Mathematical Appendix for *Is This Paper Dangerous?*

This version October 15, 2008.

This document contains a formal game-theoretic presentation of the arguments in *Is This Paper Dangerous?* The majority of the proofs are by construction: we explicate the algorithm by which the unique equilibrium can be found, in most cases via backward induction. Accordingly, rather than list propositions before proofs, we list them after. In many cases our propositions reference the algorithms that precede them. Throughout the analysis we maintain the following assumptions:

- Number of potential target sites: N.
- Payoff to G if target i is struck successfully: $-L_i$.
- Payoff to T if target i is struck successfully: A_i.
- Resources allocated to site i: r_i with $\sum r_i \leq R$.
- Information released about site i: $s_i \in [0, 1]$.
- General information released: $x \in [0, 1]$.
- Probability that an attack on site i will succeed: $\delta_i(r_i, s_i, x)$, where $-\delta_i(r_i, s_i, x)$ is supermodular.[1]

[1] Note that we assume the cross-partial derivative for each pair of variables in $-\delta_i$ is strictly positive.

- Externality that obtains only if site i is attacked (whether successfully or not): $g_i(s_i, x)$. If positive, $g_i(s_i, x)$ is increasing in each variable and displays increasing differences. If negative, it is decreasing in each variable and its negative displays increasing differences. $g_i(0, 0) = 0$. We assume that $L_i > g_i(s_i, x) \forall i, s_i, x$ as otherwise G would *want* T to attack site i.

- Site-specific externality that obtains whether or not site i is attacked: $h_i(s_i)$, which is positive and increasing in s_i. $h_i(0) = 0$.

- Externality that always obtains: $h(x)$, which is positive and increasing in x. $h(0) = 0$.

- Probability that T knows of site i: $P_i^T(x, s_i)$, which is increasing in each variable and displays increasing returns.

- Probability that G knows of site i: $P_i^G(x)$, which is increasing in x.

1 General Information—Propositions 1 and 2

To analyze the role of general information we start by considering the simple setting in which both G and T are searching for vulnerable targets. At the start of each game (period 0), G decides on an information sharing policy $x \in [0, 1]$. Each period, T either pays cost C and randomly samples $k_T(x)$ targets from a finite set of N targets, or attacks one of the targets it has already seen and receives the appropriate payoff. If T searches for new targets, G likewise samples $k_G(x)$ targets from the the same set. Both $k_T(x)$ and $k_G(x)$ are increasing, concave functions of x. There are two types of targets: high-valued ones that yield L_1 and low-valued ones that yield $L_2 < L_1$. There are $H \leq N$ high-valued targets to be found and this number is common knowledge.[2] We assume that the chance of finding a target is independent of

[2]Rather than specify only two levels of attack utility, we can instead think of H as being the number of targets that meet the terrorists' aspiration level for political impact. So long as this number is common knowledge, the analysis remains the same.

its type.[3] If a target is unknown to G, then G cannot defend it so $\delta_i(x, unknown) = 1$ and T receives the above payoffs from striking it. If a target is known to G, then G can defend it perfectly so $\delta_i(x, known) = 0$, and T receives a payoff of zero for striking it. The probability of a target's being seen is uniform across targets, and so independent of its type. No externalities are present. Proposition 1 covers the case where T must decide how many periods to search at the beginning of the game. Proposition 2 covers the case where T can decide to attack or keep searching in every period.

We consider the simpler variant first. T's strategy here is a stopping time that maximizes its expected utility looking forward from the beginning of the game. This is a function of two things: the chance that G has not identified the target T will strike, and the payoff that T would receive if an attack were to succeed. As G's search procedure is independent of target type, the probability at any t that G has not seen the target T strikes is just $(1 - k_G(x)t/N)$.

Now consider what T can expect to receive upon attacking. Since T will always choose to attack a high-value (good) target over a low-value (bad) one, we need to know the probability of drawing at least one good target by period t. This is equal to one minus the probability of drawing no good targets in t periods (the periods from 0 to $t-1$), encompassing $k_T(x)t$ draws, which is $1 - \prod_{j=0}^{k_T(x)t-1}(1 - \frac{H}{N-j}) = 1 - \prod_{j=0}^{k_T(x)t-1}(\frac{N-H-j}{N-j}) = 1 - \frac{(N-H)_{(k_T(x)t)}}{N_{(k_T(x)t)}}$, where $a_{(n)} = a(a-1)\ldots(a-n+1)$.[4] The expected utility for T if it attacks in period t is thus:

$$\left(1 - \frac{k_G(x)t}{N}\right)\left[L_1 - \left(\frac{(N-H)_{(k_T(x)t)}}{N_{(k_T(x)t)}}\right)(L_1 - L_2)\right] - Ct. \tag{1}$$

The optimal stopping time is the time $t^*(x)$ that maximizes (1).[5] Denote this time $t^*(x)$.

