Deterring and Defending Against Strategic Attackers: Deciding How Much to Spend and on What*

Robert Powell

April 2008

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

* For helpful comments, criticisms and discussion, I thank Ernesto Dal Bó, Santiago Oliveros, Larry Samuelson, and Jacob Shapiro.

Deterring and Defending Against Strategic Attackers:
Deciding How Much to Spend and on What

Abstract

How much should a defender spend on defense rather than something else like health care, education, or other social purposes? And how should a defender allocate however much it decides to spend across the multiple sites it is trying to protect? This paper characterizes the generically unique equlibrum outcome of a sequential game in which a defender first has to decide how much to spend on defense and what to spend it on. The more that a defender dedicates to hardening or protecting a specific site, the less likely an attack on that site is to succeed. After the defender moves, the attacker, e.g., a terrorist group, chooses how much effort or resources to devote to attacking the defender and how to allocate that effort to the various sites. The more effort the attacker devotes to any given site, the more likely the attack on that site is to succeed. The analysis shows that (i) the marginal effect of spending on the defender's payoff can be decomposed into the sum of a defensive effect, a deterrent effect, and a cost effect; (ii) the defender's allocation and level-of-spending problems are separable in that regardless of how much the defender decides to spend, it allocates that amount in the same general way; (iii) the optimal (equilibrium) allocation of the defender's resources minmaxes the attacker; and (iv) the defender's equilibrium level of spending minimizes a simple loss function.

Deterring and Defending Against Strategic Attackers:
Deciding How Much to Spend and On What

Defenders must deal with two related resource-allocation questions when facing a potential attacker that might strike in different places. First, how much should the defender spend on defense rather than something else like health care, education, or other social purposes? And, second, how should a defender allocate however much it decides to spend across the multiple sites it is trying to protect? Although very old questions, they have received renewed attention following the attacks of September 11, 2001 (e.g., 911 Public Discourse Project 2005, DHS 2006, Flynn 2004, Masse *et. al.* 2007, Mueller 2006, O'Hanlon *et. al.* 2003).

To address these questions, we study a sequential game in which a defender first has to decide how much to spend on defense and what to spend it on. The more that a defender dedicates to hardening or protecting a specific site, the less likely an attack on that site is to succeed. After the defender moves, the attacker, e.g., a terrorist group, chooses how much effort or resources to devote to attacking the defender and how to allocate that effort to the various sites. The more effort the attacker devotes to any given site, the more likely the attack on that site is to succeed.

This game has a (generically) unique equilibrium and five key results follow. First, the analysis shows that we can decompose the effect that additional spending has on the defender's payoff into three basic components: a *defense effect*, a *deterrent effect*, and a *cost effect*. The defense effect is the *ceteris paribus* gain the defender derives from the fact that the sites are more difficult to destroy and that an attack is less likely to succeed. The deterrent effect is the *ceteris paribus* benefit resulting from the attacker's choosing to devote less effort to attacking the defender. The cost effect is the disutility from spending less on other social ends. This decomposition and, especially, the distinction between the defense and deterrent effects offer a clarifying framework for the growing policy debate about whether it is possible to deter terrorists, a debate which often confounds deterrence with defense (e.g., Schmitt and Shanker 2008).

Second, the defender's allocation and level-of-spending decisions turn out to be separable. However much the defender chooses to spend on defense, it allocates those resources in the same general way.

Third, a simple, intuitive principle underlies the allocation rule. Suppose that a defender has decided to spend an amount R on defense but has not yet allocated any of it. Given this null allocation, the attacker will strike the most attractive target, i.e., the site offering the attacker its highest expected payoff. The defender, therefore, should start spending R by investing in hardening or protecting this site. But the more the defender spends on this site, the lower the attacker's payoff to striking this site and the less attractive it becomes. Eventually, going after this site will be no more attractive than striking what was originally the second most attractive target. At this point, the defender must invest in hardening both of these sites so that neither is a more appealing target than the other. The more the defender spends on these two sites, the lower the payoff to attacking them and the less attractive they become. Eventually, attacking either of these two sites will offer an expected payoff no higher than going after what was originally the third most attractive site. From here on, the defender will have to invest in all three of those sites so that no one of them is more attractive than any other. The defender continues to allocate R 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 brief, the defender's optimal allocation of R minmaxes the attacker. That is, the defender allocates R so as to minimize the attacker's maximum payoff.

Fourth, the optimal (equilibrium) level of spending minimizes a simple loss function and is generically unique. Because minmaxing the attacker equalizes the attacker's payoffs to the most attractive sites, the attacker will be indifferent to going after anyone of several sites. Nevertheless, equilibrium considerations imply that the attacker allocates its effort in the way that imposes the lowest expected loss on the defender. The defender's optimal level of spending minimizes these losses along with the cost of spending on defense rather than some other social end.

Subject to a significant qualification, the defender's optimal allocation equates the marginal benefit of spending more on defense with the marginal cost of having less to spend on other social ends. Importantly, the marginal gains from greater spending decrease as the number of protected sites increases because any additional spending has to be distributed across all of these sites. As the number of defended sites becomes large, the marginal gains become very small.

As for the qualification, the allocation that equates marginal gains with marginal losses is always a lower bound on the optimal level of spending and often is the optimum. But the defender may want to spend more than this amount in some circumstances. The defender can influence the site the attacker strikes to some degree. Spending more on one site and less on another makes the former a less attractive target. The defender will want to spend more than the allocation that equates marginal costs and benefits if by doing so the defender can induce the attacker to strike a site the defender values much less than the site the attacker would have hit had the defender equated marginal costs with marginal benefits.

The final result centers on comparative statics. As the attacker becomes more determined to strike the defender, the equilibrium levels of spending and effort increase. But the attacker's determination has no effect on spending and effort allocations.

The present analysis is most closely related to Powell (2007); Bier, Oliveros, and Samuelson (2007); and Bueno de Mesquita (2007). A defender in Powell's model has to decide how to allocate its resources across an arbitrarily large number of sites and the attacker has to choose which site to strike. Unlike the model developed here, both actors move simultaneously; the defender's level of total spending is exogenous; and the attacker can only decide where to strike, not how much to invest in attacking. The attacker moreover must go after something. Not attacking is not an option. Hence, it is impossible to study deterrence with that model. Bier, Oliveros, and Samuelson focus on the allocation problem when the defender is unsure of the attacker's payoffs. In their game, the defender does determine the level of spending endogenously. The defender, however, only has two sites to protect, and the attacker again must go after one of them. Deterrence therefore is also moot in that formulation. Both the defender and attacker do make allocation decisions in Bueno de Mesquita, but neither chooses the level of expenditure. Indeed, the present study model appears to be first to analyze a model in which the defender and attacker choose both how much to spend and how to allocate those resources.

The next section presents the model. The subsequent section describes the equilibrium and main results less formally and more intuitively. There follows a more rigorous analysis of the defender's allocation problem when the level of total spending is exogenously fixed. The last section shows that the allocation and level decisions are separable and characterizes the optimal level of spending.

A Model

A defender is trying to protect N sites and must decide how to much to spend on protecting against an attack and how to allocate those resources among the sites it is trying to protect. The more the defender spends on a particular site, the "harder" that site becomes and the less likely an attack on that site is to succeed. After observing the defender's allocation, an attacker decides how much effort to devote to attacking the defender and how it will allocate that effort across the N sites.

A strategy for the defender in this game is simply an allocation $r = (r_1, ..., r_N)$ where $r_j \geq 0$ is the amount spent on defending site j and $R = \sum_{j=1}^{N} r_j$ is the total spent on defense. Analogously, the attacker's strategy specifies how much effort or, equivalently, the amount of resources the attacker puts into attacking each site after observing any possible allocation r. More precisely, a strategy for the attacker is a function $e(r) = (e_1(r), ..., e_N(r))$ where $e_j(r) \geq 0$ is the effort the attacker puts into striking site j and $E(r) \equiv \sum_{j=1}^{N} e_j(r)$ is the total effort the attacker invests in striking the defender. (An equivalent interpretation is that $e_j(r)$ measures the resources the attacker dedicates to striking site j.)

Let $\lambda_j > 0$ denote the defender's loss and $\gamma_j > 0$ the attacker's gain if site j is successfully attacked. If the attack fails, the defender loses nothing and the attacker gains nothing. (We assume for simplicity that an attack either succeeds or fails.)

The more the defender spends on a site, the less likely an attack on that site is to succeed. Conversely, the more effort the attacker dedicates to going after a site, the higher the probability of success. Formally, let $V_j(r_j, e_j)$ be the probability that an attack on j succeeds with $\partial V_j/\partial r_j < 0$ and $\partial^2 V_j/\partial r_j^2 \geq 0$ as long as j is imperfectly defended (i.e., $V_j(r_j, e_j) > 0$) and $\partial V_j/\partial e_j > 0$ and $\partial^2 V_j/\partial e_j^2 \leq 0$ as long as $V_j(r_j, e_j) < 1$.

Spending on defense means diverting resources from other social ends. These costs are assumed to rise and at an increasing rate. More concretely, take the cost to devoting R to defense to be $c_D(R)$ with $c_D'(0) = 0$, $c_D'(R) > 0$ for $R > 0$, and $c_D''(R) > 0$.[1] Marshalling the effort to carry out an attack is also costly. Let $c_A(E)$ denote cost of investing E in attacking the defender, and assume that these costs are increasing at an increasing rate. Formally, $c_A'(0) \geq 0$, $c_A'(E) > 0$ for $E > 0$, $c_A''(E) > 0$, and $c_A'''(E) \leq 0$.

The substantive significance of assuming $c_A''' \leq 0$ will be discussed below. When $c_A'(0) > 0$, the defender at least in principle can fully deter the attacker. If the defender spends enough to drive the attacker's marginal return on investing in an attack below below $c_A'(0)$, then the optimal level of effort will be $E = 0$ as shown below.

Putting all of this together, the defender's loss and the attacker's gain to the allocations (r, e) are $L(r, e) = \sum_{j=1}^{N} \lambda_j V(r_j, e_j) + c_D(R)$ and $G(r, e) = \sum_{j=1}^{N} \gamma_j V(r_j, e_j) - c_A(E)$ respectively.

We now make an important assumption which greatly simplifies the analysis. The marginal effect of additional effort on the vulnerability of a site is independent of the level of effort already being exerted. This is equivalent to assuming that the vulnerability of a site is proportional to the effort devoted to attacking that site.

Assumption 1: *For all sites j, $V_j(r_j, e_j) = v_j(r_j)e_j$ where v_j is decreasing in r_j and convex ($v_j' < 0$ and $v_j'' \geq 0$).*

The significance of this assumption is that the defender's allocation fixes the attacker's marginal gain to striking each site. That is, the marginal gain to investing in attacking site j only depends on the defender's allocation r and not on e: $\partial G/\partial e_j = \partial \gamma_j V_j(r_j, e_j)/\partial e_j = \gamma_j v_j(r_j)$. After observing the defender's allocation and these marginal returns, the at-

[1] The assumption that $c_D'(0) = 0$ ensures that the there is no corner solution, i.e., that the optimal value of spending is not zero. This is assumption is made purely for convenience.

tacker will then invest its effort in going after the sites offering the highest marginal returns on that effort.

Finally, we also assume that resources are scarce in that the defender cannot defend everything perfectly. However the defender distributes R, at least one site will still be vulnerable. Formally,

ASSUMPTION 2 (SCARCE RESOURCES): *For any R and any allocation r of R, there exists at least one site j such that $v(r_j) > 0$.*

This is the simplest formulation of the idea that it is impossible to protect everything at a reasonable cost. Another formulation which is more tedious to state precisely and more difficult to work with buy yields the same results is that while it might be possible to defend everything at some very high level of spending R', the cost of sustaining this level of expenditure, $c_D(R')$, is so high that the defender would never choose to do so.

An Overview of the Results

This section describes the equilibrium and discusses the main results. It emphasizes insights and intuitions, giving the minimum of attention to technical details. Those follow in subsequent sections.

In a subgame perfect equilibrium, the attacker must play optimally after observing any possible allocation r. To maximize its payoff, the attacker allocates its effort to the sites offering the highest return. Let $M_A(r)$ be the attacker's maximum marginal return given r. That is, $M_A(r) = \max\{\partial G/\partial e_1, ..., \partial G/\partial e_N\} = \max\{\gamma_1 v_1(r_1), ..., \gamma_N v_N(r_N)\}$. Also let $\text{br}_A(r)$ be the sites offering the maximum marginal return: $\text{br}_A(r) \equiv \{j : \gamma_j v_j(r_j) = M_A(r)\}$. Then the attacker will only allocate effort to sites in $\text{br}_A(r)$ if it invests any effort at all in striking the defender. At the slight abuse of the definition, it will be convenient to refer to the sites in $\text{br}_A(r)$ as the attacker's best replies to r.

Let e be the attacker's allocation where $e_j = 0$ if $j \notin \text{br}_A(r)$. Then the attacker's payoff is $G(r, e) = \sum_{j \in \text{br}_A(r)} \gamma_j v_j(r_j) e_j + c_A(E)$ where $E = \sum_{j \in \text{br}_A(r)} e_j$. Since the marginal return is $M_A(r) = \gamma_j v_j(r_j)$ at all sites $j \in \text{br}_A(r)$, the attacker's gain simplifies to $G(r, e) = \sum_{j \in \text{br}_A(r)} M_A(r) e_j - c_A(E) = M_A(r) \sum_{j \in \text{br}_A(r)} e_j - c_A(E) = M_A(r)E - c_A(E) = \gamma_k v_k(r_k)E + c_A(E)$ for any $k \in \text{br}_A(r)$.

The attacker's optimal allocation maximizes these gains. If the marginal gain to any effort is less than the marginal cost of that effort (i.e., if $M_A(r) < c'_A(0)$), the optimal level of effort devoted to striking that defender is $S^* = 0$ and the attacker is fully deterred. If the marginal gain is higher than $c'_A(0)$, the optimal level of effort equates the marginal gain with the marginal cost and satisfies the first-order condition $M_A(r) = c'_A(S^*)$.[2] Solving for the level of effort gives $S^* = {c'_A}^{-1}(M_A(r))$.

