Defending Against Terrorist Attacks with Private Information about Vulnerability*

Robert Powell

July 2006

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 and Jacob Shapiro.

Defending Against Terrorist Attacks with Private Information about Vulnerability

Abstract

Following the attacks of September 11, 2001, the United States undertook a massive effort to secure its critical infrastructure. As part of this endeavor the Department of Homeland Security identified 1700 highest priority sites and plans to conduct vulnerability assessments of them. Suppose as a result of these studies the government learns that a site is very vulnerable. What should the government do? On the one hand, the government might be tempted to invest the resources needed to make the site more resistant to attack and thereby reduce the expected losses should that site be attacked. On the other hand, these efforts might signal to strategic terrorists that the site is relatively more vulnerable and thereby make an attack on that site more likely. This paper models this trade off as a signaling game in which a government has to decide how to allocate its resources between two sites and has private information about the vulnerability of one of them. In equilibrium, secrecy concerns swamp the vulnerability issue. Rather than allocating its resources in the way that would be optimal given the level of vulnerability, the government always "pools," i.e., the government allocates its resources in the same way regardless of the level of vulnerability. The terrorists therefore cannot infer anything about the sites' relative vulnerability after observing the government's allocation. This pooling allocation is only optimal with respect to the average level of vulnerability. It under invests in protecting highly vulnerable sites and over invests in protecting low-vulnerability sites. Even so, the government's expected loss is lower when the terrorists are uncertain about the sites' vulnerability than when there is no uncertainty.

Defending Against Terrorist Attacks with Private Information about Vulnerability

The United States undertook a massive effort to secure its critical infrastructure and key assets following the attacks of September 11, 2001. The federal government alone now spends about $20 billion dollars a year on infrastructure protection.[1] As part of this effort, the Department of Homeland Security maintains the National Asset Database of critical resources.[2] By the end of 2005, this database listed approximately 80,000 sites including "nuclear power plants, pipelines, bridges, stadiums, and locations such as Times Square" (GAO 2005, 75). Of these, the Department and identified 1700 as the highest priority and intended "to visit each of these high priority sites to assess their vulnerabilities to various forms of attack..." (Moteff 2006, 38).

Vulnerability studies can create an informational problem if the terrorists are strategic as the *National Strategy for Homeland Security* emphasizes they are. "One fact dominates all homeland security threat assessments, terrorists are strategic actors... Increasing the security of a particular type of target, such as aircraft or buildings, makes it more likely that terrorists will strike a different target. Increasing countermeasures to a particular terrorist tactic such as hijacking, makes it more likely that terrorists will favor a different tactic" (White House 2002; 7, 29).[3]

Suppose then that strategic terrorists are unsure of the vulnerability of some sites but can see how the government allocates its resources to deal with the sites' vulnerabilities. Suppose further that a vulnerability study reveals to the government that a site is very vulnerable. What should the government do? On the one hand, the government might be tempted to invest the resources needed to harden the target and reduce the expected losses should that site be attacked. On the other hand, the government's efforts to protect a site may signal to strategic terrorists that the site is relatively more vulnerable and thereby make an attack on that site more likely.

This paper models the dilemma facing the government as a signaling game in which

[1] See Moteff 2006 for an overview of these programs.
[2] The *National Infrastructure Protection Plan* (2006, 38-39) describes this inventory.
[3] Enders and Sandler (2004) find statistical support for these substitution effects.

the government has to decide how to divide its resources between two sites. The terrorists see this allocation and then decide which site to attack. The more the government spends on a site, the lower its expected losses if that site is attacked. But the government also has private information about the vulnerability of one of the sites and thus must be concerned about what the terrorists will infer about the relative vulnerability of the sites after seeing how the government allocates its resources.

An equilibrium of this game is fully separating if the government makes a different allocation for every different level of vulnerability. The terrorists therefore can infer the level of vulnerability from the government's allocation and then attack the site that offers highest expected payoff. In a sense, the vulnerability issue swamps secrecy concerns in a separating equilibrium in that the government 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 government divides its resources between the two sites in the same way regardless of what it knows about their relative vulnerability. The terrorists now cannot infer anything about the sites' relative vulnerability when they see this allocation. Secrecy concerns in this type of equilibrium swamp the vulnerability issue in that the government's optimal strategy is to reveal nothing about the level vulnerability.

In a semi-separating equilibrium, the government's allocation varies to some degree with the level of vulnerability. The terrorists are able to infer something about the relative vulnerability of the two sites cannot infer the exact level of relative vulnerability (as they do in a separating equilibrium). In a semi-separating equilibrium, neither the vulnerability issue nor the secrecy concerns dominates the other. The government reveals some information about the level of vulnerability in order to address the vulnerability issue, but the government does not reveal the exact level of vulnerability.

The analysis below shows that secrecy concerns swamp the vulnerability issue. In the unique equilibrium outcome, the government pools on the allocation that is optimal with respect to the "average" vulnerability. This allocation spends less on highly vulnerability sites and more on low-vulnerability sites than the government would spend on them if signaling were of no concern. As the terrorist perceive a site to be more vulnerable on average, the government devotes more on protecting that site.

The next section specifies the model and summarizes the equilibria when the level of vulnerability is common knowledge. The subsequent section assumes that the government has private information about the relative vulnerability of the two sites and characterizes the equilibria of that game. A final section considers the case in which the government still has private information but the terrorists are unable to observe the government's allocation. Signaling is not an issue here, but the terrorists uncertainty about the level of vulnerability still affects the government's optimal allocation.

The Complete Information Model and Equilibria

A simplified version of Powell's (2006) resource-allocation game provides a point of departure. As illustrated in Figure 1, the government begins the game by deciding how to divide R resources between the two sites it is trying to protect. That is, the government 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 terrorists attack decide which site to attack.[4]

The more the government allocates to a site, the harder that site becomes and the less likely an attack on that site is to succeed. Formally, the probability that an attack on site 2 succeeds is strictly decreasing in the resources devoted to it but with diminishing marginal effects: $\delta_2' < 0$ and $\delta_2'' > 0$. The probability that an attack on site 1 succeeds depends on the resources allocated to that site and on the inherent vulnerability v of the site. The larger v, the more resources the government has to spend on site 1 in order to reduce the probability of a successful attack to a given level. In particular, the probability of a successful attack is taken to be $\delta_1(r) + v$ where $\delta_1' < 0$ and $\delta_1'' > 0$.[5] Assume for now

[4] The model is of course a simplification. An important challenge facing the Department of Homeland Security is that much of the critical infrastructure in the United States is privately owned and the Department currently lacks regulatory authority over it and cannot directly allocate resources to hardening these sites "inside the fence" (GAO 2005, 84).

[5] Taking the probability of a successful attack to be linear in v simplifies the analysis immensely as discussed in footnote 6.

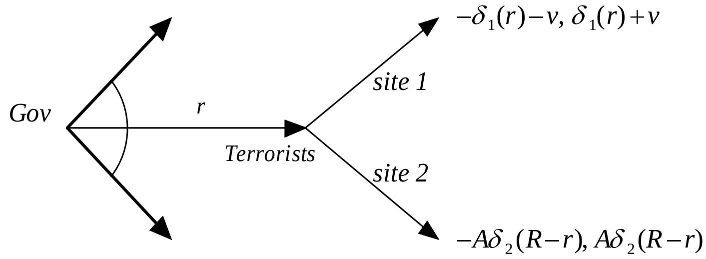

Figure 1: The signaling game.

that v along with all of the other parameters of the model are common knowledge. The government will be assumed below to know v while the terrorists are unsure of v.

The government suffers a loss of one if site 1 is destroyed and a loss of $A > 0$ if site 2 is destroyed. Thus, the government's expected loss if it allocates r to site 1 and the terrorists attack that site with probability α is $L(r, \alpha) = \alpha[\delta_1(r) + v] + (1-\alpha)A\delta_2(R-r)$. As for the terrorists' payoffs, the game is assumed to be zero sum, so the government's losses are the terrorists' gains. The terrorists receive one if they successfully strike site 1 and A if they destroy site 2. If the government allocates r to site 1 and the terrorists attack this site with probability α, the terrorists' expected payoff is $L(r, \alpha)$.[6]

The basic elements of this model are the same as those at the core of the Department of Homeland Security's risk-management approach which centers on vulnerability, threat, and consequences. "The *vulnerability* of an asset is the probability that a particular attempted attack will succeed against a particular target or class of targets" (GAO 2005,

[6] A more general formulation would not assume the probability that an attack on site 1 succeeds is linear in v and would allow this probability to be a function of r and v, $\delta_1(r, v)$, such that $\partial\delta_1/\partial r < 0$, $\partial^2\delta_1/\partial r^2 > 0$, and $\partial\delta_1/\partial v > 0$. The key issue in this more general approach is how the level of vulnerability influences the marginal effects of additional spending on government's expected losses L, i.e., on the sign of $\partial^2 L/\partial v\partial r$ or equivalently on the sign of $\partial^2\delta_1/\partial v\partial r$. The quasilinear specification treats v as a "shift" parameter which has no effect on the marginal returns to allocating more to site 1, $\partial^2\delta_1/\partial v\partial r = 0$. Suppose alternatively that more vulnerable sites were inherently easier to protect at the margin so that an increase in r reduced L by a larger amount if v were higher. This would imply $\partial^2\delta_1/\partial v\partial r < 0$ and would be the case if v entered multiplicatively so $\delta_1(r, v) = v\delta_1(r)$ and $\partial^2 L/\partial v\partial r = \partial^2\delta_1/\partial v\partial r = \delta_1'(r) < 0$. If, conversely, more vulnerable sites are harder to protect at the margin, then $\partial^2\delta_1/\partial v\partial r > 0$. These more general cases are left for future work.