[3]Relaxing this assumption, say by making the probability of being seen higher (lower) for low-valued targets, has a similar effect on the players' strategies as decreasing (increasing) H and so is not treated separately.

[4]This is known as a falling factorial.

[5]Whether or not a non-zero stopping time exists depends on the values of the parameters. Though one can differentiate (2), given its complexity and the fact that stopping times are discrete it is far easier to maximize (2) numerically via a simple line search. This is the method by which we compute the optimal stopping times in the text.

Changing the amount of information released alters three things: T's stopping point, and the rates of target discovery by both G and T. G's decision on x minimizes

$$\left(1 - \frac{k_G(x)t^*(x)}{N}\right)\left[L_1 - \left(\frac{(N-H)_{(k_T(x)t^*(x))}}{N_{(k_T(x)t^*(x))}}\right)(L_1 - L_2)\right], \tag{2}$$

which is just (1) with $C = 0$ and with $t^*(x)$ inserted for t.

Increasing k_G is strictly better for G and increasing k_T is strictly worse, so increasing the rate at which G finds targets relative to T's rate increases the range of parameter values over which it is beneficial to release information. (Though we do not explore this here, the same is true for the addition of a positive externality like $h(x)$, as it does not change T's stopping rule.) These results are summarized in the following proposition.

Proposition 1: Under a simple stopping rule, T chooses a stopping time and G chooses a level of information released according to the algorithm given above. The range of parameter values over which information is released increases as G becomes relatively more efficient at discerning targets than T is. If the optimal stopping time can be made less than one then T never searches at all, which is the best outcome for G.

Next we consider the more complex variant, in which T can decide to attack or to continue searching in every period. T's equilibrium strategy is a rule that dictates when it chooses to attack and which target it attacks for every configuration of known targets at every time. Consider first T's choice as to what target to attack. T knows that no payoff ever exceeds L_1, and that G has an equal—and increasing—chance of discovering any particular target during its search. Thus, as soon as it observes a good target it should stop searching, since it can do no better by waiting.

This is half of T's search strategy: attack a good target if seen. The other half is a rule for what to do when no high-value target has been identified. The rule simply compares the present value of attacking a worse target with the expected value of the game looking forward. Attacking a worse target at time t yields $(1 - k_G(x)t/N)L_2$. Because the game is of finite length—all the targets must eventually be found by G (at which point there is no benefit to T's searching)—T can use backward induction to find the expected value of the game.

G will have found all of the targets by period $t_{Last} + 1 = Ceiling(N/k_G)$,[6] so we know that in period t_{Last} it always pays for T to attack if it has not yet found a high-valued target. Whether T searches in period $t_{Last} - 1$ depends on the probability that T finds a good target in the next $k_T(x)$ draws. Since search is always conditioned on not having observed a good target yet, we can write this probability at time t as

$$P_{good}(t) \equiv 1 - \prod_{j=0}^{k_T(x)-1} (1 - \frac{H}{N - k_T(x)t - j}) = 1 - \frac{(N - H - k_T(x)t)_{(k_T(x))}}{(N - k_T(x)t)_{(k_T(x))}},$$

where the product is the probability the $k_T(x)$ additional draws yield only bad targets.

Thus T's payoff if it searches in time $t_{Last} - 1$ is $(1 - k_G(x)t_{Last}/N)(L_1 P_{good}(t_{Last} - 1) + L_2(1 - P_{good}(t_{Last} - 1))) - C$ which is compared to the payoff for not searching in time $t_{Last} - 1$ contingent on not having seen a good target yet, $(1 - k_G(x)(t_{Last} - 1)/N)L_2$. If the former is greater T searches in period $t_{Last} - 1$, if the latter, T does not. For completeness, we assume that T attacks if indifferent.

