



Calhoun: The NPS Institutional Archive
DSpace Repository

Conferences and Events at NPS

Conferences and Events at NPS

2008-06

How Much and on What?

Powell, Robert

<http://hdl.handle.net/10945/51785>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

HOW MUCH AND ON WHAT?

Robert Powell*

March 2008

Abstract:

How much should a defender spend on defense and how should it allocate those resources across the sites it is trying to protect? This paper analyzes a model in which a defender first has to decide how much to spend on defense and what to spend it on. The more that a defender devotes to protecting a specific site, the less likely an attack on that site is to succeed and, crucially, the lower the marginal return to investing in attacking that site. After the defender moves, the attacker decides how much effort to devote to attacking each site. Three key conclusions result: First, the questions of how much to spend and what to spend it on are “separable.” However much the defender decides to spend, it should allocate those resources in the same general way. Second, a very simple principle or algorithm determines the optimal allocation. The defender *minmaxes* the attacker’s marginal gains, i.e., allocates its resources in the way that minimizes the attacker’s maximum marginal gain from exerting additional effort. Third, the defender is in effect a Stackelberg leader. The optimal level of spending takes into account how the defender’s allocation affects the attacker’s effort and generally is that level of spending which equates the marginal benefits of additional spending with the marginal cost of diverting these resources from other social ends.

* Travers Department of Political Science, 210 Barrows Hall, UC Berkeley Berkeley, CA 94720-1950 (RPowell@Berkeley.edu).

HOW MUCH AND ON WHAT?

Two factors make the problem of defending against terrorists especially daunting. First, as the *National Strategy for Homeland Security* emphasizes, “terrorists are strategic actors” (White House 2002, 7). No one believes that hardening the levies around New Orleans affects the probability that another hurricane like Katrina will strike New Orleans again or Miami rather than New Orleans. But strategic actors do try to strike where the defense is weak and the expected gains are high. Protecting one site may 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).

Second, relative to the large number of potential targets, resources are scarce. We cannot defend everything. As Department of Homeland Security Secretary Michael Chertoff assessed the situation shortly after taking office, “Although we have substantial resources to provide security, these resources are not unlimited. Therefore, as a nation, we must make tough choices about how to invest finite human and financial capital to attain the optimal state of preparedness” (2005b). Echoing the *9/11 Commission*, Secretary Chertoff has emphasized throughout his tenure that these scarce resources should be allocated on the basis of risk. “Risk management must guide our decision making as we examine how we can best organize to prevent, respond and recover from an attack” (2005a).

This paper offers a game-theoretic framework for analyzing two related questions. How much should a defender spend on defending against a strategic attacker, i.e., a terrorist group, instead of devoting those resources to other social ends like health care or education? Second, how should a defender allocate however much it decides to spend among the multiple sites it is trying to protect?

In the model, a defender first has to decide how much to spend on defense and what to spend it on. The more that a defender devotes to protecting a specific site, the less likely an attack on that site is to succeed and, crucially, the lower the marginal return to investing in

attacking that site. After the defender moves, the attacker decides how much effort to devote to attacking each site. In order to focus on the fundamental ideas, insights, and intuitions, we simplify matters by assuming that all of the sites the defender is trying to protect are identical. But the results generalize to a setting in which some or all of the sites differ from others.

Three key conclusions follow from the analysis. First, the questions of how much to spend and what to spend it on are “separable.” However much the defender decides to spend, it should allocate those resources in the same general way.

Second, a very simple principle or algorithm determines the optimal allocation. Suppose that the defender has decided to spend a specific amount on defense but has not yet allocated it. Given this null allocation, a strategic attacker will direct its efforts to the site offering the highest marginal return on that effort. The defender therefore should invest in hardening this site and reducing the attacker’s expected return from trying to attack it. The more the defender spends on this site, the less vulnerable it becomes and the lower the expected return to an attack. Eventually, this site will be no more attractive than what was initially the second most attractive site. That is, both offer the same marginal return on the attacker’s effort to strike them. At this point, the defender must invest in protecting both sites so the neither is more attractive than the other. The more the defender spends on these two sites, the lower their vulnerability and the less attractive targets they become. At some point, these sites are no more attractive than what was originally the third most attractive site. From here on the defender must invest in guarding all three sites so that that no one site is any more attractive than the other two. The defender continues to allocate its resources in this way by spending so as to make the most attractive profile as unattractive as possible. In brief, the defender *minmaxes* the attacker’s marginal gains, i.e., allocates its resources in the way that minimizes the attacker’s maximum marginal gain from exerting additional effort.