25). This is just what $\delta_1(r_1)+v$ and $\delta_2(r_2)$ represent: the conditional probabilities that an attack on a specific site succeeds given the level of resources devoted to that site's defense. Threat is "the probability that a specific target is attacked in a specific way" (Willis *et. al.* 2005, 8) and is formalized in the probabilities α and $1-\alpha$ in the model. Finally, the expected losses of one if site 1 is successfully attacked and A if site 2 is successfully attack correspond to the consequences in risk-management which are "the expected magnitude of damage (e.g., deaths, injuries, or property damage)" resulting from a successful attack (Willis *et. al.* 2005, 9).[7]

The government has a unique equilibrium allocation. Let $\widetilde{r}(v)$ be the allocation at which the terrorists are indifferent between attacking sites 1 and 2 when the level of vulnerability is v, i.e., $\widetilde{r}(v)$ uniquely satisfies $\delta_1(r) + v = A\delta_2(R - r)$.[8] As Figure 2 illustrates, this allocation minmaxes the terrorists, i.e., $\widetilde{r}(v)$ is the allocation that minimizes the maximum payoff the terrorists can obtain after observing r. Formally, $\widetilde{r}(v)$ solves $\min_{r\in[0,R]} \max\{\delta_1(r) + v, A\delta_2(R - r)\}$, and is what the government does in any subgame perfect equilibrium.

To establish that the only equilibrium allocation is $\widetilde{r}(v)$, it suffices to show that the government can profitably deviate from any $r' \neq \widetilde{r}(v)$. Observe that the government is spending too little on site 1 if $r < \widetilde{r}(v)$ in that the terrorists' payoff to striking site 1 is greater than that of striking site 2: $\delta_1(r) + v > A\delta_2(R - r)$. The terrorists therefore attack site 1 which leaves the government with an expected loss of $\delta_1(r) + v$ whenever

[7] It is important to emphasize that expected utility analysis on which the present analysis is based assumes that the government can rank the sites in terms of which one it would least like to lose and that the terrorists can rank the sites in terms of the site it would most like to destroy. These rankings may or may not correspond to the damage an attack would do measured in terms of lives, injuries or property. For example, terrorists might rank destroying the Statue of Liberty very high although this would actually impose little in direct damages. Efforts to measure consequences in terms of lives, injuries, or dollars should really be seen as ways of trying to assess the actors' preference orderings. Finally, only relative payoffs matter in utility analysis, so the payoff to losing site 1 can be normalized to one with A measuring the loss if site 2 is destroyed relative to the loss if site 1 is destroyed. See Mas-Colell, Whinston, and Green (1995) for a discussion of utility theory.

[8] The expected payoff $\delta_1(r)$ 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.

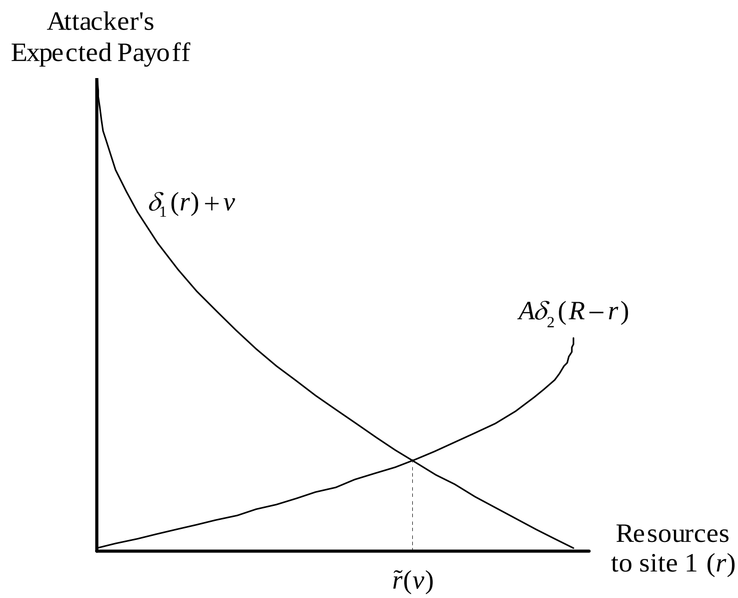

Figure 2: Minmaxing the attacker.

$r < \widetilde{r}(v)$. Hence, the government can profitably deviate to any r' between r and $\widetilde{r}(v)$ as $\delta_1(r) > \delta_1(r')$ for all $r' \in (r, \widetilde{r}(v))$.

Conversely, the government is spending too much on site 1 when $r > \widetilde{r}(v)$ and, as a result, makes site 2 the more attractive target (i.e., $A\delta_2(R-r) > \delta_1(r)+v$ when $r > \widetilde{r}(v)$). The terrorist therefore attack site 2 whenever $r > \widetilde{r}(v)$, leaving the government with an expected loss of $A\delta_2(R-r)$. But this means that the government can now profitably deviate from r to any r' between and $\widetilde{r}(v)$ and r as the government's expected loss is increasing in r.

In sum, the equilibria of the complete information game have the government allocating $\widetilde{r}(v)$ to site 1; the terrorist attacking site 1 if $r < \widetilde{r}(v)$; attacking site 2 if $r > \widetilde{r}(v)$; and attacking site 1 with any probability $\alpha \in [0,1]$ if $r = \widetilde{r}(v)$.

PROPOSITION 1: *The unique subgame perfect equilibrium allocation in the complete-information game, $\widetilde{r}(v)$, minmaxes the terrorists; i.e., $\widetilde{r}(v)$ is the unique solution to* $\min_{r \in [0,R]} \max\{\delta_1(r)+v, A\delta_2(R-r)\}$.

Powell (2006) proves this result. He also shows the government's minmaxing the terrorist group remains the unique equilibrium allocation even if the game is non-zero sum and regardless of whether the terrorists can observe the government's allocation before deciding which site to attack.

A direct implication of Proposition 1 is that the more vulnerable site 1 is, i.e., the larger v, the more the government allocates to it. Implicit differentiation of $\delta_1(\widetilde{r}(v))+v = A\delta_2(R-r)$ yields $d\widetilde{r}(v)/dv = -1/[\delta_1' + A\delta_2'] > 0$.

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

Private Information about Vulnerability

To formalize the trade off between hardening a very vulnerable site and calling attention to the vulnerability of that site, suppose the government knows the vulnerability v of site 1 but the terrorists do not. Uncertain of the exact level of vulnerability, the terrorists believe that v is distributed over $[\underline{v}, \overline{v}]$ according to the cumulative distribution Λ which has a strictly positive density over $(\underline{v}, \overline{v})$. Λ is also assumed to be common knowledge.

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

The government and terrorist group are playing a signaling game in which the government's allocation may reveal information to the terrorists about the relative vulnerability of the two sites. Suppose, for example, that the government knows that site 1 is very vulnerable (v is high). On the one hand, the high level of vulnerability of site 1 inclines the government 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 site 1 is quite vulnerable and thereby make an attack on that site more likely. The equilibria of the signaling game formalize the resolution of this trade off.

The remainder of this section characterizes the equilibria of the signaling game and shows that secrecy concerns swamp vulnerability issues. The government pools on a single allocation in the unique equilibrium outcome of the game. This allocation is optimal with respect to the average vulnerability. But this allocation spends too little on site 1if its actual vulnerability is higher than average and too much if its vulnerability is lower than average relative to what the government would spend if there were no uncertainty about v.

Formally, a pure-strategy for the government specifies the government's allocation as a function of its private information about the vulnerability: $r : [\underline{v}, \overline{v}] \to [0, R]$.[9] A strategy for the terrorists specifies the probability of attacking site 1 as a function of the allocation the terrorists observe: $\alpha : [0, R] \to [0, 1]$. As for what the terrorists believe about the vulnerability of site 1 after seeing a given allocation to that site, let Δ be the set of probability distributions over $[\underline{v}, \overline{v}]$ and let $\mu(r) \in \Delta$ for all $r \in [0, R]$ denote the terrorists' beliefs after observing allocation r. Finally, a perfect Bayesian equilibrium (PBE) is a strategy profile (r, α) and system of beliefs μ such that (i) the government can never profitably deviate from offering $r(v)$ given the terrorists' strategy $\alpha(r)$; (ii) $\alpha(r)$ maximizes the terrorists' payoff against r given $\mu(v|r)$; and (iii) $\mu(v \backslash r)$ is derived from Λ and r via Bayes' rule.[10]

Three lemmas help characterize the equilibria. The first demonstrates that in any equilibrium the terrorists are (weakly) less likely to attack site 1 if that site is more vulnerable. That is, the government allocates enough to site 1 to induce the terrorists to be (weakly) less likely to attack that site when it is more vulnerable (i.e., v is higher).

