Location and Strategies in Stackelberg Security Games with Risk Aversion

Renaud Chicoisne, Fernando Ordóñez, Daniel Castro

Location and Strategies in Stackelberg Security Games with Risk Aversion

Renaud Chicoisne, Fernando Ordóñez, Daniel Castro

September 30, 2022

Abstract

In Stackelberg security games, a leader locates security resources to protect a set of targets from strategic adversaries that aim to attack these targets after observing the leader's strategy. In this setting, the leader decision problem is to optimize an uncertain reward that can take a discrete set of values with a probability distribution that depends on the decision variable.

We show how diverse risk aversion models of the leader decision problem can be formulated as tractable optimization problems, such as imposing: a bound on the expected disutility, chance constraints, bounded distortion risk, first and second order stochastic dominance constraints, or optimizing a value-at-risk and conditional value-at-risk. We detail the resulting optimization problems and present computational results that show how the solution changes in two specific settings: 1) an entropic risk measure or value-at-risk minimization with a quantal response follower and 2) a prospect theory model with optimal follower response.

Keywords— Stackelberg Security Games, Risk aversion, Quantal Response, Convex Optimization, Mixed-Integer Programming

1 Introduction

A Stackelberg game models the strategic interaction between a leader and one or more followers, where the leader decides on a strategy to maximize its utility knowing that followers will observe this strategy when deciding their own utility maximizing action [34]. In particular, Stackelberg game models have been used in security applications to represent the interaction between defenders (that act as the leader) and attackers (corresponding to followers) [5, 6, 19]. We denote by a Stackelberg security games (SSGs) a Stackelberg game where the leader is the defender that locates security resources to protect a subset of targets that can be attacked by one or more adversaries (followers) [26, 17]. Such SSGs have been successfully deployed in real-world security applications to help locate the patrols conducted by the Los Angeles International Airport Police on the LAX airport and the US Federal Air Marshal Service on transatlantic flights [17], the LA Sheriff department on Los Angeles' subway system [12], and the US Coast Guard on the ports and waterways in Boston and New York City [1].

In an SSG both the defender and attacker receive a penalty or a reward depending whether the defender strategy locates security resources on the target attacked by the follower strategies. Therefore, the players' utility functions depend on the strategies selected by the adversaries. Assuming that players use mixed strategies, i.e. a probability distribution over possible actions, the utility of a player for a given strategy is uncertain, depending on the outcome of the combined mixed strategies. Note that this means that the uncertainty of the utility functions depends on the decision variables.

Different expressions of the uncertain utility can be considered to solve these SSG with decision variables that modify the probability distribution of the utility function. It is natural to consider that players, individually, optimize the expected value of these uncertain utility functions [25, 26, 17]. In other words, players optimize the expectation of a reward that is stochastic due to the uncertainty of the adversary's strategy. In a security setting, however, the expected reward utility function does not always provide an accurate model of player interaction, see [8]. If an expected utility model is used, the adversary response can be misrepresented which can lead to less than optimal strategies. Also, by optimizing the expected utility, the outcome of catastrophic unlikely events is not explicitly considered. Doing so can provide mixed strategy solutions that are fragile, or that have high likelihood of very bad outcomes. Both effects can be modeled with nonlinear distortion functions that transform the uncertain reward objective, such as prospect theory [18], and risk measures [2, 21].

In this work, we investigate how to efficiently formulate and solve an SSG with decision variables that influence the uncertainty distribution of the utility function. We consider a single follower and a finite set of actions for each player. In particular we focus on modeling risk-averse behavior with respect to the uncertainty due to the adversary's probability distribution over actions (i.e. its mixed strategy). We present different mathematical optimization formulations to represent chance constrained, perturbed utility functions, stochastic dominance, value at risk (VaR), and conditional value at risk (CVaR). We also present computational results for important examples that do not consider the expected reward utility function. In particular we consider Stackelberg security game models where the leader either uses an *entropic risk measure* [28] and a Quantal Response model [24], or a model that uses Prospect Theory [18]. We briefly describe these three concepts below.

An entropic risk measure amplifies the importance of outcomes that exceed a given threshold to model risk-adverse behavior against the attacker's probability over actions. The entropic risk measure of parameter $\alpha \geq 0$ of a random variable Y is defined by $\alpha \ln \mathbb{E}[e^{Y/\alpha}]$. While all outcomes are weighted, scenarios with a payoff larger than α contribute more to this measure. Therefore, the parameter α corresponds to a payoff value of risky outcomes and must be chosen carefully to tune the risk aversion level of the decision maker.

