on j given that the terrorists reach that site. In symbols, the probability that an attack on j succeeds is $\delta_j(r_j)\beta(b)$ where $\beta(b)$ is the probability that the terrorists penetrate the border if the defender allocates b to border defense. This probability decreases as more resources are spent on border defense (i.e., $\beta' < 0$).[14] As before, the terrorist group attacks sites offering the highest expected gain which are those j such that $A_j\delta_j(r_j)\beta(b) = \max\{A_1\delta_1(r_1)\beta(b), ..., A_N\delta_N(r_N)\beta(b)\}$ where $(r_1, ..., r_N, b)$ is defender's allocation.

The resource-allocation problem with border defense is closely related to the baseline case. The key to this relation as well as to the analysis of the threats with non-strategic components and uncertainty about the terrorist's priorities is the *optimal-resource-allocation path.* Suppose that the defender only had σ resources to allocate to site defense and wanted to minimize the attacker's payoff defined by $M_A(r|\sigma) \equiv \max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}$ where $r_1 + \cdots + r_N \leq \sigma$. The algorithm above defines the unique allocation $\widehat{r}(\sigma) = (\widehat{r}_1(\sigma), ..., \widehat{r}_N(\sigma))$ which minimizes $M_A(r|\sigma)$. Now let σ vary from zero to R. As σ varies, the allocation $\widehat{r}(\sigma)$ traces out a path P_B where $\widehat{r}(\sigma) = P_B(\sigma)$. By construction, each allocation $\widehat{r}(\sigma)$ along P_B optimally allocates the σ resources devoted to site defense.

The level of spending on site defense at any point along the path also implicitly defines the amount dedicated to border defense, namely, $R - \sigma$. For example, at the start of P_B (where $\sigma = 0$) the defender is allocating all of its resources to border defense and none of them to site defense. Conversely, the defender dedicates all of its resources to site defense and nothing to border defense at $\sigma = R$. Indeed, we can think of moving from $\sigma = 0$ to $\sigma = R$ along P_B as in effect shifting more and more resources out of border defense and optimally reallocating them to site defense.

The path P_B helps pin down the defender's optimal allocation in two ways. First, an equilibrium allocation must lie on the path. No matter how much the defender allocates to site defense, say σ', the defender should allocate those resources optimally. But that requirement is precisely what defines the point $P_B(\sigma')$.

Second, imagine starting at $\sigma = 0$ and moving toward $\sigma = R$ along P_B by shifting more and more resources from border defense to site defense. As the defender cuts back its spending on border defense to free up resources for site defense, the defender has to

---

[14] The present analysis makes no distinction between border defense, counter-terrorist operations, and intelligence efforts. The basic simplifying assumption is that each of these activities affects the probability of a successful attack in the same way.

spread this added investment in site defense across more and more sites so as to keep the attacker's payoffs to striking any one of these sites equal to that of attacking any other site. Having to distribute the resources it spends on site defense in this way reduces the marginal gain to site defense. An equilibrium allocation is a point on the path at which the defender's marginal loss from investing slightly less in border defense just offsets the marginal gain from spending slightly more on site defense. When this marginal condition does not hold the defender can profitably deviate from the allocation by shifting resources between border and site defense.

The fact that the defender must distribute its site defense resources across multiple sites suggests two useful approximations which can now be formalized. The first relates the defender's allocation decision to the marginal gain from protecting the easiest site to defend. Site j is easier to defend than k if investing more in defending j has a larger effect on the attacker's expected payoff to attacking j than investing in k has on the attacker's payoff to striking at k. That is, j is easier to defend than k if $A_k\delta'_k(r_k) \geq A_j\delta'_j(r_j)$, and e is the easiest site to defend (out of all the defended sites, i.e., for all sites for which $r_k > 0$) if $A_k\delta'_k(r_k) \geq A_e\delta'_e(r_e)$.[15] Then, the defender is certain to want to invest relatively more in border defense if the marginal gain to border defense is larger than $1/T$-th of the marginal gain from protecting e. This condition reduces to $\delta'_e(r_e)\beta(b)/T > \delta_e(r_e)\beta'(b)$ which depends only on the underlying technologies and not on the attacker's or defender's payoffs. (See the appendix for the derivation.)

The second approximation obtains if one of the defended sites is much harder to protect than the others. Suppose that the marginal effect of investing in protecting h is much less than in protecting other sites, $A_h\delta'_h \gg A_j\delta'_j$. Then any additional spending on site defense mostly has to go to h so as to keep the attacker's gain from striking any site the same across all defended sites. This implies that the marginal gain from site defense is approximately equal to the marginal gain from defending h, namely $A_h\delta'_h\beta$. The defender, therefore, will prefer to invest relatively more in border defense when doing so effects a larger reduction in the attacker's payoff than would investing in h, i.e., when

---

[15] Because $\delta'_j < 0$, larger marginal effects mean lower (negative) values of δ'_j.

$A_h \delta_h' \beta > A_h \delta_h \beta'$ or $\delta_h' \beta > \delta_h \beta'$.

Proposition 2 formalizes the marginal conditions defining the equilibrium allocation. Observe that the defender's expected loss at point $\widehat{r}(\sigma) = P_B(\sigma)$ given the attacker's strategy α can be written as $S_D(\sigma, \alpha)\beta(R-\sigma)$ where the site defense factor is defined by $S_D(\sigma, \alpha) = \sum_{j=1}^{N} \alpha_j L_j^D \delta_j(\widehat{r}_j(\sigma))$. Also let $T_A(r)$ be the set of targets offering the attacker its highest payoff given any allocation r. That is, $T_A(r)$ is the set of best replies for the attacker: $T_A(r) = \{j : A_j\delta_j(r_j) = \max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}\}$. Now define the attacker's "pseudo-equilibrium" strategy $\widehat{\alpha}$ against allocation $\widehat{r}$ to be a best response to $\widehat{r}$ that also rationalizes the defender's relative spending on the sites in $T_A(\widehat{r})$. As shown above, $\widehat{\alpha}_j = 0$ for all $j \notin T_A(\widehat{r})$ and $\widehat{\alpha}_j = \left(\sum_{n \in T_A(\widehat{r}(\sigma))} L_j^D \delta_j' / L_n^D \delta_n'\right)^{-1}$ for $j \in T_A(\widehat{r})$. Then the marginal condition above holds at the critical points of $S_D(\sigma, \widehat{\alpha})\beta(R-\sigma)$, i.e., where $\partial S_D(\sigma, \widehat{\alpha})/\partial\sigma \beta(R-\sigma) = S_D(\sigma, \widehat{\alpha})\beta'(R-\sigma)$. The expression on the left is the marginal reduction in the defender's expected loss from spending slightly more on site defense whereas the right side is the marginal reduction due to spending slightly more on border defense. Then:

Proposition 2: *Equilibrium allocations always exist, must lie on the path P_B, and be an endpoint or (almost everywhere) satisfy the marginal condition $\partial S_D(\sigma, \widehat{\alpha})/\partial\sigma \beta(R-\sigma) = S_D(\sigma, \widehat{\alpha})\beta'(R-\sigma)$ which is equivalent to:*[16]

$$\frac{\beta(R-\sigma)}{\beta'(R-\sigma)} = \sum_{n \in T_A(\widehat{r})} \frac{\delta_n(\widehat{r}_n)}{\delta_n'(\widehat{r}_n)}$$

There is a unique equilibrium allocation if the elasticities of site defense and border defense

---

[16] The partial derivative $\partial S_D(\sigma, \widehat{\alpha})/\partial\sigma$ is well defined everywhere except possibly at the finitely many points along P_B where sites enter the target set, i.e., at σ such that $r_j(\sigma) = 0$ for some $j \in T_A(r(\sigma))$. As the proof of the propositon in the appendix shows, the partials from the right and left are well defined but generally not equal at point where sites enter $T_A(r(\sigma))$. The marginal condition holds in this case if $S_D(\sigma, \widehat{\alpha})\beta'(R-\sigma)$ is between these values.

to additional spending are decreasing.[17]

In sum, the optimal division of resources between border and site defense does equate the marginal gain from investing in border defense with the marginal gain from investing more in site defense. But the "strategic" marginal gain from site defense presumes that the defender allocates any additional spending on site defense so as to keep the attacker's payoff to attacking any defended site the same. Spreading the additional resources in this way reduces the effect of additional spending on site defense.

Threats with Strategic and Non-Strategic Components

Some threats entail both a strategic and non-strategic component and defending against one kind of threat also protects against the other. Chemical facilities in densely populated areas may pose a risk because of accidents as well as terrorism.[18] Many public health measures taken to detect and contain contagious diseases defend against natural outbreaks as well as deliberate bio-terror attacks (Chyba 2001, 2002). Improved building codes for tall towers can help people evacuate more quickly in the event of a fire as well as a bomb (Dwyer and Lipton 2005). Better emergency communications and interoperability facilitate coordination and response to natural disasters and deliberate attacks (Mayer-Schönberger 2003). How should a defender allocate its resources when threats involve both a strategic and non-strategic component?

The answer and the way of thinking about finding the optimal allocation turn out to be quite similar to the problem of border defense. Imagine that the defender dedicated all of its resources to defending against the non-strategic threat. At this allocation some site, say t_1, offers the terrorist its highest payoff. If therefore the defender wants to defend against a terrorist attack, it has to shift resources away from all of the other sites

---

[17] The elasticity of defense at site j, $\epsilon_j = -\delta'_j(r_j)[r_j/\delta_j(r_j)]$, is the percentage change in δ_j resulting from a one percent change in r_j and similarly for the elasticity of border defense, ϵ_B. The responsiveness or sensitivity of the sites or border defense are decreasing if $d\epsilon_j/dr_j < 0$ and $d\epsilon_B/dr_b < 0$.

