DOI: 10.1017/S0003055407070244

Defending against Terrorist Attacks with Limited Resources

ROBERT POWELL *University of California, Berkeley*

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: (1) 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; (2) if the defender can also allocate resources to border defense, intelligence, or counterterrorist operations which, if successful, protect all of the sites; (3) if threats have strategic and nonstrategic components (e.g., the threat to public health from bioterror attacks and the natural outbreak of new diseases); and (4) 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.

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 (OMB 2003, 2004, 2005). But the scope of the task is vast. Prior to 9/11, the Environmental Protection Agency (EPA) 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 100,000, 600 sites threatened between 100,000 and one million people, and 123 sites threatened more than one million people (Belke 2000). By the end of 2005, the National Asset Database listed about 80,000 sites including "nuclear power plants, pipelines, bridges, stadiums, and locations such as Times Square" (GAO 2005, 75).[1] Of these, the Department of Homeland Security (DHS) had identified 1,700 highest priority sites and intended to visit each of them "to assess their vulnerabilities to various forms of attack . . ." (Moteff 2006, 38). As DHS 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. "One fact dominates all homeland security threat assessments," the *National Strategy for Homeland Security* emphasizes, "terrorists are strategic actors" (White House 2002, 7). No one thinks that hardening and rebuilding the levies around New Orleans affects the probability another hurricane like Katrina will strike that city. But strategic actors do try to strike where the defense is weak and the expected gains are high. Protecting one site may therefore shift the risk of attack to another. "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).[2]

This paper develops a framework for analyzing the problem of allocating limited resources against a strategic adversary in four settings.[3] 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, that is, 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

Robert Powell is Robson Professor, Travers Department of Political Science, 210 Barrows Hall, 1950, University of California, Berkeley, CA 94720-1950 (RPowell@Socrates.Berkeley.edu).

For helpful comments, criticisms, and discussion, I thank Daron Acemoglu, Manas Baveja, Mariano-Florentino Cuéllar, Ernesto Dal Bó, Lynn Eden, James Fearon, Ron Hassner, William Kastenberg, Lance Kim, Victoria Kingsley, Michael May, Santiago Oliveros, 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.

[1] A critical review of the database is DHS 2006b, which notes that the list of assets also includes petting zoos, fun parks, and some Wal-Marts.

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

[3] The present analysis focuses on understanding the first-best outcome, that is, the optimal allocation for a single unitary defender. Examples studying second-best outcomes due to the distorting effects of multiple actors with different interests include Bueno de Mesquita 2005 and Rosendorff and Sandler 2004.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

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 other. 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.

In other words, the defender's equilibrium allocation minmaxes the attacker; that is, it is the allocation that minimizes the attacker's maximum payoff. This allocation is remarkably robust. It remains the unique equilibrium allocation whether the baseline model is zero-sum or not and regardless of whether the attacker sees the defender's allocation before deciding which site to attack (i.e., whether the game is static or dynamic). 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. Nor it is it generally the case that the equilibrium action in a static game is the same as in a dynamic game.

The second setting extends the baseline case to allow the defender to allocate resources to border defense, counterterrorist operations, and intelligence in addition to spending on specific sites. The key idea here is that investing in these 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 nonstrategic 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 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 nonstrategic: 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, about the way that the terrorists rank potential targets. How, for example, should a defender allocate its resources if it is unsure whether its adversary is primarily interested in inflicting casualties and economic losses or in destroying symbolically significant targets like the Statue of Liberty?

The following analysis characterizes the defender's optimal (equilibrium) allocations in each of these settings. These characterizations show that although the extensions of the baseline case are substantively quite diverse, the same basic approach can be used to determine the optimal allocation in all of them.

STRATEGIC VERSUS NONSTRATEGIC TERRORISTS

Although the *National Strategy for Homeland Security* emphasizes the dominating fact that terrorists are strategic, early spending decisions on homeland security and, especially, on critical infrastructure protection have been widely criticized for being dominated by pork-barrel politics. As the *9/11 Commission Report* delicately puts it, "In a free-for-all over money, it is understandable that representatives will work to protect the interests of their home states and districts. But this issue is too important for politics as usual to prevail. Resources must be allocated according to vulnerabilities" (2003, 396).[4] Secretary Chertoff underscored the same point shortly after taking office, "Risk management must guide our decision making as we examine how we can best organize to prevent, respond and recover from an attack" (2005a, 2).

In keeping with this, the cornerstone of the *National Infrastructure Protection Plan* is "its risk management framework" (DHS 2006a, 35). This framework does take an adversary's intentions into account, but it 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). 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 the model to an arbitrary number of sites and develops a more general framework for analyzing the allocation problem.[5]

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 $v_j(r_j)$ to be the vulnerability of site j; that is, $v_j(r_j)$ is 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_1 and the probability of an attack on site 2 is α_2 where $\alpha_1 + \alpha_2 = 1$. All of this implies that the defender's expected loss to allocating r and $R - r$ to sites 1 and 2, respectively, is $L(r) = \alpha_1 L_1 v_1(r) + \alpha_2 L_2 v_2(R - r)$. The first term on the right in this equality is the expected loss from an attack on the first site, $L_1 v_1(r)$, weighted by the probability of an attack on that site. The second term is the expected loss from an attack on the second site weighted by the probability of an attack on that site.

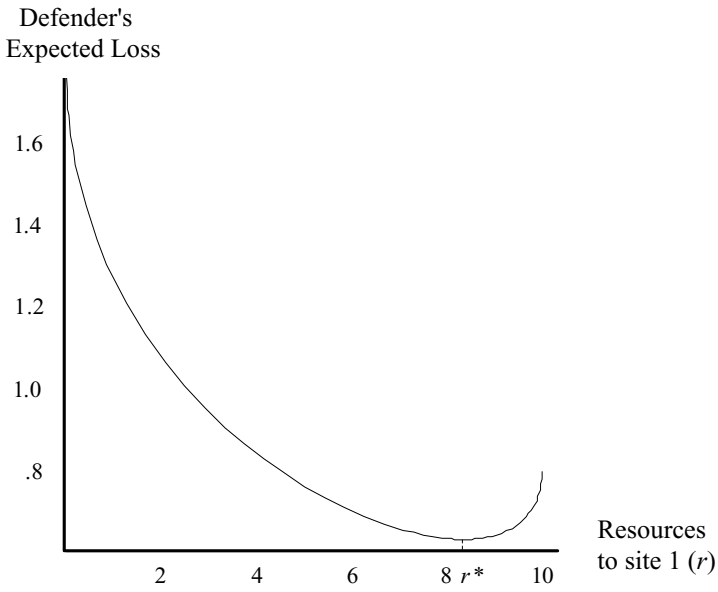
The basic elements of this model parallel the three key components of DHS's approach to risk management which are vulnerability, threat, and consequence. Vulnerability "is the probability that a particular attack will succeed against a particular target" (GAO 2005, 25), and this is precisely what v_j is. Threat "is

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

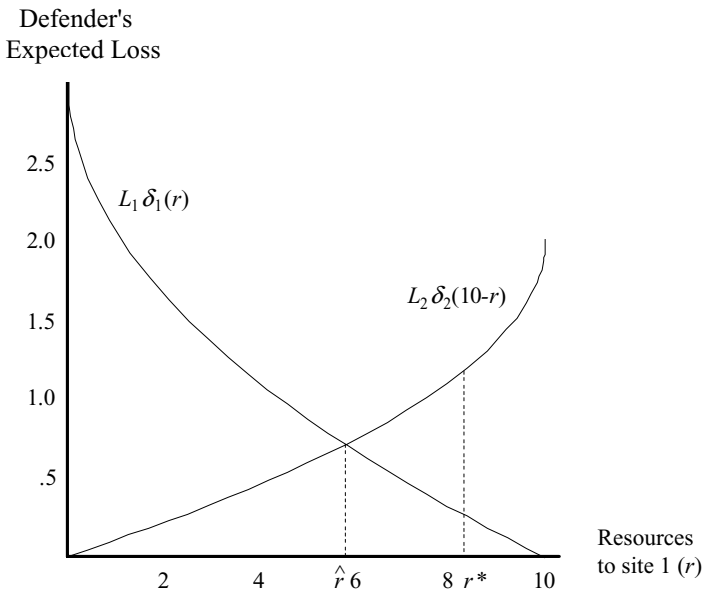
[5] Overviews and critiques of the DHS's risk-management approach are DHS 2006a, GAO 2005, Moteff 2004, 2006.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

FIGURE 1. Strategic versus Nonstrategic Terrorists



(a) Non-strategic terrorists.