Hence, the optimal level of effort for any allocation r is $S^*(r) = \max\{0, {c'_A}^{-1}(M_A(r))\}$. Since the marginal return to going after any site in $\text{br}_A(r)$ is the same as going after any other site (i.e., $\gamma_k v_k(r_k) = \gamma_j v_j(r_j)$ for all sites j, k in $\text{br}_A(r)$), it does not matter how the attacker allocates $S^*(r)$ among the sites in $\text{br}_A(r)$. Every allocation brings the same maximal payoff, namely, $M_A(r)S^*(r) + c_A(S^*(r))$.

The defender's goal is to minimize its losses given that the attacker exerts $S^*(r)$ in response to r. Proposition 3 shows that the defender's allocation and level-of-spending problems are separable. Regardless of the level of spending the defender chooses, it allocates its resources it in the same way: the defender minimizes the attacker's maximum marginal return which, given Assumption 1, is equivalent to minmaxing the attacker. Formally, let $\Delta(R)$ be the set of feasible allocations of R: $\Delta(R) = \{r : \sum_{j=1}^{N} r_j \leq R$ and $r_j \geq 0$ for all $j\}$. Then the optimal allocation of R is the minmax allocation $r^*(R)$ which solves $\min_{r \in \Delta(R)} M_A(r)$. (Lemma 1 demonstrates that this minimization problem has a unique solution.)

Determining the defender's equilibrium strategy now reduces to finding the level of spending R which when allocated according to $r^*(R)$ minimizes the defender's loses given that the attacker will exert a total level of effort $S^*(r^*(R))$. Even though the attacker is indifferent to the way that $S^*(r^*(R))$ is distributed across the sites in $\text{br}_A(r)$, the defender generally will not be indifferent. Suppose, for example, that $\lambda_k v_k(r_k) < \lambda_j v_j(r_j)$ for an j and k in $\text{br}_A(r)$. Then defender's losses would be lower if the attacker went after k rather than j: $\lambda_k v_k(r_k)S^*(r) + c_D(R) < \lambda_j v_j(r_j)S^*(r) + c_D(R)$.

Proposition 2 shows that the attacker, despite its indifference, must nevertheless allo-

[2] This follows from the assumption that $c''_A > 0$.

cate $S^*(r)$ in a specific way in equilibrium. The attacker must respond to the equilibrium allocation r by distributing $S^*(r)$ across the sites in $\text{br}_A(r)$ in the way that is most favorable to the defender. The intuition underlying this result has a parallel in the ultimatum game in which one player makes a take-it-or-leave-it offer to a second player. In the unique subgame perfect equilibrium of that game, the player receiving the offer player is indifferent to between accepting and rejecting it. Nevertheless, that player acts in the way that is most favorable to the player making the offer by accepting.[3]

To formalize what it means for the attacker to distribute $S^*(r)$ in the way that is most favorable to the defender, let $m_D(r)$ be the defender's minimum marginal loss among the sites in $\text{br}_A(r) : m_D(r) = \min\{\lambda_j v_j(r_j) : j \in \text{br}_A(r)\}$. Take $\mu(r)$ to be the sites that minimize the defender's losses $\mu(r) \equiv \{j \in \text{br}_A(r) : \lambda_j v_j(r_j) = m_D(r)\}$. Then as Proposition 2 below demonstrates, if the defender's equilibrium allocation is r, the attacker must distribute $S^*(r)$ over the sites in $\mu(r)$. That is, $e_j(r) > 0$ only if $j \in \mu(r)$.

This implies that if the defender spends R on in equilibrium, it will minmax the attacker with allocation $r^*(R)$ in response to which the attacker distributes $S^*(r^*(R))$ over the sites in $\mu(r^*(R))$. To reduce the notational burden, let $E^*(R) \equiv S^*(r^*(R))$. Then the defender's loss is $L(R) = \sum_{j=1}^{N} \lambda_j v_j(r_j^*(R)) e_j(r^*(R)) + c_D(R)$. Using $e_j(r_j^*(R)) = 0$ for all $j \notin \mu(r^*(R))$ and $\lambda_j v_j(r_j^*(R)) = m_D(r^*(R))$ for all $j \in \mu(r^*(R))$ to simplify $L(R)$ gives:

$$
\begin{aligned}
L(R) &= \sum_{j \in \mu(r^*(R))} m_D(r^*(R)) e_j(r^*(R)) + c_D(R) \\
&= m_D(r^*(R)) \sum_{j \in \mu(r^*(R))} e_j(r^*(R)) + c_D(R) \\
&= \lambda_k v_k(r_k^*(R)) E^*(R) + c_D(R) \qquad (1)
\end{aligned}
$$

where k is any element in $\mu(r^*(R))$. Moreover, $\mu(r^*(R))$ generally contains only one

[3] Although the second player is indifferent, rejecting an offer with positive probability cannot be part of an equilibrium. If the receiver followed such a strategy, then offering zero would not be an equilibrium offer as the offerer could do better by buying a "yes" for sure with an offer of slighlty more than zero. Indeed, there is no best offer and hence no equilibrium.

element, so the attacker's allocation of E^* is unique.[4]

The defender's equilibrium level of spending minimizes the losses in Eq (1). Differentiating L with respect to R yields the marginal effect of greater spending on the defender's losses. Formally,

$$\frac{dL}{dR} \;=\; \underbrace{\frac{d\lambda_k v_k(r_k^*(R))}{dR} E^*(R)}_{\text{defense effect}} \;+\; \underbrace{\lambda_k v_k(r_k^*(R))\frac{dE^*(R)}{dR}}_{\text{deterrent effect}} \;+\; \underbrace{\frac{dc_D(R)}{dR}}_{\text{cost effect}} \qquad (2)$$

Equation (2) shows that we can decompose the marginal effects of greater spending into the sum of three more basic components, the *defense effect*, the *deterrent effect*, and the *cost effect*.

The notions of deterrence and defense are often conflated, and this can lead to confusing policy discusions. For example, even if it were impossible to deter a terrorist group, there still may be substantial benefits to increased spending because the defense effect is large. Defensive measures reduce an attacker's physical ability to carry out a successful attack. A pure deterrent, by contrast, affects a state's willingness to strike but not its physical ability to do so. States typically try to deter through threats of retaliation in the event of an attack.

The clearest discussion of the distinction between deterrence and defense can be traced back to some of the early work on nuclear deterrence theory. The advent of secure, second-strike forces during the Cold War rendered defense impossible. Neither the United States nor the Soviet Union could physically protect itself from a devastating nuclear attack should the other state decide to launch one. In this technological condition of mutually assured destruction (MAD), each state necessarily had to rely on deterrence, on trying

[4] More precisely, $\mu(r^*(R))$ is generically a singleton. To see that this is the case, suppose $j \neq k$ are both in $\mu(r^*(R))$. Since these sites minimize the defender's losses among th sites in $\text{br}_A(r^*(R))$, $\lambda_j v_j(r_j^*(R)) = \lambda_k v_k(r_k^*(R))$. Furthermore, $\gamma_j v_j(r_j^*(R)) = \gamma_k v_k(r_k^*(R))$ because j and k are both in $\text{br}_A(r^*(R))$. Hence, $\gamma_j/\gamma_k = \lambda_j/\lambda_k$ whenever j and k are in $\mu(r^*(R))$, and this condition is clearly nongeneric.

to convince the other state not to do what it was physically capable of doing.[5]

The first term on the right of Eq (2) formalizes the defensive effect of greater spending. It is the marginal change in the defender's losses resulting from hardening the sites and reducing their vulnerability *given that the attacker continues to exert the same level of effort*. Lower vulnerability means that the expected losses the attacker is physically able to impose for a given level of effort is less. Hence, this term measures the marginal effects of the attacker's being physically less able to strike the defender.

But an increase in spending also has a deterrent effect which the second term represents. Greater spending if allocated optimally reduces the attacker's return on effort. As a result, the defender chooses to exert less effort. The attacker is physically capable of exerting just as much effort as it did before but is now deterred from doing so. Formally, $dE^*(R)/dR < 0$, and the second term on the right of Eq (2) is the decrease in the defender's losses resulting purely from the lower level of effort the attacker dedicates to striking the defender. (As shown formally below, the assumption that $c_A''' \leq 0$ implies that the marginal effects of greater spending on E^* are decreasing: $d^2E^*(R)/dR^2 \geq 0$.)

Finally, spending more to protect from an attack means diverting resources from other social ends. The cost effect, $c_D'(R)$, measures this loss.

Subject two qualifications, the defender's optimal level of spending, R^*, solves the first-order condition $dL(R^*)/dR = 0$ by equating the margin gain from greater spending (the sum of the defensive and deterrent effects) with the marginal cost. In equilibrium, the defender spends R^*, allocates it by minmaxing the attacker with $r^*(R^*)$, and the attacker invests $E^*(R^*)$ in going after sites in $\mu(r^*(R))$.

The first qualification is that at least in principle, the defender might spend enough to completely deter the attacker (i.e., there is a corner solution at which $E^* = 0$). If at some level below R^* the marginal return on attacking drops below the marginal cost $c_A'(0)$, the

[5] Glenn Snyder (1961,14-16) was one of the earliest to explicate the distrinction between deterrence and defense although he called these "deterrence by punishment" and "deterrence by denial," respectively. Brodie (1959) discusses the effect of nuclear weapons on the ability of states to defend themselves. See also Jervis (1984, 1989), Powell (1990, 7-12), and Schelling (1966, 1-34).

attacker will not find it worthwhile to invest any effort in striking the attacker. Let $\overline{R}$ be the level of spending at which the attacker's marginal gain to investing anything in going after the defender equals its marginal cost, i.e., $\overline{R}$ satisfies $M_A(r^*(\overline{R})) = c_A'(0)$. The the defender's equilibrium level of spending is the lesser of R^* and $\overline{R}$.

The second qualification is that $L(R)$ is not differentiable everywhere. At N values of R, the function $L(R)$ is kinked and may discontinuously jump down. As a result, the minimization of L is more complicated than solving the first-order condition and verifying that L is convex. Despite these complications, the defender's equilibrium level of spending is generically unique and bounded below by R^* (see Proposition 4).(see Proposition 4).

Facing a More Determined Attacker

What happens to the equilibrium levels of spending R^* and effort E^* and the allocations of them as the attacker's aggressiveness or determination to strike the defender rises? As did the previous section, this discusses the comparative statics with a minimum of technical detail. Proposition 5 below presents a more rigorous statement of the results.

One way to conceive of the attacker's determination is as the ratio of the marginal gain from exerting more effort to the marginal cost. The more determined the attacker, the higher this ratio and the more willing the attacker will be to exert that effort. To formalize this, recall that the attacker's payoff to exerting E and allocating it optimally is $M_A(r)E - c_A(E) = \gamma_j v_j(r_j)E - c_A(E)$ where j is any element in $\text{br}_A(r)$. Thus, the ratio of the marginal gain to the marginal cost is $\gamma_j v_j(r_j)/c_A'(E)$. Now suppose that the attacker's payoff to a successful strike on any site i is parametrized as $\delta\gamma_i$. Then the ratio of the marginal gain to the marginal cost is $\delta\gamma_j v_j(r_j)/c_A'(E)$. The larger δ, the larger the marginal gain to investing effort in attacking relative to the marginal cost at any allocation r and the more determined the attacker is.

The attacker's level of determination has no effect on the equilibrium allocations of effort and spending. The defender's optimal allocation of R minimizes the attacker's marginal return $M_A(r, \delta) \equiv \max\{\delta\gamma_1 v_1(r_1), ..., \delta\gamma_N v_N(r_N)\} = \delta M_A(r, 1)$. Thus the minmax allocation $r^*(R)$ which minimizes $M_A(r, 1)$ also minimizes $M_A(r, \delta)$. This implies that the defender's optimal allocation of R is independent of the attacker's level of determination. In symbols, $r^*(R,\delta) = r^*(R,1)$ for all $\delta > 0$.

The attacker's equilibrium allocation is also independent of its determination. If the defender allocates $r^*(R)$, the attacker will go after the site that minimizes the defender's losses, i.e., the generically unique site k such that $\lambda_k v_k(r_k^*(R)) = \min\{\lambda_j v_j(r_j^*(R)) : j \in \text{br}_A(r^*(R))\}$. The identity of this site depends on the allocation $r^*(R)$ and the defender's losses λ_j. Neither of these depend on the attacker's determination δ, and the allocations of effort across the attacker's best replies that minimize the defender's losses are independent of δ.

The attacker's level of determination only affects the equilibrium levels of spending and effort, and both of these increase as the attacker becomes more determined. Setting aside the two qualifications noted above, we focus on the case in which the equilibrium level of spending R^* equates the marginal gains from additional spending – the sum of the defense and deterrent effects – with the marginal cost (see Eq 2). If spending marginally more at R^* leads to lower losses as δ increases, then dL/dR at R^* will become negative and the defender will increase its spending in order to minimize its losses. That is, an increase in δ leads to higher equilibrium spending if $\partial^2 L(R^*)/\partial\delta\partial R < 0$.

Since $r^*(R^*)$ is independent of δ, differentiating Eq (2) with respect to δ yields

$$\frac{\partial^2 L(R^*)}{\partial\delta\partial R} = \lambda_k v_k'(r_k^*(R^*)) \frac{dr_k^*(R)}{dR} \frac{\partial E^*(R^*,\delta)}{\partial\delta} + \lambda_k v_k(r_k^*(R^*)) \frac{dr_k^*(R)}{dR} \frac{\partial^2 E^*(R^*,\delta)}{\partial\delta\partial R}. \quad (3)$$

The first term on the right is the effect of a higher δ on the defense effect; the second term is the effect of a higher δ on the deterrent effect. An increase in δ has no effect on the cost effect as this depends on the losses to the defender of diverting funds from other social priorities and not on the attacker's payoffs ($\partial^2 c_D(R)/\partial\delta\partial R = 0$).

To evaluate the sign of $\partial^2 L(R^*)/\partial\delta\partial R$, note that as the attacker becomes more determined, the marginal return to effort and the level of effort rise ($\partial E^*/\partial\delta > 0$). This implies that the first term on the right of Eq (3) is negative. Thus the greater the attacker's determination, the larger the marginal reduction in the defender's losses due to the defense effect.