The third conclusion is that the defender is in effect a Stackelberg leader. The optimal level of spending takes into account how the defender’s allocation affects the attacker’s effort and generally is that level of spending which equates the marginal benefits of additional spending with the marginal cost of diverting these resources from other social ends. In principle, the defender may be able to spend enough to induce the attacker to exert zero effort in carrying out an attack. But this may require such a high level of defense spending that it is not optimal.

The next section presents the game-theoretic model. It also links the basic components of the model to the critical elements of risk management, consequence, vulnerability, and threat. The subsequent section characterizes the defender's optimal level of spending and the attacker's optimal level of effort. There follows a discussion of the comparative statics describing how the optimal levels of spending and effort change as the underlying parameters change. The last section discusses the generality of the results and an appendix sketches a game-theoretic analysis of the model.

A Model

A defender has N identical sites to protect and must decide how much to spend on defending them and how to distribute those resources across the sites it is trying to guard. The more the defender dedicates to a given site, the "harder" that site becomes and the less likely an attack on that site is to succeed. After observing the defender's allocation, an attacker decides how much effort to devote to attacking each site. The more effort the attacker devotes to striking a specific site, the more likely the attack on that site is to succeed.

A strategy for the defender in this game simply specifies how much the defender spends on each site. In symbols, it is an allocation $r = (r_1, \dots, r_N)$ where $r_j \geq 0$ is the amount allocated to site j . The total spent on defense is implicitly defined by $R = \sum_{j=1}^N r_j$. Analogously, the attacker's strategy specifies how much effort it will put into attacking each site after observing any possible allocation r . More precisely, a strategy for the attacker is a function $e(r) = (e_1(r), \dots, e_N(r))$ where $e_j(r) \geq 0$ is the effort the attacker puts into striking site j .

Let $\lambda > 0$ denote the loss the defender suffers if a site is successfully attacked. If the attack fails, the defender's loss is zero. (We assume for simplicity that an attack either succeeds or fails.) The attacker gains of $\gamma > 0$ if a site is successfully attacked and zero if the attack fails.

The more the defender spends on a site, the less likely an attack on that site is to succeed. Formally, let $V_j(r_j, e_j)$ be the probability that an attack on site j succeeds if the defender spends r_j on hardening that site and the attacker expends effort e_j on hitting that site. $V_j(r_j, e_j)$ is increasing in r_j and decreasing in e_j .

We now make an important assumption which greatly simplifies the analysis. The vulnerability of a site is multiplicatively separable in effort. That is, we can write the vulnerability V_j as $V_j(r_j, e_j) = v_j(r_j)e_j$ where the function v_j depends solely on r_j . The substantive significance of this assumption is that the marginal effect of additional effort on the vulnerability of a site is independent of the level of effort already being exerted. That is, $\partial V_j / \partial e_j$ is independent of e_j or $\partial^2 V_j / \partial e_j^2 = 0$. Stating this assumption more formally:

ASSUMPTION 1 (SEPARABILITY): *The vulnerability of each site j can be written as $V_j(r_j, e_j) = v_j(r_j)e_j$.*

Assumption 1 is critical to the analysis. A second simplifying assumption makes the algebra easier but is not substantively critical. We assume the v_j is linear in resources, i.e., $v_j(r_j) = 1 - vr_j$. If the defender devotes nothing to site j , then $r_j = 0$, $v_j(0) = 1$, and an attack on this site is sure to succeed. The parameter v measures the marginal effect that additional resources have on the vulnerability of a site. The larger v , the greater the effect of additional spending on the vulnerability of a site.¹

Spending on defense means diverting resources from other social ends. These costs are assumed to rise and at an increasing rate as R increases. More concretely, let take the cost to devoting R to defense to be $c_D(R) = k_D R^2$. The parameter k_D measures the social opportunity cost of spending on defense rather than some other social goal. The higher k_D , the more costly defense is relative to other social priorities and the faster these costs rise as R increases.

