

Choosing What to Protect

Abstract: We study a strategic model in which a defender must allocate defensive resources to a collection of locations, and an attacker must choose a location to attack. The defender does not know the attacker's preferences, while the attacker observes the defender's resource allocation. The defender's problem gives rise to negative externalities, in the sense that increasing the resources allocated to one location increases the likelihood of an attack at other locations. In equilibrium, the defender exploits these externalities to manipulate the attacker's behavior, sometimes optimally leaving a location undefended, and sometimes preferring a higher vulnerability at a particular location even if a lower risk could be achieved at zero cost. Key results of our model are as follows: (1) the defender prefers to allocate resources in a centralized (rather than decentralized) manner; (2) as the number of locations to be defended grows, the defender can cost-effectively reduce the probability of a successful attack only if the number of *valuable* targets is bounded; (3) the optimal allocation of resources can be non-monotonic in the relative value of the attacker's outside option; and (4) the defender prefers her defensive allocation to be public rather than secret.

Key words: Security; game theory; uncertainty; resource allocation; externalities

This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number EMW-2004-GR-0112, the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number DAAD19-01-1-0502, and the National Science Foundation under grant number SES-0241506. Any opinions, findings, and conclusions or recommendations expressed herein are those of the author and do not necessarily reflect the views of the sponsors. This paper summarizes research reported in Bier et al.⁵ I would like to thank the coauthors of that paper, Santiago Oliveros and Larry Samuelson of the Department of Economics at the University of Wisconsin-Madison, for their contributions to that research, and their comments on this manuscript. Naturally, any errors in this paper are my own responsibility, not theirs.

1 Introduction

Past game-theoretic models of security investment have generally advised defenders to put all their eggs in one basket (or a small number of baskets), corresponding to those assets believed to be most vulnerable,¹ most valuable,² or most attractive to attackers.³ This is obviously unrealistic in practice; real-world decision makers will likely want to hedge their bets in case they have guessed wrong about which assets are most attractive to potential attackers. For example, nobody would recommend that the U.S. invest all its security resources in defense from smallpox, no matter how devastating one hypothesizes that a smallpox attack might be, in case potential attackers do not have access to smallpox stocks, or are not willing to risk the “blowback” of smallpox epidemics in their own communities. Thus, taking into account the defender’s uncertainty about attacker goals, valuations, and constraints would seem to be central to achieving a good security policy. In fact, Banks⁴ has suggested that intelligence may actually be more cost-effective in some situations than defending against particular attack scenarios; see also Brookings.² Uncertainty about attacker goals and values plays an important role in contemporary discussions of terrorism, and hence is important to capture in a model of the optimal defensive allocation.

Major⁵ and Woo⁶⁻⁷ achieve the more realistic result of hedging at optimality. However, they achieve this by the unrealistic assumption that attackers can observe “the marginal effectiveness of defense” at each target (which even defenders may not know accurately for defenses that have not been evaluated), but not which defensive investments have actually been implemented. Moreover, the models of Major and Woo do not explicitly consider the defender’s uncertainty about the attacker’s asset valuations, and hence do not allow one to explore how optimal asset valuations might vary in the face of greater or lesser uncertainties. Presumably, the extent of hedging in defensive investments should depend in some way on the extent of the defender’s uncertainty about likely attack strategies.

Moreover, Major⁵ assumes that defenders and attackers have exactly the same valuations for potential targets, or in other words that security is a zero-sum game, perhaps because the attacker cares only about inflicting harm on the defender. This again is unrealistic. In principle, the value to the attacker of successfully attacking a given target may depend not only on the damage inflicted on the

defender, but also on the propaganda value of the target, the cost or difficulty of mounting the attack, and other factors that the defender may not even fully comprehend. Woo⁷ has observed that “If a strike against America is to be inspirational [to al-Qaeda], the target should be recognizable in the Middle East”; thus, for example, attacks against iconic targets such as the Statue of Liberty or the Sleeping Beauty Castle at Disneyland may be disproportionately attractive to attackers relative to the economic damage and loss of life that they would cause. Similarly, in the context of computer security, Besnard and Arief⁸ note that “attackers may care less about costs than legitimate users do.”

In this paper, we describe the results of a model in which attacker and defender valuations for any given target are allowed to differ. In this model, the defender must allocate defensive resources to a collection of assets, and an attacker must choose a single asset to attack. The proposed model assumes that attackers can observe defensive investments perfectly (which is conservative, but perhaps not overly so for some types of investments—e.g., costly capital improvements), but that defenders are uncertain about the attractiveness of each possible target to the attackers. This last assumption is reasonable, in light of the fact that lack of knowledge about attacker values, goals, and motivations is precisely one of the reasons for gathering intelligence about potential attackers. The mathematical derivations involved in identifying equilibrium solution strategies for both attackers and defenders are described in Ref. [9]. This paper highlights the policy implications of that model.

We begin by presenting a simple two-asset model, in which only a single defender is responsible for security of both assets (both with and without defender budget constraints). We then discuss the policy implications of generalizations to this basic model, to account for situations with much larger numbers of assets, decentralized defenses (in which a different asset owner is responsible for defensive investment in each asset), and attacker opportunity costs. Finally, we briefly explore the issue of defender secrecy and even deception, and when these might be advantageous in practice.

Section 2 presents the simplest version of the model, with only two assets to defend and a centralized defender. Section 3 discusses how the results change when defensive decisions are decentralized. Section 4 extends the results to (arbitrarily) large numbers of assets, and Section 5

explores the effects of attacker opportunity costs. Section 6 considers several extensions of the model that allow us to develop policy implications. Finally, Section 7 gives some conclusions.

2 A Simple Two-Asset Model

We assume that the attacker can choose to attack one (and only one) of asset 1 or asset 2 (perhaps because an attack on one location would exhaust the attacker's resources, or would lead to the attacker being detected and disabled). These assets may represent different locations against which a given type of attack may be launched (such as different cities¹⁰), but the model can equally well be used to explore the relative merits of defenses against different types of attacks against a single asset (such as nuclear versus biological attacks against a given city).

