Defending Against Terrorist Attacks with Limited Resources*

Robert Powell

January 15, 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, Ron Hassner, William Kastenberg, Lance Kim, Michael May, Gerard Padro-i-Miquel, Per Peterson, Larry Samuelson, Todd Sandler, 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 also allocate resources to border defense, intelligence, or counter-terrorist operations which, if successful, protect all of the 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, an algorithm or approach to finding the optimal allocations, and relevant comparative statics. These characterizations 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). 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 shortly after taking office, "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. As the *National Strategy for Homeland Security* emphasizes, "One fact dominates all homeland security threat assessments, terrorists are strategic actors" (White House 2002, 7). They try to strike where the defense is weak and the expected gains are high. Protecting one site may therefore merely 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]

[1] For statistical evidence of these "substitution" effects see Enders and Sandler 2004.

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 tries 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 most attractive target, i.e., the site offering the highest expected payoff. The defender, therefore, should invest in hardening or protecting this site. But the more the defender spends on this site, the lower the attacker's payoff to striking this site and the less attractive it becomes. Eventually, going after this site will be no more attractive than striking what was originally the second most attractive target. At this point, the defender must invest in hardening both of these sites so that neither is a more appealing target than the other. The more the defender spends on these two sites, the lower the payoff to attacking them and the less attractive they become. Eventually, attacking either of these two sites will offer an expected payoff no higher than going after what was originally the third most attractive site. From here on, the defender will have to invest in all three of those sites so that no one of them is more attractive than any otherf. The defender continues to allocate the rest of its budget in this way by spending on and hardening more and more sites so as to make the most attractive targets as unattractive as possible.

The second setting extends the baseline case to allow the defender to allocate resources to border defense, counter-terrorist operations, and intelligence in addition to spending on specific sites. The key idea here is that investing in 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 virus like SARS or a highly

[2] The present analysis focuses on understanding the first-best outcome, i.e., the optimal allocation for a single unitary defender. Other analyses of the basic resource allocation problem include, e.g., Bier, Oliveros, and Samuelson (2005) and Brown *et. al.* (2005). Examples centering on second-best outcomes due to the distorting effects of multiple actors with different interests include Bueno de Mesquita 2005 and Rosendorff and Sandler 2004.

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 strategic and non-strategic components? 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 show that although the extensions to the baseline case are substantively quite diverse, the same basic approach can be used to determine the optimal allocation. The key to this commonality is that even though the baseline game and its extensions are generally not zero-sum games, the defender's unique equilibrium strategy in the baseline model is to minmax the attacker, i.e., the defender should allocate its resources so as to minimize the attacker's payoff. This is a surprising result. As is well known, a pair of strategies is an equilibrium of a two-actor, zero-sum game if and only if the actor's are minmaxing each other. But minmaxing is usually not an equilibrium strategy in nonzero-sum games. In the extensions, the defender does try 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 some 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).[3] Secretary Chertoff underscored the same point shortly after taking office, "Risk management must 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 it typically does not treat adversaries as 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 treating an adversary as fully strategic can lead to a significant misallocation of 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. Let r_1 represent the resources dedicated to protecting the first site and r_2 be the resources allocated to the second where $r_1+r_2 = R$. Take $\delta_j(r_j)$ to be the probability that an attack on site j succeeds if the defender allocates r_j to protecting it. Finally, the probability of an attack on site 1 is α and the probability of an attack on site 2 is $1 - \alpha$. 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.

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

[3] More direct criticisms include, e.g., the 9/11 Public Discourse Project's final report (2005) and O'Beirne 2003.

threat.[4] Consequences are the outcomes of a successful attack and correspond to L_1 and L_2 in the model. The chances of a successful attack, δ_1 and δ_2, reflect a site's 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 less to spend 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 the 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

[4] Moteff (2004) provides an overview of risk-based managment in relation to the Department of Homeland Security's critical infrastructure program. Also see the discussion the of the U.S. Army's approach to risk managment in Brown *et. al* (2005). The draft *National Infrastructure Protection Plan* of November 2005 describes the Department of Homeland Security's evolving method.

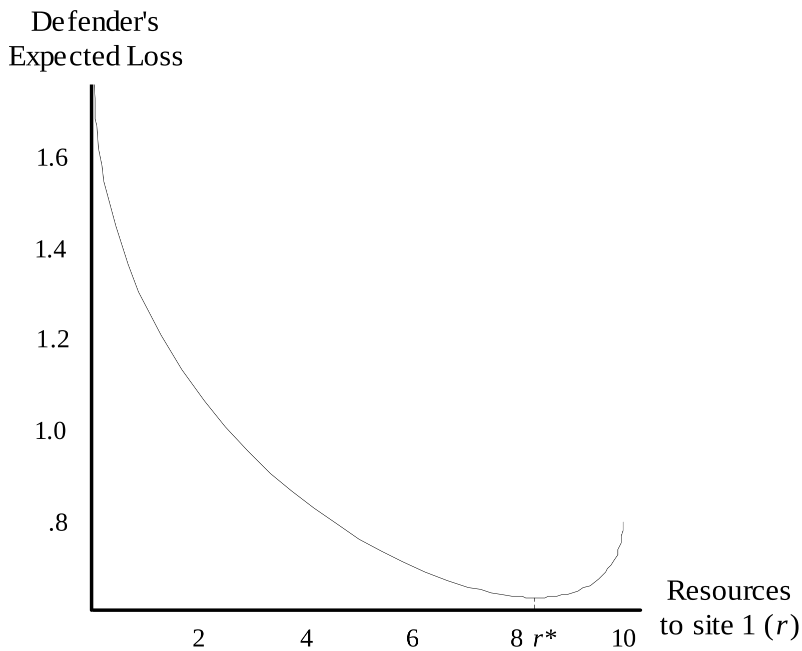

(a) Non-strategic terrorists.

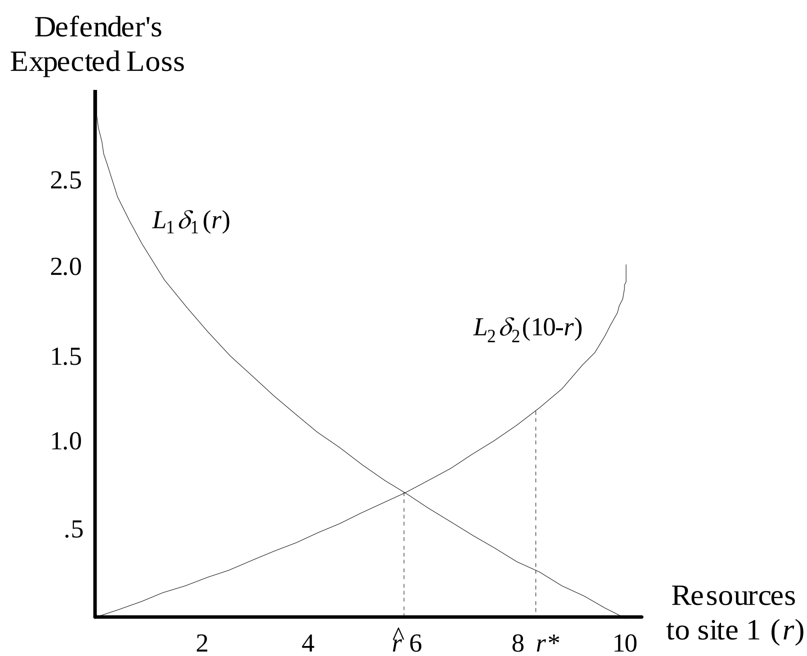

(b) Strategic terrorists.

Figure 1: Strategic versus non-strategic terrorists.

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.

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 becomes the most attractive target because it offers the attacker its highest expected payoff. If the defender spends more than $\widehat{r}$ on site 1, say r^*, site 2 becomes the most appealing target 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 on which the allocation r^* was based. 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. 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 the terrorist group attacks site j with $\alpha_1 + \cdots + \alpha_N = 1$.

To specify the player's payoffs, suppose that the defender suffers a loss $L_j \geq 0$ if site j is successfully attacked. Allocating more resources to a site lowers the probability of a successful attack but with diminishing effect. In symbols, the probability that an attack on site j succeeds is $\delta_j(r_j)$ with $\delta_j'(r_j) < 0$ and $\delta_j''(r_j) > 0$ as long as this site is imperfectly defended (i.e., as long as $\delta_j(r_j) > 0$).[5] Thus, the defender's expected loss if the terrorists strike j is $L_j\delta_j(r_j)$, and its expected loss if it plays strategy r and the terrorist group plays α is $L(r, \alpha) \equiv \sum_{j=1}^{N} \alpha_j L_j \delta_j(r_j)$. As for the attacker's payoffs, let $A_j \geq 0$ be the attacker's gain from successfully striking site j. Then the terrorist group's expected gain is $A(r, \alpha) \equiv \sum_{j=1}^{N} \alpha_j A_j \delta_j(r_j)$.