As for the deterrent effect, the marginal decline in effort due to an increase in R, $\partial E^*/\partial R$, decreases as the attacker becomes more determined. That is, $\partial^2 E^*/\partial\delta\partial R < 0$. More formally, the equilibrium level of spending $E^*(R,\delta)$ satisfies the first-order condition $\delta\gamma_k v_k(r_k^*(R)) = c_A'(E^*(R,\delta))$ for any $k \in \mathrm{br}_A(r^*(R))$. Implicit differentiation then gives $\partial E^*/\partial\delta = \gamma_k v_k(r_k^*(R))/c_A''(E^*)$ (which shows $\partial E^*/\partial\delta > 0$) and

$$\frac{\partial^2 E^*(R,\delta)}{\partial\delta\partial R} = \left[\frac{\gamma_k v_k'(r_k^*(R))}{c_A''(E^*)}\frac{dr_k^*(R)}{dR} - \frac{\delta\gamma_k v_k(r_k^*(R))c_A'''(E^*)}{[c_A''(E^*)]^2}\frac{\partial E^*}{\partial R}\right].$$

But $dr_k^*/dR > 0$, $\partial E^*/\partial R < 0$, and $c_A''' \leq 0$, so $\partial^2 E^*/\partial\delta\partial R < 0$.

This implies the second term in the expression for $\partial^2 L(R^*)/\partial\delta\partial R$ is also negative. Hence the marginal reduction in the defender's losses due to the deterrence effect is also increasing in the attacker's determination. Taking both defense and deterrent effects into account leaves $\partial^2 L(R^*)/\partial\delta\partial R < 0$ and therefore that the defender's equilibrium level of spending increases as the attacker becomes more determined.

Two opposing influences determine the attacker's equilibrium level of effort. The greater the attacker's determination, the higher the marginal return to effort, $M_A(r^*(R),\delta)$. These higher returns incline the attacker to increase its effort. But as just shown, a higher level of determination leads the defender to spend more, and this reduces the marginal return on effort. This tends to lower attacker's effort. Proposition 5 shows that the former always dominates the latter and that the equilibrium level of effort increases as the attacker's determination δ increases.

When the Defender's Resources are Exogenous

This section analyzes an allocation game in which the defender's level of resources R is exogenously specified and the defender only has to decide how to distribute R across the N sites it is trying to protect. As will be seen in the next section, the allocation rule derived in this game carries over to the game in which the defender also has to choose R. In this sense, the allocation and level-of-spending problems are separable.

Let Γ denote the game described in the previous section in which the defender chooses how much to spend and take $\Gamma_x(R)$ to be the game in which the level of resources is exogenously set at $R \geq 0$. A strategy for the defender in $\Gamma_x(R)$ is an allocation $r = (r_1,...,r_N)$ such that $r_j \geq 0$ and $\sum_{j=1}^N r_j \leq R$. As before, let $\Delta(R)$ be the set of all such strategies. A strategy for the attacker is the same in $\Gamma_x(R)$ as it was in Γ, namely, a function $e(r) = (e_1(r), ..., e_N(r))$ where $e_j(r) \geq 0$ is the effort the attacker devotes to striking site j and $E(r) \equiv \sum_{j=1}^N e_j(r)$ is the total effort the attacker invests in striking the defender. The defender's losses and the attacker's gains are $L_x(r, e) = \sum_{j=1}^N \lambda_j v_j(r_j)e_j + c_D(R)$ and $G_x(r, e) = \sum_{j=1}^N \gamma_j v_j(r_j)e_j - c_A(E)$, respectively.

A very simple intuition leads to the optimal allocation when the level of resources is exogenous. Suppose that the defender has not allocated any resources ($r = 0$). The attacker's marginal return to investing effort in hitting j is $\gamma_j v_j(0)$. Suppose that attacking k offers the highest expected marginal return on any invested effort at $r = 0$, i.e., $\gamma_k v_k(0) = M_A(0)$.[6] The attacker will then go after this site and exert a level of effort $E^*(0)$ where $E^*(0)$ satisfies the first-order condition $\gamma_k v_k(0) = c'_A(E^*(0))$. This leaves the defender with a loss of $\lambda_k v_k(0)E^*(0) + c_D(R)$.

Clearly, then the defender should begin distributing R by allocating resources to k. This both reduces the vulnerability of site k (the defensive effect) and the attacker's level of effort (the deterrent effect).[7] The more the defender spends on k, the harder and less attractive that site becomes. Eventually the marginal return to investing effort in attacking k is no higher than what was initially (i.e., at allocation $r = 0$) the second most attractive target. Call this site k'. If the defender still has resources to spend on defense, it must now allocate them to both k and k' so as to keep the attacker's marginal returns to going after these sites equal.[8]

As the defender spends more on k and k', they become less attractive targets and eventually are no more attractive than what was initially the third most attractive target, say, k''. At this point the defender must allocate resources to these three sites so as to

[6] To ease the intuitive exposition, assume site k is the only best reply.

[7] Since the level of resources is fixed at R in $\Gamma_x(R)$, there is no cost effect.

[8] Suppose the marginal returns are unequal with $\gamma_{k'} v_{k'}(r_{k'}) > \gamma_k v_k(r_k) > \gamma_j v_j(r_j)$ for $r_{k'} > 0$, $r_k > 0$, and $r_j = 0$ for $j \neq k, k'$. Then the attacker will only invest in going after site k' and the defender could have reduced its expected loss by allocating less to k and more to k'.

keep the attacker's expected marginal returns to investing in striking these sites equal to each other. And, of course, the more the defender devotes to these three sites, the less attractive they become. Eventually they are no more attractive than what was originally the fourth most attractive site, and so on. The defender continues to allocate its resources in this way, having to spread them across more and more sites, until it allocates R. Allocating resources in this way minimizes the attacker's maximum expected return.

The remainder of this section proves that the minmax allocation $r^*(R)$ is unique and the defender plays $r^*(R)$ in every subgame perfect equilibrium of $\Gamma_x(R)$. We also determine the defender's generically unique equilibrium loss.

LEMMA 1: *The minmax allocation* $r^*(R) \in \arg\min\{M_A(r) : r \in \Delta(R)\}$ *is unique.*
Proof: See the appendix.

To see that the defender plays the unique minmax allocation $r^*(R)$ in every subgame perfect equilibrium of $\Gamma_x(R)$, observe that because the attacker must play optimally after any r, the attacker can only invest positive effort in going after sites offering the highest expected marginal return, i.e., the sites in $\mathrm{br}_A(r)$. It will be useful to establish three related facts which follow from the intuition that minmaxing the attacker entails allocating resources to the most attractive target and then spreading them out over more and more sites as they become best replies for the attacker. Allocating resources in this way means, first, that the defender only allocates resources to protecting sites the attacker might actually invest in striking, i.e., $\sum_{j \in \mathrm{br}_A(r^*)} r_j^* = R$.[9] Second, if r' differs from the minmax allocation, then the set of best replies to r' is contained in the set of best replies to the minmax allocation r^*. Consequently, the amount of resources spent on the best replies to r' is less than R. Formally,

LEMMA 2: *Let* r^* *be the unique minmax allocation and suppose* $r' \neq r^*$*. Then (i)* $r_j^* = 0$ *for all* $j \notin \mathrm{br}_A(r^*)$ *and* $\sum_{j \in \mathrm{br}_A(r^*)} r_j^* = R$*; (ii)* $\mathrm{br}_A(r') \subseteq \mathrm{br}_A(r^*)$*; and (iii)* $\sum_{j \in \mathrm{br}_A(r')} r'_j < R$*.*
Proof: See the appendix.

Turning to the equilibria of the game, a subgame perfect equilibrium of $\Gamma_x(R)$ is a

9 To ease the exposition, we drop the argument R and write for example r_j^* rather than $r_j^*(R)$ when we can do so unabiguously.

strategy profile $(\widehat{r}, \widehat{e}(r))$ satisfying two conditions: (i) $\widehat{r}$ minimizes the defender's loss given that the attacker plays according to $\widehat{e}(r)$, and (ii) playing according to $\widehat{e}(r)$ maximizes the attacker's expected payoff at every $r \in \Delta(R)$. Condition (ii) is equivalent to $\widehat{e}_j(r) > 0$ only if $j \in \mathrm{br}_A(r)$ and exerting a level of effort equal to $S^*(r)$ where recall $S^*(r) = \max\{0, {c'_A}^{-1}(M_A(r))\}$.

Recall further that $\overline{R}$ is the (possibly infinite) level of spending needed to fully deter the attacker. That is, $r^*(\overline{R})$ induces the the attacker to invest nothing in going after the attacker: $S^*(r^*(\overline{R})) = E^*(\overline{R}) = 0$ since $\overline{R}$ solves $M_A(r^*(\overline{R})) = c'_A(0)$. Clearly, the defender would never spend more than $\overline{R}$ in any equilibrium of Γ where it chooses its level of spending. As for the equilibria of $\Gamma_x(R)$:

Proposition 1: *Let $0 \leq R \leq \overline{R}$. Then subgame perfect equilibria of $\Gamma_x(R)$ exist, and the defender plays the unique minmax allocation r^* in all of them.*

There is nothing to show if $R = 0$ as the defender only has one strategy $(r = 0)$ which is vacuously the unique minmax allocation and to which the attacker has a best reply. Hence, a subgame perfect equilibrium exists.

When $R \in (0, \overline{R}]$, there are two parts to the proof of Proposition 1, and each is stated as a separate lemma. Lemma 3 demonstrates existence by constructing a subgame perfect equilibrium in which the defender plays r^*. Lemma 4 shows that the defender cannot play $r' \neq r^*$ in any subgame perfect equilibrium.

The key to these results is recognizing that if $\widehat{r}$ is an equilibrium allocation then even though the attacker is indifferent to how it distributes $S^*(\widehat{r})$ over the sites in $\mathrm{br}_A(r)$, the defender will generally not be indifferent when the game is nonzero-sum. That is, $\lambda_i v_i(\widehat{r}_i)$ may not equal $\lambda_j v_j(\widehat{r}_j)$ for $i, j \in \mathrm{br}_A(\widehat{r})$. Let k be the site in $\mathrm{br}_A(\widehat{r})$ the defender would prefer the attacker to go after, i.e., $\lambda_k v_k(\widehat{r}_k) = \mu(\widehat{r})$. (To ease the exposition assume for the moment this site is unique.) Then save for some nongeneric circumstances, the attacker must dedicate all of its effort to striking k in response to an equilibrium offer of $\widehat{r}$.

To see why, suppose the contrary, i.e., the attacker hits $j \neq k$ where $j, k \in \mathrm{br}_A(\widehat{r})$ and $\lambda_k v_k(\widehat{r}_k) < \lambda_j v_j(\widehat{r}_j)$. Then the defender could deviate from $\widehat{r}$ to r' by spending slightly less on site k. That is, define $r'_k = \widehat{r}_k - \varepsilon$ and $r'_j = \widehat{r}_j$ for all $j \neq k$. Shifting resources away from k makes it a more attractive target than any other site in $\text{br}_A(\widehat{r})$. Indeed, the attacker now strictly prefers going after k to striking any other site. This leaves the defender with a loss of $\lambda_k v_k(r'_k)S^*(r')$ rather than $\lambda_k v_k(\widehat{r}_k)S^*(\widehat{r})$. Taking ε small enough ensures that r' is a profitable deviation from $\widehat{r}$, and this contradiction implies that the attacker must only go after k.

To construct an equilibrium of $\Gamma_x(R)$, define the attacker's strategy $e^*(r)$ to be $e^*_k(r) = S^*(r)$ if $k = \min \mu(r)$ and $e^*_k(r) = 0$ otherwise. In words, the attacker invests the optimal level of effort in hitting a site offering the highest marginal return, i.e., a site in $\text{br}_A(r)$. If $\text{br}_A(r)$ contains two or more sites, the attacker invests in striking the site in $\text{br}_A(r)$ that minimizes the defender's expected loss, i.e., $j \in \mu(r)$). If two or more sites in $\text{br}_A(r)$ minimize the defender's expected loss, i.e., if $\mu(r)$ has two or more elements, the attacker goes after the site with the smallest index among the sites in $\mu(r)$.[10]

It is now straightforward to show that $(r^*, e^*(r))$ is a subgame perfect.

Lemma 3: *Let* $0 < R \leq \overline{R}$ *and take* $e^*(r)$ *to be* $e^*_k(r) = S^*(r)$ *if* $k = \min \mu(r)$ *and* $e^*_k(r) = 0$ *otherwise. Then* $(r^*(R), e^*(r))$ *is a subgame perfect equilibrium of* $\Gamma_x(R)$*.*

Proof: See the appendix.

To demonstrate that the defender can never play any $r' \neq r^*(R)$ in any subgame perfect equilibrium of $\Gamma_x(R)$, consider any strategy profile $(r', e'(r))$ in which $e'_j(r) > 0$ only if $j \in \text{br}\ _A(r)$ and the total level of effort is $S^*(r)$. It suffices to show that the defender can profitably deviate from r' to some other allocation thereby establishing that $(r', e'(r))$ is not an equilibrium.

The intuition guiding the proof begins with part (iii) of Lemma 2 which ensures that if $r' \neq r^*(R)$, then $\sum_{j \in \text{br}_A(r')} r'_j < R$. This implies that r' allocates resources to one or more sites outside $\text{br}_A(r')$. These resources can be redistributed across the sites in $\text{br}_A(r')$ to create a profitable deviation from r' and thereby establish that $(r', e'(r))$ is not an equilibrium allocation.

Lemma 4: *Suppose* $0 < R \leq \overline{R}$ *and* $r' \neq r^*$*. Take* $(r', e'(r))$ *to be any strategy profile such that* $e'(r)$ *maximizes the attacker's payoff for all* $r \in \Delta(R)$*, i.e.,* $e'_j(r) > 0$ *only if*

[10] More generally, e^* could be defined as any allocation of $S^*(r)$ over $\mu(r)$.

$j \in \mathrm{br}_A(r)$ *and* $\sum_{j=1}^{N} e'_j(r) = S^*(r)$. *Then the defender can profitably deviate from* r'.
Proof: See the appendix.