[18] For example, the EPA report on the risks posed by chemical accidents (Belke 2000) is now frequently cited in connection with the dangers of terrorist attacks (e.g., Flynn 2004, GAO 2004, Kocieniewsky 2005).

and do so in a way that optimally allocates the resources that remain directed against the nonstrategic threat. As the defender allocates more resources to protecting site t_1, the attacker's payoff to striking that site decreases until some other site, say t_2, now becomes an equally attractive target given the defender's allocation. This means that as the defender shifts more resources from protecting against the non-strategic threat to defending against the strategic threat, it must begin to distribute those additional resources across t_1 and t_2 so as to keep the terrorist's payoff to attacking either of the these sites the same. The defender continues shifting resources in this way, having to spread them across more and more sites, until the marginal gain of shifting more resources to defending against the strategic threat is just offset by the marginal loss of spending less to protect against the strategic threat. These marginal conditions define the unique equilibrium allocation.

To introduce the non-strategic component, let η_j be the probability of an "accidental," i.e., non-strategic, attack on site j. Then the overall probability of an attack on j is $\eta_j + \alpha_j(1 - \eta_j)$.[19] This implies that the defender's expected loss is $L^D(r, \alpha) = \sum_{j=1}^{N} L_j^D \delta_j(r_j)[\eta_j + \alpha_j(1 - \eta_j)]$. As always, the attacker strikes at sites that offer the highest expected payoff, i.e., at any $j \in T_A(r)$. We also assume diminishing marginal returns to defense, i.e., $\delta_j''(r_j) > 0$ for all j whenever $\delta_j(r_j) > 0$.

Once again the optimal resource allocation path plays a critical role in the analysis. This path starts at the allocation that would minimize the defender's expected loss if it were only facing the non-strategic threat, and the path ends with at the allocation at which the defender is using all of its resources to defend against the strategic threat. Along the way, the defender is shifting more and more resources from the non-strategic threat to the strategic danger.

Formally, the path begins at the allocation that equates the defender's marginal benefit to protecting a given site against the non-strategic threat with the marginal benefits of protecting any other site against that threat. That is, the optimal non-strategic allocation

---

[19] The probability of no attack on j is $(1 - \eta_j)(1 - \alpha_j)$. So the probability of an attack on j is just one minus this probability or $\eta_j + \alpha_j(1 - \eta_j)$.

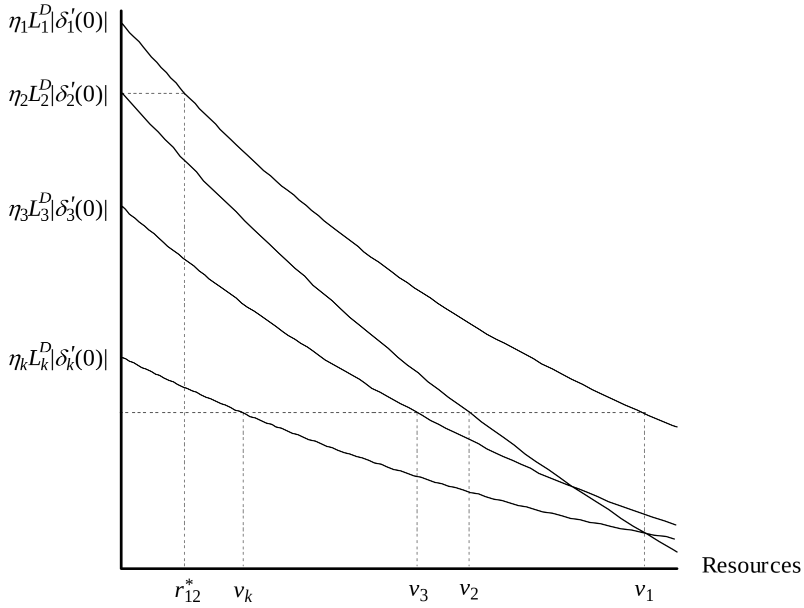

Figure 3: The optimal allocation against a non-strategic threat.

v satisfies three conditions: (i) $\eta_j L_j^D \delta_j'(v_j) = \eta_k L_k^D \delta_k'(v_k)$ whenever $v_j > 0$ and $v_k > 0$, (ii) $\eta_j L_j^D \delta_j'(v_j) \le \eta_m L_m^D \delta_m'(v_m)$ whenever $v_j > 0$ and $v_m = 0$, and (iii) $v_1 + \cdots + v_N = R$.

The algorithm described in the baseline case can also be used to derive this allocation. Assume that the defender has not allocated any of its resources and that the sites are now indexed so that the defender's marginal gain (in absolute value) to protecting site 1 against the non-strategic threat is larger than its marginal gain to protecting site 2 against a non-strategic threat which is larger than its expected gain to striking site 3 and so on. That is, $\eta_1 L_1^D |\delta_1'(0)| > \eta_2 L_2^D |\delta_2'(0)| > \cdots > \eta_N L_N^D |\delta_N'(0)|$ where, recall, $\delta_j' < 0$ and $\delta_j'' > 0$.[20]

To minimize the defender's expected loss against the non-strategic threat, the defender begins by allocating resources to the site with the largest marginal benefit. As the defender devotes more resources to defending site 1, its marginal gain declines as illustrated in Figure 3. At r_{12}^*, $\eta_1 L_1^D|\delta_1'(r_{12}^*)| = \eta_2 L_2^D|\delta_2'(0)| = \max\{\eta_1 L_1^D|\delta_1'(r_{12}^*)|, \eta_2 L_2^D|\delta_2'(0)|, ...,$

---

[20] Because $\delta_j' < 0$, the size of the marginal gain of investing in j is $\eta_j L_j^D \left|\delta_j'(0)\right|$. The inequalities are again assumed to be strict in order to simplify the exposition. Generalizing the analysis if the inequalities are weak is straighforward.

$\eta_N L_N^D|\delta'_N(0)|\}$. At this point, the defender must begin to allocate resources to both sites 1 and 2 so that $\eta_1 L_1^D|\delta'_1(r_1)| = \eta_2 L_2^D|\delta'_2(r_2)|$. Otherwise the defender could profitably reallocate resources between these two sites. The defender continues in this way with the marginal return from protecting the sites decreasing but constant across those sites until the defender has fully allocated R. This defines the optimal non-strategic allocation v.

Then a point is on the path P_v leading from v if and only if the resources guarding against the non-strategic threat are optimally allocated against that threat and those devoted to the strategic threat minimize the attacker's expected payoff. In symbols, $r \in P_v$ if and only if: (i) $\eta_j L_j^D \delta'_j(r_j) = \eta_k L_k^D \delta'_k(r_k)$ whenever $r_j > 0$, $r_k > 0$, and $j, k \notin T_A(r)$; and (ii) $\eta_j L_j^D \delta'_j(r_j) \leq \eta_m L_m^D \delta'_m(r_m)$ whenever $r_j > 0$, $r_m = 0$, and $j, m \notin T_A(r)$. These conditions guarantee that the resources devoted to the non-strategic threat, i.e., to sites $j \notin T_A(r)$, cannot be profitably reallocated among those sites. As for the resources dedicated to the strategic threat, namely to those $j \in T_A(r)$, the definition of $T_A(r)$ ensures that the resources spent on these sites minimize the attacker's payoff over these sites.[21] Finally, identify a specific point on the path with the amount of resources dedicated to the strategic threat. That is, $P_v(\sigma)$ is the point $r \in P_v$ such that $\sum_{j \in T_A(r)} r_j = \sigma$ where σ is defined over the interval $\sigma \in [\underline{\sigma}, \overline{\sigma}]$ with $\underline{\sigma} = \sum_{j \in T_A(v)} v_j$ and $\overline{\sigma} = R$.

As the defender moves along P_v it is effectively changing the mix between the strategic and non-strategic threat but optimally allocating however much it dedicates to each of these separate threats. The unique equilibrium allocation is the point along the path at which the marginal gain from allocating more to the strategic threat is just offset by the marginal loss from dedicated less to the non-strategic threat. To characterize this point formally, define the defender's expected loss from defending against the strategic threat at $r = P_v(\sigma)$ to be $L_S^D(\sigma, \alpha) \equiv \sum_{j \in T_A(r(\sigma))} L_j^D \delta_j(r_j(\sigma))[\eta_j + \alpha_j(1 - \eta_j)]$ and from

---

[21] Fix the set of sites in $T_A(r)$. Then the optimal allocation of the resources $\sigma = \sum_{j \in T_A(r)} r_j$ across these sites minmizes the attacker's payoff $\max\{A_j \delta_j(r_j)\colon j \in T_A(r)\}$. But, $A_j \delta_j(r_j) = A_k \delta_k(r_k)$ for all $j, k \in T_A(r)$ by construction. So, any reallocation $\widehat{r}$ of σ across $T_A(r)$ must give fewer resources to at least one site n. However, $\widehat{r}_n < r_n$ implies $A_n \delta_n(\widehat{r}_n) > A_n \delta_n(r_n) = \max\{A_j \delta_j(r_j) : j \in T_A(r)\}$ and consequently that $\widehat{r}$ is not an optimal allocation of σ.

the non-strategic threat to be $L_v^D(\sigma,\alpha) \equiv \sum_{n \notin T_A(r(\sigma))} L_n^D \delta_n(r_n(\sigma))[\eta_n + \alpha_n(1-\eta_n)]$. The defender's expected loss can now be written as is $L^D(r(\sigma),\alpha) = L_v^D(\sigma,\alpha) + L_S^D(\sigma,\alpha)$. Then the unique equilibrium allocation is either at one of the endpoints of the path, v or $P_v(R)$, or satisfies $-\partial L_v^D(\sigma,\widehat{\alpha})/\partial\sigma = \partial L_S^D(\sigma,\widehat{\alpha})/\partial\sigma$ where $\widehat{\alpha}$ is attacker's "pseudo-equlibrium" strategy at r. That is, $\widehat{\alpha}$ is the best-response to r that also rationalizes the defender's allocation across the sites in $T_A(r)$ and, as shown in the appendix, is given by $\widehat{\alpha}_n = 0$ for $n \notin T_A(r(\sigma))$ and

$$\eta_j + \widehat{\alpha}_j(1-\eta_j) \;=\; \left(\sum_{k \in T_A(r(\sigma))} \frac{1}{1-\eta_k}\right)\left(\sum_{k \in T_A(r(\sigma))} \frac{L_j^D \delta_j'}{L_k^D \delta_k'(1-\eta_k)}\right)^{-1} \tag{1}$$

for $j \in T_A(\sigma)$. Proposition 3 collects these results.

