DOI: 10.1017/S0003055407070530

Allocating Defensive Resources with Private Information about Vulnerability

ROBERT POWELL *University of California, Berkeley*

In many resources-allocation problems, strategic adversaries move sequentially and are likely to have private information about the effectiveness of their spending. A defender, for example, often has to allocate its defensive resources before an attacker (e.g., a terrorist group) decides where to strike. Defenders are also likely to have private information about the vulnerability of the things they are trying to protect. Sequential decisions and private information about effectiveness creates a dilemma for a defender. Allocating more to a highly vulnerable site reduces the expected losses should that site be attacked but may also signal that that site is more vulnerable and thereby increase the probability of an attack. Modeling this tradeoff as a signaling game, the analysis shows that secrecy concerns generally swamp vulnerability concerns when more vulnerable sites are weakly harder to protect on the margin. The defender pools in equilibrium, that is, allocates its resources in the same way, regardless of the level of vulnerability. If more vulnerable sites are easier to protect on the margin, vulnerability concerns may swamp secrecy concerns.

In many settings, strategic adversaries—be they political parties, lobbies, armies, or firms—have to decide how to allocate limited resources across multiple fronts. The more an actor devotes to one front, the more likely it is to prevail on that front. However, an actor's overall payoff depends on how well it does on all of the fronts. For example, political parties may compete for electoral support by offering transfers, bribes, or pork to voters or districts which in turn support the party offering the highest payment (Groseclose and Snyder 1996; Laslier and Picard 2002; Myerson 1993; Snyder 1989). Competing firms have to decide how to allocate limited R&D budgets among various projects. Or opposing armed forces have to decide how much to devote to each battlefield where the probability of winning on that battlefield depends on the relative forces committed to that front (Luce and Raiffa 1957; Shubik 1982; Shubik and Weber 1981).

Such allocation problems have long been studied in the context of Colonel Blotto games and, more recently, as all-pay auctions. In these static games, the two adversaries make their allocation decisions simultaneously or secretly. Thus neither player can condition its own allocation on the other player's allocation.[1]

However, some allocation problems are sequential. One side has to allocate its resources first and the other gets to respond to that allocation. For example, two vote-buyers (e.g., lobbies) with opposing policy preferences make sequential bids for legislators' votes in Groseclose and Snyder's (1996) and Morgan and Vardy's (2007) studies of vote-buying and supermajorities. More generally, defenders often have to decide how to allocate some or all of their resources, and an attacker then chooses where to strike after observing (perhaps imperfectly) what the defender has done. Frequently, one or both adversaries will also have private information about the effectiveness of their spending on different fronts. In particular, defenders are likely to have private information about the relative vulnerabilities of the things they are trying to protect.

For instance, the United States undertook a massive effort to secure its critical infrastructure and key assets following the attacks of September 11, 2001, with the federal government alone now spending about $20 billion dollars a year on infrastructure protection.[2] As part of this effort, the Department of Homeland Security identified 1,700 highest priority sites and intended to visit each of them "to assess their vulnerabilities to various forms of attack . . . " (Moteff 2006, 38). These assessments are likely to reveal previously hidden vulnerabilities.

Private information about vulnerabilities creates a dilemma for the defender. On the one hand, the defender might be tempted to invest the resources needed to harden a very vulnerable site and thereby reduce the expected losses should that site be attacked. On the other hand, the defender's efforts to protect a site might signal to an attacker that the site is relatively more vulnerable and thereby make an attack on that site more likely.

This paper models this dilemma as a signaling game in which a defender has to decide how to divide its resources between two sites. An attacker sees this allocation and then decides which site to strike. The more the defender spends on a site, the lower its expected losses if that site is hit. Because the defender has private information about the vulnerability of one of the sites, it must also be concerned about what the attacker will infer about the sites' relative vulnerability from the defender's allocation.

An equilibrium of this game is fully separating if each level of vulnerability leads the defender to make

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

I thank Scott Ashworth, Bruce Bueno de Mesquita, Ernesto Dal Bó, Santiago Oliveros, and Jacob Shapiro for helpful comments, criticisms and discussion.

[1] Early work on Blotto games includes Borel 1921, Tukey 1949, and Blackett 1958. More recent contributions are Golman and Page 2006 and Roberson 2006. See Kvasov 2006 on all-pay auctions.

[2] See Moteff 2006 for an overview of these programs.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

a different allocation. Because each observable allocation is associated with a unique underlying level of vulnerability in this kind of equilibrium, the attacker can infer the level of vulnerability from the defender's allocation and then strike the site that offers the highest expected payoff. In a sense, vulnerability concerns swamp secrecy concerns in a separating equilibrium in that the defender reveals the level of vulnerability in order to allocate the resources needed to deal with that level of vulnerability optimally.

In a pooling equilibrium, by contrast, the defender divides its resources between the two sites in the same way, regardless of what it knows about their relative vulnerability. The attacker now cannot infer anything about the sites' relative vulnerability when it sees this allocation. Secrecy concerns in this type of equilibrium swamp vulnerability concerns in that the defender's optimal strategy is to reveal nothing about the level vulnerability.

In a semiseparating equilibrium, the defender's allocation varies to some degree with the level of vulnerability. The attacker is able to infer something about the relative vulnerability of the two sites but cannot infer the exact level of relative vulnerability (as it does in a separating equilibrium). In a semiseparating equilibrium, neither vulnerability nor secrecy concerns dominate the other. The defender reveals some information about the level of vulnerability in order to partially address the vulnerability issue, but the defender does not reveal the exact level of vulnerability.

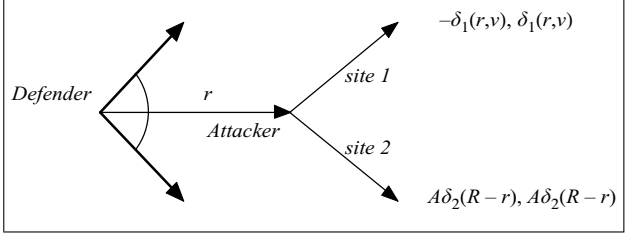
The following analysis shows that secrecy concerns generally swamp vulnerability concerns when more vulnerable sites are harder to protect on the margin. The defender pools on the allocation that is optimal with respect to an "average" vulnerability. This allocation spends less on high-vulnerability sites and more on low-vulnerability sites than the defender would spend on them if signaling were of no concern. Nevertheless, uncertainty makes the defender better off regardless of the actual level of vulnerability. As the average level of vulnerability of a site increases, the defender devotes more to protecting that site.

The next section specifies the model and summarizes the equilibria when the level of vulnerability is common knowledge. The subsequent section characterizes the equilibria when the defender has private information about the relative vulnerability of the two sites, and more vulnerable sites are also harder to protect on the margin. A final section discusses the opposite case when more vulnerable sites are easier to protect on the margin. This section shows by example that separating equilibria sometimes exist in these circumstances and explains why separation is sensitive to the effect that the level of vulnerability has on the marginal benefits of spending more.

THE COMPLETE INFORMATION MODEL AND EQUILIBRIA

A simplified version of Powell's (2007b) resource-allocation game provides a point of departure. As illustrated in Figure 1, the defender begins the game by deciding how to divide R resources between the two sites it is trying to protect. That is, the defender chooses an $r \in [0, R]$ where r and $R - r$ are the resources allocated to sites 1 and 2, respectively. After observing r, the attacker decides which site to strike.

FIGURE 1. The Signaling Game



Spending more on a site reduces the chances of a successful attack but with declining effect. Formally, let $\delta_2(r)$ be the probability that an attack on site 2 succeeds if the defender spends r. Then $\delta_2' < 0$ and $\delta_2'' > 0$. The probability that an attack on site 1 succeeds depends both on the resources allocated to that site and on the inherent vulnerability v of the site. Let $\delta_1(r, v)$ denote this probability, where $\partial\delta_1/\partial r < 0$ and $\partial^2\delta_1/\partial r^2 > 0$. The more vulnerable site 1 (the larger v), the more resources the defender has to devote to that site to reduce the probability of a successful attack to a given level: $\partial\delta_1/\partial v > 0$.

The defender suffers a loss of one if site 1 is destroyed and a loss of $A > 0$ if site 2 is destroyed. Thus, the defender's expected loss if it allocates r to site 1 and the attacker hits that site with probability α is $L(r, \alpha) = \alpha\delta_1(r, v) + (1 - \alpha)A\delta_2(R - r)$. As for the attacker's payoff, the game is assumed to be zero sum, so the defender's losses are the attacker's gains. The attacker receives one if it successfully strikes site 1 and A if it destroys site 2. If the defender allocates r to site 1 and the attacker strikes this site with probability α, the attacker's expected payoff is $L(r, \alpha)$. In the complete-information model, v along with all of the other parameters of the model are common knowledge.

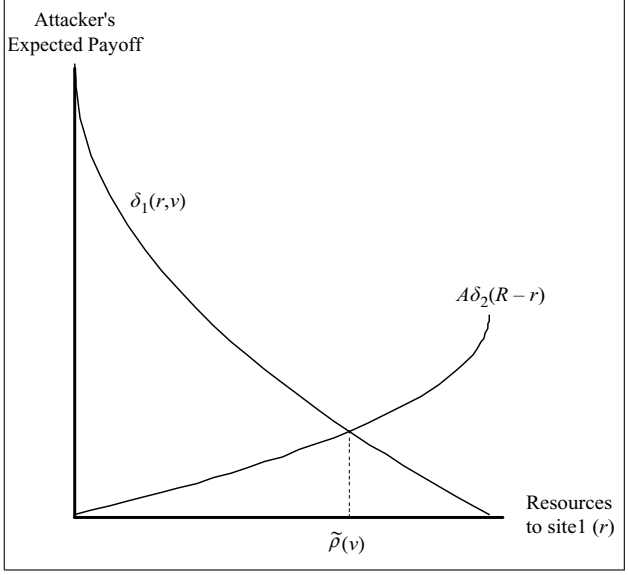
As will be seen, whether secrecy or vulnerability concerns dominate the defender's allocation strategy depends on the way that higher levels of vulnerability affect the marginal return to allocating additional resources to that site. The analysis begins with the assumption that more vulnerable sites are weakly harder to protect on the margin. That is, the more vulnerable a site (larger v), the less effect additional spending has on reducing that sites' vulnerability. (The implications of the opposite assumption will be examined below.) Formally:

Assumption 1 (More vulnerable sites are harder to defend on the margin): $\partial^2\delta_1/\delta v\delta r \geq 0$.[3]

[3] This is suggestive of the standard, single-crossing assumption in signaling games, which ensures that different types' indifference curves only cross once. But the suggestion is misleading. Each of

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

FIGURE 2. Minmaxing the Attacker



To trace some of the implications of this assumption, suppose that the defender were concerned about the risk of an accident rather than a deliberate attack. If the probabilities of an accident at sites 1 and 2 are a_1 and a_2 with a_1 independent of the level of vulnerability of the site, then the defender's expected loss to allocating r to site 1 is $a_1\delta_1(r, v) + a_2A\delta_2(R - r)$. The optimal allocation $r^*(v)$ satisfies the first-order condition by equating the marginal gains to additional spending at each site; that is, $r^*(v)$ satisfies $a_1\partial\delta_1(r^*(v), v)/\partial r = a_2A\delta_2'(R - r^*(v))$.