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 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, $A_1\delta_1(0) > A_2\delta_2(0) > \cdots > A_N\delta_N(0)$.[6] Then the most attractive target, i.e., the site offering the attacker the highest expected payoff, is site 1 because $A_1\delta_1(0) = \max\{A_1\delta_1(0), ..., A_N\delta_N(0)\}$. The defender, therefore, should begin

[5] 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 $r > \overline{r}$.

[6] 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.

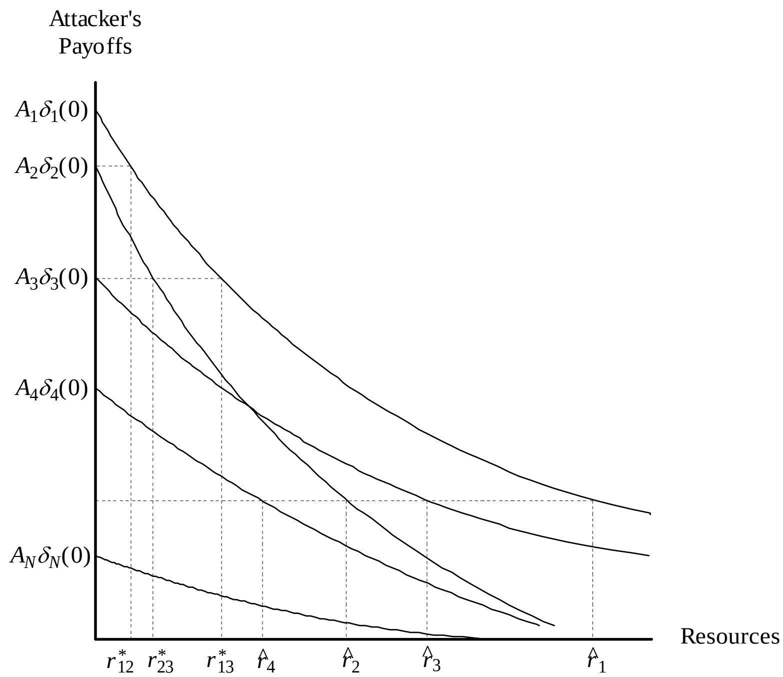

Figure 2: The optimal allocation.

by allocating its resources to this site. As the defender devotes more resources to this site, 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 $A_1\delta_1(r^*_{12}) = A_2\delta_2(0) = \max\{A_1\delta_1(r^*_{12}),\ A_2\delta_2(0),\ ...,\ A_N\delta_N(0)\}$. From this point on, the defender must allocate resources to both sites 1 and 2 in order to keep $A_1\delta_1(r_1) = A_2\delta_2(r_2)$. Otherwise the attacker would strike either site 1 or 2 depending on which offered the highest gain, and the resources 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 these sites are no more attractive than site 3, 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 allocation and the algorithm leading to it produce what may at first seem to be counter-intuitive results. Suppose a site becomes uniformly more difficult to defend at the margin, i.e., the marginal return to investing in protecting the site decreases. Formally, $\delta_j(r)$ shifts to $d_j(r)$ where $d_j(0) = \delta_j(0)$ and $d'_j(r) > \delta'_j(r)$ for all r. One might expect that this change would induce the defender to spend less on the site that is harder to secure, because this investment now does less to reduce the expected loss if the site is attacked. But this is not the case. As a site becomes more difficult to protect, the defender allocates more to that site and necessarily, less to the other sites. The attacker also becomes more likely to attack that site as shown below.

Figure 3 illustrates these results for the case in which site 3 becomes harder to protect and the original equilibrium allocation is $\widehat{r} = (\widehat{r}_1, \widehat{r}_2, \widehat{r}_3, \widehat{r}_4)$. If the original allocation remains unchanged, striking site 3 offers the attacker its highest payoff. But if the attacker is sure to go after this site, protecting any other site is a waste of resources and $\widehat{r}$ is not an equilibrium allocation. Relative to $\widehat{r}$, the defender needs to spend more on site 3 and less on the other sites in order to equalize the attacker's payoffs across these sites. This results in allocation r^* in Figure 3.

Stating the results more formally, the defender's optimal allocation minmaxes the attacker by minimizing the attacker's maximum payoff. That is, the algorithm solves $\min\{\max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}: r_1 + \cdots + r_N \leq R\}$. 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 the (generically) unique equilibrium.[7] The following proposition summarizes the results. (Proofs of the propositions are in the appendix.)

PROPOSITION 1: *In equilibrium, the defender minmaxes the attacker, i.e., plays the unique pure strategy that makes the most attractive targets as unattractive as possible by minimizing* $\max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\}$. *This optimal allocation is given by the*

[7] If the attacker at $\widehat{r}$ is indifferent between attacking an undefended site k ($r_k = 0$) and a defended site j ($r_j > 0$), the game has multiple equilibria. The defender's strategy in all of them is still $\widehat{r}$ but now there will be multiple $\widehat{\alpha}$'s for which $(\widehat{r}, \widehat{\alpha})$ is an equilibrium.

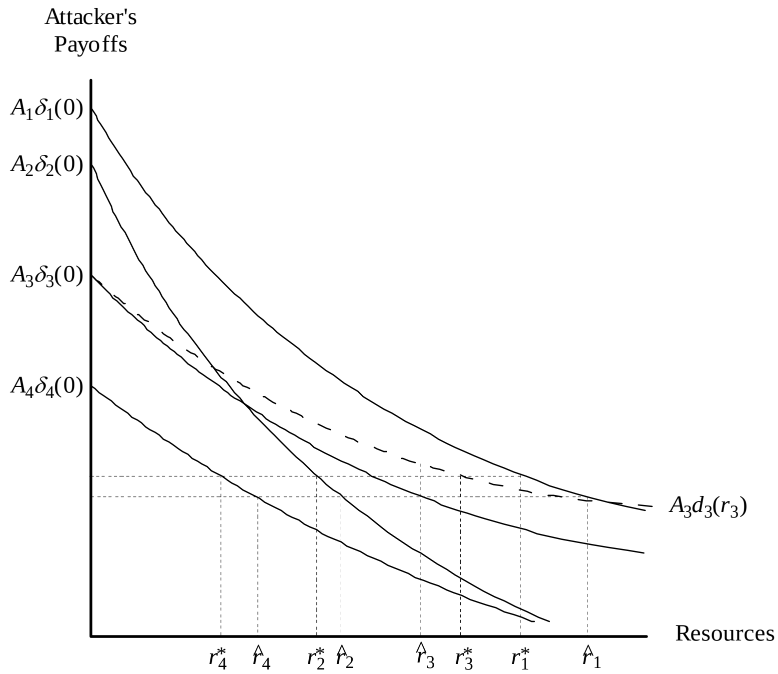

Figure 3: When a site becomes harder to defend.

algorithm above.[8]

This is a surprising result. Minmax strategies are equilibrium strategies in two-actor, zero-sum games. The players' payoffs are diametrically opposed in these games so that anything that increases one actor's payoff decreases the other's. As a result, holding an adversary's payoff down as low as possible by minmaxing it also serves to maximize a player's own payoff and mutual minmaxing is an equilibrium. But the allocation problem described above is generally not zero sum, and minmax strategies are usually not equilibrium strategies in nonzero sum games.

To see that the baseline allocation game is typically not zero sum and appreciate the significance of the zero-sum issue, it is useful to begin by noting the close relationship between this game and "Colonel Blotto" games which have long been used to study

[8] That $\delta_j'' > 0$ for all j ensures the defender never mixes in equilibrium. If $\delta_j'' \geq 0$, the allocation characterized in the proposition remains the unique, pure-strategy equilibrium allocation. But there may be equilibria in which the defender mixes.

force allocation issues (e.g., Tukey 1949, Blackett 1958, Shubik and Weber 1981). The classic blotto game is a two-actor, zero-sum game in which players I and II compete on N independent battlefields. Player I has F_I units of force to distribute among these battlefields, and player II has F_{II}. If I allocates r_I^k to battlefield k and II allocates r_{II}^k, I receives a payoff of $p_k(r_I^k, r_{II}^k)$ in the ensuing battle and II gets $-p_k(r_I^k, r_{II}^k)$. The players' total payoffs are the sums of their battlefield payoffs. When the basic allocation game is zero sum, it is also a blotto game. If, for example, the defender's loss if a site is destroyed equals the attacker's gain, i.e., if $L_j = A_j$ for all sites j, then the defender's loss at j if it allocates r_j and the attacker "allocates" α_j is $L_j\delta_j(r_j)\alpha_j = A_j\delta_j(r_j)\alpha_j$.