PROPOSITION 3: *There is a unique equilibrium allocation which lies at one of the endpoints of P_v or is the point at which the marginal gain of spending more on the strategic threat is just offset by the marginal loss of spending less on the non-strategic threat and is defined almost everywhere by $-\partial L_v^D(\sigma,\widehat{\alpha})/\partial\sigma = \partial L_S^D(\sigma,\widehat{\alpha})/\partial\sigma$ and which is equivalent to:*[22]

$$\eta_j L_j^D \delta_j'(\widehat{r}_j) = \left(\sum_{k \in T_A(\widehat{r})} \frac{1}{1-\eta_k}\right)\left(\sum_{k \in T_A(\widehat{r})} \frac{1}{L_k^D \delta_k'(\widehat{r}_k)(1-\eta_k)}\right)^{-1}$$

for any $j \notin T_A(\widehat{r})$ for which $\widehat{r}_j > 0$.

In addition to the marginal conditions that characterize the equilibrium allocation, two broader conditions or rules of thumb describe circumstances in which the defender is sure to be able to profitably reallocate its resources between strategic and non-strategic threats. Let e be the strategic site that offers the defender its highest return to additional spending. That is, $L_k^D \delta_k' \geq L_e^D \delta_e'$ for e and $k \in T_A(r)$. Then the defender can profitably reallocate resources to the non-strategic threat if $L_e^D \delta_e'(r_e) > \eta_m L_m^D \delta_m'(r_m)$ for any $m \notin T_A(r)$. Recalling that the marginal gain to defending sites not in $T_A(r)$ along the path P_v is the

---

[22] This marginal conditon pins down the equilibrium condition almost everywhere along the path. But as before with border defense, the pseudo-equilibrium strategy $\widehat{\alpha}$ and the partial derivatives may not be well defined at the finitely many points along P_v where one or more sites joins $T_A(r)$. The proof in the appendix addresses this issue.

same at each site (as long as there is a positive level of spending), this condition amounts to saying that the marginal gain in allocating more to a non-strategic site is greater (i.e., does more to lower the defender's expected loss) than the gain from investing in the easiest strategic site to defend.

Conversely, the defender can profitably invest more in strategic defense when the marginal return from dedicating more to the hardest strategic site to defend is greater than the marginal gain of spending more on any non-strategic site, i.e., when $\eta_m L_m^D \delta_m'(r_m) > L_h^D \delta_h'(r_h)$ where $m \notin T_A(r)$ and $L_h^D \delta_h' \geq L_k^D \delta_k'$ for h and any $k \in T_A(r)$. Note that neither the number of sites the defender is guarding against the strategic threat, i.e., the number of sites in $T_A(r)$, nor the number of sites the defender is protecting against the non-strategic threat, i.e., the number of defended sites not in $T_A(r)$, play a role in these two approximations.

Once again a simple algorithm yields the equilibrium allocation. The defender starts at the allocation that would be optimal were there no strategic threat. Given this allocation, a specific site or set of sites offers the attacker its highest payoff to attacking. If, therefore, the defender wants to begin to guard against the strategic threat, it needs to begin spending more on defending these sites and, necessarily, less on the other sites. As the defender devotes more and more to strategic defense, it has to distribute these resources across a larger and larger set of sites. At the equilibrium allocation, the marginal loss from having less to spend against the non-strategic threat just offsets the marginal gain from additional spending on the strategic threat.

Uncertainty about the Attacker's Preferred Targets

The previous discussion assumed that the defender knows how the attacker ranks the potential targets. But the defender may actually be uncertain about the types of targets that a terrorist group would most like to strike. How should a defender allocate its resources in the face of this uncertainty?

We simplify matters here by assuming that the defender is unsure whether it is facing one of two possible types of attacker, γ and τ. These types may get different payoffs from attacking a particular site and may even rank the sites in a different order. Attacker γ might, for example, rank a successful attack on a nuclear facility higher than on a chemical site whereas τ might rank the latter higher than the former.

The answer to the resource allocation question and the way of thinking about finding the optimal allocation turn out to be quite similar to the border-defense and non-strategic-threat problems. Imagine that the defender were certain that it was facing γ and allocated all of its resources accordingly. (The baseline algorithm can be used to find this allocation.) At this allocation some site, say t_1, offers τ its highest payoff. If therefore the defender wants to hedge against the possibility that it is facing τ, it should shift resources away from all the other sites, keeping γ's payoff to attacking any one of those sites equal to its payoff to striking any other. As the defender dedicates more resources to defending against τ by hardening t_1, τ's payoff to striking this site decreases until some other site, say t_2, now becomes an equally attractive target. This means that as the defender shifts more resources from protecting against γ to defending against τ, the defender must spread those additional resources across t_1 and t_2 so as to keep τ's payoff to attacking either of these sites the same. The defender continues shifting resources in this way, having to spread them across more and more sites, until the marginal gain to defending against τ is just offset by the marginal loss of spending less on defending against γ. These marginal conditions define the unique equilibrium allocation.

To formalize the defender's uncertainty and the argument as a whole, suppose that the defender believes that it is facing γ with probability μ and τ with probability $1-\mu$. Type γ's payoff to a successful attack on site j is A_j^γ and τ's payoff is A_j^τ. Then the attackers' expected payoffs are $A^\gamma(r,\alpha) \equiv \sum_{j=1}^N \alpha_j A_j^\gamma \delta_j(r_j)$ and $A^\tau(r,\phi) \equiv \sum_{j=1}^N \phi_j A_j^\tau \delta_j(r_j)$ where α and ϕ are γ's and τ's strategies. The defender's expected loss is $L^D(r,\alpha,\phi) \equiv \sum_{j=1}^N [\mu\alpha_j + \phi_j(1-\mu)] L_j^D \delta_j(r_j)$. As with non-strategic threats, we assume diminishing marginal returns to defense, so $\delta_j''(r_j) > 0$ for all j whenever $\delta_j(r_j) > 0$.

To specify the path defined by shifting resources from defending against γ to guarding against τ, let g and t be the defender's optimal allocation if it were sure that it was facing γ or τ respectively. Then a point is on the path P_g from g to t if the defender is allocating all of its resources against the sites that offer γ or τ their highest expected payoffs. That is, $\widehat{r} \in P_g$ if and only if $\sum_{j\in T_\gamma(\widehat{r})\cup T_\tau(\widehat{r})} \widehat{r}_j = R$ where $T_\gamma(r) = \{j : A^\gamma_j \delta_j(r_j) = \max\{A^\gamma_1 \delta_1(r_1),$ ..., $A^\gamma_N \delta_N(r_N)\}\}$ and similarly for $T_\tau(r)$. Paralleling the notation used in the case of threats with a non-strategic component, take $P_g(\sigma)$ to be the point $r \in P_g$ where the defender allocates σ against τ, i.e., $\sigma = \sum_{j\in T_\tau(\widehat{r})} r_j$, for $\sigma \in [\underline{\sigma}, \overline{\sigma}]$ with $\underline{\sigma} = \sum_{j\in T_\tau(g)} r_j$ and $\overline{\sigma} = \sum_{j\in T_\tau(t)} r_j = R$.

As the defender moves along P_g from g to t, it is shifting more and more resources away from protecting against γ to defending against τ. But it is also shifting these resources so that however much the defender dedicates to guarding against one type of attacker, the defender is optimally allocating those resources against that type. The unique equilibrium allocation is the point along this path at which the marginal gain to allocating more to defending against τ is just offset by the marginal losses of spending less against γ.

To characterize this allocation, take the defender's expected loss from defending against γ and τ at any $\widehat{r} = P_g(\sigma)$ to be $L^D_\gamma(\sigma, \alpha) \equiv \mu \sum_{j\in T_A(\widehat{r}(\sigma))-T_\tau(\widehat{r}(\sigma))} \alpha_j L^D_j \delta_j(\widehat{r}_j(\sigma))$ and $L^D_\tau(\sigma, \phi) \equiv (1-\mu) \sum_{j\in T_\tau(\widehat{r}(\sigma))} \phi_j L^D_j \delta_j(\widehat{r}_j(\sigma))$ where $T_A(\widehat{r}) = T_\gamma(\widehat{r}) \cup T_\tau(\widehat{r})$. Combining these payoffs gives the defender's total payoff $L^D(\widehat{r}(\sigma), \alpha) = L^D_\gamma(\sigma, \alpha) + L^D_\tau(\sigma, \alpha)$. Then the equilibrium allocation is either at one of the endpoints, $P_g(\underline{\sigma})$ or $P_g(\overline{\sigma})$, or satisfies $\partial L^D_\tau(\sigma, \widehat{\phi})/\partial\sigma = -\partial L^D_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ where $\widehat{\alpha}$ and $\widehat{\phi}$ are γ's and τ's pseudo-equilibrium strategies. This leaves:

PROPOSITION 4: *There is a unique equilibrium allocation which lies at one of the endpoints of P_g or is the point along the path at which the marginal gain of spending more on defending against τ is just offset by the marginal loss of spending less on defending against γ and is defined almost everywhere by $\partial L^D_\tau(\sigma, \widehat{\phi})/\partial\sigma = -\partial L^D_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ which is equivalent to:*[23]

$$(1-\mu)\left(\sum_{n\in T_\gamma(r)} \frac{1}{L^D_n \delta'_n(r_n)}\right)^{-1} = \mu\left(\sum_{n\in T_\gamma(r)} \frac{1}{L^D_n \delta'_n(r_n)}\right)^{-1}.$$

As with border defense and non-strategic threats, some looser conditions characterize

---

[23] As with border defense and threats with non-strategic components, these partial derivatives are generally not well defined at the finitely many points along the path at which sites leave $T_\gamma(r(\sigma))$ or enter $T_\tau(r(\sigma))$. This issue is discussed in the proof in the appendix.

circumstances in which the defender is sure to be able to profitably reallocate its resources. Let $h(i)$ and $e(i)$ for $i = \gamma$ or τ be the sites that offer the lowest and highest return on investing against type i. That is, $L^D_{h(i)}\delta'_{h(i)} \geq L^D_k \delta'_k \geq L^D_{e(i)}\delta'_{e(i)}$ for $k \in T_i(r)$.