The Quantal Response (QR) Equilibrium model presented in [24] assumes that human adversaries do not behave rationally, sometimes selecting actions that do not maximize their utility. In this model, followers use a logit discrete choice model to decide between n possible actions, where action i (that gives a payoff U_i) is selected with probability:

$$\mathbb{P}\left[\text{selecting action } i\right] = \frac{1}{\sum_{j=1}^{n} e^{\lambda U_j}} e^{\lambda U_i} \ ,$$

where the parameter λ represents a *degree of rationality*, with perfect rationality ($\lambda \to \infty$) or indifference ($\lambda = 0$) as special cases. The QR model has been used to model human behavior in various settings, including economics [16, 31], game theory [36],

transportation engineering [4], marketing [15], and security applications [37].

Prospect theory [18] explicitly represents player biases, modeling risk averse and risk seeking behavior. It does so by considering perturbation functions on both the reward values and the probability distribution of possible outcomes. That is, if outcome i has a probability of occurrence p_i and payoff U_i, prospect theory proposes players perceive the following expected utility

$$V(p, U) = \sum_{i=1}^{n} \pi(p_i) V(U_i) \ .$$

Where $\pi(\cdot)$ and $V(\cdot)$ are perturbation functions with specific properties that model how players perceive both payoffs and the likelihood of occurrence. Prospect theory has contributed in economics [32], politics [23], online auctions [7], and security [37] applications.

In the next section we present the SSG problem and fix the notation. Section 3 formulates an SSG problem for different risk aversion models. Section 4 presents the algorithms for computing VaR and CVaR with an uncertainty that depends on decision variables. We present some preliminary computational results in 7 and conclude the paper in Section 8.

2 Notation and Basic Assumptions

We begin introducing the Stackelberg security game considered, which is similar to the problem in [20]. The SSG assumes there is a finite set of targets denoted by $I = \{1, \dots n\}$. The attacker decides between n actions that indicate which target to attack. One of the targets can represent the decision not to attack. The defender actions determines where to locate security resources to protect or cover a subset of targets. A defender action, or pure strategy, $z \subset I$ indicates which targets are covered simultaneously and depends on physical constraints, such as number of defender resources, capacity of defender resources or target compatibility. Let Z denote the set of feasible defender actions. The payoff of each player depends only on whether the attacked target $i \in I$ is protected by the defender action $z \in Z$, denoted by $i \in z$, or not. Given actions $i \in I$ and $z \in Z$ the reward received by the defender (by the attacker) is either a reward $\bar{R}_i$ (a penalty P_i) if $i \in z$ or a penalty $\bar{P}_i$ (a reward R_i) if $i \notin z$. Here $\bar{R}_i, R_i > 0$ and $\bar{P}_i, P_i < 0$. Therefore, under actions $i \in I$ and $z \in Z$, the utilities of the defender and attacker, respectively, are:

$$u_D(i, z) = \begin{cases} \bar{R}_i & i \in z \\ \bar{P}_i & i \notin z \end{cases} \qquad u_A(i, z) = \begin{cases} P_i & i \in z \\ R_i & i \notin z \ . \end{cases}$$

We assume players decide on mixed strategies, or probability distributions over their set of actions, denoted by $y \in \mathcal{I} = \{y \in [0,1]^n : \sum_{i=1}^{n} y_i = 1\}$ and $q \in \mathcal{Z} = \{q \in [0,1]^{|Z|} : \sum_{z \in Z} q_z = 1\}$. Since player payoff only depends on whether the attacked target is protected or not, we consider the more succinct $x \in \mathcal{X} = \{x \in [0,1]^n : x_i = \sum_{z \in Z : i \in z} q_z, \ q \in \mathcal{Z}\}$. The set $\mathcal{X}$ is the projection on $[0,1]^n$ of the feasible mixed strategies of the defender and, for $x \in \mathcal{X}$, the value x_i is the frequency with which target i is protected by a mixed strategy in $\mathcal{Z}$. The players' rewards as a function of the mixed strategies, denoted by $U_D(y, x)$ and $U_A(y, x)$ for the defender and attacker respectively, are discrete random variables. For example the defender utility equals $\bar{P}_i$ with probability $y_i(1 - x_i)$ and equals $\bar{R}_i$ with probability $y_i x_i$. If Ψ and Ψ' denote

statistics for the leader and follower utilities, we can write the problem that optimizes the leader utility as the following bilevel problem:

$$\begin{array}{ll} \max & \Psi\left(U_D(y,x)\right) \\ \text{s.t.} & x \in \mathcal{X} \\ & y = \text{argmax } \Psi'(U_A(y,x)) \\ & \qquad \text{s.a. } \ y \in \mathcal{I} \ . \end{array} \tag{1}$$

The solution to this problem determines the strong Stackelberg equilibrium of the Stackelberg game, where the follower breaks ties in favor of the leader [20].