T's algorithm thus is as follows. First T decides whether it would attack or search in period $t_{Last} - 1$ according to the rule in the previous paragraph. The expected payoff from this decision—call this $Q_{t_{Last}-1}$—becomes the relevant payoff for observing no good targets in period $t_{Last} - 2$. The relevant comparison in period $t_{Last} - 2$ is therefore whether $P_{good}(t_{Last} - 2)L_1(1 - k_G(x)(t_{Last} - 1)/N) + (1 - P_{good}(t_{Last} - 2))Q_{t_{Last}-1} - C$ is greater than $(1 - k_G(x)(t_{Last} - 2)/N)L_2$. If it is, T searches in period $t_{Last} - 2$, otherwise T attacks a lesser target. Generalizing to Q_t allows us to continue this train of logic as we work up the game tree. In every period, T compares the same two terms (shifted up a period) and decides whether to attack or to search. We call the earliest period in which T would choose to attack rather than to continue searching $t^*(x)$. Thus, T's stopping rule is to attack whenever a good

[6]The ceiling function equals the next highest integer if its argument is not an integer, or the argument itself otherwise.

target is seen, or in period $t^*(x)$ if no good target is seen before then.

Since G must decide on an information release strategy before the search game takes place, G maximizes an expected utility that consists of the sum of all potential losses from good targets in the first $t^*(x)-1$ periods and the potential loss from both targets in period $t^*(x)$. Each period's potential loss is weighted by the probability of stopping at that period, which is just the probability of no good targets being observed by T in the $t-1$ periods before that multiplied by $P_{good}(t-1)$. So x minimizes G's loss:

$$P_{good}(0)\left(1-\frac{k_G(x)}{N}\right)L_1 \quad + \quad \sum_{t=2}^{t^*(x)-1}\left[\frac{(N-H)_{(k_T(x)(t-1))}}{N_{(k_T(x)(t-1))}}P_{good}(t-1)\left(1-\frac{k_G(x)t}{N}\right)L_1\right] + \frac{(N-H)_{(k_T(x)(t^*(x)-1))}}{N_{(k_T(x)(t^*(x)-1))}}\left[\left(1-\frac{k_G(x)t^*(x)}{N}\right)(P_{good}(t^*(x)-1)L_1+(1-P_{good}(t^*(x)-1))L_2)\right]. \quad (3)$$

As with the simpler stopping rule, for some parameter values the optimal amount of information released by G will be positive. Again, increasing k_G is strictly better for G and increasing k_T is strictly worse, so increasing the rate at which G finds targets relative to T's rate increases the range of parameter values over which it is beneficial to release information. Further, there will still be situations in which G can deter T from searching by releasing sufficient information. We summarize these results in the following proposition.

Proposition 2: Under a complex stopping rule, T attacks upon seeing any good target or upon reaching period $t^*(x)$, whichever comes first. G chooses the level of information that minimizes (3). The range of parameters over which information is released increases as G becomes relatively more efficient at discerning targets than T. If $t^*(x)$ is non-positive, T never searches at all, which is the optimal outcome for G.

2 Target-Specific Information—Lemmas 1 and 2

We begin our analysis with a simple model where resources are held constant and G can release varying levels of target-specific information. Consider first the case where both T and G know all of the targets and $s_i \in \{0,1\}$. G's utility is then $(-L_i + g_i(s_i))\delta_i(s_i) + h_i(s_i)$. Define $Z_i(s_i) \equiv (-L_1 + g_i(s_i))\delta_i(s_i) + h_i(s_i)$. Assume that the sites are ordered such that $(-L_1)\delta_1(0) < (-L_2)\delta_2(0) < \ldots < (-L_N)\delta_N(0)$. Thus T strictly prefers to attack site 1, then site 2, and so on. The assumptions on the functions g and h imply that this same ordering holds for G at $s_i = 0$: $Z_1(0) < Z_2(0) < \ldots < Z_N(0)$.