Were the defender to allocate more to defending against τ, it would have to spread this spending across the sites in $T_\tau(r)$, and the defender's marginal gain at each site is at least $L^D_{h(\tau)}\delta'_{h(\tau)}$. Similarly, any more spending devoted to guarding against γ must be distributed across the sites in $T_\gamma(r)$. Hence, protecting against τ is sure to offer a larger return than defending against γ whenever $(1-\mu)L^D_{h(\tau)}\delta'_{h(\tau)}/\left\|T_\tau(r)\right\| < \mu L^D_{e(\gamma)}\delta'_{e(\gamma)}/\left\|T_\gamma(r)\right\|$ where $\|S\|$ is the number of sites in a set S. Conversely, the return to protecting against γ is sure to exceed that of guarding against τ whenever $(1-\mu)L^D_{e(\tau)}\delta'_{e(\tau)}/\left\|T_\tau(r)\right\| > \mu L^D_{h(\gamma)}\delta'_{h(\gamma)}/\left\|T_\gamma(r)\right\|$.

In sum, there is a simple way to think about the defender's resource-allocation problem when it is uncertain whether it is facing one of two possible types of attacker. The defender begins at the allocation that would optimally distribute its resources if it were certain of the adversary's goals. Given this division, a specific site or set of sites will offer an attacker with different priorities its highest payoff to attacking. This means that if the defender wants to hedge against the possibility of facing this other possible attacker, it should spend more on those sites and necessarily less on the other sites. As it begins to shift resources away from guarding against one type of attacker, it has to spread the resources that it is using to guard against that type over fewer sites. Conversely, the defender has to distribute the resources it dedicating to the other type across more and more sites. The point at which the marginal benefits of spending less to protect against one type are just offset by the marginal gains of spending more against the other type defines the equilibrium allocation.

Conclusion

It turns out that we can think of the defender's allocation problem in four different substantive settings in a broadly similar way. A strategic adversary deciding where to attack attempts to strike where the defender is weak and the expected gains are large.

Anticipating this, the defender tries to minimize the attacker's highest expected payoff. But as the defender invests more in protecting the sites offering the attacker its highest payoff, that payoff declines and eventually other sites become equally attractive. At this point, the defender has to begin to devote resources to protecting those sites as well. As the defender allocates more and more to site defense, it must spread those resources across a larger and larger set of sites.

This algorithm defines the optimal allocation path if the sites are independent and the attacker only has to decide how much to invest in hardening each site. This process can also be used to define the optimal allocation path which is the key to the defender's resource allocation problem in a variety of settings. As the defender moves along this path, the defender's marginal gains from continuing decline as the it has to distribute those resources across more sites. At the equilibrium allocation, the marginal cost of moving farther along the path by spending less on border defense, non-strategic threats, or guarding against a particular type of adversary just offsets the marginal gain of spending more on site defense, protecting against strategic threats, or hedging against the possibility of facing a different type of attacker.

Appendix

Proof of Proposition 1: The algorithm described above clearly yields an allocation $\widehat{r}$ that minimizes $M_A(r|R)$. The discussion in the text after the proposition also establishes that $\widehat{r}$ is an equilibrium allocation by finding an $\widehat{\alpha}$ such that $(\widehat{r}, \widehat{\alpha})$ is an equilibrium. What remains to be shown is that $\widehat{r}$ is the unique allocation that minimizes $M_A(r|R)$ and that $\widehat{r}$ is the unique equilibrium allocation.

To see that $\widehat{r}$ uniquely minimizes $M_A(r|R)$, suppose $r^* \neq \widehat{r}$. Then there exists a site j that the defender protects in $\widehat{r}$ and that receives fewer resources in r^*, i.e., $\widehat{r}_j > 0$ and $\widehat{r}_j > r_j^*$. But $\widehat{r}_j > 0$ means $A_j\delta_j'(\widehat{r}_j) = M_A(\widehat{r})$ and $\widehat{r}_j > r_j^*$ implies $M_A(\widehat{r}) < A_j\delta_j'(r_j^*) \leq M_A(r^*)$. Hence, r^* does not minimize $M_A(r)$.

Now suppose $r^* \neq \widehat{r}$ is an equilibrium allocation. Then there must exist an α^* such that (r^*, α^*) is an equilibrium. The fact that $r^* \neq \widehat{r}$ implies $M_A(r^*|R) > M_A(\widehat{r}|R)$. It follows that there exists some site j such that $r_j^* > 0$ and $A_j\delta_j(r_j^*) > M_A(r^*|R)$. But any of the attacker's best replies to r^* only put positive probability on the sites that offer the highest payoff, i.e., those k for which $A_k\delta_k(r_k^*) < M_A(r^*|R)$. Hence, $\alpha_j^* = 0$ in any best reply α^* to r^*. The defender then can profitably deviate from r^* against any best reply to it by reallocating resources away from j. Consequently, r^* is not an equilibrium allocation. ■

Proof of Proposition 2: To establish that all equilibrium allocations lie on P_B, suppose $\widehat{r} \notin P_B$ is an equilibrium allocation. Then $\widehat{r}_k > 0$ for some $k \notin T_A(\widehat{r})$, and the defender can clearly profitably deviate from $\widehat{r}$ by reallocating $\widehat{r}_k$ across $T_A(\widehat{r})$.

To characterize the equilibrium allocations along P_B and demonstrate existence, note that the marginal gain from investing slightly more in site or point defense $\partial L_p^D(\sigma, \widehat{\alpha})/\partial\sigma$ is well defined along P_B except possibly at the finitely many point where the defender begins to allocate resources to an additional site. To see that this is so, say that a site j becomes a target by entering T_A at $r(\widehat{\sigma}) \in P_B$ if $j \in T_A(r(\widehat{\sigma}))$ and $j \notin T_A(r(\sigma))$ for $\sigma < \widehat{\sigma}$. Because there are only finitely many sites, there are only finitely many points along P_B at which sites become targets. Let E be this set along with the endpoint $P_v(0)$

where the defender invests nothing in point defense and the other endpoint $P_v(R)$ where everything is devoted to point defense.

Then $\partial L_p^D(\sigma, \widehat{\alpha})/\partial\sigma$ is well defined at $P_B - E$ because the $\widehat{\alpha}$ are well defined. Consider a $\widehat{r} = P_B(\widehat{\sigma}) \in P_B - E$. Then there exists a neighborhood around $\widehat{\sigma}$ such that the sites that maximize the attacker's payoff are the same, i.e., $T_A(r(\sigma)) = T_A(\widehat{r})$. This means that when summing over the elements of $T_A(r(\sigma))$ in the neighborhood of $\widehat{\sigma}$, the set of sites over which the summation is taken does not change.

This fact ensures that the "pseudo-equilibrium" strategy $\widehat{\alpha}$ is well defined at $\widehat{\sigma}$ where, recall, a pseudo-equilibrium strategy is the best reply to $\widehat{r}$ that also rationalizes the defender's relative allocations across the sites in $T_A(\widehat{r})$. Because $\widehat{\alpha}$ is a best response to $\widehat{r}$, $\widehat{\alpha}_j = 0$ for all $j \notin T_A(\widehat{r})$. As for the $j \in T_A(\widehat{r})$, the attacker is indifferent to striking any site in $T_A(\widehat{r})$, so any $\widehat{\alpha}_j$'s such that $\sum_{j\in T_A(\widehat{r})} \widehat{\alpha}_j = 1$ will be a best reply to $\widehat{r}$. But the $\widehat{\alpha}_j$'s must also rationalize the defender's relative spending across the sites in $T_A(\widehat{r})$ and, therefor, satisfy the marginal conditions $\widehat{\alpha}_j L_j^D \delta'_j(\widehat{r}_j) = \widehat{\alpha}_k L_k^D \delta'_k(\widehat{r}_k)$ for all j, $k \in T_A(\widehat{r})$. (That no site enters T_A at $\widehat{r}$ means that $\widehat{r}_j$ and $\widehat{r}_k$ are positive and therefore that the marginal condition must hold lest the defender have an incentive to reallocate between j and k.) Solving these equations for $\widehat{\alpha}_j$ in terms of $\widehat{\alpha}_k$, summing the $\widehat{\alpha}_j$ over $T_A(\widehat{r})$, using that fact that is sum equals one, and rearranging terms yield $\widehat{\alpha}_k = \left(\sum_{n\in T_A(\widehat{r})} L_k^D \delta'_k / L_n \delta'_n\right)^{-1}$ for $k \in T_A(\widehat{r})$.

With $\widehat{\alpha}$ defined, $\partial L_p^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ is also well defined and given by $\partial L_p^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = \sum_{j\in T_A(r)} \widehat{\alpha}_j L_j^D \delta'_j(\widehat{r}_j) d\widehat{r}_j/d\sigma$ where recall $\widehat{\sigma}$ is the amount of resources allocated to defending sites in $T_A(\widehat{r})$. That is, $\sum_{k\in T_A(\widehat{r})} \widehat{r}_j = \widehat{\sigma}$ which implies $\sum_{k\in T_A(\widehat{r})} d\widehat{r}_j/d\sigma = 1$. Substituting for $\widehat{\alpha}_j$ then leaves:

$$\begin{aligned}
\frac{\partial L_p^D(\widehat{\sigma}, \widehat{\alpha})}{\partial\sigma} &= \left(\sum_{n\in T_A(\widehat{r})} \frac{1}{L_n^D \delta'_n(\widehat{r}_n)}\right)^{-1} \sum_{k\in T_A(\widehat{r})} \frac{d\widehat{r}_j}{d\sigma} \\
&= \left(\sum_{n\in T_A(\widehat{r})} \frac{1}{L_n^D \delta'_n(\widehat{r}_n)}\right)^{-1}.
\end{aligned}$$

Now consider a $\widehat{r} = P_B(\widehat{\sigma})$ in E but not an endpoint. Because a site, say e, enters T_A at $\widehat{\sigma}$, the summations in the expression for $\partial L_p^D(\sigma, \widehat{\alpha})/\partial\sigma$ for $\sigma \in (\widehat{\sigma}, \widehat{\sigma} + \varepsilon)$ and for $\sigma \in (\widehat{\sigma}, \widehat{\sigma} - \varepsilon)$ are taken over different sets. (Only one site enters because we have assumed the values of $L_j^A \delta_j(0)$ are all distinct. This assumption simplifies the analysis but does not affect the overall results.) The continuity of the individual terms will therefore not ensure the continuity of the sums, and $\partial L_p^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma^- \equiv \lim_{\sigma\to\widehat{\sigma}^-} \partial L_p^D(\sigma, \widehat{\alpha})/\partial\sigma$ will generally not equal $\partial L_p^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma^+ \equiv \lim_{\sigma\to\widehat{\sigma}^+} \partial L_p^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$.