If more vulnerable sites are harder to protect on the margin, the marginal returns to investing in site 1 decrease as v increases. As a result, the defender shifts resources away from site 1 and to site 2 where the marginal returns are higher: $r^*(v)$ is decreasing in v. If more vulnerable sites are easier to protect on the margin ($\partial^2\delta_1/\delta v\delta r < 0$), the return to defending site 1 rises with v and the defender spends more. That is, $r^*(v)$ increases in v.[4]

The defender's equilibrium allocation in the complete-information game is unique. Let $\widetilde{\rho}(v)$ be the allocation that equalizes the attacker's payoffs to striking sites 1 and 2 when the level of vulnerability is v; that is, $\widetilde{\rho}(v)$ uniquely satisfies $\delta_1(r, v) = A\delta_2(R - r)$.[5] As Figure 2 illustrates, this allocation minmaxes the attacker; that is, $\widetilde{\rho}(v)$ is the allocation that minimizes the maximum payoff the attacker can obtain after observing r. Formally, $\widetilde{\rho}(v)$ solves $\min_{r\in[0,R]}\max\{\delta_1(r, v), A\delta_2(R - r)\}$, and is what the defender does in any subgame perfect equilibrium.

To establish that the only equilibrium allocation is $\widetilde{\rho}(v)$, observe that the defender can profitably deviate from any $r' \neq \widetilde{\rho}(v)$. Assume, first, that $r < \widetilde{\rho}(v)$. Then the defender is spending too little on site 1 in that the attacker's payoff to striking that site is greater than that of striking site 2: $\delta_1(r, v) > A\delta_2(R - r)$. The attacker therefore hits site 1, which leaves the defender with an expected loss of $\delta_1(r, v)$ whenever $r < \widetilde{\rho}(v)$. Hence, the defender can profitably deviate to any r' between r and $\widetilde{\rho}(v)$ because $\delta_1(r, v) > \delta_1(r', v)$ for all $r' \in (r, \widetilde{\rho}(v))$ as the defender's expected loss is decreasing in r.

Conversely, the defender is spending too much on site 1 when $r > \widetilde{\rho}(v)$ and, as a result, makes site 2 the more attractive target. That is, $A\delta_2(R - r) > \delta_1(r, v)$ when $r > \widetilde{\rho}(v)$. The attacker consequently strikes site 2 whenever $r > \widetilde{\rho}(v)$, leaving the defender with an expected loss of $A\delta_2(R - r)$. Now the defender can profitably deviate from r to any r' between and $\widetilde{\rho}(v)$ and r.

Thus, the only possible equilibrium allocation is $\widetilde{\rho}(v)$, and it is easy to verify that the following are equilibria of the complete information game: the defender allocates $\widetilde{\rho}(v)$ to site 1; the attacker hits site 1 if $r < \widetilde{\rho}(v)$; site 2 if $r > \widetilde{\rho}(v)$; and strikes site 1 with any probability $\alpha \in [0, 1]$ if $r = \widetilde{\rho}(v)$. More formally,

Proposition 1. *Subgame perfect equilibria of the complete-information game exist, and the defender's allocation is the same in all of them. This allocation, $\widetilde{\rho}(v)$, minmaxes the attacker; that is, $\widetilde{\rho}(v)$ is the unique solution to* $\min_{r\in[0,R]}\max\{\delta_1(r, v), A\delta_2(R - r)\}$.[6]

A direct implication of Proposition 1 is that the more vulnerable site 1 is; that is, the larger v, the more the defender allocates to it. Figure 3 illustrates the point. If $v' > v$, then $\widetilde{\rho}(v') > \widetilde{\rho}(v)$. Formally, implicit differentiation of $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R - \widetilde{\rho}(v))$ yields $d\widetilde{\rho}(v)/dv = -(\partial\delta_1/\partial v)/[\partial\delta_1/\partial r + A\delta_2'] > 0$. Hence,

Corollary 1. *The more vulnerable site 1, the more the defender allocates to that site when there is complete information about vulnerability $d\widetilde{\rho}(v)/dv = -(\partial\delta_1/\partial v)/[\partial\delta_1/\partial r + A\delta_2'] > 0$.*

Corollary 1 highlights the difference between protecting against strategic threats and protecting against accidents. As noted earlier, the optimal allocation against a stochastic, nonstrategic threat equates the marginal benefits of additional spending at each of the two sites. This implies that the optimal level of spending $r^*(v)$ is increasing in v if $\partial^2\delta_1/\partial v\partial r > 0$ and decreasing if $\partial^2\delta_1/\partial v\partial r > 0$. By contrast, the optimal allocation against a strategic threat $\widetilde{\rho}(v)$ is always increasing in v. The reason why is that the defender equates levels when defending against a strategic threat by making

v's indifference curves generally crosses $\widehat{v}$'s at least twice regardless of the sign of $\partial^2\delta_1/\delta v\delta r$.

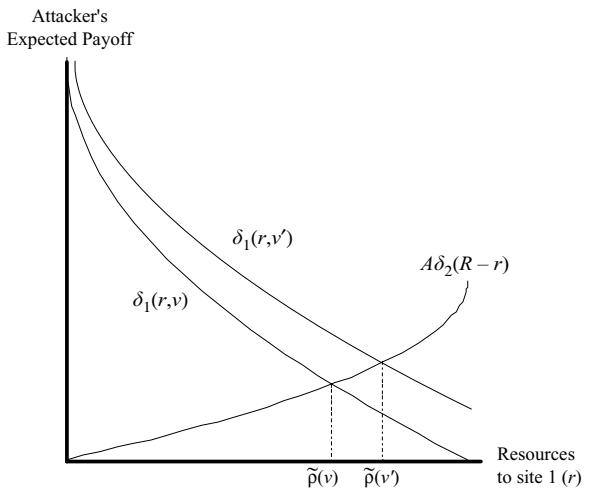
[4] Differentiating the first-order condition yields $dr^*/dv = -a_1(\partial^2\delta_1/\partial r\partial v)/[a_1\partial^2\delta_1/\partial r^2 + a_2A\delta_2'']$. So the sign of dr^*/dv is the opposite of the sign of $\partial^2\delta_1/\partial r\partial v$.

[5] The expected payoff $\delta_1(r, v)$ is strictly decreasing in r, and $A\delta_2(R - r)$ is strictly increasing. So a unique r solves $\delta_1(r, v) = A\delta_2(R - r)$ for any value of v.

[6] Powell (2007a) proves this result. He also shows the defender's minmaxing the attacker remains the unique equilibrium allocation even if the game is non-zero sum and regardless of whether the attacker can observe the defender's allocation before deciding which site to attack.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

FIGURE 3. The Minmaxing Allocation Is Increasing in *v*



the attacker indifferent between striking sites 1 and 2, whereas the defender equates margins when defending against stochastic, nonstrategic threats.

PRIVATE INFORMATION ABOUT VULNERABILITY

To formalize the tradeoff between hardening a very vulnerable site and calling attention to the vulnerability of that site, suppose the defender knows the vulnerability v of site 1 but the attacker does not. Uncertain of the exact level of vulnerability, the attacker believes that v is distributed over $[\underline{v}, \overline{v}]$ according to the cumulative distribution G which has a strictly positive density over $(\underline{v}, \overline{v})$. G is common knowledge.[7]

Assume further that if the defender invests everything in site 1, the attacker prefers to strike site 2 regardless of the value of v. Formally, $\delta_1(R, \overline{v}) < A\delta_2(0)$. Similarly, the attacker prefers hitting site 1 if the defender allocates all of its resources to site 2: $\delta_1(0, \underline{v}) > A\delta_2(R)$. These restrictions ensure that the defender faces a real allocation decision and that the equilibrium is not simply to allocate everything to one site in the knowledge that the attacker will strike that site.

The defender and attacker are playing a signaling game in which the defender's allocation may reveal information to the attacker about the relative vulnerability of the two sites. Suppose, for example, that the defender knows that site 1 is very vulnerable (v is high). On the one hand, the high level of vulnerability of site 1 inclines the defender to spend more on that site in order to reduce the expected loss if that site is attacked. On the other hand, spending more on site 1 may signal that that site is quite vulnerable and thereby make an attack on it more likely. The equilibria of the signaling game formalize the resolution of this tradeoff.

To simplify matters, the model assumes that the attacker can perfectly observe the defender's allocation r. A more realistic assumption would be that the attacker imperfectly observes r. If, for example, the defender allocates r to site 1, the attacker might observe or detect signs s that lead it to believe that the allocation is $r = s + \varepsilon$, where ε is an "observation error" distributed according to some mean-zero probability distribution F. Although more realistic, such a model would be much harder to analyze.

If, at the extreme, the attacker is unable to observe r or anything correlated with it, signaling is clearly no longer a concern. The game is now static, and the restrictions above ensuring that the defender faces a meaningful allocation problem imply that the attacker must mix in equilibrium, that is, attack site 1 with a probability α^* strictly between zero and one.[8] Type v's best response to a probability of attack α^*, $\rho^*(\alpha, v)$, is the allocation that minimizes its expected loss against α^*. That is, $\rho^*(\alpha^*, v)$ satisfies the first-order condition $\alpha^* \partial\delta_1(\rho^*, v)/\partial r - (1 - \alpha^*)A\delta_2'(R - \rho^*) = 0$.[9]

The remainder of this section characterizes the equilibria of the signaling game and shows that secrecy concerns generally swamp vulnerability concerns. The defender pools on a single allocation. This allocation is optimal with respect to an "average" vulnerability. But it devotes less to high-vulnerability sites and more to low-vulnerability sites than the defender would spend on them if signaling were of no concern.

