Dynamic Decisions and Adaptive Allocations: Robust Planning for Physical and Cyber Threat Screening Games

Sara Marie Mc Carthy, Phebe Vayanos, Milind Tambe
University of Southern California
{saramarm,phebe.vayanos,tambe}@usc.edu

Abstract

We consider the problem of dynamically allocating screening resources of different efficacies (e.g., X-ray imaging, deep packet inspection, cyber analysts) at checkpoints (e.g., at airports or ports, enterprise networks) to determine the threat level of the incoming screenees. Previously, the Threat Screening Game model was introduced to address this problem under the assumption that screenee arrival times are perfectly known. In reality, arrival times are uncertain, which severely impedes the implementability and performance of this approach. We thus propose a novel framework for dynamic allocation of threat screening resources that explicitly accounts for uncertainty in the screenee arrival times. We model the problem as a multistage robust optimization problem and propose a tractable solution approach using compact linear decision rules combined with robust reformulation and constraint randomization. We perform extensive numerical experiments which showcase that our approach outperforms *(a)* exact solution methods in terms of tractability, while incurring only a very minor loss in optimality, and *(b)* methods that ignore uncertainty in terms of both feasibility and optimality.

1 Introduction

Screening for threats is an important security challenge, be it inspecting cargo at ports, alerts generated by computer security systems, passengers at airports, or fans entering a stadium. Given a strategic adversary capable of exploiting gaps in security measures, along with a large number of screenees, it becomes critical to optimize the allocation of limited screening resources. In domains such as cyber security not only can the number of screenees can be overwhelmingly large, but there are also many different types of attacks that can occur, which in turn may each require a different screening method to detect.

Threat Screening Games (TSGs) have been previously introduced to model screening domains as bayesian Stackelberg games. These games model situations where the screener attempts to screen for threats, while a strategic attacker attempts to penetrate security. They have been used to model both physical and cyber threat screening; in the context of airport security [Brown *et al.*, 2016; Schlenker *et al.*, 2016] where there may be a terrorist trying to pass through airport screening. A sub-category of these games, known as Cyber Allocation Games (CAG)[Schlenker *et al.*, 2017] looks at screening suspicious activity alerts generated by and intrusion detection systems (IDS) to identify a hacker attempting to penetrate a computer network. Optimizing the defender (mixed) strategy in such games helps optimize the limited screening resources against a strategic adversary. TSGs provide a better model for strategic settings than ones that do not take strategic adversaries into account. [Tambe, 2011; Korzhyk *et al.*, 2010; Yin *et al.*, 2015; Balcan *et al.*, 2015; Basilico *et al.*, 2009; Letchford and Vorobeychik, 2011; Gan *et al.*, 2015; Guo *et al.*, 2016], where a defender protects a set of targets from a strategic adversary. However, TSGs differ significantly because they *(i)* do not have an explicitly modeled set of targets; *(ii)* include a large number of non-player screenees that must be screened while a single adversary attempts to pass through undetected; and *(iii)* encompass screening resources with differing efficacies and capacities that are combined to work in teams. These key differences make TSGs more appropriate for screening settings.

In previous work, despite promising results, previous work in TSG fails in its mission to realistically model real-world settings. Its fundamental limitation is its assumption of perfect fore-knowledge of screenee arrival times (e.g., arrival times of packets in networks, alerts generated by IDS, or passengers at aiports). However, in the real-world there is significant uncertainty in arrival times. Addressing this challenge is difficult, as it requires reasoning about all the possible realizations of the uncertainty and coming up with an optimal plan for each of those scenarios. When dealing with a large number of screenees, this result in millions of possible scenarios, making the planning problem extremely difficult.

To address this shortcoming, our first contribution is a new model *Robust Threat Screening Games (RTSG)*, which expresses the required uncertainty in screenee arrival times. In RTSG, we model the problem faced by a screener as a robust multistage optimization problem. We present a tractable solution approach with three key novelties that contribute to its efficiency: *(i)* compact linear decision rules; *(ii)* robust re-

formulation; and *(iii)* constraint randomization. We present extensive empirical results that show that our approach outperforms the original TSG methods that ignore uncertainty, and the exact solution methods that account for uncertainty. While dealing with uncertainty has been previously addressed in security games [Kiekintveld *et al.*, 2011; Yin *et al.*, 2011; Kiekintveld *et al.*, 2013] these novel techniques for handling uncertainty have not been previously explored and thus provide additional contribution not only to TSG, but to the general security game literature.

2 TSG Problem Formulation

2.1 The Case when Screenee Arrivals are Known

Time Windows. We consider a finite planning horizon consisting of W time windows (periods) $\mathcal{W} := \{1, \ldots, W\}$.

Screenee Categories. During each period, a *known* number of screenees arrive, each from a *known* category $\kappa := (\rho, \phi)$, $\rho \in \mathcal{P} := \{1, \ldots, P\}$, $\phi \in \mathcal{F} := \{1, \ldots, F\}$. The first (second) component of their category, ρ (ϕ), represents the *uncontrollable* (resp. *controllable*) part of the screenee's category. Thus, each screenee can decide the controllable part of their category, however, they cannot decide the uncontrollable part of their category, which stems from their inherent characteristics. For notational convenience, we let $\mathcal{K} := \mathcal{P} \times \mathcal{F}$. We assume that each screenee knows their own category. As an example, in the context of passenger screening at airports, ρ can represent the risk category of the passenger (e.g., normal boarding versus TSA pre-check), while ϕ can represent a flight type (e.g., international with given departure time) – note that both these components are known to the passenger. In CAG's where screenees correspond to alerts, ρ may be the level of severity assigned to an alert by an IDS, while ϕ can represent a choice of server or system of origin. We let N_κ^w denote the number of screenees in category κ to arrive in time window w. Since the category and arrival time of each screenee is known, the quantities N_κ^w are perfectly known. Without loss of generality, we assume that $N_\kappa^w > 0$ for all w and κ.

