

Addressing Scalability and Robustness in Security Games with Multiple Boundedly Rational Adversaries

Matthew Brown, William B. Haskell, and Milind Tambe

University of Southern California, Los Angeles, CA, USA

Abstract. Boundedly rational human adversaries pose a serious challenge to security because they deviate from the classical assumption of perfect rationality. An emerging trend in security game research addresses this challenge by using behavioral models such as quantal response (QR) and subjective utility quantal response (SUQR). These models improve the quality of the defender's strategy by more accurately modeling the decisions made by real human adversaries. Work on incorporating human behavioral models into security games has typically followed two threads. The first thread, scalability, seeks to develop efficient algorithms to design patrols for large-scale domains that protect against a single adversary. However, this thread cannot handle the common situation of multiple adversary types with heterogeneous behavioral models. Having multiple adversary types introduces considerable uncertainty into the defender's planning problem. The second thread, robustness, uses either Bayesian or maximin approaches to handle this uncertainty caused by multiple adversary types. However, the robust approach has so far not been able to scale up to complex, large-scale security games. Thus, each of these two threads alone fails to work in key real world security games. Our present work addresses this shortcoming and merges these two research threads to yield a scalable and robust algorithm, MIDAS (MaxImin Defense Against SUQR), for generating game-theoretic patrols to defend against multiple boundedly rational human adversaries. Given the size of the defender's optimization problem, the key component of MIDAS is incremental cut and strategy generation using a master/slave optimization approach. Innovations in MIDAS include (i) a maximin mixed-integer linear programming formulation in the master and (ii) a compact transition graph formulation in the slave. Additionally, we provide a theoretical analysis of our new model and report its performance in simulations. In collaboration with the United States Coast Guard (USCG), we consider the problem of defending fishery stocks from illegal fishing in the Gulf of Mexico and use MIDAS to handle heterogeneity in adversary types (i.e., illegal fishermen) in order to construct robust patrol strategies for USCG assets.

1 Introduction

Incorporating human behavioral models [11,3] into security games represents an important progression that has been demonstrated to improve the performance

of defender patrol strategies in both simulations and human subject experiments [15,19,18,13]. Behavioral models allow for the relaxation of the one of the strongest assumptions in classical game theory: namely, that the adversary is a perfectly rational utility maximizer. Instead, behavioral models, such as the quantal response (QR) model [11] and the subjective utility quantal response (SUQR) model [13], feature stochasticity in human decision making. These models are able to better predict the actions of real human adversaries and thus lead the defender to choose strategies that perform better in practice. Boundedly rational human behavioral models raise two fundamental research challenges that previous work has tried to address separately: scalability and robustness.

While perhaps counter-intuitive, modeling adversaries which behave suboptimally actually makes the defender’s optimization problem computationally more difficult. Both QR and SUQR are non-linear models and are difficult to use directly in large-scale security domains. This issue of scalability for large-scale security games with boundedly rational adversaries has received attention in the literature. [19] presented a mixed-integer linear programming (MILP) approximation for QR and SUQR models which improves tractability. Additionally, [18] introduces a cutting planes approach which can handle general patrol schedules and uses a master-slave formulation to iteratively generate deep cuts. We emphasize that the work [19,18] only allows for a single boundedly rational adversary.

However, in many domains the defender could encounter multiple different types of boundedly rational human adversaries. Thus, a separate line of security games research has focused on achieving robustness against uncertainty in the true adversary model. [17] proposed a Bayesian approach which learns a Gaussian distribution over adversary types. This approach has two potential drawbacks. First, the assumption that the adversary types are normally distributed is difficult to justify in practice. Second, even if the adversaries are normally distributed, a large amount of data is needed to learn the Gaussian distribution. Alternatively, [5] introduced a *maximin* approach which does not use a distribution over the adversary types. Instead, the defender chooses a patrol that maximizes the worst-case expected defender reward over a set of adversary types. In an effort to scale up, [17,5] focused on security games with a simplified defender strategy space that do not have complicated patrol schedules.

In this paper, we merge these two research threads for the first time by addressing scalability and robustness simultaneously. Each thread alone is impractical for important real-world security domains, such as environmental crime. Security games with complicated patrol schedules *and* multiple boundedly rational adversary types present a number of modeling and computational challenges. However, overcoming these challenges is critical as they are precisely the characteristics that define real-world security games. Our main contribution here is MIDAS (MaxImin Defense Against SUQR) which computes robust defender patrols for large-scale security games with a heterogeneous adversary population. Building off the insights of [19,18,17,5], we offer two key innovations: (i) a *robust* model that generates patrols that hedge against uncertainty about a heterogeneous population of adversaries and (ii) a *tractable* MILP approximation of

our robust problem. We develop key theoretical properties of MIDAS and also compare MIDAS against previous approaches in simulation.

In collaboration with the United States Coast Guard (USCG), we have applied MIDAS to protect fisheries in the Gulf of Mexico, where illegal, unreported, and unregulated (IUU) fishing seriously threatens the health of local fish stocks. The USCG has both surface and air assets with which to deter IUU fishing. We frame the interaction between the USCG and illegal fisherman from Mexico (henceforth called Lanchas) as a Stackelberg security game. By using historical data on Lancha sightings, we learn and construct a set of SUQR adversary types. However, there is not sufficient data to accurately construct a probability distribution over Lancha types. Generation of robust defender strategies for this domain has previously been explored in [5]. However, that work was more of a hot spot prediction model and it did not account for actual USCG schedules. In contrast, MIDAS constructs patrol schedules directly, resulting in higher quality patrol schedules for the USCG. The USCG began live testing of patrol schedules generated using MIDAS in July 2014.

2 Related Work

Game theory has been successfully applied to security problems such as the protection of networks [9,12,14] and physical infrastructure [16]. In particular, the Stackelberg game model with its leader-follower paradigm has been used extensively in security domains. Stackelberg games capture the fact that, in the real world, the defender (i.e., the security agency) must commit first to a strategy that may be observed and then exploited by adversaries. Given this first mover advantage, it is critical to understand and predict how adversaries will respond to a given strategy in order to find the best strategy. Classical game theory assumes that the adversary is perfectly rational and will always select the best available action in response to the defender's strategy. In some domains, such as network security [4,8], this assumption is reasonable as the game is played by software agents. For other domains, particularly those with human adversaries, a theoretically optimal defender strategy under standard rationality assumptions can perform poorly in practice. Under the assumption of perfect rationality, the adversary will always select just one action (the utility maximizing action). This assumption can lead to non-robust strategies for the defender.

