Defending Against *Strategic* Terrorists Over the Long Run: A Basic Approach to Resource Allocation*

Robert Powell

August 2006

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

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

Defending Against *Strategic* Terrorists Over the Long Run: A Basic Approach to Resource Allocation

Abstract

The efficient allocation of resources to defend the United States' critical infrastructure and key assets against terrorist attacks involves both short and long-run issues. The former focus on attempts to detect and disrupt the planning and execution of operations already underway. The latter focus on long-term efforts to harden sites, reduce vulnerabilities, and make attacks more difficult and less attractive. Because these are longer-term efforts, strategic terrorists will adjust and respond to these measures in order to strike where the defense is weak and the expected gains are high. Recognizing that terrorists are strategic and that resources are limited, the Department of Homeland Security emphasizes that resources must be allocated on the basis of risk. This paper shows that the current approach to risk management does not threat terrorists as fully strategic and that the failure to do so can lead to a significant misallocation of defensive resources. The paper also provides a framework for allocating resources against long-term threats.

Defending Against *Strategic* Terrorists Over the Long Run: A Basic Approach to Resource Allocation

The United States faces an immense challenge in securing the country's critical infrastructure and key assets in the aftermath of the attacks of September 11, 2001. The Department of Homeland Security (DHS) has been charged with this task, and it is vast.[1] By the end of 2005, the National Asset Database listed about 80,000 sites including "nuclear power plants, pipelines, bridges, stadiums, and locations such as Times Square" (GAO 2005, 75).[2] Of these, the Department identified 1700 highest priority sites and intended to visit each of them "to assess their vulnerabilities to various forms of attack..." (Moteff 2006, 38).

The challenge of protecting the country's critical infrastructure is all the more formidable because terrorists are strategic. As the *National Strategy for Homeland Security* emphasizes, "One fact dominates all homeland security threat assessments, terrorists are strategic actors" (White House 2002, 7). No one thinks that hardening and rebuilding the levies around New Orleans affects the probability that another hurricane like Katrina will strike that city. But, strategic actors do try to strike where the defense is weak and the expected gains are high. Protecting one site may therefore shift the risk of attack to another. "Increasing the security of a particular type of target, such as aircraft or buildings, makes it more likely that terrorists will seek a different target. Increasing countermeasures to a particular terrorist tactic, such as hijacking, makes it more likely that terrorists will favor a different tactic" (White House 2002, 29).[3]

Because resources are limited and the number of targets is huge, it is impossible to defend everything. Since assuming office, Homeland Security Secretary Chertoff has stressed that resources must be allocated according to risk. "Risk management must guide our decision making as we examine how we can best organize to prevent, respond and recover from an attack"(Chertoff 2005, 2). To this end, the Department is developing a method

[1] For recent overviews of these efforts, see GAO (2005) and Moteff (2006).

[2] A critical review of the database is DHS (2006) which notes that the list of assets also includes petting zoos, fun parks, and some Wal-Marts.

[3] For statistical evidence of these "substitution" effects see Enders and Sandler 2004.

of making comparative risk assessments of different sites within and across sectors of the economy. The Department's own efforts and most others are fundamentally grounded in what engineers broadly call "risk analysis."[4] Risk analysis has a fundamentally important role to play in the efficient allocation of defensive resources. But risk analysis generally does not treat terrorists as fully strategic actors, and this can lead to a significant misallocation of resources.[5]

Game theory, by contrast, does treat actors as fully strategic (indeed perhaps too strategic). This paper presents a basic game-theoretic framework for analyzing the problem of allocating limited resources to defend against fully strategic terrorists over the long run.[6] In contrast to shorter-term attempts to detect and disrupt the planning and execution of terrorist operations already underway the like August 2006 plot to bomb multiple transatlantic flights, long-term efforts like improved port and container security focus on hardening potential targets, reducing vulnerabilities, and making attacks generally more difficult and less attractive. Because these efforts take time, strategic terrorists are likely to be aware of and react to some or all of these measures in order to strike where the defense is weak and the expected gains are high. The goal of the present study is to identify a few general principles to guide our thinking about the long-run resource-allocation problem. Ideally, these principles and the underlying framework will help decisionmakers get their bearings as they confront the daunting task of protecting the country's infrastructure and as they try to navigate the political pressures of pork-barrel

[4] The draft *National Infrastructure Protection Plan*, *V2.0* describes the Department's efforts (NIPP 2006, especially 35-58). Haimes (2004) provides an introduction to risk analysis, and Willes *et. al.* (2005) discuss the issues of estimating terrorism risk in the context of risk analysis.

[5] On the failure of risk analysis to treat terrorists as fully strategic, see Bier 2005 and Kardes 2005.

[6] Efforts to include elements of game theory in risk analysis include Bier (2004); Bier, Nagaraj, and Abhichandani (2005); Hausken (2002); and Paté-Cornell and Guikema (2002). Fully game-theoretic studies include Bier, Oliveros, and Sammuelson (2005); Bueno de Mesquita (2005); Powell (2006a,b); Rosendorf and Sandler (2005); and Sandler (2005).

politics.[7]

Allocating resources against a strategic adversary is an old problem in game theory going back at least as far back as studies of Colonel Blotto games done in the early years after the Second World War.[8] In the games studied here, the government has to decide how to divide its resources to defend multiple sites against different types of attack or threat scenarios. A terrorist group then decides which site to strike and how to attack it. The analysis considers these decisions in four different settings: (i) the zero-sum case in which the government's and terrorists' payoffs are diametrically opposed; (ii) the nonzero-sum case; (iii) a setting in which the terrorists can observe and therefore react to the government's allocation when deciding what to attack; and (iv) a situation in which the government's allocations are secret and unobservable.[9] Remarkably, the optimal (i.e., equilibrium) allocation is the same in all of these cases, and it is to minmax the terrorists. That is, the government's optimal allocation minimizes the terrorists' maximum payoff or, in other words, imposes the lowest possible ceiling on the terrorists' payoff given the available resources.[10]

Two important policy implications follow. The first centers on the role that intelligence and basic research on terrorists and terrorist organizations play in the allocation of long-term defensive resources. At present, DHS's risk-managment approach plans to

[7] Past spending on "protecting" the country's critical infrastrcture has been widely criticized for being heavily influenced by pork-barrel politics (e.g., 911 Public Discourse Project's final report (2005)). Even if these political pressures could be contained, it is not clear that decisionmakers have yet got their bearings and have a framework of thinking about how to allocate resources against a strategic adversary.

[8] See for example Blackett 1958, Gross and Wagner 1950, and Tukey 1949. Important subsequent work includes Coughlin 1992 and Shubik and Weber 1981.

[9] Supposing the game between the government and terrorists to be zero sum does not sound like a strong assumption. But it is more demanding that it may seem at first. For example, the zero-sum assumption means that government and terrorists have exactly the same risk prospensity, i.e., the terrorists are willing to take a gamble if and only if the government is willing to take that gamble.

[10] The Minimax theorem ensures that players always minmax each other in any equilibrium in any finite, two-player, zero-sum game. The surprising result established below is that the government's equilibrium strategy continues to be to minmax the terrorists even in the nonzero-sum cases and regardless of the observability of the government's allocation. By contrast, the terrorists' strategies do vary as the setting changes.

use intelligence and basic research to formulate "threat assessments" which are the probabilities that different kinds of attack occur. The game-theoretic analysis below indicates that this focus on threats is misplaced when dealing with strategic adversaries and the long-term allocation of defensive resources. The question to be answered is not "what are there probabilities of attack?" These probabilities will change as the government's defensive efforts change the relative vulnerabilities of the possible targets. The question to ask is "what are the terrorists' goals and motivations?"

The second implication is a simple principle for allocating defensive resources: establish a "threshold of attraction" or, more accurately, a "threshold of expected terrorist gain from attack." If the terrorists' expected gain from an attack on a site exceeds this threshold, the site has to be hardened and the expected gain reduced to the threshold level. Sites with an expected gain below this threshold do not have to invest any further in longer-run defense. The effect of this policy is that resources will be dedicated to hardening sites that are most likely to be attacked while not spending on those sites much less likely to be struck. This policy is closely related to what has been a common policy recommendation, namely, hardening those sites that if attacked would impose the largest expected losses on the United States in terms of the "number of casualties, extent of economic damage, harm to key institutions, and ... symbolic significance" (O'Hanlon *et. al.* 2003, 5). The present analysis provides some theoretical support for this policy and highlights important qualifications.

The next section shows that failing to treat terrorists as fully strategic actors can lead to a significant misallocation of resources. That section also introduces some of the key elements of the model and relates them to the core concepts of consequences, threat, and vulnerability which are the basis of the Department of Homeland Security's approach to risk analysis and management. The subsequent section develops the general framework and offers an intuitive explanation for the optimality of the minmax allocation as well as a simple algorithm for finding it. There follows a discussion of the role of intelligence and threat assessments in the long-run resource-allocation problem. A final section focuses on the policy option associated with the minmax allocation, its relation to "weakest-link" policies, and some important caveats. The Appendix presents a formal equilibrium analysis.

Threat Assessments and Fully Strategic Actors

The cornerstone of the National Infrastructure Protection Plan (NIPP) is "its risk management framework" (NIPP 2006, 35). This framework does take terrorists' intentions into account, but it does not threat them as fully strategic actors "who modify their tactics and targets to exploit perceived vulnerabilities and avoid observed strengths," who "shift their focus to less protected sites" as security around other sites increases (White House 2003, viii). As this section shows, not treating adversaries as fully strategic can lead to a significant misallocation of long-term resources.

