Risk Analysis, Vol. 32, No. 5, 2012 DOI: 10.1111/j.1539-6924.2011.01702.x

Robust Allocation of a Defensive Budget Considering an Attacker's Private Information

Mohammad E. Nikoofal[1] and Jun Zhuang[2,*]

Attackers' private information is one of the main issues in defensive resource allocation games in homeland security. The outcome of a defense resource allocation decision critically depends on the accuracy of estimations about the attacker's attributes. However, terrorists' goals may be unknown to the defender, necessitating robust decisions by the defender. This article develops a robust-optimization game-theoretical model for identifying optimal defense resource allocation strategies for a rational defender facing a strategic attacker while the attacker's valuation of targets, being the most critical attribute of the attacker, is unknown but belongs to bounded distribution-free intervals. To our best knowledge, no previous research has applied robust optimization in homeland security resource allocation when uncertainty is defined in bounded distribution-free intervals. The key features of our model include (1) modeling uncertainty in attackers' attributes, where uncertainty is characterized by bounded intervals; (2) finding the robust-optimization equilibrium for the defender using concepts dealing with budget of uncertainty and price of robustness; and (3) applying the proposed model to real data.

KEY WORDS: Defender-attacker game; defense resource allocation; private information; robust optimization

1. INTRODUCTION

Terrorist threats of attack are a serious concern to be addressed for the sake of the national economy and the quality of life. The recent history of terrorist attacks on both military and civilian targets is another reason for a deeper focus on this issue. In assessing terrorist attacks, the goal of a defender is to understand the terrorist group before an attack occurs, or to implement adequate security measures to deter an attack or, at least, to decrease the expected damage of an attack. In the risk analysis community, a number of researchers have used system analysis,[1] mathematical models,[2] decision trees,[3] probabilistic risk analysis,[4] intelligent adversary methods,[5–7] and game theory[8–11] to model the strategic interactions in homeland security and counterterrorism risk management.

Deterring terrorism is generally expensive and deciding how, where, and when to allocate resources in order to protect critical infrastructure is a difficult problem, specifically when we have no accurate estimation about the attacker's attributes. The principal issues are (1) precise estimation of the probability of an attack; (2) exact assessment of the attacker's effort; and (3) identification of actual targets and the attacker's valuation of those targets. However, the most important thing to be considered by the

[1]Desautels Faculty of Management, McGill University, Montreal, Quebec, Canada.
[2]Department of Industrial and Systems Engineering, University at Buffalo, NY, USA.
*Address correspondence to Jun Zhuang, Department of Industrial and Systems Engineering, University at Buffalo, The State University of New York, Buffalo, NY 14260-2050, USA; jzhuang@buffalo.edu.

 0272-4332/12/0100-0930$22.00/1 © 2011 Society for Risk Analysis

defender among mentioned issues is the attacker's valuation of its target, since it reveals the attacker's preferences regarding the targets, which are known by the attacker but not to the defender.

A number of studies investigate the scenario where the defender knows nothing, but assumes probabilistic distributions about the attacker's preferences, while the attacker observes the defender's resource allocations. Bier *et al.*[12] study a model in which a defender allocates defensive resources to a collection of targets, and an attacker chooses a target to attack from the collection. They assume a probability distribution for the attacker's valuation of targets, and find the equilibria in which the attacker receives the valuation of the target in the case of a successful attack and zero otherwise, and the defender experiences a loss of her valuation of the target from a successful attack by the attacker and no loss otherwise. Bier *et al.*[13] provide a methodology for identifying attacker and defender equilibrium strategies in a sequential game where the defender plays first. In their model, they assume that the probability of an attack on each target is a function of the budget allocations to all targets, and also presume that the success probability of an attack on a specific target is a function of the defensive resources allocated to that target. Using the data from Willis *et al.*,[14] Bier *et al.*[13] claim that the cost effectiveness of defensive investment plays an important role in the optimal defense allocation plan. They take into account the defender's uncertainty about the attacker's target valuations using a probability distribution that affects the defender's assessment of the probability of an attack, and the success of an attack on the targets. In a more recent study by Hao *et al.*,[15] similar to Bier *et al.*,[13] a probabilistic optimization model for defense budget allocations in which the attacker might be strategic or nonstrategic with a known probability distribution is investigated. Considering probability distributions to model uncertainty in attackers' attributes is far from reality, since the optimal defender's strategy is seriously affected by such prespecified probability distributions.

There are also several studies that investigate allocating defensive resources between terrorism and natural disasters, since nature plays by chance, but is still capable of causing significant economic losses and casualties (e.g., Hurricane Andrew[16] and the Northridge earthquake[17]). The model presented by Powell[18] considers terrorism and natural disaster simultaneously. A game-theoretic approach to resource allocation for countering terrorism and natural disasters simultaneously is proposed by Zhuang and Bier[19] to identify equilibrium strategies for both the attacker and the defender; however, they do not consider any uncertainty in the attacker's behavior, and also, they do not provide real data to illustrate the use of their model. In more recent works, Golany *et al.*[20] and Levitin and Hausken[21] study the problem of allocating limited defensive resources among vulnerable sites when the damage might be caused by either probabilistic risk from nature (unintentional threat), or strategic risk from terrorists (intentional threat).

There exist some recent game-theoretical studies on defender-attacker models using complete and incomplete information. The bi- and tri-level optimization model for sequential attacker-defender and defender-attacker games, respectively, with information transparency is proposed by Brown *et al.*[22] to protect critical infrastructure against terrorist attacks. They use mixed-integer programming to define defender and attacker decisions and budgetary constraints for each of them. A model of secrecy and deception in multiple-period attacker-defender resource allocation games is proposed by Zhuang *et al.*,[23] which analyzes a finite game between a single attacker and a single defender as the defender has private information, such as target value and cost effectiveness, while the attacker does not, but the attacker can update his knowledge about the defender after observing the defender's signals, and also after observing the result of a contest. There also exists similar research on counterinsurgency operations that studies the optimal government's force allocation against insurgents. In the work by Kaplan *et al.*,[24] the authors extend a game-theoretical model to obtain the optimal government force allocation against insurgent strongholds when the government has imperfect intelligence.

While the model presented in this article assumes a sequential static game, a similar case with simultaneous moves known as the Colonel Blotto game includes situations in which there are complementarities among the targets being defended.[25] In the Colonel Blotto game, two players simultaneously distribute forces across n battlefields. Within each battlefield, the player that allocates the highest level of force wins. Roberson[26] extends the literature on the Colonel Blotto game by characterizing the equilibrium for all players.

In this article, we propose a defender-attacker game with incomplete information in which the attacker has private information about the valuation of

targets. The defender tries to maximize her payoffs using the attacker's best response, and tries to cope with uncertainty in the attacker's attributes by modeling uncertainty in distribution-free, but bounded, intervals and choosing robust-optimization equilibrium. More precisely, our proposed model is distinguishable among similar models in achieving robust solutions that consider the tradeoffs between the robustness of defense allocations and the conservatism level of the defender, taking into account the cost of having such robustness in solutions. To our best knowledge, no previous research has applied robust optimization in homeland security resource allocation.

2. WHY ROBUST OPTIMIZATION?

Addressing data uncertainty in mathematical programming models has long been recognized as an important issue in optimization. There are two principal methods proposed to address data uncertainty over the years: (1) stochastic-optimization programming, and (2) robust optimization. In stochastic-optimization problems, a common goal is to optimize the expected value of some objective function, while robust-optimization problems attempt to optimize the worst-case performance of the system. As early as the mid 1950s, Dantzig[27] introduced stochastic programming as an approach to model data uncertainty by assuming different scenarios for data to occur with different probabilities. The two main difficulties with such an approach are (1) knowing the exact distribution for the data, and (2) the size of the resulting optimization model increases drastically as a function of the number of scenarios, which poses substantial computational challenges.[28] In recent years, robust optimization has gained substantial popularity as a competing methodology in solving several types of stochastic-optimization models. It has been successful in immunizing uncertain mathematical optimization. The first step in this direction was taken by Soyster.[29] He proposed a worst-case scenario model for linear optimization. Subsequently, more elaborate uncertainty sets and computationally attractive robust-optimization methodologies were proposed by Ben-Tal and Nemirovski,[30,31] Goldfarb and Iyangar,[32] and Bertsimas and Sim.[28,33]

Robust optimization refers to the modeling of optimization problems with data uncertainty to obtain a solution that is guaranteed to be good for all, or most, of the possible realizations of the uncertain parameters. In robust-optimization approaches, the random parameters can be either continuous (which are restricted to lie in some prespecified intervals), or described by discrete scenarios (scenario-based approaches). The scenario approach has two main drawbacks. First, identifying scenarios is difficult and contaminated with speculations and inaccuracy. And second, due to the computational limitations, one tends to identify a relatively small number of scenarios. Using robust optimization in defender-attacker models appears superior since the attacker's attributes almost always contain uncertainty in regards to the defender. Defining probability distributions as a way to model uncertainty might be a tool, but applying distribution-free models in conservative situations related to human lives and the critical infrastructures of a country seems more appropriate.