Lemmas 3 and 4 establish Proposition 1: the defender minmaxes the attacker by playing $r^*(R)$ in every subgame perfect equilibrium of $\Gamma_x(R)$. Recall further that the strategy profile that was shown to be an equilibrium in Lemma 3 had the attacker allocating $S^*(r)$ over the sites in $\mathrm{br}_A(r)$ that was most favorable to the defender. Proposition 2 shows that the attacker must allocate its effort in this way in response to the equilibrium offer $r^*(R)$ except possibly at as many as N values of R. That is, the attacker after observing $r^*(R)$ must allocate $S^*(r^*(R)) = E^*(R)$ among the sites in $\mu(r^*(R))$. An immediate consequence of this is that the defender's equilibrium loss is $\lambda_k v_k(r_k^*(R))E^*(R)$ where k is any element of $\mu(r^*(R))$.

The intuition behind this result is that if the attacker invested positive effort in striking a $j \notin \mu(r^*(R))$, then there would be a profitable deviation from $r^*(R)$. In particular, the defender could allocate $\varepsilon > 0$ less than $r_k^*(R)$ to some $k \in \mu(r^*(R))$ and thereby induce the defender to devote all of its effort to going after site k. Continuity ensures that this is a profitable deviation if ε is sufficiently small and hence a contradiction.

This intuition fails (and the claim does not hold) when $r_k^*(R) = 0$ for all $k \in \mu(r^*(R))$, because it is impossible to allocate slightly less to any site in $\mu(r^*(R))$. However this difficulty arises at most at N vallues of R. To see why, observe that $r_k^*(R) = 0$ for all $k \in \mu(r^*(R)) \subseteq \mathrm{br}_A(r^*(R))$ implies $\gamma_k v_k(0) = M_A(r^*(R))$. Since $M_A(r^*(R))$ is decreasing in R, $\gamma_k v_k(0) = M_A(r^*(R))$ can hold for at most one value of $R \in [0, \overline{R}]$ for each site. Given that there are N sites, this possibility is at issue at N values of R at most.

Proposition 2: *Suppose* $0 < R \leq \overline{R}$ *and let* $(r^*, e^*(r))$ *be any subgame perfect equilibrium of* $\Gamma_x(R)$. *If* $r_k^*(R) > 0$ *for at least one* $k \in \mu(r^*(R))$, *then:*
(i) $e_j^*(r^*(R)) > 0$ *only if* $j \in \mu(r^*(R))$;
(ii) *the defender's equilibrium losses are* $\lambda_k v_k(r_k^*(R))E^*(R) + c_D(R)$.

Proof: Suppose $r_k^* > 0$ for a $k \in \mu(r^*)$. Then the support of the attacker's effort must be contained in $\mu(r^*)$. Arguing by contradiction, let $(r^*, \widehat{e}(r))$ be a subgame perfect equilibrium of Γ_x for which $\widehat{e}_i(r_i^*) > 0$ for $i \notin \mu(r^*)$. Because the attacker only strikes at sites in $\mathrm{br}_A(r^*)$ in a subgame perfect equilibrium, $i \in \mathrm{br}_A(r^*)$. The definition of $\mu(r^*)$

then ensures $\lambda_i v_i(r_i^*) > \lambda_k v_k(r_k^*)$.

Since $r_k^* > 0$, we can construct the allocation r' such that $r'_k = r_k^* - \varepsilon > 0$ for an $\varepsilon > 0$ and $r'_j = r_j^*$ for all $j \neq k$. Clearly, $\gamma_k v_k(r'_k) > M_A(r^*) \geq \gamma_j v_j(r_j^*) = \gamma_j v_j(r'_j)$ for all $j \neq k$. Hence, $\text{br}_A(r') = \{k\}$. This implies that $\widehat{e}_k(r'_k) = S^*(r')$ and consequently that the defender's loss to r' is $\lambda_k v_k(r'_k) S^*(r')$.

The defender's payoff to offering r^* is $\sum_{j \in \text{br}_A(r^*)} \lambda_j v_j(r_j^*) \widehat{e}_j(r^*) \geq \lambda_i v_i(r_i^*) \widehat{e}_i(r^*) + [S^*(r^*) - \widehat{e}_i(r^*)] \lambda_k v_k(r_k^*) > \lambda_k v_k(r_k^*) S^*(r^*)$. The strict inequality along with continuity guarantee that we can choose ε so that $\sum_{j \in \text{br}_A(r^*)} \lambda_j v_j(r_j^*) \widehat{e}_j(r^*) > \lambda_k v_k(r'_k) S^*(r')$. Hence, the defender strictly prefers deviating to r', contradicting the assumption that $(r^*, \widehat{e}(r))$. ■

In sum, the analysis of $\Gamma_x(R)$ establishes that if the defender's resources are exogenously fixed, the optimal allocation of them is to minmax the attacker. The next section shows that whatever level of resources the defender chooses in Γ, the defender allocates them by minmaxing the attacker and that the losses resulting from an attack are given by Proposition 2. The optimal R then minimizes these losses subject to the cost of $c_D(R)$ of spending R.

Deciding How Much to Spend on Defense

As the defender spends more on defense, the defense and deterrent effects reduce its losses. But the marginal benefits of spending more decrease for two reasons. First, the marginal effect of greater spending on any one site is weakly decreasing ($\lambda_j v'_j < 0$ and $\lambda_j v'' \geq 0$). Second, as the defender spends more on a given set of sites, those sites eventually become no more attractive targets than other sites. This forces the defender to spread its resources across more and more sites as its level of spending increases. As a result, the marginal gain from additional spending decreases as that spending must be spread out over more sites. Meanwhile, the marginal cost of spending more steadily increases ($c''_D(R) > 0$).

Subject to an important qualification, the generically unique subgame perfect equilibrium allocation in Γ equates the marginal benefits of spending more on defense, i.e., the sum of the defensive and deterrent effects, with the marginal cost of diverting these resources from other purposes. The qualification is that the allocation that equates marginal gains with marginal losses is always a lower bound on the optimal level of spending and often actually is the optimal level of spending. But in some circumstances the defender may want to spend more than this amount. To some extend, the defender can influence the site the attacker strikes. Spending more on one site and less on another makes the former a less attractive target. The defender will want to spend more than the allocation that equates marginal costs and benefits if by doing so the defender can induce the attacker to invest in hitting a site the defender values much less than the site the attacker would have hit had the defender equated marginal costs with marginal benefits.

The formalization of these results begins with two observations. First, Lemma 5 below shows that the defender never spends more than $\overline{R}$ on defense. This level of spending is enough to induce the attacker not to attack, so additional spending only brings added cost. Second, consider any allocation r with $R = \sum_{j=1}^{N} r_j$. If if r does not equal the minmax allocation $r^*(R)$, then r cannot be an equilibrium allocation of Γ. This follows from Lemma 4 which shows that the defender has a profitable deviation from r in $\Gamma_x(R)$ when $r \neq r^*(R)$. It is easy to see that this is also a profitable deviation in Γ, so r cannot be an equilibrium allocation in Γ either. These two observations imply that we can narrow the search for subgame perfect equilibria of Γ to values of $R \in [0, \overline{R}]$ for which there exists a strategy $e^*(r)$ such that the minmax allocation $r^*(R)$ and $e^*(r)$ are subgame perfect.

This search reduces to a simple optimization problem. Defining this problem precisely is cumbersome because the defender's losses are a kinked and possibly discontinuous function of R. But the intuition is straightforward. Proposition 2 shows that the defender's subgame perfect equilibrium payoff in $\Gamma_x(R)$ is generically equal to $\lambda_k v_k(r_k^*(R))E^*(R)$ where k is a site in $\text{br}_A(r^*(R))$ at which the defender's loss is lowest. This also holds in Γ. That is, the defender's loss from an attack is $\lambda_k v_k(r_k^*(R))E^*(R)$ if the defender plays $r^*(R)$ in Γ. An equilibrium allocation must be a best reply to the attacker's strategy and therefore must minimize these losses plus the cost of spending R on defense, i.e., an equilibrium allocation must minimize $\lambda_k v_k(r_k^*(R))E^*(R) + c_D(R)$. This minimization problem has at most finitely many solutions and, generically, a unique solution.

The first step in establishing these results formally is to show that the defender never spends more than $\overline{R}$ in any subgame perfect equilibrium of Γ.

LEMMA 5: *The defender never allocates $R > \overline{R}$ in any subgame perfect equilibrium of* Γ.

Proof: Since $R > \overline{R}$, there are allocations of R that fully deter the attacker. Thus, the defender's losses to any allocation of R are bounded below by $c_D(R)$. But deviating from any allocation of R to $r^*(\overline{R})$ also deters the attacker ($E^*(\overline{R}) = 0$) and leaves the defender with lower losses of $c_D(\overline{R})$. Hence, there is a profitable deviation and this contradiction establishes the claim. ■

Turning more directly to the minimization problem, let $k(R)$ be the site in $\text{br}_A(r^*(R))$ at which the defender's losses are minimized. That is, $k(R)$ satisfies $\lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) = \min_{j \in \text{br}_A(r^*(R))}\{\lambda_j v_j(r^*(R))\}$. (If two or more sites minimize the defender's losses, take $k(R)$ to be the site with the smallest index. In symbols, $k(R) \equiv \min \mu(r^*(R))$.) Now define the defender's loss function $L(R)$ as in Eq (1) to be:

DEFINITION 1 (THE DEFENDER'S LOSS FUNCTION): *Let* $k(R) = \min \mu(r^*(R))$, *then*

$$L(R) \equiv \begin{cases} \lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) E^*(R) + c_D(R) & \text{for } R \in [0, \overline{R}) \\ \\ c_D(R) & \text{for } R \geq \overline{R} \end{cases}.$$

Unfortunately, L' is not continuous and L may not be. As R increases, the defender has to defend more and more sites. As additional sites enter the set of defended sites, the site that minimizes the defender's loss, $k(R)$, may change and the defender's minimum losses $\lambda_{k(R)} v_{k(R)} E^*(R)$ may discontinuously jump down. If it does, L will be discontinuous. The fact that the defender must spread its resources across more sites as they enter the set of best replies also means that the marginal return to additional spending will discontinuously drop as additional sites become best replies. Hence L' will jump up, and L will be kinked even if it is continuous. However, these discontinuous jumps can only occur when sites enter the set of attacker's best replies $\text{br}_A(r^*(R))$ This implies that the number of discontinuities is bounded above by the number of sites. Everywhere else L is well-behaved: L' and L'' exist and $L'' > 0$.

To define the points of discontinuity, let Z_j be the value of R at which site j begins to offer the highest marginal return to the attacker, i.e., the value of R at which j joins $\text{br}_A(r^*(R))$. More precisely, Z_j is defined by $\gamma_j v_j(0) = M_A(r^*(Z_j))$. If j never becomes a best reply for $R \leq \overline{R}$, i.e., if $\gamma_j v_j(0) < M_A(r^*(\overline{R}))$, define Z_j to be anything larger than $\overline{R}$. For spending levels less than Z_j, j is not a best reply to $r^*(R)$ as $\gamma_j v_j(0) < M_A(r^*(R))$ for $R < Z_j$. Site j is a best reply for all $R \geq Z_j$.

It will be convenient to order the Z_j's. Take B_1 to be the smallest value of $R \leq \overline{R}$ at which a site enters $\text{br}_A(r^*(R))$, i.e., the smallest Z_j; let B_2 the next smallest value at which a site enters, and so on with B_{H-1} being the largest value of $R \leq \overline{R}$ at which a site enters. $B_1 = 0$ since $\gamma_j v_j(0) = M_A(0)$ for at least one j. Defining $B_H = \overline{R}$ then leaves $0 = B_1 < B_2 < \cdots < B_H = \overline{R}$. Take $\mathcal{B}$ to be $\overline{R}$ along with the set of ordered values of R at which sites enter $\text{br}_A(r^*(R))$: $\mathcal{B} \equiv \{B_1, ..., B_H\}$.

By construction the set of sites yielding the highest marginal return, $\text{br}_A(r^*(R))$, is constant for all $R \in (B_j, B_{j+1})$. This implies that $k(R)$ is constant and L behaves nicely for $R \notin \mathcal{B}$. To see that this is the case, differentiate L to obtain:

$$\frac{dL}{dR} = \lambda_k v_k'(r_k^*(R))\frac{dr_k^*(R)}{dR}E^*(R) + \lambda_k v_k(r_k^*(R))\frac{dE^*(R)}{dR} + c_D'(R)$$

To solve for $dr_{k(R)}^*(R)/dR$, observe that $\gamma_{k(R)} v_{k(R)}(r_{k(R)}^*(R)) = \gamma_j v_j(r_j^*(R))$ for all $j \in \text{br}_A(r^*(R))$ because the attacker is indifferent to striking any of the sites in $\text{br}_A(r^*(R))$. Differentiating implicitly now gives $dr_j^*(R)/dR = [\gamma_{k(R)} v_{k(R)}'(r_{k(R)}^*(R))]/[\gamma_j v_j'(r_j^*(R))] \cdot dr_{k(R)}^*(R)/dR$ for all $j \in \text{br}_A(r^*(R))$. Since the defender spends all of its resources on sites in $\text{br}_A(r^*(R))$, $\sum_{j \in \text{br}_A(r^*(R))} r_j^*(R) = R$ and consequently $\sum_{j \in \text{br}_A(r^*(R))} dr_j^*(R)/dR = 1$. Substituting the expressions for $dr_j^*(R)/dR$ in the previous equality yields

$$\frac{dr_{k(R)}^*(R)}{dR} = \frac{1}{\gamma_{k(R)} v_{k(R)}'(r_{k(R)}^*(R))}\left(\sum_{j \in \text{br}_A(r^*(R))} \frac{1}{\gamma_j v_j'(r_j^*(R))}\right)^{-1}. \tag{4}$$

This expression formalizes the fact that having to spread a marginal increase in spending across more sites reduces the marginal amount spend on each site. That is, the larger the number of sites in $\text{br}_A(r^*(R))$, the smaller the $dr_j^*(R)/dR$. And, the smaller $dr_j^*(R)/dR$, the smaller the marginal effect of greater spending on the defender's losses.

