Risk Analysis, Vol. 28, No. 2, 2008 DOI: 10.1111/j.1539-6924.2008.01023.x

Optimal Stopping Analysis of a Radiation Detection System to Protect Cities from a Nuclear Terrorist Attack

Michael P. Atkinson,[1] Zheng Cao,[2] and Lawrence M. Wein[3]*

We formulate and analyze an optimal stopping problem concerning a terrorist who is attempting to drive a nuclear or radiological weapon toward a target in a city center. In our model, the terrorist needs to travel through a two-dimensional lattice containing imperfect radiation sensors at some of the nodes, and decides at each node whether to detonate the bomb or proceed. We consider five different scenarios containing various informational structures and two different sensor array topologies: the sensors are placed randomly or they form an outer wall around the periphery of the city. We find that sensors can act as a deterrent in some cases, and that the government prefers the outer wall topology unless the sensors have a very low detection probability and the budget is tight (so that they are sparsely deployed).

KEY WORDS: Homeland security; Stackelberg game; stochastic dynamic programming

1. INTRODUCTION

A nuclear weapon (made of uranium or plutonium) detonated by a terrorist in a large U.S. city could kill a half-million people and cause 1 trillion dollars in direct economic damage (Bunn *et al.*, 2003). Although this threat is deemed "real and urgent" (Bunn *et al.*, 2003), a more likely scenario is for terrorists to assemble a radiological dispersal device, or so-called dirty bomb, containing radiological material such as cesium, which would inflict much less damage but would nonetheless wreak considerable havoc. Because the majority of nuclear material in the former Soviet Union remains vulnerable to theft (Bunn *et al.*, 2003) and smuggling nuclear or radiological material into a U.S. port in a shipping container is fairly easy (Flynn, 2000; Stana, 2004), this article analyzes our last line of defense, which is to detect an assembled nuclear or radiological weapon as it is driven into a city and to provide timely and effective interdiction. Indeed, the U.S. government is in the process of developing (U.S. Department of Homeland Security, 2005) and deploying (including a pilot test in New York City; Lipton, 2007) such detection-interdiction systems in its largest cities. In this article, we focus on the game-theoretic aspects of these systems by formulating and analyzing five versions of an optimal stopping problem that make different assumptions regarding the topology of the radiation sensor deployment and the amount of information available to the terrorist. As explained in Section 5, we embed our results into a broader model (which allows for optimizing over the number of sensors deployed) in a companion article (Wein & Atkinson, 2007) to analyze the entire detection-interdiction system; the goal of the companion article is to provide a rough-cut feasibility analysis of such systems.

The detection model in Section 2 is an optimal stopping problem on a finite two-dimensional lattice

[1] Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305, USA.
[2] Physics Department, Stanford University, Stanford, CA 94305, USA.
[3] Graduate School of Business, Stanford University, Stanford, CA 94305, USA.
* Address correspondence to Lawrence M. Wein, Graduate School of Business, Stanford University, Stanford, CA 94305, USA; lwein@stanford.edu.

that has imperfect radiation sensors at some of its nodes (e.g., highway exit ramps). The terrorist drives a vehicle carrying a bomb from node (0, 0) toward a target at node (N, N), and the damage from a detonated bomb at node (m, n) is an increasing linear function of $m + n$; other topologies are possible, such as a circular topology. Before entering each node, he decides whether to proceed in the hope of getting closer to the target, or to stop and detonate the bomb. If the terrorist is detected at a node then he is immediately interdicted and with a specified probability he successfully detonates the bomb before being captured. In Section 3, we analyze five different scenarios depending on the topology of the sensor array (sensors are either placed randomly or in an outer wall formation) and the information known by the terrorist (e.g., the detection probability of the sensors is known or updated in a Bayesian manner). The main goal of this analysis is to determine whether either the random topology or the outer wall topology consistently dominates the other. Numerical results for the five scenarios are presented in Section 4 and concluding remarks are offered in Section 5.

Network interdiction is an active field of study within operations research; see Morton *et al.* (2007) for a more thorough review of this literature than we present here. Many of the early problems considered maximizing an adversary's shortest path, or minimizing an adversary's maximum flow, through a deterministic network. Several authors (Wollmer, 1964; Washburn & Wood, 1995; Pan *et al.*, 2003), motivated by military operations, or smuggling of nuclear materials or drugs, allow the inspector to locate detectors at certain arcs and permit the adversary to choose a path through the network to maximize his probability of evading detection. Relative to these articles, our model makes the simplifying assumption that the detection probability of sensors is independent of location. While this assumption seems reasonable in our context because the same technology is used throughout the system, it may be violated if vehicle speeds at different highway ramps (dictated by the ramp curvature) and/or length of red lights at traffic intersections differ appreciably. Our grid is also more restrictive than the general networks considered in the network interdiction literature, but it does allow for closed-form solutions, which is helpful when embedding the results in the model of Wein and Atkinson (2007). On the other hand, our model is more complex than those in the aforementioned articles in that it considers the terrorist's optimal stopping problem and considers a variety of sensor array topologies and informational structures. However, some recent network interdiction studies consider uncertain network topology (Hemmecke *et al.*, 2003), interdictor's uncertainty about the smuggler's origin-destination pair (Morton *et al.*, 2007), and different perceptions of the two players about the network parameters (Morton *et al.*, 2007), while Bailey *et al.* (2006) models the adversary's problem as a Markov decision process. Finally, although using a much different model in a different setting—pedestrian suicide-bombers—Kaplan and Kress (2005) is the only other study besides Wein and Atkinson (2007) that we are aware of that analyzes a model that takes into account sensors, terrorist behavior, and interdiction.

2. PROBLEM DESCRIPTION

The city is represented by a two-dimensional square lattice, where the nodes are street intersections or highway entrance/exit ramps, and the edges are road segments; it may also be possible to insert sensors directly into the highway pavement, although they would have to be engineered for durability. Imperfect sensors that are capable of detecting nuclear or radiological material with a specified probability are deployed at some of the nodes, and other nodes may have phantom sensors, which are not functional but are indistinguishable from real sensors from the terrorist's viewpoint; we refer to the collection of sensors as an array. The terrorist starts at node (0, 0) and travels toward node (N, N). The damage caused by a bomb detonated at node (m, n) is assumed to be a linear function of $m + n$ (this symmetry allows us to analyze the problem in one dimension rather than two); because of the uncertainty in the exact nature of the damage function, we carried out an analogous study with a damage function that is exponential in $m + n$, and compare the results from the linear and exponential functions in Section 4.4. The use of a linear (or exponential) damage function is a gross simplification. There are four main effects of a nuclear weapon: shock and blast, thermal radiation, initial nuclear radiation, and residual nuclear radiation (Glasstone & Dolan, 1977). All four exposures are nonlinear functions of distance (thermal radiation varies inversely with distance squared, and radiation exposure varies inversely with distance squared given scattering and decreases exponentially with no scattering) and depend greatly on the yield of the bomb. Moreover, the dose-response effects and the population gradient are also nonlinear. Although the yield of a bomb detonated by a terrorist is highly uncertain, the instantaneously fatal effects are on the order of miles (e.g., for the Hiroshima bomb) to tens of miles, and

Table I. The Five Scenarios

Scenario	Topology	Terrorist's Knowledge of Sensor Array	Terrorist's Knowledge of Detection Probability
RK	Random	Known P(real sensor)	Known
RB	Random	Bayesian P(real sensor)	Bayesian
OKK	Outer wall	Observable wall	Known
OKB	Outer wall	Observable wall	Bayesian
OUB	Outer wall	Bayesian P(real sensor)	Bayesian

the residual effects (which could also be eventually fatal) are on the order of tens of miles. Although the linear damage function with a 1-to-10 scale (see Equation (20)) allows policymakers to easily internalize our results (by interpreting damage as relative distance), we believe that our results would need to be combined with detailed simulation models of the four effects of a nuclear weapon (including dose-response models and spatial population data) to provide comprehensive input to policymakers; such an effort is beyond the scope of this article.

At each node, the terrorist makes the decision of either detonating the bomb at that node or moving forward based on his estimate of the probability of reaching the next node without being detected, which depends on his perception of whether the next node contains a real sensor and his belief about the detection probability of the sensor. In this section, we make no attempt to explicitly model the interdiction process. This is done in a companion article (see Section 5). Here, we simply assume that if the terrorist is detected by the sensor at a certain node, he would try to detonate the bomb at that node, and he would succeed in doing so (before being killed or captured) with a specified probability.

In Sections 3.1–3.5, we solve the terrorist's optimal stopping problem under the five scenarios described in Table I, which vary according to the topology of the sensor array and the terrorist's knowledge about the array. The first two scenarios in Table I assume a random topology, where each node contains a real sensor with a certain probability, and phantom sensors are placed at all other nodes. In scenario RK (R = random, K = known), the terrorist knows the probability that each node has a real sensor, but not the actual location of the sensors (because each node has either a real or a phantom sensor), and he also knows the detection probability of the real sensors; we do not investigate the scenario in which the terrorist can see the randomly placed sensors (i.e., there are no phantom sensors) because then the terrorist can simply avoid the sensors if the density of real sensors is not sufficiently high (by standard results in percolation theory (Grimmett, 1999), the threshold density is 0.5 in the asymptotic, large-network limit). In scenario RB (R = random, B = Bayesian), the terrorist knows neither the fraction of real sensors nor the detection probability of real sensors, and updates his probability of successfully traversing a node in a Bayesian manner. The other three scenarios have an outer wall topology, where the outermost layers of nodes (i.e., nodes (m, n) such that m or n are near 0) are deployed with sensors that form a wall around the periphery of the city. In scenario OKK (O = outer wall, K = known topology, K = known probability), the terrorist can observe the wall (i.e., there are no phantom sensors) and knows the detection probability of the sensors. In scenario OKB (outer wall, known, Bayesian), the terrorist can observe the wall but updates the detection probability of the sensors in a Bayesian fashion. Finally, in scenario OUB (outer wall, unknown topology, Bayesian), the terrorist's information is the same as in scenario RB: he cannot observe the wall because of the presence of phantom sensors and does not know the detection probability of sensors, and he updates his probability of successfully traversing a node in a Bayesian manner.

In our view, scenarios RB and OUB possess the most realistic informational structures in the case of a mildly sophisticated terrorist, and their direct comparison allows us to determine the best (i.e., damage-minimizing) array topology. Indeed, we view this problem as a Stackelberg game (Gibbons, 1992): the government is the leader and designs the sensor array, and the terrorist, as the follower who cannot observe the array or the detection probability of sensors, solves an optimal stopping problem with Bayesian updating to maximize the expected damage. However, the analysis of these five scenarios allows us not only to solve this Stackelberg game, but also to assess the value of phantom sensors in the outer wall topology, and to compute the value to the terrorists of having

prior information about the array topology and sensor sensitivity. A sophisticated terrorist could conceivably gain this knowledge (at the risk of arousing the government's suspicion) by surveillance and by probing the network with legal shipments of radiological materials prior to the actual attack. In a somewhat different setting, Bier (2007) discusses the effects of various informational assumptions in attacker-defender games.