(b) Strategic terrorists.

the probability that a specific target is attacked in a specific way" (Willis et al. 2005, 8) and is formalized as the probability of attack α_j. Finally, the loss L_j corresponds to the consequence which "is the range of loss or damage that can be expected from a successful attack" (2006a, 41).[6] Combining these elements, the risk associated with a site is the probability of an attack (threat) times the probability that an attack succeeds (vulnerability) times the loss from a successful attack (consequence; Willis et al. 2005, 10). That is, risk = consequence × vulnerability × threat which parallels the expected loss $L_j\, v_j(r_j)\alpha_j$.

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

[6] Strictly speaking, L_j is a von Neumann-Morgenstern utility which is related to but not the same as economic losses.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

3/2 (i.e., $L_1 = 3$ and $L_2 = 2$), the probability of a successful attack on a site is $v_j(r_j) = 1 - \sqrt{r_j/10}$, and total resources are $R = 10$. 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., $\alpha_1/\alpha_2 = 3/2$, which means $\alpha_1 = 3/5$ and $\alpha_2 = 2/5$).

The defender's optimal allocation $r^* \approx 8.3$ minimizes its expected loss. At this allocation, the marginal gain from spending more on protecting the first site is just offset by the marginal loss of having less to spend on the second site. More formally, the optimal risk-management allocation r^* given the threat assessment α_1 satisfies the first-order condition $dL(r^*)/dr = 0$ or, equivalently, r^* solves $\alpha_1 L_1 v_1'(r^*) = (1 - \alpha_1) L_2 v_2'(R - r^*)$.

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.[7] The odds of an attack on the first site are not 3/2 and a spending decision based on these odds leads to a misallocation.

Put another way, the allocation r^* and attack probability $\alpha = 3/5$ are not a Nash equilibrium of the underlying game. In a Nash equilibrium each actor plays optimally against the other; that is, 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 ($r^*/\widehat{r} \approx 1.42$).

[7] This zero-sum assumption will be discussed and relaxed below.

A BASIC FRAMEWORK

This section presents a basic game-theoretic framework for 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 S sites; a terrorist group then chooses which target to attack. In the static version of the game discussed here, the attacker is unable to observe the defender's allocation. (The dynamic version in which the attacker can see how the defender allocates its resources before deciding which site to hit is discussed below.) A strategy for the defender is simply a resource allocation $(r_1, \ldots, r_S)$ such that $r_1 + \cdots + r_S \leq R$; a strategy for the terrorist group is a set of probabilities $(\alpha_1, \ldots, \alpha_S)$, where α_j is the probability the terrorist group attacks site j with $\alpha_1 + \cdots + \alpha_S = 1$.

To specify the player's payoffs, suppose that the defender suffers a loss $L_j > 0$ 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 $v_j(r_j)$ with $v_j'(r_j) < 0$ and $v_j''(r_j) > 0$ as long as this site is imperfectly defended (i.e., as long as $v_j(r_j) > 0$).[8] Thus, the defender's expected loss if the terrorists strike j is $L_j v_j(r_j)$, and its expected loss if it plays strategy r and the terrorist group plays α is $L(r, \alpha) \equiv \sum_{j=1}^{S} \alpha_j L_j v_j(r_j)$. As for the attacker's payoffs, let $G_j > 0$ be the attacker's gain from successfully striking site j. Then the terrorist group's expected gain is $G(r, \alpha) \equiv \sum_{j=1}^{S} \alpha_j G_j v_j(r_j)$.

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.

ASSUMPTION 1 (SCARCE RESOURCES). *For any allocation of R, $v_j(r_j) > 0$ for some j.*

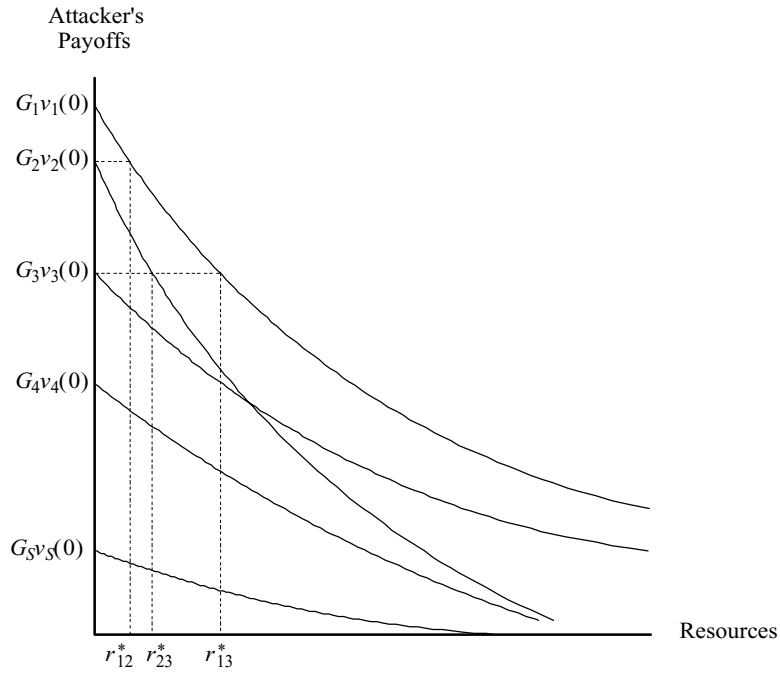
A Nash equilibrium is a pair of strategies (r, α) such that each player's strategy maximizes its payoff against the other player's strategy. If, therefore, the defender's allocation is $(r_1, \ldots, r_S)$, the terrorist group will attack the site offering the highest expected gain. That is, the attacker strikes a site j such that $G_j v_j(r_j) = M_A(r)$ where $M_A(r) \equiv \max\{G_1 v_1(r_1), \ldots, G_S v_S(r_S)\}$.

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, $G_1 v_1(0) > G_2 v_2(0) >$

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

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

FIGURE 2. The Optimal Allocation



$\cdots > G_Sv_S(0)$.[9] Then the most attractive target; that is, the site offering the attacker the highest expected payoff, is site 1, because $G_1v_1(0) = \max\{G_1v_1(0), \ldots, G_Sv_S(0)\}$. The defender, therefore, should begin 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 $G_1v_1(r_{12}^*) = G_2v_2(0) = \max\{G_1v_1(r_{12}^*), G_2v_2(0), \ldots, G_Sv_S(0)\}$. From this point on, the defender must allocate resources to both sites 1 and 2 in order to keep $G_1v_1(r_1) = G_2v_2(r_2)$. 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 two sites are no more attractive than site 3, that is, until $G_1v_1(r_{13}^*) = G_2v_2(r_{23}^*) = G_3v_3(0) = \max\{G_1v_1(r_{13}^*), G_2v_2(r_{23}^*), G_3v_3(0), \ldots, G_Sv_S(0)\}$. Now the defender allocates resources to sites 1, 2, and 3 so as to keep $G_1v_1(r_1) = G_2v_2(r_2) = G_3v_3(r_3)$. 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.