The concept of robust optimization was first introduced in game theory by Aghassi and Bertsimas.[34] They presented a distribution-free model of incomplete information games in which the players use a robust-optimization approach to deal with payoff uncertainties. In their proposed robust game, they relax the assumptions of Harsanyi's Bayesian game model,[35] which assumes prior probability distributions for players' type. In this research, regarding the concept of *robust game theory* by Aghassi and Bertsimas[34] and the seminal work by Bertsimas and Thiele,[36] we propose a robust-optimization game-theoretical approach to model the uncertainty in attacker's attributes where uncertainty is characterized on bounded intervals based on the attacker's target valuation. In our proposed model, the defender can adjust the robustness of the solution against the level of conservatism, which reflects the reality that the defense allocations have enough robustness against any relaxation of the attacker's valuation of the targets.

3. PROBLEM FORMULATION

3.1. Notations

We study a one-shot Stackelberg game in which the defender, the Stackelberg leader, chooses her defense strategy first and then the attacker, the Stackelberg follower, observes this decision and makes his own strategy choice. We define the following parameters and variables to propose the model:

V_i: defender's valuation of target i,
u_i: attacker's valuation of target i,
$\widetilde{u_i}$: attacker's valuation of target i in the defender's belief,

$[u_i^-, u_i^+]$: bounded interval of attacker's valuation of target i in the defender's belief,
A_i: attacker's budget assigned to target i,
$L_i^a(u_i, D_i, A_i)$: attacker's expected damage of an attack on target i,
$\widetilde{L}_i^d(V_i, D_i, A_i)$: defender's expected damage of an attack on target i when the attacker's target valuation is unknown to the defender,
λ: effectiveness ratio of an attack,
D_i: defender's budget allocated to target i,
D: total budget of the defender,
$\overline{D}_i$: least possible defense level required to deter an attack on target i,
$\Omega(D_i)$: attacker's best response when the defender chooses D_i for $i = 1, \ldots, N$,
ψ: defender's robust-optimization equilibrium for the sequential game.

3.2. Damage Function Specifications

The strategic attacker observes the defense distribution among the targets and then chooses the target to launch an attack on in order to maximize his payoffs. In our proposed model, we assume that if target i is attacked, the expected damage (L_i) depends on three factors: (i) the attacker's effort assigned to target i, A_i; (ii) the defense allocations to target i, D_i; and (iii) the player's valuation of target i, that is, V_i, u_i, for the defender and the attacker, respectively. To define the player's optimization problem, let us specify a damage function with the following properties: $\widetilde{L}_i^d(V_i, D_i, A_i)$ and $L_i^a(u_i, D_i, A_i)$ are twice differentiable with respect to the defender's and attacker's effort, where $D_i > 0$, $L_i^a(u_i, D_i, 0) = 0$, and $\lim_{D_i \to \infty} L_i^a(u_i, D_i, A_i) = \lim_{D_i \to \infty} \widetilde{L}_i^d(V_i, D_i, A_i) = 0$. Meaning the expected damage from an attack on target i is zero when either the attacker's effort is zero or the defender's effort goes to infinity. An appropriate candidate for the expected damage function, which nicely possesses the mentioned properties, is the cumulative exponential function, as used in Bier *et al.*[13] and Golalikhani and Zhuang.[37] However, we recognize that there are other types of damage functions, such as linear functions as used by Golany *et al.*[20] and ratio functions as used by Zhuang and Bier.[19] So, let us define the expected damage from the standpoint of the defender and the attacker, respectively, as follows:

$$\widetilde{L}_i^d(V_i, D_i, A_i) = V_i(1 - e^{-\lambda A_i / D_i}), \tag{1}$$

$$L_i^a(u_i, D_i, A_i) = u_i(1 - e^{-\lambda A_i / D_i}). \tag{2}$$

In Equations (1) and (2), λ is the effectiveness ratio of an attack. One unit increment in ratio A_i/D_i increases the expected damage on target i about $100e^{-\lambda}\%$; λ may vary for different targets, but we only investigate the effect of different levels of λ on the defender's optimal solution in the following sections. Note that the value of a target depends on the decision-making environment and may consist of different measures. For example, the psychological impact of an attack on a target may be more critical than its imposed property loss (e.g., the terrorist attack on the World Trade Center in New York on September 11, 2001). To capture different aspects of an attack, we assume that the total value of a target, that is, V_i for the defender and u_i for the attacker, has three weighted dimensions: the target's monetary value to show its property, the target's mortality value to show the number of fatalities and injuries, and the target's strategic value to indicate the target's political and psychological importance. Assuming that the attacker has private information about each element of the mentioned weights imposes uncertainty about the defender's optimization problem. Throughout this article, for analytical convenience, we assume that the uncertainty about the attacker's valuation of the target happens in the attacker's total value of the target, that is, u_i.

3.3. Defender's Optimization Model

The defender seeks to find the optimal defensive resource allocation, D_i, which minimizes the total expected damage, $\sum_{i=1}^{N} \widetilde{L}_i^d(V_i, D_i, A_i)$, and the total defensive investment costs, $\sum_{i=1}^{N} D_i$. A strategic attacker observes the defense budget that is allocated to each target, and then chooses which target to attack to maximize his payoffs. As the defender is strategic, she does not leave target i purposefully undefended; otherwise, the attacker observes such a target and may launch an attack to add V_i to his payoffs with the least possible attack budget. Bier *et al.*[12] propose a model in which the defender optimally manipulates the attacker's behavior by leaving a target undefended. Levitin and Hausken[38] then study the deployment of false targets as part of a defense strategy. The defender wishes to minimize her

defense costs and expected losses:

$$\text{Min} \sum_{i=1}^{N} D_i + \sum_{i=1}^{N} I_{i=i^*} \widetilde{L}_i^d (V_i, D_i, A_i) \qquad (3)$$

subject to

$$\sum_{i=1}^{N} D_i \leq D \qquad , \qquad (4)$$

$$D_i \geq 0, \qquad i = 1, \ldots, N, \qquad (5)$$

where, $I_{i=i^*}$ is the binary indicator that takes 1 if target i is attacked by a strategic attacker, and 0 otherwise. Constraint (4) assures that the sum of defenses allocated to all targets cannot exceed the total budget of the defender, and Constraint (5) assures that D_i is a nonnegative variable. The defender then tries to minimize her expected loss by choosing robust responses against a strategic attacker who has private information about his target valuation. We will further discuss the defender's robust-optimization equilibrium in Section 4.

3.4. Attacker's Optimization Model

Recently, Wang and Bier[39] proposed an optimization attacker model based on a multi-attribute terrorist utility. Their model explores how intelligence about terrorists' preferences can affect optimal resource allocations for infrastructure protection. Traditionally, the attacker's behavior is defined on a set of attack probabilities, based on certain targets,[13,18] which is more suitable in modeling a nonstrategic attacker because such an attacker decides regardless of the defense allocations. Assigning a probability distribution of the uncertainty of attacks to certain targets in facing a strategic attacker is limited, since a strategic attacker may choose the targets after observing the defense distribution among all targets.