Adversary Actions. One of the screenees is planning on conducting an attack using an attack method m of his choosing from the set $\mathcal{M}$. The adversary selects *a)* his attack method m, *b)* his attack window w, and *c)* the components of his category that he can control in κ, so as to cause maximum harm. We refer to such a choice as an attack (m, w, κ). In the context of airport security such an attack may be a concealed weapon, or liquid explosive. In CAG's this may refer to denial of service attacks, malware, web exploitation, or social engineering attacks.

Defender Actions. The adversary's attack can be averted by adequate screening. For this reason, the screener is operating a checkpoint comprised of T teams indexed by $t \in \mathcal{T}$ and can decide which team should screen each screenee based on their category. Each of these teams consists of various resource types. The set of all available resource types is denoted by $\mathcal{R}$. The subset of resources composing team t is denoted by $\mathcal{R}(t) \subseteq \mathcal{R}$. If a screenee is assigned to team t, then he must be screened by all resource types allocated to that team. Such resources may be physical screening devices such as x-ray machines, or walk through metal detectors in airports or human analyst assigned to resolve alerts in CAGS. Unfortunately, not all screenees can be screened by the most effective resources as each resource has a capacity C_r on the number of screenees that it can process in each time window. The attack will be averted if the attack method is identified by any one of the resources screening the attacker. We let $E_{t,m}$ denote the effectiveness (ie. probability of interception) of team t at detecting attack method m, determined by the effectiveness of each resource. Assuming independence of the effectiveness of the resources that make up each team and letting $E^{\mathrm{r}}_{r,m}$ denote the probability of detecting an attack of type m using resource r, we have $E_{t,m} = 1 - \prod_{r \in \mathcal{R}(t)}(1 - E^{\mathrm{r}}_{r,m})$.

Following the (by now standard) approach in the literature, we formalize this problem as a *Threat Screening Game*, i.e., a Stackelberg game in which the screener, as the leader, commits to mixed strategies, and the attacker acts as the follower [Brown *et al.*, 2016; Schlenker *et al.*, 2016]. The rationale is that the screener acts first by selecting a (randomized) screening strategy, i.e., a feasible assignment of screenees to teams. In response to the choice of screening strategy, the attacker (after observing the screenee allocation) selects an attack (m, w, κ). If the attack is caught, the screener receives a utility U_κ^+, which depends the category of the adversary. Accordingly, if the screener is unsuccessful at preventing the attack, he receives the (negative) utility U_κ^-. The attacker's utilities are assumed to be negative of the screener's utilities, so that the game is zero-sum. We assume that the defender knows the *probability* that the attacker's uncontrollable category is ρ, denoted by P_ρ and we have $\sum_{\rho \in \mathcal{P}} P_\rho = 1$.The objective of the screener is then to select the best randomized allocation (i.e., mixed strategy), in anticipation of the attacker's best response.

We are now ready to provide a mathematical formulation of the TSG problem in the spirit of [Brown *et al.*, 2016].

Defender Pure Strategy Set. An assignment of screenees to teams occurs at the beginning of each period $w \in \mathcal{W}$, and corresponds to a decision on the number of screenees from each category κ to allocate to each team t out of the N_κ^w screenees that arrive in that time window. Letting $\nu_{\kappa,t}^w$ denote this assignment, the defender pure strategy set is given by

$$\mathcal{S} := \left\{ \nu \,:\, \nu_{\kappa,t}^w \in \mathbb{N}_+ \ \forall t \in \mathcal{T},\ \sum_{t \in \mathcal{T}} \nu_{\kappa,t}^w = N_\kappa^w \ \forall \kappa \in \mathcal{K},\ \sum_{t : r \in \mathcal{R}(t)} \sum_{\kappa \in \mathcal{K}} \nu_{\kappa,t}^w \le C_r \ \forall r \in \mathcal{R}, w \in \mathcal{W} \right\}.$$

The first constraint in the set stipulates that the number of screenees must be a non-negative integer. The second ensures that all the screenees are allocated to a team. The third guarantees that resource capacities are not exceeded. Note that $\mathcal{S}$ has finite cardinality, i.e., there are finitely many pure strategies available to the screener. The probability of detecting an

attack (m, w, κ) given defender strategy s is given by

$$D^{w,s}_{\kappa,m} := \sum_{t\in\mathcal{T}} E_{t,m}\nu^{w,s}_{\kappa,t}/N^w_\kappa,$$

where $\nu^{w,s}_{\kappa,t}$ denotes the number of screenees in category κ screened by team t in window w according to pure strategy s.

Defender Mixed Strategies. A mixed strategy corresponds to a distribution over pure strategies, i.e., to a choice

$$q \in \mathcal{Q} := \left\{ (q_s)_{s\in\mathcal{S}} : \sum_{s\in\mathcal{S}} q_s = 1,\ q_s \geq 0 \right\}.$$

The probability of detecting an attack (m, w, κ) is given by $\sum_{s\in\mathcal{S}} q_s D^{w,s}_{\kappa,m}$.