[9] Indexing the sites in this way simplifies the exposition and has no effect on the results. The inequalities are assumed to be strict to ease the exposition. This assumption will be relaxed below.

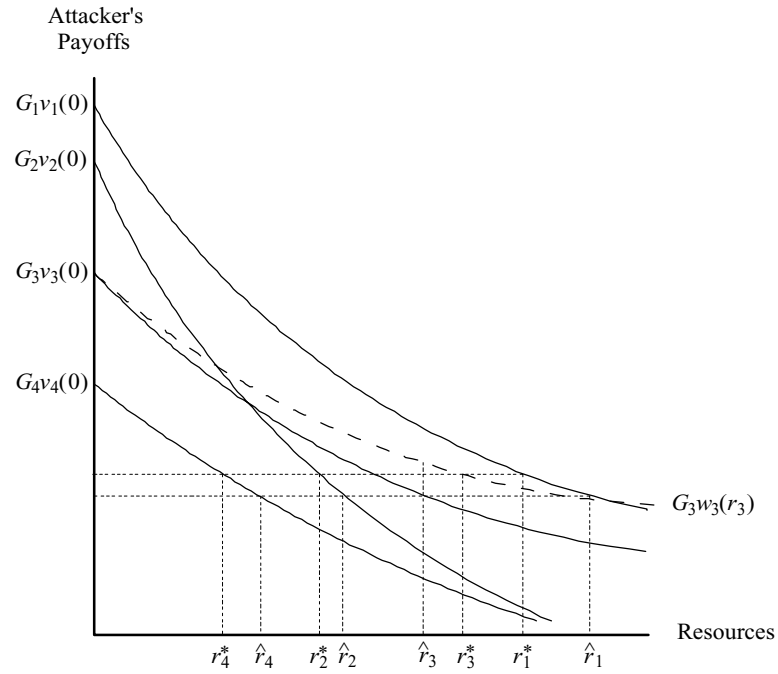
This allocation and the algorithm leading to it produce what may at first seem to be counterintuitive results. Suppose a site becomes uniformly more difficult to defend at the margin, that is, the marginal return to investing in protecting the site decreases. Formally, $v_j(r)$ shifts to $w_j(r)$, where $w_j(0) = v_j(0)$ and $w_j'(r) > v_j'(r)$ for all r. Intuitively, 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 at the margin, 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, which necessarily means spending less on the other sites. Equalizing the attacker's payoffs across the defended sties leads to allocation r^* in Figure 3.

Stating the results formally, assume more generally the strict inequalities, mentioned previously, are weak so that $G_1v_1(0) \geq G_2v_2(0) \geq \cdots \geq G_Sv_S(0)$. Then the defender's optimal allocation minmaxes the attacker by minimizing the attacker's maximum payoff. That is, the defender's unique equilibrium allocation $\widehat{r}$ solves $\min\{M_A(r)\colon r_1 + \cdots + r_S \leq R\}$. Still more formally,

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

FIGURE 3. When a Site Becomes Harder to Defend



LEMMA 1. *The defender has a unique minmax allocation.*

Proposition 1. *Nash equilibria of the static baseline game exist, and the defender plays the minmax allocation $\hat{r}$ in every Nash equilibrium.*[10]

Proofs of these claims and the propositions below are in the Appendix. In related work, Powell (2006) analyzes the dynamic version of the baseline model in which the attacker decides which site to strike after seeing how the defender allocates its resources. He shows that subgame perfect equilibria exist and that the defender plays the minmax allocation $\hat{r}$ in all of them. The minmax allocation $\hat{r}$ thus turns out to be a remarkably robust equilibrium allocation, remaining optimal in both zero-sum and nonzero-sum versions of the game as well as in both static and dynamic versions.

These results are surprising in two ways. As is well known, 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 earlier 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 the baseline model and "Colonel Blotto" games which have long been used to study force allocation issues (e.g., Blackett 1958; Shubik and Weber 1981; Tukey 1949). The classic blotto game is a two-actor, zero-sum game in which players A and B compete on S independent battlefields. Player A has F_A units of force to distribute among these battlefields, and player B has F_B. If A allocates f_A^k to battlefield k and B allocates f_B^k, A receives a payoff of $p_k(f_A^k, f_B^k)$ in the ensuing battle and B gets $-p_k(f_A^k, f_B^k)$. The players' payoffs in the game 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 when a site is destroyed equals the attacker's gain; that is, if $L_j = G_j$ for all sites j, then the defender's loss at j if it allocates r_j and the attacker "allocates" α_j is $\alpha_j L_j v_j(r_j)$ and exactly equal to the attacker's gain $\alpha_j G_j v_j(r_j)$.[11]

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

[10] In fact, the static game has a generically unique Nash equilibrium. If $\hat{r}_k > 0$ for all k for which $G_k v_k(\hat{r}_k) = M_A(\hat{r})$, the equilibrium is unique. Mulitplicity can arise if $G_n v_n(0) = M_A(\hat{r})$ for some n. The defender's equilibrium strategy in this case remains $\hat{r}$, but now there will be multiple $\hat{\alpha}$'s for which $(\hat{r}, \hat{\alpha})$ is an equilibrium.

[11] Blotto games are sometimes defined still more narrowly (e.g., Golman and Page 2006 and Roberson 2006) by assuming that a player is sure to win a battle if it allocates more to that battlefield than its adversary does: $p_k(f_A^k, f_B^k) = 1$ if $f_A^k > f_B^k$, $p_k(f_A^k, f_B^k) = 0$ if $f_A^k = f_B^k$, and $p_k(f_A^k, f_B^k) = -1$ if $f_A^k < f_B^k$. Blackett (1958), by contrast, assumes that p_k is continuous.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

concretely $G_1v_1(0) > G_2v_2(0)$ but $L_1v_1(0) < L_2v_2(0)$, 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., $L_i > L_j \Leftrightarrow G_i > G_j$), the game will not be zero sum unless the defender and attacker put precisely the same relative weights on the sites (i.e., $L_i/L_j = G_i/G_j$ for all i and 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 nonstrategic 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 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. Spending on sites that will not be attacked brings no gain and means that there is less to be spent on sites that will be attacked with positive probability. But only the sites offering the attacker the highest expected payoff will be attacked with positive probability and allocating resources solely to these sites has the effect of minimizing the attacker's maximum payoff.

The second surprising aspect of the robustness of the minmax allocation is that the defender's equilibrium allocation remains the same in both the static and dynamic versions of the game. These different settings typically lead to different equilibrium behavior (see Powell 2006 for an analysis of the dynamic model).

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 minmax allocation $\widehat{r}$. Let $T_A(r)$ be the set of sites giving the attacker its highest payoff at r. These are the sites the attacker might hit if the allocation were r. More formally, $T_A(r)$ is the set of pure-strategy, best replies to $\widehat{r}$:

$$T_A(r) = \{k : G_k v_k(r_k) = M_A(r)\}.$$

If the defender plays $\widehat{r}$, the attacker's payoff to striking any site in $T_A(\widehat{r})$ is the same as its payoff to striking any other site in $T_A(\widehat{r})$ and higher than its payoff to going after any site not in $T_A(\widehat{r})$. This implies that the probability the attacker strikes any site $j \notin T_A(\widehat{r})$ 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 $(\widehat{r}, \widehat{\alpha})$ is to be an equilibrium. 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 $i, j \in T_A(\widehat{r})$ whenever $r_j > 0$ and $r_i > 0$.[12] These first-order conditions imply $\widehat{\alpha}_1 L_1 v_1'(r_1) = \widehat{\alpha}_j L_j v_j'(r_j)$ for $j \in T_A$. Adding the $\widehat{\alpha}_j$'s for $j \in T_A$ and using the fact that these probabilities sum to one yields $\widehat{\alpha}_j = (\sum_{n \in T_A(\widehat{r})} L_j v_j'(\widehat{r}_j))/(L_n v_n'(\widehat{r}_n))^{-1}$. 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 $\widehat{\alpha}_j = (\sum_{n \in T_A(\widehat{r})} L_j v_j'(\widehat{r}_j))/(L_n v_n'(\widehat{r}_n))^{-1}$.*

It follows that the probability of an attack on a site increases as that site becomes harder to defend. As shown previously, the defender spends more on a site (and necessarily less on the other sites) as that site becomes harder to defend at the margin. 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.[13]

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 counterterrorist 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 counterterrorist 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 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

[12] Because $v_j'' > 0$ for all j, $L(r, \widehat{\alpha})$ is strictly convex in r, so these marginal conditions pin down the unique best-reply to $\widehat{\alpha}$. If $\widehat{r}_k = 0$ for some $k \in T_A(\widehat{r})$, 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; that is, $\partial L/\partial r_k \geq \partial L/\partial r_j$.