A key element of DHS's risk-management approach is to develop a systematic method of assessing the "risk" associated with terrorist attacks on different targets. "Risk" here is a term of art and not simply the probability or likelihood of an attack. Rather it is the expected loss resulting from terrorism and is a function of three more basic factors: consequence, vulnerability, and threat. Consequence "is the range of loss or damage that can be expected from a successful attack" (NIPP 2006, 41). The vulnerability of an asset "is the probability that a particular attack will succeed against a particular target" (GAO 2005, 25). Threat "is the probability that a specific target is attacked in a specific way" (Willis *et. al.* 2005, 8). Combining these elements, the expected loss from terrorism or risk associated with a site is the probability of an attack (threat) times the probability that an attack succeeds (vulnerability) times the loss from a successful attack (consequence) (Willis *et. al.* 2005, 10). That is, risk = consequence × vulnerability × threat.

One can think of the consequences and vulnerabilities associated with a specific site as akin to physical properties that have little or nothing to do with terrorists' goals or motivations. Estimating the expected loss that would result from a particular kind of attack *given* that the attack takes place is in principle much like estimating the expected loss that would result from an accidental fire, explosion, or software failure given that those incidents occur. Indeed, wide used estimates of the potential casualities from a terrorist attack on various chemical facilities come from an Environmental Protection Agency report done prior to 9/11 on the effects of the accidental release of hazardous materials.[11] Consequence and vulnerability estimates may be very hard to do, but they are part and parcel of ordinary risk analysis and management. Indeed, DHS apparently plans to have the individual owners and operators of the sites conduct most of the consequence and vulnerability assessments themselves with the Department providing a common method to ensure that risks based on these underlying assessments can be compared across different sites (NIPP 2006, 42; GAO 2005, 77). To this end, the Department is developing a set of tools called "Risk Analysis and Management of Critical Asset Protection" (RAMCAP) to "enable owners and operators to calculate potential consequences and vulnerability to an attack using a consistent system of measurements" (NIPP 2006, 42).

DHS will provide the third component of risk. Unlike consequences and vulnerabilities, the "[a]ssessment of the current threat to the United States is derived from extensive study and understanding of terrorists and terrorist organizations and is frequently dependent on analysis of classified information. DHS will provide U.S. government-coordinated assessments of potential terrorist threats..." (NIPP 2006, 47).

Using intelligence in this way is not easy and developing intelligence-based threat assessments has turned out to be one of the biggest challenges to risk management. According to a recent Government Accounting Office report, DHS officials "stated that a lack of intelligence data and law enforcement data limits their ability to develop the relative probability for various threat scenarios, and for this reason, they have focused their initial efforts on developing vulnerability and consequence data... [T]he intelligence community – including the intelligence components of DHS – has been unable to provide detailed intelligence on threats to most sectors, infrastructure, assets, or asset types" (GAO 2005, 76).

In addition to the pragmatic difficulty of developing intelligence-based threat assessments, there is a deeper, more conceptual problem with using this kind of threat analysis

[11] See Belke (2000) for the original report.

as the basis for helping to decide how to allocate resources over the long run to defend against strategic terrorists. Hardening a site and thereby reducing its vulnerability will induce strategic terrorists to “seek more accessible and less well protected facilities and events” (White House 2003, 7). Thus lowering the vulnerability of a site also lowers the threat to that site (and possibly raises the threat to other sites). That is, strategic terrorists create a feedback loop between vulnerability and threat, and this feedback complicates efforts to allocate long-term resources on the basis of comparative risk. Should resources be allocated on the basis of the original threat assessments or on the induced threats that the sites will actually face once they have been hardened? Or should some combination of these threats be used to calculate relative risks and allocate resources?

A simple vignette illustrates the complications. A terrorist group is trying to destroy a particular site and there are only two avenues of attack, through the front or through the back. The threat assessment also indicates that the probability that the terrorists will try to come through the front is 2/3. The defender, therefore, spends it resources on hardening the front. In light of these efforts, the terrorists “shift and focus on another vulnerability” (White House 2002, 7), namely, the back, and they attack there. Shortly before they strike an updated threat assessment warns that any attack is almost certain to come at the rear.

Did the defender allocate its resources optimally? Should it have used the original threat assessment, the updated estimate, some combination of them, or neither of them?

A simple model helps answer these questions and highlights the potential misallocation of resources that can result from neglecting the feedback between vulnerability and threat. The formalization also introduces key components of the more general game-theoretic model which is developed below and which treats terrorists as fully strategic actors.

The government is trying to decide how to spend R resources on protecting two sites and suffers a loss of L_1 if site 1 is successfully attacked and L_2 if site 2 is successfully attacked. Let r_1 and r_2 denote the resources allocating sites 1 and 2 respectively with $r_1 + r_2 = R$. Then the probability that an attack on site j succeeds if the defender spends r_j on defense is $v_j(r_j)$. The more the defender allocates to a site, the less likely a successful attack, so $v_j(r_j)$ decreases as r_j goes up.[12] Finally, the terrorists attack site 1 with probability τ and site 2 with probability $1 - \tau$. All of this implies that the defender's expected loss to allocating r to site 1 and $R - r$ to site 2 is $L(r) = \tau L_1 v_1(r) + (1 - \tau) L_2 v_2(R - r)$. The first term on the right in this equality is the expected loss from an attack on the first site, $L_1 v_1(r)$, weighted by the probability of an attack on that site, and the second term is the expected loss from an attack on the second site weighted by the probability of that attack.

The basic elements of this model correspond to vulnerability, threat, and consequence. The vulnerability of a site is the probability that an attack on that site succeeds, and this is precisely what $v_1(r_1)$ and $v_2(r_2)$ are. The threat to a site is the chance of an attack on that site and is formalized as τ and $1 - \tau$.

Finally, consequences which are usually measured in terms of "the expected magnitude of damage (e.g., deaths, injuries, or property damage)" resulting from a successful attack (Willis *et. al.* 2005, 9) correspond to the losses L_1 and L_2 but with important qualifications. Game-theoretic analyses are generally based on expected-utility theory which assumes that the actors can rank the possible outcomes from best to worst.[13] In the present case, this means that the terrorists can rank the sites in order from the one they would most like to destroy to the one they care least about destroying.[14] Similarly, the government can order the sites from the one it would be least willing to lose to the one that it would be most willing to lose (and obviously not losing any is the best outcome). These preference orderings formalize the actors' goals and motivations.

These rankings may or may not correspond to the damage an attack would do measured in terms of lives, injuries or property. For example, terrorists might rank destroying the Statue of Liberty very high although this would actually impose little in direct damages. Efforts to measure consequences in terms of lives, injuries, or dollars should really be seen

[12] This assumption will be maintained throughout. But as Sagan (2004) argues, this may not always be the case.

[13] On expected-utility theory, see Mas-Colell, Whinston, and Green (1995).

[14] The actors are also assumed to be able to rank risky or uncertain alternatives, e.g., a twenty-percent chance of destroying the most attractive target compared to a forty-percent chance of destroying the fifth most prefered target.

as ways of trying to assess the actors' underlying preference orderings. This qualification will be especially important below when we discuss the role of intelligence and threat assessments in allocating resources. The distinction between payoffs and consequences will also be crucial to the discussion of "weakest-link" policies.

Resources are limited, so investing more in one site means investing less in another. Figure 1(a) plots the government's expected loss and the trade off it faces when the value of the first site relative to the second is two to one (i.e., $L_1 = 2$ and $L_2 = 1$). The probability of a successful attack on a site is $v_j(r_j) = (1 - r_j/R)^2$, and the total resources is $R = 10$. The relative threats to the sites are also assumed to mirror their relative losses. That is, the odds of an attack on site 1 are the same as the sites' relative value. In symbols, $\tau/(1-\tau) = 2/1$ so the threat assessments are $\tau = 2/3$ and $1 - \tau = 1/3$ (just as they were in the vignette above).

If we ignore the feedback between vulnerability and threat by assuming the threat $\tau = 2/3$ remains the same regardless of the allocation r and its effect on the sites relative vulnerability, then the government's optimal allocation minimizes the government's expected loss. This optimal allocation is $r^* = 8$ as illustrated in Figure 1(a).[15] This allocation equates the marginal benefit of spending slightly more on site 1 with the marginal cost of having slightly less to spend on site 2. More formally, the optimal risk-management allocation r^* given the threat assessment τ satisfies the first-order condition $dL(r^*)/dr = 0$ or, equivalently, r^* solves $\tau L_1 v_1'(r^*) = (1-\tau)L_2 v_2'(R - r^*)$.

But ignoring the feedback between vulnerability and threat leads to a significant misallocation of resources against strategic adversaries. The terrorists attack the site that offers the highest expected payoff which is assumed here to be equivalent to imposing the largest expected loss on the government.[16] As Figure 1(b) shows, the terrorists' expected payoff to attacking site 2 at $r^* = 8$ is larger than their expected payoff to striking site 1. The terrorists therefore attack site 2 thereby vitiating the original threat assessment on

[15] Taking the threat τ to constant or exogenous treats the resource allocation issue as a decision-theoretic problem. Decision theory, unlike game theory, disregards the feedback between the actors' strategies.

[16] This zero-sum assumption is discussed and relaxed below.