Robust Linear Programming Formulation. Since the attacker can select his attack (m, w, κ), but cannot select the uncontrollable aspect of his category, the problem faced by the screener is expressible as the following robust optimization problem in variables z and q

$$\begin{aligned} \text{maximize} \ & \min_{w,m,\phi} \sum_{\rho\in\mathcal{P}} P_\rho[z^w_{\kappa,m}U^+_\kappa + (1 - z^w_{\kappa,m})U^-_\kappa] \\ \text{subject to} \ & z^w_{\kappa,m} = \sum_{s\in\mathcal{S}} q_s D^{w,s}_{\kappa,m} \qquad \forall\kappa,\ m,\ w \\ & q \in \mathcal{Q}. \end{aligned} \tag{1}$$

We have omitted the sets of the variables κ, m, w and ϕ to minimize notational overhead. The variable $z^w_{\kappa,m}$ is the probability of detecting an attack (m, w, κ). Accordingly, the objective function corresponds to the worst-case expected utility of the screener. The expectation is taken with respect to the uncontrollable component of the attacker's category. The minimum is taken across all choices available to the attacker.

The cardinality of the strategy set $\mathcal{S}$ (and accordingly the number of decision variables in Problem (1)) is exponential in the number of time windows and Problem (1) is $\mathcal{NP}$-hard [Brown *et al.*, 2016]. We thus consider a relaxation to Problem (1) obtained by performing the change of variables $\pi^w_{\kappa,t} := \sum_{s\in\mathcal{S}} q_s n^{w,s}_{\kappa,t}/N^w_\kappa$. The variable $\pi^w_{\kappa,t}$ can be interpreted as the (marginal) probability of allocating a screenee in category κ to team t in window w. We obtain the following robust linear problem in variables z and π whose size is polynomial in the number of time windows

$$\begin{aligned} \text{maximize} \ & \min_{w,m.\phi} \sum_{\rho\in\mathcal{P}} P_\rho[z^w_{\kappa,m}U^+_\kappa + (1 - z^w_{\kappa,m})U^-_\kappa] \\ \text{subject to} \ & z^w_{\kappa,m} = \sum_{t\in\mathcal{T}} E_{t,m}\pi^w_{\kappa,t} \qquad \forall\kappa, m, w \\ & \pi \in \mathit{\Pi}. \end{aligned} \tag{2}$$

The first constraint is a direct consequence of the first constraint in Problem (1) combined with the change of variables, and

$$\mathit{\Pi} := \left\{ \pi \ : \ \begin{array}{ll} \displaystyle\sum_{t:r\in\mathcal{R}(t)} \sum_{\kappa\in\mathcal{K}} \pi^w_{\kappa,t} N^w_\kappa \leq C_r & \forall r, w \\ \displaystyle\sum_{t\in\mathcal{T}} \pi^w_{\kappa,t} = 1 & \\ 0 \leq \pi^w_{\kappa,t} \leq 1 \quad \forall t & \end{array} \right\} \quad \forall w, \kappa. \right\}$$

denotes the set of all marginal strategies. We note that Problem (2) is equivalent to a moderately sized linear program obtained by linearizing the piecewise linear concave objective function using the standard epigraph reformulation approach.

2.2 RTSG: The Case of Uncertain Screenee Arrivals

Insofar, we have assumed that screenee arrival times are perfectly known. Unfortunately, this assumption fails to hold in most threat screening problems. Moreover, ignoring uncertainty in the screenee arrivals during optimization may yield severely suboptimal or even infeasible allocations, see Section 4. We thus develop RTSG (Robust Threat Screening Game), a novel modeling and solution framework for threat screening that is robust to uncertainty in screenee arrival times. Our framework builds upon formulation (2) which enjoys better tractability properties than Problem (1).

Model of Uncertainty. We model the number of screenees from each category to arrive in each time window as random variables that are defined on the probability space $(\Xi, \mathcal{F}, \mathbb{P})$, which consists of the sample space Ξ, the Borel σ-algebra $\mathcal{F}$ and the probability measure $\mathbb{P}$. The elements of the sample space are denoted by $\xi := (\xi_0, \xi_1, \ldots, \xi_W)$ where the subvector $\xi_w := (\xi_{w,\kappa})_{\kappa\in\mathcal{K}}$ is observed at the end of period w and $\xi_{w,\kappa}$ represents the number of people from category κ that arrive in window w. We also let $\xi^w := (\xi_0, \ldots, \xi_w)$ denote the portion of ξ that has been observed by the end of time window w. We assume that Ξ is a bounded set expressible as

$$\Xi := \{\xi \ : \ \xi_{w,k} \in \mathbb{N},\ V\xi \leq h\} \tag{3}$$

for some matrix $V \in \mathbb{R}^{\ell\times WK}$ and vector $h \in \mathbb{R}^\ell$, where ℓ corresponds to the number of constraints in the uncertainty set. Thus Ξ corresponds to the intersection of the set of all non-negative integers with a polyhedral set. Without loss of generality, we assume that $\Xi \subset \{\xi : \xi_0 = 1\}$ (since $w = 0$ is not a valid time period, we let ξ_0 be a constant, so that affine functions of $(\xi_w)_{w\in\mathcal{W}}$ can be represented compactly as linear functions of ξ). We assume that Ξ is bounded. In the spirit of robust optimization, we refer to Ξ as the *uncertainty set*. Polyhedral uncertainty sets allow for a lot of modeling flexibility and enable us to capture a wide variety of constraints of practical relevance. In particular we can model the uncertainty present in security screening at airports using such a set.