Following Powell[18] and Bier *et al.*,[13] we assume that the strategic attacker analyzes the defense distribution; thus, concentrating his attack budget to launch an attack on only one target, the target that maximizes his payoffs. The attacker gets his highest utility through following the equation:

$$\begin{aligned} & L_i^a (u_{i^*}, D_{i^*}, A_{i^*}) - A_{i^*} , \\ & i^* = \arg\max \left\{ L_k^a (u_k, D_k, A_k) - A_k \right\} \\ & \forall\, k = 1, \ldots, N. \end{aligned} \qquad (6)$$

Note that, in Equation (6), the attacker chooses his efforts based on his private information about target valuation, which subsequently contaminates the defender's optimization problem with uncertainty.

4. SEQUENTIAL MOVES: STACKELBERG EQUILIBRIUM APPROACH

4.1. Attacker's Best Response

The concept in this section is based on the work by Zhuang and Bier.[19] The defender employs the attacker's best response to estimate the attacker's effort. This is a key feature of our research; it makes it possible to define prespecified intervals to consider uncertainty in the attacker's attributes. Satisfying the optimality condition, $\frac{\partial [L_{i^*}^a (u_{i^*}, D_{i^*}, A_{i^*}) - A_{i^*}]}{\partial A_{i^*}} = 0$, yields the attacker's best response as follows:

$$\Omega(D_{i^*}) = \begin{cases} 0 & \text{if } D_{i^*} \geq \bar{D}_{i^*}, \\ \dfrac{D_{i^*}}{\lambda} \ln\left(\dfrac{\lambda u_{i^*}}{D_{i^*}}\right) & \text{otherwise.} \end{cases} \qquad (7)$$

where $i^* = \arg\max\{L_k^a(u_k, D_k, A_k) - A_k\}, \forall k = 1, \ldots, N$. Regarding Equation (7), $\bar{D}_i = \lambda u_i$ is the least possible level of defense required to deter an attack on target i^* that satisfies the $\Omega(D_{i^*}) = 0$. If the attacker's valuation of target i^* becomes small such that $u_{i^*} \leq D_{i^*}/\lambda$, then the attacker does not attack target i, which makes $\Omega(D_{i^*}) = 0$. Increasing the attacker's target valuation induces the attacker to launch an attack in a way that if the attacker's valuation of target i becomes high such that $u_{i^*} > D_{i^*}/\lambda$, then the attacker's response, $\Omega(D_{i^*})$, will be initially increasing in D_{i^*} for $0 \leq D_{i^*} < \lambda u_{i^*}/e$, and then decreasing in D_{i^*} for $\lambda u_{i^*}/e \leq D_{i^*} < \bar{D}_{i^*}$, and finally zero for $D_{i^*} \geq \bar{D}_{i^*}$, in which the attacker will be completely deterred.

As we show in Proposition 1, some of the targets may be left undefended in the optimal defense allocation. On the other hand, the attacker has a continuous-attack level and can benefit the whole value u_i of an undefended target (from Equation (2)) by incurring an infinitesimal attack level. In this article, we assume that the attacker incurs a fixed attack cost to strike the target, as studied in Bier *et al.*,[12] Zhuang *et al.*,[23] Golalikhani and Zhuang,[37] Konrad,[40] Zhuang and Bier,[41] Dighe *et al.*,[42] Zhuang,[43] and Bier and Haphuriwat.[44]

4.2. Defender Robust-Optimization Equilibrium

This section aims to show how the defender can incorporate the uncertainty in the attacker's valuation of targets in her optimal decision while no prior belief exists about such an uncertainty. Before finding the defender's robust-optimization model, let us first find out the structure of the defender's optimal strategy.

PROPOSITION 1. *In a sequential game between a defender with optimization model (3)–(5) and an attacker with optimization model (6), there is an optimal threshold for each target to be defended in equilibrium such that target* i *is defended if and only if the defender's valuation of the target is greater than its optimal threshold. Such a threshold is increasing in the attacker's valuation of the target, but it is decreasing in the defender's budget.*

Proof: See the Appendix.

REMARKS: First note that in the sequential game, some targets may be left undefended in equilibrium. This result may critically depend on our assumption that the attacker could attack at most one target. If this assumption was relaxed, the defender would defend more targets, since the attacker could destroy those economically undefended targets with negligible costs. Second, note that the defender, as the leader, tries to choose her best effort by plugging the attacker's best response (7) into her optimization model (3), but there are unknown parameters, $\tilde{u}_{i^*}$, in the attacker's best response in the defender's point of view, necessitating robust decisions by the defender. The traditional way to handle such a static incomplete information game was proposed by Harsanyi,[35] which determined the Bayes-Nash equilibria. It assumes that the players have conditional probability distributions derived from a certain probability distribution over the parameters unknown to the various players. However, the equilibria are seriously affected by prior assumptions about the probability distributions. In Aghassi and Bertsimas,[34] the authors proposed a distribution-free model of incomplete information games, which relaxes the assumptions of the Harsanyi's Bayesian model, in which the players use robust-optimization approaches to contend with payoff uncertainties.

Before determining the defender's robust-optimization equilibrium, let us assume that the defender can define a free-distribution interval for the attacker's target valuation such that $\tilde{u}_i \in [u_i^-, u_i^+]$. The assessment of the lower and higher bounds is the main obstacle in using the robust-optimization approach when uncertainty is characterized on bounded intervals, but on the other hand, the decisionmaker can extend the length of an interval to satisfy her belief about uncertain parameters. Assuming $\bar{l}_i = 1/\tilde{u}_i$, alternatively we can say $\bar{l}_i \in [1/u_i^+, 1/u_i^-]$, which is centered at its nominal value $\bar{l}_i$ and of half-length $\hat{l}_i$, but its exact value is unknown. The nominal value may be the decisionmaker's prior belief about the value of uncertain parameters, which is, in our case, the defender's prior belief about the attacker's valuation of the targets. In particular, the defender may assume her valuation of target i, which is V_i, as the nominal value for the attacker's valuation of target i, which is u_i, and then define the extreme bounds as a fraction of V_i, for example, $[0.5V_i, 1.5V_i]$. The nominal value and the half-length value can be determined by $\bar{l}_i = (\frac{1}{u_i^-} + \frac{1}{u_i^+})/2$ and $\hat{l}_i = (\frac{1}{u_i^-} - \frac{1}{u_i^+})/2$. It is a practical approach to adjust the level of conservatism in the solution so that a reasonable tradeoff between robustness and performance is achieved. We define the scaled deviation of the uncertain parameter $\hat{l}_i$ from its nominal value as $z_i = \hat{l}_i - \bar{l}_i/\hat{l}_i$. It is clear that the scaled deviation takes on a value in [–1, 1]. Moreover, following Bertsimas and Sim,[33] we impose a constraint on uncertainty in the following way: the total scaled deviation of the uncertainty parameters cannot exceed some threshold Γ, called the *budget of uncertainty*, which leads to:

$$\sum_{i \in J} |Z_i| \leq \Gamma, \tag{8}$$

where J is the set of indexes of the uncertain parameters. By taking $\Gamma = 0$ (or, $\Gamma = |J|$), we obtain the nominal (or, worst) case, respectively. Bertsimas and Sim[33] show that having the threshold Γ varied in $(0, |J|)$ allows greater flexibility to build a robust model without excessively affecting the optimal cost. In order to avoid assigning certain probability distributions to the attacker's target valuation, we consider the budget of uncertainty in the defender's robust optimization. To define the defender's robust-optimization equilibrium, consider the following set:

$$\Lambda = \left\{ \bar{l}_i \in [\bar{l}_i - \hat{l}_i, \bar{l}_i + \hat{l}_i] \forall i \in J, \sum_{i \in J} [|\bar{l}_i - \bar{l}_i|/\hat{l}_i] \leq \Gamma \right\}, \tag{9}$$

where Λ includes all uncertain variables in the model whose total scaled deviations cannot exceed the threshold Γ. We now use the concept of *robust game*

theory proposed by Aghassi and Bertsimas[34] to obtain the robust-optimization equilibrium for the defender by plugging the attacker's best response into the defender's optimization model, which creates the following equilibrium:

$$\psi = \underset{\substack{\sum_{i=1}^{N} D_i \le D \\ \tilde{l}_i \in \Lambda}}{\arg\min} \sum_{i=1}^{N} D_i + \sum_{i=1}^{N} I_{i=i^*} V_i \left(1 - \frac{\tilde{l}_i D_i}{\lambda}\right), \quad (10)$$