As for $dE^*(R)/dR$, site $k(R)$ offers the attacker its maximum return since $k(R) \in \mu(r^*(R)) \subseteq \text{br}_A(r^*(R))$. Hence, $c'_A(E^*(R)) = \gamma_{k(R)} v_{k(R)}(r^*_{k(R)}(R))$ for $R < \overline{R}$. Implicit differentiation then gives $dE^*(R)/dR = \gamma_{k(R)} v'_{k(R)}(r^*_{k(R)})/c''_A(E^*(R)) dr^*_{k(R)}(R)/dR$. Substituting for $dr^*_{k(R)}(R))/dR$ yields:

$$\begin{aligned} \frac{dE^*(R)}{dR} &= \frac{\gamma_{k(R)} v'_{k(R)}(r^*_{k(R)}(R))}{c''_A(E^*(R))} \left(\sum_{j \in \text{br}_A(r^*(R))} \frac{\gamma_{k(R)} v'_{k(R)}(r^*_{k(R)}(R))}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} \\ &= \frac{1}{c''_A(E^*(R))} \left(\sum_{j \in \text{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1}. \qquad (5) \end{aligned}$$

Using these results to simplify the expression for dL/dR for $R \notin \mathcal{B}$ gives:

$$\begin{aligned} L'(R) = \lambda_{k(R)} v'_{k(R)}(r^*_{k(R)}(R)) &\left(\sum_{j \in \text{br}_A(r^*(R))} \frac{\gamma_{k(R)} v'_{k(R)}(r^*_{k(R)}(R))}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} E^*(R) \\ &+ \frac{\lambda_{k(R)} v_{k(R)}(r^*_{k(R)})}{c''_A(E^*(R))} \left(\sum_{j \in \text{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} + c'_D(R) \\ = \frac{\lambda_{k(R)}}{\gamma_{k(R)}} &\left(\sum_{j \in \text{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} E^*(R) \\ &+ \frac{\lambda_{k(R)} v_{k(R)}(r^*_{k(R)})}{c''_A(E^*(R))} \left(\sum_{j \in \text{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} + c'_D(R). \end{aligned}$$

The first two terms on the right of the last equation are explicit representations of the defense and deterrent effects respectively. These expressions too show that the larger the number of sites at risk, i.e., the larger the number of sites in $\text{br}_A(r^*(R))$, the smaller these effects on the defender's marginal losses.

Simplifying still further yields

$$L'(R) = \lambda_{k(R)} \left[\frac{E^*(R)}{\gamma_{k(R)}} + \frac{v_{k(R)}(r^*_{k(R)})}{c''_A(E^*(R))} \right] \left(\sum_{j \in \mathrm{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} + c'_D(R)$$

Differentiating a second time shows that $L'' > 0$ (see the proof of Lemma 6(i)).

Proposition 3 below shows that $r^*(R)$ is an equilibrium allocation of Γ if and only if R minimizes the defender's losses L. The discussion above demonstrates that L is continuously differentiable and convex everywhere except at points in $\mathcal{B}$. Were L continuously differentiable and convex everywhere, a unique R^* would minimize L. Save for the possibility of a corner solution at $R = \overline{R}$, the optimal allocation R^* would equate the marginal benefits (lower losses) with the marginal cost and satisfy the first-order condition $L'(R^*) = 0$ or $-\lambda_{k(R)} v'_{k(R)}(r^*_{k(R)}(R^*))\ dr^*_{k(R)}(R^*)/dR = c'_D(R^*)$.[11]

That L can discontinuously jump down qualifies this. Suppose that the attacker would strike $k^* = k(R^*)$ were the defender to spend R^* where, allowing for the potential discontinuities, R^* is uniquely defined by $\lim_{R \uparrow R^*} L'(R) \leq 0 \leq \lim_{R \downarrow R^*} L'(R)$.[12] At least in principle, the defender may be able to lower its losses still further by spending more than R^* if doing so induces the attacker to strike a site the defender values much less than k^*. That is, it may pay the defender to bear the cost of increasing its spending from R^* to $\widehat{R}$ if the attacker then hits $\widehat{k}$ where the drop in losses $\lambda_{k^*} v_{k^*}(r^*_{k^*}(R^*))E^*(R^*) - \lambda_{\widehat{k}} v_{\widehat{k}}(0) E^*(\widehat{R})$ more than offsets the increase in cost $c_D(\widehat{R}) - c_D(R^*)$.[13] The remainder of this section establishes these results formally.

Lemma 6 summarizes the properties of L.

11 The assumption that $c'_D(0) = 0$ precludes a corner solution at $R = 0$.

12 If the losses are smooth at R^*, then $L'(R) = 0$ and R^* equates the marginal gains from greater spending with the marginal losses. If L' is discontinuous at R^*, the marginal gains are larger than the marginal costs for $R < R^*$ and less than the marginal costs for $R > R^*$.

13 Site $\widehat{k}$ enters the attacker's set of best replies at $\widehat{R}$, so $r^*_{\widehat{k}}(\widehat{R}) = 0$.

Lemma 6: *Let*

$$\sigma(R) \equiv \lambda_{k(R)} \left[\frac{E^*(R)}{\gamma_{k(R)}} + \frac{v_{k(R)}(r^*_{k(R)})}{c''_A(E^*(R))} \right] \left(\sum_{j \in \mathrm{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} + c'_D(R)$$

for $R \in [0, \overline{R})$ *and* $\sigma(R) \equiv c'_D(R)$ *for* $R \geq \overline{R}$. *Then:*

(i) $L'(R) = \sigma(R)$ *and* $L''(R) > 0$ *for* $R \notin \mathcal{B}$;

(ii) $L(R)$ *and* $\sigma(R)$ *are continuous from the right for all* R;

(iii) *If* $\widehat{R} = B_j > 0$, *then* $\lim_{R \uparrow \widehat{R}} L(R) \geq L(\widehat{R})$;

(iv) *If* $\widehat{R} = B_j > 0$, *then* $\lim_{R \uparrow \widehat{R}} L'(R) < \lim_{R \downarrow \widehat{R}} L'(R) = \sigma(\widehat{R})$.

Proof: See the appendix.

Lemma 6 ensures that L has at most two local minima over $[B_j, B_{j+1}]$. Lemma 6(ii) guarantees that L is continuous from the right and consequently $L(B_j) = \lim_{R \downarrow B_j} L(R)$. This along with the strict convexity of L over (B_j, B_{j+1}) from Lemma 6(i) implies that: (i) L has a unique minimum at B_j for $R \in [B_j, B_{j+1})$, or (ii) L has a unique minimum over $[B_j, B_{j+1})$ for some $R_j \in [B_j, B_{j+1})$ where $L'(R_j) = 0$, or (iii) L is decreasing over $[B_j, B_{j+1})$ and has no minimum. Hence, L has at most one minimum over $[B_j, B_{j+1})$ and at most two minima over $[B_j, B_{j+1}]$. Taking the minimum of these ensures that L has a global minimum and that the number of possible global minima is bounded above by $\|\mathcal{B}\| + 1 = H + 1$.

The minimum of L is also generically unique in the following sense. Suppose the defender's cost to R comes from a family of functions with slightly different marginal costs. Formally, suppose the defender's cost of spending R is $\phi c_D(R)$ where $\phi \in [1 - \varepsilon, 1 + \varepsilon]$. Then the defender's loss is $L(R, \phi) \equiv \lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) E^*(R) + \phi c_D(R)$ for $R \in [0, \overline{R})$ and $L(R, \phi) \equiv \phi c_D(R)$ for $R \geq \overline{R}$. These losses have a unique minimum at all but finitely many values of ϕ. Lemma 7 demonstrates these claims.

Lemma 7: *The defender's losses* $L(R)$ *have finitely many minima the number of which is bounded above by* $H + 1$. *The loss function* $L(R, \phi)$ *for* $\phi \in [1 - \varepsilon, 1 + \varepsilon]$ *has a unique minimum for all but finitely many values* ϕ.

Proof: See the appendix.

It now follows that R is an equilibrium allocation of Γ if and only if it minimizes $L(R)$.

Proposition 3: *Let* W *be the set of minima of* $L(R)$. *Then* $\widehat{r}$ *with* $\widehat{R} \equiv \sum_{j=1}^{N} \widehat{r}_j$ *is a* *subgame perfect equilibrium allocation of* Γ *if and only if* $\widehat{r} = r^*(\widehat{R})$ *and* $\widehat{R} \in W$*.*

Proof: See the appendix.

An immediate consequence of Proposition 3 is that if $\overline{R}$ minimizes L, then the defender's optimal allocation fully deters the attacker, $E^*(\overline{R}) = 0$.

The attacker's equilibrium allocation is generically unique as well. The defender's allocation $r^*(R)$ determines the attacker's level of effort to be $E^*(R)$. Proposition 2 ensures that that attacker only invests effort in striking sites in $\mu(r^*(R))$ except possibly at finitely many values of R, and it is easy to show that this result obtains at all values of R in Γ.[14] But the set $\mu(r^*(R))$ generically only has one element. Call it k. Then the attacker's generically unique equilibrium response to $r^*(R)$ will be to invest $E^*(R)$ in striking k.

To see that $\mu(r^*(R))$ is generically a singleton, suppose $j \neq k$ are both in $\mu(r^*)$. It follows that $\gamma_j v_j(r_j^*) = \gamma_k v_k(r_k^*)$ since $\mu(r^*) \subseteq \text{br}_A(r^*)$. That j and k are in $\mu(r^*)$ also means that they minimize the defender's losses over the sites in $\text{br}_A(r^*)$. Consequently, $\lambda_j v_j(r_j^*) = \lambda_k v_k(r_k^*)$. These equalities imply that $\lambda_j/\lambda_k = \gamma_j/\gamma_k$ which is clearly non-generic.

Finally, Proposition 4 shows that in the generically unique subgame perfect equilibrium allocation, the defender always spends at least as much as is needed to equate the marginal benefits of spending more on defense (and allocating it optimally) with the marginal costs of larger defense expenditures. That is, the defender always spends at least as much as R^* where R^* uniquely satisfies $\lim_{R \uparrow R^*} L'(R) \leq 0 \leq \lim_{R \downarrow R^*} L'(R)$.[15] In some circumstances, the defender may do better by spending more than this if doing so can induce the attacker to strike at a site that the defender values less than what the attacker would have hit at R^*. Formally,

Proposition 4: *There exists a unique allocation* R^* *such that* $\lim_{R \uparrow R^*} L'(R) \leq 0 \leq$

[14] Suppose the attacker invests positive effort in a $j \notin \mu(r^*(R_0))$ for some $R_0 \in \mathcal{B}$. Then the defender could profitably deviate by slightly increasing its expenditure and distributing this spending across the sites in $\mu(R_0)$. This would induce the attacker to strike only at the sites in $\mu(R_0)$ and would be a profitable deviation. This argument, of course, fails in the case of $\Gamma_x(R_0)$, because the defender cannot increase its spending.

[15] If L is differentiable at R^*, then R^* satisfies $L'(R^*) = 0$.

$\lim_{R \downarrow R^*} L'(R)$, *and the generically unique equilibrium allocation is bounded below by* R^*.
Proof: See the appendix.

In sum, the equilibrium path in Γ is generically unique. Balancing the gains from the defense and deterrent effects of greater spending against the cost of spending less on other social goals, the defender chooses the level of spending R and the corresponding minmax allocation $r^*(R)$ which minimizes its losses. In response, the attacker equates the marginal return to investing in hitting the most attractive sites against the costs of gathering the resources needed to mount an attack. This determines the attacker's optimal level of effort. The attacker directs this effort to the site which, among the most attractive sites to the attacker, minimizes the defender's loss.

Comparative Statics of the Attacker's Determination

To formalize the comparative statics on the attacker's determination, let $R^*(\delta)$ denote the defender's equilibrium level of spending and $r^*(R, \delta)$ be the minmax allocation of R given δ. Take $S^*(r, \delta)$, to be the attacker's optimal level of effort given r and δ, and define $E^*(R, \delta) \equiv S^*(r^*(R, \delta), \delta)$. Take $\text{br}_A(r, \delta)$ to be the attacker's best replies and $\mu(r, \delta)$ to be the set of sites in $\text{br}_A(r, \delta)$ at which the defender's losses are lowest. Finally, define $k(R, \delta) \equiv \min \mu(R, \delta)$.

Proposition 5 shows that if the defender's equilibrium level of spending is an interior solution, then an increase in the attacker's determination (i) has no effect on the defender's allocation, $\partial r^*(R, \delta)/\partial \delta = 0$; (ii) no effect on the attacker's allocation, $\text{br}_A(r, \delta) = \text{br}_A(r, 1)$ and $\mu(r, \delta) = \mu(r, 1)$ for all $\delta > 0$; (iii) leads to more spending, $\partial R^*(\delta)/\partial \delta > 0$; and (iv) more effort $\partial E^*(R^*, \delta)/\partial \delta > 0$.

Proposition 5: *If* $R^*(\delta_0) \notin \mathcal{B}$, *then (i)* $\partial r^*(R, \delta)/\partial \delta = 0$*; (ii)* $\text{br}_A(r, \delta) = \text{br}_A(r, 1)$ *and* $\mu(r, \delta) = \mu(r, 1)$ *for all* $\delta > 0$*; (iii)* $R^*(\delta)/\partial \delta > 0$*; and (iv)* $\partial E^*(R^*(\delta), \delta)/\partial \delta > 0$.
Proof: See the appendix.

Conclusion

Strategic attackers poses two related resource questions to the defender. How much should the defender spend and what should the defender spend it on? These questions turn out to be separable and have relatively straightforward answers. However much the defender decides to spend on defense, the optimal allocation of those resources minmaxes the attacker. Because minmaxing the attacker's marginal returns equalizes those returns across the most attractive sites, the attacker will be indifferent to how it allocates its optimal level of effort among those sites. Nevertheless, equilibrium considerations imply that the attacker allocates its effort among its best replies in the way that imposes the lowest expected loss on the defender. Given the attacker allocates its effort in this way, the defender's optimal level of spending minimizes the sum of the expected losses from an attack and the cost of spending on defense rather than some other social end.