Example 1 (Airport Screening). In the context of security screening at airports, the total number of people to travel in category κ on a given day, denoted by N_κ is known from the flight manifests. At the same time, passenger arrival times are conditioned by the time of their flight category ϕ. It is thus natural to assume that all passengers in category κ will arrive in some window $w \in \mathit{\Delta}_\kappa \subseteq \mathcal{W}$ (covering e.g., a couple of hours before their flight time). A suitable choice of uncertainty set is then given by

$$\Xi_{\text{AS}} := \left\{ \xi \ : \ \xi_{w,k} \in \mathbb{N}_+,\ \sum_{w\in\mathit{\Delta}_\kappa} \xi_{w,\kappa} = N_\kappa\ \forall\kappa \right\},$$

which we denote by AS for Airport Screening.

In the context of cyber security, we may only have estimates N_κ of the total number of alerts of each type κ and may not require that all the estimated screenee's or alerts arrive. We can additionally encode dependencies among alert types, so that certain alerts or attacks may only proceed from specific observed sequences of alerts.

In this paper, we take the view of a risk-averse screener that wishes to be immunized against all possible realizations of $\xi \in \Xi$. This view point is very natural for the set of applications under consideration that fall under the realm of security. This implies that the attacker can in some sense "strategize with nature" to devise a maximally harmful attack. Equivalently, it can be interpreted as the desire to be immunized against an attacker who would, by his own fortune, select the maximally harmful attack relative to uncertainty in arrivals.

Adaptive Screening. As information about screenee arrivals is revealed sequentially over time, the screener has the opportunity to adjust his screening policy in an adaptive fashion, at the beginning of each time window, in response to these observations. In particular, at the beginning of time window w, the screenee has observed the sequence of past arrivals ξ^{w-1} and can use that information to reason about uncertainty in remaining time windows and adjust his screening strategy accordingly. Mathematically, the screening decisions made at the beginning of time window w (i.e., π_w) in Problem (2) must be modeled as functions of the history of screenee arrivals ξ^{w-1}. Given a realization $\tilde{\xi}^{w-1}$ of ξ^{w-1}, the screener will allocate $\pi^w_{\kappa,t}(\tilde{\xi}^{w-1})$ percent of screenees of category κ to team t in window w. Accordingly, the probability of intercepting an attacker from category κ using attack method m in time window w (i.e., $z^w_{\kappa,m}$) also depends on the realization of ξ^{w-1} and must be modeled as a function of the history of observations, i.e., we have $z^w_{\kappa,m}(\xi^{w-1})$.

Resource Overflow. When arrivals are uncertain, the resource capacity constraint in (2) reads

$$\sum_{t:r\in\mathcal{R}(t)} \sum_{\kappa\in\mathcal{K}} \pi^w_{\kappa,t}(\xi^{w-1})\xi_{w,\kappa} \leq C_r \ \forall r \in \mathcal{R}, w \in \mathcal{W}, \xi \in \Xi.$$

It requires that *for all possible realizations of screenee arrivals*, the allocation must be such that all screenees be screened by available resources in the window in which they arrive. This may lead to highly conservative strategies that allocate most (if not all) screenees to the team with the highest capacity. To mitigate such over-conservatism, we propose to allow each resource $r \in \mathcal{R}$ to *overflow* from one time window to the next at a cost F_r per screenee that is delayed. Thus, each screenee is allocated to a team in the window in which they arrive. However, screening by some (or all) of the resources in that team may take place in a future time window if that resource is over-capacity. The higher the overflow fine F_r, the least likely that resource r will be overcapacity. We note that similarly to the screening policy, the number of screenees to overflow in each resource from time window w to time window $w+1$, denoted o^{w+1}_r, must be modeled as functions of the history of screenee arrivals, ξ^w. Under these considerations, the resource capacity constraint becomes

$$\sum_{t:r\in\mathcal{R}(t)} \sum_{\kappa\in\mathcal{K}} \pi^w_{\kappa,t}(\xi^{w-1})\xi_{w,\kappa} \leq C_r - o^w_r(\xi^{w-1}) + o^{w+1}_r(\xi^w) \tag{4}$$

and is enforced for all $r \in \mathcal{R}$, $w \in \mathcal{W}$, and $\xi \in \Xi$.

Adaptive Robust Optimization Formulation. We now formulate the screener's problem as a multi-stage robust optimization problem. We note that if the attacker chooses category κ and time window w for his attack, at least one screenee in category κ (corresponding to the attacker) must arrive in that time window, i.e., it must hold that $\xi_{w,\kappa} > 0$. The screener's problem may be formulated in epigraph form as

$$\begin{aligned}
&\text{maximize } \theta \\
&\text{subject to } \theta \leq \sum_{\rho\in\mathcal{P}} P_\rho u_\rho - \sum_{w\in\mathcal{W}} \sum_{r\in\mathcal{R}} F_r o^w_r \quad \forall \xi \\
&\qquad u_\rho \leq z^w_{\kappa,m} U^+_\kappa + (1 - z^w_{\kappa,m}) U^-_\kappa \quad \forall \xi : \xi_{w,\kappa} > 0 \\
&\qquad z^w_{\kappa,m} = \sum_{t\in\mathcal{T}} E_{t,m} \pi^w_{\kappa,t} \qquad\qquad \forall \xi, \kappa, m, w \\
&\qquad \pi \in \mathit{\Pi}_{\mathrm{o}}.
\end{aligned} \tag{$\mathscr{P}$}$$