To finesse this issue, take X to be the correspondence from points on P_B into $\mathbb{R}$ where $X(\widehat{\sigma})$ for $\widehat{r} = P_B(\widehat{\sigma})$ is the closed interval between $\sum_{n\in T_A(\widehat{r})} \delta_n(\widehat{r}_n)/\delta_n'(\widehat{r}_n)$ and $\sum_{n\in T_A(\widehat{r})-\{e\}} \delta_n(\widehat{r}_n)/\delta_n'(\widehat{r}_n)$ where the second summation is over the elements of $T_A(\widehat{r})$ less an entering site if there is one. Of course, sites can only enter at the finite points in E, so $X(\widehat{\sigma})$ is simply a singleton at all $\widehat{r} \in P_B - E$. As for the endpoints, define $X(\underline{\sigma})$ and $X(\overline{\sigma})$ to be the one-sided limits $\delta_1(0)/\delta_1'(0)$ and $\sum_{n\in T_A(r(\overline{\sigma}))-\{e\}} \delta_n(r_n(\overline{\sigma}))/\delta_n'(r_n(\overline{\sigma}))$ where, recall, the defender allocates all of its resources to border defense at $P_B(\underline{\sigma})$ (and therefore $T_A(\widehat{r}) = \{1\}$ with $\widehat{r}_1 = 0$) and none of its resources on border defense at $X(\overline{\sigma})$.

Then any $\widehat{r} = P_B(\widehat{\sigma})$ is an equilibrium allocation if $\beta(R - \widehat{\sigma})/\beta'(R - \widehat{\sigma}) \in X(\widehat{\sigma})$. To establish this suppose, first, that $\widehat{r} \notin E$. Because no sites enter at $\widehat{r}$, $\widehat{r}_j > 0$ for $j \in T_A(\widehat{r})$ and $\widehat{r}_j = 0$ otherwise. Hence, $\widehat{r}$ is a best reply to $\widehat{\alpha}$ if and only if $\partial L^D(\widehat{r}, \widehat{\alpha})/\partial r_j = 0$ for $j \in T_A(\widehat{r})$ where $\partial L^D(\widehat{r}, \widehat{\alpha})/\partial r_j = \widehat{\alpha}_j L_j^D \delta_j'(\widehat{r}_j)\beta(R - \widehat{\sigma}) - \sum_{n=1}^N \widehat{\alpha}_n L_n^D \delta_n(\widehat{r}_n)\beta'(R - \widehat{\sigma})$. (Recall that the resources dedicated to border defense are implicitly defined by $b = R - \sum_{n=1}^N \widehat{r}_n$ which also means $b = R - \widehat{\sigma}$).

That $X(\widehat{\sigma})$ is a singleton, $\beta(R - \widehat{\sigma})/\beta'(R - \widehat{\sigma}) \in X(\widehat{\sigma})$, and the expressions for $\widehat{\alpha}_j$

imply

$$
\begin{aligned}
\frac{\beta(R-\widehat{\sigma})}{\beta'(R-\widehat{\sigma})} &= \sum_{n\in T_A(\widehat{r})} \frac{\delta_n(\widehat{r}_n)}{\delta_n'(\widehat{r}_n)} \\
\beta(R-\widehat{\sigma})\left(\sum_{k\in T_A(\widehat{r})} \frac{1}{L_k^D\delta_k'(\widehat{r}_k)}\right)^{-1} &= \beta'(R-\widehat{\sigma})\sum_{n\in T_A(\widehat{r})}\left(\sum_{k\in T_A(\widehat{r})} \frac{1}{L_k^D\delta_k'(\widehat{r}_k)}\right)^{-1}\frac{\delta_n(\widehat{r}_n)}{\delta_n'(\widehat{r}_n)} \\
\widehat{\alpha}_j L_j^D\delta_j'(\widehat{r}_j)\beta(R-\widehat{\sigma}) &= \sum_{n=1}^{N}\widehat{\alpha} L_n^D\delta_n(\widehat{r}_n)\beta'(R-\widehat{\sigma})
\end{aligned}
$$

for all $j\in T_A(\widehat{r})$. Hence, $\partial L^D(\widehat{r},\widehat{\alpha})/\partial r_j = 0$ for $j\in T_A(\widehat{r})$, and $\widehat{r}$ is an equilibrium allocation because $(\widehat{r},\widehat{\alpha})$ is an equilibrium. (The other marginal conditions $\partial L^D(\widehat{r},\widehat{\alpha})/\partial r_k \geq 0$ hold trivially at $j\notin T_A(\widehat{r})$ as $\widehat{\alpha}_j = 0$.)

Now consider any $\widehat{r} = P_B(\widehat{\sigma})\in E$ for which $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma})\in X(\widehat{\sigma})$. To show that $\widehat{r}$ is an equilibrium, we must find an $\widehat{\alpha}$ such that $(\widehat{r},\widehat{\alpha})$ is an equilibrium. To this end, suppose site e enters T_A at $\widehat{r}$. Then $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma})\in X(\widehat{\sigma})$ implies $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma}) = \sum_{n\in T_A(\widehat{r})-\{e\}}\delta_n(\widehat{r}_n)/\delta_n'(\widehat{r}_n) + \theta\delta_e(\widehat{r}_e)/\delta_e'(\widehat{r}_e)$ for a $\theta\in[0,1]$.

Because e enters T_A at $\widehat{r}$, $\widehat{r}_j > 0$ for all $j\neq e$ and $\widehat{r}_e = 0$. Consequently, the marginal conditions required of $\widehat{r}_j$ for $j\neq e$ and $\widehat{r}_e$ are slack since the defender cannot transfer resources from e to some other site. More specifically, the marginal conditions that must hold if $\widehat{r}$ is to be a best reply to $\widehat{\alpha}$ are $\widehat{\alpha}_j L_j^D\delta_j(\widehat{r}_j) = \widehat{\alpha}_k L_k^D\delta_k(\widehat{r}_k)$ and $\widehat{\alpha}_j L_j^D\delta_j'(\widehat{r}_j)\leq \widehat{\alpha}_e L_e^D\delta_e'(0)$ for $j,k\in T_A(\widehat{r})-\{e\}$. The effect of this is that when we solve for $\widehat{\alpha}_j$ by summing over the marginal conditions that hold with equality, the total is now $1-\widehat{\alpha}_e$ and leaves $\widehat{\alpha}_j = (1-\widehat{\alpha}_e)\left(\sum_{n\in T_A(\widehat{r})-\{e\}} L_j^D\delta_j'(\widehat{r}_j)/\left[L_n^D\delta_n'(\widehat{r}_n)\right]\right)^{-1}$. As for $\widehat{\alpha}_e$, it must lie in the range $[0,\overline{\alpha}_e]$ where $\overline{\alpha}_e$ satisfies $\widehat{\alpha}_j L_j^D\delta_j'(\widehat{r}_j) = \overline{\alpha}_e L_e^D\delta_e'(0)$. This and the expression for $\widehat{\alpha}_j$ mean $\overline{\alpha}_e L_e^D\delta_e'(0) = (1-\overline{\alpha}_e)\left(\sum_{n\in T_A(\widehat{r})-\{e\}}\left[L_n^D\delta_n'(\widehat{r}_n)\right]^{-1}\right)^{-1}$.

Take $\theta = \left[\widehat{\alpha}_e L_e^D\delta_e'(0)/(1-\widehat{\alpha}_e)\right]\sum_{n\in T_A(\widehat{r})-\{e\}}\left[L_n^D\delta_n'(\widehat{r}_n)\right]^{-1}$. Clearly, $\theta = 0$ at $\widehat{\alpha}_e = 0$ and $\theta = 1$ at $\widehat{\alpha}_e = \overline{\alpha}_e$, so this mapping associates a unique $\widehat{\alpha}_e\in[0,\overline{\alpha}_e]$ with every $\theta\in[0,1]$. Then $\widehat{r}$ is a best reply to $\widehat{\alpha}$.