But the basic allocation game is generally not zero sum and therefore not a blotto game. The defender and attacker's payoffs are not diametrically opposed if, for example, they rank the sites differently. Suppose more concretely $A_1\delta_1(0) > A_2\delta_2(0)$ but $L_1\delta_1(0) < L_2\delta_2(0)$, then attacking site 1 maximizes the attacker's payoff but does not minimize the defender's. Indeed, even if the the attacker and defender rank the sites in the same order (i.e., $L_i > L_j \Leftrightarrow A_i > A_j$), the game will not be zero sum unless the defender and attacker put precisely the same relative weights on the sites (i.e., $L_i/L_j = A_i/A_j$).

Relaxing the zero-sum assumption is important for two reasons. First, assuming that the terrorist's gain equals the defender's loss may not sound very demanding intuitively. But its formal equivalent – that the actors' put the same relative weights on the sites – seems much more demanding and less substantively appealing. Still more fundamentally, some situations simply cannot be modeled as zero-sum games. Indeed, neither the case in which threats have non-strategic components nor the case in which the defender is uncertain of the attacker's payoffs are zero sum.

Why does the defender minmax the attacker in the basic allocation game even though it is not a zero sum? Observe that regardless of how the defender ranks the sites, the optimal use of the defender's resources is to protect only those sites that will be attacked. These are the sites offering offer the attacker the highest expected payoff and allocating resources solely to them has the effect of minimizing the attacker's maximum payoff.

The emphasis in this analysis is on the defender's allocation. But it is useful to specify the attacker's equilibrium strategy $\widehat{\alpha}$ given the optimal allocation $\widehat{r}$. Let $T_A(\widehat{r})$ be the set of sites gives the attacker its highest payoff at $\widehat{r}$. That is, $T_A(\widehat{r})$ is the set of pure-strategy, best replies to $\widehat{r}$: $T_A(\widehat{r}) = \{k : A_k\delta_k(\widehat{r}_k) = \max\{A_1\delta_1(\widehat{r}_1), ..., A_N\delta_N(\widehat{r}_N)\}\}$. Then the attacker's payoff to striking any site in T_A is the same as its payoff to striking any other site in T_A and higher than its payoff to going after any site not in T_A. This implies that the probability the attacker strikes any site $j \notin T_A$ is zero, $\widehat{\alpha}_j = 0$.

As for the probabilities of striking sites in T_A, $\widehat{r}$ must minimize the defender's expected loss $L(r, \widehat{\alpha})$ against the attacker's strategy $\widehat{\alpha}$ if $\widehat{r}$ is to be an equilibrium strategy. This requires the defender's marginal returns to investing in these sites be the same: $\partial L/\partial r_j = \partial L/\partial r_i$ for all $i, j \in T_A$.[9] These first-order conditions imply $\widehat{\alpha}_1 L_1 \delta_1'(r_1) = \widehat{\alpha}_j L_j \delta_j'(r_j)$ for $j \in T_A$. Adding the $\widehat{\alpha}_j$'s for $j \in T_A$, and using the fact that the probabilities of attack add up to one give $\widehat{\alpha}_j = \left(\sum_{n \in T_A(\widehat{r})} L_j\delta_j'/L_n\delta_n'\right)^{-1}$. Corollary 1 summarizes these results:

COROLLARY 1: *The attacker's strategy in the generically unique equilibrium is to attack site* $j \in T_A(\widehat{r})$ *with probability* $\widehat{\alpha}_j = \left(\sum_{n \in T_A(\widehat{r})} L_j\delta_j'(\widehat{r}_j)/L_n\delta_n'(\widehat{r}_n)\right)^{-1}$.

It follows that the probability of an attack on a site increases as that site becomes harder to defend. As this happens, the defender as shown above spends more on that site and necessarily less on the other sites. Because this site is both harder to defend and more is being spent on it, the marginal effect any additional investment has on reducing the probability that an attack on this site will succeed is lower. Given this lower marginal effect, the defender's larger investment in this site will only be optimal if the site is more likely to be attacked which raises the marginal expected gain from trying to protect this site enough to rationalize the defender's higher level of spending.[10]

---

[9] These equalities presume $\widehat{r}_i > 0$ and $\widehat{r}_j > 0$ for all $i, j \in T_A$. If $\widehat{r}_k = 0$ for some $k \in T$, then α_k may be positive and the marginal condition needed to rationalize $\widehat{r}$ is that the defender cannot profitably shift resources from j to k, i.e., $\partial L/\partial r_j \geq \partial L/\partial r_k$.

[10] Formally, suppose the initial equilibrium allocation is r^* and δ_1 shifts to d_1 where $d_1'(r_1) > \delta_1'(r_1)$ and $d_1(0) = \delta_1(0)$. Then at r^* and the new equilibrium allocation $\widehat{r}$ define $z_k^* \equiv L_1\delta_1'(r_1^*)/[L_k\delta_k'(r_k^*)]$ and $\widehat{z}_k \equiv L_1 d_1'(\widehat{z}_1)/[L_k\delta_k'(\widehat{r}_k)]$ for $k > 1$, and observe that $\alpha_k^* = z_k^*\alpha_1^*$, $\widehat{\alpha}_k^* = \widehat{z}_k\widehat{\alpha}_k$, and $\widehat{z}_k < z_k^*$. Summing the probabilities gives $1 = \alpha_1^*\left(1 + \sum_{k>1} z_k^*\right) = \widehat{\alpha}_1\left(1 + \sum_{k>1} \widehat{z}_k\right)$ which implies $\widehat{\alpha}_1 > \alpha_1^*$.

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 operations 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* (White House 2002). This section examines the trade off between defending individual sites as in the basic framework above and allocating resources to border defense, intelligence, or counter-terrorist operations.

The key to introducing this trade off into the baseline model is that investing in these activities 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).[11] These activities also present a puzzle.

Consider border defense. Given the difficulty of securing the borders, it might seem that the marginal gain from investing more in border defense is likely to be very small and less than the marginal gain from spending that money on hardening, say, a nuclear power plant or a chemical facility. If so, why allocate anything to border defense? Would it not be better to devote everything to site defense? But surely spending nothing on border defense cannot be the right thing to do.

The analysis below explains why spending on activities that protect all of the sites but have a low marginal return makes sense. The defender's optimal allocation does equate the marginal gain from investing in border defense with the marginal gain from site defense. But as the analysis of the baseline case shows, any additional spending on

[11] Assuming that border defense, intelligence, and counter-terrorist operations protect all sites and, in the formalization below, protect all sites equally is clearly a tremendous simplification and at best a first approximation to be refined in future work.

site defense must be spread across multiple sites in order to keep the attacker's payoff to striking any one of these sites equal to its payoff to hitting 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 is very small and eventually equal the marginal gain to border defense even if the latter is also quite small.

The allocation algorithm leading to these results is also closely related to the algorithm that yields the optimal allocation in the baseline case. Suppose that the defender allocates all of its resources to border defense and none to site defense. (To ease the exposition, "border defense" will be used broadly to include counter-terrorist and intelligence activities as well.) Then the attacker would strike the site that offers the highest expected payoff given this allocation. If, therefore, the defender wants to hedge by dedicating some of its resources to site defense, it should spend them on that site. But the more the defender takes out of border defense and spends on that site, the harder that site becomes. Eventually attacking that site will be no more appealing than going after what was originally the second most attractive site. If at this point the defender wants to invest still more in site defense, it must devote those additional resources to protecting both of those sites. At some point the payoff to attacking either of those sites will be the same as that of attacking what was originally the third most attractive site. The defender must divide any further spending on site defense to those three targets. As the defender shifts more and more resources from border defense to site defense, it has to spread the latter over more and more sites. The optimal division of resources between border and site defense occurs when the marginal gain from allocating more to site defense is just offset by the marginal loss of spending less on border defense.