An attack may be either a success or failure. The attacker receives a payoff of a_i in the event of a successful attack on asset i , and 0 in the event of an unsuccessful attack. The defender experiences a loss of d_i from a successful attack on location i , and no loss from an unsuccessful attack. We assume the attacker knows the defender's valuations of the two assets, d_1 and d_2 . The attacker's preferences (a_1, a_2) are assumed to be known by the attacker but not the defender. We thus have a game of incomplete information. The defender's uncertainty about the attacker preferences is assumed to be described by a cumulative distribution function F , which is assumed to satisfy certain reasonable properties.

In particular, the model assumes that F gives non-zero probability density to all possible attacker valuations in a set A (which itself must satisfy certain reasonable properties). Of course, the attacker may well desire to inflict losses on the defender, and hence F may attach high probability to values close to d_1 and d_2 . The model allows this, as long as there is also some non-zero probability of any other attacker asset valuation in A . Thus, this formulation allows the attacker's preferences to be linked to the defender's valuations d_1 and d_2 , as long as the defender valuations are not the only factors considered by the attacker.

Note, by the way, that while we assume the attacker observes the defender's allocation of defensive resources and adapts the choice of attack strategy accordingly, the attacker's preferences (i.e., asset valuations) are not allowed to depend on the observed defensive allocations. Thus, for example, the

model in Ref. [9] does not capture the behavior of computer hackers who delight in attacking the most secure systems just for the challenge of doing so. Similarly, we exclude the possibility that the attacker may not know the defensive valuations perfectly, and attempt to infer those valuations by observing the defender's resource allocations. This will become important later, in discussing the potential benefits of defender secrecy and/or deception.

The defender may attempt to minimize the expected loss from an attack given a fixed budget C of defensive resources to allocate between the two assets, or, in the unconstrained version of the problem, may minimize the sum of the expected loss from an attack plus the cost of any defensive investments. The probability p_i that an attack against asset i will succeed is assumed to be a continuous, convex function of the level of defensive investment in that asset. Continuity and convexity again become important in the discussion of defender secrecy and deception. We also assume that neither asset can be perfectly defended at finite cost; in other words, an attack on asset i will always have some positive (though perhaps small) probability of success, no matter what resources have been allocated to its defense. Other, more technical assumptions are discussed in Ref. [9].

The defender is assumed to choose $(p_1; p_2)$ without knowing the attacker's valuations, guided by the prior distribution F over those valuations. The attacker observes the defender's allocation, and then chooses which asset to attack, guided by the actual values of the attacker valuations, so as to maximize the expected payoff from the attack, $a_i p_i$. It can be shown that there exists an equilibrium solution to this attacker-defender game. In fact, due to the convexity of the defender's strategy set and the assumptions about the distribution F and the success probabilities p_i , this equilibrium will be unique and involve pure strategies on the part of both the attacker and the defender. In other words, for this formulation of the problem, there is no need to examine so-called "mixed" or randomized strategies, in which the attacker randomly chooses which asset to attack, and/or the defender randomly chooses the level of resources to invest in defense of each asset. (Relaxations of some of the above assumptions will be discussed later.)

Before going on to explore more interesting implications and extensions of this basic model, we first describe its general behavior, to confirm that it captures important aspects of security decision

making. For example, if the value of asset 1 to the defender increases, then in general the model will advise assigning more defensive resources to asset 1, at the expense of asset 2. As a result, the success probability of an attack against asset 1 will decrease, the success probability of an attack at location 2 will increase, and the attacker will become more likely to attack asset 2 (all else held equal). This effect is stronger in the budget-constrained problem than in the unconstrained version, since in the constrained version p_1 can be decreased only by increasing p_2 . However, the same general behavior also occurs in the unconstrained version. Moreover, if both assets receive non-zero defensive investment at optimality, then the ratio of the marginal expected losses from attacks against assets 1 and 2 must equal the ratio of the marginal costs of additional investments in those assets (as we might expect).

In the constrained version of the problem, as the total budget for defense grows, the amount allocated to each asset will in general increase, and the success probabilities of attacks on both targets will decline (provided that the budget is large enough to justify investing in the security of both targets). Moreover, at optimality, the defender will never allow resources to go unused in the constrained problem.

Since the defender is uncertain about the attacker valuations a_i (and we have not yet modeled the attacker as having an opportunity cost), the defender's probability that asset i will be attacked is given by the probability (under the distribution F) that $a_1 p_1$ will exceed $a_2 p_2$. Holding the distribution F constant, the defender will put higher probability on asset 1 being attacked as the ratio p_1/p_2 grows, since larger values of p_1 will tend to make asset 1 more attractive to the attacker. Conversely, increasing the resources allocated to defending asset 1 (i.e., reducing p_1) will reduce the defender's probability of an attack on that asset. In fact, in the absence of attacker opportunity costs, the defender's probability that there will be an attack on asset 1 depends only on the ratio p_1/p_2 , not on the success probabilities p_1 and p_2 individually.

Similarly, if the defender's distribution of the attacker's valuations changes so that attacks on asset 1 are believed to be more likely, the defender's optimal response is to shift defensive resources to that asset. However, the exact relationship is complicated, since it depends not on the defender's expected values for a_1 and a_2 , but rather on the probability that a_1 is greater than a_2 . Thus, for example, reduced uncertainty about the attacker's valuation of a particular asset can in principle lead to either

greater or lesser defensive investment in that asset, depending on whether the uncertainty is resolved in favor of a lower or higher estimate of the asset's value. (From the attacker's viewpoint, the optimal choice of attack strategies is actually a deterministic function of the known a_i and the observed p_i .)