3. ANALYSIS OF THE OPTIMAL STOPPING PROBLEM

The five scenarios described in Table I are analyzed in Sections 3.1–3.5.

3.1. RK Scenario: Random Array, Known Passing Probability

In scenario RK, sensors are deployed randomly and the terrorist has knowledge about the probability that a sensor is real and the detection probability of the real sensors. The square lattice has $(N+1)^2$ nodes indexed by $m, n = 0, \ldots, N$. Because the damage caused by a detonated bomb at node (m, n) depends only on $m + n$, by the symmetry along the diagonal line segment connecting the points $(0, 0)$ and (N, N), we can reduce the state of the system to the one-dimensional quantity $k = m + n$, which can take on the values $0, 1, \ldots, 2N$. Each node has a real sensor with probability p_s, and has a phantom sensor otherwise, independent of all other nodes. Each real sensor has a false negative rate f, i.e., its detection probability is $1 - f$, and we assume that detection at different nodes are statistically independent events, which is partially justified by the fact that neutron emissions are very bursty (Hage & Cirafelli, 1985), background noise can vary across time and space, and different nodes have different sensors. To the extent that a small amount of positive correlation may exist among sensor results, we may be slightly overestimating the efficacy of these multi-layer detection systems. False positives are introduced in Wein and Atkinson (2007), where we discuss the current capabilities of radiation sensors, and the factors that influence the sensitivity and specificity. If we let p_u denote the probability that a terrorist traverses a node without getting detected, then

$$p_u = 1 - p_s + p_s f. \tag{1}$$

By our informational assumptions, the terrorist knows the value of p_u in this scenario.

In our model, the terrorist makes his detonate vs. proceed decision just before he passes through the sensor at each node. Suppose the terrorist manages to arrive at (but not yet pass through) state k (i.e., node (m, n) where $n + m = k$) without being caught. Now he has two choices. He can either detonate the bomb at this node or move to the next node (at state $k + 1$) in an attempt to increase the damage inflicted. If he detonates the bomb, it causes damage $ak + b$, where $a > 0$, $b \geq 0$. If he instead proceeds through state k, he avoids detection with probability p_u, in which case he travels to the next node at state $k + 1$. If the terrorist is detected as he passes through state k, then with probability q he detonates the bomb before being killed or captured and causes $ak + b$ in damage. Taken together, if $V(k)$ is the optimal value function (i.e., maximum expected damage if in state k) then the terrorist's optimal stopping problem can be formulated as (e.g., Section 3.4 in Bertsekas, 1976)

$$V(k) = \max \left\{ \underbrace{ak+b}_{\text{detonate}}, \underbrace{(1-p_u)q(ak+b)}_{\text{proceed, detected}} + \underbrace{p_u V(k+1)}_{\text{proceed, undetected}} \right\} \quad \text{for} \quad k = 0, \ldots, 2N-1, \tag{2}$$

with boundary condition

$$V(2N) = 2aN + b. \tag{3}$$

The following proposition (proved in Section A of the Appendix), gives the solution to this problem.

PROPOSITION 1. *The optimal solution to Equations (2)–(3) is for the terrorist to proceed to state*

$$k^* = \min \left\{ 2N, \max \left\{ 0, \left\lceil \frac{p_u}{(1-p_u)(1-q)} - \frac{b}{a} \right\rceil \right\} \right\} \tag{4}$$

and detonate the bomb just before passing through k^* if he has yet to be caught. The expected damage under the optimal policy is

$$U = (1-q)p_u^{k^*}(ak^* + b) + aqp_u \frac{1 - p_u^{k^*}}{1 - p_u} + qb. \tag{5}$$

Proposition 1 is proved by showing that if there is a state k in which it is optimal to proceed, then it is also optimal to proceed in states $0, \ldots, k-1$. We then derive Equation (4) by solving the conditions in which it is preferable to proceed in states $0, \ldots, k^* - 1$ and to detonate in states $k^*, \ldots, 2N$.

3.2. RB Scenario: Random Array, Bayesian Passing Probability

We now turn to scenario RB, in which the terrorist updates his perception about p_u as he moves through the network. Since the event that the terrorist gets caught or not at any given node is a Bernoulli random variable, the beta-binomial conjugate prior is a natural choice for modeling the terrorist's updating process. Furthermore, we assume that the terrorist has no prior information about the government's deployment of real sensors or the detection probability of real sensors, and uses the uniform distribution as an uninformative prior to represent his initial perception. That is, just before passing through state k, the terrorist has successfully passed through states $0, \ldots, k-1$, and believes that $E(p_u) = \frac{k+1}{k+2}$ (Berger, 1985). However, if the prior is informative with beta distribution parameters α and β, then tractability is maintained and $E(p_u) = \frac{\beta+k}{\alpha+\beta+k}$ after successfully passing states $0, \ldots, k-1$ (see problem 5 on page 287 of Berger, 1985); we return to the informative prior case at the end of this subsection.

The uninformative prior yields the optimal stopping problem

$$V(k) = \max\left\{ak + b, \frac{q}{k+2}(ak+b) + \frac{k+1}{k+2}V(k+1)\right\} \text{ for } k = 0, \ldots, 2N-1,$$
$$V(2N) = 2aN + b.$$

PROPOSITION 2. *It is optimal for the terrorist to either detonate the bomb in state 0 or to proceed to state* $2N$ *and detonate it there.*

The intuition behind Proposition 2, which is proved in Section C of the Appendix, is that if the terrorist decides to proceed in state 0 and the attempt is successful, then the posterior probability of p_u becomes stochastically larger, which gives the terrorist an even stronger incentive to proceed in state 1, and this argument holds as he proceeds toward the target. We show in Section B of the Appendix that $k^* = 2N$ if

$$b \le \sum_{i=0}^{2N-1} \frac{q(ai+b)}{(i+1)(i+2)} + \frac{2aN+b}{2N+1}, \tag{6}$$

$$\approx \frac{2Nq(b-2a)}{2N+1} + aq\ln 4N + \frac{2aN+b}{2N+1}. \tag{7}$$

Isolating q in Equation (7) and using a second analytical approximation gives

$$q \ge \frac{2N\left(\frac{b}{a}-1\right)}{(2N+1)\ln 4N + 2\frac{b}{a}N - 4N}, \tag{8}$$

$$\approx \frac{\frac{b}{a}-1}{\left(\ln 4N + \frac{b}{a} - 2\right)} \text{ for large } N. \tag{9}$$

By Equation (5), the expected damage under the terrorist's optimal policy is

$$U = \begin{cases} b & \text{if } k^* = 0; \\ (1-q)p_u^{2N}(2aN+b) & \\ \quad + aqp_u\dfrac{1-p_u^{2N}}{1-p_u} + qb & \text{if } k^* = 2N, \end{cases} \tag{10}$$

where p_u continues to represent the true probability of passing through a node undetected.

In the more general case where the prior distribution is informative with parameters α and β, Proposition 2 does not hold in general because to prove it we would need to show that (generalizing inequalities (C.3) and (C.5) in Section C of the Appendix)

$$\frac{2N-1+\frac{b}{a}}{2N+\frac{b}{a}} \ge \frac{\beta+2N-1}{\alpha+\beta+2N-1-q\alpha} \tag{11}$$

implies

$$\frac{2N-2+\frac{b}{a}}{2N-1+\frac{b}{a}} \ge \frac{\beta+2N-2}{\alpha+\beta+2N-2-q\alpha}, \tag{12}$$

but there are (α, β) pairs that violate Equations (11)–(12). However, this more general case can be solved using a generic $O(N)$ dynamic programming algorithm (see Section 4.5).

3.3. OKK Scenario: Outer Wall, Known Topology, Known Detection Probability

In the OKK scenario, the terrorist observes the precise location of the real sensors and knows these sensors' detection probability. Hence, he correctly perceives that his passing probability is f at a node with a real sensor and is 1 at all other nodes. In the outer wall topology with K layers of sensors, i.e., with

thickness K, real sensors are placed at all nodes (m, n) such that $\min(m, n) \leq K - 1$. The terrorist, upon observing the wall, chooses a path through the wall that only goes through nodes satisfying $\max(m, n) \leq K$, so as to minimize the number of real sensors that are confronted. The associated optimal stopping problem with K layers of sensors is

$$\begin{aligned} V(k) &= \max\{ak + b, q(1 - f)(ak + b) \\ &\quad + fV(k + 1)\} \quad \text{for } k = 0, \ldots, 2K - 1, \\ V(2K) &= V(2N) = 2aN + b. \end{aligned}$$

The solution to this optimal stopping problem is derived in Section D of the Appendix. The optimal stopping point is

$$k^* = \max\left\{0, \left\lceil \frac{f}{(1 - f)(1 - q)} - \frac{b}{a} \right\rceil\right\} \quad \text{if } \frac{1 - q(1 - f)}{f} \geq \frac{2N + \frac{b}{a}}{2K - 1 + \frac{b}{a}}. \tag{13}$$

If the inequality in Equation (13) is violated, then it is preferable to proceed rather than detonate in state $2K - 1$, and

$$\begin{aligned} V(2K - 2) = \max\{&a(2K - 2) + b, \\ &q(1 - f)(a(2K - 2) + b) \\ &+ f[q(1 - f)(a(2K - 1) + b) \\ &+ f(2aN + b)]\}. \end{aligned} \tag{14}$$

If the optimal decision at each state is to proceed then $k^* = 2N$; however, if this is not the case, there exists a state M such that the optimal decision is to detonate in state M but the optimal decision is to proceed in states $M + 1$ and beyond. Then the optimal stopping point is

$$k^* = \min\left\{M, \max\left\{0, \left\lceil \frac{f}{(1 - f)(1 - q)} - \frac{b}{a} \right\rceil\right\}\right\}. \tag{15}$$

The expected damage under the optimal policy is given by Equation (5) if $k^* < 2N$. If $k^* = 2N$, then the expected damage is

$$\begin{aligned} U &= \sum_{i=0}^{2K-1} (1 - f) f^i q(ai + b) \\ &\quad + \left[1 - \sum_{i=0}^{2K-1} (1 - f) f^i\right](2aN + b), \\ &= (2aN + b) f^{2K} + aqf \frac{1 - f^{2K}}{1 - f} \\ &\quad + qb - q(2aK + b) f^{2K}. \end{aligned} \tag{16}$$