To incorporate border defense in the basic framework formally, suppose that the defender has to decide how much to direct to the N individual sites and to border defense which has a low marginal return but protects all of the sites. 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 probability an attack on j succeeds is the probability that the terrorist group penetrates the border 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 with diminishing returns ($\beta' < 0$ and $\beta'' > 0$). Formalizing the notion that it is much easier to protect one site than secure the entire border, if the defender had only had one site to protect, it would always do better by spending on that site rather than on border defense: $\partial\beta(b - r_j)\delta_j(r_j)/\partial r_j < 0$ and $\partial^2\beta(b - r_j)\delta_j(r_j)/\partial r_j^2 > 0$ for all j. The defender's and attacker's payoffs are then $L((r, b), \alpha) = \sum_{n=1}^{N} \alpha_n L_n \beta(b)\delta_n(r_n)$ and $A((r, b), \alpha) = \sum_{n=1}^{N} \alpha_n A_n \beta(b)\delta_n(r_n)$. 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.

As will be seen, the same fundamental formal approach can be used to characterize the defender's optimal allocation when the defender can invest in border defense, faces threats with non-strategic components, or is uncertain of the terrorist's priorities. The intuition underlying this approach is also straightforward. The algorithm derived in the baseline case can be used to define a curve through the space of possible allocations where movement along the curve defines the fundamental trade off the defender faces in the particular substantive case: spending more on border defense and less on site defense, more on the strategic component of the threat and less on the non-strategic component, or more to protect against one type of terrorist and less on any other. The equilibrium allocations lie along the curve where the marginal gain of spending more is just offset by the marginal loss of spending less.

In the case of border defense, take allocation $r = (r_1, ..., r_N)$ with $\sum_{j=1}^{N} r_j \leq R$ to be a point in $\mathbb{R}^N$ at which the defender is spending $\sigma = \sum_{j=1}^{N} r_j$ on site defense and $b = R - \sigma$ on border defense. Now define the curve P_B through $\mathbb{R}^N$ parameterized by $\sigma \in [0, R]$. At $P_B(\sigma)$, the defender spends σ on site defense with those resources optimally allocated against the attacker. That is, $r(\sigma)$ is the unique allocation that minimizes $M_A(r|\sigma) \equiv \max\{A_1\delta_1(r_1), ..., A_N\delta_N(r_N)\colon r_1 + \cdots + r_N \leq \sigma\}$ and can be found with the algorithm used to solve the baseline model.

We can now think of moving along P_B from $P_B(0)$ to $P_B(R)$ as the allocations that result from the defender's spending less on border defense in order to spend more on site defense while being sure to optimally allocate however much it does spend on site defense. It follows that any equilibrium allocations must lie on P_B; otherwise the resources dedicated to site defense would not be allocated optimally. Moreover, the equilibria must lie at critical points along P_B where the marginal gain from dedicating slightly more to site defense is just offset by the marginal loss of spending slightly less on border defense.

Proposition 2 formalizes these marginal conditions. 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 $S_D(\sigma, \alpha) = \sum_{j=1}^{N} \alpha_j L_j \delta_j(\widehat{r}_j(\sigma))$. Now define the attacker's "pseudo-equilibrium" strategy $\widehat{\alpha}$ against any allocation $\widehat{r}$ on P_B to be a best response to $\widehat{r}$ that also rationalizes the defender's relative spending on the sites in $T_A(\widehat{r})$ where, recall, $T_A(r)$ is set of attacker's pure-strategy, best replies to r. This implies $\widehat{\alpha}_j = 0$ for all $j \notin T_A(\widehat{r})$, $L_j\delta_j'(\widehat{r})\widehat{\alpha}_j = L_k\delta_k'(\widehat{r})\widehat{\alpha}_k$ whenever $j, k \in T_A(\widehat{r})$ with $r_j > 0$ and $r_k > 0$, and $L_j\delta_j'(\widehat{r})\widehat{\alpha}_j \leq L_k\delta_k'(\widehat{r})\widehat{\alpha}_k$ whenever $j, k \in T_A(\widehat{r})$ with $r_j \geq 0$ and $r_k = 0$. Then the equilibrium allocations lie on P_B where $\beta(R - \sigma)\partial S_D(\sigma, \widehat{\alpha})/\partial\sigma = \beta'(R - \sigma)S_D(\sigma, \widehat{\alpha})$. 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: *The defender always plays a pure strategy in equilibrium. These allocations lie on P_B where the marginal loss of investing less in border defense just offsets the marginal gain of spending more on site defense and is defined almost everywhere by 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:*

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

If, moreover, β and δ_j are log convex ($d^2 \ln \beta/db^2 > 0$ and $d^2 \ln \delta_j/d\delta_j^2 > 0$), the equilib-

rium allocation is unique.

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. 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). 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).

Recognizing these complementarities, the formula by which the Department of Homeland security makes grants to cities and urban areas has recently begun to take the overlap between natural disasters and terrorist attacks into account. Announcing the 2006 Urban Areas Security Initiative grants, Secretary Chertoff explained

> sometimes there are capabilities you need in a natural disaster that can also be relevant in a terrorism-created disaster. That can involve hurricanes, it can involve fires, which may sometimes be caused by a natural event, sometimes by a terrorist. Therefore, capabilities such as evacuation, capabilities such as the ability to move rescue aid very quickly, which may have some value in a non-terrorism type of case, will be funded as long can be a demonstrable connection to a terrorism-related crisis as well (Chertoff 2006, 7).

Although the measures taken to defend against a natural disaster and a terrorist attack may overlap, there is a fundamental difference between protecting against a strategic or non-strategic threat. Hardening a chemical facility does not increase the risk of an accident at a nuclear power plant, but it may induce a strategic actor to look for other targets and thereby increase the risk of a deliberate attack on those targets. 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 and draw heavily on the analysis of the baseline case. Suppose that the defender completely discounts the possibility of a terrorist attack and allocates all of its resources optimally against the non-strategic threat. This allocation equates the marginal gain from spending a little more on site j with the marginal gain of spending a bit more on k where these margins depend on the chances of an accident at j and k but not on the prospects of a deliberate attack. Spending a bit more or less on either site will not affect the non-strategic threat to the other site.

But equating the margins to optimize spending against a non-strategic threat means that a terrorists' payoffs will generally not be equal across the sites. Suppose that the attacker's payoff to striking site t_1 is higher than its payoff to going after any other site. Then this site is sure to be attacked at the optimal allocation against the non-strategic threat. If therefore the defender wants to hedge against a terrorist attack, it has to shift resources to t_1 and necessarily away from the other sites, doing so in a way that keeps the marginal gains of protecting those other sites against the non-strategic threat equal. As the defender allocates more resources to protecting 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. If the defender shifts still more resources away from protecting against the non-strategic threat, it must begin to distribute those 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 spending more on the strategic threat is just offset by the marginal loss of spending on the non-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 $\eta_j +$

$\alpha_j(1-\eta_j)$.[12] This implies that the defender's expected loss is $L(r,\alpha) = \sum_{j=1}^{N} L_j \delta_j(r_j)[\eta_j + \alpha_j(1-\eta_j)]$. As always, the attacker strikes sites offering the highest expected payoff, i.e., at any site in $T_A(r)$.

As for the comparative statics, anything that increases the marginal gain of defending sites not subject to deliberate attack (i.e., those $j \notin T_A(r)$) leads (weakly) to greater spending on those sites and vice versa.[13] Suppose, for example, that site k is only subject to a non-strategic threat and will be struck with probability η_k (as $\alpha_k = 0$) because the attacker's payoff to hitting this site is lower than at other sites ($k \notin T_A(r)$). If the threat η_k increases, the marginal gain to protecting against the non-strategic threat increases. As a result, the defender (weakly) shifts resources away from sites under risk of terrorist attack (i.e., those $j \in T_A(r)$) and allocates those resources to guarding against the non-strategic danger. If, alternatively, site k were to become harder to protect at the margin, then the marginal cost of protecting against the non-strategic threat rises and the defender (weakly) shifts resources toward the strategic threat by spending more on the sites in $T_A(r)$. Finally, suppose a site subject to deliberate attack becomes more difficult to guard at the margin or the non-strategic threat against that site η_j declines. Then the marginal cost of strategic defense goes up, and the defender shifts resources away from it by spending less on the sites in $T_A(r)$. (See the proof of Proposition 3.)

As with border defense, we characterize the equilibrium allocations formally by describing a curve through the space of possible allocations. In the case of border defense, the defender moved along the curve as it shifted resources from border defense to site defense, optimally allocating the latter among the sites. In the present case, the defender moves along the curve by shifting resources from defending against the non-strategic threat to protecting against the strategic threat where the resources dedicated to each type of threat are optimally allocated against that threat. The defender's unique equilibrium allocation lies on this curve at the point where the marginal gain from spending

[12] The probability of no attack on j is $(1-\eta_j)(1-\alpha_j)$. So the probability of an attack is $1-(1-\eta_j)(1-\alpha_j) = \eta_j + \alpha_j(1-\eta_j)$.