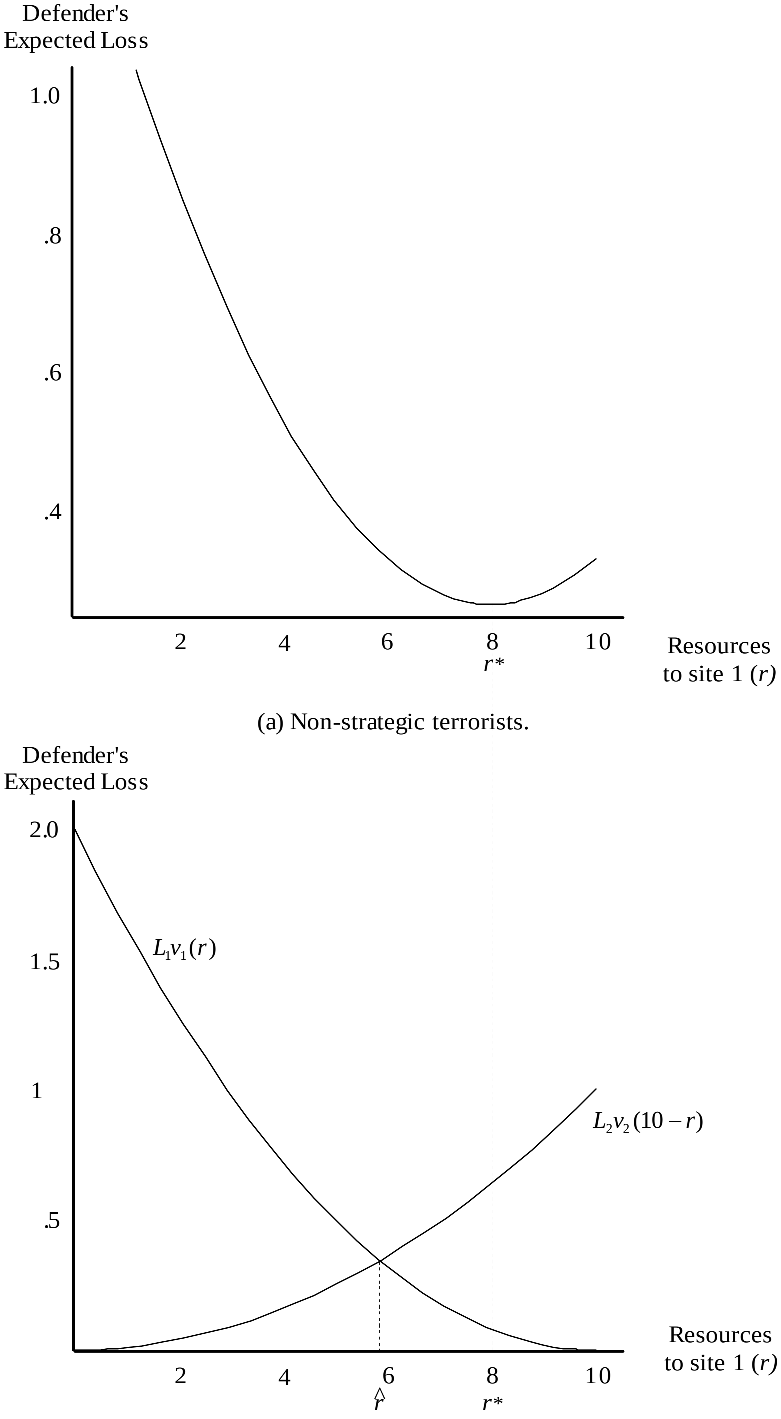

(a) Non-strategic terrorists.

(b) Strategic terrorists.

Figure 1: Strategic and non-strategic terrorists.

the basis of which the government allocated its resources.

In game-theoretic terms, the allocation $r^* = 8$ and the threat $\tau = 2/3$ are not an equilibrium of the underlying game. Equilibria incorporate strategic feedback effects by requiring each actor to play optimally against each other, i.e., each chooses a strategy which maximizes its payoff (or minimizes its loss) given the other actor's strategy. The allocation r^* and threat τ are only "half" an equilibrium. The government is playing optimally against the terrorists' strategy τ, but the terrorists are not playing optimally against the government. They maximize their payoff against r^* by hitting site 2 ($\tau = 0$), not by going after site 1 with probability $\tau = 2/3$.

What is the optimal allocation against a strategic terrorist in this example? It is the allocation that minimizes the terrorists' expected gain (or, equivalently, the government's expected loss) given that the terrorists anticipate and respond to the allocation. Allocation $\widehat{r} \approx 5.9$ in Figure 1(b) does this. If the government spends more than this on site 1, the expected payoff to striking site 2 is larger than the expected payoff to hitting site 1. The terrorists therefore attack site 2 and the government's expected loss is larger than it would have been at $\widehat{r} = 5.9$. If, by contrast, the government spends less than this on site 1, that site becomes the more attractive target and the terrorists go after it. This again leaves the government with a higher loss than it would have had at $\widehat{r} \approx 5.9$.

Failing to take the feedback between threat and vulnerability into account can lead to a significant misallocation of resources. Even in this simple two-site example, the government overspends on site 1 by more than 35% ($r^*/\widehat{r} = 1.37$) based on the threat assessment of $\tau = 2/3$.

The GAO suggests DHS may be inclined to assume threats are equally likely in light of the Department's difficulty in developing intelligence-based threat assessments. But, the GAO warns, assigning "equal likelihood to various threat scenarios would mean ... risk assessments will not include key threat data... And because data on the relative likelihood of threat scenarios are not included, the assessments will emphasize high-consequence events that may have a low probability of occurring. This approach is bound to result in potentially unreliable or incomplete data on where to establish priorities" (GAO 2005,

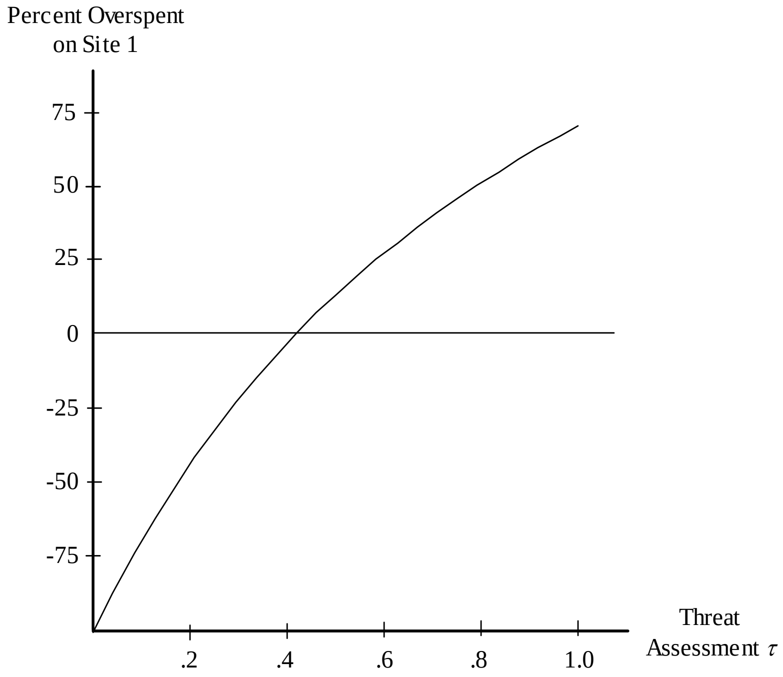

Figure 2: Threat assessments and the misallocation of resources.

76). Equally likely threats ($\tau = 1/2$) yield an "optimal" allocation of $r^* \approx 6.7$ and a misallocation of about 15% in this case ($r^*/\widehat{r} = 1.14$).[17]

Figure 2 plots the percentage of resources misallocated as a function of the assumed threat τ. Let $r^*(\tau)$ be the "optimal" allocation against threat τ. That is, the marginal value of spending slightly more on site 1 is just offset by the marginal loss of having just a bit less to spend on site 2 at $r^*(\tau)$ if the threat level is τ. Formally, $r^*(\tau)$ solves $dL(r)/dr = 0$ with $r^*(2/3) = 8$ and $r^*(1/2) \approx 6.7$ in the examples above. Then the level of over spending on site 1 relative to the actual optimal $\widehat{r}$ is $r^*(\tau)/\widehat{r} - 1$ and is plotted in Figure 2. The spends too much at threat assessments above $\tau \approx .41$ and too little at threat assessment below this level.

A Basic Framework for Allocating Resources Against Strategic Terrorists

In the example above, the government only has two sites to protect. More generally, suppose that a defender has to protect S sites against T possible types of attack or

[17] That is, $dL/dr = 0$ at $r^* \approx 6.7$ and $\tau = 1/2$.

"threat scenarios." The defender therefore has to guard against roughly $S \times T$ attack profiles.[18] The more the defender spends to protect against a given profile, the lower the vulnerability and the less attractive that profile is to the terrorists. How should the defender allocate R resources across the attack profiles given that the terrorists will take this allocation and its effects on relative vulnerabilities into account in deciding which site to attack and how to attack it?

The answer turns out to be very simple at least in principle. A strategic terrorist will pursue the most attractive attack profile, i.e., the site-threat combination offering the highest expected payoff. The defender therefore should invest in hardening against that site-threat combination. But the more the defender spends, the less vulnerable that site becomes to that kind of attack and the lower the expected payoff to that attack profile. Eventually, this profile will be no more attractive than what was initially the second most attractive attack profile. At this point, the defender must invest in protecting against both profiles so the neither is more attractive than the other. The more the defender spends on these two site-threat combinations, the lower the vulnerability of the respective sites and the less attractive the profiles become. Eventually, they are no more attractive than what was originally the third most attractive attack profile. From here on the defender must invest in guarding against all three kinds of attack so that no one is any more attractive than the other two. The defender continues to allocate its resources in this way by spending on and hardening against more and more attack profiles so as to make the most attractive profile as unattractive as possible.