Begin with site N. The decision as to whether or not to allocate information boils down to a comparison between $(-L_N + g_N(1))\delta_N(1) + h_N(1)$ and $(-L_N)\delta_N(0)$. If the former is greater G releases information; if the latter is greater, G does not. Define the greater of these two Q_N and consider site $N-1$. There are two possibilities here. One, setting $s_{N-1} = 1$ might not change T's ordering on the targets. In this case, the decision at this site is self-contained, and is again a comparison between $(-L_{N-1} + g_{N-1}(1))\delta_{N-1}(1) + h_{N-1}(1)$ and $(-L_{N-1})\delta_{N-1}(0)$. Two, setting $s_{N-1} = 1$ might change T's ordering. Since δ_i is decreasing in s_i, the only ordering change possible is for site $N-1$ to flip with site N. Allocating information to site $N-1$ then causes T to target site N, yielding a payoff of Q_N to G. This is always beneficial for G, as $Q_N \geq (-L_N)\delta_N(0) > (-L_{N-1})\delta_{N-1}(0)$, so whenever T's ordering would shift as a result, it is beneficial for G to release information. This logic can be continued up the chain of sites, and at each site it is either always beneficial to release information, if T's ordering would change, or beneficial whenever $(-L_i + g_i(1))\delta_i(1) + h_i(1) > (-L_i)\delta_i(0)$. Rearranged, this yields $g_i(1) + \frac{h_i(1)}{\delta_i(1)} > -L_i \frac{\delta_i(0)-\delta_i(1)}{\delta_i(1)}$, which provides a sufficient condition on the magnitude of externalities for the release of information for T's preferred target (which will be site 1 unless the ordering has changed upon the release of information).[7] For all other targets—those which T will not attack in equilibrium—the condition is simpler: G should release information whenever $h_i(1) > 0$, as the other factors in the above inequality hold only when a target is attacked.

Lemma 1: If resources are fixed, all players know the existence of all targets, and

[7]Note that this is not a necessary condition: there will be some circumstances—those in which T's ordering changes—where it will be beneficial to release information even when the condition does not hold.

$s_i \in \{0,1\}$, then G releases information about site i whenever $g_i(1) + \frac{h_i(1)}{\delta_i(1)} > -L_i \frac{\delta_i(0)-\delta_i(1)}{\delta_i(1)}$ if site i will be attacked by T in equilibrium, or whenever $h_i(1) > 0$ otherwise.

We now assume that G does not know what T knows. We also eliminate the externality h_i for simplicity, as we want to focus on the role of $P_i^T(s_i)$.[8] G's expected utility from an attack on site i is now $(-L_i + g_i(s_i))\delta_i(s_i) \equiv Z_i(s_i)$. Again assume the same ordering over sites for T: $(-L_1)\delta_1(0) < (-L_2)\delta_2(0) < \ldots < (-L_N)\delta_N(0)$, and again start with site N. As releasing information about N cannot change T's ordering, G releases information if $Z_N(1) > P_N^T(0)Z_N(0)$. Since the Z's are all negative, this condition is the same as $P_N^T(0) > \frac{Z_N(1)}{Z_N(0)}$, putting a lower bound on the prior probability that site N could be attacked.

Now denote the greater (less negative) of $(Z_N(1), P_N^T(0)Z_N(0))$ $P_N^Q Q_N$ and consider site $N-1$. As in Lemma 1, we have two possibilities: allocating information to this site changes T's ordering, or it does not. If it does not, then releasing information ensures that target $N-1$ is always struck (assuming attack modes for the previous $N-2$ sites are not known by T), since T then knows perfectly how to attack it. G allocates information whenever: (a) $Z_{N-1}(1) > P_{N-1}^T(0)Z_{N-1}(0) + (1 - P_{N-1}^T(0))P_N^Q Q_N$. If it does change T's ordering then it is possible that site N is attacked; in this case G releases information whenever: (b) $P_N^Q Q_N + (1-P_N^Q)Z_{N-1}(1) > P_{N-1}^T(0)Z_{N-1}(0) + (1-P_{N-1}^T(0))P_N^Q Q_N$. As $Z_{N-1}(0) < P_N^Q Q_N$ by the assumption on the ordering, condition (a) reduces to $P_N^T(0) > \frac{Z_{N-1}(1) - P_N^Q Q_N}{Z_{N-1}(0) - P_N^Q Q_N}$. Similarly, condition (b) reduces to $P_N^T(0) > \frac{(1-P_N^Q)Z_{N-1}(1)}{Z_{N-1}(0) - P_N^Q Q_N}$. In each case, information is released when the prior probability exceeds a certain threshold. This procedure can be continued up the chain of sites, yielding a lower bound on the prior probability for each site. This leads to Lemma 2.

Lemma 2: If resources are fixed and $s_i \in \{0,1\}$, then G releases information about site i whenever the prior probability $P_i^T(0)$ is sufficiently high.