Resource are scarce for the attacker too. Let $E = \sum_{j=1}^N e_j$ denote the total effort expended on attacking. Then the cost of exerting this effort is assumed to be $c_A(E) = k_A E^2 / 2$ where k_A measures the relative difficulty the attacker has in exerting the effort needed to carry out an attack.

¹ We assume v is small enough that $v_j(r_j) > 0$ over the substantive relevant range of resource allocations.

In light of all of this, the defender's expected loss if it allocates r and the attacker replies with $e(r)$ is $L(r, e(r)) = \sum_{j=1}^N \lambda v_j(r_j) e_j(r) + c_D(R)$. The attacker's payoff is

$$G(r, e(r)) = \sum_{j=1}^N \gamma v_j(r_j) e_j(r) + c_A(E).$$

The basic elements of this model broadly correspond to the three key components of risk-management which are vulnerability, threat, and consequence. Vulnerability "is the probability that a particular attack will succeed against a particular target" (GAO 2005, 25), and this is what $v_j(r_j)$ is in the model. Threat "is the probability that a specific target is attacked in a specific way" (Willis et. al. 2005, 8). In this formulation, the amount of effort the attacker puts into hitting a site serves as a proxy for the probability of an attack on that site. Finally, λ formalizes the defender's "range of loss or damage that can be expected from a successful attack" (NIPP 2006a, 41).²

Note, however, that nothing in the risk-management framework corresponds to the cost of spending on defense rather than something else, i.e., nothing corresponds to $c_D(R)$ in the model. At its best, risk-management provides guidance on how one should allocate a fixed amount of resources. It says little or nothing about how to determine the optimal amount to spend on defense.

The Optimal Levels of Resources and Effort

This section describes the intuitions underlying the equilibrium outcome. The appendix offers a more detailed game-theoretic discussion of the equilibrium. The fact that the sites are identical suggests the defender will distribute however much it decides to spend evenly across the N sites. In symbols, $r_j = R/N$. This leaves the defender with losses of

$$L = \sum_{j=1}^N \lambda [1 - vR/N] e_j + c_D(R) = \lambda [1 - vR/N] \sum_{j=1}^N e_j + c_D(R) = \lambda [1 - vR/N] E + c_D(R). \text{ The}$$

² Strictly speaking, λ is a von Neumann-Morgenstern utility which is related to economic losses but is not the same thing.

attacker's payoffs are $G = \sum_{j=1}^N \gamma [1 - vR/N] e_j + c_A(E) = \gamma [1 - vR/N] E - c_A(E)$. Note that the only choice left to determine is how much the defender spends and level of effort E .

The defender will choose R partly based on the defender's anticipation of how the attacker will react. To determine this, consider the attacker's decision after seeing that the defender has allocated R to defense and spread these resources evenly across the N sites. The attacker chooses E to maximize its gain G given this allocation. Taking the derivative of G with respect to E and setting it equal to zero gives the first-order condition:

$$0 = \frac{\partial G}{\partial E}$$

$$0 = \gamma \left[1 - \frac{vR}{N} \right] - k_A E$$

$$E = \frac{\gamma}{k_A} \left[1 - \frac{vR}{N} \right]$$

where recall $c_A(E) = k_A E^2 / 2$.³ Thus, for any given allocation R , the attacker's optimal level of effort is $E^*(R) \equiv (\gamma / k_A) [1 - vR/N]$. This level of effort equates the marginal gain from additional effort, $\gamma [1 - vR/N]$, with the marginal cost $c'_A(E) = k_A E$ (see the second equality above). As expected, there is an inverse relation between the defender's spending and the attacker's effort. As R increases, $E^*(R)$ declines.