[13] Formally, suppose the initial equilibrium allocation is $\widehat{r}$ and v_1 shifts to w_1, where $w_1'(r_1) > v_1'(r_1)$ and $w_1(0) = v_1(0)$. Then at $\widehat{r}$ and the new equilibrium allocation r^* define $\widehat{z}_k \equiv L_1 w_1'(\widehat{z}_1)/[L_k v_k'(\widehat{r}_k)]$ and $z_k^* \equiv L_1 v_1'(r_1^*)/[L_k v_k'(r_k^*)]$ for $k > 1$, and observe that $\widehat{\alpha}_k = \widehat{z}_k \widehat{\alpha}_1$, $\alpha_k^* = z_k^* \alpha_1^*$, and $\widehat{z}_k > z_k^*$. Summing the probabilities gives $1 = \widehat{\alpha}_1(1 + \sum_{k>1} \widehat{z}_k) = \alpha_1^*(1 + \sum_{k>1} z_k^*)$, which implies $\widehat{\alpha}_1 < \alpha_1^*$.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

defenses" (O'Hanlon et al. 2003, 2).[14] 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 very 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 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 equals 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 counterterrorist 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 among 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.

[14] Assuming that border defense, intelligence, and counterterrorist 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.

To incorporate border defense in the basic framework formally, suppose that the defender has to decide how much to direct to the S 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 $v_j(r_j)\beta(b)$, 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 $\beta'' > 0$). To capture the idea that it is much easier to protect one site than secure the entire border, assume that if the defender only had one site to protect, it would always do better by spending on that site rather than on border defense, but the marginal returns to such spending are decreasing. Formally,

ASSUMPTION 2 (BORDERS ARE HARD TO DEFEND). $\partial\beta(b-r_j)v_j(r_j)/\partial r_j < 0$ *and* $\partial^2\beta(b-r_j)v_j(r_j)/\partial r_j^2 > 0$ *for all sites* j*, any* $b \in [0, R]$*, and any* $r \leq b$.

The defender's and attacker's payoffs are now $L((r,b),\alpha) = \sum_{n=1}^{S}\alpha_n L_n\beta(b)v_n(r_n)$ and $G((r,b),\alpha) = \sum_{n=1}^{S}\alpha_n G_n\beta(b)v_n(r_n)$. As before, the terrorist group attacks sites offering the highest expected gain; that is, those j such that $G_j v_j(r_j)\beta(b) = \max\{G_1v_1(r_1)\beta(b), \ldots, G_Sv_S(r_S)\beta(b)\} = \beta(b)M_A(r)$, where $r = (r_1, \ldots, r_S)$ is the defender's allocation to site defense.

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 nonstrategic 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 nonstrategic 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 $r = (r_1, \ldots, r_S)$ with $\sum_{j=1}^{S} r_j \leq R$ to be a point in $\mathbb{R}^S$ at which the defender is spending $\sigma = \sum_{j=1}^{S} r_j$ on site defense and $b = R - \sigma$ on border defense. Now define the curve P_B through $\mathbb{R}^S$ 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 $M_A(r)$ subject to $r_1 + \cdots + r_S \leq \sigma$ and can be found with the algorithm used to solve the baseline model.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

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 allocate however much it does spend on site defense optimally. 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.

To formalize these marginal conditions, define a *pseudo-equilibrium attack strategy* $\widehat{\alpha}(r)$ to be a best response to r that also rationalizes the defender's relative spending on the sites in $T_A(r)$, where, recall, $T_A(r)$ is set of attacker's pure-strategy, best replies to r. In the case of border defense, this means that $\widehat{\alpha}_j = 0$ for all $j \notin T_A(r)$; $L_j v_j'(r_j)\widehat{\alpha}_j = L_k v_k'(r_k)\widehat{\alpha}_k$ whenever $j, k \in T_A(r)$ with $r_j > 0$ and $r_k > 0$; and $L_j v_j'(r_j)\widehat{\alpha}_j \leq L_k v_k'(r_k)\widehat{\alpha}_k$ whenever $j, k \in T_A(r)$ with $r_j > 0$ and $r_k = 0$.

Now observe that the defender's expected loss at point $\widehat{r}(\sigma) = P_B(\sigma)$ can be written as $\Pi_D(\sigma, \alpha)\beta(R - \sigma)$, where $\Pi_D(\sigma, \alpha) \equiv \sum_{j=1}^{S} \alpha_j L_j v_j(\widehat{r}_j(\sigma))$. It follows that the equilibrium allocations lie on P_B, where $\beta(R - \sigma)\partial\Pi_D(\sigma, \widehat{\alpha})/\partial\sigma = \beta'(R - \sigma)\Pi_D(\sigma, \widehat{\alpha})$. 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.

This marginal condition is quite intuitive if all of the sites are identical; that is, $L_j = L_1$, $G_j = G_1$, and $v_j = v_1$ for $1 \leq j \leq S$. In these circumstances, the defender simply divides the resources spent on site defensive evenly across all of the sites, and the attacker strikes each site with equal probability. The defender's payoff therefore can be written as $\Pi_D(\sigma, \widehat{\alpha})\beta(R - \sigma) = \sum_{j=1}^{S} \alpha_j L_j v_j(\widehat{r}_j(\sigma))\beta(R - \sigma) = S(1/S)L_1 v_1(\sigma/S)\beta(R - \sigma) = L_1 v_1(\sigma/S)\beta(R - \sigma)$. The first-order condition is $L_1 v_1'(\sigma/S)\beta(R - \sigma)/S = L_1 v_1(\sigma/S)\beta'(R - \sigma)$. Hence, the marginal benefit of spending more on site defense (the left side of the previous equation) is the marginal gain of spending more on a site divided by the number of sites. If the number of sites is large, this marginal gain to site defense is small.

Proposition 2 states the results in the more general case when the sites are not identical.

Proposition 2. *Nash equilibria exist, and the defender always plays a pure strategy lying on P_B where the marginal loss of investing less on border defense just offsets the marginal gain of spending more on site defense. This condition generically satisfies $\partial\Pi_D(\sigma, \widehat{\alpha})/\partial\sigma\beta(R - \sigma) = \Pi_D(\sigma, \widehat{\alpha})\beta'(R - \sigma)$ which is equivalent to*:[15]

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

[15] That is, the set of equilibrium allocations at which this condition does not hold has zero measure.

THREATS WITH STRATEGIC AND NONSTRATEGIC COMPONENTS

Some threats entail both a strategic and nonstrategic 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 bioterror 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 Department of Homeland Security intends 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 crisis, will be funded as long as there 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 and protecting against a nonstrategic 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, thereby increasing the risk of a deliberate attack on those targets. How should a defender allocate its resources when threats involve both a strategic and a nonstrategic component?

The answer and the way of 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 nonstrategic 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 does not affect the nonstrategic threat to the other site.