As for Lemma 2, recall that $\widetilde{r}(v)$ is the allocation that makes v indifferent between an

9 Only strategy profiles and equilibria in which the government plays a pure strategy will be considered.

10 Defining PBE's with a continuum of types also raises an additional technical issue. Consider, for example, a separating equilibrium in which each type 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 terrorist's beliefs following any allocation. Bayes' rule does not require the terrorists to believe they are 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 terrorists' beliefs following r' must be contained in the closure of $\{v : r(v) = r'\}$. See Ramey (1996) for a definition of a sequential or perfect Bayesian equilibrium with a continuum of types.

attack on site 1 or 2 (i.e., $\widetilde{r}(v)$ satisfies $\delta_1(\widetilde{r}(v)) + v = A\delta_2(R - \widetilde{r}(v))$). Lemma 2 then ensures that v's equilibrium loss is more than $A\delta_2(R - \widetilde{r}(v))$ which it can guarantee itself by offering $\widetilde{r}(v)$.

Lemma 3 shows that if v' separates, i.e., no other v'' allocates $r(v')$, then v' must be spending $\widetilde{r}(v')$. As Proposition 1 shows, $\widetilde{r}(v')$ is also what the government would allocate to site 1 if there were no uncertainty about the level of vulnerability. Lemma 3 thus shows that if the government foregoes all of the benefits of secrecy by acting in a way that fully discloses the vulnerability of site 1, then the government allocates its resources optimally with respect to the vulnerability issue. That is, the government allocates its resources as it would have had site 1's vulnerability been common knowledge at the outset.

Proposition 2 shows that secrecy concerns swamp the vulnerability issue in that the government always pools in equilibrium. All v allocate $\widetilde{r}(v^*)$ to site 1 where v^* is the average vulnerability, i.e., $v^* \equiv \int v d\Lambda$. Because the government pools on the same allocation regardless of the actual level of v, the terrorists cannot infer anything about v after observing the government's allocation of $\widetilde{r}(v^*)$.

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

To gain some intuition for the forces undermining separation, suppose $v < v' < v''$ did separate. Lemmas 2 and 3 imply that these types respectively allocate $\widetilde{r}(v)$, $\widetilde{r}(v')$, $\widetilde{r}(v'')$ and suffer equilibrium losses of $A\delta_2(R-\widetilde{r}(v))$, $A\delta_2(R-\widetilde{r}(v'))$, and $A\delta_2(R-\widetilde{r}(v''))$. Suppose further that the terrorists attack site 1 with probability α' after observing allocation $\widetilde{r}(v')$.

This probability cannot be too small if high-vulnerability types are to be deterred from mimicking v' by allocating $\widetilde{r}(v')$ to site 1. If, that is, α' is too small, then high-vulnerability types would be willing to "buy" a much lower probability of attack on the very vulnerable site 1 in return for the cost of a somewhat smaller allocation to that site. Conversely, the probability that the terrorists strike site 1 cannot be too high if low vulnerability types are to be deterred from deviating to $\widetilde{r}(v')$. If α' is too high, then low-vulnerability types will want to buy the correspondingly low probability of attack on the relatively more vulnerable site 2. Hence, higher values of α' discourage high-vulnerability types from deviating but encourage low-vulnerability types to deviate and vice-versa. As it happens, no α' can satisfies both of these constraints simultaneously and consequently there is no separation.

That the government must pool in equilibrium immediately implies that it offers $\widetilde{r}(v^*)$. Suppose the government pools on $\widehat{r} \in (\widetilde{r}(\underline{v}), \widetilde{r}(\overline{v}))$.[11] If $\alpha(\widehat{r}) = 1$, Lemma 2 guarantees that any v such that $\widetilde{r}(v) > \widehat{r}$ can profitably deviate to $\widetilde{r}(v)$ as $\delta_1(\widehat{r}) + v > \delta_1(\widetilde{r}(v)) + v$ where, recall, these are losses so lower is better. This contradiction implies $\alpha(\widehat{r}) < 1$. If $\alpha(\widehat{r}) = 0$, any v for which ensures $\widetilde{r}(v) < \widehat{r}$ can profitably deviate to $\widetilde{r}(v)$ since $A\delta_2(R - \widehat{r}) > A\delta_2(R - \widetilde{r}(v))$. This contradiction leaves $0 < \alpha(\widehat{r}) < 1$.

That the terrorists mix after observing $\widehat{r}$ implies that the attacker is indifferent between attacking sites 1 and 2. Pooling also means that the terrorists' posterior beliefs are unchanged from their prior beliefs. These facts imply $\int[\delta_1(\widehat{r})) + v]d\Lambda = \int A\delta_2(R - \widehat{r}))d\Lambda$. Rewriting this equation gives $v^* = \int v d\Lambda = A\delta_2(R - \widehat{r}) - \delta_1(\widehat{r}))$ and $\widehat{r} = \widetilde{r}(v^*)$.

Comparative statics follow immediately from this. The higher the average vulnerability of site 1, the more the government allocates to that site ($\widetilde{r}(v^*)$ is increasing in v^*) and the more likely that site is to be attacked.

Propositions 1 and 2 also show that uncertainty makes the government better off. The former shows that v allocates $\widetilde{r}(v)$ to site 1 and suffers an expected loss of $A\delta_2(R - \widetilde{r}(v))$ when there is no uncertainty about v. When there is, v allocates $\widetilde{r}(v^*)$ to this site and it is easy to show that this allocation results in an expected loss strictly less than $A\delta_2(R - \widetilde{r}(v))$ as long as $v \neq v^*$.

[11] The proof of Proposition 2 shows that the government's equilibrium allocation always lies in this range.

Formalizing the claims:[12]

LEMMA 1: *Let $\{r(v), \alpha(r), \mu\}$ be a PBE of the signaling game. If $v < v'$, then $\alpha(r(v)) \geq \alpha(r(v'))$.*

Proof: Suppose $v' < v''$, and let $r' = r(v')$, $r'' = r(v'')$, $\alpha' = \alpha(r(v'))$, and $\alpha'' = \alpha(r(v''))$. Incentive compatibility requires that neither v' nor v'' can gain by mimicking the other's strategy implies:

$$\alpha'[\delta_1(r') + v'] + (1 - \alpha')A\delta_2(R - r') \leq \alpha''[\delta_1(r'') + v'] + (1 - \alpha'')A\delta_2(R - r'')$$

$$\alpha''[\delta_1(r'') + v''] + (1 - \alpha'')A\delta_2(R - r'') \leq \alpha'[\delta_1(r') + v''] + (1 - \alpha')A\delta_2(R - r')$$

where, recall, these are expected losses so lower is better. Subtracting the first inequality from the second yields $\alpha''(v'' - v') \leq \alpha'(v'' - v')$. Hence, $\alpha(r(v')) \geq \alpha(r(v'))$ and $\alpha(r(v))$ is weakly decreasing in v. ■

LEMMA 2: *Let the allocation $\widetilde{r}(v)$ be defined by $v = A\delta_2(R - \widetilde{r}(v)) - \delta_1(\widetilde{r}(v))$. Then type v's expected loss in any PBE is less than or equal to its loss of offering $\widetilde{r}(v)$ which is $A\delta_2(R - \widetilde{r}(v)) = \delta_1(\widetilde{r}(v)) + v$.*

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

LEMMA 3: *Let $\{r(v), \alpha(r), \mu\}$ be any PBE in which v' separates, i.e., v' alone offers $r(v')$. Then $r(v') = \widetilde{r}(v')$.*

Proof: Arguing by contradiction, suppose v' offers $r' > \widetilde{r}(v')$. Because v' separates, the terrorists infer they are facing v' after observing r'. Given this inference and r', the terrorists strictly prefer attacking site 2. That is, $r' > \widetilde{r}(v')$ implies $A\delta_2(R - r') > A\delta_2(R - \widetilde{r}(v')) = \delta_1(\widetilde{r}(v')) + v' > \delta_1(r') + v'$. Hence, the terrorists strike site 2, leaving v' with an expected loss of $A\delta_2(R - r')$. But, as Lemma 2 shows, v' can always obtain an expected loss of $A\delta_2(R - \widetilde{r}(v'))$ by playing $\widetilde{r}(v')$, and $r' > \widetilde{r}(v')$ implies that this loss is less than $A\delta_2(R - r')$. This contradiction ensures that v' never separates with an allocation

[12] Readers less interested in the formal statements and proofs can skip the rest of this section.

greater than $\widetilde{r}(v')$. A similar argument shows that the terrorists hit site 1 after observing r' and that v' could once more profitably deviate to $\widetilde{r}(v')$ if $r' < \widetilde{r}(v')$. Thus v' must offer $\widetilde{r}(v')$ if it separates. ∎

These lemmas lead to:

Proposition 2: *Let $v^* \equiv \int v d\Lambda$ be the average vulnerability. Then perfect Bayesian equilibria equilibria exist and all of them share the same equilibrium path, namely, all v pool on $\widetilde{r}(v^*)$ and the terrorists attack site 1 with probability $\alpha(\widetilde{r}(v^*)) = A\delta_2'(R - \widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))]$.*

Proof: See the appendix.