As such, human behavioral models are becoming an increasingly important aspect of security games research. [19] was the first to address human adversaries in security games by incorporating the quantal response (QR) model [10] from the social psychology literature. QR predicts a probability distribution over adversary actions where actions with higher utility have a greater chance of being chosen. By anticipating possible adversary deviation from the optimal action, strategies computed with QR are more robust to uncertainty in human decision making. [7] generalized the QR model to be robust against all adversary models satisfying monotonicity (i.e., higher utility actions are selected more frequently than lower utility actions), but this approach struggles to scale up to

larger security games. [13] extended the QR model by proposing that humans use “subjective utility”, a weighted linear combination of factors (such as defender coverage, adversary reward, and adversary penalty), to make decisions. [13] proposes the subjective utility quantal response (SUQR) model which was shown to outperform QR in predicting the actions of participants of human subject experiments, thus leading to better defender strategies.

Building off that foundation, [18] presented an efficient cutting planes approach for solving security games with a large defender strategy space and a single adversary following a QR model. Meanwhile, two approaches have emerged for handling security games with multiple human adversary types. [17] utilized a Bayesian approach which learns a distribution over a set of SUQR types from available data. This distribution was assumed to be normal so as to minimize the number of parameters that need to be learned. Alternatively, [5] developed a robust version of [17] and applies it to the fishery protection domain where only limited data about the adversaries is available. Borrowing from the robust optimization literature [1,2], a *maximin* approach is used to optimize defender expected utility against the worst-case type from the set of possible adversary types. However, [18] handles only one adversary type, while [17] and [5] both fail to scale up. Neither of these two threads of research is individually able to handle the needs of security game applications in real-world domains such as environmental crime.

Most security problems do not feature static deployments, but rather have dynamic deployments that evolve in time and space. Thus, it is imperative to consider the capabilities of and restrictions on security resources such as personnel, cars, boats, and aircraft. Additionally, the adversaries in most physical security domains are likely to be humans, who have biases and limitations in their decision making process. This bounded rationality makes it difficult to predict the actions of the adversary and in turn for the defender to optimize their strategy. As a further complication, rather than a single adversary type there is usually a set of potential adversary types that may be encountered and it is critical to be robust against uncertainty in adversary type. Prior work on boundedly rational adversaries in security games has addressed only one of the challenges of scalability and robustness.

In this paper, we propose MIDAS which improves upon prior work by providing a holistic model that better captures the practicalities of large-scale, real-world security domains. More specifically, MIDAS enhances the incremental cut generation technique for solving large-scale security games with a single boundedly rational adversary type from [18] by using a robust *maximin* formulation for handling the uncertainty posed by multiple potential boundedly rational adversary types. Additionally, the QR model used in [18] for modeling boundedly rational adversary types is replaced with the SUQR model. Thus, MIDAS addresses the challenges of both scalability and robustness simultaneously, representing the first and only approach for solving security games with patrols schedules *and* multiple boundedly rational adversary types.

3 Background

We consider a Stackelberg security game (SSG) where the defender uses M available resources to protect a set of targets $T = \{1, \dots, |T|\}$ from a set of boundedly rational adversaries Ω . For the remainder of this paper we will focus on the SUQR behavioral model and treat $\omega \in \Omega$ as an SUQR adversary type. SUQR outperforms QR and other human behavioral models in human subject experiments. As a result, SUQR is widely considered to be the state of the art for modeling boundedly rational adversaries in security games.

Each target $t \in T$ is assigned a set of payoffs $\{R_t^a, P_t^a, R_t^d, P_t^d\}$: R_t^a is the reward earned by an adversary if they successfully attack target t , while P_t^a is the penalty received by an adversary for an unsuccessful attack on target t . Conversely, if the defender assigns a resource to protect target t and an adversary attacks target t , the defender receives a reward R_t^d . If an adversary attacks target t and the defender has not assigned a resource to protect target t , the defender receives a penalty P_t^d . It should be noted that the payoffs for all adversary types in Ω are identical, it is the parameters of the SUQR behavioral model that distinguish between types in Ω .

The defender commits to a mixed strategy that the adversaries are able to observe and then respond to by choosing a target to attack (Korzhyk, Conitzer, and Parr 2010; Basilico, Gatti, and Amigoni 2009). We denote the j^{th} defender pure strategy as A_j , which is an assignment of all the security resources. A_j is represented as a column vector $A_j = \langle A_{tj} \rangle^T$, where A_{tj} indicates whether target t is covered by A_j . For example, in an SSG with 4 targets and 2 resources, $A_j = \langle 1, 1, 0, 0 \rangle$ represents the pure strategy of assigning one resource to target 1 and another to target 2. Let $\mathcal{A} = \{A_j\}$ be the collection of feasible assignments of resources, i.e., the set of defender pure strategies. The defender's mixed strategy can then be represented as a vector $\mathbf{a} = \langle a_j \rangle$, where $a_j \in [0, 1]$ is the probability of choosing A_j . For large-scale security games, the number of pure strategies can grow so large that \mathcal{A} cannot be represented explicitly in practice making it impossible to optimize \mathbf{a} directly. However, there is a more compact "marginal" representation for defender strategies. Let \mathbf{x} be the marginal strategy, where $x_t = \sum_{A_j \in \mathcal{A}} a_j A_{tj}$ is the probability that target t is covered. The set of all feasible marginal distributions is

$$\mathcal{X}_f = \left\{ \mathbf{x} : x_t = \sum_{A_j \in \mathcal{A}} a_j A_{tj}, t \in T, \sum_{A_j \in \mathcal{A}} a_j = 1, \mathbf{a} \geq 0 \right\}.$$

We treat $\omega \in \Omega$ as an SUQR adversary type with the weight vector $\omega = \{\omega_1, \omega_2, \omega_3\}$ which encodes the relative importance of x_t , R_t^a , and P_t^a , respectively, in the decision making process of the adversary. Recall that the SUQR model selects a probability distribution over adversary actions rather than deterministically selecting the utility maximizing adversary action. Given defender strategy \mathbf{x} , the probability that adversary ω will attack target t is

$$q_t(\omega | \mathbf{x}) = \frac{e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a}}{\sum_{t'} e^{\omega_1 x_{t'} + \omega_2 R_{t'}^a + \omega_3 P_{t'}^a}}.$$

If an adversary chooses to attack target t , then for a given defender strategy \mathbf{x} , the defender's expected utility is

$$U_t(\mathbf{x}) = x_t R_t^d + (1 - x_t) P_t^d.$$

Against a known adversary type $\omega \in \Omega$, the defender's optimization problem is then

$$\max_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x} | \omega) \triangleq \sum_t U_t(\mathbf{x}) q_t(\omega | \mathbf{x}), \quad (1)$$

which can be solved for a defender mixed strategy \mathbf{a} . However, in this paper we consider an entire population of heterogeneous adversaries in Ω . Thus, the optimization problem above is inadequate.