The function $E^*(R)$ describes how the attacker alters its level of effort in response to varying levels of defense spending. Anticipating that the attacker will respond in this way, the defender's losses to R are $L = \lambda [1 - vR/N] E^*(R) + c_D(R)$. The optimal allocation R minimizes these losses. To solve for this, differentiate L with respect to R to obtain:

³ This critical point is sure to maximize G since $\partial^2 G / \partial E^2 = -k_A < 0$.

$$\frac{\partial L}{\partial R} = \underbrace{\frac{\partial \lambda [1 - vR/N]}{\partial R} E^*(R)}_{\text{defensive effect of increasing } R} + \underbrace{\lambda [1 - vR/N] \frac{\partial E^*(R)}{\partial R}}_{\text{deterrent effect of increasing } R} + \underbrace{c'_D(R)}_{\text{cost effect of increasing } R}$$

The expressions on the right side of this equality offer a useful decomposition of the effects of an increase in defense spending into the *defensive effect*, the *deterrent effect*, and the *cost effect*. The first term, the defensive effect of an increase in R , is the effect that spending more on hardening the sites has on the defender's expected losses *given that the attacker's level of effort remains the same*. The second term might be thought of as the deterrent effect. This is the decrease in the defender's losses resulting from the attacker's decision to invest less effort in mounting an attack. Finally, the third term is the increase in losses due to the greater expenditure on defense.

Substituting the expressions for $E^*(R)$ and $c_R(R)$ and then solving for the optimal allocation R^* gives:⁴

$$R^* = \frac{\gamma \lambda v N}{\gamma \lambda v^2 + k_A k_D N^2}.$$

This then implies that the optimal level of effort $E^*(R^*)$ is

$$E^* = \frac{\gamma k_D N^2}{\gamma \lambda v^2 + k_A k_D N^2}$$

The defender's losses are:

$$L^* = \frac{\gamma \lambda k_D N^2}{\gamma \lambda v^2 + k_A k_D N^2}.$$

⁴ This critical point is sure to minimize L since $\partial^2 L / \partial R^2 > 0$.

In sum, when the defender anticipates how the attacker responds to the defender's actions, the optimal level of spending is R^* , the attacker exerts E^* , and the defender's expected loss is L^* .

Comparative Statics

How do the optimal level of spending and the defender's losses vary with the parameters of the model? Suppose, for example, that the opportunity cost of spending on defense increases (i.e., k_D goes up). This makes defense spending more costly and, intuitively, seems likely to result in lower spending and higher losses. Moreover, these higher losses will be due in part to the fact that the attacker will exert more effort to carrying out an attack. Formally, the effect of an increase in k_D on L^* is:

$$\frac{\partial L^*}{\partial k_D} = \underbrace{\frac{\partial \lambda [1 - vR^* / N]}{\partial k_D} E^*}_{\text{direct effect on losses from an attack due to lower spending}} + \underbrace{\lambda [1 - vR^* / N] \frac{\partial E^*}{\partial k_D}}_{\text{indirect effect on losses from an attack due to changes in effort}} + \underbrace{\frac{\partial c_D(R^*)}{\partial k_D}}_{\text{cost effect of } k_D}$$

Inspection of the expression for R^* shows that the level of defense spending decreases as the cost of diverting those resources from other social purposes k_D increases ($\partial R^* / \partial k_D < 0$). Hence the direct effect of an increase in k_D is positive. Spending goes down, sites are not more vulnerable, and the defender's losses from an attack rise.

The same is true of the indirect effect. As k_D increases, the defender's spending decreases, and this induces the attacker to increase its effort E^* . Finally, the cost effect by itself is ambiguous as a larger k_D makes any given level of spending more costly but the higher k_D also reduces the level of spending R^* . Nevertheless, the first two effects swamp the potentially ambiguous third effect and the defender's losses increase as k_D increases (the expression for L^* is clearly increasing in k_D).

Similarly, the defender's losses are increasing in (i) the losses λ the defender suffers if a site is destroyed, (ii) the gains γ the attacker gets from destroying a site (since this induces greater effort), and (iii) the number of sites N . Defense spending R^* is increasing in the gains γ and losses λ . It decreases as the costs k_A and k_D rise. Finally, the attacker's effort is increasing in the attacker's gains γ and decreasing in its costs k_A .