This result immediately gives:

Corollary 2: *The higher the average vulnerability v^*, the more the government allocates to site 1 and the more likely the terrorists are to attack that site: $d\widetilde{r}(v^*)/dv^* > 0$ and $d\alpha(\widetilde{r}(v^*))/dv^* > 0$.*

Propositions 1 and 2 also imply:

Corollary 3: *Uncertainty about v makes the government better off: v's expected loss to allocating $\widetilde{r}(v^*)$ to site 1 is strictly less than its loss to allocating $\widetilde{r}(v)$ for all $v \neq v^*$.*

Proof: Proposition 1 shows that v allocates $\widetilde{r}(v)$ to site 1 and this results in an expected loss of $A\delta_2(R - \widetilde{r}(v))$ when there is complete information. In the asymmetric information game, v allocates $\widetilde{r}(v)$ and does strictly better at this allocation if

$$A\delta_2(R - \widetilde{r}(v)) \quad > \quad \alpha(\widetilde{r}(v^*))[\delta_1(\widetilde{r}(v^*) + v] + [1 - \alpha(\widetilde{r}(v^*)]A\delta_2(R - \widetilde{r}(v^*))$$

$$\alpha(\widetilde{r}(v^*))(v^* - v) \quad > \quad A\delta_2(R - \widetilde{r}(v^*)) - A\delta_2(R - \widetilde{r}(v)).$$

Let v' be any level of vulnerability less than v^* and v'' be any level greater than v^*. Then the previous inequality implies:

$$\frac{A\delta_2(R - \widetilde{r}(v^*)) - A\delta_2(R - \widetilde{r}(v'))}{v^* - v'} < \alpha(\widetilde{r}(v^*)) < \frac{A\delta_2(R - \widetilde{r}(v'')) - A\delta_2(R - \widetilde{r}(v^*))}{v'' - v^*}.$$

The function $A\delta_2(R - \widetilde{r}(v))$ is convex in v: $d^2 A\delta_2(R - \widetilde{r}(v))/dv^2 > 0$. So the bounds on $\alpha(\widetilde{r}(v^*))$ hold for all $v \neq v^*$ only if $\alpha(\widetilde{r}(v^*)) = A\delta_2'(R - \widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))]$

which is precisely what Proposition 2 demonstrates it to be. (See the proof of Proposition 2 for the details of the convexity argument.) ■

Private Information about Vulnerability with Unobservable Allocations

The preceding assumed that the terrorists were able to observe the government's allocation before deciding which site to attack. The government, therefore, had to be concerned about what its allocation might signal about the vulnerability of site 1. Suppose instead that the government's allocation is unobservable. How does the fact that the terrorists remain uncertain of the vulnerability of site 1 affect the government's allocation?

Observability has no effect on the government's allocation as long as there is no uncertainty about the level of vulnerability (see Proposition 1 and Powell 2006). The government minmaxes the terrorists by allocating $\widetilde{r}(v)$ to site 1. As the vulnerability increases, so does the government's allocation (Corollary 1).

What happens if the terrorists are uncertain of v but believe it to be distributed according to Λ as assumed above? Because the terrorists can no longer observe the government's allocation r, they cannot condition their decisions on it. The terrorists' strategy is simply a probability of attacking site 1, $a \in [0, 1]$, rather than a function specifying a possibly different probability for every value of r. A Bayesian Nash equilibrium is an allocation strategy, $r(v)$, and a probability of attack a such that for all $v \in [\underline{v}, \overline{v}]$ allocating $r(v)$ to site 1 minimizes v's expected loss given a, and attacking with probability a maximizes the terrorists' expected payoff given $r(v)$ and their beliefs Λ.

The government's expected loss to allocating r against a given a level of vulnerability v is $L(r, a) = a[\delta_1(r) + v] + A\delta_2(R - r)$. Differentiation then implies that the optimal allocation r^* satisfies the first-order condition $a\delta_1'(r^*) = A\delta_2'(R - r^*)$.[13] The solution to this equation does not depend on v, and this implies that all v allocate r^* to site 1 even though the allocations are unobservable. (The insensitivity of the optimal allocation to v is a direct consequence of the assumption that the vulnerability is linear in v and therefore that changes in v shift the level of vulnerability but have no effect on the

[13] The second-order condition, $\partial^2 L/\partial r^2 = a\delta_1'' + (1 - a)\delta_2'' > 0$, ensures r^* minmizes L.

margins, i.e., $\partial^2\delta_1(r,v)/\partial v\partial r = \partial^2[\delta_1(r)+v]/\partial v\partial r = 0$. If changes in v also affect the margins ($\partial^2\delta_1(r,v)/\partial v\partial r \neq 0$), then the government's allocation will vary with v when its allocation is unobservable.[14])

The terrorists must also play a mixed strategy, i.e., $0 < a < 1$, if the government's allocation is unobservable. Arguing by contradiction, assume that $a = 1$. Then v's expected loss to r is $L(r,1) = \delta_1(r) + v$ which v will minimize by spending all of its resources on site 1. That is, $r(v) = R$ if $a = 1$. But if all v allocate all of their resources to site 1, the terrorists strictly prefer to attack site 2. The terrorists' expected payoff to striking site 1 with $r(v) = R$ is $\int[\delta_1(R)+v]d\Lambda$ and $A\delta_2(0)$ to hitting site 2. But, $\int[\delta_1(R)+v]d\Lambda = \delta_1(R)+v^* < \delta_1(R)+\overline{v} < A\delta_2(0)$ where, recall, the last inequality is part of the specification of the game and ensures that there is a meaningful allocation problem facing the government. Thus, the terrorists would have a profitable deviation if $a = 1$ which is impossible in equilibrium and means $a < 1$. If $a = 0$, all v allocate 0 to site 1, and a similar argument shows that the terrorists will want to deviate to $a = 1$. Hence, $0 < a < 1$.

That the terrorists play a mixed strategy implies that they must be indifferent between attacking site 1 and site 2. This along with the fact that all v allocate r^* to site 1 implies $\int[\delta_1(r^*)+v]d\Lambda = \int A\delta_2(R-r^*)d\Lambda$. Hence $v^* = A\delta_2(R-r^*) - \delta_1(r^*)$ and $r^* = \widetilde{r}(v^*)$. The allocation $r^* = \widetilde{r}(v^*)$ must also be a best response to a, i.e., minimize $L(r,\alpha)$. Consequently, $\widetilde{r}(v^*)$ must satisfy the first-order condition $a\delta_1'(\widetilde{r}(v^*)) = (1-a)A\delta_2'(R - \widetilde{r}(v^*))$. Solving for a gives $a = A\delta_2'(R-\widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R-\widetilde{r}(v^*))]$.

Proposition 3: *If the government's allocation is unobservable, the terrorists mix in equilibrium:* $0 < a < 1$. *If the vulnerability of site 1 is linear in* v, *then all* v *allocate* $\widetilde{r}(v^*)$ *to site 1 and the terrorists attack this site with probability* $a = A\delta_2'(R-\widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*))+$

[14] The more general first-order condition that the optimal allocation must satisfy is $\alpha\partial\delta_1(r^*(v),v)/\partial r = (1-\alpha)A\delta_2'(R-r^*(v))$. Implicit differentiation then shows $dr^*(v)/dv = -\alpha(\partial^2\delta_1/\partial v\partial r)[\alpha\partial^2\delta_1/\partial v\partial r + (1-\alpha)A\delta_2'']^{-1}$ and $\operatorname{sgn}(dr^*(v)/dv) = -\operatorname{sgn}(\partial^2\delta_1/\partial v\partial r)$.

$A\delta_2'(R - \widetilde{r}(v^*))]$.[15]

As was the case with observable allocations, the terrorists' uncertainty about the level of vulnerability leads the government to spend less on relatively more vulnerable sites than it would were the terrorists certain of the level of vulnerability. That is, $v > v^*$ implies $\widetilde{r}(v) > r^* = \widetilde{r}(v^*)$. Conversely, the government spends more on relatively less vulnerable sites if the terrorists are uncertain: $\widetilde{r}(v) < r^* = \widetilde{r}(v^*)$ if $v < v^*$.