Subject to some qualifications, the defender's optimal level of spending equates the marginal gain from spending with the marginal cost diverting those resources from other social goals. The analysis also shows that the marginal gain from additional spending can be decomposed into the sum of a defense effect and a deterrent effect. The former is the *ceteris paribus* gain the defender derives from the fact that the sites are physically more difficult to destroy and therefore that an attack is less likely to succeed. The deterrent effect is the *ceteris paribus* benefit resulting from the attacker's choosing to devote less effort to attacking the defender.

Finally, both the equilibrium levels of spending and effort are increasing in the attacker's determination. But the attacker's determination has no effect on the defender' allocation or on the attacker's.

Appendix

Proof of Lemma 1: There is nothing to show if $R = 0$. The only feasible allocation is $r = 0$ which because it is the only feasible allocation is also the unique minmax allocation.

Now suppose $0 < R$ and note that at least one minmax allocation is sure to exist because the $v_j(r_j)$ 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 $r^* \neq r'$ both minmax the attacker. Then $M_A(r') = M_A(r^*) > 0$ where the inequality follows from the fact that resources are scarce.

That $r^* \neq r'$ implies $r_j^* \neq r_j'$ for at least one j. Without loss of generality suppose $r_j^* < r_j'$. If $v_j(r_j^*) > 0$, then v_j is imperfectly defended and greater than $v_j(r_j')$. This leaves $M_A(r') = M_A(r^*) \geq \gamma_j v_j(r_j^*) > \gamma_j v_j(r_j')$. Continuity now ensures that there is an $\varepsilon > 0$ such that $M_A(r') > \gamma_j v_j(r_j' - \varepsilon)$. This ε of resources can be distributed across the sites $k \neq j$ to form the allocation $\widehat{r}$ where $\widehat{r}_j = r_j' - \varepsilon$ and $\widehat{r}_k = r_k' + \varepsilon/(N-1)$ for all $k \neq j$. But this yields the contradiction that $M_A(\widehat{r}) < M_A(r')$.

If $v_j(r_j^*) = 0$, then there exists an $\varepsilon > 0$ such that $r_j^* > r_j' - \varepsilon$ which implies $v_j(r_j^*) = v_j(r_j' - \varepsilon) = 0$. Once again ε resources can be distributed across the sites $k \neq j$ to form the allocation $\widehat{r}$ where $\widehat{r}_j = r_j' - \varepsilon$ and $\widehat{r}_k = r_k' + \varepsilon/(N-1)$ for all $k \neq j$. This too yields the contradiction that $M_A(\widehat{r}) < M_A(r')$. ■

Proof of Lemma 2: To see that (i) holds, suppose that $r_j^* > 0$ for some $j \notin \text{br}_A(r^*)$. Now redistribute $\varepsilon > 0$ from r_j^* across all of the sites in $\text{br}_A(r^*)$. That is, define $\widehat{r}$ as $\widehat{r}_k = r_k^* + \varepsilon/\|\text{br}_A(r^*)\|$ for all $k \in \text{br}_A(r^*)$, $\widehat{r}_k = r_k^*$ for all $k \notin \text{br}_A(r^*)$ and $k \neq j$, and $\widehat{r}_j = r_j^* - \varepsilon$. Since $M_A(r^*) = \gamma_k v_k(r_k^*) > \gamma_j v_j(r_j^*)$ for all $k \in \text{br}_A(r^*)$, continuity ensures we can take ε sufficiently small that $\gamma_k v_k(r_k^* + \varepsilon/\|\text{br}_A(r^*)\|) > \gamma_j v_j(r_j^* - \varepsilon)$ for all $k \in \text{br}_A(r^*)$. This leaves $M_A(\widehat{r}) < M_A(r^*)$ and contradicts the assumption that r^* minimizes $M_A(r)$.

Turning to (ii), let $j \notin \text{br}_A(r^*)$. Then (i) implies $r_j^* = 0$. Now observe that $\gamma_j v_j(r_j') \leq \gamma_j v_j(0) = \gamma_j v_j(r_j^*) < M_A(r^*) < M_A(r')$ where the last inequality holds because r^* is the unique minmax allocation. But $\gamma_j v_j(r_j') < M_A(r')$ means $j \notin \text{br}_A(r')$. Hence, $j \notin \text{br}_A(r^*) \Rightarrow j \notin \text{br}_A(r')$. The is equivalent to $j \in \text{br}_A(r') \Rightarrow j \in \text{br}_A(r^*)$ or, alternatively, $\text{br}_A(r') \subseteq \text{br}_A(r^*)$.

As for (iii), $\gamma_j v_j(r'_j) = M_A(r') > M_A(r^*)$ for all $j \in \text{br}_A(r')$ where the strict inequality follows from the fact that r^* is the unique minmax allocation. Further, $M_A(r^*) = \gamma_j v_j(r^*_j)$ for all $j \in \text{br}_A(r')$ as $\text{br}_A(r') \subseteq \text{br}_A(r^*)$. This gives $\gamma_j v_j(r'_j) > \gamma_j v_j(r^*_j)$ or $r'_j < r^*_j$ for all $j \in \text{br}_A(r')$. Summing over $\text{br}_A(r')$ yields $\sum_{j \in \text{br}_A(r')} r'_j < \sum_{j \in \text{br}_A(r')} r^*_j \leq R$. ■

Proof of Lemma 3: By construction, $e^*(r)$ maximizes the attacker's payoff following any r as the attacker invests the optimal level of effort $S^*(r)$ in attacking a site in $\text{br}_A(r)$. It therefore suffices to show that the defender cannot profitably deviate to any $r' \neq r^*$. To this end, let $k^* = \min \mu(r^*)$ and $k' = \min \mu(r')$. It follows that the defender's payoffs to playing r^* and r' against e^* are $\lambda_{k^*} v_{k^*}(r^*_{k^*}) S^*(r^*)$ and $\lambda_{k'} v_{k'}(r^*_{k'}) S^*(r')$ respectively.

Lemma 2 now gives $k' \in \text{br}_A(r') \subseteq \text{br}_A(r^*)$. That $k' \in \text{br}_A(r^*)$ implies $\gamma_{k'} v_{k'}(r^*_{k'}) = M_A(r^*)$. The fact that r^* is the unique minmax allocation also means $M_A(r^*) < M_A(r')$. Moreover, $M_A(r') = \gamma_{k'} v_{k'}(r'_{k'})$ because $k' \in \text{br}_A(r')$. Combining these relations gives $\gamma_{k'} v_{k'}(r^*_{k'}) = M_A(r^*) < M_A(r') = \gamma_{k'} v_{k'}(r'_{k'})$ which yields $v_{k'}(r^*_{k'}) < v_{k'}(r'_{k'})$. Hence, $\lambda_{k^*} v_{k^*}(r^*_{k^*}) \leq \lambda_{k'} v_{k'}(r^*_{k'}) < \lambda_{k'} v_{k'}(r'_{k'})$ where the weak inequality follows from the fact that k^* is a site in $\text{br}_A(r^*)$ at which the defender's loss is minimized.

The optimal level of effort $S^*(r) = \max\{0, {c'_A}^{-1}(M_A(r))\}$ is strictly increasing in the maximum marginal returns $M_A(r)$ as long as $M_A(r) > M_A(r^*(\overline{R}))$. So $M_A(r^*) < M_A(r')$ implies $S^*(r^*) < S^*(r')$. Hence, $\lambda_{k^*} v_{k^*}(r^*_{k^*}) S^*(r^*) < \lambda_{k'} v_{k'}(r'_{k'}) S^*(r')$ which shows that the defender's expected loss to playing r^* is strictly less than that of playing r'. ■

Proof of Lemma 4: Let $k = \min \mu(r')$. Then the defender's loss to playing r' is bounded below by $\lambda_k v_k(r'_k) S^*(r')$. Lemma 2(iii) ensures $\sum_{j \in \text{br}_A(r')} r'_j < R$ since $r' \neq r^*$. Hence there exists a site $n \notin \text{br}_A(r')$ such that $r'_n > 0$. Observe that the following holds for all $j \in \text{br}_A(r')$: $\gamma_j v_j(r'_j) = M_A(r') > \max\{\gamma_s v_s(r'_s) : s \notin \text{br}_A(r')\} \geq \gamma_n v_n(r'_n)$.

Now construct the allocation $\widehat{r}$ by shifting a small amount of resources from site n and distributing them across all of the sites in $\text{br}_A(r')$ so that attacker's marginal returns to striking every site in $\text{br}_A(r')$ decline slightly but that the return to striking k declines the least. Constructing the deviation in this way ensures that site k uniquely offers the highest expected return following $\widehat{r}$, i.e., $\text{br}_A(\widehat{r}) = \{k\}$.

Formally, the previous inequalities along with continuity guarantees that we can take

$\varepsilon > 0$ small enough that $\gamma_n v_n(r'_n - \varepsilon) < [M_A(r') + \max\{\gamma_s v_s(r'_s) : s \notin \text{br}_A(r')\}]/2 \equiv B$. Now let $\varepsilon_j \equiv \varepsilon / \|\text{br}_A(r')\|$ and $\widehat{r}_j \equiv r'_j + \varepsilon_j$ for all $j \neq k$, $j \in \text{br}_A(r')$. Let $\widehat{r}_j \equiv r'_j$ for all $j \neq n$, $j \notin \text{br}_A(r')$. Choose $\delta_k > 0$ so that $\gamma_k v_k(r'_k + \delta_k) > \max_{j \neq k, j \in \text{br}_A(r')}\{\gamma_j v_j(r'_j + \varepsilon_j), B\}$. The fact that $\gamma_k v_k(r'_k) = \gamma_j v_j(r'_j)$ for all $j \in \text{br}_A(r')$ ensures that such an ε_k exists. Define $\varepsilon_k \equiv \min\{\delta_k, \varepsilon / \|\text{br}_A(r')\|\}$ and let $\widehat{r}_k \equiv r'_k + \varepsilon_k$. Finally, set $\varepsilon_n \equiv \varepsilon_k + \sum_{j \neq k, j \in \text{br}_A(r')} \varepsilon_j \leq \varepsilon$ and define $\widehat{r}_n \equiv r'_n - \varepsilon_n$.

This construction gives $\text{br}_A(\widehat{r}) = \{k\}$. To see why, note that $\gamma_k v_k(\widehat{r}_k) = \gamma_k v_k(r'_k + \varepsilon_k) \geq \gamma_k v_k(r'_k + \delta_k) > \max_{j \neq k, j \in \text{br}_A(r')}\{\gamma_j v_j(r'_j + \varepsilon_j), B\} = \max_{j \neq k, j \in \text{br}_A(r')}\{\gamma_j v_j(\widehat{r}_j), B\}$. These inequalities and the definition of B imply that the attacker prefers going after k to any $j \neq n$. That $\gamma_k v_k(\widehat{r}_k) > B > \gamma_n v_n(r'_n - \varepsilon) \geq \gamma_n v_n(r'_n - \varepsilon_n)$ ensures that the attacker prefers investing in hitting k to n. Hence $\text{br}_A(\widehat{r}) = \{k\}$.

The attacker therefore invest $S^*(\widehat{r})$ in striking k after observing $\widehat{r}$, leaving the defender with a loss of $\lambda_k v_k(\widehat{r}_k) S^*(\widehat{r})$. By construction, $\widehat{r}_k > r'_k$, so $\lambda_k v_k(\widehat{r}_k) < \lambda_k v_k(r'_k)$. The definition of $\widehat{r}$ also implies $M_A(\widehat{r}) < M_A(r')$ which means $S^*(\widehat{r}) < S^*(r')$. Combining the previous inequalities gives $\lambda_k v_k(\widehat{r}_k) S^*(\widehat{r}) < \lambda_k v_k(r'_k) S^*(r')$. But $\lambda_k v_k(r'_k) S^*(r')$ is a lower bound on the defender's payoff to r'. Hence, $\widehat{r}$ is a profitable deviation from r'. ■

Proof of Lemma 6: The first claim follows trivially for $R > \overline{R}$, and discussion preceding the statement of the lemma establishes $L'(R) = \sigma(R)$ for $R \leq \overline{R}$.

As for L'', rewrite L' as

$$\frac{dL}{dR} = \lambda_{k(R)} v'_{k(R)}(r^*_{k(R)}(R)) \frac{dr^*_{k(R)}(R)}{dR} E^*(R) + \lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) \frac{dE^*(R)}{dR} + c'_D(R)$$

and differentiate again to obtain

$$\frac{d^2L}{dR^2} = \lambda_{k(R)} v''_{k(R)}(r^*_{k(R)}(R)) \left(\frac{dr^*_{k(R)}(R)}{dR} \right)^2 E^*(R)$$

$$+ 2\lambda_{k(R)} v'_{k(R)}(r^*_{k(R)}(R)) \frac{dr^*_{k(R)}(R)}{dR} \frac{dE^*(R)}{dR}$$

$$+ c''_D(R) + \lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) \frac{d^2E^*(R)}{dR^2}$$

The first three terms on the right are clearly positive as $v''_j \geq 0$, $dr^*_{k(R)}(R)/dR \geq 0$, $dE^*(R)/dR \leq 0$, and $c''_D > 0$. Equation (5) and the assumption that $c'''_A \leq 0$ ensures that $d^2E^*(R)/dR^2 \leq 0$ and $L'' > 0$.

Claim (ii) follows trivially for $R \geq \overline{R}$ as $L(R) = c_D(R)$ and $\sigma(R) = c'_D(R)$. For $R < \overline{R}$, there exists an $\varepsilon > 0$ and a site j such that $B_j \leq R < R + \varepsilon < B_{j+1}$. Over $[R, R + \varepsilon)$, $k(R)$ and $\text{br}_A(r^*(R))$ are constant, and continuity from the right follows immediately.

Turning to (iii), suppose the set of sites e enter $\text{br}_A(r^*(B_n))$ and define $\overline{\rho}$ and $\underline{\rho}$ to be the limit of $\lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R))$ as R goes to B_n from the left and right respectively. Clearly, $\overline{\rho} \geq \underline{\rho} > 0$ as the minimization defining $k(R)$ for $R \geq B_n$ is over at least as many sites as it is for $R < B_n$ since sites e enter $\text{br}_A(r^*(R))$ at B_n. Then, $\lim_{R \uparrow B_n} L(R) = \overline{\rho} E^*(B_n) + c_D(B_n) \geq \underline{\rho} E^*(B_n) + c_D(B_n) = \lim_{R \downarrow B_n} L(R) = L(B_n)$ where the last equality follows from the fact that L is continuous from the right.