Formally, a pure strategy for the defender specifies the defender's allocation as a function of its private information about the vulnerability: $\rho : [\underline{v}, \overline{v}] \to [0, R]$.[10] A strategy for the attacker specifies the probability of attacking site 1 as a function of the allocation the attacker observes: $\alpha : [0, R] \to [0, 1]$. As for what the attacker believes about the vulnerability of site 1 after seeing a given allocation to that site, let M be the set of probability distributions over $[\underline{v}, \overline{v}]$ and let $\mu(v \mid r) \in M$ for all $r \in [0, R]$ denote the attacker's

[7] In the present analysis, the defender has private information about the sites' relative vulnerability, but the actors' preferences are common knowledge. Bier, Oliveros, and Samuelson (2005) and Powell (2007b) consider the allocation problem when the defender is unsure of the attacker's payoffs. Analyzing situations in which a defender has private information about vulnerability and is unsure of the attacker's payoffs is open issue for future research.

[8] To see that the attacker must randomize between hitting sites 1 and 2, suppose the attacker strikes site 1 for sure, $\alpha^* = 1$. Then v 's best reply is to allocate everything to that site. This leaves the atacker with an expected payoff of $\int_{\underline{v}}^{\overline{v}} \delta_1(R, v)dG$. But this is less than the payoff to going after site 2: $\int_{\underline{v}}^{\overline{v}} \delta_1(R, v)dG < \int_{\underline{v}}^{\overline{v}} \delta_1(R, \overline{v})dG < \int_{\underline{v}}^{\overline{v}} A\delta_2(0)dG = A\delta_2(0)$ where the second inequality follows from the assumption that $\delta_1(R, \overline{v}) < A\delta_2(0)$. This contradiction ensures $\alpha^* < 1$, and a similar argument shows $\alpha^* > 0$.

[9] Convexity ensures that v's best response is unique. To pin down the equilibrium value of α^*, note that the attacker must be indifferent between hitting sites 1 and 2 because it is mixing. Hence, $\int_{\underline{v}}^{\overline{v}} \delta_1(\rho^*(\alpha^*, v), v)dG = \int_{\underline{v}}^{\overline{v}} A\delta_2(R - \rho^*(\alpha^*, v))dG$. Implicit differentiation of the first-order condition defining ρ^* yields $\partial\rho^*/\partial\alpha^* > 0$. Consequently, the integral on the left of the previous equality is decreasing in α^*, and the integral on the right is increasing. The equilibrium value of α^* is the unique value at which equality holds. Powell (2006) explicitly solves for α^* in the separable case in which $\partial^2\delta_1/\partial r\partial v = 0$.

[10] Only strategy profiles and equilibria in which the defender plays a pure strategy are considered.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

beliefs about v after observing allocation r. Then, a perfect Bayesian equilibrium (PBE) is a strategy profile (ρ, α) and system of beliefs μ, such that (i) the defender can never profitably deviate from $\rho(v)$ given the attacker's strategy $\alpha(r)$; (ii) $\alpha(r)$ maximizes the attacker's payoff against r given $\mu(v \mid r)$; and (iii) $\mu(v \mid r)$ is derived from G and r via Bayes' rule.[11]

Two lemmas help characterize the equilibria. Recall that $\widetilde{\rho}(v)$ equalizes the attacker's payoff to striking the two sites if the level of vulnerability is v. That is, $\widetilde{\rho}(v)$ satisfies $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R - \widetilde{\rho}(v))$. Because the game is zero sum, this allocation also equalizes the defender's expected losses to an attack on sites 1 and 2. Lemma 1 ensures that v's equilibrium loss is no more than $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R - \widetilde{\rho}(v))$, which v can guarantee itself by devoting $\widetilde{\rho}(v)$ to site 1.

Lemma 1. *Let the allocation $\widetilde{\rho}(v)$ be defined by $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R - \widetilde{\rho}(v))$. Then type v's loss in any PBE is no more than its loss to allocating $\widetilde{\rho}(v)$ to site 1 which is $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R - \widetilde{\rho}(v))$.*

Proof. Trivially, v can always allocate $\widetilde{\rho}(v)$ to site 1, which brings an expected loss of $\alpha(\widetilde{\rho}(v))\delta_1(\widetilde{\rho}(v), v) + [1 - \alpha(\widetilde{\rho}(v))]A\delta_2(R - \widetilde{\rho}(v)) = A\delta_2(R - \widetilde{\rho}(v))$. Hence, v's equilibrium loss can be no greater than $A\delta_2(R - \widetilde{\rho}(v))$. ■

Lemma 2 shows that if v' separates; that is, no other v'' allocates $\rho(v')$, then v' must be spending $\widetilde{\rho}(v')$ on site 1 and, as a result, suffers a loss of $\delta_1(\widetilde{\rho}(v'), v') = A\delta_2(R - \widetilde{\rho}(v'))$. As Proposition 1 shows, $\widetilde{\rho}(v')$ is also what the defender would allocate to site 1 if there were no uncertainty about the level of vulnerability. Lemma 2 thus demonstrates that if the defender foregoes all of the benefits of secrecy by acting in a way that fully discloses the vulnerability of site 1, then the defender allocates its resources optimally with respect to the vulnerability issue. That is, the defender allocates its resources exactly as it would have had site 1's vulnerability been common knowledge at the outset.

Lemma 2. *Let $\{\rho(v), \alpha(r), \mu\}$ be any PBE in which v' separates; that is, v' alone offers $\rho(v')$. Then $\rho(v') = \widetilde{\rho}(v')$.*

Proof. Arguing by contradiction, suppose v' allocates $r' > \widetilde{\rho}(v')$. Because v' separates, the attacker infers it is facing v' after observing r'. Given this inference and r', the attacker strictly prefers attacking site 2. That is, $r' > \widetilde{\rho}(v')$ implies $A\delta_2(R - r') > A\delta_2(R - \widetilde{\rho}(v')) = \delta_1(\widetilde{\rho}(v'), v') > \delta_1(r', v')$. Hence, the attacker strikes site 2, leaving v' with an expected loss of $A\delta_2(R - r')$. But, as Lemma 1 shows, v' can always obtain an expected loss of $A\delta_2(R - \widetilde{\rho}(v'))$ by playing $\widetilde{\rho}(v')$, and $r' > \widetilde{\rho}(v')$ implies that this loss is less than $A\delta_2(R - r')$. This contradiction ensures that v' never separates with an allocation greater than $\widetilde{\rho}(v')$. A similar argument shows that if $r' < \widetilde{\rho}(v')$, then the attacker hits site 1 after observing r' and v' could once more profitably deviate to $\widetilde{\rho}(v')$. Thus v' must offer $\widetilde{\rho}(v')$ if it separates. ■

Propositions 2 and 3 show that secrecy concerns swamp vulnerability concerns in that the defender generally pools in equilibrium. All v allocate r^* to site 1 where r^* uniquely satisfies $\int \delta_1(r^*, v)dG = A\delta_2(R - v^*)$. In effect, the defender plays as if it were common knowledge that the level vulnerability was the "average" $v^* \equiv \widetilde{\rho}^{-1}(r^*)$.

Because the defender pools on the same allocation regardless of the actual level of v, the attacker cannot infer anything about v after observing the defender's allocation of $\widetilde{\rho}(v^*)$. But not revealing anything about v comes at a cost. The defender is not allocating its resources to deal with the vulnerability issue in the optimal way. As Proposition 1 shows, v would allocate $\widetilde{\rho}(v)$ were the level of vulnerability common knowledge and signaling not an issue. Signaling concerns therefore lead the defender to reduce its spending on high-vulnerability sites ($v > v^*$) from $\widetilde{\rho}(v)$ to $\widetilde{\rho}(v^*)$. Conversely, signaling concerns lead the defender to spend more on low-vulnerability sites than it otherwise would (i.e., $\widetilde{\rho}(v) < \widetilde{\rho}(v^*)$ if $v < v^*$).

To gain some intuition for the forces undermining separation, suppose an interval of types between v' and v'' did separate, that is, $v \in (v', v'')$ separate. Lemma 1 implies that each v allocates $\widetilde{\rho}(v)$ to site 1. Assume further that the attacker strikes site 1 with probability $\widehat{\alpha}$ after observing allocation $\widetilde{\rho}(\widehat{v})$ for some $\widehat{v} \in (v', v'')$.

This probability cannot be too small if higher vulnerability types ($v > \widehat{v}$) are to be deterred from mimicking $\widehat{v}$ by allocating $\widetilde{\rho}(\widehat{v})$ to site 1. If, that is, $\widehat{\alpha}$ is too small, then higher vulnerability types will want to "buy" the much lower probability of attack $\widehat{\alpha}$ on the more vulnerable site 1 in return for the cost of allocating somewhat less to that site, that is, $\widetilde{\rho}(\widehat{v})$ rather than $\widetilde{\rho}(v)$. Conversely, the probability that the attacker strikes site 1 cannot be too high if lower vulnerability types ($v < \widehat{v}$) are to be deterred from deviating to $\widetilde{\rho}(\widehat{v})$. If $\widehat{\alpha}$ is too high, then lower vulnerability types will want to buy the correspondingly lower probability of attack on the relatively more vulnerable site 2. Hence, higher values of $\widehat{\alpha}$ discourage higher vulnerability types from deviating but encourage lower vulnerability types to deviate and vice-versa. As show formally below, a unique $\widehat{\alpha}$ deters types arbitrarily close to $\widehat{v}$ from deviating to $\widetilde{\rho}(\widehat{v})$.