To see that this is so, substitute the expression for θ in β/β':

$$\begin{aligned}
\frac{\beta(R-\widehat{\sigma})}{\beta'(R-\widehat{\sigma})} &= \sum_{n\in T_A(\widehat{r})-\{e\}} \frac{\delta_n(\widehat{r}_n)}{\delta'_n(\widehat{r}_n)} + \frac{\widehat{\alpha}_e L_e^D \delta'_e(0)}{1-\widehat{\alpha}_e} \sum_{k\in T_A(\widehat{r})-\{e\}} \frac{1}{L_k^D \delta'_k(\widehat{r}_k)} \frac{\delta_e(0)}{\delta'_e(0)} \\
\frac{\beta(R-\widehat{\sigma})}{\beta'(R-\widehat{\sigma})} &= \sum_{n\in T_A(\widehat{r})-\{e\}} \left(\frac{\widehat{\alpha}_n L_n^D \delta'_n(\widehat{r}_n)}{\widehat{\alpha}_j L_j^D \delta'_j(\widehat{r}_j)}\right) \frac{\delta_n(\widehat{r}_n)}{\delta'_n(\widehat{r}_n)} + \frac{\widehat{\alpha}_e L_e^D \delta'_e(0)}{\widehat{\alpha}_j L_j^D \delta'_j(\widehat{r}_j)} \frac{\delta_e(0)}{\delta'_e(0)} \\
\widehat{\alpha}_j L_j^D \delta'_j(\widehat{r}_j) \beta(R-\widehat{\sigma}) &= \beta'(R-\widehat{\sigma}) \sum_{n\in T_A(\widehat{r})} \widehat{\alpha}_n L_n^D \delta_n(\widehat{r}_n)
\end{aligned}$$

where the second equation follows from the first via the marginal conditions $\widehat{\alpha}_j L_j^D \delta'_j(\widehat{r}_j) = \widehat{\alpha}_k L_k^D \delta'_k(\widehat{r}_k)$ and the expression for $\widehat{\alpha}_j$ in terms of $\widehat{\alpha}_e$ for $j,k \in T_A(\widehat{r})-\{e\}$. But the last equation means $\partial L^D(\widehat{r},\widehat{\alpha})/\partial r_j = 0$ for $j \in T_A(\widehat{r})-\{e\}$ and $\partial L^D(\widehat{r},\widehat{\alpha})/\partial r_e \geq 0$. Therefore, $\widehat{r}$ is a best reply to $\widehat{\alpha}$, $(\widehat{r},\widehat{\alpha})$ is an equilibrium, and $\widehat{r}$ is an equilibrium allocation.

To see that an equilibrium allocation must exist, i.e., that $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma})$ must intersect X, for some σ along P_B, consider the start of the path $\underline{r}$ where the defender allocates all of its resources to border defense. Then $\underline{r}_j = 0$ for all j and site 1 offers the attacker its highest expected payoff, $T_A(\underline{r}) = \{1\}$. The attacker's best reply to this allocation is $\underline{\alpha}_1 = 1$, so $\underline{r}$ is an equilibrium allocation only if it is a best response to $\underline{\alpha}$, i.e., only if $L_1^D \delta'_1(0)\beta(R) \geq L_1^D \delta_1(0)\beta'(R)$. If, therefore, $\underline{r}$ is not an equilibrium allocation, $\delta_1(0)\beta'(R) > \delta'_1(0)\beta(R)$ or $X(\underline{\sigma}) > \beta(\underline{\sigma})/\beta'(\underline{\sigma})$.

At the end of the path $\overline{r} = P_B(\overline{\sigma})$, the defender allocates nothing to border defense. This is an equilibrium if the defender cannot profitably reallocate any resources from $j \in T_A(\overline{r})-\{e\}$ to border defense, i.e., if $L_j^D \delta'_j(\overline{r}_j)\beta(0) \leq L_j^D \delta_j(\overline{r}_j)\beta'(0)$. If, therefore, $\overline{r}$ is not an equilibrium allocation, there must be a $k \in T_A(\overline{r})-\{e\}$ such that $\delta'_k(\overline{r}_k)\beta(0) > \delta_k(\overline{r}_k)\beta'(0)$ or $\beta(0)/\beta'(0) > \delta_k(\overline{r}_k)/\delta'_k(\overline{r}_k)$. Summing over the elements of $T_A(\overline{r})$ gives $\beta(0)/\beta'(0) > \sum_{n\in T_A(\overline{r})-\{e\}} \delta_n(\overline{r}_n)/\delta'_n(\overline{r}_n)$. So, $\beta(\overline{\sigma})/\beta'(\overline{\sigma}) > \max X(\overline{\sigma})$.

With $X(\underline{\sigma}) > \beta(\underline{\sigma})/\beta'(\underline{\sigma})$ and $\beta(\overline{\sigma})/\beta'(\overline{\sigma}) > \max X(\overline{\sigma})$ if neither endpoint is an equilibrium, the fact that the graph of $X(\sigma)$ is connected and β/β' is continuous imply that β/β' and X must intersect. Hence, an equilibrium allocation exists. Finally, if the elasticities $\epsilon'_j = -\delta'_j r_j/\delta_j$ and $\epsilon'_B = -\beta' b/\beta$ are decreasing, then β'/β and δ'_j/δ_j are decreasing.

So, $\beta(R-\sigma)/\beta'(R-\sigma)$ is increasing in σ and $X(\sigma)$ is decreasing, and their intersection is unique. ■

Turning to the first approximation, the defender can profitably reallocate resources to border defense when $\beta/\beta' > \sum_{n\in T_A(r)-\{e\}} \delta_n/\delta'_n$. But the fact that e is the easiest site to defend implies

$$\sum_{n\in T_A(r)-\{e\}} \frac{\delta_n}{\delta'_n} = \sum_{n\in T_A(r)-\{e\}} \frac{A_n\delta_n}{A_n\delta'_n} \leq \sum_{n\in T_A(r)-\{e\}} \frac{A_n\delta_n}{A_e\delta'_e} < \frac{A_e\delta_e T}{A_e\delta'_e} = \frac{\delta_e T}{\delta'_e}$$

where T is the number of elements in $T_A(r)-\{e\}$ and the strict inequality follows from the fact that $A_e\delta_e = A_n\delta_n$ for all $n \in T_A$. Hence, $\beta A_e\delta'_e/T < \beta' A_e\delta_e$ ensures that the defender can profitably reallocate resources to border defense.

Proof of Proposition 3: There are three steps. Observe, first, that the defender's set of pure strategies and the attacker's set of mixed strategies are nonempty, convex, and compact. Each actor's payoff, $-L^D(r,\alpha)$ or $A(r,\alpha)$, is also continuous in r and α and quasiconcave in each actor's own strategy. This ensures that an equilibrium exists in which the defender plays a pure strategy (Fudenberg and Tirole 1991, 34). In fact, the defender's payoffs are strictly concave which ensures that it never plays a mixed strategy in equilibrium. For the defender to mix, it must be indifferent between the strategies over which it is randomizing. But concavity implies that it would strictly prefer a convex combination of these strategies and thus that the defender would have a profitable deviation from the mixed strategy.

The second step is to show that any equilibrium allocation must lie on P_v. Suppose not. Then there would be an equilibrium allocation $\widehat{r} \notin P_v$ such that one of the following two conditions does not hold: (i) $L^D_j\delta'_j(\widehat{r}_j)\eta_j = L^D_k\delta'_k(\widehat{r}_k)\eta_k$ whenever $\widehat{r}_j > 0$, $\widehat{r}_k > 0$, and $j,k \notin T_A(\widehat{r})$ and (ii) $L^D_j\delta'_j(\widehat{r}_j)\eta_j \leq L^D_m\delta'_m(\widehat{r}_m)\eta_m$ whenever $\widehat{r}_j > 0$, $\widehat{r}_m = 0$, and $j,m \notin T_A(\widehat{r})$. If the set of sites not in $T_A(\widehat{r})$ is empty or a singleton, then these conditions are satisfied trivially, so $\widehat{r}$ would lie on P_v. This implies that there must be at least two distinct sites $j,k \notin T_A(\widehat{r})$ at which at least one of the previous conditions does not hold. But if either (i) or (ii) does not hold, then the defender can clearly reduce its expected loss by reallocating its resources between j and k. This contradiction leaves $\widehat{r} \in T_A(\widehat{r})$.

The third step characterizes the equilibrium allocation on P_v and shows that it is unique. The marginal gain from investing slightly more in defending against the strategic threat $\partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$ is well defined along P_v except possibly at the finitely many points where sites enter T_A. As before, let E be this set along with the two endpoints $P_v(\underline{\sigma})$ and $P_v(\overline{\sigma})$.

Then $\partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$ is well defined at $P_v - E$ because $\widehat{\alpha}$ is well defined in a neighborhood of $\widehat{\sigma}$. Because $\widehat{\alpha}$ is a best response to $\widehat{r}$, $\widehat{\alpha}_j = 0$ for all $j \notin T_A(\widehat{r})$. As for the $j \in T_A(\widehat{r})$, rationalizing the defender's relative allocations across the sites in $T_A(\widehat{r})$ means satisfying the marginal conditions $L_j^D \delta_j'(\widehat{r}_j)[\eta_j + \widehat{\alpha}_j(1-\eta_j)] = L_k^D \delta_k'(\widehat{r}_k)[\eta_k + \widehat{\alpha}_k(1-\eta_k)]$ for all j, $k \in T_A(\widehat{r})$. (That $\widehat{r} \in P_v - E$ means $\widehat{r}_j > 0$ for all $j \in T_A(\widehat{r})$ because this can only happen when $\widehat{r}_j$ enters T_A. And $\widehat{r}_j > 0$ implies that the marginal conditions used to pin down $\widehat{\alpha}$ must satisfy hold strictly.)

Solving these equations for $\widehat{\alpha}_j$ in terms of $\widehat{\alpha}_k$, summing the $\widehat{\alpha}_j$ over $T_A(\widehat{r})$, using that fact that the sum equals one, and rearranging terms yield the expression for $\eta_k + \widehat{\alpha}_k(1-\eta_k)$ in equation (1) above.

With $\widehat{\alpha}$ defined, $\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ is also well defined and given by $\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = \sum_{j \in T_A(\widehat{r})}[\eta_j + \widehat{\alpha}_j(1-\eta_j)] L_j^D \delta_j'(\widehat{r}_j) d\widehat{r}_j/d\sigma$ where recall $\widehat{\sigma}$ is the amount of resources allocated to defending sites in $T_A(\widehat{r})$. Substituting for $\widehat{\alpha}_j$ leaves:

$$\frac{\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})}{\partial\sigma} = \left(\sum_{k \in T_A(\widehat{r})} \frac{1}{1-\eta_k}\right)\left(\sum_{k \in T_A(\widehat{r})} \frac{1}{L_k^D \delta_k'(\widehat{r}_k)(1-\eta_k)}\right)^{-1} \sum_{k \in T_A(\widehat{r})} \frac{d\widehat{r}_j}{d\sigma}$$

But, $\sum_{k \in T_A(\widehat{r})} \widehat{r}_j = \widehat{\sigma}$ implies $\sum_{k \in T_A(\widehat{r})} d\widehat{r}_j/d\sigma = 1$ which leads to:

$$\begin{aligned}\frac{\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})}{\partial\sigma} &= \left(\sum_{k \in T_A(\widehat{r})} \frac{1}{1-\eta_k}\right)\left(\sum_{k \in T_A(\widehat{r})} \frac{1}{L_k^D \delta_k'(\widehat{r}_k)(1-\eta_k)}\right)^{-1} \\ &= [\eta_k + \widehat{\alpha}_k(1-\eta_k)]\, L_k^D \delta_k'(\widehat{r}_k)\end{aligned}$$

for any $\widehat{r} \in P_v - E$ and $k \in T_A(\widehat{r})$.