4 Adversary Uncertainty

4.1 Bayesian Estimation

If we have a distribution \mathbb{P} over the set Ω of all possible types, then the expected utility maximizing problem is

$$\max_{\mathbf{x} \in \mathcal{X}_f} \int_{\Omega} F(\mathbf{x} | \omega) \mathbb{P}(d\omega). \quad (2)$$

Problem (2) maximizes the expected defender utility, where the expectation is over the adversary types. In practice Problem (2) requires \mathbb{P} to be estimated from sample data. Estimation of \mathbb{P} presents two potential issues: first, it assumes that the types in Ω are normally distributed in order to use convenient update rules; second, large amounts of data are required. This method is referred to as Bayesian SUQR [17].

4.2 Maximin

Robust optimization offers up remedies for the shortcomings of Bayesian SUQR. *Maximin* does not require large amounts of data, but it can still utilize data when it is available even if only in small quantities. It is also less sensitive to assumptions about the nature of the underlying data, for instance the assumption that \mathbb{P} is a normal distribution.

We treat Ω as an uncertainty set in line with robust optimization. For convenience, we assume that Ω is finite. This assumption is reasonable in practice since we will only ever have finitely many observations of the adversary. Then we solve the robust optimization problem

$$\max_{\mathbf{x} \in \mathcal{X}_f} \min_{\omega \in \Omega} F(\mathbf{x} | \omega) \quad (3)$$

to get a patrol for the defender, where again $F(\mathbf{x}|\omega)$ is the expected utility corresponding to type ω . Problem (3) is a nonlinear, nonconvex, nonsmooth optimization problem. For easier implementation, we transform Problem (3) into the constrained problem

$$\max_{\mathbf{x} \in \mathcal{X}_f, s \in \mathbb{R}} \{s : s \leq F(\mathbf{x}|\omega), \forall \omega \in \Omega\}, \quad (4)$$

by introducing a dummy variable $s \in \mathbb{R}$ to replace the nonsmooth objective with a collection of smooth constraints.

5 Mixed-Integer Linear Programming

By considering a human behavior model such as SUQR, Problem (4) becomes a nonlinear nonconvex optimization problem. In the general case, this problem class has been shown to be NP-hard to solve to optimality. Our idea in this section is to introduce a tractable MILP approximation scheme.

An approximate approach for solving Problem (1) with a single boundedly rational adversary was presented in [19,18]. This approach is based on a piecewise linear approximation that leads naturally to an MILP. In this section, we generalize this approach to create MIDAS, an algorithm for solving the robust Problem (4) with a set of boundedly rational adversaries.

First notice that, $F(\mathbf{x}|\omega)$, the defender's payoff against a single adversary type $\omega \in \Omega$ can be written out as

$$F(\mathbf{x}|\omega) = \sum_t U_t(\mathbf{x}) q_t(\omega|\mathbf{x}) = \frac{\sum_t ((R_t^d - P_t^d)x_t + P_t^d) e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a}}{\sum_t e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a}}$$

which is a fractional function $N(\mathbf{x}|\omega)/D(\mathbf{x}|\omega)$ where

$$N(\mathbf{x}|\omega) = \sum_t ((R_t^d - P_t^d)x_t + P_t^d) e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a}$$

and $D(\mathbf{x}|\omega) = \sum_t e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a}$. The goal in this section is to estimate the optimal value, which we will denote s^* , of Problem (4), i.e., the defender receives a payoff of at least s^* against every adversary type $\omega \in \Omega$. We use a binary search to compute s^* by updating a parameter r . We know that $r \leq s^*$ if there exists some $\mathbf{x} \in \mathcal{X}_f$ such that

$$r \leq \frac{N(\mathbf{x}|\omega)}{D(\mathbf{x}|\omega)}, \forall \omega \in \Omega.$$

Equivalently, we can rearrange the terms to require

$$r D(\mathbf{x}|\omega) - N(\mathbf{x}|\omega) \leq 0, \forall \omega \in \Omega.$$

Therefore, to check if $r \leq s^*$, we solve

$$\min_{\mathbf{x} \in \mathcal{X}_f, \xi \in \mathbb{R}} \{\xi : \xi \geq r D(\mathbf{x}|\omega) - N(\mathbf{x}|\omega), \forall \omega \in \Omega\}. \quad (5)$$

If the optimal value of the above problem is less than or equal to zero, then $r \leq s^*$; otherwise, $r > s^*$; then r is adjusted appropriately. However, Problem (5) is still nonlinear and nonconvex. Thus, we need to find a tractable approximation to implement this scheme.

5.1 Linear Approximation

The nonlinearity and nonconvexity of Problem (5), whose objective function is a summation of nonlinear functions in \mathbf{x} , can be overcome by approximating each nonlinear function with a piecewise linear function with K pieces. The functions $rD(\mathbf{x}|\omega) - N(\mathbf{x}|\omega)$ in the constraints of Problem (5) can be approximated with piecewise linear functions $L(\mathbf{x}|\omega)$ of the form:

$$L(\mathbf{x}|\omega) = \sum_{t \in T} (r - P_t^d) \left(f_t(0|\omega) + \sum_{k=1}^K \gamma_{\omega tk} x_{tk} \right) - \sum_{t \in T} (R_t^d - P_i^d) \sum_{k=1}^K \mu_{\omega tk} x_{tk}$$

where $\gamma_{\omega tk}$ is the slope of the function $f_t(x_t|\omega)$ in the k^{th} segment while $\mu_{\omega tk}$ is the corresponding slope of $x_t f_t(x_t|\omega)$. With this approximation, we then solve the feasibility check problem

$$\min_{\mathbf{x}, \xi} \xi \tag{6}$$

$$\text{s.t. } \xi \geq L(\mathbf{x}|\omega), \quad \forall \omega \in \Omega, \tag{7}$$

$$0 \leq x_{tk} \leq 1/K, \quad \forall t, \quad k = 1 \dots K, \tag{8}$$

$$z_{tk}/K \leq x_{tk}, \quad \forall t, \quad k = 1 \dots K - 1, \tag{9}$$

$$x_{t(k+1)} \leq z_{tk}, \quad \forall t, \quad k = 1 \dots K - 1, \tag{10}$$

$$z_{tk} \in \{0, 1\}, \quad \forall t, \quad k = 1 \dots K - 1, \tag{11}$$

$$x_t = \sum_{A_j \in \mathcal{A}} a_j A_{tj}, \quad \forall t, \tag{12}$$

$$\sum_{A_j \in \mathcal{A}} a_j = 1, \tag{13}$$

$$\mathbf{x}, \mathbf{a} \geq 0. \tag{14}$$

5.2 Column Generation

In this subsection we produce a tractable scheme for solving Problem (6) - (14). First, we derive a relaxation of Problem (6) - (14). Second, we show how to iteratively improve this approximation via a network flow problem: to that end Problem (6) - (14) is used to add new constraints to the relaxed version of the problem, and column generation is used in service of solving Problem (6) - (14) which then uses the network flow representation. Our network flow problem

differs substantially from earlier work, which focused on aviation security and environmental crime, because of the generality of our formulation.