Even this simplistic formulation of the game, with only two components, has interesting and counterintuitive implications. For example, changing the success probability of an attack on asset i makes that asset more secure, but also increases the likelihood (from the defender's point of view) that the other asset will be attacked. For this reason, it is possible for a given asset to be "too secure"; the success probability of an attack on that asset may be so low as to deflect too much risk onto the other asset. In other words, it might sometimes be desirable to increase one of the p_i , even if there is no compensating decrease in the success probability of an attack on the other asset. Thus, if the defender cannot improve security of a poorly defended asset, it might be better off throwing away resources than further defending a highly secure asset (even in the constrained problem). For example, when p_1 is extremely small relative to p_2 , the attacker will likely attack asset 2, which might be more valuable to the defender than asset 1. Increasing p_1 (by reducing the level of defensive investment in asset 1) would divert some of the attack probability from asset 2 to asset 1. If the probability of success of an attack against asset 1 is not allowed to get too large, this may be beneficial to the defender, since the expected loss from diverting attacks to asset 1 will be small enough to ensure a net reduction in the defender's total expected loss.

Note that this is a general phenomenon in the model. In other words, there will always be some level of investment in any asset (no matter how important) that is "too much" in comparison to the defensive investment in other assets; the result does not rely on the assumption that either the cost (or the marginal cost) of perfect protection is large. More to the point, however, the result shows that spending too much on defense of assets that are not highly valuable hurts the defender in two ways—not only by wasting resources on defense of assets that are unlikely to be attacked in any case, but also by increasing the likelihood of a more valuable asset being attacked! Similarly, the defender would not want to reduce the success probability of an attack against any one asset to zero unless this could be done for both assets.

It can also often be optimal to leave one or more assets undefended. Even if marginal reductions of the success probability of an attack are costless, one asset may be left undefended at equilibrium in the constrained model if its value is sufficiently small. (In the unconstrained model, it may be optimal to leave both assets undefended if defending them costs more than it is worth.) This is true even if an attack on an undefended asset is guaranteed to succeed. Intuitively, if the values of the two assets are sufficiently different, then the less valuable asset may be quite unlikely to be attacked, in which case defensive resources should be allocated only to the more valuable asset. The larger the defender's budget, the more likely it is that both assets will be defended. Thus, it is more likely optimal to leave assets undefended when the defender is budget constrained, and the values of the assets differ widely.

Of course, this is exactly the situation in which we find ourselves in the real world. This suggests that it may well be optimal not to invest in much additional protection of smaller states, etc. By contrast, recent security funding has been subject to the criticism that "states like Wyoming...get more per capita in terrorism grants than New York".¹¹ Stephen Flynn, director of the Hart-Rudman task force on homeland security, has stated: "At the end of the day, blowing off New York and L.A. so that you can make sure Wyoming is safe just makes no sense".¹² However, skewed funding priorities of this type have been mandated by the Congress in a formula that "guaranteed each state 0.75% of the total amount appropriated to DHS for state terrorism preparedness grants...[with] 40 percent of the total pot of money being divided up equally...regardless of size, risk, or need".¹³ Zycher¹⁴ coins the concept of "efficient" pork to reflect situations in which such subsidies are necessary to create "sufficient political support" to undertake more urgent and cost-effective investments. However, Brunet¹⁵ has noted that the 0.75% minimum for each state "is larger than most minimum amounts found in existing federal grant programs," casting doubt on the idea that current funding levels for small states are either necessary or appropriate.

The same type of approach sometimes occurs at the state level as well. Thus, the House Select Committee on Homeland Security¹⁶ states that "many States follow the Federal government's example by providing a base amount to each county, [sometimes] with an additional amount based on population. In fact, almost one-third of our Nation's States distributed their Federal first responder funds...by formulas

that did not account for either need or risk (other than population).” For example, Carafano¹⁷ reports that “California distributes its federal grants in base-amounts of \$5,000 to each county.” As a result of this type of policy, a government official from one rural county in the state of Washington stated that “We’re getting stuff we won’t use. This equipment could have gone to Seattle where the real threat is”.¹⁶ The Urban Area Security Initiative was initially intended to address this problem by funneling resources to the largest and highest-risk cities. However, within about a year, the list of cities to receive funding grew from seven to 80, including several with populations of less than half a million people.¹³

Excessive allocations of security investments to protection of relatively low-risk targets not only divert defensive resources from more important targets, but can actually be harmful to overall security, since such sub-optimal investment strategies can deflect attacks to alternative targets that were initially less attractive to the attackers, but are also more damaging to the defenders. For example, making particular targets less vulnerable to attack could lead terrorists to adopt attack strategies that are more costly or difficult for them to implement, or would yield less publicity benefit to the attackers, but are also more lethal. This could be important in light of observed past substitution effects.¹⁸ In fact, occasional small or moderate attacks could be a sign of a successful defensive strategy, not a failed one!

3 Decentralized Defensive Decisions

Should strategic defensive decisions be centralized or decentralized? In the absence of diseconomies of scale or other costs of centralization, economic theory would normally predict that decentralizing the choice of p_1 and p_2 (so that a different decision maker chooses the level of defensive investment in each asset) cannot make things better. We investigate this here, focusing on the symmetric unconstrained case.

In particular, rather than a single centralized decision maker (corresponding, for example, to a federal or state government), we assume that there are two separate defenders, where defender i chooses p_i and suffers the consequence d_i if asset i is successfully attacked. We can view this situation as involving a game between the two defenders, with each defender choosing its level of defensive investment to minimize its own payoff, taking into account the other defender’s choice and the attacker’s

optimal behavior. Ref. [9] shows that the equilibrium for a centralized decision maker in the game between an attacker and a single defender will in general involve higher success probabilities of attack (and lower defensive investments) than the equilibrium of the decentralized game with two defenders. This is because of the negative externality between different assets, with a decrease in the success probability of an attack on one asset making it more likely that the other asset will be attacked, and inducing the defender of that asset to redouble its defensive efforts. The result is a “security race” that culminates in inefficiently excessive defensive measures. Thus, the decentralized game will in general yield a smaller total equilibrium payoff to the defenders, with both agents in the decentralized game having stronger incentives to invest in security than in the centralized case (although even decentralized decision makers may choose to leave both assets undefended if the assets are low in value, in which case the total equilibrium payoffs will be the same as in the centralized game).