But equating the margins to optimize spending against a nonstrategic threat means that a terrorists' payoffs will generally not be equal across the sites.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

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 nonstrategic 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 nonstrategic 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 nonstrategic threat, it must begin to distribute those resources across t_1 and t_2 so as to keep the terrorists' payoff to attacking either of the those 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 less on the nonstrategic threat.

To introduce the non-strategic component, let η_j be the probability of an "accidental," that is, non-strategic, attack on site j. Then the overall probability of an attack on j is $\eta_j + \alpha_j(1-\eta_j)$.[16] The defender's expected loss is thus $L(r,\alpha) = \sum_{j=1}^{S} L_j v_j(r_j)[\eta_j + \alpha_j(1-\eta_j)]$. As always, the attacker strikes sites offering the highest expected payoff, i.e., sites 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.[17] Suppose, for example, that site k is only subject to a non-strategic threat, because the attacker's payoff to hitting this site is lower than at other sites ($k \notin T_A(r)$). Then the probability this site will be struck is η_k as $\alpha_k = 0$. If the threat η_k increases, the marginal gain to protecting against the nonstrategic threat increases. As a result, the defender (weakly) shifts resources away from sites at risk of terrorist attack (i.e., those $j \in T_A(r)$) and allocates those resources to guarding against the nonstrategic danger. If, alternatively, site k were to become harder to protect at the margin (v_k' increases), then the marginal cost of protecting against the nonstrategic 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. Then the marginal cost of strategic defense goes up, and the defender shifts resources away from it by spending less on the sites in $T_A(r)$.

As with border defense, the equilibrium allocations can be characterized 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 nonstrategic threat to protecting against the strategic threat where the resources dedicated to each type of threat are optimally allocated against that threat. An equilibrium allocation lies on this curve at the point where the marginal gain from spending more to defend against the strategic threat is just offset by the marginal loss of having less to spend on defending against the nonstrategic threat.

To describe this curve analytically, suppose that the defender completely discounts the strategic threat. The resulting allocation δ equates the marginal benefit to protecting a given site against the nonstrategic threat with the marginal benefits of protecting any other site. That is, δ is the unique allocation minmizing $\sum_{k=1}^{S} \eta_k L_k v_k(\delta_k)$ and, therefore, satisfying three conditions: (i) $\eta_j L_j v_j'(\delta_j) = \eta_k L_k v_k'(\delta_k)$ whenever $\delta_j > 0$ and $\delta_k > 0$, (ii) $\eta_j L_j v_j'(\delta_j) \leq \eta_m L_m v_m'(\delta_m)$ whenever $\delta_j > 0$ and $\delta_m = 0$, and (iii) $\delta_1 + \cdots + \delta_S = R$.

The path P_η begins at δ, and a point is on P_η if and only if the resources guarding against the nonstrategic threat are optimally allocated against that threat and those devoted to the strategic threat minimize the attacker's expected payoff. In symbols, $r \in P_\eta$ 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 nonstrategic threat, that is, 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 those sites minimize the attacker's payoff over these sites.[18] Finally, identify a specific point on the path with the amount of resources dedicated to the strategic threat. That is, $P_\eta(\sigma)$ is the point $r \in P_\eta$, 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 $\underline{\sigma} = \sum_{j \in T_A(\delta)} \delta_j$ and $\overline{\sigma} = R$. As the defender moves along P_η, it is effectively changing the mix between the strategic and the nonstrategic threat but optimally allocating however much it dedicates to each of these separate threats.

To identify an 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 nonstrategic threat, define the defender's expected loss from guarding against the strategic threat at $r = P_\eta(\sigma)$ to be $L_D(\sigma,\alpha) \equiv \sum_{j \in T_A(r(\sigma))} L_j v_j(r_j(\sigma))[\eta_j + \alpha_j(1-\eta_j)]$ and from the nonstrategic threat to be $L_\eta(\sigma,\alpha) \equiv \sum_{k \notin T_A(r(\sigma))} L_k v_k(r_k(\sigma))[\eta_k + \alpha_k(1-\eta_k)]$. The defender's expected loss can now be written as is $L(r(\sigma),\alpha) = L_\eta(\sigma,\alpha) + L_D(\sigma,\alpha)$. Then an equilibrium

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

[17] As elsewhere, a change in the marginal gain may not lead to a reallocation if the marginal conditions do not bind, and they may not when at m if $A_m v_m(0) = M_A(r)$.

[18] 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 $\max\{G_j v_j(r_j)\colon j \in T_A(r)\}$. But, $G_j v_j(r_j) = G_k v_k(r_k)$ 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 $G_n v_n(\widehat{r}_n) > G_n v_n(r_n) = \max\{G_j v_j(r_j) : j \in T_A(r)\}$ and consequently that $\widehat{r}$ is not an optimal allocation of σ.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

allocation lies on P_η and generically satisfies $-\partial L_\eta(\sigma, \widehat{\alpha})/\partial\sigma = \partial L_D(\sigma, \widehat{\alpha})/\partial\sigma$, where $\widehat{\alpha}$ is the attacker's pseudo-equlibrium strategy at r.[19]

Before stating the results in the general case, consider the simpler case in which there are H identical high-value sites for the attacker and $S - H$ identical low-value sites. The latter are so much less attractive that the attacker would never go after them.[20] Then the defender's expected loss to evenly dividing σ across the identical high-value sites and $R - \sigma$ across the identical low-value sites is $H[\eta_h + \alpha_h(1 - \eta_h)]L_h v_h(\sigma/H) + (S - H)\eta_l L_l v_l((R - \sigma)/(S - H))$. The attacker strikes each high-value site with probability $\alpha_h = 1/H$, so the first-order condition is $[\eta_h + (1/H)(1 - \eta_h)]L_h v_h'(\sigma/H) = \eta_l L_l v_l'((R - \sigma)/(S - H))$. The left side of this equation is the marginal gain to spending more against strategic threats; the right side is the marginal gain to spending more on protecting against nonstrategic threats.

Observe that if the number of potential targets subject to strategic threat (H) is very large, the probability of a strategic attack on any of them ($1/H$) is very small. The probability of a attack on this site is swamped by the nonstrategic component, and the equilibrium condition approximates $\eta_h L_h v_h'(\sigma/H) = \eta_l L_l v_l'((R - \sigma)/(S - H))$. *But this is just what the defender would do if it completely ignored the terrorist threat and simply allocated its resources optimally against the non-strategic threat.*

Turning to the general case:

Proposition 3. *Nash equilibria exist, and the defender always plays a pure strategy lying on P_η where the marginal loss of spending less on the non-strategic threat just offsets the marginal gain of spending more on the strategic threat. This condition generically satisfies $-\partial L_\eta(\sigma, \widehat{\alpha})/\partial\sigma = \partial L_D(\sigma, \widehat{\alpha})/\partial\sigma$ which is equivalent to*

$$\eta_k L_k v_k'(\widehat{r}_k) = \left(1 + \sum_{n \in T_A(\widehat{r})} \frac{\eta_n}{1 - \eta_n}\right) \times \left(\sum_{n \in T_A(\widehat{r})} \frac{1}{(1 - \eta_n)L_n v_n'(\widehat{r}_n)}\right)^{-1}$$

for any $k \notin T_A(\widehat{r})$ for which $\widehat{r}_k > 0$.[21]

UNCERTAINTY ABOUT THE ATTACKER'S PREFERRED TARGETS

The previous discussion assumed that the defender knows how the attacker ranks potential targets. But what if the defender is 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.[22]

The answer to the resource allocation question and the approach to finding the optimal allocation are quite similar to the border-defense and nonstrategic threat 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 this allocation 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 has 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 an 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 $G_j^\gamma > 0$ and τ's payoff is $G_j^\tau > 0$. Then the attackers' expected payoffs are $G^\gamma(r, \alpha) \equiv \sum_{j=1}^S \alpha_j G_j^\gamma v_j(r_j)$ and $G^\tau(r, \phi) \equiv \sum_{j=1}^S \phi_j G_j^\tau v_j(r_j)$, where α and ϕ are γ's and τ's strategies. The defender's expected loss is

$$L(r, \alpha, \phi) \equiv \sum_{j=1}^S [\mu\alpha_j + \phi_j(1 - \mu)]L_j v_j(r_j).$$

The comparative static results parallel those in the case of threats with nonstrategic components. Anything that increases the marginal gain to defending

[19] That is, the set of equilibrium allocations at which this condition does not hold has zero measure.