Some Generalizations

The formal analysis has centered on a model in which the sites are identical. But many of the results generalize beyond this. The critical assumption is the separability assumption which recall is that the vulnerability of every site j can be written as $V_j(r_j, e_j) = v_j(r_j)e_j$. As long as this holds, the results go through.⁵ More precisely, let λ_j and γ_j be the defender's loss and the attacker's gain if site j is successfully attacked. Then the results described above hold even if these losses and gains differ across the sites (i.e., λ_j need not equal λ_k and γ_j need not equal γ_k), the defender's losses differ from the attacker's gain (i.e., λ_j need not equal γ_k), and the functions relating vulnerability to resources, $v_j(r_j)$, differ from site to site.⁶

To outline the analysis in the more general case, recall that the marginal return the attacker obtains from investing effort in attacking site j is $\gamma_j v_j(r_j)$. Thus the attacker will only invest effort in attacking the sites offering the highest return on this investment, namely those sites k such that $\gamma_k v_k(r_k) = \max\{\gamma_j v_j(r_j)\}$. Given that the marginal return to effort is $\max\{\gamma_j v_j(r_j)\}$, the attacker exerts the level of effort E^{**} that equates the marginal return on this

⁵ Some mild technical assumptions are also needed. The loss function L is kinked and possibly discontinuous at finitely many value of R . The needed technical conditions ensure that $\partial^2 L / \partial R^2 > 0$ everywhere else.

⁶ See Powell (2008) for an analysis of a more general game that allows each site to differ from the others. The attacker in Powell's model chooses the probability of attacking rather than the level of effort. But the separability assumption ensures that these two formulations are essentially equivalent.

effort to the marginal cost, i.e., $\max\{\gamma_j v_j(r_j)\} = c'_A(E^{**})$. It follows that if the defender decides to spend R on defense, it will allocate those resources so as to minimize the attacker's maximum marginal return to effort. That is, the defender distributes R in the way that minimizes $\max\{\gamma_j v_j(r_j)\}$. The defender now chooses the allocation R that minimizes the defender's losses in light of this reaction.

Conclusions

The *National Strategy for Homeland Security* emphasizes that terrorists are strategic, and this poses at least two questions. When allocating scarce resources to defend against strategic attacker's, how much should the defender spend on defense and how should it allocate those resources across the sites it is trying to protect? Strategic interaction often makes resource-allocation problems extraordinarily difficult to analyze, but that turns out not to be the case here. Taking the effects of strategic interaction is relatively straightforward and yields three key findings.

First, the defender's level and allocation problems are separable. However much the defender decides to spend, it should allocate those resources in the same general way. Second, the defender should allocate however much it decides to spend so as to minmax the attacker's return on its effort. Finally, the defender's strategic position is analogous to that of a Stackelberg leader. Taking into account how the defender's allocation will affect the attacker's effort, the optimal level defense spending generally equates the marginal benefits of additional spending with the marginal cost of diverting these resources from other social ends.

Appendix

This appendix sketches a game-theoretic analysis of the model. A subgame perfect equilibrium is a strategy profile $(r^*, e^*(r))$ such that (i) the effort allocation $e^*(r)$ maximizes the attacker's payoff $G(r, e(r)) = \sum_{j=1}^N \gamma[1 - vr_j]e_j(r) + k_A E^2 / 2$ for every resource allocation r , and (ii) the resource allocation r^* minimizes the defender's loss

$$L(r, e(r)) = \sum_{j=1}^N \lambda[1 - vr_j]e_j(r) + k_D R^2 \text{ given that the attacker plays according to } e^*(r).$$

Solving the game by starting with the last decision and working up the game tree to the first decision, consider the attacker's decision following any allocation r . It wants to choose e_j so as to maximize $\sum_{j=1}^N \gamma[1 - vr_j]e_j + k_A E^2 / 2$ where $E = \sum_{j=1}^N e_j$. The separability assumption plays a crucial role at this point. Given that the attacker's marginal return to increasing e_j , i.e., $\gamma[1 - vr_j]$, is independent of e_j , this maximization problem has a very simple solution.