It is tempting to view the lower success probabilities provided by the solution to the decentralized game as an advantage. However, these levels are inefficiently low. The centralized defender achieves a higher level of utility (or, equivalently, a lower overall loss) than the sum of the utilities achieved by the decentralized defenders. This is because decentralized decision-makers end up investing not only when doing so reduces the total overall societal loss (taking into account both cost of investment and expected damage from attacks), but also in some cases when their investment merely deflects some of the attack risk to another defender. This is obviously not a net societal benefit, but is still beneficial to the defender who makes the investment (in a reverse “tragedy of the commons”).

One need not assume that the decentralized defenders explicitly want to deflect risk onto each other, only that they care (substantially) less about each other’s losses than about their own. Thus, for example, individual communities in the same state may not care whether their defensive investments increase the risks to neighboring communities, while the governor of the state might take such effects into account in making statewide decisions about security investments. Note that the inefficiencies of decentralized decision making may appear in a wide variety of situations. For instance, the measures undertaken by the U.S. Postal Service to make the mail system more secure may well have deflected

some risk onto private mail carriers such as Federal Express; similarly, measures undertaken to make the aviation system more secure may have deflected some risk onto other modes of transportation, etc.

Of course, there is no reason a priori to expect a private decision maker (such as Federal Express), or even a quasi-private agency (like the Postal Service), to make decisions that are socially optimum (rather than just in the interests of that particular organization). However, even a nominally centralized decision maker (e.g., the federal government) may in fact operate in a decentralized manner. For example, individual federal agencies may take into account only costs and benefits within their own areas of responsibility (such as aviation), and not consider the implications of their actions for other sectors of society (e.g., other modes of transportation).

Moreover, due to the imperfect nature of the political process, even decisions made by a single agency may end up looking like decentralized rather than centralized solutions. Thus, small states with powerful congressional representation can end up receiving nominally centralized federal security investments that are not proportionate to the risks they actually face.¹¹ For example, Ripley¹⁹ observes that “When asked about relative risks...officials [of small states] talk about relative worth and the right of their citizens to get the same kind of protection that they are afforded in other places in the country.” Along the same lines, de Rugy¹³ points out that “When first responder programs are funded at the federal [rather than state] level, a Congressman from Wyoming has no incentive [for] admitting that his state is not a likely target or that if it ever were a target, the level of damages would be limited.”

In addition, private companies that would not find it in their interests to invest in their own security may nonetheless find it advantageous to lobby for public investment in security. Finally, solution providers can lobby to have proprietary technologies adopted or recommended for use, even when those solutions are not effective (or cost-effective).¹⁹⁻²⁰ These perspectives are summed up in the observation that “pressure groups—e.g. first responders, state officials and/or specific industries like the airline industry—may have an incentive to lobby lawmakers to try to grab a bigger share of the funding allocated to homeland security programs and /or to transfer their responsibilities to the federal government”.¹³

3.1 Example: Public versus Private Investment in the Construction Industry

It has been argued that terrorism insurance may deter building owners from implementing supposedly cost-effective security improvements. For example, Orszag²¹ cites the argument that “Firms and individuals who have insurance against terrorism would appear to lack incentives to take appropriate precautions against an attack.” While it is true that provision of insurance can reduce the incentives for self-protection (so-called “moral hazard”), the model suggests that in the absence of insurance, the incentives faced by individual building owners may well result in excessive security investment relative to the social optimum, since many types of security measures may simply deflect risk to other buildings. Since insurance agencies are exposed to a larger fraction and spectrum of total terrorism losses than most individual building owners, insurance may serve to make the benefits of potential security improvements to building owners more closely approximate the net societal benefit of such improvements. In this case, insurance might help by “internalizing some of the externalities and interdependencies associated with terrorism risk”²² and thereby preventing the type of over-investment that might lead to a “race to the top.”

For technologies that are costly, and could therefore be deployed only in selected buildings, the cost-effectiveness of investing in such technologies would depend strongly on attacker goals, motivations, and constraints. If attackers are interested in only a few “signature” buildings—for example, icons of economic domination (in the case of al-Qaeda) or environmental evil (in the case of hypothetical environmental terrorists)—then it may be cost-effective from a societal point of view to defend those few buildings, and reduce or eliminate that particular threat. This would be a strong argument for public funding in the development (and possibly deployment) of such technologies.

However, if terrorists are willing to adjust their choice of strategy to target less valuable buildings that are also less well defended, then defending a few large “signature” buildings would yield a private benefit to the owners and occupants of those buildings, but little or no societal or public benefit. For example, consider a hypothetical technology that could make buildings invulnerable to terrorist attack. Such a technology might be sufficiently costly that its use would be limited—e.g., to buildings of more than 100 stories. If attackers simply shifted their targeting strategies to attack buildings of only 99 stories,

then implementation of the technology would reduce total societal risk by only a small percentage. This would argue against public funding of such technologies—not only their implementation or deployment, but possibly also their development. Thus, de Rugy has suggested that the private sector should be responsible for the security of “skyscrapers and individual houses”.¹³ Owners of large “signature” buildings might nonetheless have an incentive to push for public funding of security improvements, however, even if they would yield primarily a private benefit, with little or no societal benefit.