The decision variables of Problem ($\mathscr{P}$) are $\theta \in \mathbb{R}$, $u_\rho(\xi)$, $o^w_r(\xi^{w-1})$, $z^w_{\kappa,m}(\xi^{w-1})$, $\pi^w_{\kappa,t}(\xi^{w-1}) \in \mathbb{R}$, and

$$\mathit{\Pi}_{\mathrm{o}} := \left\{ \begin{array}{l} \exists o \text{ with } o^w_r \geq 0 : \text{Constraint (4) } \forall \xi, r, w \\ \left\{ \begin{array}{ll} \pi & : \sum_{t\in\mathcal{T}} \pi^w_{\kappa,t} = 1 \\ & 0 \leq \pi^w_{\kappa,t} \leq 1 \quad \forall t \end{array} \right\} \quad \forall \xi, w, \kappa \end{array} \right\}.$$

We omit the dependence on ξ to minimize notational overheard. The variables $u_\rho(\xi)$ express the utility of the screener in scenario ξ when the uncontrollable category of the screener is ρ. The remaining variables admit the same interpretation as in Section 2.1. In the present setting they are however *adaptive*. The first set of constraints is used to linearize the piecewise linear concave objective function. The second set of constraints determines the worst-case value of $u_\rho(\xi)$ for each scenario ξ. For any given choice of (ϕ, w) by the attacker, this constraint is only enforced over those $\xi \in \Xi$ for which $\xi_{w,\kappa} > 0$ since at least one screenee must arrive in the attacker's chosen category and attack window. The following Proposition establishes correctness of the above formulation by showing equivalence of Problem ($\mathscr{P}$) and an appropriately constructed robust dynamic program.

Proposition 1. *The multi-stage robust optimization problem ($\mathscr{P}$) computes the optimal defender screening strategy, which maximizes his worst-case expected utility when screenee arrivals are uncertain. It is always feasible.* [1]

[1] All proofs can be found in the appendix at: https://www.dropbox.com/s/jmwoont986wu0p2/appendix.pdf?dl=0

Complexity. Since Ξ is discrete and bounded, Problem ($\mathscr{P}$) is equivalent to a deterministic linear program obtained by enumerating all possible realizations of $\xi \in \Xi$ and imposing appropriate non-anticipativity constraints, in the spirit of scenario-based stochastic programming [Birge and Louveaux, 1997]. While the numbers of decision variables and constraints in that problem is linear in the number of scenarios, the number of scenarios (cardinality of Ξ) can grow very large. In particular for the airport security setting, we show that the number of decision variables and constraints grow exponentially with the number of categories and time windows.

Complexity of Airport Screening. Consider the uncertainty set 3 corresponding to the uncertainty in passenger arrivals. For any fixed screenee category κ, the number of possible ways in which these screenees may arrive is

$$g := \binom{N_\kappa + |\Delta_\kappa| - 1}{N_\kappa}.$$

For fixed $|\Delta_\kappa|$ this quantity is $\mathcal{O}(N_\kappa^{|\Delta_\kappa|})$; and for fixed N_κ, it is $\mathcal{O}(|\Delta_\kappa|^{N_\kappa})$. Since passenger arrivals are independent across different categories, the cardinality of Ξ is given by $g^{|\mathcal{K}|}$ and is thus exponential in the number of categories. In the context airport screening, the number of scenarios is thus exponential in the number of flight categories. In addition, both the number of flight categories and corresponding number of passengers are generally linear in the number of time windows. This implies that the size of the corresponding scenario problem is exponential in the number of time windows.

3 Proposed Solution Approach

Problem ($\mathscr{P}$) can become computationally expensive to solve for realistic size instances where the cardinality of Ξ is exponential in the number of time windows, see Example 2.2. We thus propose a solution approach that results in a tractable problem even when Ξ has exponentially many scenarios. In what follows, we describe our approach and main results. The proofs can be found in the Appendix.[2]

3.1 Linear Decision Rule Approximation

Information Aggregation. In Problem ($\mathscr{P}$), the decision variables π^w are modeled as functions of the entire vector of past arrival realizations ξ^{w-1}. As a first step to obtain a tractable problem we propose to reduce information available to the screener and only allow his screening policy to adapt to the *aggregate number of screenees* that have arrived in past windows. Thus, we model the screening policy π^w for time window w as a function of the aggregate information $\zeta_{w-1} := \{\zeta_{w-1,\kappa}\}_{\kappa\in\mathcal{K}}$, where $\zeta_{w,\kappa} := \sum_{w'=1}^{w} \xi_{w',\kappa}$. The following proposition shows that this results in a conservative approximation to the optimal screening policy, since the restricted policy lies within the space of feasible policies.

Proposition 2. *Restricting the adaptive decision variables π^w and z^w for each time window $w \in \mathcal{W}$ to be functions of the aggregate information vector ζ_{w-1} provides a lower bound on the optimal objective value of Problem ($\mathscr{P}$).*

However, even when restricting π to be functions of the aggregate arrival ζ, o^w and u_ρ are still functions of the full passenger arrival ξ^{w-1}. The overflow in time window w is a function of not only ξ_w but all ξ_i $\forall i \le w$ since $o^w \ge \sum_{i=1}^{w}(\pi^i_{\kappa,t}\xi_{i,\kappa} - C_r)$. Additionally, u_ρ depends on ζ_w for all w, which is equivalent to knowing the actual passenger arrival ξ. Since restricting o^w and u_ρ to be functions of ζ^{w-1} would result in further loss of optimality we avoid it here.