[8] This is the worst case for information sharing, if these externalities are positive.

3 Resource Allocation—Propositions 3a and 3b

Because it will help explain the interaction between information release and resource allocation, we must first examine the case where G solely allocates resources and all targets are known.[9] Assume without loss of generality that sites are indexed so that $L_1\delta_1(0) > L_2\delta_2(0) > \cdots > L_N\delta_N(0)$ and that $A_{z_1}\delta_{z_1}(0) > A_{z_2}\delta_{z_2}(0) > \cdots > A_{z_N}\delta_{z_N}(0)$, where $z_i \in 1, \ldots, N$. We solve the game using backward induction.

Consider first the zero-sum case, so L_i=A_i. T picks the optimal target to attack, which given the assumed ordering is target 1. Knowing this will happen, G minimizes its expected loss by putting resources into target 1, which is strictly optimal until such time as T becomes indifferent between attacking targets 1 and 2. This occurs after putting r_a resources into defending site 1, with $L_1\delta_1(r_a) = L_2\delta_2(0)$. At this point G must defend both targets 1 and 2 and thus splits resources so as to keep T indifferent between attacking either of these two targets. Again, this continues until T is now indifferent between attacking sites 1, 2, or 3, at which point G must begin defending target 3 as well. This logic continues as new sites become desirable targets for T, until such time as G runs out of resources and the chain ends. The number of sites—i_R—to which G makes T indifferent is determined by the resource total R, the L_i, and the functions $\delta_i(r_i)$. As it is indifferent between these sites, we assume that T attacks one at random in equilibrium. Thus we have:

Proposition 3a: Assume $L_i = A_i$ for all i. In equilibrium, government utilizes all resources so as to make the terrorists indifferent between attacking the first i_R sites, where i_R depends on R, the L_i, and the functions $\delta_i(r_i)$. Terrorists choose one of these first i_R sites to attack at random.

Now consider the case where G and T may value targets differently. We assume a minimal resource level in this case, so that, $\forall i$, $r_i \geq r_{min} > 0$.[10] r_{min} can be arbitrarily small, but must be strictly greater than zero. Here T's ordering over targets is what matters, which

[9] As noted in the text, Proposition 3a essentially restates the analysis in Powell (2007a).

[10] For federal resources in the United States this amount is a dollar, the smallest denomination that appears

leads to a complication: making T indifferent between two targets might introduce into the indifference set a target that yields a strictly worse expected outcome for G than the other targets in the set did.[11] To avoid this, G adds additional r_{min} resources to this new, more valued target, which removes it from T's indifference set, causing T not to attack that target in equilibrium.[12] Further resource additions to this and other targets are made so as to maintain this relationship—any targets that G values more than T's ordering are kept out of T's indifference set. Other than this complication, the game is the same. As the extra resources added to the more costly targets are small, G does weakly better when the game is not zero-sum, since T is now choosing among targets that matter weakly less to G.[13] So:

Proposition 3b: In equilibrium, government utilizes all resources so as to make the terrorists indifferent between attacking a subset of the first i_R sites, where this subset depends on R, the L_i and A_i, and the functions δ_i. Terrorists choose one of this subset to attack at random, and G does weakly better in terms of expected utility than in the zero-sum case.

4 Interaction of Information and Resources—Proposition 4

We now turn to the case where G can both release target-specific information and allocate resources. Since a zero-sum game represents the worst-case scenario for G, as discussed at the end of the last section, we will also henceforth assume that $L_i = A_i$ for all i; if it pays to

in spending bills.

[11]The existence of a minimum resource unit renders the procedure of allocating to achieve indifference somewhat problematic. We assume that if for some cases the discreteness of resources renders indifference impossible to achieve G makes T indifferent between the sites that produce the least loss for G.

[12]A greater amount of resources can be added if G does not believe that T is quite this rational.

[13]Note that having a minimal resource unit r_{min} is necessary for the result. If resources were continuous the equilibrium from Powell (2007a) would result in which G allocates resources to minimize T's maximum gain and T attacks the target from his indifferent set that does the least damage to G.

release information in the zero-sum case, it pays in the case of differing valuations as well. A strategy for G is an $N \times 2$-tuple that describes a resource allocation, r_i, and an information policy, s_i, for each target i, with $s_i \in [0,1]$ and $\sum r_i \leq R$. As discussed in the text, we are considering the zero-sum case for this and subsequent propositions, so that $L_i = A_i \; \forall \; i$. Further, henceforth we will be considering only the case of a negative backlash if a site is struck, so that $g_i(s_i) \leq 0 \; \forall \; i, s_i$. Though we do this for simplicity, this assumption will make it more difficult to show that information will be released in equilibrium, as per Lemmas 1 and 2, and so we are again treating the worst-case scenario for openness.