[20] The payoff to striking the high-value targets is $G_h v_h(R/H)$ if the defender evenly distributes its resources across these sites. As long as this payoff is larger than that of attacking an undefended low-value site; that is, as long as $G_h v_h(R/H) > G_l v_l(0)$, then the attacker will never strike one of these sites and $\alpha_l = 0$.

[21] That is the set of equilibrium allocations at which this condition does not hold has zero measure.

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

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

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 nonstrategic components, the defender moved along the curve by shifting resources between defending against the nonstrategic and strategic 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 the optimal allocation against γ and against τ. An equilibrium allocation lies on this curve 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 $T_\gamma(r') \cup T_\tau(r')$.

An equilibrium allocation equates the marginal loss of spending less on guarding against γ by increasing σ with the marginal gain to spending more defending against τ. To characterize these allocations, take the defender's expected loss from defending against γ and τ at any $\widehat{r} = P_g(\sigma)$ to be $L_\gamma(\sigma, \alpha) \equiv \mu \sum_{j \in T_A(\widehat{r})} \alpha_j L_j v_j(\widehat{r}_j(\sigma))$ and $L_\tau(\sigma, \phi) \equiv (1-\mu) \sum_{j \in T_\tau(\widehat{r})} \phi_j L_j v_j(\widehat{r}_j(\sigma))$. 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 lies on P_g and generically 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. *Bayesian Nash equilibria exist, and the defender always plays a pure strategy lying 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. This condition generically satisfies* $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma = -\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$ *which is equivalent to*:

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

CONCLUSION

Protecting a 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 an assessment of risks and vulnerabilities" (2004, 396), and Secretary Chertoff has 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 one 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 to equalize the payoff to attacking this site with the payoff to attacking the other defended sites.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

In the cases of border defense, threats with nonstrategic 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 nonstrategic 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. If the risk of a nonstrategic "attack"—a natural disaster or an accident—increases, the defender allocates more 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 adversary's 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 Lemma 1. Recalling that $M_A(r) \equiv \max\{G_1 v_1(r_1), \ldots, G_S v_S(r_S)\}$, then $\widehat{r}$ is a minmax allocation if $M_A(\widehat{r}) \leq M_A(r)$ for every allocation r. At least one minmax allocation is sure to exist because the $v_j(r)$ are continuous in r and the set of possible allocations is compact.

Arguing by contradiction to show that only one minmax allocation exists, assume that $\widehat{r} \neq r'$ both minmax the attacker. Because $\widehat{r} \neq r'$, there exists a site j such that $\widehat{r}_j \neq r'_j$. Without loss of generality suppose $\widehat{r}_j < r'_j$. Then $M_A(r') = M_A(\widehat{r})$ because both r' and $\widehat{r}$ minmax the terrorists, and $G_j v_j(\widehat{r}_j) > G_j v_j(r'_j)$ because $\widehat{r}_j < r'_j$. Hence, $M_A(r') = M_A(\widehat{r}) \geq G_j v_j(\widehat{r}_j) > G_j v_j(r'_j)$. Continuity ensures that there is an $\varepsilon > 0$ such that $M_A(r') > G_j v_j(r'_j - \varepsilon)$. This ε of resources can now be distributed across the other sites to form the allocation r^*, where $r^*_j = r'_j - \varepsilon$ and $r^*_k = r'_k + \varepsilon/(S-1)$ for all $k \neq j$. Clearly, $M_A(r^*) < M_A(r')$, contradicting the assumption that r' minimizes $M_A(r)$ and thereby establishing the claim that the minmax allocation is unique. ■

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 $G(r,\alpha)$, 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 see that the minmax allocation $\widehat{r}$ is the unique equilibrium allocation, observe first that the discussion prior to the statement of Corollary 1 in the text shows that there exists an attack profile $\widehat{\alpha}$ such that $(\widehat{r}, \widehat{\alpha})$ is a Nash equilibrium. All that remains to be shown is that the defender always plays $\widehat{r}$ in any equilibrium.

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 $M_A(r^*) > M_A(\widehat{r})$. It follows that there exists some site j such that $r^*_j > 0$ and $G_j v_j(r^*_j) < M_A(r^*)$. But any of the attacker's best replies to r^* only put positive probability on the sites that offer the highest payoff; that is, those k for which $G_k v_k(r^*_k) = M_A(r^*)$. 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 Corollary 1. The claim follows directly from the discussion in the text and a more detailed proof is omitted.

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 $\partial^2 B(b - r_j)v_j(r_j)/\partial r_j^2 > 0$.) 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, recall that $r = P_B(\sigma)$ if and only if r minimizes $M_A(r')$ subject to $r'_1 + \cdots + r'_S \leq \sigma$. Now suppose $\widehat{r} \notin P_B$ is an equilibrium allocation. Then $\widehat{r}_k > 0$ for some $k \notin T_A(\widehat{r})$, and the defender can clearly profitably deviate from $\widehat{r}$ by reallocating $\widehat{r}_k$ across $T_A(\widehat{r})$.

To see that an equilibrium allocation satisfies the condition stated in Proposition 2 except possibly at a set of points along P_B with zero measure, say that j becomes a target (i.e., a best reply) 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 and define E^+ to be E along with the endpoints $P_B(0)$ and $P_B(R)$. Then if $\widehat{r} = P_\eta(\widehat{\sigma}) \notin E^+$ is an equilibrium allocation, it satisfies the condition described in Proposition 2.[23]

Since $\widehat{r}$ is an equilibrium allocation, there must be a strategy $\widehat{\alpha}$ such that $(\widehat{r}, \widehat{\alpha})$ is Nash. 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 $j \in T_A(\widehat{r})$, $\widehat{r}_j > 0$ because $\widehat{r} \notin E$. Hence, $\widehat{\alpha}_j L_j v'_j(\widehat{r}_j) = \widehat{\alpha}_i L_i v'_i(\widehat{r}_i)$ for all $i, j \in T_A$. Solving these equations for $\widehat{\alpha}_i$ in terms of $\widehat{\alpha}_j$, summing the $\widehat{\alpha}_i$ over $T_A(\widehat{r})$, and using the fact that the sum equals one yield $\widehat{\alpha}_j = (\sum_{i \in T_A(\widehat{r})} L_j v'_j / L_i v'_i)^{-1}$.

Since $\widehat{r}$ is not an endpoint, positive amounts are being spent on both site and border defense. So $\widehat{r}$ must satisfy the marginal condition $\partial L/\partial r_j = \partial L/\partial b$ or $\widehat{\alpha}_j L_j v'_j(\widehat{r}_j)\beta(R - \widehat{\sigma}) = \beta'(R - \widehat{\sigma}) \sum_{n \in T_A} \widehat{\alpha}_n L_n v_n(\widehat{r}_n)$ for $j \in T_A(\widehat{r})$. Substituting for the $\widehat{\alpha}_n$'s gives $\beta(R - \sigma)/\beta'(R - \sigma) = \sum_{n \in T_A(\widehat{r})} v_n(\widehat{r}_n)/v'_n(\widehat{r}_n)$.