3.2 Other Possible Effects of Decentralization

The above discussion has focused on cases with negative externalities, where defensive investment by one entity increases the risk to others. However, there are also cases with positive externalities (e.g., investment in computer security, which may prevent viruses spreading to other computer users).²³ Such situations create incentives for some potential victims to “free ride” on defensive investment by others, and thus lead to inefficiently low investments in defense. Arce and Sandler²⁴ show that treaties (i.e., binding agreements) may solve the free-riding problem, while Lakdawalla and Zanjani²⁵ show that insurance can be used to coordinate the behavior of defenders facing incentives to free ride.

Keohane and Zeckhauser²⁶ consider both positive and negative externalities among defenders. In particular, they note that centralized actions to reduce risk may be of little benefit in the context of decentralized risk-management decisions by members of the public. For example, as a particular location (e.g., New York City) becomes safer due to government investment, the number of people it attracts will increase, making it more attractive to potential attackers. According to Keohane and Zeckhauser, under certain conditions, “exposure will increase...so that everyone will be exactly as badly off, in expectation, as they were before the government’s action...Government measures that are successful in averting attacks or ameliorating their effects...may nevertheless not improve expected welfare”.²⁶ While the assumptions of their model are perhaps idealized and unrealistic, the same phenomenon occurring to a lesser extent could still diminish if not eliminate the benefits of government investment in security.

4 Large Numbers of Assets

One of the more vexing aspects of any proposed response to a possible terrorist attack is the sheer number of potential targets. The simple two-target model thus appears to neglect some important aspects of the problem. Therefore, in this section we examine a model with an arbitrary number of assets, N . We focus on the unconstrained version of the problem (in which defenders wish to minimize the sum of expected attack losses plus defensive investments), and pay particular attention to the case of large N .

4.1 Symmetric Assets

We initially assume that all assets are equally valuable to the defender, equally costly to defend, and believed to be equally valuable to the attacker (letting the a_i be independent, with identical marginal distributions). (We relax some of these in the next subsection.) Under these assumptions, the N assets are essentially identical from the viewpoint of a defender. Note that the various assets need not (and in general will not) be identical to the attacker, since the actual values of the a_i will typically differ; we assume only that any defenders do not know which assets are more valuable to the attacker.

For a single centralized defender, as the number of (symmetric) defender assets gets large, the optimal defense in the unconstrained version of the problem is not to defend any of them; all targets are left undefended in equilibrium. The intuition behind this result is that since the defender cannot distinguish the various targets, the defender's probability that any particular asset will be attacked is simply given by $1/N$. As the number of targets N grows, the probability that any given target is attacked will converge to zero, while the cost of defending that target will presumably be roughly constant in the number of targets. In the limit, it is simply not worth attempting to defend all N targets (or, given their symmetry, any of them). (Note also that under some reasonable conditions, this limiting behavior can be reached even for finite values of N , not only in the limit as N gets large.)

We now assume that the N assets are identical, but are owned by N decentralized defenders, each of which wishes to minimize the sum of expected attack losses plus defensive investment associated with its own asset. As in the case with only two assets, we still find (not surprisingly) that the decentralized

defenders will use more defensive resources (in total) than a single centralized defender. This reflects the fact that the decentralized defenders are assumed to ignore the negative externalities they impose on other defenders when deciding whether to increase defensive investments.

In the limit as N becomes arbitrarily large, the decentralized defenders know that their defensive resources will almost certainly not be useful in equilibrium, since any one defender is vanishingly unlikely to be attacked. However, unlike in the centralized case, they may still sometimes choose positive allocations of defensive resources as N becomes large. The ability to shift the attacker toward other assets makes it valuable for a decentralized defender to invest in defense, even when there are so many identical assets that any individual one is almost certain not to be attacked. In fact, the overall cost of decentralized defensive investments can grow without bound as N increases, leading to results that are arbitrarily far from the social optimum of no defensive investment for large N .

4.2 Asymmetric Assets

The case of symmetric assets is obviously not sufficiently general to be useful in the real world. In practice, some assets are much more valuable to defenders than others. Therefore, it is important to explore whether this affects the conclusion that it is not optimal to invest in defense as N gets large. Of course, once we allow asset values to be unequal, there are infinitely many possible distributions of asset values. Rather than trying to explore a wide range of possible cases, we assume for simplicity that there are N symmetric assets, and an additional one (asset 0) that is more valuable to either the defender or the attacker (or both) than the others. As before, we assume that the costs of defending these $N+1$ assets are identical. If the defender values asset 0 more than the others, that asset may also be more valuable to the attacker. We allow for this possibility, but do not require it. If the defender values asset 0 more than the others, then as the attacker's valuations become more closely related to the defender's valuations, the chance that the attacker will prefer to target asset 0 grows. However, in the model, the defender can never be completely sure of the attacker's preferences, so for any finite N , there will still be a non-zero probability of attack on any other asset (which will be the same for all N symmetric assets).

Arguments analogous to the symmetric case ensure that for sufficiently large N , the defender should not defend any of the N symmetric assets, and may or may not invest in defense of asset 0 (depending on its value to the attacker). As before, due to the large number of symmetric assets, the optimal policy is not to defend any of them. However, we can now make two additional observations. First, if the attacker is more likely to attack asset 0 than any other asset, then the optimal defensive investment approaches the solution to a problem in which the defender cares only about asset 0. In other words, even if the absolute probability of an attack on asset 0 is small, the asymmetry still causes the defender to act as if asset 0 is the only valuable asset as N gets large. Interestingly, this can occur even if asset 0 is less valuable to the defender than the other assets, as long as it is more likely to be attacked.

As a hypothetical example, consider a defender nation that contains both mosques and synagogues. If the attacker is much more likely to attack a synagogue, then the optimal policy for the defender will often be to defend only the synagogues (possibly even if the defender cares less about synagogues than about mosques!). Similarly, consider a nation with one large water system (say, in the nation's capital city) and numerous small water systems, each serving a small town. Even if the attacker is indifferent between attacking the larger system (due to its publicity value) or a smaller system (due to the lesser difficulty of the attack), the chance of an attack on any particular small system will be small. Thus, small targets may be inherently more secure than large ones, not just because they are less valuable to attackers (although that may also be true), but simply because there are likely to be so many of them.