We begin with a lemma that will be useful throughout the analysis of Proposition 4.

Lemma 3: G's expected utility at each site displays increasing differences in r_i, s_i.

Proof: To see this, let $F_i(s_i) \equiv P_i^T(s_i)(-L_i + g_i(s_i))$; then

$$\frac{\partial^2}{\partial r_i \partial s_i}\left[P_i^T(s_i)(-L_i + g_i(s_i))\delta_i(r_i, s_i) + h_i(s_i)\right] = \frac{\partial F_i(s_i)}{\partial s_i}\frac{\partial \delta_i(r_i, s_i)}{\partial r_i} + F_i(s_i)\frac{\partial^2 \delta_i(r_i, s_i)}{\partial r_i \partial s_i}.$$

Note that $\frac{\partial F_i(s_i)}{\partial s_i} = (-L_i + g_i(s_i))P_i^{T\prime}(s_i) + g_i(s_i)' P_i^T(s_i) < 0$, $\delta(r_i, s_i)$ and $F_i(s_i)$ are decreasing in r_i and s_i respectively, $F_i \leq 0$, and $-\delta_i(r_i, s_i)$ displays increasing differences. Therefore each of the two terms on the right hand side is positive, proving the lemma.

G does not know which targets T knows, so constructing an equilibrium is more complicated than in Propositions 3a and 3b. We begin by writing G's utility unconditional on T's actions, which is $U^G(\mathbf{r}, \mathbf{s}) = \sum h_i(s_i) + \mathbf{1}_i Z_i(s_i, r_i)$, where $\mathbf{1}_i$ is an indicator function that takes the value 1 when T attacks target 1 and 0 otherwise, and $Z_i(s_i, r_i) \equiv (-L_i + g_i(s_i))\delta_i(r_i, s_i)$. As T attacks the site that yields it the highest value in expectation, G must maximize its expected utility conditional on T's choice of where to attack. Assuming that sites are ordered such that $L_1\delta_1(s_1^*(0), 0) > L_2\delta_2(s_2^*(0), 0) > \ldots > L_N\delta_N(s_N^*(0), 0)$, where $s_i^*(0)$ is the optimal level of information released at resource level 0, this expected utility is:

$$E[U^G(\mathbf{r},\mathbf{s})|T] = \sum_i h_i(s_i) + (P_1(s_1))Z_1(s_1,r_1) + (1-P_1(s_1))P_2(s_2)Z_2(s_2,r_2) + \ldots + \prod_{i=1}^{N-1}(1-P_i(s_i))P_N(s_N)Z_N(s_N,r_N). \quad (4)$$

We will denote this EU^G. By Lemma 3, $Z_i(r_i, s_i)$ has increasing differences for all i. Accordingly, since only $Z_i(r_i, s_i)$ depends on r_i, EU^G has increasing differences in the pair $(r_i, s_i)\ \forall\ i$. The equilibrium $N \times 2$-tuple can be found by maximizing (4) under the constraint $\sum r_i \leq R$. As the maximization is over a compact subset of $\Re^{2N-1}$, an equilibrium exists as long as EU^G is bounded and continuous in a region of strictly positive measure surrounding the maximum of the function.

This latter condition creates complexity. Because T's ordering over sites can change as G increases information and resources at a site, G's expected utility function may not be continuous over the entire domain. To show this, consider a typical algorithm for solving this maximization problem that is very similar to the proof of Proposition 4a. First, G adds resources to site 1, which increases G's utility (by assumption) and also weakly increases the optimal level of information released at site 1 (due to increasing differences). Increasing s_1 has two distinct effects: (1) it decreases G's expected loss at site 1, even taking into account the increase in the chance that T knows of site 1; and (2) it decreases the likelihood that T attacks any other site. The net result is that it becomes even more beneficial for G to invest in site 1's defense. However, at some point this increased investment may cause T's ordering to switch so that now site 2 is preferred to site 1. As T can only target one site, G's conditional expected utility changes form: all 1 and 2 subscripts switch places. This can cause EU^G to be discontinuous at the point of the switch.