[13] As elsewhere, a change in the marginal gain may not lead to a reallocation if a theatened site is unprotected at that allocation (i.e., site m enters $T_A(r)$ so $r_m = 0$ and $\alpha_m > 0$).

more to defend against the strategic threat is just offset by the marginal loss of having less to spend on defending against the non-strategic threat.

To describe this curve analytically, suppose that the defender completely discounts the strategic threat. The resulting allocation v equates the marginal benefit to protecting a given site against the non-strategic threat with the marginal benefits of protecting any other site. That is, v satisfies three conditions: (i) $\eta_j L_j \delta_j'(v_j) = \eta_k L_k \delta_k'(v_k)$ whenever $v_j > 0$ and $v_k > 0$, (ii) $\eta_j L_j \delta_j'(v_j) \leq \eta_m L_m \delta_m'(v_m)$ whenever $v_j > 0$ and $v_m = 0$, and (iii) $v_1 + \cdots + v_N = R$.

The path P_v begins at v, and a point is on the path P_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 conditions (i) and (ii) hold for all sites not in $T_A(r)$. These conditions guarantee that resources devoted to the non-strategic threat, i.e., to $j \notin T_A(r)$, cannot be profitably reallocated among those sites. As for the resources dedicated to the strategic threat, the definition of $T_A(r)$ ensures that the resources spent on these sites minimize the attacker's payoff over these sites.[14] 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.

To identify the unique equilibrium allocation where 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, define the defender's expected loss from guarding against the strategic threat at $r = P_v(\sigma)$ to be $L_S(\sigma, \alpha) \equiv \sum_{j \in T_A(r(\sigma))} L_j \delta_j(r_j(\sigma))[\eta_j + \alpha_j(1 - \eta_j)]$ and

[14] 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.

from the non-strategic threat to be $L_v(\sigma,\alpha) \equiv \sum_{n \notin T_A(r(\sigma))} L_n \delta_n(r_n(\sigma))[\eta_n + \alpha_n(1-\eta_n)]$. The defender's expected loss can now be written as is $L(r(\sigma),\alpha) = L_v(\sigma,\alpha) + L_S(\sigma,\alpha)$. Then the unique equilibrium allocation is either at one of the endpoints of the path, v or $P_v(R)$, or satisfies $-\partial L_v(\sigma,\widehat{\alpha})/\partial\sigma = \partial L_S(\sigma,\widehat{\alpha})/\partial\sigma$ where $\widehat{\alpha}$ is the 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 $T_A(r)$, i.e., $\widehat{\alpha}$ satisfies $[\eta_j + \widehat{\alpha}_j(1-\eta_j)]L_j\delta'_j(r_j) = [\eta_k + \widehat{\alpha}_k(1-\eta_k)]L_k\delta'_k(r_k)$ for $j,k \in T_A(r)$. Proposition 3 collects these results.

Proposition 3: *The defender always plays the same pure-strategy in any equilibrium of the game. This allocation lies on P_v where the marginal loss of spending less on the non-strategic threat just offsets the marginal gain of spending more on the strategic threat and is defined almost everywhere by $-\partial L_v(\sigma,\widehat{\alpha})/\partial\sigma = \partial L_S(\sigma,\widehat{\alpha})/\partial\sigma$ which is equivalent to:*

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

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

Uncertainty about the Attacker's Preferred Targets

The previous discussion assumed that the defender knows how the attacker ranks potential targets. But the defender may actually be uncertain about the 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.[15]

The answer to the resource allocation question and the way of thinking about finding the optimal allocation are quite similar to the border-defense and non-strategic-threat

[15] See Bier, Oliveros, and Samuelson (2005) for a complementary analysis that allows for a continuum of types of attacker but assumes that the defender only has two sites to protect.

problems. Suppose the defender were certain it was facing γ and allocated all of its resources accordingly. (The baseline algorithm can be used to find this allocation.) This allocation equalizes γ's payoff to attacking across the defended sites. But it generally is not optimal against a terrorist like τ with different priorities. The optimal allocation against γ will usually not minmax τ by equalizing its payoffs across all of the defended sites.

Suppose attacking site t_1 maximizes τ's payoff at the optimal allocation against γ. Then the defender, if it wants to hedge against the prospect of facing τ, should shift resources to t_1 and necessarily away from the other sites while keeping γ's payoff to attacking those sites equal. 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, becomes equally attractive. At this point allocating still more resources to protect against τ means that the defender will have to spend those resources on both t_1 and t_2 so as to keep τ's payoff to attacking them equal while at the same time keeping γ's payoff to attacking the other sites equal. The defender continues shifting resources in this way, having to spread them across more and more sites, until the marginal loss of spending less on defending against γ just offsets the marginal gain from allocating more to defending against τ. 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 $A_j^\gamma \geq 0$ and τ's payoff is $A_j^\tau \geq 0$. 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 $L(r,\alpha,\phi) \equiv \sum_{j=1}^N [\mu\alpha_j + \phi_j(1-\mu)] L_j \delta_j(r_j)$.

The comparative static results parallel those in the case of threats with non-strategic components. Anything that increases the marginal gain to defending against one type induces the defender to spend (weakly) more guarding against that type. If, for example, the defender is more confident it is facing γ (μ increases), the defender spends (weakly) more on the sites γ threatens (the $j \in T_\gamma(r)$ which are the sites offering γ its highest payoff). Similarly, if a site γ threatens becomes harder to protect at the margin, the defender shifts resources away from defending against γ and invests more in protecting against τ by reallocating from sites in $T_\gamma(r)$ to sites in $T_\tau(r)$.

As before, we characterize the defender's equilibrium allocation by describing a curve through the space of possible allocations. In the case of border defense, the defender moved along the curve by shifting resources between border defense and site defense while optimally allocating the latter among the sites. In the case of threats with non-strategic components, the defender moved along the curve by shifting resources between defending against the non-strategic and strategic components threats where the resources dedicated to each type of threat are optimally allocated against that threat. In the present case, the defender moves along the curve by shifting resources between their optimal allocation against γ and against τ. The defender's unique equilibrium allocation lies on this curve at the point where the marginal loss of having less to spend on defending against γ just offsets the marginal gain from dedicating more to defending against τ.

To specify the path, 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 payoffs. That is, $\widehat{r} \in P_g$ if and only if $\sum_{j \in T_\gamma(\widehat{r}) \cup T_\tau(\widehat{r})} \widehat{r}_j = R$. Parameterize movement along P_g in terms of σ, the resources the defender allocates against τ, by defining $P_g(\sigma)$ to be the point $r \in P_g$ where $\sigma = \sum_{j \in T_\tau(r)} r_j$ for $\sigma \in [\underline{\sigma}, \overline{\sigma}]$ with $\underline{\sigma} = \sum_{j \in T_\tau(g)} r_j$ and $\overline{\sigma} = \sum_{j \in T_\tau(t)} r_j = R$.

Clearly any equilibrium allocation must lie on P_g. If $r' \notin P_g$, then $\sum_{j \in T_\gamma(r') \cup T_\tau(r')} r'_j < R$. Consequently, there exists a site k such that $r'_k > 0$ and $k \notin T_\gamma(r') \cup T_\tau(r')$. Thus, neither γ nor τ attack k with positive probability, so the defender can profitably deviate from r' by reallocating r'_k across the sites in $T_\gamma(r') \cup T_\tau(r')$.