Second, if the attacker is equally likely to attack any of the $(N+1)$ assets (i.e., does not prefer to target asset 0), then under certain conditions, for large N the optimal policy will be for the defender to treat all $(N+1)$ assets equally and defend none of them. This can be true even if the defender values asset 0 much more than the others. Intuitively, one way to think about this is that, for any finite value of asset 0 (no matter how large), there will be a value of N large enough that the $1/(N+1)$ probability of an attack on asset 0 is too small to justify its defense.

To summarize, it is a hopeless task to defend large numbers of individual assets. It is optimal for the defender to invest in security only if those investments can be focused on a relatively small number of

attractive targets, with the remainder viewed as so unlikely to be attacked as not to merit investment. This further emphasizes the benefits of good intelligence about attacker goals and motivations.

The difficulty of defending extremely large numbers of assets also suggests that psychological factors (i.e., risk perceptions) may play an important role in achieving a sensible security strategy.²⁷ If the public demands protection against any possible terrorist attack (even attacks that are unlikely or not terribly severe), then security investment may come to have a significant detrimental effect on the health of the economy. While not strictly implied by the results of the model discussed here, part of a successful defense strategy may be to create asymmetries in the values of different assets by reshaping public perceptions. If some types of attacks come to be viewed as either largely unavoidable or else less than catastrophic, this may make it possible to focus defensive resources on the most serious risks.

5 Attacker Opportunity Costs

Until now, we have not yet modeled the cost of an attack to the attacker. Since attackers were viewed as having no opportunity cost for launching an attack, the model discussed above predicts that the attacker will invariably attack some asset, with the only uncertainty being which asset will be targeted. In practice, however, one objective of a sensible defensive policy is obviously to deter attacks if possible.

To model this, we assume the attacker has an opportunity cost K . An attack will be launched on asset i only if its expected value to the attacker, $a_i p_i$, exceeds K (in addition to asset i being more attractive than other possible targets). In other words, for an attack on asset i , a_i must be large not only relative to the other asset values, but also relative to K (with the difference depending on the p_i). Of course, K could reflect the direct cost of resources that have to be expended to mount an attack. However, since K is an opportunity cost, it could also reflect the expected cost of retaliation by the defender after either an attempt or a successful attack, the cost of withdrawn goodwill or foreign aid, or the potential value of attacking some other defender (although some of these might involve repeated games over time).

Once the attacker has an opportunity cost, some valuation profiles will now lead the attacker not to attack the defender at all. Moreover, the probability of an attack on any given asset will now depend

not only on the ratio of the p_i , but also on their magnitudes. Thus, decreasing the success probability of an attack on some particular asset now has two effects—encouraging the attacker to target another asset in some cases, but also discouraging the attacker from attacking at all in other cases.

This would seem to suggest that the existence of an attacker opportunity cost will increase the benefits of security investment, and therefore lead to higher defensive investments (at least in the unconstrained version of the problem). However, this situation is more complex than that, since increases in K can either increase or decrease the optimal level of protection. In particular, when K is initially small, then an increase in K has the effect of making security investments more valuable, in order to deter attacks. However, for large values of K , the probability of an attack will already be quite low. As a result, in such cases, there is relatively little value to defensive investments, since they are unlikely to be used. In this regime, the optimal level of defensive investment is decreasing in K .

After some reflection, this result makes sense. For example, one can think of the events of September 11, 2001, as revealing that al-Qaeda had a smaller opportunity cost than was previously believed. Since the opportunity cost had previously been believed to be high, this change resulted in significantly increased resources being allocated to defense. Similarly, countries do not spend much on defense against potential attacks by their allies. As countries become more closely allied (and hence face higher opportunity costs for attacking each other), they spend less on defenses against each other.

In practice, of course, K will not be an exogenous constant. Rather, the defender may be able to affect K as well as the p_i ; e.g., through the choice of strategies regarding retaliation, or through foreign policy that encourages other nations to impose sanctions in response to an attack. Depending on the circumstances, an increase in K may involve a combination of carrots and sticks—e.g., including both enhanced military preparedness (to more effectively retaliate after an attack) and enhanced aid (to reward those that do not attack). Note also that the above discussion treats these various types of opportunity costs as being comparable. However, this is most likely not true in practice; see for example Ref. 28.

Section 4 showed that allocating resources to hardening individual targets will become ineffective as the number of valuable targets gets large. In such cases, efforts to increase the opportunity costs faced

by potential attackers may be more effective than investments in target hardening. In the absence of some technological breakthrough, attempts to guard every bridge or inspect every container crossing the border are likely to prove futile. By contrast, steps to increase the opportunity costs of mounting an attack (perhaps through better intelligence gathering that allows attacks to be more readily interdicted, or policies that engage potential terrorists in mutually beneficial peaceful interactions) can still be valuable.²⁹

6 Effects of Secrecy

We have assumed until now that the attacker can observe the defender's allocation of resources. This may be a better approximation in some settings than others. For example, defensive plans for short-lived events may be difficult to observe until shortly before the events occur, making it difficult for attackers to adapt their choice of which event to attack. Moreover, the extent to which the attacker can observe defensive investments may be at least partially under the control of the defender. Even for a short-lived event, the defender could choose to announce its security measures publicly, or attempt to conceal them.

Conceptually, we can model a situation in which the attacker cannot observe the defender's choices as a simultaneous-move game (even if the two players do not actually move simultaneously). In a simultaneous game, the attacker's optimal choice of which asset to attack will now depend only on the a_i , not on the p_i (assumed to be unobservable). (Of course, the attacker will attempt to anticipate the defender's choice of the p_i from what they do know about the defender's preferences and constraints.)