Linear Decision Rule. In Problem ($\mathscr{P}$), the decision variables of the problem are arbitrary (bounded) functions of the uncertain parameter realizations. As a second step to obtain a tractable problem, we propose to restrict the space of feasible adaptive decisions to those that exhibit affine dependence on the data in the spirit of [Ben-Tal *et al.*, 2004]. Thus, we let

$$\begin{aligned}
\pi^w_{\kappa,t}(\xi^{w-1}) &= (\pi^w_{\kappa,t})^\top \zeta_{w-1} && \forall \kappa, t, w, \xi \\
z^w_{\kappa,m}(\xi^{w-1}) &= (z^w_{\kappa,m})^\top \zeta_{w-1} && \forall \kappa, m, w, \xi \\
o^w_r(\xi^{w-1}) &= (o^w_r)^\top \xi^{w-1} && \forall r, w, \xi \\
u_\rho(\xi) &= u_\rho^\top \xi && \forall \rho, \xi
\end{aligned}$$

where the vectors $\pi^w_{\kappa,t}$, $z^w_{\kappa,m} \in \mathbb{R}^K$, $o^w_r \in \mathbb{R}^{K(w-1)}$ and $u_\rho \in \mathbb{R}^{KW}$ represent the new decision variables of the problem. Following the decision rule approximation, the number of decision variables of the problem is polynomial in the number of time windows, categories, resources, and teams. Also, it is independent of the number of scenarios. Since the linear functions lie in the space of all feasible functions the decision rule results in a conservative approximation. We denote the resulting conservative approximation by ($\mathscr{P}_1$).

Proposition 3. *Problem ($\mathscr{P}_1$) provides a lower bound on the optimal objective value of problem ($\mathscr{P}$).*

3.2 Robust Counterpart

Problem ($\mathscr{P}_1$)exhibits only a moderate number of decision variables but still a very large number of constraints. In what follows, we propose to mitigate the number of constraints by using techniques inspired from modern robust optimization [Ben-Tal *et al.*, 2004]. The key observation is that under the linear decision rule approximation, all constraints in the problem (except from (4)) are linear in ξ, thus being expressible in the form $a(x)^\top \xi \le 0\ \forall \xi \in \Xi$, for some linear function a that maps the collection of all decision rule coefficients (denoted by x) to coefficients of ξ. The following proposition enables us to reformulate these constraints in a compact fashion.

Proposition 4. *For any $y \in \mathbb{R}^k$, define:*

i) $y^\top \xi \le 0 \quad \forall \xi \in \Xi$

ii) $\exists \lambda \in \mathbb{R}^\ell$ *with* $\lambda \ge 0$, $V^\top \lambda \ge y$, *and* $h^\top \lambda \le 0$.

Then $ii)$ implies $i)$.

Applying the above result to each constraint in Problem ($\mathscr{P}_1$)(except constraint (4)), we are able to represent

[2] http://teamcore.usc.edu/papers/2017/smc17.Appendix.pdf

each of these constraints efficiently. We denote the resulting problem by $(\mathscr{P}_{\mathrm{1-rc}})$. For the *airport security setting* where Ξ is defined as set (3), then the multi-stage robust optimization problem $(\mathscr{P})$ is efficiently solvable.

Proposition 5. *Suppose that the uncertainty set is defined as in (3). Then, statements $i)$ and $ii)$ in Proposition 4 are equivalent. Moreover, Problem $(\mathscr{P})$ (equivalently $(\mathscr{P}_{\mathrm{1-rc}})$) is equivalent to a linear programming problem whose size is polynomial in the number of time windows, categories, resources, and teams.*

3.3 Constraint Randomization

Although we were able to obtain an exact tractable reformulation of Problem $(\mathscr{P})$ under the uncertainty set from Example 1, this is not the case for general uncertainty sets. Indeed, for general Ξ, Problem $(\mathscr{P}_{\mathrm{1-rc}})$still involves constraint (4) enforced over a set Ξ of potentially very large cardinality. We obtain a tractable approximation to $(\mathscr{P}_{\mathrm{1-rc}})$by replacing Ξ with subsets $\Xi^N \subset \Xi$ of cardinality N. We denote the resulting problem by $(\mathscr{P}^N_{\mathrm{1-rc}})$The following theorem shows that a randomly sampled subset Ξ^N of moderate cardinality N will lead a good approximation.

Theorem 1 ([Campi and Garatti, 2008]). *Suppose that $(\mathscr{P}_{\mathrm{1-rc}})$is feasible and accommodates n decision variables. For a prespecified violation probability $\epsilon \in (0,1)$ and confidence $\beta \in (0,1)$, define*

$$N(\epsilon,\beta) := \min\left\{ N \in \mathbb{N} \;:\; \sum_{i=0}^{n-1} \binom{i}{N} \epsilon^i (1-\epsilon)^{N-i} \leq \beta \right\}$$

Then, the probability mass of all $\xi \in \Xi$ whose associated constraints are violated by an optimal solution of $(\mathscr{P}_{\mathrm{1-rc}})$, for $N \geq N(\epsilon,\beta)$, does not exceed ϵ with confidence $1-\beta$.