To begin, we approximate the constraint $\mathbf{x} \in \mathcal{X}_f$ with a linear relaxation

$$\{\mathbf{x} : \hat{H} \mathbf{x} \leq \hat{h}\},$$

which represents a subset of linear boundaries of \mathcal{X}_f . Then we solve the relaxation

$$\max_{\mathbf{x}, s \in \mathbb{R}} \left\{ s : s \leq F(\mathbf{x} | \omega), \forall \omega \in \Omega, \hat{H} \mathbf{x} \leq \hat{h} \right\} \quad (15)$$

using the binary search method, i.e. Problem (6) - (14).

Given a candidate $\tilde{\mathbf{x}}$, we check if $\tilde{\mathbf{x}} \in \mathcal{X}_f$ by solving the projection problem

$$\min_{z \in \mathbb{R}^{|\mathcal{T}|}, \mathbf{a} \in \mathbb{R}^J} \sum_{t \in \mathcal{T}} z_t \quad (16)$$

$$\text{s.t. } A \mathbf{a} - \tilde{\mathbf{x}} \leq z, \quad (17)$$

$$-z \leq A \mathbf{a} - \tilde{\mathbf{x}}, \quad (18)$$

$$\sum_{A_j \in \mathcal{A}} a_j = 1, \quad (19)$$

$$\mathbf{a} \geq 0. \quad (20)$$

Problem (16) - (20) finds the best 1-norm approximation of \mathbf{x} in \mathcal{X}_f , and returns the optimal value zero if $\mathbf{x} \in \mathcal{X}_f$. Otherwise, we find a violated constraint which we add to the approximation $\hat{H} \mathbf{x} \leq \hat{h}$.

Problem (16) - (20) has a large number of variables since \mathcal{A} is exponentially large. We solve (16) - (20) using a column generation method similar to the one introduced in [6]. We solve a restriction of Problem (16) - (20) with a subset of columns $\hat{A} \subset A$ where a is now understood as a vector in $a \in \mathbb{R}^{|\hat{A}|}$, with $a_j = 0$ for all j with $A_j \notin \hat{A}$. Then we check for columns A_j to add to \hat{A} by computing the reduced costs of variables a_j with $A_j \notin \hat{A}$ via the dual problem.

The dual to Problem (16) - (20) is

$$\max_{y, u} \tilde{\mathbf{x}}^T y + u \quad (21)$$

$$\text{s.t. } A^T y + u \leq 0, \quad (22)$$

$$-1 \leq y \leq 1, \quad (23)$$

which has a large number of constraints due to the presence of the matrix A . For a subset of columns $\hat{A} \subset A$ (abusing notation since these are matrices), we have the relaxation of the dual

$$\max_{y, u} \tilde{\mathbf{x}}^T y + u \quad (24)$$

$$\text{s.t. } \hat{A}^T y + u \leq 0, \quad (25)$$

$$-1 \leq y \leq 1, \quad (26)$$

$$g \geq 0. \quad (27)$$

We are looking for a column A_j such that

$$A_j^T y + u \leq 0$$

is violated. So, we solve the slave problem

$$\max_{A_j \in \mathcal{A}} \{y^T A_j\} + u \tag{28}$$

and identify a violated constraint if the optimal value of this problem is positive. Specifically, we solve Problem (28) using the technique in [6], i.e. we use a maximum reward network flow problem (since Problem (28) is a maximization problem).

To setup this network flow problem, we create a source node with supply 1, and a sink node with demand 1. We have a fixed time horizon, $n = 0, 1, \dots, N$ stages, so we create a node (n, t) for every target and every time. The variables in this problem are the flow between nodes,

$$\mu^{(t,n), (t',n+1)}$$

which indicate a transition in the asset from target t at time n to target t' at time $n + 1$ in the next period. Effectively, we are taking a transition graph representation on the state space T^{N+1} . This formulation has the advantage of allowing us to express constraints on feasible patrols. The maximum reward network flow problem is then of the form

$$\max_{\mu} \left\{ \sum_{n \in \mathbb{N}} y_t \sum_{n,t} \mu^{(t,n), (t',n+1)} : \text{network flow constraints on } \mu \right\}.$$

The preceding network flow problem is a linear programming problem. This problem class is well studied and many efficient solution algorithms (such as the Simplex algorithm) exist that can obtain an exact optimal solution. We also point out that the preceding network flow problem can be solved efficiently for any underlying network topology.

6 Problem Properties

This section summarizes some key problem properties. The main points are to better understand our approximation scheme, to confirm that our cut generation scheme produces deep cuts, and to see how the standard Bayesian estimation approach relates to our robust approach.

6.1 MILP Approximation Error

Our underlying approach is a piecewise linear approximation to a nonconvex problem. We want to better understand the error bound for this approximation and the resulting solution quality of the corresponding MILP. We will show that

all of the nonconvex functions we are approximating have bounded Lipschitz constants. Thus, since their variability is bounded, we have an upper bound on the piecewise linear approximation error as a function of the fineness of the discretization.

Recall that we are approximating the feasibility check problem, which solves

$$\min_{\mathbf{x} \in \mathcal{X}_f} \max_{\omega \in \Omega} \{r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)\},$$

by linearly interpolating the functions $r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)$ for all $\omega \in \Omega$. The first step in our approximation analysis is to estimate the Lipschitz constant of $r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)$ for fixed $\omega \in \Omega$.