where $\tilde{l}_i = 1/\tilde{u}_i$. Note that Equilibrium (10) is a linear robust-optimization model with respect to the defender's effort where uncertainty is defined on bounded intervals. Bertsimas and Sim[28] show that uncertain linear programming problems can be solved as a more complex linear programming problem by reformulating their robust linear counterparts. Moreover, Bertsimas and Thiele[36] propose the *budget of uncertainty* to flexibly adjust the level of conservatism of robust solutions. To characterize the defender's robust-optimization equilibrium stated in Equilibrium (10), we use the robust-optimization approach presented by Bertsimas and Thiele,[36] and we also describe how the concept of the *budget of uncertainty* can be employed to adjust the robustness of the defender's decision against the uncertainty in the attacker's attributes. The attacker chooses the only target that maximizes the expected damage while the defender tries to minimize damage, which enables us to have the equivalent defender's robust-optimization problem as follows:

$$\text{Min} \sum_{i=1}^{N} D_i + z \quad (11)$$

subject to

$$V_i \left(1 - \frac{\tilde{l}_i D_i}{\lambda}\right) \le z, \qquad i = 1, \ldots, N, \quad (12)$$

$$\sum_{i=1}^{N} D_i \le D, \quad (13)$$

$$D_i \ge 0, \; i = 1, \ldots, N, \quad (14)$$

$$\tilde{l}_i \in \Lambda. \quad (15)$$

Note that, in the above optimization model, constraint set (12) assures that the expected damage of target i is counted in the objective function only if the attacker chooses target i to launch his attack on. The model presented in (11)–(15) is a robust linear optimization problem, where uncertainty incurs in the constraints' coefficients. It can be reformulated to its robust linear counterpart by linearization techniques presented by Bertsimas and Thiele.[36] We refer the reader to Bertsimas and Thiele[36] to see how the robust linear counterparts of our above model can be obtained.

4.3. The Role of Budget of Uncertainty

To adjust the robustness of defensive budget allocations to the defender's level of conservatism, we impose a budget of uncertainty, Γ, that is the total scaled deviation of the uncertain parameters, $\sum_{i \in J} |z_i|$, which cannot exceed in Γ. In particular, Γ is an exogenous variable in the model that controls the uncertainty in decision making. We will show that the robustness of the solution incurs more cost to the decisionmaker, which is called the *price of robustness*. As the budget of uncertainty controls the total scaled deviation of uncertain parameters, the upper value of Γ is equal to the number of uncertain parameters. The idea here is that the most perturbation in the model happens when all uncertain parameters take on the values in their upper or lower bounds in prespecified intervals, and in this case $z_i = 1$ or $z_i = -1, \forall i \in J$, respectively, and consequently $\sum_{i \in J} |z_i|$ equals the number of members of l, where l includes all uncertain parameters.

Bertsimas and Sim[33] show that in the model with N constraints, each of them has l_n uncertain parameters, and one can control the uncertainty in constraint n by considering $0 \le \Gamma_n \le |J_n|$, where Γ_n is the budget of uncertainty related to the nth constraint. In the constraint set (12) in the defender's robust-optimization equilibrium, we have N constraints each having one uncertain parameter, which means $J_n = 1$ for $n = 1, \ldots, N$ and $0 \le \Gamma_n \le 1$.In the next section, we apply the proposed game to real data to see how the defender is able to trade off between the robustness of his solution and the level of his conservatism that imposes the price of robustness.

5. APPLYING ROBUST GAME TO REAL DATA

We apply our approach to data from the FY 2004 Grant Allocations, and also Willis *et al.*,[14] which provides estimates on the expected annual terrorism losses to the 10 most valuable urban areas of the United States (Table I). According to Section 3.2, the target value consists of three dimensions. In Table I, columns 2, 3, and 4 correspond to

Table I. Expected Property Losses, Expected Fatalities and Injuries, UASI Budget Allocations for the 10 Urban Areas with the Highest Losses ($ Million),[14] Air Departures from U.S. Department of Transportation (DOT)[13]

Urban Area	Expected Property Losses	Expected Fatalities & Injuries	Air Departures (Major & Minor Airports)	FY 2004 UASI Grant Allocations ($ Million)
New York	413	5350	23599	47
Chicago	115	1212	39949	34
San Francisco	57	472	19142	26
Washington, DC-MD-VA-WV	36	681	17253	29
Los Angeles-Long Beach	34	402	28816	40
Philadelphia, PA-NJ	21	199	13640	23
Boston, MA-NH	18	225	11625	19
Houston	11	160	20979	20
Newark	7.3	74	12827	15
Seattle-Bellevue-Everett	6.7	88	13578	17
Total	719	8863	201408	270

different dimensions of a target value. Specifically, the expected property losses correspond to the monetary value, the total number of fatalities and injuries show the mortality value, and the total air departures from both major and minor airports indicate the political value. The total target valuation is $719 million and the total defensive budget is $270 million.

We study how the defender's optimal strategy is affected by the uncertainty on the attacker's attributes by solving the defender's robust-optimization equilibrium (11)–(15) in different values of Γ (budget of uncertainty). The length of the bounded interval that the defender assumes about the attacker's valuation of target i is also a critical factor in the defenses allocated to target i, which leads us to determine the defender's optimal solution for different lengths of bounded intervals $(u_i^+ - u_i^-)$. We also study the effect of robustness of the defender's solution on the objective value of nominal problems using the concept of the price of robustness. Finally, the impact of the effectiveness ratio of attack, λ, on the defense allocation is studied. While the goal of this extensive example is to study the uncertainty in the attacker's value of targets on defense allocations, simplified later, we assume that the target value equals the expected property loss. One can find the results when the target value equals other measures in the Appendix.

For the effectiveness ratio of attack $\lambda = 0.05$, Fig. 1 shows how the defensive budget allocations may vary among targets for different values of Γ, and also when the defender extends the extreme bounds of the intervals that she assumes concerning the attacker's valuation on certain targets. The figures on the left panel shows the portion (%) of each target from the budget allocated to all targets, and the right panel shows the absolute value ($ million) of the budget allocated to each target.

Regarding the left panel of Fig. 1, when the defender becomes more uncertain about the attacker's valuation of targets, that is, when Γ increases, the defense allocation highly depends on the length of the interval that the defender assumes about the attacker's value of the targets. Extending the extreme bounds of the interval, when each uncertain parameter can violate more than 80% of its scaled deviation, the defensive budget is mostly allocated to more valuable targets. In particular, when the defender assumes $\tilde{u}_i \in [0.1V_i, 10V_i]$, then most of the defensive budget is allocated to New York and Chicago as the two most valuable targets. On the other hand, when the defender shortens the length of the intervals, when $\tilde{u}_i \in [0.8V_i, 1.25V_i]$, then she faces less uncertainty about the extreme bounds on the intervals, which results in the defensive budget distribution among the targets. From the right panel figures, it is intuitively inferred that the total defensive budget increases in the budget of uncertainty. Indeed, the defender increases the defense level when she becomes more uncertain about the attacker's value of the targets.