Proposition 3 also shows that uncertainty makes the government better off when the government's allocation is unobservable. If the level of vulnerability v were known, then v would spend $\widetilde{r}(v)$ on site 1 rather than $\widetilde{r}(v^*)$. But v's expected loss to $\widetilde{r}(v^*)$ is strictly less than its loss from $\widetilde{r}(v)$ for all $v \neq v^*$.

COROLLARY 4: *Uncertainty makes the government better off when the government's allocation in unobservable: v's expected loss to $\widetilde{r}(v^*)$ is strictly less than its expected loss to $\widetilde{r}(v^*)$ for all $v \neq v^*$.*

Proof: The argument is identical to the proof of Corollary 3.

Conclusion

Private information about the vulnerability of a site poses a trade off for the government if it is facing strategic terrorists who can observe the government's efforts to deal with a site's vulnerability. The more the government allocates to a site, the lower the expected loss if that site is attacked. But higher allocations may also signal to the terrorists that the site is more vulnerable and thereby increase the probability of an attack

[15] Propositions 2 and 3 show that the equilibrium paths of the game with observable and unobservable allocations are the same. All v allocate $\widetilde{r}(v^*)$ and $a = \alpha(v)$. That these paths are identical is a direct consequence of the assumption that the probability of an attack on site 1 is linear in v and will not be the case in general. As noted above, the government's allocation when it is unobservable will vary with v whenever $\partial^2\delta_1(r, v)/\partial v\partial r \neq 0$. In this more general case, the equilibrium strategies $r^*(v)$ and a satisfy the first-order condition, $a\partial\delta_1'(r^*(v), v)/\partial r = (1 - a)A\delta_2'(R - r^*(v))$ and $\int \delta_1(r^*(v), v)d\Lambda = \int A\delta_2(R - r^*(v))d\Lambda$. If, by contrast, the government pools on $\widehat{r}$ if its allocation is observable, then this allocation uniquely satisfies $\int \delta_1(\widehat{r}, v)d\Lambda = A\delta_2(R - \widehat{r})$. Recalling that $\operatorname{sgn}(dr^*(v)/dv) = -\operatorname{sgn}(\partial^2\delta_1/\partial v\partial r)$ (see footnote 14) and assuming that $\partial^2\delta_1/\partial v\partial r$ is either positive or negative for all v and r, at most only one v satisfies $r^*(v) = \widehat{r}$ and makes the same allocation regardless of the observability of the government's allocation.

on that site.

In equilibrium, secrecy concerns swamp the vulnerability issue. The government pools on a common allocation and therefore signals nothing to the terrorists. This allocation is optimal at the average level of vulnerability. Thus in order to avoid revealing anything to the terrorists the government allocates "too little" to site 1 if that site is very vulnerable and too much to that site if it is not very vulnerable. Even so, uncertainty makes the government better off. As the average level of vulnerability increases, the government spends more on that site.

Appendix

This section proves Proposition 2. Lemma 1A shows that there are no separating PBEs. Lemma 2A ensures that if v and v'' pool on the same allocation, then all $v' \in (v, v'')$ pool on the same allocation. Lemma 3A uses this result to demonstrate that no semi-separating equilibria exist. Hence, all PBEs of the game must be pooling. The final step in the proof is to construct a family of pooling equilibria, characterize their equilibrium paths, and show that all pooling PBE have the same equilibrium path.

LEMMA 1A: *There are no separating PBEs.*

Proof: Arguing by contradiction, assume a separating PBE $\{r(v), \alpha(r), \mu\}$ exists, i.e., $v \neq v'$ implies $r(v) \neq r(v')$. Lemma 3 then implies $r(v) = \widetilde{r}(v)$ for all $v \in [\underline{v}, \overline{v}]$ where, recall, $\widetilde{r}(v)$ is uniquely defined by $\delta_1(\widetilde{r}(v)) + v = A\delta_2(R - \widetilde{r}(v))$.

Let $v' \in (\underline{v}, \overline{v})$, $\alpha' \equiv \widetilde{r}(v')$, and $v'' > v' > v$. That v'' cannot profitably mimic the less vulnerable type v' implies:

$$A\delta_2(R - \widetilde{r}(v'')) \leq \alpha'[\delta_1(\widetilde{r}(v') + v''] + (1 - \alpha')A\delta_2(R - \widetilde{r}(v'))$$

$$A\delta_2(R - \widetilde{r}(v'')) \leq A\delta_2(R - \widetilde{r}(v')) + \alpha'[\delta_1(\widetilde{r}(v') + v' - A\delta_2(R - \widetilde{r}(v'))] + \alpha'(v'' - v')$$

$$\alpha' \geq \frac{A\delta_2(R - \widetilde{r}(v'')) - A\delta_2(R - \widetilde{r}(v'))}{v'' - v'}.$$

That v cannot profitably mimic the more vulnerable v' means α' must satisfy $A\delta_2(R - \widetilde{r}(v)) \leq \alpha'[\delta_1(\widetilde{r}(v') + v] + (1 - \alpha')A\delta_2(R - \widetilde{r}(v'))$ which implies $\alpha' \leq [A\delta_2(R - \widetilde{r}(v')) - A\delta_2(R - \widetilde{r}(v)]/(v' - v)$. Hence,

$$\frac{A\delta_2(R - \widetilde{r}(v'')) - A\delta_2(R - \widetilde{r}(v'))}{v'' - v'} \leq \alpha' \leq \frac{A\delta_2(R - \widetilde{r}(v')) - A\delta_2(R - \widetilde{r}(v))}{v' - v}.$$

No α' can satisfy both of these constraints and consequently no separating equilibria

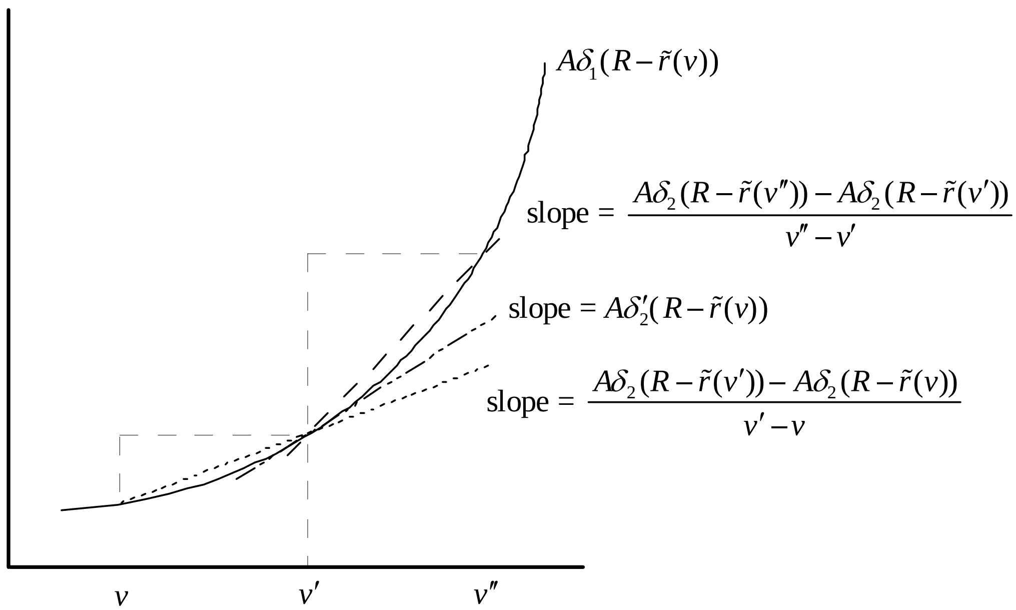

Figure 3: The implications of the convexity of $A\delta_2(R-\widetilde{r}(v))$.

exist. The function $A\delta_2(R-\widetilde{r}(v))$ is convex $(d^2A\delta_2(R-\widetilde{r}(v))/dv^2>0)$. Hence,

$$\frac{A\delta_2(R-\widetilde{r}(v'))-A\delta_2(R-\widetilde{r}(v))}{v'-v}<\frac{dA\delta_2(R-\widetilde{r}(v'))}{dv}<\frac{A\delta_2(R-\widetilde{r}(v''))-A\delta_2(R-\widetilde{r}(v'))}{v''-v'}.$$

as illustrated in Figure 3. Thus, α' cannot simultaneously satisfy the previous bounds on α'. ■

LEMMA 2A: *Let* $\{r(v),\alpha(r),\mu\}$ *be a PBE. If* $v<v''$ *and* $r(v)=r(v'')$*, then all* $v'\in(v,v'')$ *pool on* $r\equiv r(v)$*.*

Proof: Suppose v' offers $r'\neq r$. Lemma 1 implies $\alpha\equiv\alpha(r(v))\geq\alpha(r(v'))\geq\alpha(r(v''))=\alpha$. The terrorists therefore attack site 1 with the same probability after observing r or r'.