Now consider a $\widehat{r} = P_v(\widehat{\sigma})$ in E but not an endpoint. Because sites enter T_A at $\widehat{\sigma}$, the summations in the expression for $\partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$ for $\sigma \in (\widehat{\sigma}, \widehat{\sigma}+\varepsilon)$ and for $\sigma \in (\widehat{\sigma}, \widehat{\sigma}-\varepsilon)$ are taken over different sets. The continuity of the individual terms will therefore not ensure the continuity of the sums, and $\partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$ $\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma^- \equiv \lim_{\sigma\to\widehat{\sigma}^-} \partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$ will generally not equal $\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma^+ \equiv \lim_{\sigma\to\widehat{\sigma}^+} \partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$.

As before, finesse this issue by filling the gap between these limits by defining Y to be the correspondence from points on P_v into $\mathbb{R}$ where, except at the endpoints of P_v, $Y(\widehat{\sigma})$ is the closed interval between $\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma^-$ and $\partial L_S^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma^+$. (This interval is simply the single point at all $\widehat{r} \in P_v - E$.) At the endpoints, define $Y(\underline{\sigma}) \equiv \partial L_S^D(\underline{\sigma}, \widehat{\alpha})/\partial\sigma^+$ and $Y(\overline{\sigma}) \equiv \partial L_S^D(\overline{\sigma}, \widehat{\alpha})/\partial\sigma^-$.

Turning to the marginal loss from investing slightly less against the non-strategic threat at $\widehat{r}(\widehat{\sigma}) \in P_v - E$, $\partial L_v^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma \equiv \sum_{j\notin T_A(r)} \eta_j L_j^D \delta_j'(\widehat{r}_j\,(\widehat{\sigma}))d\widehat{r}_j/d\sigma$ where the total resources dedicated to non-strategic defense are $\sum_{j\notin T_A(r)} \widehat{r}_j = R - \widehat{\sigma}$. By construction, $\eta_j L_j^D \delta_j'(\widehat{r}_j) = \eta_k L_k^D \delta_k'(\widehat{r}_k)$ whenever $r_j > 0$, $r_k > 0$, and $j, k \notin T_A(\widehat{r})$. This and the fact that $\sum_{j\notin T_A(r)} d\widehat{r}_j/d\sigma = -1$ give $\partial L_v^D(\sigma, \widehat{\alpha})/\partial\sigma = -\eta_k L_k^D \delta_k'(\widehat{r}_k)$ for any $r_k > 0$ with $k \notin T_A(\widehat{r})$. This expression for $\partial L_v^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ does not involve summations and we can simply take $\partial L_v^D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ at $\widehat{r}(\widehat{\sigma}) \in E$ to be $\eta_k L_k^D \delta_k'(\widehat{r}_k)$ for any $\widehat{r}_k > 0$ with $k \notin T_A(\widehat{r})$.

The correspondence $Y(\sigma)$ representing the marginal gain of investing more in strategic defense and the marginal gain from dedicating more to non-strategic defense $-\partial L_v^D(\sigma, \widehat{\alpha})/\partial\sigma$ intersect at most once. Observe first that the $-\partial L_v^D(\sigma, \widehat{\alpha})/\partial\sigma = \eta_k L_k^D \delta_k'(\widehat{r}_k)$ is strictly decreasing in σ because $d\widehat{r}_k/d\sigma < 0$ and $\delta_k'' > 0$ for all k. The correspondence Y is strictly increasing, i.e., if $\sigma' > \sigma$, then $\min Y(\sigma') > \max Y(\sigma)$. Again, $\delta_k'' > 0$ for all k implies that that all of the terms in the second summation in the expression for $\partial L_S^D(\sigma, \widehat{\alpha})/\partial\sigma$ are decreasing. These terms are also negative so summing over more terms also reduces the sum. Hence, the inverse of the sum is increasing. The first summation is also weakly increasing as it jumps up when more sites enter T_A. Because Y is increasing and $-\partial L_v^D(\sigma, \widehat{\alpha})/\partial\sigma$ is decreasing, they intersect in at most one point along P_v.

If this intersection occurs at $\widehat{r} = P_v(\widehat{\sigma}) \in P_v - E$, then $\widehat{r}$ is an equilibrium allocation as there exists an α, namely, $\widehat{\alpha}$, such that $(\widehat{r}, \widehat{\alpha})$ is an equilibrium. By construction, $\widehat{\alpha}$ is a best response to $\widehat{r}$. And, by construction, $\widehat{r}_j$ is a best response to $\widehat{\alpha}$ across the sites in $T_A(\widehat{r})$ and in $N - T_A(\widehat{r})$. Further, $\partial L^D_S(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = -\partial L^D_v(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ implies $[\eta_k + \widehat{\alpha}_k(1 - \eta_k)]\, L^D_k \delta'_k(\widehat{r}_k) = \eta_j L^D_j \delta'_j(\widehat{r}_j)$ for any $\widehat{r}_k > 0$ with $k \in T_A(\widehat{r})$ and $\widehat{r}_j > 0$ with $j \notin T_A(\widehat{r})$.

If Y and $-\partial L^D_v(\sigma, \widehat{\alpha})/\partial\sigma$ intersect at some $\widehat{r} = P_v(\widehat{\sigma}) \in E$, the marginal conditions between sites in $T_A(r')$ and in $N - T_A(r')$ cannot be satisfied at any $r' = \widehat{r}$. The defender could always profitably reallocate its resources between the strategic and non-strategic threats, so these points cannot be equilibrium allocations. This and the fact that an equilibrium exists means that $\widehat{r}$ is an equilibrium allocation. And, clearly, if Y and $-\partial L^D_v(\sigma, \widehat{\alpha})/\partial\sigma$ do not intersect, the existence of an equilibrium implies one endpoint or the other is the unique equilibrium.[24] ■

Proof of Proposition 4: The proof closely parallels that of Proposition 3 and will only be sketched. As before, an equilibrium in which the defender is playing a pure strategy is sure to exist. It is also clear that any equilibrium allocation must lie on P_g. If $r \notin P_g$, then there is a $k \notin T_\gamma(r) \cup T_\tau(r)$ such that $r_k > 0$. But, this means that the defender can profitably deviate from r by allocating r_k to sites in $T_\gamma(r) \cup T_\tau(r)$.

To characterize the equilibrium allocation on P_g and demonstrate that it is unique, take E to be the two endpoints, $P_g(\underline{\sigma})$ and $P_g(\overline{\sigma})$, and the set of points points where sites leave T_γ or enter T_τ. Then the partial derivatives $\partial^D_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ and $\partial^D_\tau(\sigma, \widehat{\phi})/\partial\sigma$are well defined at any $r(\sigma) \in P_g - E$ because the $\widehat{\alpha}$ and $\widehat{\phi}$ are well defined with $\widehat{\alpha}_j = 0$ for $j \notin T_\gamma(r(\sigma))$, $\widehat{\alpha}_j = \left(\sum_{n \in T_\gamma(r(\sigma))} L^D_j \delta'_j(r_j)/[L^D_n \delta'_n(r_n)]\right)^{-1}$ for $j \in T_\gamma(r(\sigma))$, $\widehat{\phi}_j = 0$ for $j \notin T_\tau(r(\sigma))$, and $\widehat{\phi}_j = \left(\sum_{n \in T_\tau(r(\sigma))} L^D_j \delta'_j(r_j)/[L^D_n \delta'_n(r_n)]\right)^{-1}$ for $j \in T_\tau(r(\sigma))$. Substituting for $\widehat{\alpha}$ and $\widehat{\phi}$ in

---

[24] When $\widehat{r} \in E$ along P_v, the proof of Proposition 2 constructed the best response to $\widehat{r}$. That approach relied on $\widehat{r}_e = 0$ when a stie entered $T_A(\widehat{r})$ and, therefore, that the marginal conditions involving $\widehat{r}$ were slack. However, $\widehat{r}_e$ need not be zero when it becomes a strategic target, so the marginal conditions must hold strictly and the constructive approach taken in the proof of Proposition 2 does not go through.

the expressions for the derivatives gives:

$$
\begin{aligned}
\frac{\partial^D_\gamma(\sigma,\widehat{\alpha})}{\partial\sigma} &= \mu \sum_{n\in T_\gamma(r)} \widehat{\alpha}_n L^D_n \delta'_n(r_n)\frac{dr_n}{d\sigma} \\
&= -\mu\left(\sum_{n\in T_\gamma(r)} \frac{1}{L^D_n\delta'_n(r_n)}\right)^{-1} \\
&= -\mu\widehat{\alpha}_n L^D_n\delta'_n(r_n)
\end{aligned}
$$

and

$$
\begin{aligned}
\frac{\partial^D_\tau(\sigma,\widehat{\phi})}{\partial\sigma} &= (1-\mu)\left(\sum_{n\in T_\gamma(r)} \frac{1}{L^D_n\delta'_n(r_n)}\right)^{-1} \\
&= (1-\mu)\widehat{\phi}_n L^D_n\delta'_n(r_n)
\end{aligned}
$$

for any $r \in P_g - E$ where, recall, $\sum_{n\in T_\gamma(r)} r_n = R - \sigma$ and $\sum_{n\in T_\tau(r)} r_n = \sigma$.