Fig. 2 shows how much extra loss the defender would have to suffer in comparison with the baseline case, if she wants to assure a specific level of robustness. We call this extra loss the *price of robustness*. Assuming larger values of the budget of uncertainty gives more robust solutions. However, it is clearly inferred from Fig. 2 that, for any length of interval, the more robust the solution is, the greater the price of robustness the defender would endure. The

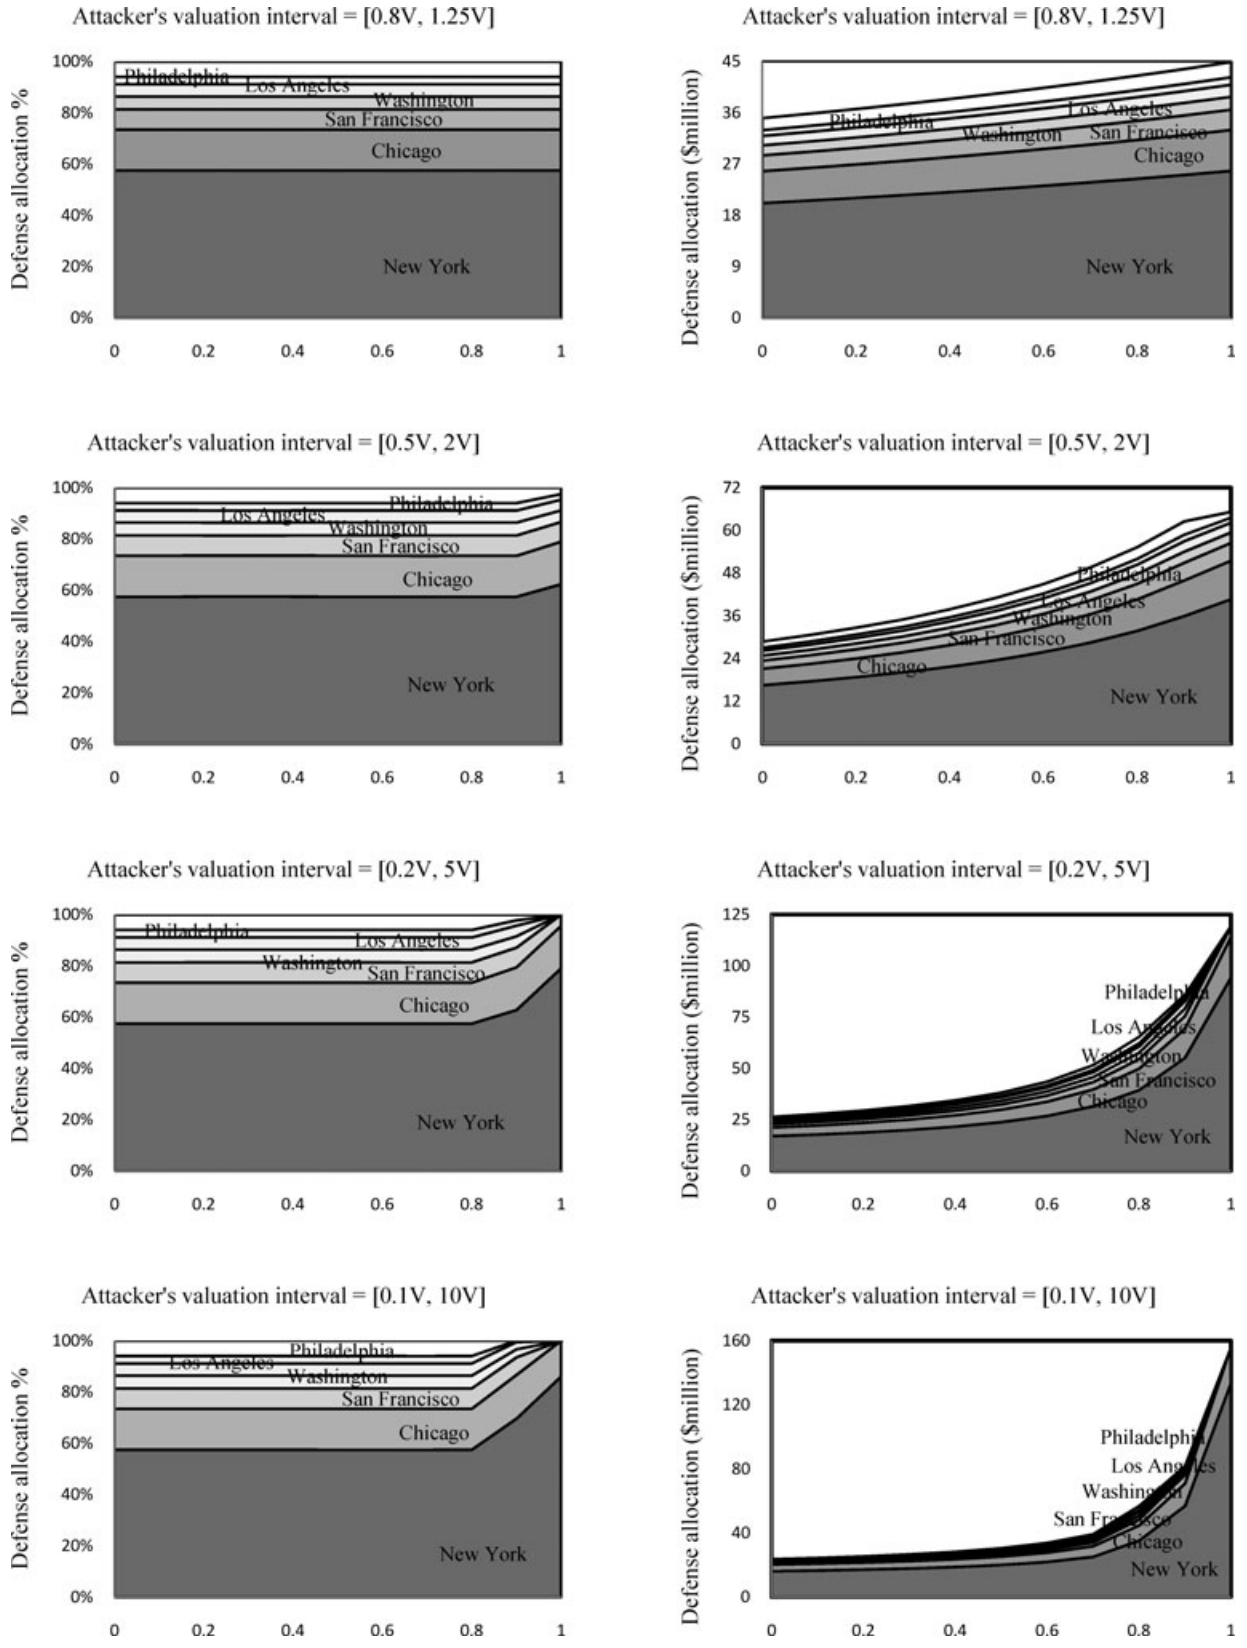

Fig. 1. Optimal defensive budget allocations among the seven most valued cities in the United States as a function of the budget of uncertainty (x-axis) when $\lambda = 0.05$, V: defender valuation of the target.

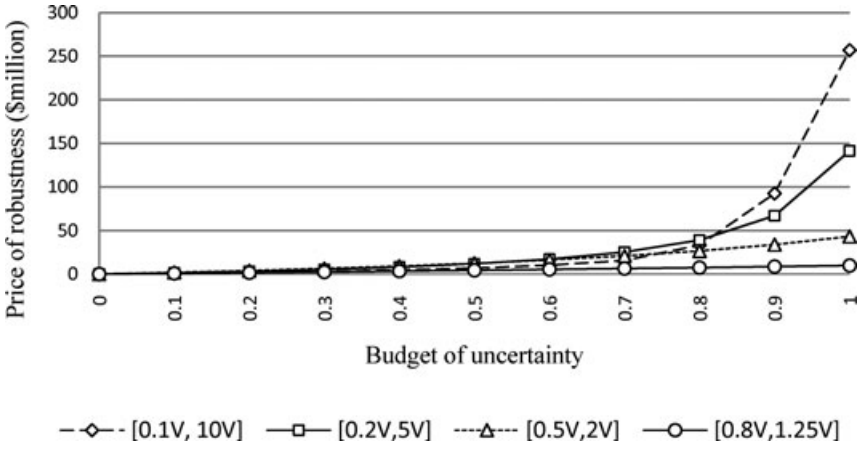

Fig. 2. The tradeoff between cost and robustness: price of robustness as a function of the budget of uncertainty when the effectiveness ratio of attack $\lambda = 0.05$.

increasing trend for the price of robustness as a function of Γ is more significant for longer intervals, which means that for a given level of robustness, the defender would suffer more of a loss when she extends the extreme bounds of the interval.

It is also inferred from Fig. 2 that increasing the robustness of the solution is more costly when the defender assumes longer intervals about uncertain parameters. For example, in Fig. 2, the extra cost to increase the robustness of the solution from the case when uncertain parameters are allowed to violate in 80% of their scaled deviations to the case that they are allowed to violate in 90% of those, is \$1.164 million for $[0.8V_i, 1.25V_i]$, \$7.214 million for $[0.5V_i, 2V_i]$, \$27.91 million for $[0.2V_i, 5V_i]$, and \$59.424 million for $[0.1V_i, 10V_i]$.