The allocation resulting from this algorithm minmaxes the terrorists. That is, the defender's optimal allocation minimizes the terrorists' maximum payoff. This section formalizes the model, illustrates the optimal (equilibrium) allocation graphically, and discusses some qualifications and limitations. The Appendix presents a technical equilibrium analysis of the game.

To specify the actors' payoffs, suppose the defender suffers a loss $L_{st} > 0$ if the terrorists

[18] Some threats against some kinds of targets are likely to be nonsensical, e.g., a truck bomb aimed at bringing an airliner down in midflight.

successfully strike site s along the lines of threat scenario t. The vulnerability of this site to this type of attack, i.e., the probability that attacking this site in this way succeeds, is $v_{st}(r_{st})$ where r_{st} are the resources devoted to protecting this site from this type of attack. Spending more reduces vulnerability, so $v_{st}(r_{st})$ goes down as r_{st} goes up.[19] The threat or probability that the terrorists strike site s using attack type t is τ_{st}. Then the expected loss at site s due to threat scenario t is $L_{st}v_{st}(r_{st})\tau_{st}$.[20]

As discussed above, efforts to assess the consequence of a given type of attack on a particular site should really be seen as attempts to assess payoffs or preferences. To the extent that consequences approximate payoffs, $L_{st}v_{st}(r_{st})\tau_{st}$ is the risk associated with a particular attack profile, i.e., the consequences of a successful attack times the probability that an attack succeeds times the probability of an attack.

Summing over all of the site-threat combinations gives the defender's expected loss to a specific resource allocation against a particular threat profile. In symbols, the defender's expected loss is $\Sigma_{s=1}^{S}\Sigma_{t=1}^{T}L_{st}v_{st}(r_{st})\tau_{st}$ where the sum of the resources spent on all of the profiles is no more than R and the sum of the probabilities across all of the profiles is one, i.e., $\Sigma_{s=1}^{S}\Sigma_{t=1}^{T}r_{st} \leq R$ and $\Sigma_{s=1}^{S}\Sigma_{t=1}^{T}\tau_{st} = 1$. As for the terrorists' payoffs, their gain to successfully striking site s using attack type t is $G_{st} > 0$. The terrorists' expected payoff from a specific threat profile against a given allocation is $\Sigma_{s=1}^{S}\Sigma_{t=1}^{T}G_{st}v_{st}(r_{st})\tau_{st}$.

Figure 3 depicts the optimal allocation and helps explain why it is optimal. The curve $G_Av_A(r)$ plots the terrorists expected payoff to site-threat combination $A = (s_A, t_A)$ as the defender invests more resources in hardening site s_A against threat scenario t_A. The more the defender spends, the lower the terrorists' payoff. Similarly, the terrorists' payoffs to the profiles B, C, and D are $G_Bv_B(r)$, $G_Cv_C(r)$, and $G_Dv_D(r)$.

To derive the optimal allocation, suppose that the defender has not yet allocated anything so $r_{st} = 0$ for all site-threat combinations. Given this initial null allocation, one profile will be the most attractive to the terrorists, one the second most attractive,

[19] Formally, $v'_{st}(r_{st}) < 0$.

[20] As noted above (see note 18), some of these site-threat combinations are likely to be nonsensical in which case they can be disregarded.

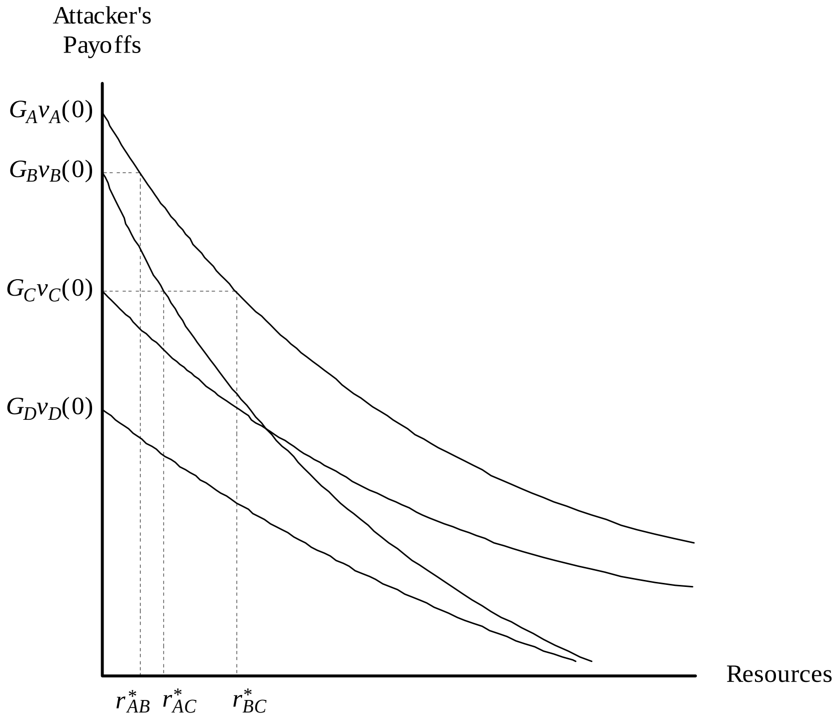

Figure 3: The optimal allocation.

one third and so on.[21] In Figure 3, attack profile A is the most attractive at the null allocation, B is next, then C and D.

The terrorist use the attack profile offering the highest expected payoff which is A at the null allocation. The defender therefore should protect against this kind of attack by spending on A. The more the defender spends, the farther the terrorists' expected payoff moves down $G_A v_A(r)$ is Figure 3. At $r = r^*_{AB}$, the terrorists' payoff to profile A equals the payoff to what was initially the second most attractive profile B. That is, $G_A v_A(r^*_{AB}) = G_B v_B(0)$.

If the defender continues to spend solely to protect against profile A, strategic terrorists will maximize their payoff by attacking with profile B and the additional resources spent on A above r^*_{AB} will have been wasted. Once A and B are equally attractive, the defender must divide any further spending between them so that the expected payoffs to these

[21] We assume there are no ties to simplify the exposition. Ties complicate the technical analysis but have no substantive effects.

profiles remain equal and neither is more attractive than the other. The more the defender allocates to A and B, the lower the terrorists' expected payoffs to attacking in these two ways. Eventually these two profiles are no more attractive than what was initially the third most attractive kind of attack. This occurs in Figure 3 when the defender is spending r^*_{AC} on A and r^*_{BC} on B so that $G_A v_A(r^*_{AC}) = G_B v_B(r^*_{BC}) = G_C v_C(0)$. From here on additional spending must be divided across these three profiles so that neither is any more attractive than the other two.

The defender continues in this way, making the most attractive site-threat combinations less and less attractive, until it has fully allocated R resources. This requires the defender to distribute its resources across more and more profiles so that no one is any more attractive than any other. The resulting allocation minimizes the terrorists' maximum payoff. This minmax allocation imposes the lowest possible ceiling on the terrorists' expected payoff given the available resources. Three key points follow.

First, the minmax allocation is optimal whether or not the terrorists can see how the defender allocates its resources before deciding which site to strike and how to attack it. The discussion above presumes that the terrorists can observe the defender's efforts to harden against potential attacks, adjust their plans accordingly, and then attack using the profile offering the highest expected payoff. But suppose that the defender can keep some or all of these efforts secret? In these circumstances, the terrorists pursue the attack profile they believe to offer the highest expected payoff. Anticipating this, the defender still wants to minimize this maximum expected payoff and the resulting optimal allocation is the same regardless of the observability of the defender's allocation.[22]

Second, the threat τ plays no role in determining the optimal allocation of long-run defensive resources. This is a consequence of treating the terrorists as fully strategic actors. If they are, then what they do depends on how the defender allocates its resources (and thereby affects the underlying vulnerabilities) or on how the terrorists expect the defender to have allocated its resources if the actual allocation is unobservable. Hence,

[22] Powell (2006b) characterizes the equilibria of the game when the defender's allocation is unobservable, showing that the defender's equilibrium strategy in this setting is also the minmax allocation.

the optimal allocation cannot depend on what the terrorists do. In game-theoretic terms, explaining the terrorists' optimal strategy τ is as much a part of the equilibrium analysis of the game as is determining the defender's optimal allocation.

Third, the optimal allocation depends on the terrorists' goals and motivations as formalized in their payoffs G_{st} but not on the defender's losses L_{st}. The reason is that the terrorists do not pursue an attack profile unless it offers the highest expected payoff, so resources spent on other profiles are wasted. And, the terrorists' expected payoff depends on their payoffs or preferences, not the defender's. The defender, therefore, should invest in protecting what terrorists see as high priority targets, not in what the defender sees as high-priority targets except in so far as these are indicative of what the terrorists value.

In a zero-sum setting, the defender's and terrorists' payoffs are mirror images of each other. The defender's losses are identical to the terrorists' gains. (More formally, the relative value of any two sites for the defender is the same as it is for the terrorists: $L_{st}/L_{s't'} = G_{st}/G_{s't'}$ for all sites s and s' and threat scenarios t and t'). High priority targets for the defender therefore are necessarily high priority sites for the terrorists. But mirror imaging is notoriously perilous. The analysis above shows that (i) minmaxing the terrorists gives the optimal allocation regardless of the zero-sum, mirror-imaging assumption, and (ii) underscores the importance of investing in hardening what terrorists see as high expected-payoff opportunities.

Two important limitations need to be emphasized. The defender is assumed above to be certain of the terrorists' payoffs. Bier, Oliveros, and Samuelson (2005) and Powell (2006b) analyze cases in which the defender is unsure of the terrorists' payoffs.