3.4. OKB Scenario: Outer Wall, Known Topology, Bayesian Detection Probability

In scenario OKB, the terrorist can see the outer wall of real sensors, but does not know the false negative probability f of these sensors. As in Section 3.2, the terrorist uses a beta-binomial conjugate pair with an uninformative prior to update the likelihood that he can pass through a node that has a real sensor. Once inside the K layers of the outer wall, he knows there are no sensors and that he can travel freely. The optimal stopping formulation in this scenario is

$$\begin{aligned} V(k) &= \max\left\{ak + b, \frac{q}{k+2}(ak + b) \right. \\ &\quad \left. + \frac{k+1}{k+2} V(k + 1)\right\} \quad \text{for } k = 0, \ldots, 2K - 1, \\ V(2K) &= V(2N) = 2aN + b. \end{aligned}$$

The analysis of this scenario closely mimics the analysis in Section 3.2, and the details are omitted. The optimal stopping point is either the first or last node, i.e., Proposition 2 carries over to this scenario. An analysis similar to Equations (6)–(7) implies that $k^* = 2N$ if

$$\begin{aligned} b &\leq \sum_{i=0}^{2K-1} \frac{q(ai + b)}{(i + 1)(i + 2)} + \frac{2aN + b}{2K + 1}, \\ &\approx \frac{2Kq(b - 2a)}{2K + 1} + aq \ln 4K + \frac{2aN + b}{2K + 1}, \end{aligned} \tag{17}$$

which simplifies to

$$\frac{b}{a} \leq \frac{q}{1 - q}\left(\frac{2K + 1}{2K} \ln 4K - 2\right) + \frac{N}{K(1 - q)}. \tag{18}$$

In the case where K is small compared to N, i.e., the wall is thin, inequality (18) is approximated by

$$K \leq \frac{aN}{b(1 - q)}, \tag{19}$$

and the right-hand side of Equation (19) is likely to be an upper bound on the true threshold. The expected

damage is $U = b$ if $k^* = 0$ and is given by Equation (16) if $k^* = 2N$.

3.5. OUB Scenario: Outer Wall, Unknown Topology, Bayesian Passing Probability

The terrorist's optimal stopping problem in scenario OUB is the same as that in Section 3.2: in both cases, the terrorist has the same information about the sensor array, and whether sensors are deployed randomly or in an outer wall does not affect his optimal stopping policy. However, the resulting damage is affected by the array topology. In this case, the damage is given by Equation (16), although the condition under which this equation applies is given in Section 3.2.

4. COMPUTATIONAL STUDY OF THE OPTIMAL STOPPING PROBLEM

The experimental design for our computational study is described in Section 4.1, the optimal stopping policies for the terrorist in the five scenarios are computed and discussed in Section 4.2, and the scenarios are compared in Section 4.3. A comparison of the results under linear vs. exponential damage appears in Section 4.4, and a brief outline of how to extend our analysis to more general damage functions and more general sensor placements is provided in Section 4.5.

4.1. Experimental Design

The government's decision variable is the probability that a node has a real sensor (p_s) in the random array topology, and the wall thickness (K) in the outer wall topology. The remaining parameters in our model are the network size (N), the damage parameters (a and b), the probability that the bomb will be detonated during interdiction (q), and the false negative probability of the real sensors (f). We consider two networks: a large network with $N = 50$ and a small network with $N = 5$; see Table II for a succinct description of our experimental design. The network size relates to how ambitious the deployment is and on the actual topology of a city. Cities that are laid out on a grid, or contain a sprawl of highways, may have a large value of N. Highway systems with very few checkpoints (e.g., cities near waterways) would have a small value of N. We set the damage parameters so that the damage equals 1 in state 0 and 10 at the target state $2N$, i.e.,

$$b = 1, \quad a = \frac{9}{2N}. \tag{20}$$

A terrorist possessing a nuclear weapon would likely be capable of detonating the bomb from the driver's seat with a remote detonation device. A suicide bomber would only be stopped if he is killed or otherwise physically prevented from detonating the device. We consider two values of q, which are 0.5 and 0.9, with the latter value probably being more practical in the absence of persuasive intelligence information that the truckdriver is a suicide bomber. As explained in Section 5, the false negative probability of real sensors depends on a variety of factors, and can essentially vary from 0 to 1. Because the optimal stopping probability strictly depends on f in only one of the five scenarios, we need not specify a value for f at this point. Hence, we consider four cases in Section 4.2: large and small networks, and $q = 0.5$ and 0.9.

Table II. The Experimental Design

Parameters	Large Network	Small Network
N	50	5
Number of nodes	2601	36
State space	$0, \ldots, 100$	$0, \ldots, 10$
Damage slope a	0.09	0.9
Damage intercept b	1	1
Detonation probability q	0.5, 0.9	0.5, 0.9

4.2. Calculation of Optimal Policies

In this subsection, the exact (i.e., not using the approximations to simplify the conditions separating $k^* = 0$ from $k^* = 2N$) optimal policies for the parameter values in Table II are summarized in Table III, and are discussed and compared to our analytical approximations.

RK Scenario. Proposition 1 implies that the optimal policy in the RK scenario satisfies

$$k^* = \begin{cases} 0 \quad \text{if } p_u \le 1 - \dfrac{1}{12.1 - 11.1q} \\ \qquad = \begin{cases} 0.847 & \text{if } q = 0.5; \\ 0.526 & \text{if } q = 0.9; \end{cases} \\ 100 \quad \text{if } p_u \ge 1 - \dfrac{1}{112 - 111q} \\ \qquad = \begin{cases} 0.982 & \text{if } q = 0.5; \\ 0.917 & \text{if } q = 0.9; \end{cases} \end{cases} \tag{21}$$

Scenario	Detonation Probability q	Large Network	Small Network
RK	0.5	$k^* = 0$ if $p_u \le 0.847$	$k^* = 0$ if $p_u \le 0.357$
RK	0.5	$k^* = 2N$ if $p_u \ge 0.982$	$k^* = 2N$ if $p_u \ge 0.835$
RK	0.9	$k^* = 0$ if $p_u \le 0.526$	$k^* = 0$ if $p_u \le 0.100$
RK	0.9	$k^* = 2N$ if $p_u \ge 0.917$	$k^* = 2N$ if $p_u \ge 0.503$
RB		$k^* = 2N$ if $q \ge 0.705$	$k^* = 2N$ if $q \ge 0.048$
OKK		Figure 1	Figure 1
OKB	0.5	$k^* = 2N$ if $K \le 10$	$k^* = 2N$
OKB	0.9	$k^* = 2N$	$k^* = 2N$
OUB		$k^* = 2N$ if $q \ge 0.705$	$k^* = 2N$ if $q \ge 0.048$

Table III. Summary of the Exact Optimal Policies.

for the large network, and

$$k^* = \begin{cases} 0 & \text{if } p_u \le 1 - \dfrac{1}{2.1 - 1.1q} \\ & = \begin{cases} 0.357 & \text{if } q = 0.5; \\ 0.100 & \text{if } q = 0.9; \end{cases} \\ 10 & \text{if } p_u \ge 1 - \dfrac{1}{11.1 - 10.1q} \\ & = \begin{cases} 0.835 & \text{if } q = 0.5; \\ 0.503 & \text{if } q = 0.9; \end{cases} \end{cases} \tag{22}$$

for the small network. In Equations (21)–(22), $k^* \in (0, 2N)$ if p_u is between the two threshold values. Hence, a sufficiently small passing probability p_u, which itself depends on the probability p_s a sensor is real and a real sensor's false-negative probability f, deters the terrorist from proceeding to the target. As expected, this deterrent is stronger (i.e., the p_u threshold for $k^* = 0$ is larger) for larger networks (because the terrorist has more sensors to traverse before getting close to the target) and for smaller values of the detonation probability q. The passing probability needs to be well above 0.5 for the terrorist to proceed to the target in all cases. Hence, the sensor array can be far from perfect and still provide a deterrent in the RK scenario.

RB and OUB Scenarios. The optimal policy in the RB and OUB scenarios (recall that they have identical solutions) does not depend on the true value of p_u, which is assumed unknown in these Bayesian scenarios. In these scenarios, the optimal policy can be described as a threshold in terms of q: the terrorist proceeds to the target as long as the detonation probability is sufficiently high. Hence, there is no need to explicitly consider the specific values of $q = 0.5$ and 0.9. By Equation (9), the optimal policy in the RB scenario is approximated by

$$k^* = 2N \quad \text{if } q \ge \frac{\dfrac{b}{a} - 1}{\ln 4N + \dfrac{b}{a} - 2} = \begin{cases} 0.701 \text{ for the large network;} \\ 0.053 \text{ for the small network.} \end{cases} \tag{23}$$

These threshold values are very close to the corresponding exact values in Table III. As in the RK scenario, the large network provides a stronger deterrent than the small network in Equation (23).

OKK Scenario. Among the five scenarios, the optimal solution is most complicated in the OKK scenario, in which the terrorist has full information about the array and the detection probability. There are two cases to consider, depending on whether the inequality in Equation (13) is satisfied or violated. Moreover, it is possible for $k^* \in (0, 2N)$ as in the RK scenario, and this stopping point is a function of q and f. Fig. 1 shows the indifference curves for $k^* = 0$ and $k^* = 2N$ in terms of f and K, in four cases (two values of q and for the large and small networks). In Fig. 1, the terrorist chooses $k^* = 0$ (i.e., detonates the bomb at the first node) in the region in the graph below the bottom curve, chooses $k^* = 2N$ (i.e., moves to the last node and detonates) in the region above the top curve, and chooses an intermediate node $k^* \in (0, 2N)$ in the region between the two curves in Fig. 1. These plots show that the terrorist is more apt to proceed to the target when interdiction is ineffective (q is large), the detection probability is low (f is large), or the network is small. Three of these four curves (all except the $q = 0.9$, small network case) are highly concave, implying that the wall thickness initially plays a strong deterrent role, but has decreasing value once the wall reaches a certain thickness (approximately $K = 10$ layers in the large network).