The parameter ϵ describes the probability that an optimal solution to $(\mathscr{P}_{\mathrm{1-rc}})$violates the overflow constraint. A violation of the overflow constraint implies that the overflows are calculated incorrectly for some samples so that the part of the objective associated with overflow is calculated incorrectly. The theorem states that such miscalculations are rare. Moreover, the size of the resulting sampled problem is polynomial in the number of time windows, categories, resources, and teams, see [Vayanos *et al.*, 2012]. Since the number of samples required is often still large, in order to solve the resulting problem more efficiently we employ a cutting plane method, in the spirit of [Fischetti and Monaci, 2012].

4 Evaluation

We evaluate our framework on airport passenger screening problems with uncertainty set Ξ_{AS}.

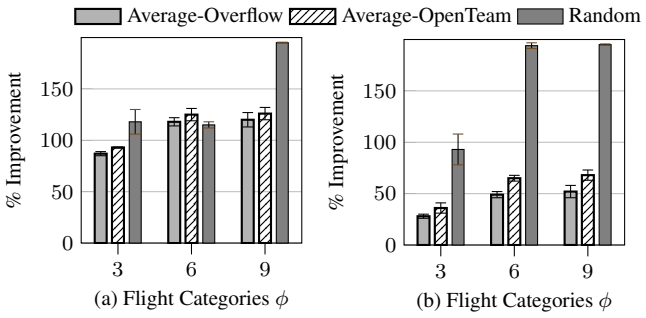

Fig. 1: Utility improvement over averaged sample and random uniform in (a) worst case and (b) average case.

4.1 Solution Quality

The optimal objective of our solution gives us the performance on the training set of samples we use. We evaluate the solution quality *out of sample* (both on average and in the worst case) by generating a large test set. We also use the test set to compute an experimental *violation probability*. We assume that the arrival of passengers is normally distributed in the range Δ_κ. Each data point is averaged over 30 trials, each with randomized parameter settings, with error bars giving the 90% confidence intervals. For each of these trials we generate 10,000 samples from the distribution of passenger arrivals, and evaluate the computed strategy on each sample so that each data point corresponds to 300,000 evaluations.

Uncertainty model vs Averaged Model. We compare our solution method to the TSG model for problems with increasing numbers of flights with $W = 10$. The TSG model optimizes against only the average ξ, so there will be many scenarios where the strategy becomes infeasible. We consider two heuristics to adjust an infeasible strategy: *(1) Overflow Heuristic:* add excess passengers to the existing overflow queue, or *(2) Open-Team Heuristic:* send excess passengers to any team with available capacity. Figure 1 summarizes our results. Against both heuristics, we outperform the TSG in worst case (average) by more than 100% (50%). The average violation probability was $98 \pm 2\%$ for the averaged sample solutions and $0.5 \pm 0.02\%$ for the solution to $(\mathscr{P}^N_{\mathrm{1-rc}})$.

Uncertainty Model vs Uniform Random. We compare to a baseline where passengers are assigned to teams uniformly at random. Figure 1 shows our results. In both the average and worst cases, the solution quality of random screening can be arbitrarily bad– we reach around 200% improvement.

Full Stochastic Program. We also compare the quality of the solution of $(\mathscr{P}^N_{\mathrm{1-rc}})$ to that of the optimal solution to the full stochastic program associated with $(\mathscr{P})$. Because the full program is exponential in the number of categories, we can only solve for very small problem instances. We fix the number of time windows, with an arrival period of 2 time windows for any flight, and show runtime and solution quality for a small range of categories. The results are shown in Table 2 where near-optimal performance is exhibited.

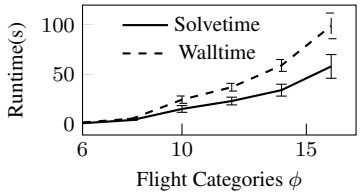

Fig. 2: Solve & Wall time with increasing number of flights.

ϕ	Violation Prob (%)	Decision Vars (10^2)
10	0.27(0.03)	40(0.5)
12	0.18(0.03)	54(0.5)
14	0.14(0.02)	70(0.6)
16	0.19(0.04)	100(2)

Table 1: Experimental violation probability with increasing problem size.

(ϕ, ρ)	% Diff Soln	Solve Time (ms)		Wall Time (ms)	
		$(\mathscr{P}^N_{\text{1-rc}})$	$(\mathscr{P})$	$(\mathscr{P}^N_{\text{1-rc}})$	$(\mathscr{P})$
(1,2)	1.3(0.1)	7.4(0.1)	13.1(0.4)	22.8(0.2)	79(9)
(1,3)	0.29(0.1)	40(1)	320(10)	110(10)	2500(230)
(1,4)	-	110(40)	-	640(10)	-
(2,1)	1.1(0.03)	2.5(0.07)	7(0.2)	10.9(0.4)	44.9(0.6)
(2,2)	0.8(0.01)	87(1)	2130(90)	260(10)	70500(90)
(2,3)	-	340(50)	-	2700(100)	-

Table 2: Comparing the $(\mathscr{P}^N_{\text{1-rc}})$ to full stochastic program $(\mathscr{P})$. Blank entries correspond to instances where the full stochastic program could not be held in memory.

4.2 Scalability

Figure 2 shows total solve and wall times for problems with increasing number of flight categories. We are able to efficiently solve for a very large number of flight categories, with polynomial growth with respect to flight categories. Table 1 summarizes our findings. We see that even for very large problems, where the cardinality of Ξ_{AS} is very large, the computed strategies have very low violation probability.