The preceding also assumed that the attack profiles are independent. Spending on one profile has no direct effect on any other. The only effect is indirect; there is less to spend on others. This is at best a first-approximation. Harder defenses "outside the fence" should make all of the threat scenarios associated with what is inside the fence less attractive. Nevertheless, this first approximation appears to parallel what has been done in practice. For example, the Coast Guard's risk assessment tools were designed to evaluate the security "in and around a building or vessel, but... the baseline established by the tools excludes areawide actions the Coast Guard has taken to reduce vulnerabilities in and around ports, such as conducting more patrols, creating operational centers, or establishing security zones in and around key ports" (GAO 2005, 40).

To relax this independence assumption, we could think of perimeter defense or "buffer zone protection" as the probability of reaching a site in order to attack that site in a specific way.[23] The more the defender spends on perimeter defense, the less likely every type of attack on that site is to succeed. Powell (2006b) formalizes this approach in the context of border defense and shows that minmaxing the attacker remains the optimal strategy (although the sites' interdependence changes what the minmax allocation is).[24]

Finally, an extensive review of the related formal literature is out of place here but a brief discussion is in order. The game underlying the analysis above and formalized in the Appendix is a two-actor dynamic game of perfect information in which the defender moves first by allocating its resources. The terrorists move second by choosing their attack profile after observing the defender's allocation. This specification contrasts with many set ups in two ways. First, many allocation games are zero-sum as is, for example, the standard Colonel Blotto game and, second, static rather than dynamic.[25]

In the present context, the zero-sum assumption means that not only do the defender and terrorists rank the sites in the same order (i.e., sites the terrorists most value destroying are the sites the defender most values protecting), but the defender's and terrorists' willingness to run risks is presumed to be the same. This is the substantive implication of the zero-sum requirement that $L_{st}/L_{s't'} = G_{st}/G_{s't'}$ for all sites s and s' and threat

[23] On DHS's Buffer Zone Protection Program, see GAO (2005, 84), DHS (2005), and Moteff (2006, 23).

[24] To sketch the approach, let $\pi_s(p_s)$ be the probability that the terrorists reach site s if the defender spends p_s on area defense for s. Then the expected loss from threat scenarios t and t' on site s are $L_{st}v_{st}(r_{st})\pi_s(p_s)\tau_t$ and $L_{st'}v_{st'}(r_{st'})\pi_s(p_s)\tau_{t'}$. Spending more on perimeter defense in this formulation (increasing p_s) reduces the vulnerability to both types of attack

[25] A game is static if no player can see what any other player does before deciding what to do. A game might be static either because the actors make their decisions simultaneously or they make them at different times but no player can see what any other player has done. Both of these situations are formally equivalent. To my knowledge, no other analysis focuses on the sequential, non zero-sum problem studied here.

scenarios t and t'.

Assuming the game to be zero sum means that the Minimax Theorem can be invoked.[26] This theorem ensures that players always minmax each other in any equilibrium of a two-person, zero-sum games. This theorem also implies that it makes no difference if the decisions are made simultaneously as in the standard Blotto game or sequentially as is assumed here. The optimal allocation will be the same in both cases.

Minmaxing however is generally not equilibrium behavior in nonzero-sum settings, and the equilibrium outcomes of dynamic interactions typically differ from the analogous static interaction. The surprising result established here is that the minmax allocation is the equilibrium allocation regardless of the zero-sum assumption and whether decisions are made sequentially or simultaneously.

The Role of Intelligence and Threat Assessments in Resource Allocation

The framework developed above and the results derived from it raise an important question about the way DHS's risk-management approach incorporates intelligence and basic research on terrorists and terrorist organizations. As described above, these enter DHS's analysis through threat assessments. By contrast, one of the key formal results is that the optimal allocation of long-run defensive resources does not depend on the threat. Indeed, the reverse is true when dealing with strategic terrorists: threats depend on the allocation and its effects on the relative attractiveness of the attack profiles. But if threats do not affect the optimal allocation, does this mean that intelligence and "extensive study and understanding of terrorists and terrorist organizations" (DHS 2006, 47) are irrelevant too? The answer of course is no. They are critically important. But the way they enter the analysis appears to be fundamentally different from the way they enter DHS's approach.

Recall that the risk associated with a specific attack profile is: risk = consequence × vulnerability × threat. However, data limitations have so far prevented DHS from being able to estimate the threats or probabilities associated with different attack profiles.

[26] The Minimax theorem is an old result in game theory back at least as far as Von Neuman and Morgenstern's 1944 classic *Theory of Games and Economic Behavior.* More recent treatments include Owen (1982) and Moulin (1986).

As noted above, Department officials report "that a lack of intelligence data and law enforcement data limits their ability to develop the relative probability for various threat scenarios" (GAO 2005, 76).

The perspective implicit in these comments is that estimating threats is essentially a data and pattern-recognition problem. This perspective may be appropriate for estimating the danger from hurricanes, fires or accidents. But it is inappropriate for threats from strategic actors.

A useful parallel can be found in, of all places, macroeconomic policy during the 1960s. At that time many macroeconomists believed in the Phillips Curve which depicted a relatively stable trade off between lower unemployment and higher inflation. The "price" for pursuing lower unemployment was higher inflation and vice versa. Macroeconomic policy was about picking a point on the Phillips Curve, with Democrats typically preferring points higher on the curve where unemployment is lower but inflation is higher and Republicans usually preferring points further down where inflation is lower but unemployment is higher.

In work that would ultimately lead to a Nobel prize, Robert Lucas showed that there was a problem with this argument, a problem which had profound policy implications. The Phillips Curve which economists were trying to estimate ever more precisely was the joint result of the interaction between monetary and fiscal authorities' policies and the way that firms and other economic actors respond to those policies in their efforts to achieve the their ends. A change in policy would induce a different response, and often an offsetting response, from these economic actors as they continued to try to achieve their unchanged ends in the new policy environment. The net effect of a change in policy would be a "shift" in the Phillips Curve, not a movement up or down a stable curve. The Phillips Curve, according to the Lucas critique, did not represent a stable trade off policymakers could exploit. As soon as they tired, it would shift.[27]

Efforts to estimate threat probabilities parallel efforts to estimate the Phillips Curve.

[27] See for example Lucas (1973). Blanchard (2006) and Romer (2006) provide overviews of the critique.

Even if data limitations could be overcome so that these probabilities could be estimated with some precision, these probabilities reflect the vulnerabilities the attackers faced at the time. Change those vulnerabilities by reallocating resources and the threats will change too as strategic actors respond to the new allocation.

What then does the model analyzed above suggest about the way that intelligence and more basic research on terrorists should be used? If the right question to ask when allocating resources for the long run is not "What are the threat probabilities?" what is the right or at least a better question to pose?

The distinction between payoffs and consequences becomes critically important here. Although the optimal, minmax allocation does not depend on threat probabilities, it does depend very much on the terrorists' goals and motivations as formalized in their payoffs G_j. The defender cannot minimize the terrorists' maximum expected payoff without knowing the terrorists' payoffs. However, consequences and payoffs are not the same thing.

Consequences may be difficult to estimate, but they are "objective." They are akin to physical properties. This kind of attack on that kind of target would kill so many people, injure so many others, and cost so much. The defender will see consequences as losses and the terrorists may see them as gains. But there is no reason to believe that the defender and terrorists *disagree* about what the consequences are.

Payoffs or preferences are different. There is no reason to believe the defender and terrorists *agree* about the sites ranking or share the same attitudes toward risk. The defender's preference ordering over the sites may or may not be the same as the terrorists, and whether it is is an important open question.

All of this suggests for the longer-term allocation of defensive resources, intelligence and basic research on terrorists should focus less on assessing threats and more on assessing payoffs. What is the terrorists' preference ordering over the sites? Which site would they most like to destroy, which site would be their second choice, which their third, and so on? Answering this question may be as hard if not harder than estimating threat probabilities. But at least it is the right question to be trying to answer.

A Simple Policy: A Threshold of Attraction

That the optimal strategy is to minmax the terrorists suggests a simple policy for allocating long-term defensive resources: establish a "threshold of attraction" or, more precisely, a "threshold of expected terrorist gain." If the terrorists expected gain from a site-threat combination exceeds this threshold, the site must be hardened and the vulnerability reduced in order to lower the terrorists' expected gain to the threshold level. Attack profiles with expected payoffs below this threshold would not be in line for additional hardening and more spending. This threshold policy in effect minmaxes the terrorists with the level of the threshold equal to the terrorists' minmax payoff. The more the defender spends, the lower this threshold can be set.

If one were willing to assume that the terrorists' payoffs are equal to the consequences of an attack, then the threshold policy would be equivalent to minimizing the expected consequences of the worst kinds of attack. Thus, the game-theoretic analysis above can be seen as providing some theoretical support for what has been a common policy recommendation. O'Hanlon *et. al.*, for example, argue that "policymakers should focus primarily on those targets at which an attack would involve large numbers of casualties, would entail significant economic costs, or would certainly damage sites of high national significance" (2003, 66). But the game-theoretic analysis also makes it clear that the soundness of this policy prescription depends very much on the degree to which consequences coincide with payoffs.

Finally, protecting against the consequences of the worst kind of attack are sometimes called "weakest-link" policies (e.g., Bier 2005). But one must be careful about the connotation implicit in the image of the weakest link of a chain. It is easy to associate "weakness" with vulnerability, i.e., with the probability that an attack succeeds. The more vulnerable a target, the weaker that link is. But if weakness means vulnerability, terrorists do not go after the weakest link. They go after the highest expected payoff which depends on both the payoff and probability of getting it. The most attractive profile might entail a lower probability of success, i.e., a stronger link, but a much larger payoff if the attack succeeds and the link breaks. Making the most attractive targets less attractive is akin to hardening the weakest link only if we measure weakness in terms of terrorists' expected payoff.