The correspondence $Z_\gamma(\sigma) = \{z : z \in \left[\partial L^D_\gamma(\sigma,\widehat{\alpha})/\partial\sigma^+, \partial L^D_\gamma(\sigma,\widehat{\alpha})/\partial\sigma^-\right]\}$ is strictly decreasing in σ whereas $Z_\tau(\sigma) = \{z : z \in [\partial L^D_\gamma(\sigma,\widehat{\phi})/\partial\sigma^-, \partial L^D_\gamma(\sigma,\widehat{\phi})/\partial\sigma^+]\}$ is strictly increasing. Hence, they intersect at most at one $\widehat{\sigma}$ along P_g.

If they do intersect, $\widehat{r} = P_g(\widehat{\sigma})$ is the unique equilibrium allocation. If $\widehat{r} = P_g(\widehat{\sigma}) \notin E$, then $\widehat{r}$ is clearly an equilibrium allocation. By construction, $\widehat{\alpha}$ is a best response to $\widehat{r}$ and the defender has no incentive to reallocate resources across $T_\gamma(\widehat{r})$ given $\widehat{\alpha}$. Similarly, $\widehat{\phi}$ is a best response to $\widehat{r}$ and the defender has no incentive to reallocate across sites in $T_\tau(\widehat{r})$. And, finally, the defender cannot profitably reallocate resources between sites in $T_\gamma(\widehat{r})$ and $T_\tau(\widehat{r})$ because $-\partial L^D_\gamma(\sigma,\widehat{\alpha})/\partial\sigma = \partial L^D_\tau(\sigma,\widehat{\alpha})/\partial\sigma$ implies $\mu\widehat{\alpha}_j L^D_j\delta'_j(r_j) = (1-\mu)\widehat{\phi}_n L^D_n\delta'_n(r_n)$ for all $j \in T_\gamma(\widehat{r})$ and $k \in T_\tau(\widehat{r})$.

Now suppose $\widehat{r} = P_g(\widehat{\sigma}) \in E$. Then $\mu\widehat{\alpha}_j L^D_j\delta'_j(r_j) > (1-\mu)\widehat{\phi}_n L^D_n\delta'_n(r_n)$ for any $\sigma < \widehat{\sigma}$, $\widehat{\alpha}$ that rationalizes $r_j(\sigma)$ across $T_\gamma(r(\sigma))$, and $\widehat{\phi}$ that rationalizes $r_k(\sigma)$ across sites in $T_\tau(r(\sigma))$. The defender, therefore, could profitably reallocate resources by investing more in protecting against τ. Conversely, $\mu\widehat{\alpha}_j L^D_j\delta'_j(r_j) < (1-\mu)\widehat{\phi}_n L^D_n\delta'_n(r_n)$ for any $\sigma > \widehat{\sigma}$, any $\widehat{\alpha}$ that rationalizes $r_j(\sigma)$ across $T_\gamma(r(\sigma))$, and $\widehat{\phi}$ that rationalizes $r_k(\sigma)$ across sites in $T_\tau(r(\sigma))$. Existence then ensures that $\widehat{r}$ is the unique equilibrium. $\blacksquare$

References

9/11 Commission Report. 2003. "Final Report of the National Commission on Terrorist Attacks upon the Unites States," New York: W.W. Norton.

Belke, James C. 2000. "Chemical Accident Risks in U.S. Industry: A Preliminary Analysis of Accident Risk Data from U.S. Hazardous Chemical Facilities." Washington, D.C.: U.S. Environmental Protection Agency, September 25.

Bueno de Mesquita, Ethan. 2005. "Politics and the Suboptimal Provision of Counterterror." Manuscript, Department of Political Science, Washington University, St. Louis. Available at: http://bdm.wustl.edu/papers.html.

Chyba, Christopher F. 2002. "Toward Biological Security," *Foreign Affairs* 81(May/June):122-136.

Chyba, Christopher F. 2001. "Biological Terrorism and Public Health," *Survival* 43(Spring): 93-106.

Chertoff, Michael. 2005a. "Remarks for Secretary Michael Chertoff U.S. Department of Homeland Security George Washington University Homeland Security Policy Institute," Washington, D.C. Department of Homeland Security, March 16. Available at: http://www.dhs.gov/dhspublic/display?content=4391

Chertoff, Michael. 2005b. "Secretary Michael Chertoff U.S. Department of Homeland Security Second Stage Review Remarks." Washington, D.C. Department of Homeland Security, July 13. Available at: http://www.dhs.gov/dhspublic/display?content=4597.

Dwyer, Jim and Eric Lipton. 2005. "3-Year Federal Study of 9/11 Urges Rules for Safer Towers," *New York Times*, June 22, A1.

Enders, Walter and Todd Sandler. 1993. "The Effectiveness of Anti-Terrorism Policies: Vector-Autoregression-Intervention Analysis." *American Political Science Review* 87:829–844.

Enders, Walter and Todd Sandler. 2002. "Patterns of Transnational Terrorism, 1970–99: Alternative Time Series Estimates." *International Studies Quarterly* 46:145–165.

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

Enders, Walter and Todd Sandler. 2004. "What Do We Know About the Substitution Effect in Transnational Terrorism?" In Andrew Silke and G. Ilardi, eds., *Researching Terrorism Trends, Achievments, Failures.* London: Frank Cass.

Enders, Walter, Todd Sandler and Jon Cauley. 1990. "UN Conventions, Technology, and Retaliation in the Fight Against Terrorism: An Econometric Evaluation." *Terrorism and Political Violence* 2:83–105.

Ervin, Clark Kent. 2005. "A To-Do List for Chertoff," *Washington Post*, February 7, A21.

Flynn, Stephen. 2004. *America the Vulnerable.* New York: Harper Collins.

Fudenberg, Drew and Jean Tirole. 1991. *Game Theory.* Cambridge, MA.: MIT Press.

GAO. 2004. "Homeland Security: Federal Action Needed to Address Security Challenges at Chemical Facilities" (GAO-04-482T). Washington, D.C.: General Accounting Office.

Guikema, Seth and Elisabeth Paté-Cornell. 2002. "Probabilistic Modeling of Terrorist Threats: A Systems Analysis Approach to Setting Priorities Among Counter-

Measures," *Military Operations Research* 7(4):5-19.

Heyman, David and James Jay Carafano. 2004. "DHS 2.0: Rethinking the Department of Homeland Security," Washington, D.C.: Center for Strategic and International Studies. Available at: http://www.csis.org/hs/041213_dhsv2.pdf.

Im, Eric Iksoon, Jon Cauley and Todd Sandler. 1987. "Cycles and Substitutions in Terrorist Activities: A Spectral Approach." *Kyklos* 40(2):238–255.

Kettl, Donald F. 2004. "Overview." In Donald Kettl, *The Department of Homeland Security's First Year: A Report Card*. New York: Century Foundation, 1-28.

Khademain, Anne M. 2004. "Strengthening State and Local Terrorism Prevention and Response." In Donald Kettl, ed., *The Department of Homeland Security's First Year: A Report Card*. New York: The Century Foundation, 97-118.

Kocieniewski, David. 2005. "Row of Loosely Guarded Targets Lies Just Outside New York City," *New York Times*, A1, May 9.

Lakdawalla, Darius and Eric Talley. 2005. "Optimal Liability for Terrorism." Manusript, USC Law School. Available at: http://www.gtcenter.org/Downloads/WLE/Talley104.pdf.

Los Angeles Times. 2005. "Port Insecurity," *Los Angeles Times*, February 26, B18.

Mas-Colell, Andreu, Michael D. Winston, and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford University Press.

Mayer-Schönberger, Viktor. 2003. "Emergency Communications: The Quest for Interoperability in the United States and Europe." In Arnold M. Howitt and Robyn L. Pangi, eds., *Countering Terrorism*. Cambridge, MA.: MIT Press.

Moteff, John. 2004. "Risk Management and Critical Infrastructure Protection: Assessing, Integrating, and Managing Threats, Vulnerablities and Consequences." Washington, D.C.: Congressional Research Service.

New York Times. 2005a. "Editorial: Follow the Port Security Money," *New York Times*, February 28, A18.

New York Times. 2005b. "Editorial: Security Loses; Pork Wins," *New York Times*, July 14, A24.

O'Beirne, Kate. 2003. "Introducing Pork-Barrel Homeland Security: A Little Here, A Lot There," *National Review,* August 11, 2003.

O'Hanlon, Michael *et. al.* 2003. *Protecting the American Homeland: One Year on.* Washington, D.C.: Brookings.

OMB. 2003. "2003 Report to Congress on Combating Terrorism." Washington, D.C.: Office of Management and Budget, September. Available at: http://www.whitehouse.gov/omb/legislative/index.html.

OMB. 2004. "Analytical Perspectives, Budget of the United States Government, Fiscal Year 2005." Washington, D.C.: Office of Management and Budget.

OMB. 2005. "Analytical Perspectives, Budget of the United States Government, Fiscal Year 2006." Washington, D.C.: Office of Management and Budget.

Rosendorff, Peter and Todd Sandler. 2004. "Too Much of a Good Thing? The Proactive Response Dilemma," *Journal of Conflict Resolution*, 48(October): 657-671.

Sandler, Todd and Daniel G. Arce M. 2003. "Terrorism and Game Theory," *Simulation and Gaming* 3434(3):319-37.

Sandler, Todd and Harvey E. Lapan. 1988. "The Calculus of Dissent: An Analysis of Terrorists' Choice of Targets." *Snythese* 76(2):245–261.

Sandler, Todd and Kevin Siqueira. 2003. "Global Terrorism: Deterrence versus Preemption." Manuscript, School of International Relations, University of Southern California.

USCG 2003. *Federal Register* 68(126), July 1, 39245.

White House. 2002, "National Strategy for Homeland Security." Washington, D.C.: White House, July. Available at: http://www.whitehouse.gov/homeland/book/index.html.

White House. 2003. "National Strategy for the Physical Protection of Critical Infrastructures and Key Assets." Washington, D.C.: White House, February. Available at: http://www.whitehouse.gov/pcipb/physical.html

Washington Post. 2005. "Editorial: Risk vs. Pork," *Washington Post*, May 17, A20.