We begin by considering the case of a single defender. Ref. 9 shows that, under the assumptions of the model discussed here, the defender will in general be better off in the sequential than in the simultaneous game. In particular, in the sequential game, the defender could simply choose the success probabilities that are optimal for the simultaneous game, or could use the first-mover advantage to choose something different. Thus, the defender is guaranteed to do at least as well in the sequential as in the simultaneous game, and should prefer to play the sequential game (i.e., to announce its defensive investments publicly). In theory, by announcing that the most valuable targets have been defended, the defender in the sequential game can deflect the attacker toward less damaging attack strategies (or

possibly deter the attacker from attacking at all, if the attacker has an opportunity cost). Thus, perhaps counter-intuitively, transparency can be an ally in strategic defense. By contrast, if the defender chose to allocate its resources in a non-optimal manner (e.g., investing in the security of low-value assets, and leaving valuable assets undefended), it might well prefer to keep those investments secret, to avoid deflecting attacks toward the more valuable targets.

However, there are obviously limits to the generality of this model. Both intuition and anecdotal evidence suggest that secrecy and even deception have an important role to play in security. We have not yet analyzed all of these situations rigorously, but here discuss some reasons why that might be the case.

First, the model has assumed that defensive investments are continuous rather than discrete, and moreover that the success probability of an attack on asset i , p_i , is a convex function of the defensive investment in that asset. Thus, the first dollar of investment is assumed to yield the most benefit, with decreasing marginal returns. When investments are indivisible, secrecy may be more important. For example, maintaining secrecy about which flights have air marshals on them is viewed as important.³⁰ If air marshals were infinitely divisible (so that flights could have fractional numbers of air marshals), and the success probability of an attack was a convex function of the (fractional) number of air marshals, then it might well be better to put some fraction of an air marshal on every flight. Similarly, if we could afford to put air marshals on every flight, then having them be visibly armed might be a better deterrent than keeping their identities secret. However, since air marshals are not divisible, and if budget constraints prevent us from putting marshals on every flight, then their effectiveness in deterring attacks would seem to be enhanced by keeping secret which flights are protected.

(Note, by the way, that this effect may not depend on the discrete nature of investment in air marshals. As long as the success probability of an attack is a non-convex function of the defensive investment in some regions, there may be optimal “randomized strategies” for the defender. Discrete investments simply satisfy non-convexity because the success probability of an attack is a step function.)

Second, the model has assumed that the attacker has no uncertainty about the value of any given asset. However, when this is not the case, secrecy or even deception can in principle be advantageous to

defenders. For example, computer hackers attempting to steal proprietary information may not know which machines contain the information they want. The use of “honey pots” in computer security takes advantage of this uncertainty by attempting to deceive attackers into believing that non-valuable targets actually contain valuable information.³¹ Secrecy and deception can also work in situations where the attacker’s goal is to damage those assets judged to be most valuable by the defender (as in Ref. [32]). In such cases, an attacker who did not know the defender’s valuations could attempt to learn about them by observing which assets the defender found worth protecting. A defender faced with this situation might want to keep its defensive investments secret (in order to avoid tipping off the attacker), or invest resources in defending assets that it did not consider valuable (in order to deceive the attacker).

Finally, in one of the examples mentioned in passing early in this paper, some types of attackers (e.g., computer hackers) may specifically put greater value on damaging those assets that are highly defended. If the attacker’s valuation of a given asset is not constant, but rather is decreasing in the success probability of an attack, then reducing the success probability through defensive investment could paradoxically make a target more attractive to the attacker. For example, computer hackers take great pride in penetrating well-defended computer systems (e.g., at the Department of Defense). Similarly, terrorists may find greater publicity value in successful attacks against well-defended targets (such as the Pentagon) than against soft targets, independent of the inherent values of the assets.

Thus, the result that public announcements of defensive investments are optimal in the current model should not be taken to imply that secrecy and deception have no role in security. We believe that there can be important benefits in some cases from announcing defensive investments (thereby deterring attacks, and/or deflecting them to less damaging targets). However, models that capture the role of secrecy and deception would be a useful complement to the work in Ref. 9, to inform on decisions about which defensive investments to announce and which to keep secret.

7 Discussion and Conclusions

The basic conclusion of this work is that, when facing the threat of an intentional attack, it is important to model the strategic behavior of the attacker. Concentrating on the last target struck by the attacker (e.g., transportation security) may be effective if the attacker continues to concentrate on transportation, but may be of no avail if the attacker switches to targeting the food supply. Pre-screening containers from ports shipping 80% of the containers entering this country (as done by the Container Security Initiative³³) might significantly enhance security if attackers did not alter which ports they used, but may be of no avail if attackers can shift their activities to those ports shipping the remaining 20% of all containers, unless coupled with effective screening of “high-risk shipments that have not been prescreened”.³⁴

The model of Ref. [9] provides a framework for examining this attacker-defender interaction. One especially significant implication of the model for current security policy involves the results concerning problems with large numbers of targets. Even a centralized decision maker who perfectly manages the externalities upon which much of the literature has been focused has little hope of defending a large number of important targets. The defender’s only alternatives in such cases are to focus security investment on the targets most likely to be attacked, or to decrease the overall attractiveness of attacks by making them more costly to attackers. Thus, this work supports the recent efforts by the Department of Homeland Security to focus on a few of the most severe threats to security and make security funding more risk-based.³⁵ These would seem to be highly desirable in light of criticism that currently, “most of the money is allocated on a political basis rather than a sound cost benefit analysis”.¹³ An effective terrorism defense must either involve hard choices about what not to defend, or change the incentives faced by potential terrorists.

References

[1] Bier, V. M., and V. Abhichandani, “Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries.” *Risk-Based Decisionmaking in Water Resources X*, pp. 59-76, Santa Barbara, California, November 3-8, 2002.