Conclusion

The United States' *National Strategy for Homeland Security* emphasizes that terrorists are strategic actors and that policy must be predicated on this. "One fact dominates all homeland security threat assessments: terrorists are strategic actors. They choose their targets deliberately based on the weaknesses they observe in our defenses and our preparations. They can balance the difficulty in successfully executing a particular attack against the magnitude of the loss it might cause. They can monitor our media and listen to our policymakers as our Nation discusses how to protect itself – and adjust their plans accordingly. Where we insulate ourselves from one form of attack, they can shift and focus on another exposed vulnerability" (White House 2002, 7).

However, current efforts to develop a systematic method for allocating long-term defensive resources based on risk do not treat terrorists as fully strategic actors. Terrorist attacks are treated more like accidents – fires, equipment failures, software glitches – rather than an adversary's determined efforts to strike where the defense is weak and the expected gains are high. Failing to treat terrorists as fully strategic can lead to a significant misallocation of resources.

Unlike risk analysis, game theory does treat actors as fully strategic (and possibly too strategic), and the present analysis suggests a general framework for allocating long-term defensive resources. Terrorists pursue attack profiles that offer the highest expected payoffs given their goals, motivations and capabilities. The defender, therefore, should focus its resources on defending against the most attractive attack. Hardening against this kind of attack makes it less and less attractive. Eventually, what was initially the most attractive attack profile will be no more appealing that what was the second most attractive attack profile. At this point the defender has to focus its resources on both of these profiles so than neither is more attractive than the other. Allocating resources in this way minimizes the terrorists' maximum expected payoff and yields the minmax allocation.

Minmaxing terrorists is at best a guiding principle. At most it provides some bearings when trying to think about the immensely complicated problem of allocating defensive resources against strategic terrorists over the long haul. As such, the minmax principle casts doubt on the way DHS plans to integrate intelligence into its risk assessments. In that plan, intelligence and basic research on terrorist organizations enter the risk calculation through threat estimates of the relative probabilities of different kinds of attack. These estimates are to be based on intelligence and law enforcement data, and that is the problem. Previous attacks also reflect past vulnerabilities. Change those vulnerabilities – which is the goal spending on defense – and the threats change too.

The minmax principle also emphasizes the fundamental distinction between payoffs and consequences. Consequences are more or less objective. Although difficult to measure, they are the physical losses resulting from a successful attack. Payoffs or preferences, by contrast, are subjective. They reflect goals and motivations. The key to minmaxing terrorists is assessing their payoffs. This may be just as hard if not harder than estimating threats from historical data. But at least it is the right question to try to answer for the long-run allocation of defensive resources.

Appendix

The appendix analyzes the game when the terrorists can see how the government allocates its resources. After observing these efforts, the terrorists decide which site to attack and how to attack it. Powell (2006b) studies the interaction when the terrorists cannot observe the defender's allocation and, therefore, can only choose their attack profile on the basis of what they expect the government to do.

The first step in characterizing the equilibria is to simplify the subscript notation which although mnemonically helpful is very cumbersome. Think of the $S \times T$ site-threat combinations as A different attack profiles where $A = S \times T$. Then relabel each site-threat (s, t) as attack profile $j = (s-1)S + t$. The effect of this re-indexing is that the notationally tedious two-dimensional subscripts referring to the site-threats (s, t) for all s and t can now be written as a one-dimensional subscript j where $j = 1, ..., A$.

Using this simpler notation, a strategy for the defender is an allocation of resources across the A attack profiles such that the amount spent hardening against each profile is non-negative and total spent spent on all of the profiles is no more than than R. Formally, a pure strategy is an allocation $r = (r_1, ..., r_A)$ such that $r_j \geq 0$ and $\Sigma_{j=1}^{A} r_j \leq R$. The set of all pure strategies is the set allocations satisfying these constraints. The set of mixed strategies for the defender, Δ, is the set of probability distributions over the set of pure strategies.

A strategy for the terrorists specifies which attack profile they use as a function of the defender's allocation. A mixed strategy is a function $\alpha(r) = (\alpha_1(r), ..., \alpha_A(r))$ where $\alpha_j(r)$ is the probability the terrorists use profile j after observing allocation r. Because the $\alpha_j(r)$ are probabilities, $0 \leq \alpha_j(r) \leq 1$ and $\Sigma_{j=1}^{A} \alpha_j(r) = 1$ for all allocations r.

A subgame perfect equilibrium is a pair of strategies $\widehat{\delta} \in \Delta$ and $\widehat{\alpha}(r)$ such that (i) $\widehat{\delta}$ minimizes the defender's expected loss given that the terrorists are playing according to $\widehat{\alpha}(r)$ and (ii) if the terrorists observe any allocation r, $\widehat{\alpha}(r)$ maximizes their payoff.

Finally, let r^* be the minmax allocation. That is, r^* minimizes the terrorists' maximum payoff. Formally, r^* solves $\min_r \max\{G_j v_j(r_j) : r_1 + \cdots + r_A \leq R\}$.

Then, the formal result underlying the graphical analysis depicted in Figure 3 is:

PROPOSITION 1: *Subgame perfect equilibria exist, and the defender plays r^* with probability one in every subgame perfect equilibrium.*

There are three parts to the proof of this claim and each is stated as a separate lemma. The first establishes that there is only one minmax allocation, i.e., r^* is unique. The second demonstrates existence by constructing a subgame perfect equilibrium in which the defender plays r^*. The third shows if $r' \neq r^*$, then there is an allocation $\widehat{r}(r')$ such that the defender's expected loss to $\widehat{r}(r')$ is strictly less than it is to r'. This ensures that there is a profitable deviation from any strategy that does not put probability one on r^*, namely, play $\widehat{r}(r')$ rather than r' for all r' in the support of the defender's strategy and not equal to r^*.

Proof: Let $M_T(r)$ be the terrorists' highest expected payoff given r, i.e., $M_T(r) \equiv \max\{G_1 v_1(r_1), ..., G_A v_A(r_A)\}$. Then r^* is a minmax allocation if $M_T(r^*) \leq M_T(r)$ for every allocation r. At least one minmax allocation is sure to exist because the $v_j(r)$ are continuous in r and the set of possible allocations is compact.

Arguing by contradiction to show that only one minmax allocation exists, assume that $r^* \neq r'$ both minmax the attacker. Because $r^* \neq r'$, there exists an attack profile j such that $r_j^* \neq r_j'$. Without loss of generality suppose $r_j^* < r_j'$. Then $M_T(r') = M_T(r^*)$ because both r' and r^* minmax the terrorists, and $G_j v_j(r_j^*) > G_j v_j(r_j')$ because $r_j^* < r_j'$. Hence, $M_T(r') = M_T(r^*) \geq G_j v_j(r_j^*) > G_j v_j(r_j')$. Continuity ensures that there is an $\varepsilon > 0$ such that $M_T(r') > G_j v_j(r_j' - \varepsilon)$. This ε of resources can now be distributed across the other attack profiles to form the allocation $\widehat{r}$ where $\widehat{r}_j = r_j' - \varepsilon$ and $\widehat{r}_k = r_k' + \varepsilon/(A-1)$ for all $k \neq j$. Clearly, $M_T(\widehat{r}) < M_T(r')$, contradicting the assumption that r' minimizes $M_T(r)$ and thereby establishing the claim that the minmax allocation is unique. ■

Lemmas 2 and 3 require some additional notation. Take $\mathrm{br}_T(r)$ to be the set of attack profiles offering the terrorists their highest expected payoff given allocation r. In other words, these are the terrorists' best replies to r: $\mathrm{br}_T(r) \equiv \{j : G_j v_j(r_j) = M_T(r)\}$. Let $\mu_D(r, J)$ be the defender's minimum loss given allocation r if the terrorists attack a site in J: $\mu_D(r, J) \equiv \min\{L_j v_j(r_j) : j \in J\}$. Finally, define $\Theta(r, J)$ to be the sites in J at which the defender's losses are minimized: $\Theta(r, J) \equiv \{j : L_j v_j(r_j) = \mu_D(r, J) \text{ and } j \in J\}$.