As for (iv), rewrite $\sigma(R)$ as

$$\sigma(R) = \lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) \left[\frac{E^*(R)}{\gamma_{k(R)} v_{k(R)}(r^*_{k(R)}(R))} + \frac{1}{c''_A(E^*(R)} \right] \times \left(\sum_{j \in \text{br}_A(r^*(R))} \frac{1}{\gamma_j v'_j(r^*_j(R))} \right)^{-1} + c'_D(R).$$

Since $\gamma_j v_j(r^*_j(R)) = \gamma_i v_i(r^*_i(R))$ for all $i, j \in \text{br}_A(R)$, the function $\gamma_{k(R)} v_{k(R)}(r^*_{k(R)}(R))$ is continuous in R even though $k(R)$, the site minimizing the defender's losses, may vary with R. This along with the continuity of v_j, v_j', and c_D' gives:

$$\begin{aligned}
\lim_{R\uparrow B_n} L'(R) &= \overline{\rho}\left[\frac{E^*(B_n)}{\gamma_{k(B_n)}v_{k(B_n)}(r^*_{k(B_n)}(B_n))} + \frac{1}{c_A''(E^*(B_n))}\right] \\
&\quad \times \left(\sum_{j\in \mathrm{br}_A(r^*(B_n))-e} \frac{1}{\gamma_j v_j'(r_j^*(B_n))}\right)^{-1} + c_D'(B_n) \\
&\leq \underline{\rho}\left[\frac{E^*(B_n)}{\gamma_{k(B_n)}v_{k(B_n)}(r^*_{k(B_n)}(B_n))} + \frac{1}{c_A''(E^*(B_n))}\right] \\
&\quad \times \left(\sum_{j\in \mathrm{br}_A(r^*(B_n))-e} \frac{1}{\gamma_j v_j'(r_j^*(B_n))}\right)^{-1} + c_D'(B_n) \\
&< \underline{\rho}\left[\frac{E^*(B_n)}{\gamma_{k(B_n)}v_{k(B_n)}(r^*_{k(B_n)}(B_n))} + \frac{1}{c_A''(E^*(B_n))}\right] \\
&\quad \times \left(\sum_{j\in \mathrm{br}_A(r^*(B_n))} \frac{1}{\gamma_j v_j'(r_j^*(B_n))}\right)^{-1} + c_D'(B_n)
\end{aligned}$$

where the strict inequality results from summing over more sites than in the previous line. Note further that the expression on the right of the last line equals $\lim_{R\downarrow B_n} L'(R)$. Hence, $\lim_{R\uparrow B_n} L'(R) < \lim_{R\downarrow B_n} L'(R)$. Since L' is continuous from the right, $\lim_{R\downarrow B_n} L'(R) = \sigma(B_n)$. ■

Proof of Lemma 7: Consider the interval $R \in [B_j, B_{j+1}]$ and let m_j be the set of allocations that minimize $L(R)$ over this interval. Lemma 6 ensures $L''(R) > 0$ for $R \in (B_j, B_{j+1})$, $L(B_j) = \lim_{R\downarrow B_j} L(R)$, and $\lim_{R\uparrow B_{j+1}} L(R) \geq L(B_{j+1})$. All of this implies that if $\lim_{R\downarrow B_j} L'(R) \geq 0$, then $m_j \subset \{B_j, B_{j+1}\}$. If $\lim_{R\downarrow B_j} L'(R) < 0$, $m_j \subset \{R_j^*, B_{j+1}\}$ where R_j^* satisfies $L'(R_j^*) = 0$ if such an R_j^* exists in $(B_j, B_{j+1}]$.

Observe further that $L'(R) = 0$ for at most one allocation. To see why, note that L' is increasing in R over every interval (B_j, B_{j+1}) and L' only jumps up at B_{j+1}. This implies $\lim_{R\downarrow B_j} L'(R) < L'(\widetilde{R}) < \lim_{R\uparrow B_{j+1}} L'(R) < \lim_{R\downarrow B_{j+1}} L'(R)$ for every interval $[B_j, B_{j+1}]$ and every $\widetilde{R} \in (B_j, B_{j+1})$. Hence, $L'(R)$ is strictly increasing at allocations at which it is defined: $L'(R) < L'(\widehat{R})$ whenever $R < \widehat{R}$ and $L(R)$ and $L'(\widehat{R})$ are defined. Consequently, $L'(R) = 0$ for at most one allocation.

It follows that the global minima of $L(R)$ must be a subset of $\bigcup_{j=1}^{H-1} m_j = \{B_1, ..., B_H, R^*\}$ where $L'(R^*) = 0$ if such a point exists. The number of global minima is therefore bounded above by $H + 1$.

Now let $\varepsilon > 0$ and $L(R, \phi) = \lambda_{k(R)} v_{k(R)}(r^*_{k(R)}(R)) + \phi c_D(R)$ for $\phi \in [1 - \varepsilon, 1 + \varepsilon] \equiv \Phi$. To see that $L(R, \phi)$ has a unique minimum for all but at most finitely many values of ϕ, let $R^*(\phi)$ satisfy $\partial L(R^*(\phi), \phi)/\partial R = 0$ if such an $R^*(\phi)$ exists. Observe further that the allocations at which sites become best replies for the attacker are independent of ϕ, because ϕ only affects the defender's payoffs. The parameter ϕ has no affect on the attacker's payoffs and consequently no affect on the attacker's best replies. Since $B_j \in$B are independent of ϕ, the argument above ensures that the minimum of $L(R, \phi)$ occur at a subset of $\{B_1, ..., B_H, R^*(\phi)\}$.

Now observe that $L(R, \phi) = L(R, 1) + (\phi - 1)c_D(R)$ and let $B_j, B_k \in$B and $j \neq k$. Then $L(B_j, \phi) - L(B_k, \phi) = L(B_j, 1) - L(B_k, 1) + (\phi - 1)[c_D(B_j) - c_D(B_k)]$. This difference must be zero if both B_j and B_k minimize $L(R, \phi)$ and this clearly can happen at only one value of ϕ, say θ_{jk}. Let Θ be the set of all such θ_{jk} for B_j and B_k with $j \neq k$. Then taking $\phi \in \Phi \backslash \Theta$ ensures that $L(R, \phi)$ has at most two minima and that if it does have two, one of them must be $R^*(\phi) \notin \mathcal{B}$ where $\partial L(R^*(\phi), \phi)/\partial R = 0$.

To show that this can occur at only finitely many values of ϕ and therefore that $L(R, \phi)$ has a generically unique minimum, it suffices to demonstrate that $L(R^*(\phi), \phi) = L(B_j, \phi)$ at only finitely many values of ϕ for every $B_j \in$B. To this end, let (B_n, B_{n+1}) be any interval for which there exists a a $\phi \in \Phi$ and an $R^*(\phi)$ such that $R^*(\phi) \in (B_n, B_{n+1})$ and $\partial L(R^*(\phi), \phi)/\partial R = 0$. Take F to be the family of all such intervals, and consider the difference $\eta_{nj}(\phi) \equiv L(R^*(\phi), \phi) - L(B_j, \phi)$ for any $B_j \in$B and $R^*(\phi) \in (B_n, B_{n+1})$.

The set $\text{br}_A(r^*(R))$ remains constant over (B_n, B_{n+1}) since $\text{br}_A(r^*(R)) = \text{br}_A(B_n)$. As a result, η_{nj} is continuously differentiable and

$$\frac{d\eta_{nj}}{d\phi} = \frac{\partial L(R^*(\phi), \phi)}{\partial R} \frac{dR^*(\phi)}{d\phi} + \frac{\partial L(R^*(\phi), \phi)}{\partial \phi} - c'_D(B_j).$$

Exploiting the fact that $\partial L(R^*(\phi),\phi)/\partial R = 0$ leaves $d\eta_{nj}/d\phi = c'_D(R^*(\phi)) - c'_D(B_j)$.

Now consider any $B_j \leq B_n$. If follows that $\eta_{nj}(\phi) = L(R^*(\phi),\phi) - L(B_j,\phi)$ is strictly increasing since $R^*(\phi) \in (B_n, B_{n+1})$ implies that $\eta'_{nj} > 0$. Thus $\eta_{nj}(\phi) = L(R^*(\phi),\phi) - L(B_j,\phi) = 0$ at most at one value of ϕ. Call this value $\widehat{\theta}_{nj}$ if it exists. For any $B_j \geq B_{n+1}$, $\eta'_{nj} < 0$, and, again, $\eta_{nj}(\phi) = L(R^*(\phi),\phi) - L(B_j,\phi) = 0$ at most at one value of ϕ, say $\widehat{\theta}_{nj}$. Now define $\widehat{\Theta} = \{\widehat{\theta}_{nj} : (B_n, B_{n+1}) \in F, B_j \in \text{B}\}$.

It follows that $L(R^*(\phi),\phi)$ has a unique minimum if $\phi \in \Phi \backslash (\Theta \cup \widehat{\Theta})$. To establish this, suppose the contrary, that $L(R^*(\phi),\phi)$ has two or more minima. Since $\phi \notin \Theta$, $L(R^*(\phi),\phi)$ can have at most two minima and one must be at $R^*(\phi')$ for some $\phi' \in \Phi \backslash (\Theta \cup \widehat{\Theta})$. That is, there must be an n and j such that $R^*(\phi') \in (B_n, B_{n+1})$ and $L(R^*(\phi'),\phi') = L(B_j,\phi')$. But this cannot be as $L(R^*(\phi),\phi) \neq L(B_j,\phi)$ for all $\phi \in \Phi \backslash \widehat{\Theta}$. ■

Proof of Proposition 3: Suppose $\widehat{r} = r^*(\widehat{R})$ and $\widehat{R} \in W$. To establish that $\widehat{r}$ is an equilibrium allocation in Γ consider the profile $(r^*(\widehat{R}), e^*(r))$ where e^* is defined as it was in Lemma 3, i.e., $e^*_k(r) = S^*(r)$ if $k = \min \mu(r)$ and $e^*_k(r) = 0$ otherwise. It now suffices to show that $(r^*(\widehat{R}), e^*(r))$ is subgame perfect in Γ.

Because $e^*(r)$ is a best reply to any r, the only way that $(r^*(\widehat{R}), e^*(r))$ can fail to be subgame perfect is if there is a profitable deviation from $r^*(\widehat{R})$. Arguing by contradiction, assume such a deviation, say r', exists where $R' = \sum_{j=1}^N r'_j \leq \overline{R}$. (If $R' > \overline{R}$, it is straight forward to show that $r^*(\overline{R})$ would also be profitable, and we use this as the profitable deviation in the argument below.)

The definition of e^* implies that the defender's loss to $r^*(\widehat{R})$ is $\lambda_k v_k(r^*_k(\widehat{R}))E^*(\widehat{R}) + c_D(\widehat{R}) = L(\widehat{R})$. As for the losses to r', Lemma 3 ensures that $(r^*(R'), e^*(r))$ is subgame perfect in $\Gamma_x(R')$. Hence, the defender's loss to r' in $\Gamma_x(R')$ is bounded below by its loss to $r^*(R')$ in $\Gamma_x(R')$. Moreover, the defender's losses to playing r' and $r^*(R')$ in $\Gamma_x(R')$ are the same as its losses to playing r' and $r^*(R')$ in Γ, the latter of which is $L(R')$. This along with the fact that r' is a profitable deviation in Γ implies that $r^*(R')$ is also a profitable deviation in Γ. Hence, $L(\widehat{R}) > L(R')$. But this contradicts the assumption that $\widehat{R}$ minimizes L (i.e., $\widehat{R} \in W$), and this contradiction establishes the claim.

Now assume that $(\widehat{r}, \widehat{e}(r))$ is a subgame perfect equilibrium of Γ. It follows that

$\widehat{r} = r^*(\widehat{R})$. To establish this, note trivially that if $\widehat{r} = 0$, then $\widehat{R} = 0$ and $r^*(\widehat{R}) = \widehat{r}$.

Now let $\widehat{r}$ be any allocation such that $\widehat{r} \neq r^*(\widehat{R})$ and $0 < \sum_{j=1}^{N} \widehat{r}_j \equiv \widehat{R}$. Arguing by contradiction, assume $(\widehat{r}, \widehat{e}(r))$ is a subgame perfect equilibrium of Γ and consider the game $\Gamma_x(\widehat{R})$. Lemma 5 ensures that $\widehat{R} \leq \overline{R}$. That $(\widehat{r}, \widehat{e}(r))$ is subgame perfect also requires $\widehat{e}_j(r) > 0$ only if $j \in \mathrm{br}_A(\widehat{r})$.

Since $\widehat{r} \neq r^*(\widehat{R})$, Lemma 4 implies that there is a profitable deviation from $\widehat{r}$ against $\widehat{e}(r)$ in $\Gamma_x(\widehat{R})$. Let r' denote this deviation. Then, $\sum_{j=1}^{N} \lambda_j v_j(r'_j)\widehat{e}_j(r') + c_D(\widehat{R}) < \sum_{j=1}^{N} \lambda_j v_j(\widehat{r}_j)\widehat{e}_j(\widehat{r}) + c_D(\widehat{R})$ where these are losses so lower is better. Since r' is a feasible strategy in $\Gamma_x(\widehat{R})$, the total allocation $R' = \sum_{j=1}^{N} r'_j$ is bounded above by $\widehat{R}$ which means $c_D(R') \leq c_D(\widehat{R})$. Thus r' is a profitable deviation from $\widehat{r}$ against $\widehat{e}(r)$ in Γ as $\sum_{j=1}^{N} \lambda_j v_j(r'_j)\widehat{e}_j(r') + c_D(R') < \sum_{j=1}^{N} \lambda_j v_j(\widehat{r}_j)\widehat{e}_j(\widehat{r}) + c_D(\widehat{R})$. This contradicts the assumption that $(\widehat{r}, \widehat{e}(r))$ is an equilibrium and thereby establishes that $\widehat{r} = r^*(\widehat{R})$.