Lemma 1. *The Lipschitz constant of $r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)$ for any $\omega \in \Omega$ is bounded above by*

$$\sum_t e^{1+\max_t R_t^a + \max_t P_t^a} + \sum_t (R_t^d - P_t^d) e^{1+\max_t R_t^a + \max_t P_t^a}.$$

Proof. By direct computation, $r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)$ is equal to

$$r \sum_t e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a} - \sum_t ((R_t^d - P_t^d) x_t + P_t^d) e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a}.$$

So

$$\begin{aligned} & |r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega) - r D(\mathbf{x}' | \omega) + N(\mathbf{x}' | \omega)| \\ & \leq \sum_t |e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a} - \sum_t ((R_t^d - P_t^d) x_t + P_t^d) e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a} \\ & \quad - e^{\omega_1 x'_t + \omega_2 R_t^a + \omega_3 P_t^a} - \sum_t ((R_t^d - P_t^d) x'_t + P_t^d) e^{\omega_1 x'_t + \omega_2 R_t^a + \omega_3 P_t^a}| \\ & \leq \sum_t |e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a} - e^{\omega_1 x'_t + \omega_2 R_t^a + \omega_3 P_t^a}| \\ & \quad + \sum_t |((R_t^d - P_t^d) x_t + P_t^d) e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a} - ((R_t^d - P_t^d) x'_t + P_t^d) e^{\omega_1 x'_t + \omega_2 R_t^a + \omega_3 P_t^a}|. \end{aligned}$$

We have

$$|e^{\omega_1 x_t + \omega_2 R_t^a + \omega_3 P_t^a} - e^{\omega_1 x'_t + \omega_2 R_t^a + \omega_3 P_t^a}| \leq e^{\omega_2 R_t^a + \omega_3 P_t^a} e^{\omega_1} |x_t - x'_t|.$$

Additionally,

$$\begin{aligned} |x_t e^{\omega_1 x_t} - x'_t e^{\omega_1 x'_t}| & \leq |x_t e^{\omega_1 x_t} - x_t e^{\omega_1 x'_t}| + |x_t e^{\omega_1 x'_t} - x'_t e^{\omega_1 x'_t}| \\ & \leq x_t e^{\omega_1} |x_t - x'_t| + e^{\omega_1} |x_t - x'_t| \\ & \leq 2e^{\omega_1} |x_t - x'_t|. \end{aligned}$$

Now use the fact that $e^{\omega_2 R_t^a + \omega_3 P_t^a} e^{\omega_1}$ is bounded above by

$$e^{1+\max_t P_t^a + \max_t R_t^a},$$

and $2e^{\omega_1}$ is bounded above by $2e$. Using Lemma 2 and the triangle inequality, for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}_f$ we compute

$$\begin{aligned} & \left| \max_{\omega \in \Omega} \{r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)\} - \max_{\omega \in \Omega} \{r D(\mathbf{x}' | \omega) - N(\mathbf{x}' | \omega)\} \right| \\ & \leq r \max_{\omega \in \Omega} |D(\mathbf{x} | \omega) - D(\mathbf{x}' | \omega)| + \max_{\omega \in \Omega} |N(\mathbf{x} | \omega) - N(\mathbf{x}' | \omega)|. \end{aligned}$$

We can expand on the previous Lipschitz computation to produce an error estimate for the overall piecewise linear approximation, by using the following fact to bound the Lipschitz constant of

$$\max_{\omega \in \Omega} \{r D(\mathbf{x} | \omega) - N(\mathbf{x} | \omega)\}.$$

Lemma 2. *Let X be a given set, and $f_1 : X \rightarrow \mathbb{R}$ and $f_2 : X \rightarrow \mathbb{R}$ be two real-valued functions on X . Then,*

- (i) $|\inf_{x \in X} f_1(x) - \inf_{x \in X} f_2(x)| \leq \sup_{x \in X} |f_1(x) - f_2(x)|$, and
- (ii) $|\sup_{x \in X} f_1(x) - \sup_{x \in X} f_2(x)| \leq \sup_{x \in X} |f_1(x) - f_2(x)|$.

Proof. To verify part (i), note

$$\begin{aligned} \inf_{x \in X} f_1(x) &= \inf_{x \in X} \{f_1(x) + f_2(x) - f_2(x)\} \\ &\leq \inf_{x \in X} \{f_2(x) + |f_1(x) - f_2(x)|\} \\ &\leq \inf_{x \in X} \left\{ f_2(x) + \sup_{y \in Y} |f_1(y) - f_2(y)| \right\} \\ &\leq \inf_{x \in X} f_2(x) + \sup_{y \in Y} |f_1(y) - f_2(y)|, \end{aligned}$$

giving

$$\inf_{x \in X} f_1(x) - \inf_{x \in X} f_2(x) \leq \sup_{x \in X} |f_1(x) - f_2(x)|.$$

By the same reasoning,

$$\inf_{x \in X} f_2(x) - \inf_{x \in X} f_1(x) \leq \sup_{x \in X} |f_1(x) - f_2(x)|,$$

and the preceding two inequalities yield the desired result. Part (ii) follows similarly.

6.2 Projection

The feasible region of our problem, \mathcal{X}_f , is exactly the same as the one found in [18]. Thus, the results of the cut generation algorithm are unchanged and we obtain deep cuts. The results are repeated here for completeness.

Lemma 3. (i) *If $\tilde{\mathbf{x}} \notin \mathcal{X}_f$, let $(\mathbf{y}^*, \mathbf{g}^*, u^*)$ be the dual variables at the optimal solution of Problem ((16)) - ((20)). Then the hyperplane $(\mathbf{y}^*)^T \mathbf{x} - (\mathbf{g}^*)^T \mathbf{b} + u^* = 0$ separates $\tilde{\mathbf{x}}$ and \mathcal{X}_f .*

(ii) *Furthermore, $(\mathbf{y}^*)^T \mathbf{x} - (\mathbf{g}^*)^T \mathbf{b} + u^* = 0$ is a deep cut.*

As in [18], we now consider a modified norm minimization problem. The idea is that we weight the norm towards an optimal solution using local rate of change information about the objective. In our case, the objective $G(\mathbf{x}) = \min_{\omega \in \Omega} F(\mathbf{x} | \omega)$ is a nondifferentiable function, so we use the subgradient instead of the gradient. The subgradient is

$$\partial G(\mathbf{x}) = \text{conv} \{ \nabla_{\mathbf{x}} F(\mathbf{x} | \omega) : F(\mathbf{x} | \omega) = G(\mathbf{x}) \}.$$

For a subgradient $\mathbf{s} \in \partial G(\mathbf{x})$, we use the objective $\sum_t (\mathbf{s}_t + \xi) z_t$ where $\xi > 0$ is chosen so that $\mathbf{s}_t + \xi > 0$ for all t .