LEMMA 2: *There are subgame perfect equilibria in which the defender plays the unique minmax allocation r^*.*

Proof: To construct an equilibrium in which the defender plays r^*, consider the playing strategy $\alpha^*(r)$ for the attacker: $\alpha^*_k(r_k) = 1$ if $k = \min\{\Theta(r, \mathrm{br}_T(r))\}$ and $\alpha^*_k(r_k) = 0$ otherwise. In words, the terrorists in equilibrium must maximize their expected payoff which means following an attack profile $j \in \mathrm{br}_T(r)$. If two or more attack profiles give the terrorists their highest payoff (i.e., if $\mathrm{br}_T(r)$ contains two or more profiles), then the terrorists break this indifference by playing the profile that in $\mathrm{br}_T(r)$ that minimizes the defender's expected loss, i.e., $j \in \Theta(r, \mathrm{br}_T(r))$. If two or more attack profiles in $\mathrm{br}_T(r)$ minimize the defender's expected loss, i.e., if $\Theta(r, \mathrm{br}_T(r))$ has two or more elements, the terrorists use the attack profile with the smallest index among the profiles in $\Theta(r, \mathrm{br}_T(r))$. (Any tie breaking rule among the sites in $\Theta(r, \mathrm{br}_T(r))$ would also work.) In brief, the terrorists strike using the profile $j = \min\{\Theta(r, \mathrm{br}_T(r))\}$. The probability of striking with any other attack profile is zero: $\alpha_k(r_k) = 0$ if $k \neq \min\{\Theta(r, \mathrm{br}_T(r))\}$.[28]

Clearly $\alpha^*(r)$ maximizes the terrorists payoff following any r as the terrorists play an attack profile which is an element of $\mathrm{br}_T(r)$ and therefore maximizes their payoff after r. Consequently, $(r^*, \alpha^*(r))$ is a subgame perfect equilibrium if r^* is a best reply to α^*. To see that it is, consider any $r' \neq r^*$ and let $k^* \equiv \min\{\Theta(r^*, \mathrm{br}_T(r^*))\}$ and $k' = \min\{\Theta(r', \mathrm{br}_T(r'))\} \in \mathrm{br}_T(r')$. Then the defender's expected losses to r^* and r' are, respectively, $L_{k^*} v_{k^*}(r_{k^*})$ and $L_{k'} v_{k'}(r'_{k'})$. It follows that the defender's expected loss to playing r' against $\alpha^*(r)$ is strictly greater than its payoff to playing r^*: $L_{k'} v_{k'}(r'_{k'}) > L_{k^*} v_{k^*}(r_{k^*})$.

[28] The terrorists' playing an element of $\Theta(r, \mathrm{br}_T(r))$ when they have more than one best reply to r is akin to what happens in the unique subgame perfect equilibrium of the Rubinstein (1982) bargaining game. In that equilibrium, offers always leave the bargainer receiving the offer indifferent between accepting and rejecting. Despite this indifference, this player does what is most favorable to the other player, namely, accepting the offer. (If a bargainer, say 2, rejects with positive probability when indifferent, then player 1's payoff discontinuously jumps down at an offer of zero. This discontinuity means 1 does not have a best reply to 2's accepting with probability less than one, and thus no equilibrium exists.)

To establish this inequality, observe that $k' \in \mathrm{br}_T(r') \subseteq \mathrm{br}_T(r^*)$. Because r^* minmaxes the terrorists, $r_j^* = 0$ for all $j \notin \mathrm{br}_T(r^*)$. (Otherwise the defender could redistribute some of r_j^* across the sites in $\mathrm{br}_T(r^*)$ and thereby reduce the terrorists' maximum expected payoff still further.) Hence, $G_j v_j(r_j') \leq G_j v_j(0) = G_j v_j(r_j^*) < M_T(r^*) < M_T(r')$ for all $j \notin \mathrm{br}_T(r^*)$ where the last inequality holds because r^* is the unique minmax allocation. But $G_j v_j(r_j') < M_T(r')$ for all $j \notin \mathrm{br}_T(r^*)$ gives $j \notin \mathrm{br}_T(r^*) \Rightarrow j \notin \mathrm{br}_T(r')$. This is equivalent to $j \in \mathrm{br}_T(r') \Rightarrow j \in \mathrm{br}_T(r^*)$ or, alternatively, $\mathrm{br}_T(r') \subseteq \mathrm{br}_T(r^*)$. Thus, $k' \in \mathrm{br}_T(r^*)$.

That $k' \in \mathrm{br}_T(r^*)$ implies $G_{k'} v_{k'}(r_{k'}^*) = M_T(r^*)$. That r^* is the unique minmax allocation also means $M_T(r^*) < M_T(r')$. Moreover, $M_T(r') = G_{k'} v_{k'}(r_{k'}')$ because $k' \in \mathrm{br}_T(r')$. Combining these relations gives $G_{k'} v_{k'}(r_{k'}^*) = M_T(r^*) < M_T(r') = G_{k'} v_{k'}(r_{k'}')$ which yields $v_{k'}(r_{k'}^*) < v_{k'}(r_{k'}')$. Hence, $L_{k^*} v_{k^*}(r_{k^*}) \leq L_{k'} v_{k'}(r_{k'}^*) < L_{k'} v_{k'}(r_{k'}')$ where the weak inequality follows from the fact that k^* is a site in $\mathrm{br}_T(r^*)$ at which the defender's loss is minimized. The strict inequality shows that the defender's expected loss to playing r^* is less than that of playing r'. Hence, r^* is a best response to $\alpha^*(r)$. ■

LEMMA 3: *If $r' \neq r^*$, then there is an allocation $\widehat{r}$ such that the defender's expected loss to $\widehat{r}$ is strictly less than it is to r'.*

The argument takes three steps. The first shows that there is an allocation with a lower expected loss than r' if the terrorists reply to r' with a best response which does not also minimize the defender's expected loss over the set of best replies to r'. That is, there is an allocation with an expected loss less than that of r' if $\alpha_n(r_n') > 0$ for some $n \notin \Theta(r', \mathrm{br}_T(r'))$.

The second step shows that if the terrorists only use attack profiles in $\Theta(r', \mathrm{br}_T(r'))$, i.e., if $\alpha_n(r_n') = 0$ for all $n \notin \Theta(r', \mathrm{br}_T(r'))$, and if $r' \neq r^*$, then r' must be dedicating some resources to attack profiles the terrorists will not use. That is, $r_k' > 0$ for some $k \notin \mathrm{br}_T(r')$. The third step shows that these resources can be reallocated to form an allocation $\widehat{r}$ which gives the defender a lower expected loss that r' does.

Step 1: If $\alpha_n(r_n') > 0$ for some $n \notin \Theta(r', \mathrm{br}_T(r'))$, then there is sure to be an $\widehat{r}$ with a lower expected loss than r'.

Suppose $\alpha_n(r_n') > 0$ for some $n \notin \Theta(r', \mathrm{br}_T(r'))$. Then $n \in \mathrm{br}_T(r')$ because $\alpha_j(r_j') = 0$

for all $j \notin \mathrm{br}_T(r')$. To see that there an allocation with an expected loss less than that from r', take $k' = \min\{\Theta(r', \mathrm{br}_T(r'))\}$. Then $L_{k'}v_{k'}(r'_{k'}) \leq L_k v_k(r'_k)$ for all $k \in \mathrm{br}_T(r')$, and $L_{k'}v_{k'}(r'_{k'}) < L_n v_n(r'_n)$ since $n \in \mathrm{br}_T(r')$ and $n \notin \Theta(r', \mathrm{br}_T(r'))$.

That $\alpha_n(r'_n) > 0$ implies that defender's expected loss $\sum_{j=1}^{A} \alpha_j(r'_j) L_j v_j(r_j)$ is strictly larger than $L_{k'}v_{k'}(r'_{k'})$. Because $\alpha_j(r'_j) = 0$ if $j \notin \mathrm{br}_T(r')$, $\sum_{j=1}^{A} \alpha_j(r'_j) L_j v_j(r_j) = \sum_{j \in \mathrm{br}_T(r')} \alpha_j(r'_j) L_j v_j(r_j) \geq [1 - \alpha_n(r'_n)] L_{k'}v_{k'}(r'_{k'}) + \alpha_n(r'_n) L_n v_n(r'_n) > L_{k'}v_{k'}(r'_{k'})$ where the strict inequality is sure to hold because $\alpha_n(r'_n) > 0$ and $L_{k'}v_{k'}(r'_{k'}) < L_n v_n(r'_n)$.

By deviating to an allocation that devotes slightly less than $r'_{k'}$ to site k', the defender can induce the terrorist to attack k' for sure. This gives the defender an expected loss of slightly more than $L_{k'}v_{k'}(r'_{k'})$ but still less than its expected loss from playing r'. Formally, suppose the defender plays $\widehat{r}$ where $\widehat{r}_{k'} = r_{k'} - \varepsilon$, $\widehat{r}_n = r'_n + \varepsilon$, $\widehat{r}_j = r'_j$ for all $j \neq k', n$. The attacker's unique best reply to $\widehat{r}$ is to attack k' for sure which imposes an expected loss of $L_{k'}v_{k'}(r_{k'} - \varepsilon)$. Continuity ensures we can take an ε small enough to guarantee $\sum_{j \in \mathrm{br}_T(r')} \alpha_j(r'_j) L_j v_j(r_j) \geq [1 - \alpha_n(r'_n)] L_{k'}v_{k'}(r'_{k'}) + \alpha_n(r'_n) L_n v_n(r'_n) > L_{k'}v_{k'}(r_{k'} - \varepsilon)$. Thus, the expected loss to $\widehat{r}$ is strictly less than to r'. ■

It remains to be shown that the defender can do better than r' if the terrorists only play attack profiles in $\Theta(r', \mathrm{br}_T(r'))$, i.e., if $\alpha_n(r'_n) = 0$ for all $n \notin \Theta(r', \mathrm{br}_T(r'))$.

Step 2: If the terrorists only play attack profiles in $\Theta(r', \mathrm{br}_T(r'))$ *after* $r' \neq r^*$, *i.e., if* $\alpha_n(r'_n) = 0$ *for all* $n \notin \Theta(r', \mathrm{br}_T(r'))$, *then* r' *devotes resources to defending against attack profiles outside* $\mathrm{br}_T(r')$*:* $\sum_{j \in \mathrm{br}_T(r')} r'_j < R$.