Deployment to Saturation Ratio. In Figure 3 we explore the space in which the decision problem becomes difficult by comparing the linear decision rule to a constant decision rule, where we make the same decisions regardless of the past arrival of passengers. It is a known phenomenon in security games, that the problem difficulty increases as the deployment to saturation ratio (ratio of defender resources to targets) approaches 0.5 [Jain *et al.*, 2012].

We measure the ratio by comparing the number of passengers to the capacity, for a single flight, so that the maximum number of passengers which can be screened in any time window is clearly defined. Figure 3 shows that as the problem difficulty increases, the gap in solution quality becomes large and the adaptive screening greatly outperforms the constant strategy.

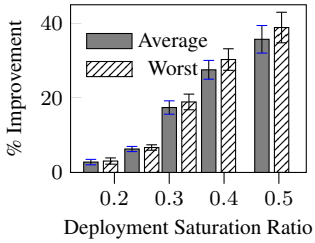

Fig. 3: Utility improvement using adaptive decision rules.

5 Conclusion and Future Work

We address a significant limitation in TSG, where the previous unrealistic assumption of complete certainty renders its solution unusable in real-world settings. We provide a novel framework which is scalable, provides good solutions quality and works for generalized models of uncertainty.

Acknowledgements

This research was supported by MURI Grant W911NF-11-1-0332

References

[Balcan *et al.*, 2015] Maria-Florina Balcan, Avrim Blum, Nika Haghtalab, and Ariel D Procaccia. Commitment without regrets: Online learning in stackelberg security games. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 61–78. ACM, 2015.

[Basilico *et al.*, 2009] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 57–64. International Foundation for Autonomous Agents and Multiagent Systems, 2009.

[Ben-Tal *et al.*, 2004] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[Birge and Louveaux, 1997] John R. Birge and Franois Louveaux. *Introduction to stochastic programming*. Springer series in operations research. Springer, New York, 1997.

[Brown *et al.*, 2016] Matthew Brown, Arunesh Sinha, Aaron Schlenker, and Milind Tambe. One size does not fit all: A game-theoretic approach for dynamically and effectively screening for threats. In *AAAI conference on Artificial Intelligence (AAAI)*, 2016.

[Campi and Garatti, 2008] Marco C Campi and Simone Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

[Fischetti and Monaci, 2012] Matteo Fischetti and Michele Monaci. Cutting plane versus compact formulations for uncertain (integer) linear programs. *Mathematical Programming Computation*, 4(3):239–273, 2012.

[Gan *et al.*, 2015] Jiarui Gan, Bo An, and Yevgeniy Vorobeychik. Security games with protection externalities. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 914–920. AAAI Press, 2015.

[Guo *et al.*, 2016] Qingyu Guo, Bo An, Yevgeniy Vorobeychik, Long Tran-Thanh, Jiarui Gan, and Chunyan Miao. Coalitional security games. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 159–167. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[Jain *et al.*, 2012] Manish Jain, Kevin Leyton-Brown, and Milind Tambe. The deployment-to-saturation ratio in security games. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1362–1370. AAAI Press, 2012.

[Kiekintveld *et al.*, 2011] Christopher Kiekintveld, Janusz Marecki, and Milind Tambe. Approximation methods for infinite bayesian stackelberg games: Modeling distributional payoff uncertainty. *In AAMAS*, 2011.

[Kiekintveld *et al.*, 2013] Christopher Kiekintveld, Towhidul Islam, and Vladik Kreinovich. Security games with interval uncertainty. In *AAMAS*, 2013.

[Korzhyk *et al.*, 2010] D. Korzhyk, V. Conitzer, and R. Parr. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proceedings of the 24th AAAI conference on Artificial Intelligence (AAAI)*, pages 805–810, 2010.

[Letchford and Vorobeychik, 2011] J. Letchford and Y. Vorobeychik. Computing randomized security strategies in networked domains. In *AARM Workshop In AAAI*, 2011.

[Schlenker *et al.*, 2016] Aaron Schlenker, Matthew Brown, Arunesh Sinha, and Milind Tambe. Get me to my gate on time: Efficiently solving general-sum bayesian threat screening games. In *European Conference on AI (ECAI)*, 2016.

[Schlenker *et al.*, 2017] Aaron Schlenker, Haifeng Xu, Mina Guirguis, Chris Kiekintveld, Arunesh Sinha Milind Tambe, Solomon Sonya, Darryl Balderas, and Noah Dunstatter. Dont bury your head inwarnings: A game-theoretic approach for intelligent allocation of cyber-security alerts. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

[Tambe, 2011] Milind Tambe. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press, 2011.

[Vayanos *et al.*, 2012] Phebe Vayanos, Daniel Kuhn, and Berç Rustem. A constraint sampling approach for multi-stage robust optimization. *Automatica*, 48(3):459–471, 2012.

[Yin *et al.*, 2011] Zhengyu Yin, Manish Jain, Milind Tambe, and Fernando Ordonez. Risk-averse strategies for security games with execution and observational uncertainty. *In AAAI*, 2011.

[Yin *et al.*, 2015] Yue Yin, Haifeng Xu, Jiarui Gain, Bo An, and Albert Xin Jiang. Computing optimal mixed strategies for security games with dynamic payoffs. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 681–687. AAAI Press, 2015.