However, an equilibrium still exists. For each ordering over targets for T, G's expected utility function is bounded and continuous on a closed and bounded domain. Thus, for every ordering there exists a maximum. As the number of possible orderings is finite ($N!$) and every allocation of s_i and r_i results in one of these orderings, the maximal expected utilities for each of these orderings can themselves be ordered. There must exist a maximum of this ordered list of utilities (though there can be more than one), and the allocation that results in this maximum will be the equilibrium of the game.

This shows that there is an equilibrium for any set of functions $g_i(\cdot)$, $h_i(\cdot)$, $P_i^T(\cdot)$, and $\delta_i(\cdot,\cdot)$ and any R. To prove the second part of the proposition, note that since EU^G has increasing differences, increasing the resources at a site must weakly increase the equilibrium level of information released at that site, implying that the sum of all information released, $\sum s_i$ is weakly increasing in R, as all available resources get assigned to a site in equilibrium.

Proposition 4: Assume $L_i = A_i$ for all i. For any set of functions $g_i(\cdot)$, $h_i(\cdot)$, $P_i^T(\cdot)$, and $\delta_i(\cdot,\cdot)$ and any R there exists an equilibrium of the game in which G allocates resources and releases target-specific information across known sites and T attacks the site it knows that provides it the highest expected utility. The sum total of all information released, $\sum s_i$, is weakly increasing in the total amount of resources available, R.

5 General and Target-Specific Information—Proposition 5

The final model we treat in the paper uses all of the assumptions listed at the beginning of this document. Government begins the game by releasing a level of general information, $x \in [0,1]$. This information influences the number of targets that G discovers, and we denote the set of realized targets Y_j. The set of all Y_j, which we will call Y, is the set of all possible arrangements of known and unknown targets. As x determines $P_i^G(x)$ for all targets i, x also determines the probability that any set of targets Y_j is realized.[14] For example, the probability of realizing the set in which the first three targets are known by G and the last

[14]Note that we have abstracted away from the two search games discussed previously. The game we present here is equivalent to one in which first both players search, then G allocates resources and releases target-specific information, and finally T attacks.

$N-3$ are not is $P_1^G(x)P_2^G(x)P_3^G(x)(1-P_4^G(x))\dots(1-P_N^G(x))$. We denote each of these probabilities by $P_j^{G,Y}(x)$.

Once G realizes a set of targets Y_j, the game follows exactly the form discussed prior to Proposition 4, but several functions now depend on x. For any given Y_j, therefore, there exist equilibrium strategies for G and T, given by Proposition 4, that fully specify $s_i^*(x)$ and $r_i^*(x)$. Denote the expected utility to G from this equilibrium $U_j^{G,Y}(x)$, which is equal to the equilibrium utility from Proposition 4 (for a given value of x) plus the global externality $h(x)$. The equilibrium of this full game thus can be solved via backward induction: Proposition 4 dictates the outcome of the game once G has chosen x^*. To specify G's full strategy we must determine the optimal amount of general information to release. The optimal x^* is given by:

$$argmax_x \sum_{Y_j \in Y} P_j^{G,Y}(x) U_j^{G,Y}(x). \tag{5}$$

The maximum is taken over a compact set, so we just need continuity near the optimum to prove an equilibrium. For any Y_j, $P_j^{G,Y}(x)$ is continuous in x and the sum adds no discontinuity. Thus we need only show that $U_j^{G,Y}(x)$ is continuous near the optimum x^*. The same logic as in the proof of Proposition 4 applies here as well: equation (4) above (with the functional dependence on x added) is continuous except for the points at which T's ordering over targets changes. Each x induces a distribution over Y_j, each element of which has an optimal $s_i^*(x)$ and $r_i^*(x)$ and so a specific associated ordering. Each ordering has a maximum, and each x gets mapped to one of these orderings. As there are a finite set of orderings, we can choose a global maximum from this set of local maxima, and so the game has an equilibrium. Further, the equilibrium level of general information is weakly increasing in the total amount of resources available, R, due to supermodularity. This gives us our final result:

Proposition 5: In the full game, G follows the algorithm given above and T attacks the target seen that yields the greatest expected utility. The amount of information released by

G in equilibrium is weakly increasing in the total amount of resources, R.