The equilibrium on P_g is the allocation at which the marginal loss of spending less on guarding against γ by increasing σ just offsets the marginal gain to defending against τ. 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 $L_\gamma(\sigma, \alpha) \equiv \mu \sum_{j \in T_A(\widehat{r}(\sigma)) - T_\tau(\widehat{r}(\sigma))} \alpha_j L_j \delta_j(\widehat{r}_j(\sigma))$ and

$L_\tau(\sigma, \phi) \equiv (1-\mu) \sum_{j \in T_\tau(\widehat{r}(\sigma))} \phi_j L_j \delta_j(\widehat{r}_j(\sigma))$ where $T_A(\widehat{r}) = T_\gamma(\widehat{r}) \cup T_\tau(\widehat{r})$. Combining these payoffs gives the defender's total payoff $L(\widehat{r}(\sigma), \alpha) = L_\gamma(\sigma, \alpha) + L_\tau(\sigma, \alpha)$. Then the equilibrium allocation is either at one of the endpoints, $P_g(\underline{\sigma})$ or $P_g(\overline{\sigma})$, or satisfies $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma = -\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ where $\widehat{\alpha}$ and $\widehat{\phi}$ are γ's and τ's pseudo-equilibrium strategies. This leaves:

PROPOSITION 4: *The defender always plays the same pure-strategy in any Bayesian-Nash equilibrium of the game. This allocation lies on P_g where the marginal loss of spending less on defending against γ just offsets the marginal gain of spending more protecting against τ and is defined almost everywhere by $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma = -\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ which is equivalent to:*

$$(1-\mu)\left(\sum_{n \in T_\tau(r)} \frac{1}{L_n \delta_n'(r_n)}\right)^{-1} = \mu \left(\sum_{n \in T_\gamma(r)} \frac{1}{L_n \delta_n'(r_n)}\right)^{-1}.$$

Conclusion

Protecting the country's critical infrastructure from a terrorist attack is an immense challenge. Although the United States is spending billions on this effort, resources are scarce when compared to the number of possible targets. The country cannot completely protect everything and choices have to be made about what to defend and how much to spend on trying to defend it. One of the central recommendations of the 911 Commission was that these decisions "should be based strictly on as assessment of risks and vulnerabilities," (2004, 396) and Secretary Chertoff has repeatedly emphasized the importance of a risk-management approach throughout his tenure. To this end, the Department of Homeland Security is developing a National Asset Database of critical infrastructure and key resources (CI/KR) and an accompanying methodology "that will allow the prioritization of CI/KR assets necessary for the optimal allocation of finite CI/KR protection resources" (DHS 2005, 125). But this risk-management approach appears to treat terrorists as only partially and not fully strategic actors. This is likely to lead to a significant misallocation of resources even if the distortions of pork-barrel politics can be overcome.

This paper offers a framework for analyzing the defender's allocation problem in a baseline case and three substantively different extensions. The framework shows that these problems share a common formal structure and that we can think about the allocation problem in each of them 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, the payoff to striking those sites 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. The defender continues in this way until it has fully allocated its resources.

This algorithm defines the optimal allocation if the sites are independent and the attacker only has to decide how much to invest in hardening each site. The algorithm also highlights a key contrast between treating actors as fully or only partially strategic. If a site becomes harder to defend and an adversary is assumed to be only partially strategic, the defender would generally spend less on that site. Against a fully strategic actor, however, the defender spends more in order to equalize the payoff to attacking this site with the payoff to attacking the other defended sites.

In the cases of border defense, threats with non-strategic components, and uncertainty about the attacker's priorities, the baseline algorithm case can also be used to trace a curve through the space of possible allocations where movement along the curve reflects the fundamental trade off defining the substantive case: spending more on border defense and less on site defense, more on the strategic component of the threat and less on the non-strategic component, or more to protect against one type of terrorist and less on the other. The equilibrium allocations lie along the curve where the marginal gain of spending more is just offset by the marginal loss of spending less.

This conceptualization leads to the relevant comparative statics. Anything increasing the marginal gain of moving along the curve in one direction induces the defender to move in that direction. If the risk of a non-strategic "attack" – a natural disaster or an accident – increases, the defender allocates more against to guarding against this type of threat and less to defending against the prospects of a deliberate attack. If the defender is uncertain about its adversaries priorities but becomes more confident that it is facing one type of terrorist, the defender allocates more to the sites this type would strike.

Appendix

Proof of Proposition 1: Observe first that equilibria always exist and that the defender always plays a pure strategy in equilibrium. This follows from the fact 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, $-L(r, \alpha)$ 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.

To pin down the defender's equilibrium allocations, 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 the defender always plays $\widehat{r}$ any equilibrium.

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

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: Observe first that L is strictly convex.in (r, b) for a fixed α. (Algebra shows the Hessian is positive definite given $\partial^2 B(b - r_j)\delta_j(r_j)/\partial r_j^2 > 0$.) The argument at the start of the proof of Proposition 1 then ensures that equilibria exist and that the defender never mixes in equilibrium.

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 $T_A(\widehat{r})$.

Given that the defender only plays pure strategies in equilibrium, these exist, and must lie on P_B, the only remaining task is to characterize the equilibrium allocations along P_B. Begin by noting that the marginal gain from investing in site defense $\partial S_D(\sigma, \widehat{\alpha})/\partial\sigma$ is well defined along P_B except possibly at the finitely many point where the defender starts to allocate resources to an additional site. To see that this is so, say that 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.

Then $\partial S_D(\sigma, \widehat{\alpha})/\partial\sigma$ 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 no site enters T_A at $\widehat{r}$, $\widehat{r}_j > 0$ and $\widehat{r}_k > 0$ and therefore $\widehat{\alpha}_j L_j \delta'_j(\widehat{r}_j) = \widehat{\alpha}_k L_k \delta'_k(\widehat{r}_k)$ for all j, $k \in T_A(\widehat{r})$. The discussion preceding Corollary 1 establishes $\widehat{\alpha}_k = \left(\sum_{n \in T_A(\widehat{r})} L_k \delta'_k / L_n \delta'_n\right)^{-1}$ for $k \in T_A(\widehat{r})$ and $\widehat{\alpha}_k = 0$ for $k \notin T_A(\widehat{r})$. Indeed, with $\widehat{r}_j > 0$ and $\widehat{r}_k > 0$ for all j, $k \in T_A(\widehat{r})$, $\widehat{\alpha}$ is the unique best response to $\widehat{r}$ that rationalizes the defenders allocation across $T_A(\widehat{r})$.

With $\widehat{\alpha}$ defined, $\partial S_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ is well defined and equals $\sum_{j \in T_A(r)} \widehat{\alpha}_j L_j \delta'_j(\widehat{r}_j) d\widehat{r}_j/d\sigma$

where recall $\widehat{\sigma}$ is the amount of resources allocated to defending sites in $T_A(\widehat{r})$. Because $\sum_{k\in T_A(\widehat{r})}\widehat{r}_j = \widehat{\sigma}$, $\sum_{k\in T_A(\widehat{r})} d\widehat{r}_j/d\sigma = 1$. Substituting for $\widehat{\alpha}_j$ gives $\partial S_D(\widehat{\sigma},\widehat{\alpha})/\partial\sigma = \left(\sum_{n\in T_A(\widehat{r})}[L_n\delta_n'(\widehat{r}_n)]^{-1}\right)^{-1}\sum_{k\in T_A(\widehat{r})} d\widehat{r}_j/d\sigma = \left(\sum_{n\in T_A(\widehat{r})}[L_n\delta_n'(\widehat{r}_n)]^{-1}\right)^{-1}$.

Now consider a $\widehat{r} = P_B(\widehat{\sigma})$ in E. Because a site, say e, enters T_A at $\widehat{\sigma}$, the summations in the expression for $\partial S_D(\sigma,\widehat{\alpha})/\partial\sigma$ for $\sigma \in (\widehat{\sigma},\widehat{\sigma}+\varepsilon)$ and for $\sigma \in (\widehat{\sigma},\widehat{\sigma}-\varepsilon)$ are taken over different sets. (Only one site enters because the values of $L_j\delta_j(0)$ are assumed to be distinct.) The continuity of the individual terms will therefore not ensure the continuity of the sums, and $\partial S_D(\widehat{\sigma},\widehat{\alpha})/\partial\sigma^- \equiv \lim_{\sigma\to\widehat{\sigma}^-}\partial S_D(\sigma,\widehat{\alpha})/\partial\sigma$ will generally not equal $\partial S_D(\widehat{\sigma},\widehat{\alpha})/\partial\sigma^+ \equiv \lim_{\sigma\to\widehat{\sigma}^+}\partial S_D(\widehat{\sigma},\widehat{\alpha})/\partial\sigma$.

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 $\widehat{r} \in P_B - E$.

Then any $\widehat{r} = P_B(\widehat{\sigma})$ is an equilibrium allocation if and only if $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma}) \in X(\widehat{\sigma})$. Suppose, first, that $\widehat{r} \notin E$. L's strict convexity ensures that an allocation minimizes L if it satisfies the relevant first-order conditions. Because no site enters 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 $\partial L(\widehat{r},\widehat{\alpha})/\partial r_j = 0$ for $j \in T_A(\widehat{r})$ where $\partial L(\widehat{r},\widehat{\alpha})/\partial r_j = \widehat{\alpha}_j L_j\delta_j'(\widehat{r}_j)\beta(R-\widehat{\sigma}) - \sum_{n=1}^N \widehat{\alpha}_n L_n\delta_n(\widehat{r}_n)\beta'(R-\widehat{\sigma})$. (Recall that the resources dedicated to border defense are implicitly defined by $b = R - \sum_{n=1}^N \widehat{r}_n$ which also means $b = R - \widehat{\sigma}$. The other marginal conditions $\partial L(\widehat{r},\widehat{\alpha})/\partial r_k \geq 0$ hold trivially at $j \notin T_A(\widehat{r})$ as $\widehat{\alpha}_j = 0$.)