6.3 Duality

Here we comment on the relationship of our approach to Bayesian estimation. Bayesian estimation is a classical and widespread tool for incorporating information under uncertainty. To reveal this relationship, we compute the dual of the constrained variant of Problem (3) which we reprint here for convenience:

$$\max_{\mathbf{x} \in \mathcal{X}_f, \mathbf{s} \in \mathbb{R}} \{ \mathbf{s} : \mathbf{s} \leq F(\mathbf{x} | \omega), \forall \omega \in \Omega \}.$$

The constraints above cause Lagrange multipliers to appear; so we can compute the standard Lagrangian dual. To proceed we first introduce the Lagrange multipliers which lie in $\mathbb{R}^{|\Omega|}$ (since there are only finitely many adversary types). We let $\mathbb{R}_+^{|\Omega|}$ denote the set of nonnegative vectors in $\mathbb{R}^{|\Omega|}$.

Let

$$\mathcal{P}(\Omega) \triangleq \left\{ \Lambda \in \mathbb{R}_+^{|\Omega|} : \sum_{\omega \in \Omega} \Lambda(\omega) = 1 \right\}$$

be the space of probability measures on Ω , it is a subset of $\mathbb{R}^{|\Omega|}$. We will see shortly that these probability measures are the decision variables in the dual to Problem (4).

Theorem 1. *The dual to Problem (4) is*

$$\min_{\Lambda \in \mathcal{P}(\Omega)} \left\{ d(\Lambda) \triangleq \max_{\mathbf{x} \in \mathcal{X}_f} \sum_{\omega \in \Omega} F(\mathbf{x} | \omega) \Lambda(\omega) \right\}. \quad (29)$$

Proof. Let $\Lambda \in \mathbb{R}_+^{|\Omega|}$ be the Lagrange multiplier for the constraint $\mathbf{s} \leq F(\mathbf{x} | \omega)$ for all $\omega \in \Omega$. We obtain the Lagrangian

$$L(\mathbf{x}, \mathbf{s}, \Lambda) = \mathbf{s} + \sum_{\omega \in \Omega} [F(\mathbf{x} | \omega) - \mathbf{s}] \Lambda(\omega).$$

The Lagrangian dual problem is then

$$\min_{\Lambda \in \mathbb{R}_+^{|\Omega|}} \max_{\mathbf{x} \in \mathcal{X}_f, \mathbf{s} \in \mathbb{R}} \{ L(\mathbf{x}, \mathbf{s}, \Lambda) \}.$$

We see that the inner maximization problem $d(\Lambda)$ yields the implied constraint $\int_{\Omega} \Lambda(d\omega) = 1$ via

$$\max_{s \in \mathbb{R}} s \left(1 - \sum_{\omega \in \Omega} \Lambda(\omega) \right),$$

which is equal to infinity unless the equality $\sum_{\omega \in \Omega} \Lambda(\omega) = 1$ holds. Thus, we have the dual problem

$$\min_{\Lambda \in \mathbb{R}_+^{|\Omega|}} \left\{ \max_{x \in \mathcal{X}_f} \sum_{\omega \in \Omega} F(\mathbf{x} | \omega) \Lambda(\omega) : \sum_{\omega \in \Omega} \Lambda(\omega) = 1 \right\}.$$

We emphasize that the dual decision variables are prior distributions on the set of types. Notice that for any fixed $\Lambda \in \mathcal{P}(\Omega)$, we see that we have a Bayesian problem since we can treat Λ as a prior distribution. For Λ , we can then perform Bayesian estimation as usual. Thus, we see that the dual problem is a search for the “best” prior distribution. As a corollary, we reason that standard Bayesian estimation gives us an upper bound on the optimal value to Problem (3).

Corollary 1. (i) $\max_{x \in \mathcal{X}_f} \min_{\omega \in \Omega} F(\mathbf{x} | \omega) \leq \min_{\Lambda \in \mathcal{P}(\Omega)} d(\Lambda)$.

(ii) Let $\Lambda \in \mathcal{P}(\Omega)$ be any prior distribution, then $\max_{x \in \mathcal{X}_f} \min_{\omega \in \Omega} F(\mathbf{x} | \omega) \leq d(\Lambda)$.

Proof. Follows from weak duality for Problem (4),

$$\max_{x \in \mathcal{X}_f, s \in \mathbb{R}} \{s : s \leq F(\mathbf{x} | \omega), \forall \omega \in \Omega\} \leq \min_{\Lambda \in \mathcal{P}(\Omega)} \max_{x \in \mathcal{X}_f} \sum_{\omega \in \Omega} F(\mathbf{x} | \omega) \Lambda(\omega)$$

which gives

$$\max_{x \in \mathcal{X}_f} \min_{\omega \in \Omega} F(\mathbf{x} | \omega) \leq \min_{\Lambda \in \mathcal{P}(\Omega)} \max_{x \in \mathcal{X}_f} \sum_{\omega \in \Omega} F(\mathbf{x} | \omega) \Lambda(\omega)$$

since

$$\max_{x \in \mathcal{X}_f, s \in \mathbb{R}} \{s : s \leq F(\mathbf{x} | \omega), \forall \omega \in \Omega\} = \max_{x \in \mathcal{X}_f} \min_{\omega \in \Omega} F(\mathbf{x} | \omega).$$

7 Evaluation

In this section, we evaluate MIDAS in the fishery protection domain, where the USCG must patrol the Gulf of Mexico to prevent Mexican fishermen (Lanchas) from entering the United States Exclusive Economic Zone (EEZ) and fishing illegally. The zero-sum Stackelberg game we consider is played on a square grid, where each grid cell is a potential target. The defender (USCG) commits to a mixed strategy over fixed length patrols, where each target can be visited at most once. Additionally, all patrols must start and end in the first row of the grid. Meanwhile, the Lanchas select their mixed strategies over targets based on the SUQR behavioral model where each adversary has a unique weight vector ω .

For our experiments, the game payoffs are randomly generated with R_t^a uniformly distributed in $[1,10]$ and P_t^d uniformly distributed in $[-10,-1]$. The remaining game payoffs, R_t^d and P_t^a , are fixed at 10 and -10, respectively. Note that R_t^a and P_t^a are the same for all adversaries. All the adversary types $\omega \in \Omega$ used in the experiments were learned from USCG data. The default settings for each experiment are: five piecewise linear segments, a set of ten adversary types (i.e., $|\Omega| = 10$), and a patrol length equal to half the number of targets rounded down (i.e. $\lfloor \frac{|T|}{2} \rfloor$). We varied the dimensions of the square grid from 5×5 to 8×8 and created thirty randomly generated game instances for each grid size.