If this value of $\widehat{\alpha}$ also deters types farther away from $\widehat{v}$ from wanting to deviate to $\widetilde{\rho}(\widehat{v})$, separation may be possible. But when more vulnerable sites are harder to defend on the margin, types farther away from $\widehat{v}$ will want to deviate and separation breaks down. This leaves:

[11] Defining PBE's with a continuum of types also raises an additional technical issue. Consider, for example, a separating equilibrium in which each v makes a distinct allocation. With a continuum of types, the prior probability of any one type is zero, so Bayes' rule places no restriction on the attacker's beliefs following any allocation. Bayes' rule does not require the attacker to believe it is facing the one type who was to make this allocation in equilibrium. To finesse this issue, assume that if the nonempty set of types allocating r' to site 1 has zero measure, then the support of the attacker's beliefs following r' must be contained in the closure of $\{v : \rho(v) = r'\}$. See Ramey 1996 for a definition of a sequential or perfect Bayesian equilibrium with a continuum of types.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

Proposition 2. *No open interval of types (v', v'') separate and hence no separating equilibria exist.*

Proof. See the Appendix.

Now define a strategy to be *simple* if whenever v' and v'' allocate the same amount to site 1, all types between v' and v'' also allocate that amount to site 1.

Definition 1. *An allocation strategy $\rho(v)$ is* simple *if $\rho(v) = \rho(v')$ for all $v \in (v', v'')$ whenever $\rho(v') = \rho(v'')$.*

Simplicity is weaker than monotonicity. The strategy $\rho(v)$ is monotonic if ρ is either weakly increasing or decreasing in v. A monotonic strategy is therefore simple, but a simple strategy may not be monotonic. For example, $\rho(v) = r$ for $v \in [\underline{v}, v']$, $\rho(v) = r - \varepsilon$ for $v \in (v', v'')$, and $\rho(v) = r - \frac{1}{2}\varepsilon$ for $v \in [v'', \overline{v}]$ is simple but not monotonic.

A semiseparating equilibrium using a simple strategy is a set of cut points $\underline{v} = k_0 \leq k_1 \leq \cdots \leq \overline{v}$, a set of allocations such that all types between adjacent cut points pool on the same allocation and no type outside this interval allocates that amount to site 1, and there are at least two distinct allocations. Proposition 3 shows that no such equilibria exist.[12]

Proposition 3. *If the defender uses a simple strategy, there are no semiseparating equilibria.*

Proof. See the Appendix.

Security concerns therefore swamp vulnerability concerns in that the defender must pool on a common allocation. Moreover, all pooling equilibria have the same equilibrium path. To see why, observe that pooling implies that the attacker uses a mixed strategy, that is, strikes site 1 with a probability strictly between zero and one. To establish this, suppose the defender pools on $\widehat{r} \in (\widetilde{\rho}(\underline{v}), \widetilde{\rho}(\overline{v}))$.[13] If the attacker is certain to attack site 1 in response to $\widehat{r}$ (i.e., if $\alpha(\widehat{r}) = 1$), then Lemma 1 guarantees that any v such that $\widetilde{\rho}(v) > \widehat{r}$ can profitably deviate by allocating the larger amount $\widetilde{\rho}(v)$ to site 1. Formally, the allocation $\widetilde{\rho}(v)$ brings $\alpha(\widehat{r})\delta_1(\widetilde{\rho}(v), v) + (1 - \alpha(\widehat{r}))A\delta_2(R - \widetilde{\rho}(v)) = \delta_1(\widetilde{\rho}(v), v)$. But $\widetilde{\rho}(v) > \widehat{r}$ implies $\delta_1(\widehat{r}, v) > \delta_1(\widetilde{\rho}(v), v)$ where, recall, these are losses so lower is better. Hence, v can profitably deviate to $\widetilde{\rho}(v)$, and this contradiction implies $\alpha(\widehat{r}) < 1$.

If the attacker is sure to hit site 2 after seeing $\widehat{r}$ ($\alpha(\widehat{r}) = 0$), any v for which $\widetilde{\rho}(v) < \widehat{r}$ can profitably deviate by allocating less to site 1, namely, $\widetilde{\rho}(v)$ rather than $\widehat{r}$, and consequently more to site 2. That is, $\widetilde{\rho}(v) < \widehat{r}$ implies $A\delta_2(R - \widehat{r}) > A\delta_2(R - \widetilde{\rho}(v))$. This second contradiction along with the first ensures that the attacker mixes after observing $\widehat{r}$: $0 < \alpha(\widehat{r}) < 1$.

Because the attacker mixes, it must be indifferent between going after site 1 and going after site 2. Pooling also means that the attacker's posterior beliefs are unchanged from its prior beliefs. These facts imply $\int \delta_1(\widehat{r}, v)dG = \int A\delta_2(R - \widehat{r})dG = A\delta_2(R - \widehat{r})$. A unique $\widehat{r}$ satisfies this condition as the expression on the right is increasing in $\widehat{r}$ and the expression on the left is decreasing.[14] Let r^* denote this solution and $v^* \equiv \widetilde{\rho}^{-1}(r^*)$.

As for the equilibrium probability of attack, pooling on r^* means that the v close to v^* must not be able to profitably deviate from $r^* = \widetilde{\rho}(v^*)$ to $\widetilde{\rho}(v)$. A single $\alpha(r^*)$ does this and is given by $\alpha(\widetilde{\rho}(v^*)) = \Phi(v^*)$, where

$$\Phi(v) \equiv \frac{A\delta_2'(R - \widetilde{\rho}(v))}{\partial\delta_1(\widetilde{\rho}(v), v)/\partial r + A\delta_2'(R - \widetilde{\rho}(v))} = \left[1 + \frac{\partial\delta_1(\widetilde{\rho}(v), v)/\partial r}{A\delta_2'(R - \widetilde{\rho}(v))}\right]^{-1}.$$

and, clearly, $0 < \Phi(v) < 1$.

All of this leaves:

Proposition 4. *Pooling perfect Bayesian equilibria exist and all of them share the same equilibrium path. All v pool on $\widetilde{\rho}(v^*)$ and the attacker strikes site 1 with probability $\Phi(v^*)$ where v^* uniquely satisfies $\int \delta_1(\widetilde{\rho}(v^*), v)dG = A\delta_2(R - \widetilde{\rho}(v^*))$.*

Proof. See the Appendix.

Comparative statics follow immediately. Loosely, as the average vulnerability of site 1 increases, the defender allocates weakly more to that site which in turn is weakly more likely to be attacked. More precisely, suppose that the prior distribution over v is H rather than G, where H first-order stochastically dominates G.[15] Assume further that the defender pools on r_h when v is distributed according to H and r^* when beliefs follow G. Then $r_h \geq r^*$ and $\alpha(r_h) \geq \alpha(r^*)$.

Corollary 2. *Suppose the defender pools on r_h when when v is distributed according to H and r^* when v is distributed according to G where H first-order stochastically dominates G. Then $r_h \geq r^*$ and $\alpha(r_h) \geq \alpha(r^*)$.*

Proof. Because $\delta_1(r, v)$ is increasing in v, first-order stochastic dominance implies that $\int \delta_1(r^*, v)dH \geq \int \delta_1(r^*, v)dG$. And, r^* satisfies $\int \delta_1(r^*, v)dG = A\delta_2(R - r^*)$. So, $\int \delta_1(r^*, v)dH \geq A\delta_2(R - r^*)$. The left side is decreasing in r and the right side is increasing. So the

[12] It appears that there are no semiseparating equilibria even if the defender does not use a simple strategy, but I have only been able to show this in three cases. Each of these cases implies that the difference between v's loss to playing r' and r with $r' > r$ is single peaked in v (or flat if $\partial^2\delta_1/\partial r\partial v = 0$). That is, there is a $\widetilde{v}$ such that $\partial[L(r', \alpha, v) - L(r, \alpha, v)]/\partial v \geq 0$ for $v < \widetilde{v}$ and $\partial[L(r', \alpha, v) - L(r, \alpha, v)]/\partial v \leq 0$ for $v < \widetilde{v}$. This means that the gains from playing r' rather than r increase in the distance between v and $\widetilde{v}$, and this ensures that even partial separation breaks down. The three cases are (i) $\partial^2\delta_1/\partial r\partial v = 0$, (ii) $\partial^2\delta_1/\partial v^2 \leq 0$ and $\partial(\partial\delta_1^2/\partial v^2)/\partial r \geq 0$, and (iii) $\partial[(\partial\delta_1(r', v)/\partial v) / (\partial\delta_1(r, v)/\partial v)]/\partial v > 0$ when $r' > r$.

[13] The proof of Proposition 4 shows that the defender's equilibrium allocation always lies in this range.

[14] Existence is assured because $\int \delta_1(\widetilde{\rho}(\underline{v}), v)dG > \delta_1(\widetilde{\rho}(\underline{v}), \underline{v}) = A\delta_2(R - \widetilde{\rho}(\underline{v}))$ and $\int \delta_1(\widetilde{\rho}(\overline{v}), v)dG < \delta_1(\widetilde{\rho}(\overline{v}), \overline{v}) = A\delta_2(R - \widetilde{\rho}(\overline{v}))$.

[15] The distribution H first-order stochastically dominates G if for any nondecreasing function u, $\int u(v)dH \geq \int u(v)dG$. In the specific case in which $u(v) = v$, $\int vdH \geq \int vdG$ which shows that the average value of v under H is at least as large as under G. See Mas-Colell, Winston, and Green 1995, 195 on stochastic dominance.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

unique r, namely, r_h, at which $\int \delta_1(r,v)dH = A\delta_2(R-r)$ must be greater than r^* if the previous inequality is strict and $r_h = r^*$ if it holds with equality.