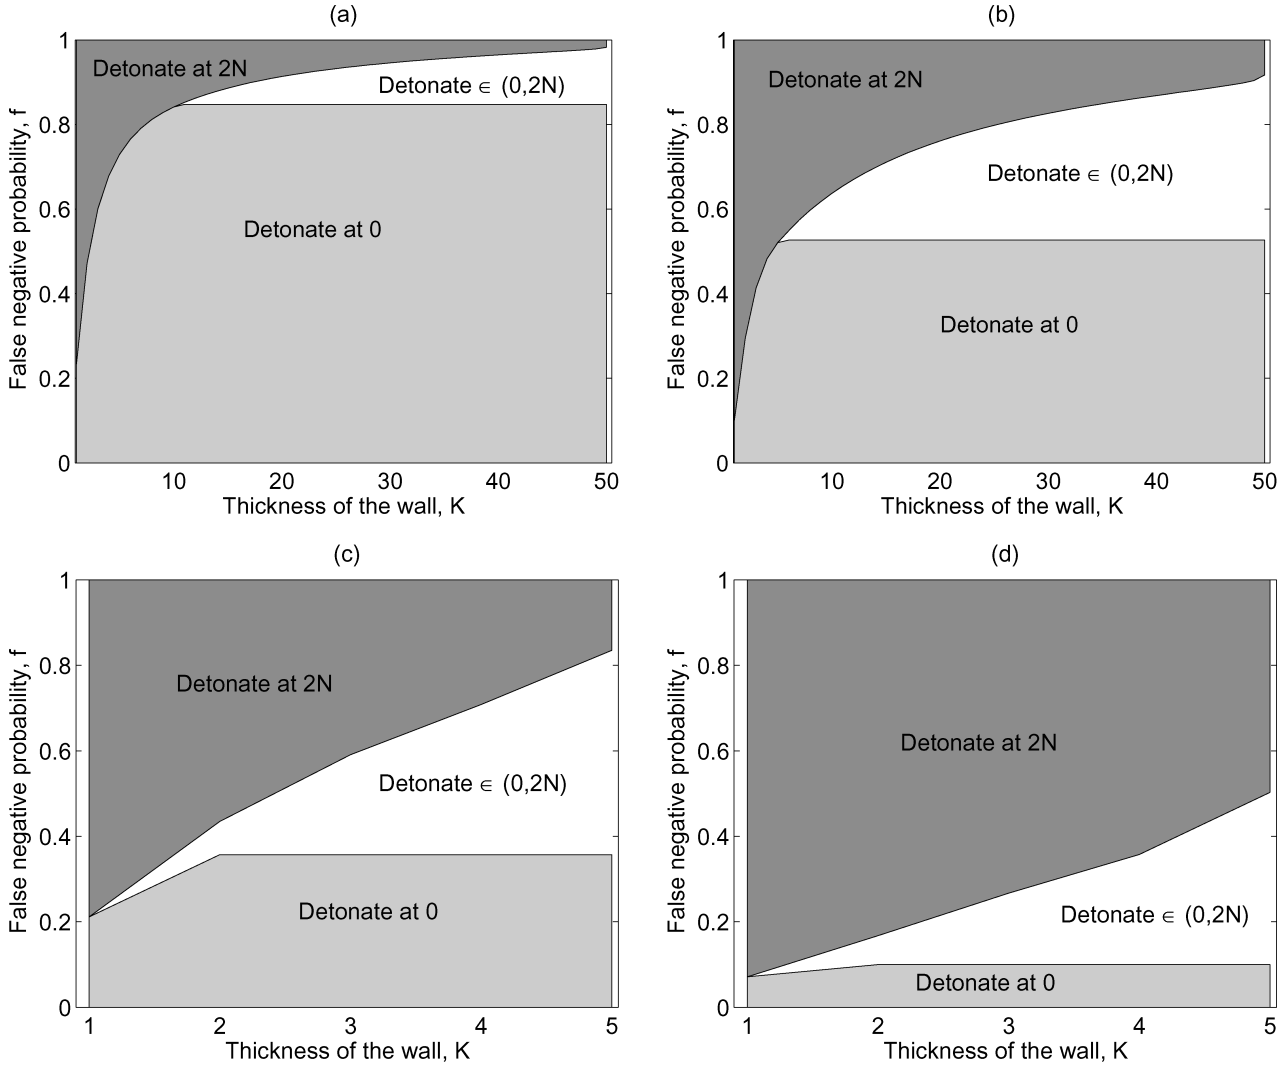

Fig. 1. The terrorist's optimal stopping point (k^*) in the OKK scenario, as a function of the false negative probability f and the wall thickness K, for the case of (a) detonation probability $q = 0.5$ and a large network ($N = 50$); (b) $q = 0.9$, $N = 50$; (c) $q = 0.5$ and a small network ($N = 5$); (d) $q = 0.9$, $N = 5$.

OKB Scenario. In scenario OKB, the Bayesian assumption implies that the optimal policy does not depend on f, and hence can be expressed as a threshold with respect to K (i.e., if the wall is sufficiently thin then proceed to the target) for the two values of q. The numerical solution to approximation (18) coincides with the exact values in Table III, and the thin-wall approximation (19) applied to the large network yields

$$k^* = 2N \quad \text{if} \quad K \leq \frac{4.5}{1-q} = \begin{cases} 9 & \text{if} \quad q = 0.5; \\ 45 & \text{if} \quad q = 0.9. \end{cases} \tag{24}$$

These two values are close to the true values in Table III, confirming the accuracy of the thin-wall approximation. Hence, when the detonation probability $q = 0.9$, the wall does not deter the terrorist.

4.3. Comparison of Scenarios

In this subsection, we compare the expected damage that is incurred across scenarios. We investigate five comparisons: RK vs. OKK and RB vs. OUB allow us to compare the two array topologies under two different informational structures, OKB vs. OUB permits us to assess the value of phantom sensors in the outer wall design, and RK vs. RB and OKK vs.

Scenario Comparison	Focus of Comparison	Linear Damage		Exponential Damage	
		Large Network	Small Network	Large Network	Small Network
RK vs. OKK	Network topology	15.66%	21.28%	3.79%	7.67%
RB vs. OUB	Network topology	11.80%	27.19%	4.75%	13.62%
OKB vs. OUB	Phantom sensors	28.95%	0.00%	28.90%	0.00%
RK vs. RB	Probing network	24.27%	5.85%	15.71%	7.86%
OKK vs. OKB	Probing network	6.66%	15.07%	6.00%	19.69%

Table IV. Percentage Increase in (Linear and Exponential) Damage of the First Scenario Relative to the Second Scenario, Averaged Over 30 Cases

OKB allow us to quantify the value to the terrorist of probing the network before the attack for the two array topologies. For each of these five comparisons and for large and small networks, we report in Table IV the percentage difference in expected damage averaged over 30 scenarios, which are the six possibilities of $f = 0.1, 0.5, 0.9$ and $q = 0.5, 0.9$, multiplied times either five values of K (1, 12, 24, 36, 48 for large networks and 1, 2, 3, 4, 5 for small networks) if two outer wall scenarios are being compared or five values of p_s (0.1, 0.3, 0.5, 0.7, 0.9) if two random array scenarios are being compared.

To meaningfully compare a random array scenario to an outer wall scenario, we need the average number of real sensors deployed in each scenario to be the same. For an outer wall with K layers, the fraction of nodes that have real sensors is $\frac{2K(N+1)-K^2}{(N+1)^2}$. Hence, we substitute this quantity in for p_s in Equation (1) to compute the passing probability p_u at a node in the corresponding random array scenario, which gives

$$p_u = \frac{(N-K+1)^2 + fK[2(N+1)-K]}{(N+1)^2}. \quad (25)$$

When comparing a random array scenario to an outer wall scenario, we choose the five values of K in the previous paragraph and then use Equation (25) to find the corresponding values of p_u for the random array scenario.

RK Scenario vs. OKK Scenario. Fig. 2 compares the expected damage of the RK and OKK scenarios as a function of the outer wall thickness for large networks. Corresponding plots for small networks are qualitatively similar and appear in Fig. 3. Fig. 2 shows that neither of these two array designs always dominates the other under this informational structure, in which the terrorist, by probing the network prior to the attack, knows the detection probability of the real sensors, the fraction of sensors in the random array that are real, and the precise location of the outer wall. From the government's point of view, the outer wall is preferable (i.e., there is less damage) when the real sensors are reasonably effective (detection probability equals 0.5 or 0.9). When the sensors have a detection probability of 0.1, the random array usually (but not when $q = 0.9$) leads to less damage if the sensor deployment is sparse (small K and p_u), and the outer wall is preferable if the sensor deployment is dense. The intuition is that if the real sensors have a high detection probability, then the terrorist will have a hard time penetrating an outer wall (even if it is relatively thin), whereas he may make some progress—and hence cause more damage—in a random array design. Conversely, an outer wall of sensors with a low detection probability does not provide much of a deterrent to the terrorist. When the detection probability is 0.9, the sensor array is a good deterrent regardless of the value of q, and the two outer wall scenarios give the same expected damage.

Computational experience shows, as expected, that better interdiction (smaller q) leads to less expected damage in Fig. 2. The expected damage can be maintained at a relatively low level (e.g., 3 on the 1-to-10 scale) even with low interdiction ($q = 0.9$) and poor sensors ($f = 0.9$), by compensating with a high density of real sensors (Fig. 2c).

RB Scenario vs. OUB Scenario. Fig. 4 is similar to Fig. 2, but assumes that the terrorist cannot observe the array topology or the detection probability of the sensors. For the large network in Fig. 4, the outer wall leads to less damage than the random array when $q = 0.9$ (poor interdiction). When $q = 0.5$ in the large network, the terrorist always detonates the bomb at the first node in both scenarios, resulting in the same damage. In these two information-poor scenarios, a larger q does not necessarily lead to more damage (particularly when the detection probability is high), because the bolder play induced by larger q may lead the terrorist further away from the optimal decision under perfect information (which is to detonate at the first node if the detection probability is high). In the small network (see the corresponding plots in Fig. 5),