Because no site enters $\widehat{r}$, $X(\widehat{\sigma})$ is a singleton and $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma}) \in X(\widehat{\sigma})$ if and only if $\beta(R-\widehat{\sigma})/\beta'(R-\widehat{\sigma}) = \sum_{n\in T_A(\widehat{r})}\delta_n(\widehat{r}_n)/\delta_n'(\widehat{r}_n)$. This can be rewritten as $\beta(R-\widehat{\sigma})\left(\sum_{k\in T_A(\widehat{r})}[L_k\delta_k'(\widehat{r}_k)]^{-1}\right)^{-1} = \beta'(R-\widehat{\sigma})\sum_{n\in T_A(\widehat{r})}\left(\sum_{k\in T_A(\widehat{r})}[L_k\delta_k'(\widehat{r}_k)]^{-1}\right)^{-1}\delta_n(\widehat{r}_n)/\delta_n'(\widehat{r}_n)$. The expressions for $\widehat{\alpha}_j$ finally yield $\widehat{\alpha}_j L_j\delta_j'(\widehat{r}_j)\beta(R-\widehat{\sigma}) = \sum_{n=1}^N \widehat{\alpha} L_n\delta_n(\widehat{r}_n)\beta'(R-\widehat{\sigma})$ for all $j \in T_A(\widehat{r})$. Hence, $\partial L(\widehat{r},\widehat{\alpha})/\partial r_j = 0$ for $j \in T_A(\widehat{r})$, and $(\widehat{r},\widehat{\alpha})$ is an equilibrium. Note further that $\beta(R-\widehat{\sigma})\left(\sum_{k\in T_A(\widehat{r})}[L_k\delta_k'(\widehat{r}_k)]^{-1}\right)^{-1} = \beta(R-\widehat{\sigma})\partial S_D(\widehat{\sigma},\widehat{\alpha})/\partial\sigma$ and $\beta'(R-\widehat{\sigma})\sum_{n\in T_A(\widehat{r})}\left(\sum_{k\in T_A(\widehat{r})}[L_k\delta_k'(\widehat{r}_k)]^{-1}\right)^{-1}\delta_n(\widehat{r}_n)/\delta_n'(\widehat{r}_n) = \beta'(R-\widehat{\sigma})S_D(\widehat{\sigma},\widehat{\alpha})$.

Now turn to any $\widehat{r} = P_B(\widehat{\sigma})$ where a site e enters $T_A(\widehat{r})$. Then $\widehat{r}_e = 0$ and $\widehat{r}_j > 0$ for all $j \in T_A(\widehat{r})$ and $j \neq e$. 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. The effect of this is that when one solves for $\widehat{\alpha}_j$ by summing over the marginal conditions that hold with equality, the total is now $1 - \widehat{\alpha}_e$ and leaves $\widehat{\alpha}_j = (1 - \widehat{\alpha}_e) \left(\sum_{n \in T_A(\widehat{r}) - \{e\}} L_j \delta_j'(\widehat{r}_j) / \left[L_n \delta_n'(\widehat{r}_n) \right] \right)^{-1}$. As for $\widehat{\alpha}_e$, it must lie in the range $[0, \overline{\alpha}_e]$ where $\overline{\alpha}_e$ satisfies $\widehat{\alpha}_j L_j \delta_j'(\widehat{r}_j) = \overline{\alpha}_e L_e \delta_e'(0)$. This and the expression for $\widehat{\alpha}_j$ mean $\overline{\alpha}_e L_e \delta_e'(0) = (1 - \overline{\alpha}_e) \left(\sum_{n \in T_A(\widehat{r}) - \{e\}} \left[L_n \delta_n'(\widehat{r}_n) \right]^{-1} \right)^{-1}$. It is straightforward to show that $\widehat{r}$ is a best response to one of these $\widehat{\alpha}_e$ if and only if $\beta(R - \widehat{\sigma})/\beta'(R - \widehat{\sigma}) \in X(\widehat{\sigma})$.

Finally, if the β and δ_j are log convex, $d^2 \ln \beta / db^2 > 0$ and $d^2 \ln \delta_j / d\delta_j^2 > 0$, β'/β and δ_j'/δ_j are increasing. So, $\beta(R - \sigma)/\beta'(R - \sigma)$ is increasing in σ and $X(\sigma)$ is decreasing, and their intersection is unique. ■

Proof of Proposition 3: The argument in the proof of Proposition 1 again ensures that equilibria exist and that the defender never mixes in equilibrium. To show that any equilibrium allocation must lie on P_v, suppose the contrary. Then there would be an equilibrium allocation $\widehat{r} \notin P_v$ such that one of the following two conditions does not hold: (i) $L_j \delta_j'(\widehat{r}_j)\eta_j = L_k \delta_k'(\widehat{r}_k)\eta_k$ whenever $\widehat{r}_j > 0$, $\widehat{r}_k > 0$, and $j, k \notin T_A(\widehat{r})$ and (ii) $L_j \delta_j'(\widehat{r}_j)\eta_j \leq L_m \delta_m'(\widehat{r}_m)\eta_m$ 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, 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 $\widehat{r} \in T_A(\widehat{r})$.

It remains to characterize the equilibrium allocation on P_v and show 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 $\partial L_S(\sigma, \widehat{\alpha})/\partial\sigma$ 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 $L_j\delta_j'(\widehat{r}_j)[\eta_j+\widehat{\alpha}_j(1-\eta_j)] = L_k\delta_k'(\widehat{r}_k)[\eta_k+\widehat{\alpha}_k(1-\eta_k)]$ for all j, $k \in T_A(\widehat{r})$. (That $\widehat{r} \in P_v - E$ means $\widehat{r}_j > 0$ for all $j \in T_A(\widehat{r})$ as $\widehat{r}_j = 0$ for $j \in T_A(\widehat{r})$ can only happen when $\widehat{r}_j$ enters T_A. And $\widehat{r}_j > 0$ implies that condition (i) holds for all sites in T_A.)

Solving these equations for $\widehat{\alpha}_j$ in terms of $\widehat{\alpha}_k$, summing the $\widehat{\alpha}_j$ over $T_A(\widehat{r})$, using the fact that the sum equals one, and rearranging terms yield $\eta_j+\widehat{\alpha}_j(1-\eta_j) = \left(\sum_{k\in T_A(r(\sigma))}(1-\eta_k)^{-1}\right)$ $\left(\sum_{k\in T_A(r(\sigma))}[L_j\delta_j'/L_k\delta_k'(1-\eta_k)]\right)^{-1}$ for $j \in T_A(\widehat{r}(\sigma))$ and, of course, $\widehat{\alpha}_n = 0$ for $n \notin T_A(\widehat{r}(\sigma))$.

With $\widehat{\alpha}$ defined, $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma$ is well defined and equal to $\sum_{j\in T_A(\widehat{r})}[\eta_j+\widehat{\alpha}_j(1-\eta_j)]L_j\delta_j'(\widehat{r}_j)d\widehat{r}_j/d\sigma$ where recall $\widehat{\sigma}$ is the amount of resources allocated to defending sites in $T_A(\widehat{r})$. Substituting for $\widehat{\alpha}_j$ leaves $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma = \left(\sum_{k\in T_A(r(\sigma))}(1-\eta_k)^{-1}\right)$ $\left(\sum_{k\in T_A(r(\sigma))}[L_j\delta_j'/L_k\delta_k'(1-\eta_k)]\right)^{-1}\sum_{k\in T_A(\widehat{r})}d\widehat{r}_j/d\sigma$. But, $\sum_{k\in T_A(\widehat{r})}\widehat{r}_j = \widehat{\sigma}$, so $\sum_{k\in T_A(\widehat{r})}d\widehat{r}_j/d\sigma = 1$ and $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma = \left(\sum_{k\in T_A(r(\sigma))}(1-\eta_k)^{-1}\right)$ $\left(\sum_{k\in T_A(r(\sigma))}[L_k\delta_k'(1-\eta_k)]^{-1}\right)^{-1} = [\eta_k+\widehat{\alpha}_k(1-\eta_k)]\,L_k\delta_k'(\widehat{r}_k)$ 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 $\partial L_S(\sigma,\widehat{\alpha})/\partial\sigma$ 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 $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma^- \equiv \lim_{\sigma\to\widehat{\sigma}^-}\partial L_S(\sigma,\widehat{\alpha})/\partial\sigma$ will generally not equal $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma^+ \equiv \lim_{\sigma\to\widehat{\sigma}^+}\partial L_S(\sigma,\widehat{\alpha})/\partial\sigma$.