This expression is moreover equivalent to $\partial \Pi_D(\sigma, \widehat{\alpha})/\partial\sigma\beta(R - \sigma) = \Pi_D(\sigma, \widehat{\alpha})\beta'(R - \sigma)$. Differentiation yields $\partial \Pi_D(\sigma, \widehat{\alpha})/\partial\sigma = \sum_{j \in T_A(\widehat{r})} \widehat{\alpha}_j L_j v'_j(\widehat{r}_j(\widehat{\sigma}))d\widehat{r}_j/d\sigma$. Substituting for $\widehat{\alpha}_j$ gives $\partial \Pi_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = (\sum_{n \in T_A(\widehat{r})}[L_n v'_n(\widehat{r}_n)]^{-1})^{-1} \sum_{j \in T_A(\widehat{r})} d\widehat{r}_j/d\sigma$. But $\sum_{j \in T_A(\widehat{r})} \widehat{r}_j = \widehat{\sigma}$, so $\sum_{j \in T_A(\widehat{r})} d\widehat{r}_j/d\sigma = 1$ and consequently $\partial \Pi_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = (\sum_{n \in T_A(\widehat{r})}[L_n v'_n(\widehat{r}_n)]^{-1})^{-1}$.

Now multiply $\beta(R - \sigma)/\beta'(R - \sigma) = \sum_{n \in T_A(\widehat{r})} v_n(\widehat{r}_n)/v'_n(\widehat{r}_n)$ by $\partial \Pi_D/\partial\sigma$ to obtain $\beta(R - \widehat{\sigma})\partial \Pi_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = \beta'(R - \widehat{\sigma}) \sum_{n \in T_A(\widehat{r})}(\sum_{k \in T_A(\widehat{r})}[L_k v'_k(\widehat{r}_k)]^{-1})^{-1} v_n(\widehat{r}_n)/v'_n(\widehat{r}_n)$. The expressions for $\widehat{\alpha}_j$ then yield $\beta(R - \widehat{\sigma})\partial \Pi_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = \sum_{n=1}^{S} \widehat{\alpha} L_n v_n(\widehat{r}_n)\beta'(R - \widehat{\sigma})$ or $\beta(R - \widehat{\sigma})\partial \Pi_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = \beta'(R - \widehat{\sigma})\Pi_D(\widehat{\sigma}, \widehat{\alpha})$. ■

[23] If a site, say e, enters T_A at $\widehat{\sigma}$, the summations in the expression for $\partial \Pi_D(\sigma, \widehat{\alpha})/\partial\sigma$ for $\sigma \in (\widehat{\sigma}, \widehat{\sigma} + \varepsilon)$ and for $\sigma \in (\widehat{\sigma}, \widehat{\sigma} - \varepsilon)$ are taken over different sets. The continuity of the individual terms in Π_D will therefore not ensure the continuity of the sums, and $\lim_{\sigma \to \widehat{\sigma}^-} \partial \Pi_D(\sigma, \widehat{\alpha})/\partial\sigma$ will generally not equal $\lim_{\sigma \to \widehat{\sigma}^+} \partial \Pi_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma$.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

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_η, suppose the contrary. Then there would be an equilibrium allocation $\widehat{r} \notin P_\eta$ such that one of the following two conditions does not hold: $L_j v_j'(\widehat{r}_j)\eta_j = L_k v_k'(\widehat{r}_k)\eta_k$ whenever $\widehat{r}_j > 0, \widehat{r}_k > 0$, and $j, k \notin T_A(\widehat{r})$ and, second, $L_j v_j'(\widehat{r}_j)\eta_j \leq L_m v_m'(\widehat{r}_m)\eta_m$ 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_η. 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 of these conditions fails to hold, the defender can clearly reduce its expected loss by reallocating its resources between j and k. This contradiction leaves $\widehat{r} \in T_A(\widehat{r})$.

To see that an equilibrium allocation satisfies the condition stated in Proposition 3 except possibly at a set of points along P_η with zero measure, let Z be the set of points where a site k enters T_A at zero allocation; that is, $r_k = 0$ when k enters. Take Z^+ to be Z along with the endpoints $P_\eta(\underline{\sigma})$ and $P_\eta(R)$. Then if $\widehat{r} = P_\eta(\widehat{\sigma}) \notin Z^+$ is an equilibrium allocation, it satisfies the condition described in Proposition 3.

Since $\widehat{r}$ is an equilibrium allocation, there must be a strategy $\widehat{\alpha}$ such that $(\widehat{r}, \widehat{\alpha})$ is Nash. 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 $j \in T_A(\widehat{r})$, $\widehat{r} \notin Z$ means $\widehat{r}_j > 0$ for all $j \in T_A(\widehat{r})$. Hence, $L_i v_i'(\widehat{r}_i)[\eta_i + \widehat{\alpha}_i(1-\eta_i)] = L_j v_j'(\widehat{r}_j)[\eta_j + \widehat{\alpha}_j(1-\eta_j)]$ for all $i, j \in T_A(\widehat{r})$; otherwise the defender could profitably reallocate across sites in $T_A(\widehat{r})$. Solving these equations for $\widehat{\alpha}_j$ in terms of $\widehat{\alpha}_i$, summing the $\widehat{\alpha}_j$ over $T_A(\widehat{r})$, using the fact that the sum equals one, and rearranging terms yield $\eta_j + \widehat{\alpha}_j(1-\eta_j) = (1 + \sum_{i \in T_A(\widehat{r})} \eta_i/(1-\eta_i))\,(\sum_{i \in T_A(\widehat{r})}[L_j v_j'/((1-\eta_i)L_i v_i')])^{-1}$ for $j \in T_A(\widehat{r}(\sigma))$.

Since $\widehat{r}$ is not an endpoint, $\widehat{\sigma} < R$, so there exists a site $k \notin T_A(\widehat{r})$ such that $\widehat{r}_k > 0$. Again because $\widehat{r}$ is a best response to $\widehat{\alpha}$, $L_j v_j'(\widehat{r}_j)[\eta_j + \widehat{\alpha}_j(1-\eta_j)] = \eta_k L_k v_k'$. Otherwise the defender could profitably reallocate between j and k. Substituting the expression for $\eta_j + \widehat{\alpha}_j(1-\eta_j)$ gives $\eta_k L_k v_k'(\widehat{r}_k) = (1 + \sum_{i \in T_A(\widehat{r})} \eta_i/(1-\eta_i))\,(\sum_{i \in T_A(\widehat{r})}[(1-\eta_i)L_i v_i')])^{-1})^{-1}$.

This expression is equivalent to $-\partial L_\eta(\sigma, \widehat{\alpha})/\partial\sigma = \partial L_D(\sigma, \widehat{\alpha})/\partial\sigma$. Since $L_i v_i'(\widehat{r}_i)[\eta_i + \widehat{\alpha}_i(1-\eta_i)] = L_j v_j'(\widehat{r}_j)[\eta_j + \widehat{\alpha}_j(1-\eta_j)]$ for all $i, j \in T_A(\widehat{r})$, $\partial L_\eta(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = \sum_{j \in T_A(\widehat{r})}[\eta_j + \widehat{\alpha}_j(1-\eta_j)]L_j v_j'(\widehat{r}_j)d\widehat{r}_j/d\sigma = [\eta_j + \widehat{\alpha}_j(1-\eta_j)]L_j v_j'(\widehat{r}_j)\sum_{j \in T_A(\widehat{r})} d\widehat{r}_j/d\sigma$. And, $\sum_{k \in T_A(\widehat{r})} \widehat{r}_j = \widehat{\sigma}$, so $\sum_{k \in T_A(\widehat{r})} d\widehat{r}_j/d\sigma = 1$. This along with the expression for $[\eta_j + \widehat{\alpha}_j(1-\eta_j)]$ yield $\partial L_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma = [1 + \sum_{k \in T_A(\widehat{r})} \eta_k/(1-\eta_k)]\;[\sum_{k \in T_A(\widehat{r})}[(1-\eta_k)L_k v_k']^{-1}]^{-1}$.