Type v', therefore, must be indifferent between allocating r or r' to site 1. Arguing by contradiction, assume v' strictly prefers r'. Then,

$$\alpha\left[\delta_1(r')+v'\right]+(1-\alpha)A\delta_2(R-r')<\alpha\left[\delta_1(r)+v'\right]+(1-\alpha)A\delta_2(R-r).$$

Rewriting this gives

$$\alpha\left[\delta_1(r') + v\right] + (1-\alpha)A\delta_2(R-r') < \alpha\left[\delta_1(r) + v\right] + (1-\alpha)A\delta_2(R-r).$$

Hence v too strictly prefers r' to its equilibrium allocation r. This contradiction ensures v' is indifferent between r and r'.

This indifference implies

$$\alpha\left[\delta_1(r') + v'\right] + (1-\alpha)A\delta_2(R-r') = \alpha\left[\delta_1(r) + v'\right] + (1-\alpha)A\delta_2(R-r)$$

$$\alpha = \frac{A\delta_2(R-r') - A\delta_2(R-r)}{A\delta_2(R-r') - \delta_1(r') - [A\delta_2(R-r) - \delta_1(r)]}$$

$$\alpha = \frac{A\delta_2(R-r') - A\delta_2(R-r)}{\widetilde{v}(r') - \widetilde{v}(r)} \equiv \Phi(r', r)$$

where $\widetilde{v}(s)$ is the type that is indifferent between an attack on site 1 or 2 given allocation s, i.e., $\widetilde{v}(s) = A\delta_2(R-s) - \delta_1(s)$.

Observe further that $\Phi(r', r) = \Phi(r, r')$ and $\Phi(r', r)$ is strictly increasing in r'. To establish the latter, note that $\widetilde{v}(s)$ is increasing in s and recall that $A\delta_2(R - \widetilde{r}(v))$ is convex. Then, $r'' > r'$ implies

$$\frac{A\delta_2(R - \widetilde{r}(\widetilde{v}(r''))) - A\delta_2(R - \widetilde{r}(\widetilde{v}(r)))}{\widetilde{v}(r'') - \widetilde{v}(r)} > \frac{A\delta_2(R - \widetilde{r}(\widetilde{v}(r'))) - A\delta_2(R - \widetilde{r}(\widetilde{v}(r)))}{\widetilde{v}(r') - \widetilde{v}(r)}.$$

But $\widetilde{r}(\widetilde{v}(r'')) = r''$ and $\widetilde{r}(\widetilde{v}(r')) = r'$, so $\Phi(r'', r) > \Phi(r', r)$.

Type $\widehat{v}$ will not play r in equilibrium if it can do better by allocating $\widetilde{r}(\widehat{v})$, i.e., if

$$A\delta_2(R - \widetilde{r}(\widehat{v})) \quad < \quad \alpha\left[\delta_1(r) + \widehat{v}\right] + (1-\alpha)A\delta_2(R-r)$$

$$\alpha[\widehat{v} - \widetilde{v}(r)] \quad > \quad A\delta_2(R - \widetilde{r}(\widehat{v})) - A\delta_2(R-r)$$

If $\widetilde{r}(\widehat{v}) > r$, then $\widehat{v} = \widetilde{v}(\widetilde{r}(\widehat{v})) > \widetilde{v}(r)$ and $\widehat{v}$ strictly prefers playing $\widetilde{r}(\widehat{v})$ to r when $\alpha > \Phi(\widetilde{r}(\widehat{v}), r)$ But $\alpha = \Phi(r', r)$, so $\widehat{v}$ will not play r in equilibrium if $\Phi(r', r) > \Phi(\widetilde{r}(\widehat{v}), r)$ or equivalently if $r' > \widetilde{r}(\widehat{v})$. Thus, no $\widehat{v}$ such that $r < \widetilde{r}(\widehat{v}) < r'$ ever allocates r in equilibrium as it could do strictly better by playing $\widetilde{r}(\widehat{v})$.

Similarly, no $\widehat{v}$ such that $r < \widetilde{r}(\widehat{v}) < r'$ plays r'. Paralleling the argument in the previous paragraph, $\widehat{v}$ will not play r' in equilibrium if it can do better by allocating $\widetilde{r}(\widehat{v})$, i.e., if $\alpha[\widehat{v} - \widetilde{v}(r')] > A\delta_2(R - \widetilde{r}(\widehat{v})) - A\delta_2(R - r')$. If $\widetilde{r}(\widehat{v}) < r'$, then $\alpha < [A\delta_2(R - r') - A\delta_2(R - \widetilde{r}(\widehat{v}))]/[\widetilde{v}(r') - \widehat{v}]$. So $\widehat{v}$ never plays r' when $\Phi(r', r) = \alpha < \Phi(r', \widetilde{r}(\widehat{v}))$. But $\Phi(r', r) < \Phi(r', \widetilde{r}(\widehat{v}))$ is equivalent to $\Phi(r, r') < \Phi(\widetilde{r}(\widehat{v}), r')$ which then implies $r < \widetilde{r}(\widehat{v}))$. Thus, no $\widehat{v}$ such that $r < \widetilde{r}(\widehat{v}) < r'$ ever allocates r' in equilibrium.

Rewriting α as $\alpha = [1 + (\delta_1(r) - \delta_1(r'))/(A\delta_2(R - r') - A\delta_2(R - r))]^{-1}$ shows $0 < \alpha < 1$. That the terrorists' mix means that they are indifferent between attacking site 1 and site 2 after observing r or r'. This yields $\int [\delta_1(r) + v] d\Lambda_r = \int A\delta_2(R - r) d\Lambda_r$ and $\int [\delta_1(r') + v] d\Lambda_{r'} = \int A\delta_2(R - r') d\Lambda_{r'}$ where Λ_r and $\Lambda_{r'}$ are the conditional distributions over v given r and r'. Let v_r^* and $v_{r'}^*$ be the means of these distributions. Then, $v_r^* = A\delta_2(R - r) - \delta_1(r)$ and $v_{r'}^* = A\delta_2(R - r') - \delta_1(r')$ which leave $\widetilde{r}(v_r^*) = r$ and $\widetilde{r}(v_{r'}^*) = r'$.

That $v_r^* = \int v d\Lambda_r$ and $v_{r'}^* = \int v d\Lambda_{r'}$ also imply that Λ_r puts positive weight on some $v \leq v_r^*$ and Λ_r puts positive weight on some $v \geq v_{r'}^*$. Thus there exist a w and w' such that $w \leq v_r^*$ and $w' \geq v_{r'}^*$ with $r(w) = r$ and $r(w') = r'$.

Now take $\widehat{w} \in (v_r^*, v_{r'}^*)$ which means $w < \widehat{w} < w'$. Lemma 1 then gives $\alpha = \alpha(r) = \alpha(r(w)) \geq \alpha(r(\widehat{w})) \geq \alpha(r(w')) = \alpha(r') = \alpha$. So, $\alpha(r(\widehat{w})) = \alpha$.

That $\widehat{w} \in (v_r^*, v_{r'}^*)$ also ensures $\widehat{w}$ strictly prefers not to offer r, so

$$\alpha [\delta_1(\widehat{r}) + \widehat{w}] + (1 - \alpha) A\delta_2(R - \widehat{r}) < \alpha [\delta_1(r) + \widehat{w}] + (1 - \alpha) A\delta_2(R - r)$$

where $\widehat{r} \equiv r(\widehat{w})$. But this expression implies all v strictly prefer $\widehat{r}$ to r as

$$\alpha [\delta_1(\widehat{r}) + v] + (1 - \alpha) A\delta_2(R - \widehat{r}) < \alpha [\delta_1(r) + v] + (1 - \alpha) A\delta_2(R - r)$$

which contradicts the assumption that v and v'' pool on r. This contradiction ensures that all $v' \in (v, v'')$ allocate r whenever v and v'' do. ■

A PBE is semi-separating if at least two types allocate the same amount to site 1 and there are at least two distinct allocations. Lemma 3A shows that there are no semi-separating PBEs.

LEMMA 3A: *The signaling game has no semi-separating PBEs.*

Proof: Arguing by contradiction, assume a semi-separating equilibrium exists and take $\widehat{r}$ to be an allocation on which there is partial pooling, i.e., there are v' and v'' such that $v' \neq v''$ and $\widehat{r} \equiv r(v') = r(v'')$. Lemma 2A then implies $r(v) = \widehat{r}$ for all $v \in (v', v'')$. Let $\underline{w} = \inf\{v : r(v) = \widehat{r}\}$ and $\overline{w} = \sup\{v : r(v) = \widehat{r}\}$ where $\underline{w} \leq v' < v'' \leq \overline{w}$. Also define $w^* = A\delta_2(R - \widehat{r}) - \delta_1(\widehat{r})$ which implies $\widetilde{r}(w^*) = \widehat{r}$.