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 $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma^-$ and $\partial L_S(\widehat{\sigma},\widehat{\alpha})/\partial\sigma^+$. (This interval is simply a single point at all $\widehat{r} \in P_v - E$.) At the endpoints, define $Y(\underline{\sigma}) \equiv \partial L_S(\underline{\sigma},\widehat{\alpha})/\partial\sigma^+$ and $Y(\overline{\sigma}) \equiv \partial L_S(\overline{\sigma},\widehat{\alpha})/\partial\sigma^-$.

Turning to the marginal loss from investing slightly less against the non-strategic threat at $\widehat{r}(\widehat{\sigma}) \in P_v - E$, $\partial L_v(\widehat{\sigma},\widehat{\alpha})/\partial\sigma \equiv \sum_{j\notin T_A(r)}\eta_jL_j\delta_j'(\widehat{r}_j\,(\widehat{\sigma}))d\widehat{r}_j/d\sigma$ where the total resources dedicated to non-strategic defense are $\sum_{j\notin T_A(r)}\widehat{r}_j = R-\widehat{\sigma}$. By construction,

$\eta_j L_j \delta_j'(\widehat{r}_j) = \eta_k L_k \delta_k'(\widehat{r}_k)$ 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 $\partial L_v(\sigma, \widehat{\alpha})/\partial\sigma = -\eta_k L_k \delta_k'(\widehat{r}_k)$ for any $r_k > 0$ with $k \notin T_A(\widehat{r})$. This expression for $\partial L_v(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ does not involve summations and we can simply take $\partial L_v(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ at $\widehat{r}(\widehat{\sigma}) \in E$ to be $-\eta_k L_k \delta_k'(\widehat{r}_k)$ for any $\widehat{r}_k > 0$ with $k \notin T_A(\widehat{r})$.

The correspondence $Y(\sigma)$ and $-\partial L_v(\sigma, \widehat{\alpha})/\partial\sigma$ intersect at most once. Observe first that the $-\partial L_v(\sigma, \widehat{\alpha})/\partial\sigma = \eta_k L_k \delta_k'(\widehat{r}_k)$ 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 $\partial L_S(\sigma, \widehat{\alpha})/\partial\sigma$ are decreasing. These terms are also negative so summing over more terms reduces the sum still further. 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 $-\partial L_v(\sigma, \widehat{\alpha})/\partial\sigma$ is decreasing, they intersect in at most one point along P_v.

If this intersection occurs at $\widehat{r} = P_v(\widehat{\sigma}) \in P_v - E$, $\widehat{r}$ 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 $\widehat{r}$. To see that $\widehat{r}$ is a best response to $\widehat{\alpha}$, observe $\delta_n'' > 0$ for all n ensures that $\widehat{r}$ is a best reply if it satisfies the requisite first order conditions. By construction, those are satisfied at any $j, k \in T_A(\widehat{r})$ and at any $j, k \notin T_A(\widehat{r})$. Further, $\partial L_S(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = -\partial L_v(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$ implies $[\eta_k + \widehat{\alpha}_k(1 - \eta_k)]\, L_k \delta_k'(\widehat{r}_k) = \eta_j L_j \delta_j'(\widehat{r}_j)$ 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 $-\partial L_v(\sigma, \widehat{\alpha})/\partial\sigma$ 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 $-\partial L_v(\sigma, \widehat{\alpha})/\partial\sigma$ do not intersect, the existence of an equilibrium implies one endpoint or the other is the unique equilibrium.

To sketch the comparative-static results, an increase in η_k or δ_k' for $k \notin T_A(r)$ shifts $\eta_k L_k \delta_k'(r_k)$ down and weakly moves the intersection to the right. An increase in δ_j' for

$j \in T_A(r)$ shifts Y up and weakly moves the intersection to the left. ■

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, and any allocation satisfying the relevant first order conditions also meets the second order conditions. Any equilibrium allocation must clearly lie on P_g.

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 $\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ and $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma$ 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 $j \notin T_\gamma(r(\sigma))$, $\widehat{\alpha}_j = \left(\sum_{n\in T_\gamma(r(\sigma))} L_j\delta'_j(r_j)/[L_n\delta'_n(r_n)]\right)^{-1}$ for $j \in T_\gamma(r(\sigma))$, $\widehat{\phi}_j = 0$ for $j \notin T_\tau(r(\sigma))$, and $\widehat{\phi}_j = \left(\sum_{n\in T_\tau(r(\sigma))} L_j\delta'_j(r_j)/[L_n\delta'_n(r_n)]\right)^{-1}$ for $j \in T_\tau(r(\sigma))$. Substituting for $\widehat{\alpha}$ and $\widehat{\phi}$ in the expressions for the derivatives gives $\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma = \mu \sum_{n\in T_\gamma(r)} \widehat{\alpha}_n L_n \delta'_n(r_n) dr_n/d\sigma = -\mu \left(\sum_{n\in T_\gamma(r)}[L_n\delta'_n(r_n)]^{-1}\right)^{-1} = -\mu\widehat{\alpha}_n L_n\delta'_n(r_n)$ and $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma = (1-\mu)\left(\sum_{n\in T_\gamma(r)}[L_n\delta'_n(r_n)]^{-1}\right)^{-1} = (1-\mu)\widehat{\phi}_n L_n\delta'_n(r_n)$ and 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 $Z_\gamma(\sigma) = \{z : z \in [-\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma^+, -\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma^-]\}$ is strictly decreasing in σ whereas $Z_\tau(\sigma) = \{z : z \in [\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma^-, \partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma^+]\}$ 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 $-\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma = \partial L_\tau(\sigma, \widehat{\alpha})/\partial\sigma$ implies $\mu\widehat{\alpha}_j L_j\delta'_j(r_j) = (1-\mu)\widehat{\phi}_n L_n\delta'_n(r_n)$ 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 $\mu\widehat{\alpha}_j L_j\delta'_j(r_j) > (1-\mu)\widehat{\phi}_n L_n\delta'_n(r_n)$ 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, $\mu\widehat{\alpha}_j L_j\delta'_j(r_j) < (1-\mu)\widehat{\phi}_n L_n\delta'_n(r_n)$ 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.

The comparative statics follow directly. An increase in μ shifts Z_γ down, Z_τ up, and weakly lowers the σ at which they intersect. Analogous results follow if sites become harder to defend at the margin. ■

References

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

9/11 Public Discourse Project Final Report. 2005. "Final Report on the 9/11 Commission Recommendations." Available at www.9-11pdp.org. Accessed on December 29, 2005.

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.

Bier, Vicki, Santiago Oliveros, and Larry Samuelson. 2005. "Choosing What to Protect." Manuscript, Department of Economics, University of Wisconsin.

Blackett, D.W. 1958. "Pure Strategy Solutions to Blotto Games," *Naval Research Logistics Quarterly* 5:107-109.

Brown, Gerald, Matthew Carlyle, Javier Salmerón, and Kevin Wood. 2005. "Defending Critical Infrastructure." Manuscript, Department of Operations Research, Naval Postgraduate School, Monterey, CA.

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.

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: 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: www.dhs. gov/dhspublic/display?content=4597.

Chertoff, Michael. 2006. "Remarks by Homeland Security Secretary Michael Chertoff at a Press Conference on the Fiscal Year 2006 Urban Areas Security Initiative Grants." Washington, D.C. Department of Homeland Security, July 13. Available at: www.dhs. gov/dhspublic/display? content=5319 (accessed January 4, 2006).

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

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

DHS. 2005. "Draft National Infrastructure Protection Plan." Washington, D.C.: Department of Homeland Security. Available at: www.fas.org /irp/agency/dhs/nipp110205.pdf. Accessed on January 4, 2006.

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. 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.

Flynn, Stephen. 2004. *America the Vulnerable.* New York: Harper Collins.
Fudenberg, Drew and Jean Tirole. 1991. *Game Theory.* Cambridge, MA.: MIT 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.
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: 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.
Shubik, Martin and Robert Weber. 1981. "Systems Defense Games," *Naval Research Logistics Quarterly* 28:281-87.
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: 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: www.whitehouse. gov/pcipb/physical.html