The attacker will only invest effort in the site or sites offering the highest marginal return on that investment. That is, the attacker will only invest effort in going after k if $\gamma[1 - vr_k] = \max\{\gamma[1 - vr_j] : j = 1, \dots, N\}$. Let $T(r)$ denote the set of the sites offering the attacker its highest expected marginal return: $T(r) = \{k : \gamma[1 - vr_k] = \max\{\gamma[1 - vr_j] : j = 1, \dots, N\}\}$. Then the attacker invests no effort in attacking sites outside T , i.e., $e_i = 0$ for $i \notin T(r)$. This implies that the attacker's payoff reduces to:

$$\begin{aligned} G(r, e(r)) &= \sum_{k \in T(r)} \gamma[1 - vr_k]e_k - \frac{k_A E^2}{2} \\ &= \gamma[1 - vr_k] \sum_{k \in T(r)} e_k - \frac{k_A E^2}{2} \\ &= \gamma[1 - vr_k]E - \frac{k_A E^2}{2} \end{aligned}$$

where the second line follows from the first because $\gamma[1 - vr_k] = \gamma[1 - vr_j]$ for all $j, k \in T(r)$.

Differentiating the previous expression with respect to E yields shows that the optimal level of effort given allocation any r is $E^*(r) = (\gamma / k_A)[1 - vr_k]$ where k is any element in $T(r)$. If there are two or more sites that offer the maximum return on the attacker's effort, i.e., if $T(r)$ contains two or more sites, the allocation $E^*(r)$ across the sites in $T(r)$ is arbitrary because the return on every allocation is the same. Since the attacker's allocation across the sites in $T(r)$ is arbitrary, let $t(r)$ be the site in $T(r)$ with the smallest index, that is, $t(r) = \min\{j : j \in T(r)\}$, and suppose that the attacker allocates all of $E^*(r)$ to $t(r)$.

Turning to the defender's strategy, the problem for the defender is to select the allocation $r^* = (r_1^*, \dots, r_N^*)$ given the attacker's strategy of allocating $E^*(r)$ to $t(r)$. Formally, the defender wants to choose r^* to minimize $L(r, e^*(r)) = \lambda[1 - vr_{t(r)}]E^*(r) + k_D R^2$. This is equivalent to selecting the allocation r that solves $\min_r \{ \max_{j=1, \dots, N} \{ (\lambda\gamma / k_A)(1 - vr_j)^2 \} + k_D R^2 \}$. Because all of the sites are identical, the minmax allocation of R is to distribute R evenly across the N sites. That is, the optimal distribution of R is to set $r_j = R / N$. This means that the defender's losses reduce to $(\lambda\gamma / k_A)(1 - vR / N)^2 + k_D R^2$. Differentiating with respect to R yields the optimal allocation $R^* = \gamma\lambda vN / (\gamma\lambda v^2 + k_A k_D N^2)$.

In sum, the subgame perfect equilibrium allocation entails a level of spending R^* spread evenly across the N sites. The attacker's total level of effort is $E^*(R^*) = \gamma k_D N^2 / (\gamma\lambda v^2 + k_A k_D N^2)$.

References

- 9/11 Commission Report. 2003. *Final Report of the National Commission on Terrorist Attacks upon the United States*, New York: Norton.
- Chertoff, Michael. 2005a. "Remarks for Secretary Michael Chertoff U.S. Department of Homeland Security George Washington University Homeland Security Policy Institute," Washington, DC Department of Homeland Security, March 16. Available at www.dhs.gov/dhspublic/display?content=4391.
- Chertoff, Michael. 2005b. "Secretary Michael Chertoff U.S. Department of Homeland Security Second Stage Review Remarks." Washington, DC: Department of Homeland Security, July 13. Available at www.dhs.gov/dhspublic/display?content=4597.
- GAO. 2005. "Risk Management." Washington, DC: Government Accountability Office. Available at: www.gao.gov/cgi-bin/getrpt?GAO-06-91.
- 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.
- Powell, Robert. 2008. "Defending Against Strategic Attackers: Deciding How Much to Spend and on What," Manuscript, Department of Political Science, UC Berkeley.
- White House. 2002, "National Strategy for Homeland Security." Washington, DC: White House, July. Available at www.whitehouse.gov/homeland/book/index.html.
- Willis, Henry H., Andrew R. Morral, Terrance K. Kelly, and Jamison J. Medby. 2005. *Estimating Terrorism Risk*. Santa Monica, CA: Rand Corporation.