To demonstrate that $\alpha(r_h) \geq \alpha(r^*)$, let $v^+ \equiv \widetilde{\rho}^{-1}(r_h)$. Then $v^+ \geq v^*$ and it suffices to show that $\Phi(v)$ is increasing, since $\alpha(r_h) = \Phi(v^+)$ and $\alpha(r^*) = \Phi(v^*)$. Inspection of the expression for $\Phi(v)$ shows that it is increasing if $\partial\delta_1(\widetilde{\rho}(v), v)/\partial r[A\delta_2'(R-\widetilde{\rho}(v))]^{-1}$ is decreasing. To show that this expression is decreasing, implicitly differentiate $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R-\widetilde{\rho}(v))$ to obtain:

$$-A\delta_2'(R-\widetilde{\rho}(v))\frac{d\widetilde{\rho}(v)}{dv} = \frac{\partial\delta_1(\widetilde{\rho}(v), v)}{\partial r}\frac{d\widetilde{\rho}(v)}{dv} + \frac{\partial\delta_1(\widetilde{\rho}(v), v)}{\partial v}$$

$$\frac{d\widetilde{\rho}(v)}{dv} = -\frac{\partial\delta_1(\widetilde{\rho}(v), v)/\partial v}{\partial\delta_1(\widetilde{\rho}(v), v)/\partial r + A\delta_2'(R-\widetilde{\rho}(v))}$$

Thus, $d\widetilde{\rho}(v)/dv > 0$. Now differentiate $\partial\delta_1(\widetilde{\rho}(v), v)/\partial r\,[A\delta_2'(R-\widetilde{\rho}(v))]^{-1}$.

$$\frac{\partial}{\partial v}\left(\frac{\partial\delta_1(\widetilde{\rho}(v), v)/\partial r}{A\delta_2'(R-\widetilde{\rho}(v))}\right) = \left(\frac{\partial\delta_1^2(\widetilde{\rho}(v), v)}{\partial r^2}\frac{d\widetilde{\rho}(v)}{dv} + \frac{\partial\delta_1^2(\widetilde{\rho}(v), v)}{\partial r\partial v}\right) \times \left[A\delta_2'(R-\widetilde{\rho}(v))\right]^{-1} + \frac{\partial\delta_1(\widetilde{\rho}(v), v)}{\partial r} \times \left[A\delta_2'(R-\widetilde{\rho}(v))\right]^{-2} A\delta_2''(R-\widetilde{\rho}(v))\frac{d\widetilde{\rho}(v)}{dv}.$$

The expression on the right is clearly negative, since $\partial^2\delta_1/\partial r\partial v \geq 0$. Hence, $d\alpha(\widetilde{\rho}(v))/dv = \Phi'(v) > 0$. ■

Propositions 1 and 4 combine to show that uncertainty makes the defender better off. The former shows that v allocates $\widetilde{\rho}(v)$ to site 1 and suffers an expected loss of $A\delta_2(R-\widetilde{\rho}(v))$ when there is no uncertainty about v. When there is uncertainty, v allocates $\widetilde{\rho}(v^*)$ to this site, and it is easy to show that this allocation results in an expected loss strictly less than $A\delta_2(R-\widetilde{\rho}(v))$ as long as $v \neq v^*$.

Corollary 3. Uncertainty about v makes the defender better off: For all $v \neq v^*$, v's loss to allocating $\widetilde{\rho}(v^*)$ to site 1 is strictly less than $A\delta_2(R-\widetilde{\rho}(v))$ which is what v would obtain if there were no uncertainty.

Proof. Proposition 1 shows that v allocates $\widetilde{\rho}(v)$ to site 1, and this results in an expected loss of $A\delta_2(R-\widetilde{\rho}(v))$ when there is complete information. In the asymmetric information game, v allocates $\widetilde{\rho}(v^*)$ and does strictly better at this allocation if $A\delta_2(R-\widetilde{\rho}(v)) > \Phi(v^*)\delta_1(\widetilde{\rho}(v^*), v) + (1-\Phi(v^*))A\delta_2(R-\widetilde{\rho}(v^*))$ where $\alpha(\widetilde{\rho}(v^*)) = \Phi(v^*)$ by Proposition 4.

Define $\Delta(v|v^*)$ to be the difference between v's loss to playing $\widetilde{\rho}(v)$ and $\widetilde{\rho}(v^*)$: $\Delta(v|v^*) \equiv A\delta_2(R-\widetilde{\rho}(v)) - [\Phi(v^*)\delta_1(\widetilde{\rho}(v^*), v) + (1-\Phi(v^*))A\delta_2(R-\widetilde{\rho}(v^*))]$. Then, $\Delta(v^*|v^*) = 0$ and it suffices to show that $\partial\Delta(v|v^*)/\partial v > 0$ for $v > v^*$ and $\partial\Delta(v|v^*)/\partial v < 0$ for $v < v^*$. Suppose $v > v^*$. Differentiation gives

$$\frac{\partial\Delta(v|v^*)}{\partial v} = -A\delta_2'(R-\widetilde{\rho}(v))\frac{d\widetilde{\rho}(v)}{dv} - \Phi(v^*)\frac{\partial\delta_1(\widetilde{\rho}(v^*), v)}{\partial v}$$
$$= \Phi(v)\frac{\partial\delta_1(\widetilde{\rho}(v), v)}{\partial v} - \Phi(v^*)\frac{\partial\delta_1(\widetilde{\rho}(v^*), v)}{\partial v}$$
$$= \Phi(v)\left[\frac{\partial\delta_1(\widetilde{\rho}(v), v)}{\partial v} - \frac{\partial\delta_1(\widetilde{\rho}(v^*), v)}{\partial v}\right] + \frac{\partial\delta_1(\widetilde{\rho}(v^*), v)}{\partial v}\left[\Phi(v) - \Phi(v^*)\right].$$

The proof of Corollary 2 shows that $\Phi(v)$ is increasing. That $v > v^*$ and $\partial^2\delta_1/\partial r\partial v \geq 0$ guarantees that the factor in brackets is nonnegative. Hence, the expression in the last line is positive and $\partial\Delta(v|v^*)/\partial v > 0$ for $v > v^*$. An analogous argument shows that $\partial\Delta(v|v^*)/\partial v < 0$ for $v < v^*$. ■

WHEN MORE VULNERABLE SITES ARE EASIER TO PROTECT ON THE MARGIN

Propositions 2, 3, and 4 show that if more vulnerable sites are harder to protect on the margin, secrecy concerns generally swamp vulnerability concerns. This section shows by example that separating equilibria sometimes exist if more vulnerable sites are easier to protect on the margin. The discussion also explains why this result depends on the sign of the $\partial^2\delta_1/\partial r\partial v$.

As Lemma 2 shows, v must offer $\widetilde{\rho}(v)$ in a separating equilibrium. Moreover, no $v \neq \widehat{v}$ can profitably deviate from $\widetilde{\rho}(v)$ to $\widetilde{\rho}(\widehat{v})$, and, as noted above, two conditions must be satisfied to ensure that no type can profitably deviate in this way. The first is that the "local" incentive compatibility constraints must hold. That is, types arbitrarily close to $\widehat{v}$ must be deterred from deviating. Second, the "global" incentive compatibility constraints must also be satisfied; that is, types far away from $\widehat{v}$ must be deterred. Deterring those close by pins down a unique probability of attacking. Whether this probability also deters types far away from $\widehat{v}$ from wanting to deviate to $\rho(\widehat{v})$ is related to the sign of $\partial^2\delta_1/\partial r\partial v$.

Suppose more formally that all $v \in (v', v'')$ separate with $\rho(v) = \widetilde{\rho}(v)$ for all types in this interval. If $w'' > \widehat{v}$, then the fact that w'' cannot profitably deviate from $\widetilde{\rho}(w'')$ to $\widetilde{\rho}(\widehat{v})$ means:

$$A\delta_2(R-\widetilde{\rho}(w'')) \leq \widehat{\alpha}\delta_1(\widetilde{\rho}(\widehat{v}), w'') + (1-\widehat{\alpha})A\delta_2(R-\widetilde{\rho}(\widehat{v}))$$
$$\widehat{\alpha} \geq \frac{A\delta_2(R-\widetilde{\rho}(w'')) - A\delta_2(R-\widetilde{\rho}(\widehat{v}))}{w''-\widehat{v}} \times \frac{w''-\widehat{v}}{\delta_1(\widetilde{\rho}(\widehat{v}), w'') - \delta_1(\widetilde{\rho}(\widehat{v}), \widehat{v})},$$

where the second line follows from the first because $A\delta_2(R-\widetilde{\rho}(\widehat{v})) = \delta_1(\widetilde{\rho}(\widehat{v}), \widehat{v})$. Taking the limit as w'' goes to $\widehat{v}$ gives:

$$\widehat{\alpha} \geq \frac{-A\delta_2(R-\widetilde{\rho}(\widehat{v}))d\widetilde{\rho}(\widehat{v})/dv}{\partial\delta_1(\widetilde{\rho}(\widehat{v}), \widehat{v})/\partial v}.$$

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

A similar argument for $w' < \widehat{v}$ shows that $\widehat{\alpha}$ must be less than or equal to the expression on the right of the previous inequality. Hence, $\widehat{\alpha} = -[A\delta_2(R-\widetilde{\rho}(\widehat{v}))d\widetilde{\rho}(\widehat{v})/dv]\ [\partial\delta_1(\widetilde{\rho}(\widehat{v}),\widehat{v})/\partial v]^{-1}$. Differentiating $\delta_1(\widetilde{\rho}(v), v) = A\delta_2(R-\widetilde{\rho}(v))$ implicitly to obtain $d\widetilde{\rho}(v)/dv$ (see the proof of Corollary 2 for this expression) and then substituting this back into the expression for $\widehat{\alpha}$ yields $\widehat{\alpha} = \alpha(\widetilde{\rho}(\widehat{v})) = \Phi(\widehat{v})$. Thus, deterring w close to $\widehat{v}$ from wanting to deviate from $\widetilde{\rho}(w)$ to $\widetilde{\rho}(\widehat{v})$ pins down a unique $\alpha(\widetilde{\rho}(\widehat{v}))$.

Now consider w's payoff to deviating if w is farther away from $\widehat{v}$. Type w prefers deviating from $\widetilde{\rho}(w)$ to $\widetilde{\rho}(\widehat{v})$ if $\Delta(w|\widehat{v}) = A\delta_2(R-\widetilde{\rho}(w)) - [\Phi(\widehat{v})\delta_1(\widetilde{\rho}(\widehat{v}), w) + (1-\Phi(\widehat{v}))A\delta_2(R-\widetilde{\rho}(\widehat{v}^*))] > 0$. (Recall that the defender's payoffs are losses so lower is better.) Since $\Delta(\widehat{v}|\widehat{v}) = 0$, all $w > \widehat{v}$ prefer to deviate if $\partial\Delta(w|\widehat{v})/\partial w > 0$ for all $w > \widehat{v}$. Indeed, the gains from deviation grow as $w - \widehat{v}$ does.