7.1 Linear Approximation

In MIDAS, we use a linear approximation to estimate the nonlinear SUQR behavioral model. The classic tradeoff when using approximation techniques is between solution quality and runtime. Thus, it is important to understand how the granularity of the approximation affects the performance of MIDAS. Figure 1(a) shows how varying the number of segments (5, 10, and 20) used in the linear approximations impacts the defender’s utility. The x -axis indicates the size of the grid, while the y -axis is the *maximin* utility obtained by the defender mixed strategy computed by MIDAS. For all grid sizes, we observe that increasing the number of segments results in higher utility for the defender as we would expect. In particular, going from 5 to 10 segments has a significant impact on the defender utility, whereas going from 10 to 20 segments produces diminishing returns and a much smaller improvement.

The other half of the tradeoff is how the number of segments impacts the runtime of MIDAS. Increasing the number of segments increases the number of variables and constraints in MIDAS, leading to a larger optimization problem which presumably would take longer to solve. The results from varying the number of segments used in the linear approximation are shown in Figure 1(b). The x -axis again indicates the size of the grid, while the y -axis is now the runtime of MIDAS in seconds. For grid sizes 5×5 through 7×7 , we see that the runtime increases as the number of segments is increased. However, for the 8×8 grid, MIDAS actually runs faster for 10 and 20 segments than it does with 5 segments. One possible explanation is that while each iteration of MIDAS algorithm takes longer to compute with more segments, the quality of the cuts generated by the separation oracle improves as the feasible marginal space is represented with higher granularity. Closer examination of the data for the 8×8 grid suggests that this is indeed the case as MIDAS with 5 segments averages with 125 calls to the separation oracle and patrol generation slave, while 10 and 20 segments average 82 and 70, respectively.

In practice, it is up to the end user to determine the right tradeoff between approximation quality and runtime. Our numerical experiments here offer guidance in this regard.

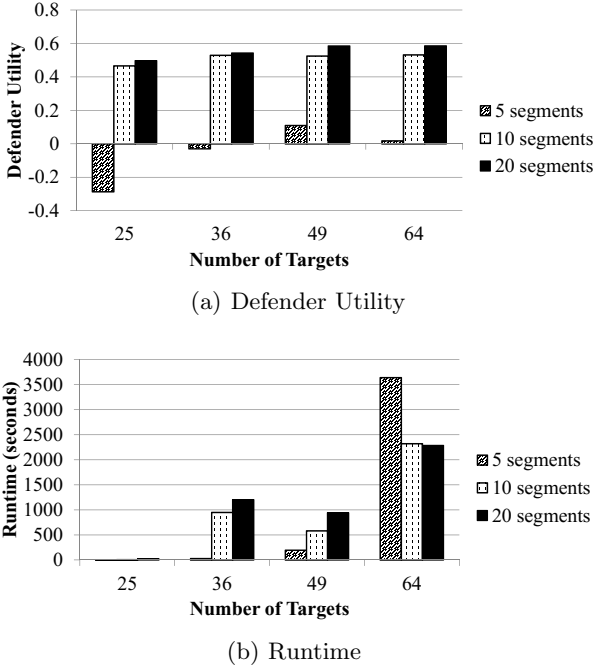




Fig. 1. Effect of the number of piecewise linear segments on MIDAS

7.2 Adversary Types

The primary purpose of MIDAS is to provide a scalable approach for generating game-theoretic patrols protecting against a set of adversaries with complex human behavior models such as SUQR. Therefore, we want to evaluate the effect of the number of adversary types on MIDAS to ensure that it serves its intended function. In Figure 2(a), we present the results for the defender *maximin* utility obtained by varying the number of adversary types on different grid sizes. Given that MIDAS computes a robust *maximin* strategy, we would expect that the defender utility monotonically decreases as the set of adversary types expands, as each additional type could present a new possible worst case for the defender. While overall this trend holds, we occasionally observe that the defender utility increases as the size of Ω is increased. One possible explanation may be the interaction between the linear approximation and the robust maximin formulation. Using 5 piecewise segments may be leading to a coarse approximation in which the monotonicity properties no longer hold. As with the number of piecewise linear segments, we would expect that increasing the number adversary types would also lead to an increase in the runtime. In Figure 2(b), we present the runtime results for MIDAS as the size of Ω is increased, which fall in line with our expectations. In particular, for the 8×8 grid we see a significant runtime increase as Ω is expanded. However, we also see that the runtimes are relatively constant for a small number of targets.



(a) Defender Utility



(b) Runtime

(b) Runtime

Fig. 2. Effect of the number of adversary types on MIDAS

7.3 Approach Comparison

Thus far, we have evaluated the performance of MIDAS as the scale of security games is increased with respect to size of the grid or the size of Ω . Now we want to compare how well MIDAS performs against other approaches that have introduced for solving security games with multiple boundedly rational adversaries. The first approach we will compare against is *Average*, in which a single adversary type ω_{avg} is constructed by averaging the weight vectors of the adversary types in Ω . After obtaining ω_{avg} , we can use MIDAS to solve the security game for $\Omega = \{\omega_{avg}\}$. The second approach we will compare against is *Marginal*, which is the robust *maximin* formulation from [5] that ignores resource assignment constraints to produce a marginal coverage distribution over the targets. To compute the *Marginal* strategy, we run MIDAS for a single iteration which produces a marginal defender strategy without considering resource assignment constraints that is then mapped into a probability distribution over patrols using the one-norm projection. The third approach is *Robust* which involves running the MIDAS algorithm to completion.

In Figure 3(a), we compare the worst case defender utility of the three approaches against sets of varying numbers of boundedly rational adversaries. The x-axis shows the number of adversary types in Ω , while the y-axis indicates the worst case defender utility of the strategies computed by the different approaches

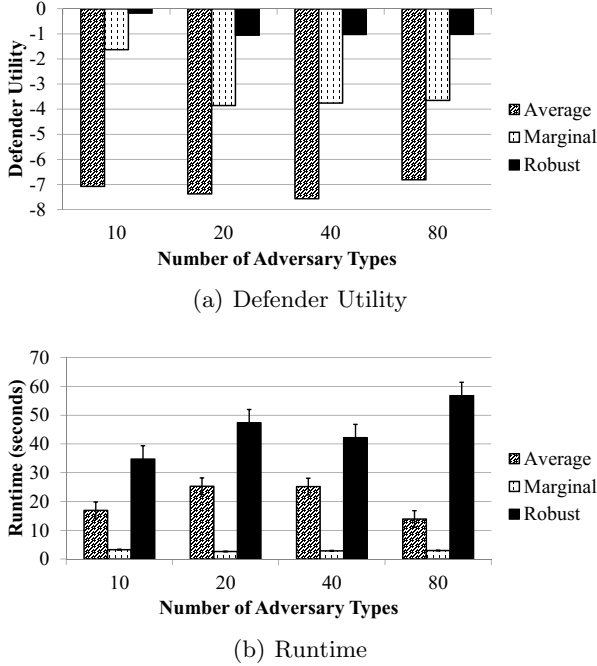


Fig. 3. Comparison of three approaches for handling multiple adversary types

against Ω . Perhaps unsurprisingly, the *Average* approach performs the worst out of the three across all sizes of Ω . The defender is optimizing against an artificially constructed adversary type ω_{avg} that is not in the set Ω . By not considering the extreme points in Ω , the resulting defender’s strategy is highly susceptible to being exploited by at least one adversary type which would define the worst case defender utility. The *Marginal* approach shows improvement by being robust against all the types in Ω , even while it initially ignores the resource assignment constraints. Finally, *Robust* uses MIDAS to its full potential and shows additional benefit of considering resource assignment constraints by outperforming *Marginal* for all sizes of $|\Omega|$.