The terrorists must mix after $\widehat{r}$, i.e., $\alpha(\widehat{r}) \in (0, 1)$. To establish this, assume $\alpha(\widehat{r}) = 1$. Then v could only play $\widehat{r}$ if its loss were no more than that of playing $\widetilde{r}(v)$, i.e., if $A\delta_2(R - \widetilde{r}(v)) = \delta_1(\widetilde{r}(v)) + v \geq \alpha(\widehat{r})[\delta_1(\widehat{r}) + v] + (1 - \alpha(\widehat{r}))A\delta_2(R - \widehat{r}) = \delta_1(\widehat{r}) + v$. Hence, $\delta_1(\widetilde{r}(v)) \geq \delta_1(\widehat{r}) = \delta_1(\widetilde{r}(w^*)) \Leftrightarrow \widetilde{r}(v) \leq \widetilde{r}(w^*) \Leftrightarrow w^* \geq v$.

The last inequality implies $w^* \geq \sup\{v : r(v) = \widehat{r}\} = \overline{w}$. Let $\Lambda_{\widehat{r}}$ be the posterior of Λ given $\widehat{r}$. Because Λ has a strictly postive density over $(\underline{v}, \overline{v})$ and $\underline{w} < \overline{w}$, $\int v d\Lambda_{\widehat{r}} < w^* = A\delta_2(R - \widehat{r}) - \delta_1(\widehat{r})$ and therefore that $\int A\delta_2(R - \widehat{r}) d\Lambda_{\widehat{r}} > \int [\delta_1(\widehat{r}) + v] d\Lambda_{\widehat{r}}$. This inequality means that the terrorists payoff to striking site 2 is strictly greater than that of attacking site 1 given $\widehat{r}$. Hence $\alpha(\widehat{r})$ must be zero in any PBE. But this contradicts the assumption that $\alpha(\widehat{r}) = 1$ and establishes that $\alpha(\widehat{r}) < 1$. A similar argument shows that $\alpha(\widehat{r}) = 0$ also leads to a contradiction, which leaves $\alpha(\widehat{r}) \in (0, 1)$.

That the terrorists mix after seeing $\widehat{r}$ implies that that they are indifferent between attacking site 1 and site 2. Thus,

$$\begin{aligned} \int [\delta_1(\widehat{r}) + v] d\Lambda_{\widehat{r}} &= \int A\delta_2(R - \widehat{r}) d\Lambda_{\widehat{r}} \\ \int v d\Lambda_{\widehat{r}} &= A\delta_2(R - \widehat{r}) - \delta_1(\widehat{r}) \\ \int v d\Lambda_{\widehat{r}} &= w^*. \end{aligned}$$

Finally, note that Lemma 2 ensures that all $w \in (\underline{w}, \overline{w})$ pool on $\widehat{r}$. This along with the fact that Λ has a strictly positive density over $(\underline{v}, \overline{v})$ guarantees $\underline{w} < w^* < \overline{w}$.

In order to deter the w in a small neighborhood of w^* from deviating from $\widehat{r}$ to $\widetilde{r}(w)$, it must be that

$$\alpha(\widetilde{r}(w^*)) = \frac{A\delta_2'(R - \widetilde{r}(w^*))}{\delta_1'(\widetilde{r}(w^*)) + A\delta_2'(R - \widetilde{r}(w^*))}.$$

To establish this, note that w's loss from $\widehat{r}$ can be no more than its loss from $\widetilde{r}(w^*)$. Taking $\widehat{\alpha} \equiv \alpha(\widetilde{r}(w^*)) = \alpha(\widehat{r})$,

$$\begin{aligned} A\delta_2(R - \widetilde{r}(w)) &\geq \widehat{\alpha}[\delta_1(\widehat{r}) + w] + (1 - \widehat{\alpha})A\delta_2(R - \widehat{r}) \\ &\geq A\delta_2(R - \widehat{r}) + \widehat{\alpha}(w - w^*) \\ \widehat{\alpha}(w^* - w) &\geq A\delta_2(R - \widetilde{r}(w^*)) - A\delta_2(R - \widetilde{r}(w)) \end{aligned}$$

If $w < w^*$, $\widehat{\alpha} \geq [A\delta_2(R - \widehat{r}) - A\delta_2(R - \widetilde{r}(w^*))]/(w^* - w)$. The expression on the right of the inequality is increasing in w because $A\delta_2(R - \widetilde{r}(w))$ is convex in w. Hence the previous inequality will hold for all $w < w^*$ only if

$$\begin{aligned} \widehat{\alpha} &\geq \lim_{w \uparrow w^*} \frac{A\delta_2(R - \widetilde{r}(w^*)) - A\delta_2(R - \widetilde{r}(w))}{w^* - w} \\ &\geq A\delta_2'(R - \widetilde{r}(w^*))\frac{d\widetilde{r}(w^*)}{dw} \\ \alpha(\widetilde{r}(w^*)) &\geq \frac{A\delta_2'(R - \widetilde{r}(w^*))}{\delta_1'(\widetilde{r}(w^*)) + A\delta_2'(R - \widetilde{r}(w^*))}. \end{aligned}$$

An analogous argument shows that $\alpha(\widetilde{r}(w^*)) \leq A\delta_2'(R - \widetilde{r}(w^*))/[\delta_1'(\widetilde{r}(w^*)) + A\delta_2'(R - \widetilde{r}(w^*))]$ if $w > w^*$. The two bounds on $\alpha(\widetilde{r}(w^*))$ give $\alpha(\widetilde{r}(w^*)) = A\delta_2'(R - \widetilde{r}(w^*))/[\delta_1'(\widetilde{r}(w^*)) + A\delta_2'(R - \widetilde{r}(w^*))]$.

Because pooling is only partial, there exists a $t \in [\underline{v}, \overline{v}]$ such that $\tau \equiv r(t) \neq \widehat{r}$. Assume $t \geq \overline{w}$. Lemma 2A ensures that the set $\{v : r(v) = \tau\}$ is either an interval or a singleton. Suppose it to be an interval and let $t^* \equiv \int v d\Lambda_\tau$. Repeating the argument above shows that $\tau = \widetilde{r}(t^*)$, $t^* > \overline{w}$, and

$$\alpha(\tau) = \alpha(\widetilde{r}(t^*)) = \frac{A\delta_2'(R - \widetilde{r}(t^*))}{\delta_1'(\widetilde{r}(t^*)) + A\delta_2'(R - \widetilde{r}(t^*))}.$$

The expression on the right is increasing in t^*, so $t \geq \overline{w} > w^*$ implies $\alpha(r(t)) = \alpha(\widetilde{r}(t^*)) > \alpha(\widetilde{r}(w^*)) = \alpha(r(w^*))$. This, however, contradicts Lemma 1 which shows that $\alpha(r(v))$ is nonincreasing in v.

Suppose now that $\{v : r(v) = \tau\}$ is a singleton. Hence, t separates, so $r(t) = \widetilde{r}(t)$ by Lemma 3. Because t cannot profitably deviate to $\widehat{r} = \widetilde{r}(w^*)$,

$$\begin{aligned} A\delta_2(R - \widetilde{r}(t)) &\leq \alpha(\widetilde{r}(w^*))\,[\delta_1(\widetilde{r}(w^*)) + t] + [1 - \alpha(\widetilde{r}(w^*))]A\delta_2(R - \widetilde{r}(w^*)) \\ \alpha(\widetilde{r}(w^*)) &\geq \frac{A\delta_2(R - \widetilde{r}(t)) - A\delta_2(R - \widetilde{r}(w^*))}{t - w^*}. \end{aligned}$$

This again yields a contradiction as

$$\alpha(\widetilde{r}(w^*)) \geq \frac{A\delta_2(R - \widetilde{r}(t)) - A\delta_2(R - \widetilde{r}(w^*))}{t - w^*} > \frac{A\delta_2'(R - \widetilde{r}(w^*))}{\delta_1'(\widetilde{r}(w^*)) + A\delta_2'(R - \widetilde{r}(w^*))} = \alpha(r(w^*))$$

where the strict inequality follows from the convexity of $A\delta_2(R - \widetilde{r}(v))$.

In sum, the assumption that $t \geq \overline{w}$ and $r(t) \neq \widehat{r}$ leads to a contradiction. An analogous argument shows that $t \leq \underline{w}$ and $r(t) \neq \widehat{r}$ also leads to a contradiction. Hence, no allocations differ from $\widehat{r}$ and no semi-separating equilibria exist. ■

Lemmas 1A and 3A imply that any PBE of the signaling game must be pooling. It remains to be shown that pooling equilibria actually exists and that the equilibrium paths in all pooling equilibria are identical.