As shown in the proof of Corollary 3,

$$\frac{\partial\Delta(w|\widehat{v})}{\partial w} = \Phi(w)\left[\frac{\partial\delta_1(\widetilde{\rho}(w), v)}{\partial v} - \frac{\partial\delta_1(\widetilde{\rho}(\widehat{v}), v)}{\partial v}\right] + \frac{\partial\delta_1(\widetilde{\rho}(\widehat{v}), v)}{\partial v}\left[\Phi(w) - \Phi(\widehat{v})\right].$$

The function $\Phi(w)$ is strictly increasing, so the second term is positive. The assumption $\partial^2\delta_1/\partial r\partial v \geq 0$ also ensures that the first term is nonnegative and consequently that $\partial\Delta/\partial w > 0$ for $w > \widehat{v}$.

Similarly, if $\partial\Delta/\partial w < 0$ for all $w < \widehat{v}$, these w prefer to deviate with the gains from doing so growing as w moves away from $\widehat{v}$. And, the expression for $\partial\Delta/\partial w$ shows that it is negative for $w < \widehat{v}$ if $\partial^2\delta_1/\partial r\partial v \geq 0$.

Hence $\partial^2\delta_1/\partial r\partial v \geq 0$ is sufficient to ensure that separation breaks down because satisfying the local incentive compatibility constraints means violating the global incentive compatibility constraints. If, alternatively, $\partial^2\delta_1/\partial r\partial v < 0$, then the sign of $\partial\Delta/\partial w$ may be reversed in which case the gains from deviation decrease as the distance between w and $\widehat{v}$ increases. Separation in these circumstances will be possible.

An example establishes that a separating equilibrium sometimes exists if $\partial^2\delta_1/\partial r\partial v < 0$.[16] Assume that the value of the sites is the same ($A = 1$), $\delta_2(r) = \delta(r)$, and $\delta_1(r, v) = \delta(r - v)$, where $\delta' < 0$ and $\delta'' > 0$. More vulnerable sites are easier to protect on the margin in this example: $\partial^2\delta_1/\partial r\partial v = -\delta''(r-v) < 0$.

Now consider the assessment $\{\rho(v), \alpha(r), \mu(v|r)\}$ where the attacker's strategy is $\alpha(r) = 1$ if $r < \rho(\underline{v})$; $\alpha(r) = \frac{1}{2}$ if $r \in [\rho(\underline{v}), \rho(\overline{v})]$; and $\alpha(r) = 0$ if $r > \rho(\overline{v})$. The defender's separating strategy $\rho(v)$ is the allocation at which $\delta'(r-v) = \delta'(R-r)$ which implies $\rho(v) = \frac{1}{2}(v+R)$. The attacker's beliefs put probability one on $\underline{v}$ if $r < \rho(\underline{v}) = \frac{1}{2}(\underline{v}+R)$, probability one on $v = 2r - R$ for $r \in [\rho(\underline{v}), \rho(\overline{v})]$; and probability one on $\overline{v}$ for $r > \rho(\overline{v})$.

The attacker is indifferent between attacking sites 1 and 2 following any equilibrium offer r given its beliefs. Hence $\alpha(r) = \frac{1}{2}$ is a best reply. The attacker's beliefs are also consistent with Bayes' Law (and the support restriction in footnote 11). And by construction, $\rho(r)$ solves the first-order condition $\partial L(\alpha, r)/\partial r = \frac{1}{2}\delta'(r-v) - \frac{1}{2}\delta'(R-r) = 0$. Convexity assures that this allocation minimizes the defender's loss. It is also clear that v cannot profitably deviate to any $r \notin [\rho(\underline{v}), \rho(\overline{v})]$.Hence $\{\rho(v), \alpha(r), \mu\}$ is a separating equilibrium.

CONCLUSION

Many allocation problems are sequential. Defenders, for example, generally have to commit some or all of their resources before an attacker decides where to strike. Adversaries are also often likely to have private information about the effectiveness of their allocations. In a sequential setting, this private information creates a dilemma.

In particular, a defender must be concerned that allocating more to a highly vulnerable site in order to reduce the expected losses should that site be attacked may also signal to the attacker that the site is more vulnerable and thereby increase the probability of an attack on that site.

Secrecy concerns turn out to swamp vulnerability concerns when more vulnerable sites are also harder to protect on the margin. The defender pools on a common allocation to avoid signaling anything to the attacker. This allocation is optimal against the average level of vulnerability but devotes too little to that site if it is very vulnerable and too much to that site if it is not very vulnerable. Even so, uncertainty makes the defender better off. When more vulnerable sites are easier to protect on the margin, vulnerability concerns may swamp secrecy concerns and separating equilibria sometimes exist.

APPENDIX

This section proves Propositions 2, 3, and 4. The argument begins with Lemma 1A, which demonstrates that if higher vulnerability types allocate more to site 1, then the equilibrium probability of an attack must be weakly decreasing in the vulnerability.

Lemma 1A. *Let $\{\rho(v), \alpha(r), \mu\}$ be a PBE of the signaling game. If $v < v'$ and $\rho(v) \leq \rho(v')$, then $\alpha(\rho(v)) \geq \alpha(\rho(v'))$.*

Proof. Suppose $v < v'$, $r \equiv \rho(v) \leq \rho(v') \equiv r'$, $\alpha \equiv \alpha(\rho(v))$, and $\alpha' \equiv \alpha(\rho(v'))$. Incentive compatibility requires:

$$\begin{aligned}\alpha\delta_1(r, v) + (1-\alpha)A\delta_2(R-r) &\leq \alpha'\delta_1(r', v) + (1-\alpha')A\delta_2(R-r')\\ \alpha'\delta_1(r', v') + (1-\alpha')A\delta_2(R-r') &\leq \alpha\delta_1(r, v') + (1-\alpha)A\delta_2(R-r).\end{aligned}$$

[16] I am grateful to Scott Ashworth for very helpful comments leading to this example.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

Subtracting the first inequality from the second yields

$$\begin{aligned}\alpha'[\delta_1(r',v') - \delta_1(r',v)] &\leq \alpha[\delta_1(r,v') - \delta_1(r,v)]\\ (\alpha'-\alpha)[\delta_1(r',v') - \delta_1(r',v)] &\leq \alpha[(\delta_1(r,v') - \delta_1(r,v))\\ &\quad - (\delta_1(r',v') - \delta_1(r',v))].\end{aligned} \tag{A1}$$

The factor in brackets on the left is positive because δ_1 is increasing in v, and the factor in brackets on the right is weakly negative because $\partial^2\delta_1/\partial v\partial r \geq 0$. Hence, $\alpha' - \alpha \leq 0$. ■

Proposition 2. *No nonempty open interval of types separate.*

Proof. Arguing by contradiction, assume there is a PBE $\{\rho(v), \alpha(r), \mu\}$ and an open interval $(\underline{z}, \overline{z})$ such that all $z \in (\underline{z}, \overline{z})$ separate. Lemma 2 implies $\rho(v) = \widetilde{\rho}(z)$ for all $z \in (\underline{z}, \overline{z})$, and it suffices to show that that some $z \in (\underline{z}, \overline{z})$ can profitably deviate from $\widetilde{\rho}(z)$.
Let $z' > z$ and $\alpha \equiv \alpha(\widetilde{\rho}(z))$. That z' cannot profitably mimic z means:

$$\begin{aligned}A\delta_2(R - \widetilde{\rho}(z')) &\leq \alpha\delta_1(\widetilde{\rho}(z), z') + (1-\alpha)A\delta_2(R - \widetilde{\rho}(z))\\ A\delta_2(R - \widetilde{\rho}(z')) &\leq \alpha[\delta_1(\widetilde{\rho}(z), z') - \delta_1(\widetilde{\rho}(z), z)]\\ &\quad + A\delta_2(R - \widetilde{\rho}(z)).\end{aligned}$$

where the second line follows from the first because $A\delta_2(R - \widetilde{\rho}(z)) = \delta_1(\widetilde{\rho}(z), z)$. Rewriting this expression and taking the limit as z' goes to z gives:

$$\begin{aligned}\alpha &\geq \frac{A\delta_2(R - \widetilde{\rho}(z')) - A\delta_2(R - \widetilde{\rho}(z))}{z' - z}\\ &\quad \times \frac{z' - z}{\delta_1(\widetilde{\rho}(z), z') - \delta_1(\widetilde{\rho}(z), z)}\\ \alpha(\widetilde{\rho}(z)) &\geq \frac{-A\delta_2'(R - \widetilde{\rho}(z))d\widetilde{\rho}(z)/dz}{\partial\delta_1(\widetilde{\rho}(z), z)/\partial z}.\end{aligned}$$

Differentiating $A\delta_2(R - \widetilde{\rho}(z)) = \delta_1(\widetilde{\rho}(z), z)$ implicitly to obtain $d\widetilde{\rho}(z)/dz$ and substituting this expression into the previous inequality gives $\alpha(\widetilde{\rho}(z)) \geq \Phi(z)$. (See the proof of Corollary 2 for the expression for $d\widetilde{\rho}(z)/dz$.) A similar argument for $z' < z$ yields $\alpha(\widetilde{\rho}(z)) \leq \Phi(z)$ and hence that $\alpha(\widetilde{\rho}(z)) = \Phi(z)$.
A contradiction follows immediately. The defender's offers $\widetilde{\rho}(z)$ are increasing in z, so Lemma 1A implies $\alpha(\widetilde{\rho}(z))$ must be weakly decreasing in z. But differentiation shows that $\Phi(z)$ is increasing (see the proof of Corollary 2). ■

A PBE is semiseparating if at least two types allocate the same amount to a site and there are at least two distinct allocations. Proposition 3 shows that there are no semiseparating PBEs using simple strategies.

Proposition 3. *The signaling game has no semiseparating PBEs in simple strategies.*

Proof. Suppose that $\{\rho(v), \alpha(r), \mu(v|r)\}$ is a PBE and that ρ is simple. Hence, two types, v' and v'' with $v' < v''$, pool on an allocation, say $\widehat{r}$. Arguing by contradiction, Proposition 2 ensures that this equilibrium is not fully separating. It therefore suffices to show that all $v \in [\underline{v}, \overline{v}]$ pool on $\widehat{r}$ and consequently that there is complete pooling.
Let $\underline{w} = \inf\{v : \rho(v) = \widehat{r}\}$ and $\overline{w} = \sup\{v : \rho(v) = \widehat{r}\}$, where $\underline{w} \leq v' < v'' \leq \overline{w}$. Because ρ is simple, $\rho(v) = \widehat{r}$ for all $v \in (\underline{w}, \overline{w})$. Now define $\widehat{w} = \widetilde{\rho}^{-1}(\widehat{r})$, which implies $\delta_1(\widehat{r}, \widehat{w}) = A\delta_2(R - \widehat{r})$. It follows that $\widehat{w} \in (\underline{w}, \overline{w})$.