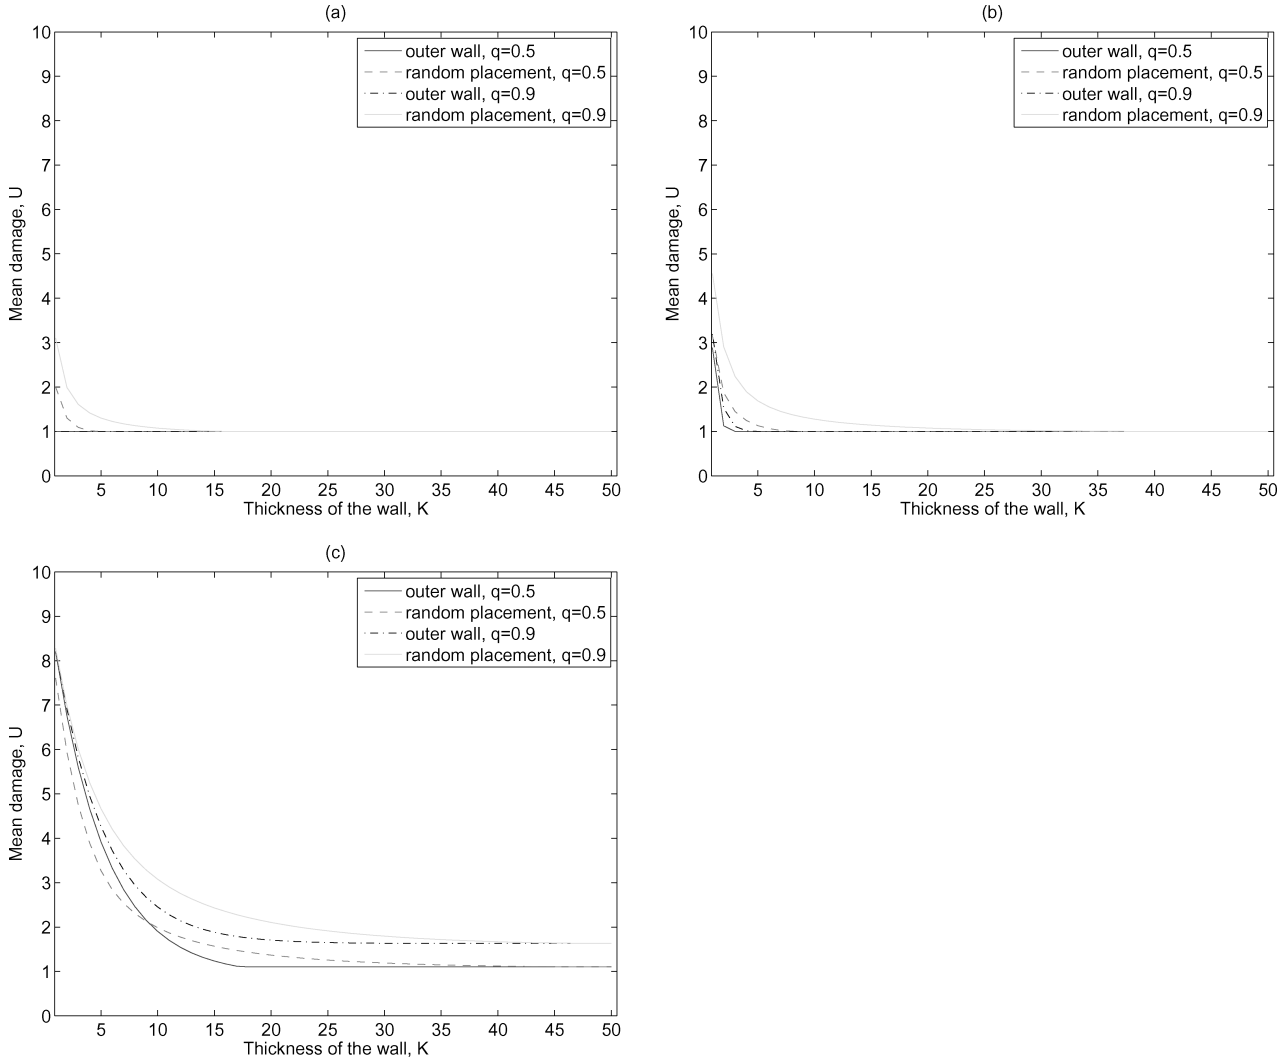

Fig. 2. RK scenario vs. OKK scenario for large networks when the false negative probability is (a) $f = 0.1$, (b) $f = 0.5$, (c) $f = 0.9$.

larger q leads to more damage because the terrorist always moves toward the target, and the outer wall is preferable in five of the six cases: in the ($f = 0.9$, $q = 0.5$) case, the random array leads to less damage for sufficiently sparse deployments.

Overall, the outer wall topology is better than the random topology for the great majority of scenarios in Figs. 2–5, and in the few cases in which the random topology is preferred, neither design is very effective. Consequently, we consider an outer wall topology in Wein and Atkinson (2007).

OKB Scenario vs. OUB Scenario. The average increase in damage as a result of the knowledge of the outer wall location is 28.95% for large networks (Table IV). However, for small networks, the terrorist proceeds to the target regardless of whether or not he can observe the outer wall, which generates the zeroes in Table IV. In some scenarios, knowledge of the wall, coupled with the uncertainty in the detection probability, causes the terrorist to play too boldly, thereby decreasing the damage. Phantom sensors appear to enhance the performance of the outer wall design for large networks.

RK Scenario vs. RB Scenario. Table III shows that the terrorist is more apt to proceed to the target in the RB scenario than in the RK scenario. The average damage in the RK scenario is 24.27% higher than in the RB scenario for large networks, and 5.85% larger for small networks (Table IV). These increases reflect

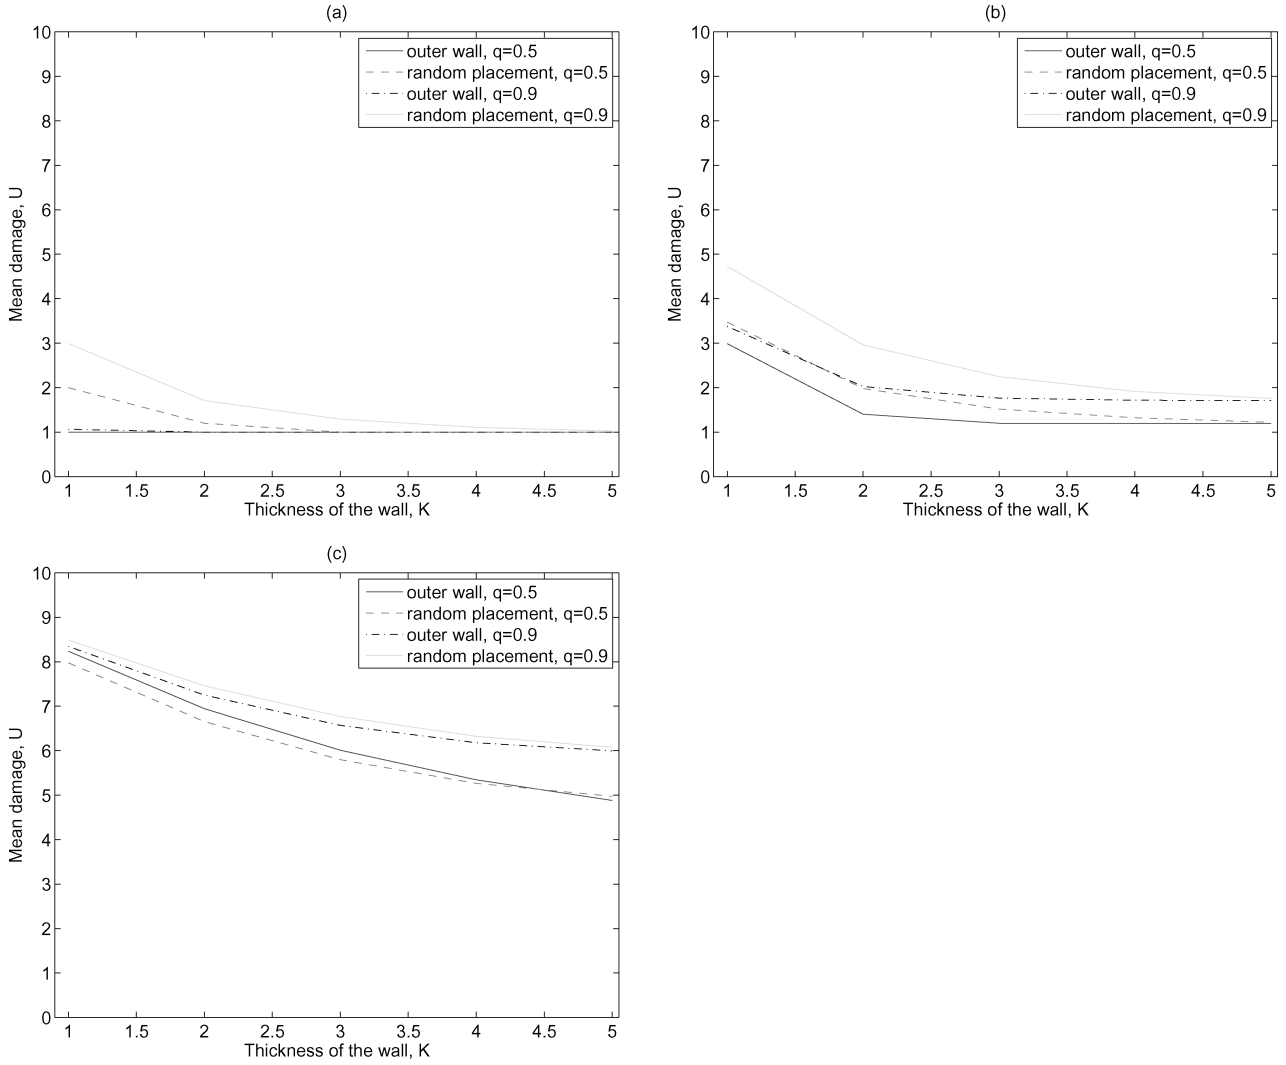

Fig. 3. RK scenario vs. OKK scenario for small networks when the false negative probability is (a) $f = 0.1$, (b) $f = 0.5$, (c) $f = 0.9$.

the fact that the terrorist's optimal Bayesian decision leads to less damage due to lack of information.

OKK Scenario vs. OKB Scenario. As in the previous paragraph, Table III reveals that the terrorist's decisions are erroneously bold under the Bayesian informational structure, where in this case he is uncertain about the detection probability of the real sensors. The average damage in the OKK scenario is 6.66% higher than in the OKB scenario for large networks and 15.07% larger for small networks (Table IV). These damage increases are smaller for large networks and larger for small networks than in the analogous comparison under the random array design (i.e., RK vs. RB). In both comparisons, it is worthwhile for terrorists to probe the network prior to the attack.

Finally, although some of the percentage increases in Table IV are modest, Figs. 2 and 4 reveal that the percentage increase can be very large for sparse deployments; i.e., the similarity of performance in the scenarios under dense deployments suppresses the numbers in Table IV.