To see that $\widehat{R} \in W$ when $(r^*(\widehat{R}), \widehat{e}(r))$ is subgame perfect in Γ, assume the contrary. Then there exists an $R' \in [0, \overline{R}]$ such that $R' \neq \widehat{R}$ and $L(R') < L(\widehat{R})$. The defender's loss to playing $r^*(\widehat{R})$ is bounded below by $L(\widehat{R})$. To establish this bound, observe that $L(r^*(\widehat{R}), \widehat{e}(r)) = \sum_{j \in \mathrm{br}_A(r^*(\widehat{R}))} \lambda_j v_j(r^*_j(\widehat{R}))\widehat{e}_j(r^*(\widehat{R})) + c_D(\widehat{R}) \geq \lambda_{k(\widehat{R})} v_{k(\widehat{R})}(r^*_{k(\widehat{R})}(\widehat{R}))E^*(\widehat{R}) + c_D(\widehat{R}) = L(\widehat{R})$.

Now consider the defender's loss to playing $r^*(R')$. Theere are two cases to examine.

Case (i): $R' < \overline{R}$. Assume without loss of generality that $R' \notin$B. (Lemma 6(ii) ensures that $L(R)$ is continuous from the right. This implies that if $R' = B_j \leq \overline{R}$ then there would exist an $\varepsilon > 0$ and an allocation $R'' \notin$B such that $R'' \equiv B_j + \varepsilon < B_{j+1}$ and $L(R'') = L(B_j + \varepsilon) < L(\widehat{R})$.) As will be seen, the defender's loss to $r^*(R')$ against $\widehat{e}(r)$ in Γ is $L(R')$. But by assumption, $L(\widehat{R}) > L(R')$. This and the fact that the defender's loss to $\widehat{R}$ is bounded below by $L(\widehat{R})$ imply that $r^*(R')$ is a profitable deviation which contradicts the assumption that $(r^*(\widehat{R}), \widehat{e}(r))$ is subgame perfect in Γ. This contradiction will establish that $\widehat{R} \in W$ in case (i).

To see that the defender's loss to $r^*(R')$ against $\widehat{e}(r)$ in Γ is $L(R')$ note that the profile $(r^*(R'), \widehat{e}(r))$ is subgame perfect in $\Gamma_x(R')$. To wit, a profitable deviation from $r^*(R')$ in $\Gamma_x(R')$ would be profitable in Γ, and this would contradict the assumption that

$(r^*(\widehat{R}), \widehat{e}(r))$ is subgame perfect in Γ.

Since $R' \notin$B, no sites enter $\text{br}_A(r)$ at $r^*(R')$. Consequently, $r_j^*(R') > 0$ for all $j \in \text{br}_A(r^*(R))$ and, in particular, for $k(R')$. Proposition 2 now implies that the defender's equilibrium loss to playing $r^*(R')$ against $\widehat{e}(r)$ in $\Gamma_x(R')$ is $\sum_{j \in \text{br}_A(r^*(R'))} \lambda_j v_j(r_j^*(R'))$ $\widehat{e}(r^*(R')) + c_D(R') = \lambda_{k(R')} v_{k(R')}(r^*_{k(R')}(R'))E^*(R') + c_D(R') = L(R')$. But the expression on the left is also the defender's loss to playing $r^*(R')$ against $\widehat{e}(r)$ in Γ. Hence the defender's loss to $r^*(R')$ when $R' \leq \overline{R}$ is $L(R')$.

Case (ii): $R' = \overline{R}$. As shown above, the defender's loss to $\widehat{r} = r^*(\widehat{R})$ against $\widehat{e}(r)$ is bounded below by $L(\widehat{R})$. Since $L(\widehat{R}) > L(R') = L(\overline{R})$ and, by Lemma 6(ii), L is continuous from the right, there exists an $\varepsilon > 0$ such that $L(\widehat{R}) > L(\overline{R} + \varepsilon)$. Now consider the allocation $\widetilde{r}_j = r_j^*(\overline{R}) + \varepsilon/N$. The defender's loss to playing this is clearly $c_D(\overline{R} + \varepsilon) = L(\overline{R} + \varepsilon)$. Thus $\widetilde{r}$ is a profitable deviation from $\widehat{r}$, and this contradiction implies $\widehat{R} \in W$ in case (ii). ■

Proof of Proposition 4: Lemma 6(i) ensures $L'' > 0$ for all $R \notin$B. It follows that $L'(R) = \sigma(R)$ is strictly increasing at all $R \notin$B. Lemma 6(iv) establishes that σ jumps up at allocations in B. Hence, at most one allocation R^* satisfies $\lim_{R \uparrow R^*} L'(R) \leq 0 \leq \lim_{R \downarrow R^*} L'(R)$.

Let B^* be the be the generically unique equilibrium allocation. To see that that $B^* \geq R^*$, observe that the proof of Lemma 7 ensures that the unique minima of $L(R)$ satisfies $L'(R) = 0$ or is an element of B. This implies that if $B^* \notin$B, then $L(B^*) = 0$ and consequently $B^* = R^*$.

Now suppose $B^* = B_j \in$B for some j with $B^* < R^*$. Observe that $R^* \leq \overline{R}$ as $L'(R) = \sigma(R) > 0$ for $R > \overline{R}$. This leaves $B^* < R^* \leq \overline{R} = B_H$. Note further that $\sigma(B^*) = \lim_{R \downarrow B^*} L'(R) < \lim_{R \uparrow R^*} L'(R) \leq 0$ which implies $\sigma(B^*) < 0$. Lemma 6 also guarantees that σ and L are continuous from the right everywhere and continuous in the interval (B^*, B_{j+1}). Consequently, there exists an $\varepsilon > 0$ such that $\sigma(R) < 0$ for all $R \in (B^*, B^* + \varepsilon)$. L, therefore, is decreasing in this neighborhood. Continuity from the right ensures $L(B^*) = \lim_{R \downarrow B^*} L(R)$. Thus, $L(B^*) > L(R)$ for $R \in (B^*, B^* + \varepsilon)$. This contradicts the assumption that B^* minimizes L and leaves $B^* \geq R^*$. ■

Proof of Proposition 5: Claim (1) follows trivially. Since $M_A(r, \delta) = \delta M_A(r, 1)$, the solutions to $\min_{r \in \Delta(R)} M_A(r, \delta)$ and $\min_{r \in \Delta(R)} M_A(r, 1)$ are identical. Hence, $r^*(R, \delta) = r^*(R, 1)$ and, consequently, $\partial r^*(R, \delta)/\partial \delta = 0$.

Turning to (ii), $\mathrm{br}_A(r, \delta) = \{j : \delta \gamma_j v_j(r_j) = M_A(r, \delta)\} = \{j : \gamma_j v_j(r_j) = M_A(r, 1)\} = \mathrm{br}_A(r, 1)$. Letting $m_D(r, \delta) = \min\{\lambda_j v_j(r_j) : j \in \mathrm{br}_A(r, \delta)\}$, then $\mu(r, \delta) = \{j \in \mathrm{br}_A(r, \delta) : \lambda_j v_j(r_j) = m_D(r, \delta)\}$. But $\mathrm{br}_A(r, \delta) = \mathrm{br}_A(r, 1)$, so $m_D(r, \delta) = m_D(r, 1)$ and $\mu(r, \delta) = \{j \in \mathrm{br}_A(r, 1) : \lambda_j v_j(r_j) = m_D(r, 1)\} = \mu(r, 1)$.

As for (iii), note that $k(R, \delta) = \min \mu(r^*(R, \delta), \delta) = \min \mu(r^*(R, 1), 1) = k(R, 1)$. This means that $L(R, \delta) \equiv \lambda_{k(R,\delta)} v_{k(R,\delta)}(r^*_{k(R,\delta)}(R, \delta)) E^*(R, \delta) = \lambda_{k(R,1)} v_{k(R,1)}(r^*_{k(R,1)}(R, 1))$ $E^*(R, \delta)$. That $R^*(\delta) \notin \mathcal{B}$ implies $k(R^*(\delta), 1)$ is constant in a neighborhood of δ. The proof of Lemma 7 shows that $R^*(\delta) \notin \mathcal{B}$ also implies that $R^*(\delta)$ satisfies $\partial L(R^*(\delta), \delta)/\partial R = 0$. Implict differentiaion then yields $dR^*(R)/d\delta = -[\partial^2 L/\partial \delta \partial R]/[\partial^2 L/\partial R^2]$. Lemma 6(ii) ensures that the denominator is positive. Hence, it suffices to shows that $\partial^2 L/\partial \delta \partial R < 0$.

To obtain an expression for $\partial L(R^*(\delta), \delta)/\partial R$, multiply every γ_j by δ in the expression for $\sigma(R)$. This leaves:

$$\frac{\partial L(R^*(\delta), \delta)}{\partial R} = \delta \lambda_k D \left[\frac{E^*(R^*(\delta), \delta)}{\gamma_k} + \frac{v_k(r^*_k(R^*(\delta), 1)}{c''_A(E^*(R^*(\delta), \delta))} \right] + c'_D(R)$$

where $D = (\sum_{j \in \mathrm{br}_A(r^*(R^*(\delta),1))} [\gamma_j v'_j(r^*_j(R))]^{-1})^{-1}$ and the arguments of k have been supressed to simplify the notation. Differentiating with respect to δ and using the fact that $\partial E^*/\partial \delta = \gamma_k v_k(r^*(R^*(\delta), 1)/c''_A(E^*)$ yields

$$\frac{\partial^2 L(R^*(\delta), \delta)}{\partial \delta \partial R} = \frac{\lambda_k v_k}{c''_A} \left[2 - \frac{\delta \gamma_k v_k c'''_A}{(c''_A)^2} \right] D$$

Since $D < 0$ and $c'''_A \leq 0$, the expression on the right is negative.

To demonstrate (iv), it suffices to show that the marginal return on effort $F(\delta) = \delta \gamma_{k(R,\delta)} v_{k(R,\delta)}(r^*_{k(R,\delta)}(R^*(\delta), \delta))$ is increasing in δ. Suppressing the arguments on k, $F' > 0$ when

$$\gamma_k v_k(r^*_k(R^*(\delta), 1)) > \delta \gamma_k v'_k(r^*_k(R^*(\delta), 1)) \frac{dr^*_k(R^*(\delta), 1)}{dR} \frac{dR^*(\delta).}{d\delta}$$

Recalling that $dR^*/d\delta = -[\partial^2 L/\partial\delta\partial R]/[\partial^2 L/\partial R^2]$ and using Eq (4) gives

$$\frac{\partial^2 L}{\partial R^2} > \frac{\delta}{\gamma_k v_k} D \frac{\partial^2 L}{\partial\delta\partial R} = \frac{\lambda_k v_k}{c''_A}\left[\frac{2}{\gamma_k} - \frac{\delta v_k c'''_A}{(c''_A)^2}\right] D^2. \tag{6}$$

To verify that this inequality always holds, differentiate the expression for $\partial L/\partial R$ in the proof of (iii). Disregarding some positive terms leaves

$$\frac{\partial^2 L}{\partial R^2} > 2\lambda_k v'_k \frac{dr^*}{dR}\frac{\partial E^*}{\partial R} + \lambda_k v_k \frac{\partial^2 E^*}{\partial R^2}. \tag{7}$$

Differentiating $\delta\gamma_k v_k(r^*(R,1)) = c'_A(E^*(R,\delta))$ gives $\partial E^*/\partial R = [\delta\gamma_k v'_k/c''_A]dr^*/dR$ and

$$\frac{\partial^2 E^*}{\partial R^2} > \delta\gamma_k v'_k \frac{d^2 r^*}{dR^2} - \frac{\delta\gamma_k v'_k c'''_A}{(c''_A)^2}\frac{dr^*}{dR}\frac{\partial E^*}{\partial R}.$$

Substituting for d^2r^*/dR^2 and $\partial E^*/\partial R$ now yield $\partial^2 E^*/\partial R^2 > -\delta^2 c'''_A D^2/(c''_A)^2$. Using this inequality and the expressions for dr^*/dR and $\partial E^*/\partial R$ in inequality (7) establishes inequality (6). ■

References

911 Public Discourse Project Final Report. 2005. "Final Report on the 9/11 Commission Recommendations." Available at www.9-11pdp.org. Accessed on April 8, 2008.

Bier, Vicki, Santiago Oliveros, and Larry Samuelson. 2005. "Choosing What to Protect," *Journal of Public Economic Theory* 9(August):563-87.

Brodie, Bernard. 1959. *Strategy in the Missile Age.* Princeton: Princeton University Press.

Bueno de Mesquita, Ethan. 2007. "Politics and the Suboptimal Provision of Counterterror," *International Organization* 61(1):9-36.

DHS. 206. *National Infrastructure Protection Plan.* Washington, D.C.: Department of Homeland Security. Available at: http://www.dhs.gov/xlibrary/assets/NIPP_Plan.pdf. Accessed April 18, 2008.

Jervis, Robert. 1984. *The Illogic of American Nuclear Strategy.* Ithaca, NY: Cornell University Press.

Jervis, Robert. 1989. *The Meaning of the Nuclear Revolution.* Ithaca, NY: Cornell University Press.

Masse, Todd, Siobhan O'Neil, and John Rollins. 2007. "The Department of Homeland Security's Risk Assessment Methodology: Evolution, Issues, and Options for Congress." Washington, D.C.: Congressional Research Service.

Mueller, John. 2006. *Overblown.* New York: Free Press.

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

Powell, Robert. 1990. *Nuclear Deterrence Theory.* Princeton: Princeton University Press.

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

Schelling, Thomas. 1966. *Arms and Influence.* New Haven: Yale University Press.

Schmitt, Eric and Thom Shanker. 2008. "U.S. Adapts Cold-War Idea to Fight Terrorists," *New York Times,* March 18. Available at: http://www.nytimes.com/2008/03/18/washington/18terror.html?_r=1&sq=cold-war&st=nyt&adxnnl=1&oref=slogin&scp=2&adxnnlx=1207875823-NE7svYZIXzmfYK3azdPhbw# . Accesssed April 8, 2008.

Synder, Glenn H. 1961. *Deterrence and Defense.* Princeton: Princeton University Press.