Now let us study the impact of the effectiveness ratio of an attack λ on the defender's optimal solution for different levels of robustness. In Fig. 3, we describe how the defensive budget distribution among the targets varies for different values of λ when the defender assumes that the attacker's valuation of target i belongs to $[0.2V_i, 5V_i]$.

From the left panel of Fig. 3, it is derived that the defensive budget is mainly distributed among the most valuable targets when the effectiveness ratio of attack λ goes to high values. Indeed, for high values of λ, the marginal benefit from hardening the most valuable targets is greater than that of defending low value targets. In particular for $\lambda = 0.02$ and $\lambda = 0.2$, as the low values for λ, the defensive budget is distributed among all targets; for $\lambda = 0.5$, it is distributed among the five most valuable targets: New York, Chicago, San Francisco, Washington, and Los Angeles; and for $\lambda = 1$, the defensive budget is distributed between only two of the most valuable targets: New York and Chicago. From the right panel figures, it is also inferred that, for a fixed value of Γ, the defender increases the total defensive budget among all targets when λ increases.

In Fig. 4, we study the price of robustness in terms of the budget of uncertainty in different values of λ.

Regarding Fig. 4, for a given level of robustness, when uncertain parameters violate less than 78% of their scaled deviations, increasing the robustness of the solution is more expensive when λ rises. For example, in Fig. 4, given the robustness for the defense allocations, when uncertain parameters have 60% freedom to violate, it costs \$6.866 million for $\lambda = 0.02$, \$60.52 million for $\lambda = 0.2$, \$108.68 million for $\lambda = 0.5$, and \$155.55 million for $\lambda = 1$. Note that, for big values of λ, the price of robustness no longer changes when uncertain parameters have high perturbation. For example, for $\lambda = 1$, when uncertain parameters have the freedom to violate more than 70% of their scaled deviations, the price of robustness is fixed to \$196.77 million. In fact, for large values of λ, when perturbation in uncertain parameters is high, then the defensive budget is mainly allocated to the most valuable targets, and, in this case, there is no more chance to increase the robustness of the solution.

6. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

In this article, we studied a sequential defender-attacker game with incomplete information in which the defender plays first, and then the strategic

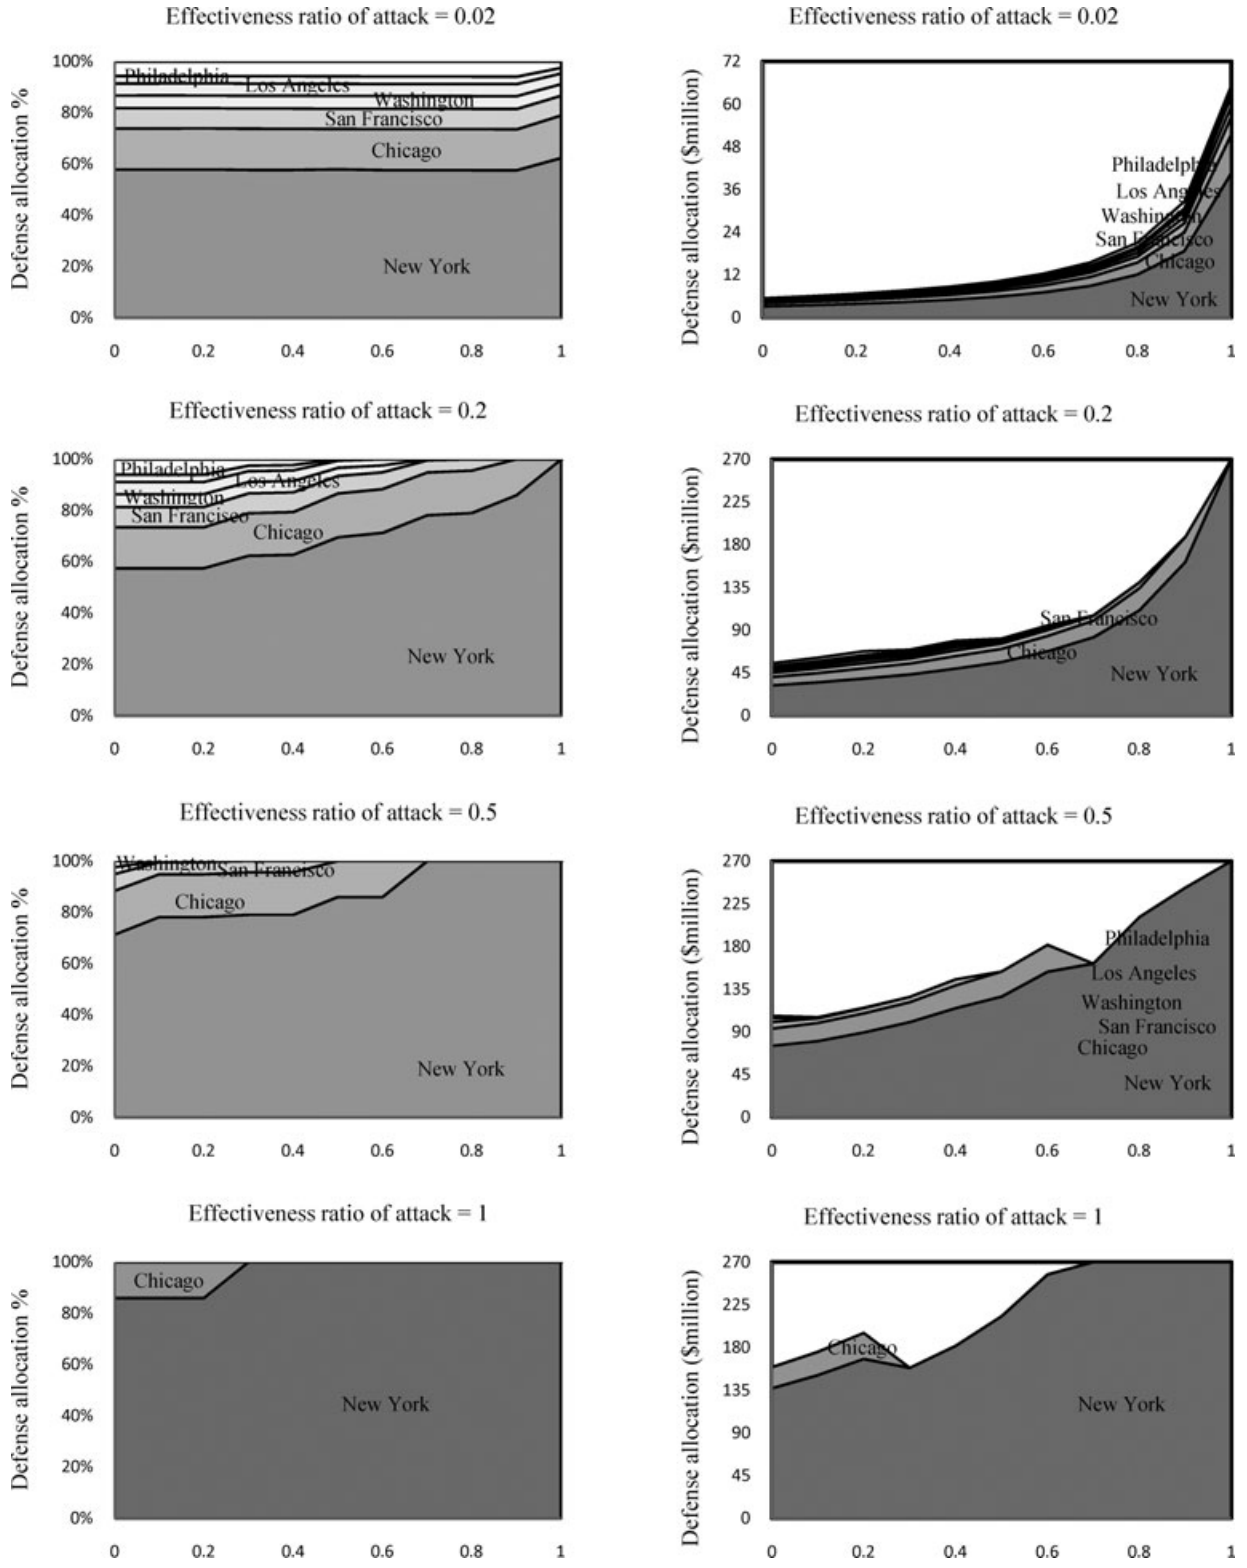

Fig. 3. Optimal defensive budget allocations among seven of the most valuable cities in United States as a function of the budget of uncertainty (x-axis) when $\tilde{u}_i \in [0.2V_i, 5V_i]$.