- [2] O'Hanlon, M., P. Orszag, I. Daalder, M. Destler, D. Gunter, R. Litan, and J. Steinberg. *Protecting the American Homeland*. Washington, DC: Brookings Institution, 2002.
- [3] Bier, V. M., A. Nagaraj, and V. Abhichandani, "Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries," *Reliability Engineering and System Safety*, Vol. 87, pp. 313-323, 2005.
- [4] Banks, D., and S. Anderson. "Game-theoretic risk management for counterterrorism." Technical Report, Office of Biostatistics and Epidemiology, Center for Biologics Evaluation and Research, U.S. Food and Drug Administration, 2003.
- [5] Major, J. A., "Advanced techniques for modelling terrorism risk," *Journal of Risk Finance*, Vol. 4, pp. 15-24, 2002.
- [6] Woo, G., "Quantitative terrorist risk assessment," *Journal of Risk Finance*, Vol. 4, pp. 7-14, 2002.
- [7] Woo, G., "Insuring against Al Qaeda," NBER Insurance Project Workshop presentation (<http://www.nber.org/confer/2003/insurance03/woo.pdf>), 2003.
- [8] Besnard, D., and B. Arief, "Computer security impaired by legitimate users," *Computers and Security*, Vol. 23, pp. 253-264, 2004.
- [9] Bier, V. M., S. Oliveros, and L. Samuelson, "Choosing what to protect: Strategic defensive allocation against an unknown attacker," submitted to *Journal of Public Economic Theory*, 2005.
- [10] Willis, H. H., A. R. Morral, T. K. Kelly, and J. J. Medby, "Estimating terrorism risk," Rand Corporation, MG-388-IR&D, May 2005.
- [11] Lipton, E., "Big cities will get more in antiterrorism grants," *New York Times*, December 22, 2004, pg. A20.
- [12] Ripley, A., "How we got homeland security wrong: the fortification of Wyoming, and other strange tales from the new front line," *Time*, March 29, 2004.
- [13] de Rugy, V., "What does homeland security spending buy?" American Enterprise Institute, April 2005.

- [14] Zycher, B., *A Preliminary Benefit/Cost Framework for Counterterrorism Public Expenditures*. Rand Corporation, MR-1693-RC, 2003.
- [15] Brunet, A., "Grant funding to state and local governments and systematic assessment of vulnerability," Center for Risk and Economic Analysis of Terrorism Events, University of Southern California, June 2005.
- [16] House Select Committee on Homeland Security, "An analysis of first responder grant funding," April 2004.
- [17] Carafano, J. J., "Homeland security dollars and sense #1: Current spending formulas waste aid to states," The Heritage Foundation, <http://www.heritage.org/Research/HomelandDefense/wm508.cfm> (accessed August 16, 2005), May 2004.
- [18] Enders, W., and T. Sandler, "What do we know about the substitution effect in transnational terrorism?" In A. Silke and G. Ilardi (editors), *Researching Terrorism: Trends, Achievements, Failures*. London: Frank Cass, 2004.
- [19] Lipton, E., "U.S. to spend billions more to alter security systems," *New York Times*, May 8, 2005, pg. A1.
- [20] Marek, A. C., "Security at any price? Homeland protection isn't just Job 1 in Washington; it's more like a big old government ATM," *U.S. News and World Report*, May 30, 2005.
- [21] Orszag, P. R., "Statement of Peter R. Orszag to the National Commission on Terrorist Attacks Upon the United States," November 19, 2003.
- [22] Kunreuther, H., and E. Michel-Kerjan, "Dealing with extreme events: New challenges for terrorism risk coverage in the U.S.," Center for Risk Management and Decision Processes, Wharton School, University of Pennsylvania, June 2004.
- [23] Kunreuther, H., and G. Heal, "Interdependent security," *The Journal of Risk and Uncertainty*, Vol. 26, pp. 231-249, 2004.
- [24] Arce, D. G., and T. Sandler, "Transnational public goods: Strategies and institutions," *European Journal of Political Economy*, Vol. 17, pp. 493-516, 2001.

- [25] Lakdawalla, D., and G. Zanjani, "Insurance, self protection, and the economics of terrorism," National Bureau of Economic Research, Working Paper 9215, 2002.
- [26] Keohane, N. O., and R. J. Zeckhauser, "The ecology of terror defense." *Journal of Risk and Uncertainty*, Vol. 26, pp. 201-229, 2003.
- [27] Mueller, J., "A false sense of insecurity?" *Regulation*, pp. 42-46, Fall 2004.
- [28] Frey, B. S., and S. Luechinger, "How to fight terrorism: Alternatives to deterrence," *Defence and Peace Economics*, Vol. 14, pp. 237-249, 2003.
- [29] Haimes, Y. Y., "Roadmap for modeling risks of terrorism to the homeland," *Journal of Infrastructure Systems*, Vol. 8, No. 2, pp. 35-41, June 2002.
- [30] Hudson, A., "Air marshals' secrecy ruined by dress code," *The Washington Times*, July 9, 2004.
- [31] Martin, W. W., "Honey pots and honey nets—Security through deception," The SANS (SysAdmin, Audit, Network, Security) Institute, Bethesda, Maryland, May 2001.
- [32] Hendricks, K., and R. McAfee, "Feints," accepted for publication in *Journal of Economics and Management Strategy*, 2005.
- [33] Myers, M., "Extending the nation's zone of security," *Customs and Border Protection Today*, U.S. Bureau of Customs and Border Protection, March 2004.
- [34] U.S. Bureau of Customs and Border Protection, "Container Security Initiative fact sheet," July 2005, http://www.cbp.gov/linkhandler/cgov/border_security/international_activities/csi/csi_fact_sheet.ctt/csi_fact_sheet.doc.
- [35] Lipton, E., "U.S. report lists possibilities for terrorist attacks and likely toll," *New York Times*, March 16, 2005.