4.4. Comparison of Linear and Exponential Damage

Because of the questionable appropriateness of a linear damage function, we performed a parallel analysis (and derived closed-form solutions to the optimal stopping problem) to that in Sections 3–4.3, but using an exponential damage function e^{ak+b}. In our numerical study, we set $b = 0$ and $a = \frac{\ln 10}{2N}$, so as to maintain

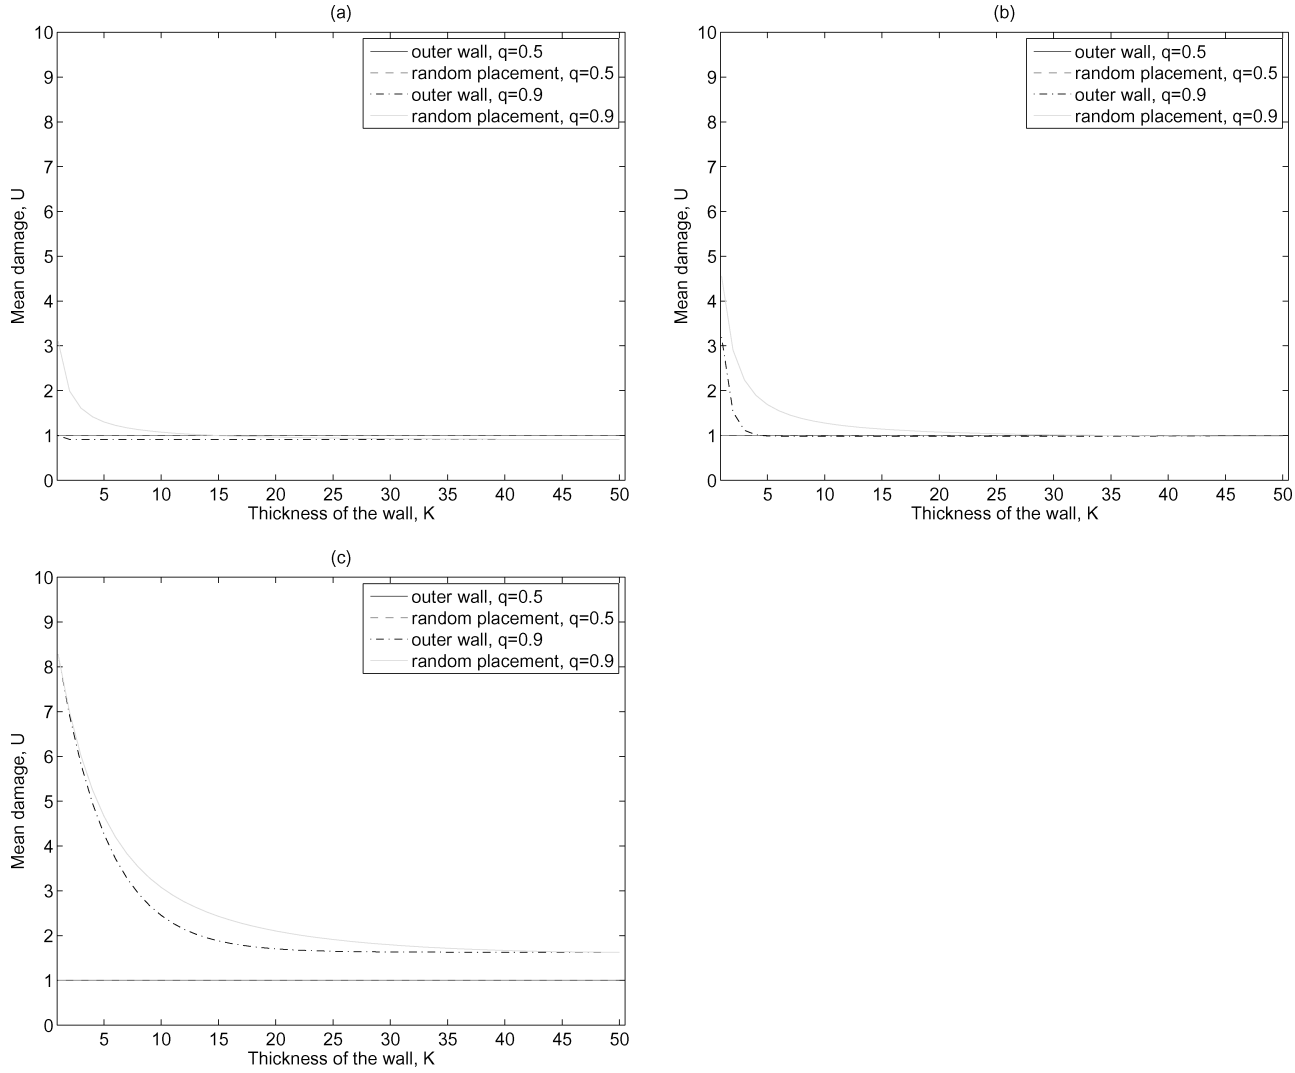

Fig. 4. RB scenario vs. OUB scenario for large networks when the false negative probability is (a) $f = 0.1$, (b) $f = 0.5$, (c) $f = 0.9$.

the 1-to-10 scale implied by Equation (20). Due to space considerations, we briefly summarize our results here but do not show any of the mathematical or computational details of the exponential damage case.

Overall, the results in the exponential damage case are qualitatively similar to the results in the linear damage case (Table IV), which suggests that our results are somewhat robust. Here, we mention the minor differences in the two cases. In the RK scenario, the relatively slow increase in damage as the terrorist proceeds through the network in the exponential case causes him to forgo an intermediate stopping point. He is more likely to proceed all the way to the target in the exponential case, but is also more likely to detonate the bomb at (0, 0). In the RB and OUB scenarios, the thresholds in the exponential case are slightly larger than the thresholds in the linear case, meaning that the terrorist is bolder (i.e., more likely to proceed to the target) in the linear case. In contrast to the RB and OUB scenarios, in the OKB scenario the terrorist moves more boldly in the exponential case than the linear case because of the lower damage gradient in the exponential case as the terrorist moves through the outer wall. Unlike the solution to the linear case in Fig. 1, the solution in the OKK scenario in the exponential case is either $k^* = 0$ or $2N$.

The RK vs. OKK comparison is similar for the linear and exponential cases, although the average

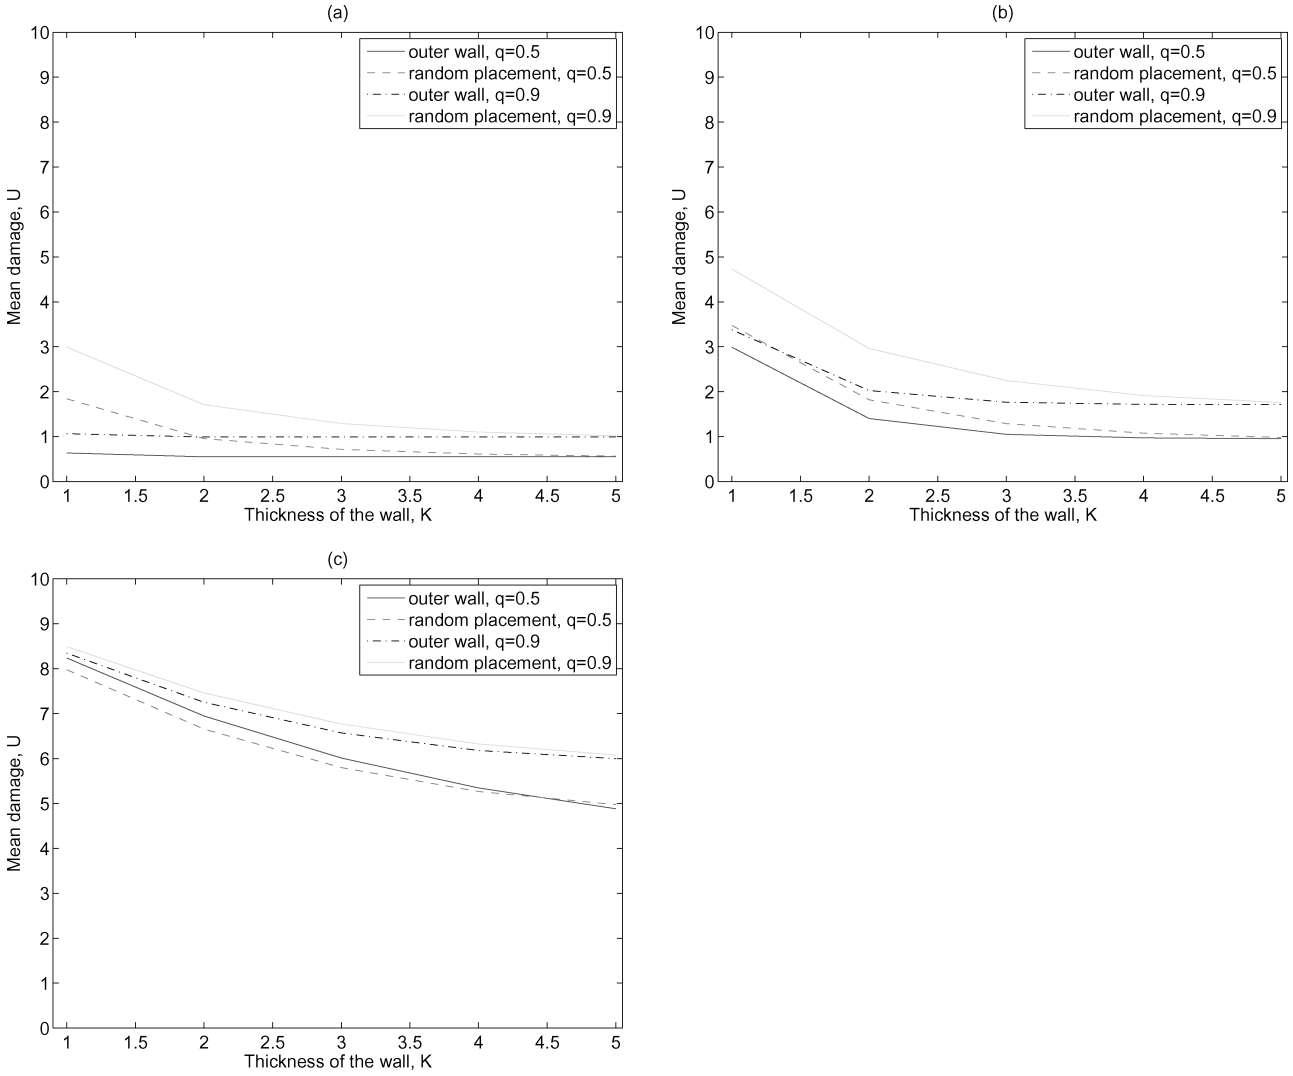

Fig. 5. RB scenario vs. OUB scenario for small networks when the false negative probability is (a) $f = 0.1$, (b) $f = 0.5$, (c) $f = 0.9$.

percentage difference in damage is 5.73% in the exponential case compared with 18.47% for the linear case. In the exponential case, the random array is preferable to the outer wall in a large network with poor interdiction ($q = 0.9$) and poor detection ($f = 0.9$), and in a small network with poor detection. The RB vs. OUB comparison is similar for both cases, but the difference in damage between the two designs is smaller (9.19% vs. 19.50%) in the exponential case than the linear case. The policies in the OKB and OUB scenarios coincide in the small network case under exponential damage (as they do in linear case), and the terrorist's knowledge of the outer wall provides a similar large increase in damage in the large network for the exponential and the linear case. The RK vs. RB and OKK vs. OKB comparisons are similar under both damage functions.