PROPOSITION 2: *Perfect Bayesian equilibria equilibria exist and all of them share the same equilibrium path, namely, all v pool on $\widetilde{r}(v^*)$ and the terrorists attack site 1 with probability $\alpha(\widetilde{r}(v^*)) = A\delta_2'(R - \widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))]$.*

Proof: Suppose the government pools on $\widehat{r}$. Then $\alpha(\widehat{r}) \in (0, 1)$. If $\alpha(\widehat{r}) = 1$, any v such that $\widetilde{r}(v) > \widehat{r}$ could profitably deviate to $\widetilde{r}(v)$ as $\delta_1(\widetilde{r}(v)) + v < \delta_1(\widehat{r}) + v$. Because all

v pool, it must be that $\delta_1(\widetilde{r}(v)) \geq \delta_1(\widehat{r})$ for all v. Hence, $\widehat{r} \geq \widetilde{r}(\overline{v})$. But the terrorists payoff to striking site 2 is now strictly larger than to hitting site 1: $\int[\delta_1(\widehat{r}) + v]d\Lambda < \delta_1(\widetilde{r}(\overline{v})) + \overline{v} = A\delta_2(R - \widetilde{r}(\overline{v})) \leq A\delta_2(R - \widehat{r})$. The terrorists best reply to $\widehat{r}$ therefore is $\alpha(\widehat{r}) = 0$. This contradiction leaves $\alpha(\widehat{r}) < 1$. An analogous argument shows $\alpha(\widehat{r}) > 0$.

Because the terrorists mix in response to $\widehat{r}$, they must be indifferent to attacking either site. This means $\int[\delta_1(\widehat{r}) + v]d\Lambda = \int A\delta_2(R - \widehat{r})d\Lambda$ and, consequently, $v^* = \int v d\Lambda = A\delta_2(R - \widehat{r}) - \delta_1(\widehat{r})$ and $\widehat{r} = \widetilde{r}(v^*)$.

Define $\widehat{\alpha} \equiv \alpha(\widehat{r})$. Then in order to deter v from deviating to $\widetilde{r}(v)$:

$$A\delta_2(R - \widetilde{r}(v)) \geq \widehat{\alpha}[\delta_1(\widehat{r}) + v] + (1 - \widehat{\alpha})A\delta_2(R - \widehat{r})$$

$$\widehat{\alpha}(v^* - v) \geq A\delta_2(R - \widetilde{r}(v^*)) - A\delta_2(R - \widetilde{r}(v)).$$

If $v < v^*$, $\widehat{\alpha} \geq [A\delta_2(R - \widetilde{r}(v^*)) - A\delta_2(R - \widetilde{r}(v))]/(v^* - v)$. Repeating the argument in the proof of Lemma 3A, the expression on the right is increasing in v because $A\delta_2(R - \widetilde{r}(v))$ is convex. That the inequality holds for all $v < v^*$ then implies

$$\begin{aligned} \widehat{\alpha} &\geq \lim_{v \uparrow v^*} \frac{A\delta_2(R - \widetilde{r}(v^*)) - A\delta_2(R - \widetilde{r}(v))}{v^* - v} \\ &\geq A\delta_2'(R - \widetilde{r}(v^*))\frac{d\widetilde{r}(v^*)}{dv} \\ &\geq \frac{A\delta_2'(R - \widetilde{r}(v^*))}{\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))}. \end{aligned}$$

An analogous argument demonstrates that deterring all $v > v^*$ from deviating to $\widetilde{r}(v)$ requires $\widehat{\alpha} \leq A\delta_2'(R - \widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))]$. This leaves $\alpha(\widehat{r}) = A\delta_2'(R - \widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))]$.

Now define the strategy $\alpha(r) = 1$ if $r < \widetilde{r}(\underline{v})$; $\alpha(r) = \widehat{\alpha}$ if $r \in [\widetilde{r}(\underline{v}), \widetilde{r}(\overline{v})]$; and $\alpha(r) = 0$ if $r > \widetilde{r}(\overline{v})$. Let Δ be the set of probability distributions with support $[\underline{v}, \overline{v}]$ and Δ_r be the subset of distributions such that $\mu(r) \in \Delta_r$ if and only if $\int v d\mu(r) = A\delta_2(R - r) - \delta_1(r)$.

The definition of Δ_r ensures that the terrorists are indifferent between attacking site 1 and site 2 following r if its beliefs are in Δ_r, i.e., if $\mu(r) \in \Delta_r$, then $\int v d\mu(r) = A\delta_2(R-r) - \delta_1(r)$ which is equivalent to $\int[\delta_1(r)+v]d\mu(r) = \int A\delta_2(R-r)d\mu(r)$. Now define the family of terrorists' beliefs $\mu : [0,R] \to \Delta$ to be any $\mu(r) \in \Delta$ for $r < \widetilde{r}(\underline{v})$ and $r > \widetilde{r}(\overline{v})$; $\mu(\widehat{r}) = \Lambda$; and $\mu(r) \in \Delta_r$ if $r \in [\widetilde{r}(\underline{v}), \widetilde{r}(\overline{v})]$ and $r \neq \widehat{r}$ (e.g., $\mu(r)$ puts probability one on $v = A\delta_2(R-r) - \delta_1(r)$).

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 $\mu(r)$. The terrorists strictly prefer attacking site 1 regardless of their beliefs if $r < \widetilde{r}(\underline{v})$. So $\alpha(r) = 1$ is the unique best reply to $r < \widetilde{r}(\underline{v})$ for any $\mu(r) \in \Delta$. Similarly, the terrorists strictly prefer hitting site 2 if $r > \widetilde{r}(\overline{v})$, and α again has the terrorists playing their unique best reply to r when $r > \widetilde{r}(\overline{v})$ for any $\mu(r) \in \Delta$.. If $r \in [\widetilde{r}(\underline{v}), \widetilde{r}(\overline{v})]$ and $r \neq \widehat{r}$, then $\mu(r) \in \Delta_r$ and the terrorists by construction are indifferent between attacking site 1 and 2. Hence any $\alpha(r)$ is a best response and inparticular $\alpha(r) = \widehat{\alpha}$ is. The beliefs $\mu(r)$ are also clearly consistent with Bayes' law. If all v pool on $r(v) = \widehat{r}$, then the only restriction that Bayes' law imposes is the terrorists posterior beliefs following $\widehat{r}$ is the prior Λ which $\mu(r)$ satisfies.

Therefore the assessment $(\widehat{r}, \alpha(r), \mu)$ will be a PBE if $\widehat{r}$ is a best response to $\alpha(r)$. The construction of $\widehat{\alpha}$ ensures that no v can profitably deviate from $\widehat{r}$ to $\widetilde{r}(v)$. It follows that no v can profitably deviate to any other $r' \neq \widetilde{r}(v)$. Let r' be any allocation in $[\widetilde{r}(\underline{v}), \widetilde{r}(\overline{v})]$. Then $r' = \widetilde{r}(v')$ for some $v' \in [\underline{v}, \overline{v}]$. To see that no $v \in [\underline{v}, \overline{v}]$ can profitably deviate from $\widehat{r}$ to r', note that $\alpha(r) = \widehat{\alpha}$ for all $r \in [\widetilde{r}(\underline{v}), \widetilde{r}(\overline{v})]$ and that v' cannot profitably deviate from $\widehat{r}$ to $r' = \widetilde{r}(v')$. These facts imply

$$\widehat{\alpha}[\delta_1(\widehat{r}) + v'] + (1-\widehat{\alpha})A\delta_2(R-\widehat{r}) \leq \widehat{\alpha}[\delta_1(r') + v'] + [1-\widehat{\alpha}]A\delta_2(R-r')$$

$$\widehat{\alpha}[\delta_1(\widehat{r}) + v] + (1-\widehat{\alpha})A\delta_2(R-\widehat{r}) \leq \alpha(r'))[\delta_1(r') + v] + [1-\alpha(r')]A\delta_2(R-r')$$

where the latter inequality holds for all v. Hence, no v can profitably deviate to r', so allocating $\widehat{r}$ is a best response to $\widehat{\alpha}$.

This completes the construction of a family of pooling equilibria corresponding to different beliefs μ, thereby showing that multiple pooling equilibria exist. However, the fact that the terrorists must mix in any pooling equilibrium ensures that the equilibrium path in any pooling equilibrium is $\widehat{r} = \widetilde{r}(v^*)$ and $\alpha(\widehat{r}) = A\delta_2'(R - \widetilde{r}(v^*))/[\delta_1'(\widetilde{r}(v^*)) + A\delta_2'(R - \widetilde{r}(v^*))]$. ■

References

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

GAO. 2005. "Risk Management: Further Refinements Needed to Assess Risks and Prioritize Protective Measures at Ports and Other Critical Infrastructure." (GAO-06-91). Washington, D.C.: General Accounting Office.

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

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

Powell, Robert. 2006. "Defending Against Terrorist Attacks with Limited Resources." Manuscript, Department of Political Science, UC Berkeley.

Ramey, Gary. 1996. "D1 Signaling Equilibria with Mulitple Signals and a Continuum of Types," *Journal of Economic Theory* 69 (May): 508-531.

Willis, Henry H., *et. al.* 2005. *Estimating Terrorism Risk.* Santa Monica, CA: Rand Corporation.

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