Turning to the marginal loss from investing less against the non-strategic threat, $\partial L_D(\widehat{\sigma}, \widehat{\alpha})/\partial\sigma \equiv \sum_{j \notin T_A(r)} \eta_j L_j v_j'(\widehat{r}_j(\widehat{\sigma}))d\widehat{r}_j/d\sigma$. By construction, $\eta_j L_j v_j'(\widehat{r}_j) = \eta_k L_k v_k'(\widehat{r}_k)$ along P_η whenever $r_j > 0$, $r_k > 0$, and $j, k \notin T_A(\widehat{r})$. The fact that $\sum_{j \notin T_A(r)} d\widehat{r}_j/d\sigma = -1$ then implies $\partial L_D(\sigma, \widehat{\alpha})/\partial\sigma = -\eta_k L_k v_k'(\widehat{r}_k)$. ■

Proof of Proposition 4. The proof closely parallels that of Proposition 3 and will only be sketched. Once again the argument in the proof of Proposition 1 ensures that equilibria exist and that the defender never mixes in equilibrium.

Any equilibrium allocation must clearly 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 $T_\gamma(r') \cup T_\tau(r')$.

To see that an equilibrium allocation satisfies the condition stated in Proposition 4 except possibly at a set of points along P_g with zero measure, let Z be the set of points points where sites leave T_γ or enter T_τ. Take Z^+ to be Z and the endpoints, $P_g(\underline{\sigma})$ and $P_g(\overline{\sigma})$. Then if $\widehat{r} = P_\eta(\widehat{\sigma}) \notin Z^+$ is an equilibrium allocation, it satisfies the condition described in Proposition 4.

Since $\widehat{r}$ is an equilibrium allocation, there must be strategies $\widehat{\alpha}$ and $\widehat{\phi}$ such that $(\widehat{r}, \widehat{\alpha}, \widehat{\phi})$ is Nash. Since $\widehat{r} \notin Z, \widehat{r}_j > 0$ for all $j \in T_\gamma(\widehat{r})$ and $\widehat{r}_k > 0$ for all $k \in T_\tau(\widehat{r})$. What is more, $\widehat{r} \notin Z$ implies $T_\gamma(\widehat{r}) \cap T_\tau(\widehat{r}) = \varnothing$. Hence, $\widehat{\alpha}_j = 0$ for $j \notin T_\gamma(\widehat{r})$, $\widehat{\alpha}_j = (\sum_{n \in T_\gamma(\widehat{r})} L_j v_j'(r_j)/[L_n v_n'(r_n)])^{-1}$ for $j \in T_\gamma(r(\sigma))$, $\widehat{\phi}_j = 0$ for $j \notin T_\tau(r(\sigma))$, and $\widehat{\phi}_j = (\sum_{n \in T_\tau(r(\sigma))} L_j v_j'(r_j)/[L_n v_n'(r_n)])^{-1}$ for $j \in T_\tau(r(\sigma))$.

That $\widehat{r}$ is not an endpoint also implies there are sites $j \in T_\gamma(\widehat{r})$ and $k \in T_\tau(\widehat{r})$ such that $r_j > 0$ and $r_k > 0$. Because $\widehat{r}$ is a best reply to $\widehat{\alpha}$ and $\widehat{\phi}$, $\partial L/\partial r_j = \partial L/\partial r_k$ or $\mu\widehat{\alpha}_j L_j v_j'(\widehat{r}_j) = (1-\mu)\widehat{\alpha}_k L_k v_k'(\widehat{r}_k)$. Substituting the expressions for $\widehat{\alpha}_j$ and $\widehat{\phi}_k$ gives $\mu(\sum_{n \in T_\gamma(r(\sigma))}(L_n v_n')^{-1})^{-1} = (1-\mu)(\sum_{n \in T_\tau(r(\sigma))}(L_n v_n')^{-1})^{-1}$ as stated in the proposition.

The marginal condition $\mu\widehat{\alpha}_j L_j v_j'(\widehat{r}_j) = (1-\mu)\widehat{\alpha}_k L_k v_k'(\widehat{r}_k)$ also implies $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma = -\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma$. Substituting for $\widehat{\alpha}$ yields $\partial L_\gamma(\sigma, \widehat{\alpha})/\partial\sigma = \mu \sum_{n \in T_\gamma(r)} \widehat{\alpha}_n L_n v_n'(r_n) dr_n/d\sigma = -\mu(\sum_{n \in T_\gamma(r)}[L_n v_n'(r_n)]^{-1})^{-1} = -\mu\widehat{\alpha}_n L_n v_n'(r_n)$. Similarly, $\partial L_\tau(\sigma, \widehat{\phi})/\partial\sigma = (1-\mu)\;(\sum_{n \in T_\gamma(r)}[L_n v_n'(r_n)]^{-1})^{-1} = (1-\mu)\,\widehat{\phi}_n L_n v_n'(r_n)$. ■

REFERENCES

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

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, DC 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, DC: 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, January 3. Available at: http://www.dhs.gov/xnews/releases/press_release_0825.shtm (accessed November 3, 2006).

Chyba, Christopher F. 2002. "Toward Biological Security," *Foreign Affairs* (May/June): 122–36.

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

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

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press

DHS. 2006a. "Draft National Infrastructure Protection Plan, v2.0." Washington, DC: Department of Homeland Security. Available at: www.fas.org /irp/agency/dhs/nipp110205.pdf.
DHS. 2006b. "Progress in Developing the National Asset Database." Washington, DC: Department of Homeland Security, Office of Inspector General (June).
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 *Researching Terrorism Trends, Achievments, Failures*, ed. Andrew Silke and G. Ilardi. 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.
GAO. 2005. "Risk Management." Washington, DC: Government Accountability Office. Available at: www.gao.gov/cgi-bin/getrpt?GAO-06-91.
Golman, Russell, and Scott E. Page. 2006. "General Blotto: Games of Allocative Strategic Mismatch," Manuscript, Center for Complex Systems, University of Michigan. Accessed November 3, 2006 at http://www.cscs.umich.edu/~spage/papers_files/blotto.pdf.
Mayer-Schönberger, Viktor. 2003. "Emergency Communications: The Quest for Interoperability in the United States and Europe." In *Countering Terrorism*, ed. Arnold M. Howitt and Robyn L. Pangi. Cambridge, MA: MIT Press.
Moteff, John. 2004. "Risk Management and Critical Infrastructure Protection: Assessing, Integrating, and Managing Threats, Vulnerablities and Consequences." Washington, DC: Congressional Research Service.
Moteff, John D. 2006. "Critical Infrastructures: Background, Policy, and Implementation." Washington, DC: 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, DC: Brookings.
OMB. 2003. "2003 Report to Congress on Combating Terrorism." Washington, DC: 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, DC: Office of Management and Budget.
OMB. 2005. "Analytical Perspectives, Budget of the United States Government, Fiscal Year 2006." Washington, DC: Office of Management and Budget.
Powell, Robert. 2006. "Defending Against *Strategic* Terrorists Over the Long Run." Manuscript, Department of Political Science, U.C. Berkeley.
Roberson, Brian. 2006. "The Colonel Blotto Game," *Economic Theory* 29: 1–24.
Rosendorff, Peter, and Todd Sandler. 2004. "Too Much of a Good Thing? The Proactive Response Dilemma." *Journal of Conflict Resolution* 48 (October): 657–71.
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.
Tukey, John W. 1949. "A Problem of Strategy." *Econometrica* 17: 53.
White House. 2002, "National Strategy for Homeland Security." Washington, DC: 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, DC: White House, February. Available at www.whitehouse.gov/pcipb/physical.html
Willis, Henry H., Andrew R. Morral, Terrance K. Kelly, and Jamison J. Medby. 2005. *Estimating Terrorism Risk*. Santa Monica, CA: Rand Corporation.

https://doi.org/10.1017/S0003055407070244 Published online by Cambridge University Press