4.5. General Damage Functions and Linear Sensor Locations

Thus far, to maintain analytical tractability, we have idealized the problem to linear or exponential damage functions and equally spaced sensors. In this subsection, we briefly outline how to formulate and solve more general problems in which the damage $g(x)$ is a function of distance to the city center and sensors can be placed at generic locations $(x_0, \ldots, x_n)$, where $0 \le x_0 < x_1 < \cdots < x_n$ and the target is located at x_n. For the RK scenario, the optimal stopping problem generalizes to

$$V(x_k) = \max\{g(x_n - x_k), (1 - p_u)qg(x_n - x_k) + p_u V(x_{k+1})\} \quad \text{for} \quad k = 0, \ldots, n-1,$$
$$V(x_n) = g(0). \tag{26}$$

Generalizing the other four scenarios requires changing $V(k)$ to $V(x_k)$, $ak + b$ to $g(x_n - x_k)$, and the boundary condition to $g(0)$. While it may no longer be possible to compute analytical conditions for the decision variable k^* as we did in Section 3, solving this optimal stopping problem numerically is a straightforward application of dynamic programming techniques (e.g., Section 3.4 in Bertsekas, 1976). Using backward induction we can solve for $V(x_{n-1})$, then for $V(x_{n-2})$, and continuing until we compute $V(x_0)$. The optimal stopping node is then given by

$$k^* = \min\{k \,:\, V(x_k) = g(x_n - x_k)\}, \tag{27}$$

and the expected damage for the RK scenario is

$$U = \sum_{i=0}^{k^*-1} p_u^i (1 - p_u) qg(x_n - x_i) + \left(1 - \sum_{i=0}^{k^*-1} p_u^i (1 - p_u)\right) g(x_n - x_{k^*}), \tag{28}$$

with similar modifications made to the expected damage expressions for the other scenarios. Implementing this algorithm only involves a few lines of code, and is an $O(n)$ algorithm because we only need to compute the value function $V(x_k)$ at n states and each calculation involves a constant number of operations and comparisons. Computing k^* requires no extra effort and the damage in Equation (28) also only requires $O(n)$ computations. Thus, the total number of calculations to determine the optimal policy and expected damage scales linearly with the number of sensor locations.

Other extensions, such as the informed priors described in Section 3.2 or node-dependent detection and interdiction probabilities, can be handled using this backward induction technique.

5. CONCLUDING REMARKS

Our analysis generates three main results, which are conditional on rapid interdiction because the model implicitly assumes instantaneous interdiction. The first result is that the sensor array can deter a terrorist if it is densely deployed (i.e., a large network with many sensors) and interdiction is effective (i.e, the detonation probability is $q = 0.5$). In our view, this result should not be taken too literally because it is very difficult to understand a terrorist's mindset (we assume he is risk-neutral, although he could be risk-seeking or risk-averse), and it seems likely, particularly if he has a device to detonate the bomb from the driver's seat, that he views the probability q as very large. Hence, we believe it is prudent to assume that a terrorist will proceed directly to the target, although the possibility of deploying a dense set of fake (i.e., inoperable) sensors in addition to the real sensors should be investigated.

Our other two main results, in contrast, are quite robust: (i) even if interdiction and the sensors are ineffective (i.e., $q = 0.9$ and the false negative probability $f = 0.9$), the mean damage can be maintained at a relatively moderate level (i.e., 3 on a 1-to-10 scale) by a dense deployment, and (ii) the outer wall leads to less mean damage than the random deployment except in the case of poor sensors ($f = 0.9$) and sparse deployment, in which case neither deployment is effective.

The random topology and the outer wall topology studied here are not the only possibilities. We also tested (results not shown) a nonrandom strategy that deploys sensors at a fixed proportion of nodes in a uniform manner (e.g., locate the sensors in a checkerboard pattern if half of the nodes have real sensors) and found that it led to slightly less damage (although essentially the same damage for large networks) than the random topology if the terrorist knows the sensors' detection probability but does not have any information about the sensor layout. This nonrandom strategy generated more damage than the outer wall topology, which also lacks randomness. Although random strategies are more difficult for a terrorist to probe, overall our analysis suggests that it is preferable to use an outer wall to keep the terrorist as far from the target as possible. This led us to consider an interdiction model with an outer wall design in Wein and Atkinson (2007).

Perhaps the biggest shortcoming of the model is that the spatial grid is merely a caricature of an actual highway system. While a spatial lattice seems appropriate for crudely comparing an outer wall and a random deployment, as we have done here, more refined insights would require the modeling of a specific city's highway structure (see Section 4.5 for a start in this direction), or possibly the modeling of highways by a percolation process (Grimmett, 1999), which would be much more difficult to analyze. Other worthwhile generalizations to our model, alluded to earlier, are noninstantaneous interdiction (which is pursued in Wein & Atkinson, 2007 via a spatial queueing model), terrorists that are risk-averse or risk-seeking, and a

detailed simulation analysis of the four effects of a nuclear weapon, which could perhaps be summarized by the generic damage function $g(x)$ in Section 4.5.

In a companion article (Wein & Atkinson, 2007), we embed three models into a Stackelberg game: a sensor model first developed in Wein *et al.* (2006), which determines the detection probability and the false positive probability as a function of the neutron threshold level of the sensor, the OKB scenario of the optimal stopping problem analyzed here, and a spatial interdiction model that incorporates scarce interdiction resources (i.e., it is a spatial queueing model). In this game, the U.S. government (as the leader) chooses the neutron threshold level, the thickness of the wall sensors (i.e., how many sensors the terrorist needs to pass through), and the number of interdiction vehicles to minimize the expected damage inflicted by a terrorist subject to a budget constraint on the annual cost of sensors and interdiction vehicles, and the terrorist (as the follower) observes the wall thickness and solves the optimal stopping problem with the goal of maximizing the expected damage.

ACKNOWLEDGMENTS

LMW thanks Tom Edmunds and Richard Wheeler for helpful conversations. We thank the reviewers for their valuable comments. This research was primarily supported by Lawrence Livermore National Laboratory, Project B529238. The first and third authors were also supported, respectively, by an Abbott Laboratories Stanford Graduate Fellowship and the Center for Social Innovation, Graduate School of Business, Stanford University.

APPENDIX

Sections A and C contain the proofs of Propositions 1 and 2, respectively. The analysis of scenarios RB and OKK appear in Sections B and D, respectively.

APPENDIX A: PROOF OF PROPOSITION 1

We first show that if it is optimal to proceed in state $2N-1$ then it is optimal to proceed in states $k = 0, \ldots, 2N-2$, in which case $k^* = 2N$, i.e., it is optimal for the terrorist to proceed directly to his goal. By Equation (2), if it is optimal to proceed in state $2N-1$ then

$$a(2N-1)+b \le (1-p_u)q[a(2N-1)+b] + p_u(2aN+b), \quad \text{(A.1)}$$

and at state $2N-2$,

$$\begin{aligned} V(2N-2) = \max\{&a(2N-2)+b, \\ &(1-p_u)q[a(2N-2)+b] \\ &+ p_u[(1-p_u)q[a(2N-1)+b] \\ &+ p_u(2aN+b)]\}. \end{aligned} \quad \text{(A.2)}$$

By writing the last term in (A.1) as $p_u[a(2N-1)+a+b]$, we can reexpress this inequality as

$$2N-1+\frac{b}{a} \le \frac{p_u}{(1-p_u)(1-q)}. \quad \text{(A.3)}$$

Turning to the decision at state $2N-2$, we have

$$\begin{aligned} 2N-2+\frac{b}{a} &< 2N-1+\frac{b}{a}, \\ &\le \frac{p_u}{(1-p_u)(1-q)} \quad \text{by (A.3)}, \end{aligned} \quad \text{(A.4)}$$

and the approach used to derive Equation (A.3) from Equation (A.1) can be applied to reexpress Equation (A.4) as

$$a(2N-2)+b \le (1-p_u)q[a(2N-2)+b] + p_u[a(2N-1)+b]. \quad \text{(A.5)}$$

But by Equation (A.1) we know that the right side of Equation (A.5) satisfies

$$\begin{aligned} &(1-p_u)q[a(2N-2)+b] + p_u[a(2N-1)+b] \\ &\quad \le (1-p_u)q[a(2N-2)+b] + p_u[(1-p_u) \\ &\qquad \times q[a(2N-1)+b] + p_u(2aN+b)], \end{aligned} \quad \text{(A.6)}$$

and Equations (A.5) and (A.6) imply that the second term inside the maximum of Equation (A.2) dominates the first term, i.e., it is optimal to proceed in state $2N-2$. This same backward induction argument applies to states $2N-3, \ldots, 0$. Hence, if Equation (A.1) holds then it is optimal to proceed to state $2N$.

Now assume that Equation (A.1) is violated, so that it is preferable to detonate rather than proceed in state $2N-1$. Furthermore, suppose that the detonation term (i.e., the first term in the maximum in Equation (2) in the main text) also dominates in states $k^*, \ldots, 2N-2$, but that the proceed term dominates in state k^*-1. Then by an argument identical to that in the previous paragraph, it follows that the proceed term also dominates in states $0, \ldots, k^*-2$, and hence it is optimal to proceed to state k^* and detonate the bomb there. Under this set of assumptions, we have

$$ak^* + b > (1 - p_u)q(ak^* + b) + p_u[a(k^* + 1) + b], \tag{A.7}$$

$$a(k^* - 1) + b < (1 - p_u)q[a(k^* - 1) + b) + p_u[ak^* + b]. \tag{A.8}$$

Equation (4) in the main text follows from inequalities (A.7) and (A.8) and the restriction that k^* needs to be between 0 and $2N$.

Before the terrorist reaches state k^*, he may be detected as he passes through states $0, \ldots, k^* - 1$, in which case he successfully detonates the bomb with probability q. Hence, the expected damage under the optimal policy is

$$U = \sum_{i=0}^{k^*-1} p_u^i (1 - p_u) q(ai + b) + \left(1 - \sum_{i=0}^{k^*-1} p_u^i (1 - p_u)\right)(ak^* + b), \tag{A.9}$$

which simplifies to Equation (5) in the main text.

APPENDIX B: ANALYSIS OF THE RB SCENARIO

Because the perceived p_u is no longer constant in the Bayesian case, the proof of Proposition 2 in Section C does not yield a necessary and sufficient condition for $k^* = 0$ or $2N$. Rather, it first shows that if it is optimal to detonate in state $2N - 1$ then it is optimal to detonate in all earlier states, which implies that $k^* = 0$. But if it is optimal to proceed in state $2N - 1$, it is no longer true that it is optimal to proceed in states $0, \ldots, 2N - 2$. However, in this case, we can still show that if there exists a state $k < 2N - 1$ where it is preferable to detonate then it is preferable to detonate in states $0, \ldots, k - 1$, implying $k^* = 0$, and if there does not exist such a state then $k^* = 2N$.