Suppose that $\widehat{w} \geq \overline{w}$. Then $A\delta_2(R - \widehat{r}) = \delta_1(\widehat{r}, \widehat{w}) > \int \delta_1(r, v)d\mu(v|\widehat{r})$. The inequality implies that the attacker strictly prefers going after site 2. Hence, $\alpha(\widehat{r}) = 0$ and w's payoff is $A\delta_2(R - \widehat{r})$. But w can always obtain $A\delta_2(R - \widetilde{\rho}(w))$ by playing $\widetilde{\rho}(w)$. Moreover, $\widetilde{\rho}(w) < \widetilde{\rho}(\widehat{w}) = \widehat{r}$ for all $w \in (\underline{w}, \overline{w})$. This means $A\delta_2(R - \widetilde{\rho}(w)) < A\delta_2(R - \widehat{r})$, so all $w \in (\underline{w}, \overline{w})$ strictly prefer to deviate from $\widehat{r}$. This contradiction ensures $\widehat{w} < \overline{w}$.
Now assume $\widehat{w} \leq \underline{w}$. Hence $A\delta_2(R - \widehat{r}) = \delta_1(\widehat{r}, \widehat{w}) < \int \delta_1(r, v)d\mu(v|\widehat{r})$. The attacker therefore strikes site 1 after observing $\widehat{r}$: $\alpha(\widehat{r}) = 1$. As a result, $w \in (\underline{w}, \overline{w})$ obtains an equilibrium payoff of $\delta_1(\widehat{r}, w)$. That $\widehat{w} \leq \underline{w}$ implies $\widehat{r} = \widetilde{\rho}(\widehat{w}) < \widetilde{\rho}(w)$ for all $w \in (\underline{w}, \overline{w})$. But all $w \in (\underline{w}, \overline{w})$ can also play $\widetilde{\rho}(w)$ which yields a loss of $\delta_1(\widetilde{\rho}(w), w)$. This loss is strictly less than $\delta_1(\widehat{r}, w)$ and consequently a profitable deviation. This contradiction means $\widehat{w} > \underline{w}$ and leaves $\widehat{w} \in (\underline{w}, \overline{w})$.
That $\widehat{w} \in (\underline{w}, \overline{w})$ implies $\alpha(\rho(\widehat{w})) = \Phi(\widehat{w})$. To establish this, let $\widehat{\alpha} \equiv \alpha(\rho(\widehat{w}))$ and $w'' > \widehat{w}$. That w'' cannot profitably deviate from $\widetilde{\rho}(w'')$ to $\rho(\widehat{w})$ yields:

$$\begin{aligned}A\delta_2(R - \widetilde{\rho}(w'')) &\leq \widehat{\alpha}\delta_1(\widetilde{\rho}(\widehat{w}), w'') + (1-\widehat{\alpha})A\delta_2(R - \widetilde{\rho}(\widehat{w}))\\ \widehat{\alpha} &\geq \frac{A\delta_2(R - \widetilde{\rho}(w'')) - A\delta_2(R - \widetilde{\rho}(\widehat{w}))}{w'' - \widehat{w}}\\ &\quad \times \frac{w'' - \widehat{w}}{\delta_1(\widetilde{\rho}(\widehat{w}), w'') - \delta_1(\widetilde{\rho}(\widehat{w}), \widehat{w})},\end{aligned}$$

where the second line follows from the first because $A\delta_2(R - \widetilde{\rho}(\widehat{w})) = \delta_1(\widetilde{\rho}(\widehat{w}), \widehat{w})$. Letting $w'' \to \widehat{w}$ gives $\widehat{\alpha} \geq \Phi(\widehat{w})$. A similar argument for $w' < \widehat{w}$ shows that $\widehat{\alpha} \leq \Phi(\widehat{w})$ and consequently $\widehat{\alpha} = \alpha(\rho(\widehat{w})) = \Phi(\widehat{w})$. Furthermore, the attacker must be mixing over sites 1 and 2 because $\Phi(\widehat{w}) \in (0, 1)$.
The defender cannot pool on any other allocation. Suppose that it did; that is, there exists an r^* and two types $z' \neq z''$, such that $\rho(z') = \rho(z'') = r^*$. Define $\underline{z} = \inf\{v : \rho(v) = r^*\}$ and $\overline{z} = \sup\{v : \rho(v) = r^*\}$, where $\underline{z} \leq z' < z'' \leq \overline{z}$. Also define $z^* = \widetilde{\rho}^{-1}(r^*)$. Repeating the argument above establishes that $z^* \in (\underline{z}, \overline{z})$ and $\alpha(\rho(z^*)) = \Phi(z^*)$. Assume without loss of generality that $\widehat{w} < z^*$. Then $\widehat{r} = \widetilde{\rho}(\widehat{w}) < \widetilde{\rho}(z^*) = r^*$, so Lemma 1A ensures $\alpha(\rho(\widehat{w})) \geq \alpha(\rho(z^*))$. But this is a contradiction, since $\Phi(w)$ is increasing in w as shown in the proof of Corollary 2. This contradiction establishes that the defender cannot pool on any $r \neq \widehat{r}$.
An immediate consequence is that $\underline{w} = \underline{v}$ and $\overline{w} = \overline{v}$. Suppose $\underline{w} > \underline{v}$. As just shown, $w \in (\underline{v}, \underline{w})$ cannot pool and therefore must separate. But this contradicts Proposition 2 and leaves $\underline{w} = \underline{v}$. A similar argument yields $\overline{w} = \overline{v}$.
Finally, suppose $\rho(\underline{v}) \neq \widehat{r}$. Then $\underline{v}$ must separate and therefore allocate $\widetilde{\rho}(\underline{v})$. Further, $(\underline{w}, \overline{w}) = (\underline{v}, \overline{v})$ implies that the attacker's posterior $\mu(v|\widehat{r})$ is just its prior G. This and the fact that the attacker mixes leaves $\int \delta_1(\widetilde{\rho}(\widehat{w}), v)dG = A\delta_2(R - \widetilde{\rho}(\widehat{w}))$. But v^* uniquely satisfies $\int \delta_1(\widetilde{\rho}(\widehat{w}), v)dG = A\delta_2(R - \widetilde{\rho}(v^*))$. Hence, $\widehat{w} = v^*$. A contradiction results if $\underline{v}$ strictly prefers $\widehat{r}$ to $\widetilde{\rho}(\underline{v})$.
Type $\underline{v}$ does prefer $\widehat{r}$ to $\widetilde{\rho}(\underline{v})$ if $\Delta(v|v^*) > 0$, where recall $\Delta(v|v^*)$ is the loss to playing $\widetilde{\rho}(v)$ less the loss to playing $\widetilde{\rho}(v^*)$: $\Delta(v|v^*) = A\delta_2(R - \widetilde{\rho}(v)) - [\Phi(v^*)\delta_1(\widetilde{\rho}(v^*), v) + (1 - \Phi(v^*))A\delta_2(R - \widetilde{\rho}(v^*))] > 0$. As shown in the proof of Corollary 3, $\Delta(v^*|v^*) = 0$, $\partial\Delta(v|v^*) > 0$ for $v > v^*$, and $\partial\Delta(v|v^*) < 0$ for $v < v^*$ if $\partial^2\delta_1(r, v)/\partial r\partial v \geq 0$. Thus, $\underline{v}$ prefers to deviate, and this contradiction leaves $\rho(\underline{v}) = \widehat{r}$. The analogous argument demonstrates that $\rho(\overline{v}) = \widehat{r}$. Hence all v pool on $\widehat{r}$. ■

It remains to be shown that pooling equilibria actually exist and that the equilibrium paths in all pooling equilibria are identical.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

Proposition 4. *Pooling perfect Bayesian equilibria exist, and all of them share the same equilibrium path. All v pool on $\widetilde{\rho}(v^*)$ and the attacker strikes site 1 with probability $\alpha(\widetilde{\rho}(v^*)) = \Phi(v^*)$.*

Proof. Suppose the defender pools on $\widehat{r}$. Then $\alpha(\widehat{r}) \in (0, 1)$. If $\alpha(\widehat{r}) = 1$, the attacker must weakly prefer attacking site 1 to striking site 2. Moreover, $\alpha(\widehat{r}) = 1$ gives v an equilibrium payoff of $\delta_1(\widehat{r}, v)$. But no v such that $\widetilde{\rho}(v) > \widehat{r}$ would play $\widehat{r}$ as it could profitably deviate to $\widetilde{\rho}(v)$ since $\delta_1(\widetilde{\rho}(v), v) < \delta_1(\widehat{r}, v)$. Hence, $\widetilde{\rho}(v) \leq \widehat{r}$ for all v which means $\widetilde{\rho}(\overline{v}) \leq \widehat{r}$. This, however, leads to the contradiction as the attacker strictly prefers hitting site 2:

$$\int \delta_1(\widehat{r}, v)dG < \int \delta_1(\widetilde{\rho}(v), v)dG$$

$$= \int A\delta_2(R - \widetilde{\rho}(v))dG < A\delta_2(R - \widehat{r}).$$

This contradiction leaves $\alpha(\widehat{r}) < 1$. An analogous argument shows $\alpha(\widehat{r}) > 0$.

Because the attacker mixes in response to $\widehat{r}$, it must be indifferent to attacking either site. This means $\int \delta_1(\widehat{r}, v)dG = \int A\delta_2(R - \widehat{r})dG = A\delta_2(R - \widehat{r})$. Recall further that this condition pins down a unique $\widehat{r}$ and the corresponding $v^* = \widetilde{\rho}^{-1}(\widehat{r})$. That $\alpha(\widehat{r}) \in (0, 1)$ also implies $v^* \in (\underline{v}, \overline{v})$.

Let $\widehat{\alpha} \equiv \alpha(\widehat{r})$. Then $\widehat{\alpha}$ must equal $\Phi(v^*)$ in order to deter the v in a neighborhood of v^* from deviating to $\widetilde{\rho}(v)$. Hence, the defender must allocate $\widetilde{\rho}(v^*)$ to site 1 and the attacker must strike this site with probability $\Phi(v^*)$ in any pooling PBE.