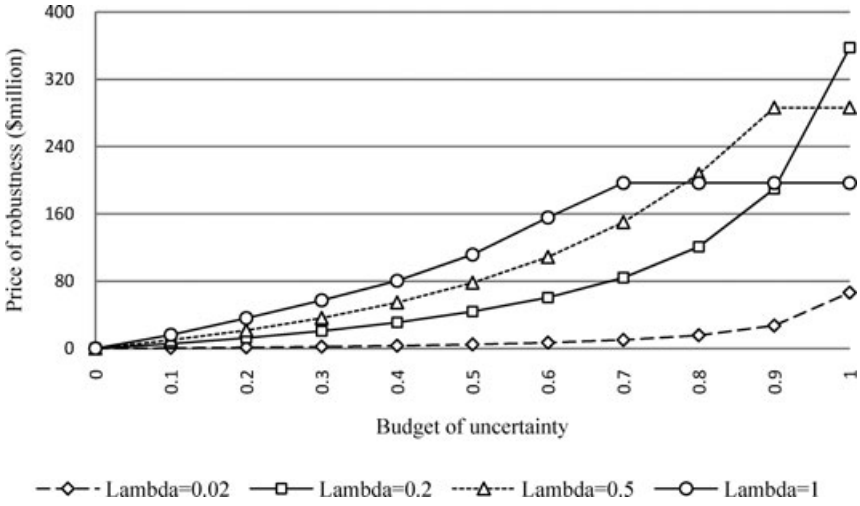

Fig. 4. The tradeoff between cost and robustness: the price of robustness as a function of the budget of uncertainty when $\tilde{u}_i \in [0.2V_i, 5V_i]$.

attacker, who has private information about his target valuation, observes the defenses, and centralizes his effort to launch only one attack. Modeling the attacker's attributes as uncertain parameters on bounded intervals, we proposed robust-optimization equilibrium for the defender, which gives her the flexibility to adjust the robustness of the solution to the level of her conservatism. By applying our proposed approach to real data, we studied the effect of the defender's assumptions of the extreme bounds of the uncertain parameters on the robustness of her solution, and also the impacts of the effectiveness ratio of attack on the robustness of the defender's solution.

In practice, it is difficult to assess attacker private information, including target valuation, justifying the robust approach proposed in this article. For reference on methods of compiling information about how the attacker and defender value targets, see Keeney.[45]

This article can be extended by considering a scenario when the attacker tries to attack on multiple targets in a simultaneous game, where the value of a set of targets is more than the sum of each individual target. Such a model would be a robust version of Colonel Blotto games. In our model, the uncertainty is defined as intervals with known upper and lower bounds. Such a way to model uncertainty is applicable only when we handle an uncertain parameter with continuous levels. One can extend the proposed model in this article to focus on facing with nonstrategic attackers, as nonstrategic attacking is considered exogenously (e.g., he may only strike the most valuable target, regardless of the observed defense levels).

ACKNOWLEDGMENTS

This research was supported by the U.S. Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the U.S. Department of Homeland Security, or CREATE. We thank Prof. Vicki M. Bier (University of Wisconsin–Madison), Area Editor Dr. Tony Cox, and three anonymous referees for their helpful comments. We also thank Ms. Elizabeth Newell (University at Buffalo) for editorial help. Author names were listed alphabetically by last name.

APPENDIX

Proof of Proposition 1. In a sequential game on hand, we need first extract the attacker's best response function (presented as Equation (7)) and then plug it into the defender's optimization problem (3)–(5) to obtain her best response. Consider the defender's optimization problem as follows:

$$\text{Min} \sum_{i=1}^{N} D_i + \sum_{i=1}^{N} L_i^d(V_i, D_i, A_i)$$

subject to

$$\sum_{i=1}^{N} D_i \le D$$
$$D_i \ge 0, \qquad i = 1, \ldots, N.$$

To satisfy the optimality condition, first we need to relax the budget constraint. The common way to

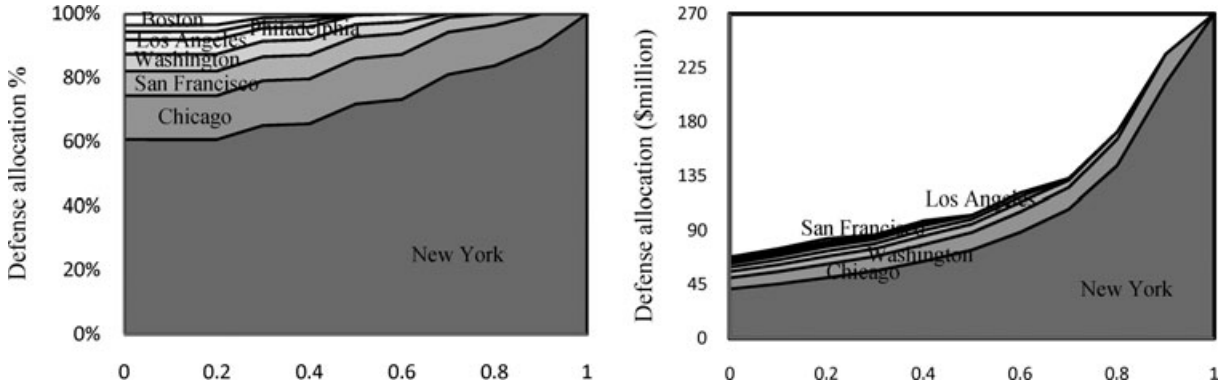

Fig. A1. Optimal defensive budget allocations among the most valuable cities in the United States as a function of the budget of uncertainty when target value corresponds for the mortality value (total number of Fatalities and injuries).

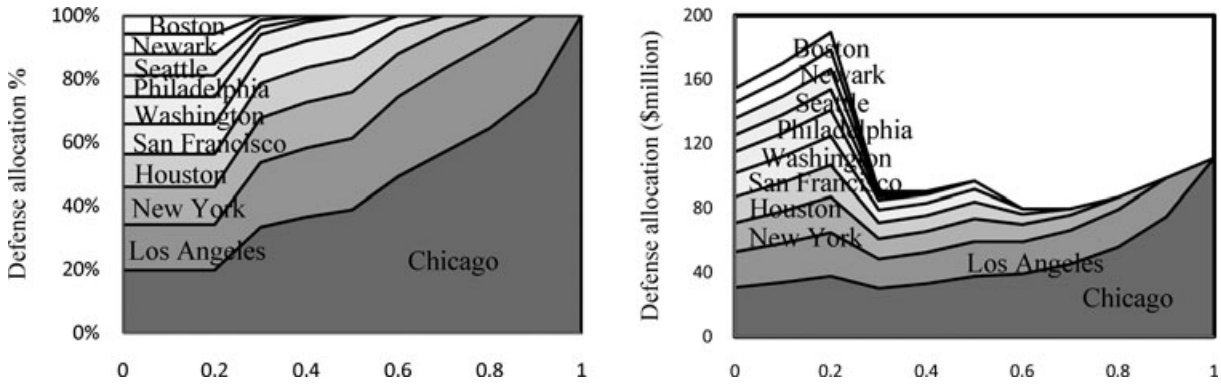

Fig. A2. Optimal defensive budget allocations among the most valuable cities in the United States as a function of the budget of uncertainty when target value corresponds to the political value (total air departures).

deal with such a condition is to penalize the objective function by $\mu(N\sum_{i=1}^{N} D_i - D)$, where $\mu \geq 0$ is the Lagrange multiplier. The Lagrangian function might be defined as follows:

$$L(D_i, \mu) = \sum_{i=1}^{N} D_i + \sum_{i=1}^{N} L_i^d(V_i, D_i, A_i) + \mu\left(\sum_{\iota=1}^{\mathrm{N}} D_i - D\right). \quad \text{(A1)}$$

Plugging the attacker's best response function into the Lagrangian function above gives the defender's optimization problem as follows:

$$L(D_i, \mu) = \sum_{i=1}^{N} D_i\left[1 - \frac{V_i}{\lambda u_i} + \mu\right] + \sum_{i=1}^{\mathrm{N}} V_i - \mu D. \quad \text{(A2)}$$

It is clear that in optimal solution of the above Lagrange function $D_i \geq 0$ if its coefficient is less than zero. In words, target i is defended if and only if its valuation is greater than a threshold, that is, $V_i > \lambda u_i(1 + \mu)$.