Therefore, the proof of Proposition 2 implies that the necessary and sufficient condition for $k^* = 2N$ is that it is optimal to proceed in state 0, given that it is optimal to proceed in states $1, \ldots, 2N - 1$. That is, if we let $\bar{V}(k)$ denote the value function in state k if it is optimal to proceed in states $k + 1, \ldots, 2N - 1$, then $k^* = 2N$ if $\bar{V}(0) \geq b$ and $k^* = 0$ otherwise. We have by induction that

$$\bar{V}(k) = \frac{q}{k+2}(ak + b) + \sum_{i=k+1}^{2N-1} \frac{q(k+1)(ai + b))}{(i+1)(i+2)} + \frac{(k+1)(2aN + b)}{2N+1},$$

and hence $k^* = 2N$ if

$$b \leq \sum_{i=0}^{2N-1} \frac{q(ai + b)}{(i+1)(i+2)} + \frac{2aN + b}{2N+1}, \tag{B.1}$$

$$= \sum_{i=0}^{2N-1} \frac{q(b - 2a)}{(i+1)(i+2)} + aq \sum_{i=0}^{2N-1} \frac{(i+2)}{(i+1)(i+2)} + \frac{2aN + b}{2N+1},$$

$$= \frac{2Nq(b - 2a)}{2N+1} + aq \sum_{i=1}^{2N} \frac{1}{i} + \frac{2aN + b}{2N+1},$$

$$\approx \frac{2Nq(b - 2a)}{2N+1} + aq \ln 4N + \frac{2aN + b}{2N+1}. \tag{B.2}$$

The approximation in Equation (B.2) comes from the approximation $\sum_{i=1}^{2N} \frac{1}{i} \approx \ln 2N + \gamma \approx \ln 2N + \ln 2 \approx \ln 4N$, where $\gamma \approx 0.5772$ is the Euler-Mascheroni constant.

APPENDIX C: PROOF OF PROPOSITION 2

We first prove that if it is preferable to detonate the bomb in state $2N - 1$ then it is preferable to detonate in states $0, \ldots, 2N - 2$, implying that $k^* = 0$. If it is preferable to detonate in state $2N - 1$ then

$$a(2N - 1) + b \geq \frac{q}{2N+1}[a(2N - 1) + b] + \frac{2N}{2N+1}(2aN + b), \tag{C.1}$$

$$V(2N-2) = \max\left\{a(2N-2) + b, \frac{q}{2N}[a(2N-2) + b] + \frac{2N-1}{2N}[a(2N-1) + b]\right\}. \tag{C.2}$$

Manipulating Equation (C.1) yields

$$a(2N - 1) + b \geq \frac{q}{2N+1}[a(2N - 1) + b] + \frac{2N}{2N+1}(2aN + b),$$

$$\iff \frac{2N+1-q}{2N+1}[a(2N-1) + b] \geq \frac{2N}{2N+1}(2aN + b),$$

$$\iff \frac{2N - 1 + \frac{b}{a}}{2N + \frac{b}{a}} \geq \frac{2N}{2N+1-q}. \tag{C.3}$$

Now we combine the inequality

$$\left(2N - 1 + \frac{b}{a}\right) + 2N + 1 - q < 2N + \frac{b}{a} + 2N \tag{C.4}$$

with Equation (C.3) to get

$$\left(2N-1+\frac{b}{a}\right)(2N+1-q)-\left(2N-1+\frac{b}{a}\right)+2N+1-q>\left(2N+\frac{b}{a}\right)2N-\left(2N+\frac{b}{a}+2N\right),$$

$$\iff \left(2N-1+\frac{b}{a}-1\right)(2N+1-q-1)>\left(2N+\frac{b}{a}-1\right)(2N-1),$$

$$\iff \frac{2N-2+\frac{b}{a}}{2N-1+\frac{b}{a}}>\frac{2N-1}{2N-q}, \tag{C.5}$$

$$\iff a(2N-2)+b>\frac{q}{2N}[a(2N-2)+b]+\frac{2N-1}{2N}[a(2N-1)+b], \tag{C.6}$$

which implies that it is preferable to detonate in state $2N-2$ in (C.2). Applying the same analysis to states $2N-3,\ldots,0$ shows that it is preferable to detonate in every state, so that $k^*=0$.

If condition (C.1), and hence Equation (C.3), is violated, then inequality (C.4) is in the wrong direction to derive an inequality in the opposite direction of Equation (C.6). That is, it is not true that if it is preferable to proceed in state $2N-1$ then it is preferable to proceed in states $0,\ldots,2N-2$. However, suppose that for some k it is optimal to proceed in states $k+1,\ldots,2N-1$, and optimal to detonate in state k, i.e.,

$$ak+b\geq\frac{q}{k+2}(ak+b)+\frac{k+1}{k+2}V(k+1)>\frac{q}{k+2}(ak+b)+\frac{k+1}{k+2}[a(k+1)+b],$$

$$V(k-1)=\max\left\{a(k-1)+b,\frac{q}{k+1}[a(k-1)+b]+\frac{k}{k+1}(ak+b)\right\}.$$

Using an analysis similar to Equations (C.4) and (C.6), we can show that it is preferable to detonate in states $0,\ldots,k-1$, and hence the optimal stopping point is still $k^*=0$. But if it is preferable to proceed in every state from 0 to $2N-1$, then the optimal stopping point is $k^*=2N$.

APPENDIX D: ANALYSIS OF THE OKK SCENARIO

We begin by deriving Equation (13) in the main text. At state $2K-1$ of the OKK optimal stopping problem, we have

$$V(2K-1)=\max\{a(2K-1)+b,\ q(1-f)(a(2K-1)+b)+f(2aN+b)\}.$$

If the first term is larger, i.e.,

$$\frac{1-q(1-f)}{f}\geq\frac{2N+\frac{b}{a}}{2K-1+\frac{b}{a}}, \tag{D.1}$$

then at state $2K-2$,

$$V(2K-2)=\max\{a(2K-2)+b,\ q(1-f)\times(a(2K-2)+b)+f(a(2K-1)+b)\}.$$

In this case, the optimal stopping problem is identical in structure to the one in scenario RK (see Equation (A.2)), which implies that the optimal stopping point is

$$k^*=\min\left\{2K-1,\max\left\{0,\left\lceil\frac{f}{(1-f)(1-q)}-\frac{b}{a}\right\rceil\right\}\right\}.$$

Moreover, condition (D.1) implies that

$$2K-1>\max\left\{0,\left\lceil\frac{f}{(1-f)(1-q)}-\frac{b}{a}\right\rceil\right\},$$

and hence

$$k^*=\max\left\{0,\left\lceil\frac{f}{(1-f)(1-q)}-\frac{b}{a}\right\rceil\right\}\quad\text{if Equation (D.1) is satisfied.} \tag{D.2}$$

We now derive Equation (15) in the main text. If Equation (D.1) is violated, then it is preferable to proceed rather than detonate in state $2K-1$, and

$$V(2K-2)=\max\{a(2K-2)+b,\ q(1-f)(a(2K-2)+b)+f[q(1-f)\times(a(2K-1)+b)+f(2aN+b)]\}. \tag{D.3}$$

Suppose there is a state M such that the optimal decision is to detonate in state M but the optimal decision is to proceed in states $M+1$ and beyond. It follows that

$$V(M+1) = \max\Bigg\{a(M+1)+b,\; q(1-f)\sum_{i=M+1}^{2K-1}(ai+b)\,f^{i-M-1} + f^{2K-M-1}(2aN+b)]\Bigg\} \quad \text{(D.4)}$$

and the second term inside the maximum is larger than the first one, whereas in state M,

$$V(M) = \max\Bigg\{aM+b, q(1-f)\sum_{i=M}^{2K-1}(ai+b)\,f^{i-M} + f^{2K-M}(2aN+b)]\Bigg\} \quad \text{(D.5)}$$

and the first term inside the maximum dominates the second one. Under conditions (D.4)–(D.5), the problem reduces again to the one in scenario RK, for at state $M-1$,

$$V(M-1) = \max\{a(M-1)+b,\; q(1-f)(a(M-1)+b) + f(aM+b)\}.$$

Therefore, Proposition 1 implies that the optimal stopping point is

$$k^* = \min\left\{M, \max\left\{0, \left\lceil \frac{f}{(1-f)(1-q)} - \frac{b}{a} \right\rceil\right\}\right\} \quad \text{if Equation (D.1) is violated.} \quad \text{(D.6)}$$

REFERENCES

Bailey, M. D., Shechter, S. M., & Schaefer, A. J. (2006). SPAR: Stochastic programming with adversarial recourse. *Operations Research Letters*, *34*, 307–315.

Berger, J. (1985). *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag.

Bertsekas, D. P. (1976). *Dynamic Programming and Stochastic Control*. New York: Academic Press.

Bier, V. M. (2007). Choosing what to protect. *Risk Analysis*, *27*, 607–620.

Bunn, M., Wier, A., & Holdren, J. P. (2003). *Controlling Nuclear Warheads and Materials: A Report Card and Action Plan*, Project on Managing the Atom, John F. Kennedy School of Government, Harvard University, Cambridge, MA.

Flynn, S. E. (2000). Beyond border control. *Foreign Affairs*, *79*, 57–65.

Gibbons, R. (1992). *Game Theory for Applied Economists*. Princeton, NJ: Princeton University Press.

Glasstone, S., & Dolan, P. J. (1977). *The Effects of Nuclear Weapons*. U.S. Dept. of Defense, Energy Research and Development Administration.

Grimmett, G. (1999). *Percolation*. Berlin: Springer-Verlag.

Hage, W., & Cirafelli, D. M. (1985). Correlation analysis with neutron count distributions in randomly or signal triggered time intervals for assay of special fissile materials. *Nuclear Science and Engineering*, *89*, 159–176.

Hemmecke, R., Schultz, R., & Woodruff, D. L. (2003). Interdicting stochastic networks with binary interdiction effort. In D. L. Woodruff (Ed.), *Network Interdiction and Stochastic Programming*, Chapter 4. Boston, MA: Kluwer.

Kaplan, E. H., & Kress, M. (2005). Operational effectiveness of suicide-bomber-detector schemes: A best-case analysis. *PNAS*, *102*, 10399–10404.

Lipton, E. (2007). New York to test ways to prevent nuclear terror. *NY Times*, p. A1 (February 9).

Morton, D. P., Pan, F., & Saeger, K. J. (2007). Models for nuclear smuggling interdiction. *IIE Transactions*, *39*, 3–14.

Pan, F., Charlton, W. S., & Morton, D. P. (2003). A stochastic program for interdicting smuggled nuclear material. In D. L. Woodruff (Ed.), *Network Interdiction and Stochastic Programming*, Chapter 1. Boston, MA: Kluwer.

Stana, R. M. (2004). *Summary of Challenges Faced in Targeting Oceangoing Cargo Containers for Inspection*. U.S. General Accounting Office Report GAO-04-557T, March 31.

U.S. Department of Homeland Security. (2005). *Fact Sheet: Domestic Nuclear Detection Office*. http://www.dhs.gov/dhspublic/display?theme=43&content=4474&print=true, accessed on September 23.

Washburn, A., & Wood, K. (1995). Two-person zero-sum games for network interdiction. *Operations Research*, *43*, 243–251.

Wein, L. M., & Atkinson, M. P. (2007). The last line of defense: Designing and managing radiation sensor arrays around cities. *IEEE Transactions on Nuclear Science*, *54*, 654–669.

Wein, L. M., Wilkins, A. H., Baveja, M., & Flynn, S. E. (2006). Preventing the importation of illicit nuclear materials in shipping containers. *Risk Analysis*, *26*, 1377–1393.