By definition, $G_j v_j(r'_j) = M_T(r')$ for all $j \in \mathrm{br}_T(r')$. The fact that r^* is the unique minmax allocation also implies $M_T(r') > M_T(r^*)$. Hence, $G_j v_j(r'_j) > M_T(r^*)$ for all $j \in \mathrm{br}_T(r')$. Further, $M_T(r^*) = G_j v_j(r^*_j)$ for all $j \in \mathrm{br}_T(r^*)$. As shown above in the proof of Lemma 2, $r' \neq r^*$ implies $\mathrm{br}_T(r') \subseteq \mathrm{br}_T(r^*)$ which means $M_T(r^*) = G_j v_j(r^*_j)$ for all $j \in \mathrm{br}_T(r')$. Combining these results yields $G_j v_j(r'_j) > G_j v_j(r^*_j)$ for all $j \in \mathrm{br}_T(r')$. These inequalities leave $v_j(r'_j) > v_j(r^*_j)$ and therefore $r^*_j > r'_j$ for all $j \in \mathrm{br}_T(r')$. Summing resources yields $R \geq \sum_{j \in \mathrm{br}_T(r')} r^*_j > \sum_{j \in \mathrm{br}_T(r')} r'_j$. Hence, $r'_k > 0$ for some $k \notin \mathrm{br}_T(r')$. ■

Step 3: That $r'_k > 0$ *for some* $k \notin \mathrm{br}_T(r')$ *ensures that there is an* $\widehat{r}$ *with a lower expected loss than* r'*.*

Because the terrorists only play an attack profiles in $\Theta(r', \mathrm{br}_T(r'))$, the defender's expected loss to r' is $\mu_D(r', \mathrm{br}_T(r'))$ where, recall, $\mu_D(r', \mathrm{br}_T(r')) = \min\{L_j v_j(r'_j) : j \in \mathrm{br}_T(r')\}$. Recall further that in order to maximize their payoff after observing r', the terrorists only use attack profiles in $\mathrm{br}_T(r')$ which means $\alpha_k(r'_k) = 0$ for all $n \notin \mathrm{br}_T(r')$. This along with Step 2 guarantees that there is a profile k such that $k \notin \mathrm{br}_T(r')$, $\alpha_k(r'_k) = 0$, and $r'_k > 0$.

To construct an allocation with a lower expected loss than r', distribute some of r'_k over the sites in $\mathrm{br}_T(r')$ to obtain r'' without changing the set of best replies, i.e., so that $\mathrm{br}_T(r') = \mathrm{br}_T(r'')$. Continuity ensures this can be done. Because $r''_j > r'_j$ for all $j \in \mathrm{br}_T(r') = \mathrm{br}_T(r'')$, it follows that $\mu_D(r' \,\mathrm{br}_T(r')) > \mu_D(r' \,\mathrm{br}_T(r'')$. Let $k' = \min\{\Theta(r', \mathrm{br}_T(r'))\}$ and $k'' = \min\{\Theta(r'', \mathrm{br}_T(r''))\}$. Then $L_{k'} v_{k'}(r'_{k'}) = \mu_D(r', \mathrm{br}_T(r')) > \mu_D(r'', \mathrm{br}_T(r'')) = L_{k''} v_{k''}(r''_{k''})$.

Now suppose the defender plays $\widehat{r}$ where $\widehat{r}$ is derived from r'' by investing slightly less in k'' and thereby inducing the attacker to strike k''. Formally, let $\widehat{r}_{k''} = r_{k''} - \varepsilon$, $\widehat{r}_n = r''_n + \varepsilon_n$ for all $n \in \mathrm{br}_T(r'')$, $n \neq k''$ with $\varepsilon = \sum_{n \in \mathrm{br}_T(r''), n \neq k''} \varepsilon_n$. Then, $G_{k''} v_{k''}(\widehat{r}_{k''}) > G_{k''} v_{k''}(r''_{k''}) = G_n v_n(r''_n) > G_n v_n(\widehat{r}_n)$ for all $n \in \mathrm{br}_T(r'')$. Hence, the attacker's best reply to $\widehat{r}$ is to attack k'', i.e., $\mathrm{br}_T(\widehat{r}) = \{k''\}$ so that $\alpha_{k''}(\widehat{r}_{k''}) = 1$. The defender's expected loss to $\widehat{r}$ is $L_{k''} v_{k''}(r_{k''} - \varepsilon)$. Continuity and the fact that $L_{k'} v_{k'}(r'_{k'}) > L_{k''} v_{k''}(r''_{k''})$ guarantee we can take ε small enough so that $L_{k'} v_{k'}(r'_{k'}) > L_{k''} v_{k''}(r_{k''} - \varepsilon)$. But the defender's payoff to r' is $\mu_D(r', \mathrm{br}_T(r')) = L_{k'} v_{k'}(r'_{k'})$ which means that the expected loss to $\widehat{r}$ is lower than it is from r'. ■

This completes the proof of Lemma 3 which, along with Lemmas 1 and 2, establishes Proposition 1.

References

9/11 Public Discourse Project Final Report. 2005. "Final Report on the 9/11 Commission Recommendations." Available at www.9-11pdp.org. Accessed on December 29, 2005.

Belke, James C. 2000. "Chemical Accident Risks in U.S. Industry: A Preliminary Analysis of Accident Risk Data from U.S. Hazardous Chemical Facilities." Washington, D.C.: U.S. Environmental Protection Agency, September 25.

Bier, Vicki. 2005. "Game-Theoretic and Reliability Methods in Counter-Terrorism and Security," In Sallie Keller-McNulty, Alyson Wilson, Yvonne Armijo, eds., *Modern Statistical and Mathematical Methods in Reliability.* Hackensack, NJ: World Scientific Publishing Co.

Bier, Vicki, Aniruddha Nagaraj, Vinod Abhichandani. 2004. "Protection of Simple Series and Parallel Systems with Components of Different Values," *Reliability Engineering and System Safety* 87:313-23 .

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

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

Blanchard, Oliver. 2006. *Macroeconomics*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall. pp 177-80.

Bueno de Mesquita, Ethan. 2005. "Politics and the Suboptimal Provision of Counterterror," *International Organization* (forthcoming).

Chertoff, Michael. 2005. "Remarks for Secretary Michael Chertoff U.S. Department of Homeland Security George Washington University Homeland Security Policy Institute," Washington, D.C. Department of Homeland Security, March 16. Available at: www.dhs.gov/dhspublic/ display?content=4391

Coughlin, P.J. 1992. "Pure Strategy Equilibria in a Class of Systems Defense Games," *International Journal of Game Theory* 20:195-210.

DHS. 2005. "Fiscal Year 2005 Buffer Zone Protection Program." Washington, D.C.: Department of Homeland Security, Office of Inspector General.

DHS. 2006. "Progress in Developing the National Asset Database." Washington, D.C.: Department of Homeland Security, Office of Inspector General (June).

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

GAO. 2005. "Risk Management." Washington, D.C.: Government Accountability Office. Available at: www.gao.gov/cgi-bin/getrpt?GAO-06-91.

Gross, Alfred O. and R. A. Weber. 1950. "A Continuous Colonel Blotto Game," Rand Memorandum. Available at: www.rand.org/pubs/research_memoranda/2006/RM408.pdf

Haimes, Yacov. Y. 2004. *Risk Modeling, Assessment, and Management*, 2nd ed. Hoboken, NJ: Wiley.

Hausken, Kjell. 2002. "Probabilistic Risk Analysis and Game Theory," *Risk Analysis* 22:17-27.

Kardes, Erim. 2005. "Survey of Literature on Strategic Decision Making in the Presence of Adversaries," Center for Risk and Economic Analysis of Terrorism Events, University of Southern California, report 05-006. March 15. Available at: http://www.usc.edu /dept/create/reports/Report05006.pdf.

Lucas, Robert. 1973. "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review,* 63(June):326-34.

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

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

Moulin, Herve. 1986. *Game Theory for the Social Sciences*, 2nd ed. New York: New York University Press.

NIPP. 2006. DHS. "Draft National Infrastructure Protection Plan, v2.0." Washington, D.C.: Department of Homeland Security. Available at: www.fas.org /irp/agency/dhs/ nipp110205.pdf.

Owen, Guillermo. 1982. *Game Theory*, 2nd ed. Orlando. FL: Academic Press.

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

Paté-Cornell, Elisabeth and Seth Guikema. 2002. "Probabilistic Model of Terrorist Threats," *Military Operations Research* 7:5-20.

Powell, Robert. 2006a. "Allocating Resources to Defend Against Terrorist Attacks with Private Information abut Vulnerability." Manuscript, Department of Political Science, UC Berkeley.

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

Romer, David. 2006. *Advanced Macroeconomics*, 3rd ed. Boston: McGraw-Hill.

Rosendorff, Peter and Todd Sandler. 2004. "Too Much of a Good Thing? The Proactive Response Dilemma," *Journal of Conflict Resolution*, 48(October): 657-671.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model," *Econometrica* 50:97-109.

Sagan, Scott D. 2004. "The Problem of Redundancy Problem: Why More Nuclear Security Forces May Produce Less Nuclear Security," *Risk Analysis* 24:935-46.

Sandler, Todd. 2005. "Collective Versus Unilateral Response to Terrorism," *Public Choice* July ???.

Shubik, Martin and Robert James Weber. 1981. "Systems Defense Games: Colonel Blotto, Command and Control," *Naval Research Logistics Quarterly* 28(2): 281-87.

Tukey, John W. 1949. "A Problem of Strategy," *Econometrica* 17: 73.

Von Neumann, John and Oskar Morgenstern. 1944. *Games and Economic Behavior.* Princeton: Princeton University Press.

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

White House. 2003. "National Strategy for the Physical Protection of Critical Infrastructures and Key Assets." Washington, D.C.: White House, February. Available at: www.whitehouse. gov/pcipb/physical.html

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