Optimal Defense Allocation for Different Dimensions of Target Value

Figs. A1 and A2 show the optimal defensive budget allocations among the most valuable cities in the United States as a function of the budget of uncertainty when target value, respectively, corresponds to the mortality value (total number of fatalities and injuries) and political value (total air departures). We assume that the attacker's valuation of target i belongs to $[0.2V_i, 5V_i]$, and the effectiveness ratio of attack $\lambda = 0.2$.

REFERENCES

1. Paté-Cornell E, Guikema S. Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. Military Operations Research, 2002; 7(4):5–23.

2. Harris B. Mathematical methods in combatting terrorism. Risk Analysis, 2004; 24(4):985–988.
3. Bakir NO. A decision tree model for evaluating countermeasures to secure cargo at United States southwestern ports of entry. Decision Analysis, 2008; 5(4):230–248.
4. Ezell BC, Bennett SP, Von Winterfeldt D, Sokolowski J, Collins AJ. Probabilistic risk analysis and terrorism risk. Risk Analysis, 2010; 30(4):575–589.
5. Rios Insua D, Rios J, Banks D. Adversarial risk analysis. Journal of the American Statistical Association, 2009; 104(486):841–854.
6. Parnell GS, Smith CM, Moxley FI. Intelligent adversary risk analysis: A bioterrorism risk management model. Risk Analysis, 2010; 30(1):32–48.
7. Merrick J, Parnell GS. A comparative analysis of PRA and intelligent adversary methods for counterterrorism risk management. Risk Analysis, 2011. Available at: http://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924-2011.01590.x/full.
8. Cox LAT. Jr. Game theory and risk analysis. Risk Analysis, 2009; 29(8):1062–1068.
9. Hall JR Jr. The elephant in the room is called game theory. Risk Analysis, 2009; 29(8):1061–1061.
10. Hausken K. Probabilistic risk analysis and game theory. Risk Analysis, 2002; 22(1):17–27.
11. Levitin G, Hausken K. Resource distribution in multiple attacks against a single target. Risk Analysis, 2010; 30(8):1231–1239.
12. Bier V, Oliveros S, Samuelson L. Choosing what to protect: Strategic defensive allocation against an unknown attacker. Journal of Public Economic Theory, 2007; 9(4):563–587.
13. Bier VM, Haphuriwat N, Menoyo J, Zimmerman R, Culpen AM. Optimal resource allocation for defense of targets based on differing measures of attractiveness. Risk Analysis, 2008; 28(3):763–770.
14. Willis HH, Morral AR, Kelly TK, Medby JJ. Estimating Terrorism Risk [Internet], 2005. Available at: http://www.rand.org/pubs/monographs/MG388.html. Accessed September 20, 2011.
15. Hao M, Jin S, Zhuang J. Robustness of optimal defensive resource allocations in the face of less fully rational attacker. Pp. 886–891 in: Proc. 2009 Indust. Engrg. Res. Conf., IIE, Norcross, GA, 2009.
16. Major JA. Advanced techniques for modeling terrorism risk. Journal of Risk Finance, 2002; 4(1):15–24.
17. Mooney SF. Are Terrorism Risks Really Uninsurable? National Underwriter Property: Casualty-Risk & Benefits Management, 2001.
18. Powell R. Defending against terrorist attacks with limited resources. American Political Science Review, 2007; 101(03):527–541.
19. Zhuang J, Bier VM. Balancing terrorism and natural disasters defensive strategy with endogenous attacker effort. Operations Research, 2007; 55(5):976–991.
20. Golany B, Kaplan EH, Marmur A, Rothblum UG. Nature plays with dice—Terrorists do not: Allocating resources to counter strategic versus probabilistic risks. European Journal of Operational Research, 2009; 192(1):198–208.
21. Levitin G, Hausken K. Redundancy vs. protection in defending parallel systems against unintentional and intentional impacts. Reliability, IEEE Transactions on, 2009; 58(4):679–690.
22. Brown G, Carlyle M, Salmeron J, Wood K. Defending critical infrastructure. Interfaces, 2006; 36(6):530–544.
23. Zhuang J, Bier VM, Alagoz O. Modeling secrecy and deception in a multiple-period attacker-defender signaling game. European Journal of Operational Research, 2010; 203(2):409–418.
24. Kaplan EH, Kress M, Szechtman R. Confronting entrenched insurgents. Operations Research, 2010; 58(2):329–341.
25. Shubik M, Weber RJ. Systems defense games: Colonel Blotto, command and control. Naval Research Logistics Quarterly, 1981; 28(2):281–287.
26. Roberson B. The Colonel Blotto game. Economic Theory, 2006; 29(1):1–24.
27. Dantzig GB. Linear programming under uncertainty. Pp. 1–11 in Infanger G (ed). Stochastic Programming. New York Springer, 2011.
28. Bertsimas D, Sim M. Robust discrete optimization and network flows. Mathematical Programming, 2003; 98(1):49–71.
29. Soyster AL. Convex programming with set-inclusive constraints and applications to inexact linear programming. Operations Research, 1973; 21(5):1154–1157.
30. Ben-Tal A, Nemirovski A. Robust convex optimization. Mathematics of Operations Research, 1998; 23(4):769–805.
31. Ben-Tal A, Nemirovski A. Robust solutions of uncertain linear programs. Operations Research Letters, 1999; 25(1):1–14.
32. Goldfarb D, Iyengar G. Robust convex quadratically constrained programs. Mathematical Programming, 2003; 97(3):495–515.
33. Bertsimas D, Sim M. The price of robustness. Operations Research, 2004; 52(1):35–53.
34. Aghassi M, Bertsimas D. Robust game theory. Mathematical Programming, 2006; 107(1):231–273.
35. Harsanyi JC. Games with incomplete information played by "Bayesian" players, I–III. Part I. The basic model. Management Science, 1967; 14(3):159–182.
36. Bertsimas D, Thiele A. A robust optimization approach to inventory theory. Operations Research, 2006; 54(1):150–168.
37. Golalikhani M, Zhuang J. Modeling arbitrary layers of continuous-level defenses in facing with strategic attackers. Risk Analysis, 2010; 31(4):533–547.
38. Levitin G, Hausken K. False targets efficiency in defense strategy. European Journal of Operational Research, 2009; 194(1):155–162.
39. Wang C, Bier VM. Impact of intelligence on target-hardening decisions. Pp. 373–380 in HST'09.IEEE Conference on Technologies for Homeland Security. Boston, MA, 2009.
40. Konrad KA. The investment problem in terrorism. Economica, 2004; 71(283):449–459.
41. Zhuang J, Bier V. Secrecy and deception at equilibrium, with applications to anti-terrorism resource allocation. Defence and Peace Economics, 2011; 22(1):43–61.
42. Dighe NS, Zhuang J, Bier VM. Secrecy in defensive allocations as a strategy for achieving more cost-effective attacker deterrence. International Journal of Performability Engineering, 2009; 5(1):31–43.
43. Zhuang J. Impacts of subsidized security on stability and total social costs of equilibrium solutions in an N-player game with errors. Engineering Economist, 2010; 55(2):131–149.
44. Bier VM, Haphuriwat N. Analytical method to identify the number of containers to inspect at US ports to deter terrorist attacks. Annals of Operations Research, 2011; 187(1):137–158.
45. Keeney RL. Modeling values for anti-terrorism analysis. Risk Analysis, 2007; 27(3):585–596.