In addition to defender utility, runtime can provide another point of comparison between the three approaches, which we analyze in Figure 3(b). Here the x-axis again indicates the number of adversary types in Ω , while the y-axis is now the runtime needed to generate the defender’s strategy using each approach. One would expect that *Average*, considering one adversary type, would run faster than *Robust*, considering $|\Omega|$ adversary types. By considering more types, the defender’s optimization becomes larger with more variables and constraints. Indeed, we observe that *Robust* takes longer than *Average* for all sizes of Ω . The gap between the two approaches seems to grow as the number of adversaries is increased, particularly for $|\Omega| = 80$. However, the runtime improvement of *Average* is likely not enough to make up for the poor solution quality in real-world

domains. Meanwhile, *Marginal* produces an essentially fixed runtime by solving only a single iteration of MIDAS and thus requires the least amount of runtime between the three approaches. Given the high stakes of real-world security domains, it is easy to imagine scenarios where security agencies would prefer the improved solution quality of *Robust* over the improved runtime of *Marginal*.

8 Conclusion

The use of bounded rationality models like QR and SUQR in security games is becoming increasingly popular in order to generate strategies that perform better against real human adversaries. These models raise two main research challenges: (i) scalability when handling resource assignment constraints and (ii) robustness when handling multiple boundedly rational adversaries. Up to this point, previous work has addressed these challenges individually. This paper addresses both scalability and robustness simultaneously by introducing a new algorithm, MIDAS. The key feature of MIDAS is the combination of incremental cut generation with a robust *minimax* formulation. Our experiments demonstrate that MIDAS can scale up to security games with complex resource allocation constraints in the form of spatio-temporal patrols. Additionally, MIDAS outperforms previous approaches for protecting against multiple adversaries by providing better solution quality guarantees in terms of worst-case performance. The overall performance of MIDAS suggests that it represents the state of the art for complex security game with boundedly rational adversaries.

Acknowledgments: This research was supported by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001 and MURI grant W911NF-11-1-0332.

References

1. Ben-Tal, A., Nemirovski, A.: Robust optimization—methodology and applications. *Mathematical Programming* 92(3), 453–480 (2002)
2. Bertsimas, D., Brown, D.B., Caramanis, C.: Theory and applications of robust optimization. *SIAM Review* 53(3), 464–501 (2011)
3. Camerer, C.: Behavioral game theory: Experiments in strategic interaction. Princeton University Press (2003)
4. Clark, A., Zhu, Q., Poovendran, R., Başar, T.: Deceptive routing in relay networks. In: Grossklags, J., Walrand, J. (eds.) *GameSec 2012*. LNCS, vol. 7638, pp. 171–185. Springer, Heidelberg (2012)
5. Haskell, W., Kar, D., Fang, F., Cheung, S., Denicola, E., Tambe, M.: Robust protection of fisheries with compass. In: *Innovative Application of Artificial Intelligence*, IAAI (2014)
6. Jain, M., Kardes, E., Kiekintveld, C., Ordóñez, F., Tambe, M.: Security games with arbitrary schedules: A branch and price approach. In: *AAAI* (2010)

7. Jiang, A.X., Nguyen, T.H., Tambe, M., Procaccia, A.D.: Monotonic maximin: A robust stackelberg solution against boundedly rational followers. In: Das, S.K., Nita-Rotaru, C., Kantarcioglu, M. (eds.) *GameSec 2013*. LNCS, vol. 8252, pp. 119–139. Springer, Heidelberg (2013)
8. Lu, W., Xu, S., Yi, X.: Optimizing active cyber defense. In: Das, S.K., Nita-Rotaru, C., Kantarcioglu, M. (eds.) *GameSec 2013*. LNCS, vol. 8252, pp. 206–225. Springer, Heidelberg (2013)
9. Manshaei, M.H., Zhu, Q., Alpcan, T., Bacşar, T., Hubaux, J.-P.: Game theory meets network security and privacy. *ACM Computing Surveys (CSUR)* 45(3), 25 (2013)
10. McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. *Games and Economic Behavior* 2, 6–38 (1995)
11. McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1), 6–38 (1995)
12. Nguyen, K.C., Alpcan, T., Basar, T.: Stochastic games for security in networks with interdependent nodes. In: *International Conference on Game Theory for Networks, GameNets 2009*, pp. 697–703. IEEE (2009)
13. Nguyen, T.H., Yang, R., Azaria, A., Kraus, S., Tambe, M.: Analyzing the effectiveness of adversary modeling in security games. In: *AAAI* (2013)
14. Píbil, R., Lisý, V., Kiekintveld, C., Bošanský, B., Pěchouček, M.: Game theoretic model of strategic honeypot selection in computer networks. In: Grossklags, J., Walrand, J. (eds.) *GameSec 2012*. LNCS, vol. 7638, pp. 201–220. Springer, Heidelberg (2012)
15. Pita, J., Jain, M., Ordonez, F., Tambe, M., Kraus, S.: Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence Journal* 174(15), 1142–1171 (2010)
16. Tambe, M.: *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, New York (2011)
17. Yang, R., Ford, B., Tambe, M., Lemieux, A.: Adaptive resource allocation for wildlife protection against illegal poachers. In: *International Conference on Autonomous Agents and Multiagent Systems, AAMAS* (2014)
18. Yang, R., Jiang, A.X., Tambe, M., Ordóñez, F.: Scaling-up security games with boundedly rational adversaries: A cutting-plane approach. In: *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pp. 404–410. AAAI Press (2013)
19. Yang, R., Ordonez, F., Tambe, M.: Computing optimal strategy against quantal response in security games. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, vol. 2, pp. 847–854. International Foundation for Autonomous Agents and Multiagent Systems (2012)