All that remains to be shown is that a PBE actually exists. Define the strategy $\alpha(r) = 1$ if $r < \widetilde{\rho}(\underline{v})$; $\alpha(r) = \Phi(\widetilde{\rho}^{-1}(r))$ if $r \in [\widetilde{\rho}(\underline{v}), \widetilde{\rho}(\overline{v})]$; and $\alpha(r) = 0$ if $r > \widetilde{\rho}(\overline{v})$. As for the attacker's beliefs, $\mu(v|\widehat{r})$ is the prior G. As for any $r \neq \widehat{r}$, $\mu(v|r)$ puts probability one on $\widetilde{\rho}^{-1}(r)$ for any $r \in [\widetilde{\rho}(\underline{v}), \widetilde{\rho}(\overline{v})]$, probability one on $\underline{v}$ if $r < \widetilde{\rho}(\underline{v})$, and probability one on $\overline{v}$ if $r > \widetilde{\rho}(\overline{v})$. Then, $\{\widehat{r}, \alpha(r), \mu\}$ is a PBE.

To see that this is so, observe first that $\alpha(r)$ is a best response given $\mu(v|r)$. The attacker strictly prefers going after site 1 regardless of its beliefs if $r < \widetilde{\rho}(\underline{v})$. So $\alpha(r) = 1$ is the unique best reply to $r < \widetilde{\rho}(\underline{v})$. Similarly, the attacker strictly prefers hitting site 2 if $r > \widetilde{\rho}(\overline{v})$, and α again has the attacker playing its unique best response to r when $r > \widetilde{\rho}(\overline{v})$. If $r \in [\widetilde{\rho}(\underline{v}), \widetilde{\rho}(\overline{v})]$, then the attacker is indifferent between attacking site 1 and attacking site 2. Hence any $\alpha(r)$ is a best response, and in particular $\alpha(r) = \Phi(\widetilde{\rho}^{-1}(r))$ is.

The beliefs $\mu(v|r)$ are also clearly consistent with Bayes' law. If all v pool on $\rho(v) = \widehat{r}$, then the only restriction that Bayes' law imposes is that the attacker's posterior beliefs following $\widehat{r}$ are its prior beliefs G which $\mu(v|r)$ satisfies.

Therefore the assessment $\{\widehat{r}, \alpha(r), \mu\}$ will be a PBE if $\widehat{r}$ is a best response to $\alpha(r)$. The argument takes two steps. The first demonstrates that no v can profitably deviate from $\widehat{r}$ to $\widetilde{\rho}(v)$. Recall from the proof of Corollary 3 that the difference between v's payoff to $\widetilde{\rho}(v)$ and $\widetilde{\rho}(v^*)$ is $\Delta(v|v^*) = A\delta_2(R - \widetilde{\rho}(v)) - [\Phi(v^*)\delta_1(\widetilde{\rho}(v^*), v) + (1 - \Phi(v^*))A\delta_2(R - \widetilde{\rho}(v^*))]$ and satisfies $\Delta(v^*|v^*) = 0$, $\partial\Delta(v|v^*)/\partial v \geq 0$ if $v > v^*$, and $\partial\Delta(v|v^*)/\partial v \leq 0$ if $v < v^*$. Hence, no $v \neq v^*$ prefers to deviate from $\widetilde{\rho}(v^*)$ to $\widetilde{\rho}(v)$.

It follows that no v can profitably deviate from $\widehat{r} = \widetilde{\rho}(v^*)$ to any other $\widetilde{\rho}(v') \in [\widetilde{\rho}(\underline{v}), \widetilde{\rho}(\overline{v})]$. Type v's loss to $\widetilde{\rho}(v')$ is $\Pi(v'|v) \equiv \Phi(v')\delta_1(\widetilde{\rho}(v'), v) + (1 - \Phi(v'))A\delta_2((R - \widetilde{\rho}(v'))$. Then v has no profitable deviation from $\widetilde{\rho}(v^*)$ if $\Pi(v'|v) - \Pi(v^*|v) \geq 0$ for all $v' \in [\underline{v}, \overline{v}]$. As will be seen, $\widetilde{\rho}(v)$ minimizes $\Pi(v'|v)$; that is, $\Pi(v'|v) \geq \Pi(v|v)$ for all $v' \in [\underline{v}, \overline{v}]$. Hence $\Pi(v'|v) - \Pi(v^*|v) \geq \Pi(v|v) - \Pi(v^*|v)$ which means that it suffices to show $\Pi(v|v) - \Pi(v^*|v) \geq 0$. But $\Pi(v|v) = \Phi(v)\delta_1(\widetilde{\rho}(v), v) + (1 - \Phi(v))A\delta_2(R - \widetilde{\rho}(v)) = A\delta_2(R - \widetilde{\rho}(v))$. So $\Delta(v|v^*) = \Pi(v|v) - \Pi(v^*|v) = A\delta_2(R - \widetilde{\rho}(v)) - \Pi(v^*|v)$ and, as shown above, $\Delta(v|v^*) \geq 0$.

To see that $\widetilde{\rho}(v)$ minimizes $\Pi(v'|v)$ differentiate $\Pi(v'|v)$:

$$\frac{\partial\Pi(v'|v)}{\partial v'} = \Phi'(v')\left[\delta_1(\widetilde{\rho}(v'), v) - A\delta_2(R - \widetilde{\rho}(v'))\right]$$

$$+ \left[\Phi(v')\frac{\partial\delta_1(\widetilde{\rho}(v'), v)}{\partial r} - (1 - \Phi(v'))A\delta_2'(R - \widetilde{\rho}(v'))\right]$$

$$\times \frac{d\widetilde{\rho}(v')}{dv'}.$$

But, $\Phi(v')\partial\delta_1(\widetilde{\rho}(v'), v')/\partial r - (1 - \Phi(v'))A\delta_2'(R - \widetilde{\rho}(v')) = 0$. Using this to simplify the second term gives

$$\frac{\partial\Pi(v'|v)}{\partial v'} = \Phi'(v')\left[\delta_1(\widetilde{\rho}(v'), v) - A\delta_2(R - \widetilde{\rho}(v'))\right]$$

$$+ \Phi(v')\left[\frac{\partial\delta_1(\widetilde{\rho}(v'), v)}{\partial r} - \frac{\partial\delta_1(\widetilde{\rho}(v'), v')}{\partial r}\right]\frac{d\widetilde{\rho}(v')}{dv'}.$$

Clearly, $\partial\Pi(v'|v)/\partial v' = 0$ at $v = v'$. If $v > v'$, the factors in brackets are positive as are Φ', Φ, and $d\widetilde{\rho}(v')/dv$. Hence, $\partial\Pi(v'|v)/\partial v > 0$ if $v > v'$. Conversely, $\partial\Pi(v'|v)/\partial v < 0$ if $v < v'$. Thus, v minimizes its loss by deviating to $\widetilde{\rho}(v)$ which is not a profitable deviation from $\widehat{r}$.

As for deviating to $r < \widetilde{\rho}(\underline{v})$, v's loss is $\delta_1(r, v)$ which is greater than $\delta_1(\widetilde{\rho}(v), v)$, and so not a profitable deviation from $\widehat{r}$. And, deviating to $r > \widetilde{\rho}(\overline{v})$ is clearly not profitable. Hence, PBEs exist. ■

REFERENCES

Bier, Vicki, Santiago Oliveros, and Larry Samuelson. 2005. "Choosing What to Protect." *Journal of Public Economics* (forthcoming). Available at: http://www.ssc.wisc.edu/~ larrysam/ publications.htm.

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

Borel, Emile. 1921[1953]. "The Theory of Play and Integral Equations with Skew Symmetric Kernels." *Econometrica* 21: 97–100.

Golman, Russell, and Scott E. Page. 2006. "General Blotto: Games of Allocative Strategic Mismatch." Manuscript, Center for Complex Systems, University of Michigan. Accessed November 3, 2006 at http://www.cscs.umich.edu/~spage/papers_files/ blotto.pdf.

Groseclose, Tim, and James M. Snyder. 1996. Buying Supermajorities." *American Political Science Review* 90 (June): 303–15.

Kvasov, Dmitriy. 2007. "Contests with Limited Resources." *Journal of Economic Theory* 136: 738–48.

Laslier, Jean-Francois, and Nathalie Picard. 2002. "Distributive Politics and Electoral Competition." *Journal of Economic Theory* 103: 106–30.

Luce, R. Duncan, and Howard Raiffa. 1957. *Games and Decisions*. New York: Wiley.

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

Morgan, John, and Felix Vardy. 2007. "On the Buyability of Voting Bodies." Haas School of Business, UC Berkeley. Available at: http://faculty.haas.berkeley.edu/rjmorgan/Buyability.pdf. Accessed on May 15, 2007.

Moteff, John. 2006. "Critical Infrastructures: Background, Policy, and Implementation." Washington, DC: Congressional Research Service.

Myerson, Roger. 1993. "Incentives to Cultivate Favored Minorities under Alternative Electoral Systems." *American Political Science Review* 87: 856–69.

Powell, Robert. 2006. "Defending Against Terrorist Attacks with Private Information about Vulnerability." Manuscript, Department of Political Science, UC Berkeley.

Powell, Robert. 2007a. "Allocating Defensive Resources Prior to Attack." Manuscript, Department of Political Science, UC Berkeley.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press

Powell, Robert. 2007b. "Defending Against Terrorist Attacks with Limited Resources." *American Political Science Review* (forthcoming).
Ramey, Gary. 1996. "D1 Signaling Equilibria with Mulitple Signals and a Continuum of Types." *Journal of Economic Theory* 69 (May): 508–531.
Roberson, Brian. 2006. "The Colonel Blotto Game." *Economic Theory* 29: 1–24.
Shubik, Martin. 1982. *Game Theory in the Social Sciences*, vol. 1. Cambridge: MIT Press.
Shubik, Martin and Robert Weber. 1981. "Systems Defense Games." *Naval Research Logistics Quarterly* 28: 281–87.
Snyder, James. 1989. "Election Goals and the Allocation of Campaign Resources." *Econometrica* 57 (May): 637–60.
Tukey, John W. 1949. "A Problem of Strategy." *Econometrica* 17: 73.

https://doi.org/10.1017/S0003055407070530 Published online by Cambridge University Press