For any mixed strategy $x \in \mathcal{X}$, we let $y(x)$ denote the follower's best response, given by the solution to the subproblem in (1). Then the leader's disutility $D(x) = -U_D(y(x), x)$ is a discrete random variable that takes the value $-\bar{R}_i$ with probability $x_i y_i(x)$ and $-\bar{P}_i$ with probability $(1 - x_i)y_i(x)$. All the possible disutilities $\{-\bar{P}_i, -\bar{R}_i\}_{i\in\{1,\ldots,n\}}$ can be referred to as $\{V_v\}_{v\in\mathcal{V}}$, with $|\mathcal{V}| = 2n$ outcomes that do not depend on the decision variables x. Without loss of generality we assume these values are sorted in increasing order: $V_1 \leq V_2 \leq ... \leq V_{2n}$. However, the probabilities of these discrete outcomes $p_v(x) := \mathbb{P}[D(x) = V_v]$ depend on x.

Different forms of the best response $y(x)$ are due to the specifics of the subproblem being solved. In the classic Stackelberg setting, the statistic for the subproblem Ψ' is the expectation, making the subproblem a linear optimization problem, which has optimal pure strategies. Non-linear statistics, such as variance or distortion functions – as in prospect theory – can generate a mixed strategy best response. A quantal response (QR) model of the follower replaces the second level problem with the assumption that a follower selects an alternative following the probability distribution

$$y_i(x) = \frac{e^{\lambda U_A(i,x)}}{\sum_{j=1}^{n} e^{\lambda U_A(j,x)}}.$$

If we assume that the utility statistic of the leader is the expected value, then $\Psi[-U_D(y(x), x)] = \mathbb{E}[D(x)] = \sum_{v\in\mathcal{V}} V_v p_v(x)$. We can then express the leader's optimization problem as

$$\min_{x\in\mathcal{X}} \sum_{v\in\mathcal{V}} V_v p_v(x) \ .$$

We show in the next Section that, under reasonable conditions, this kind of problem and generalizations of the form

$$\min_{x\in\mathcal{X}} \{f_0(x) : f(x) \leq 0\} \tag{2}$$

can be tackled efficiently. The generalization considered is able to represent different methods to handle and model the uncertainty present in the leader's utility including chance constraints, risk distortion functions, and stochastic dominance constraints.

3 Efficient Leader Problem Formulations

Here we present reformulations of (1) in the situation where, there is a known follower best response $y(x)$, the disutility function $D(x)$ takes $2n$ values that do not depend on x with probabilities that depend on x.

The formulations considered will aim to either maintain some risk measure of the disutility $D(x)$ under a given treshold – translated by some constraints $f(x) \leq 0$ – or minimize a risk measure of $D(x)$, which translates into minimizing some function $f_0(x)$. We will transform these different problem formultations to constraints over the set of decision variables $x \in \mathcal{X}$ of the form

$$\sum_{v \geq \bar{v}} p_v(x) \xi_v \leq \Xi \ , \tag{3}$$

for a real valued vector $(\xi_v)_{v \in \mathcal{V}}$ such that $\xi_1 \leq \xi_2 \leq ... \leq \xi_{|\mathcal{V}|}$, some index $\bar{v} \in \mathcal{V}$, and a right-hand side $\Xi \in \mathbb{R}$.

Notice that we can assume that $\xi_v \geq 0$ for $v \geq \bar{v}$. If this is not the case, simply define $\zeta := \max_{v \geq \bar{v}}(-\xi_v)_+$ and construct the following non-negative vector $\xi'_v = \xi_v + \zeta$ if $v \geq \bar{v}$ and $\xi'_v = \zeta$ for $v \leq \bar{v} - 1$. Then constraint (3) is equivalent to

$$\Xi + \zeta \geq \sum_{v \geq 1} p_v(x) \xi'_v \ .$$

Constraints of the form (3) are easy to solve if the dependency of x through the probability functions $p_v(x)$ form convex constraints on $\mathcal{X}$. We now show situations where enforcing bounded risk of the leader can be modeled with type (3) constraints, for different choices of $\bar{v}$, ξ and Ξ.

3.1 Maximum expected disutility

Given a reference disutility $\mathbb{E}[D(\widetilde{x})]$ coming from some known solution $\widetilde{x} \in \mathcal{X}$ we want to find some $x \in \mathcal{X}$ having an expected disutility that is no worse than the reference disutility from $\widetilde{x}$. In other words, x must satisfy the following constraint: $\mathbb{E}[D(x)] \leq \mathbb{E}[D(\widetilde{x})]$, which is by definition equivalent to the generic constraint (3) with $\Xi := \mathbb{E}[D(\widetilde{x})]$, $\xi_v := V_v$ for every $v \in \mathcal{V}$ and $\bar{v} := 1$, i.e.

$$\sum_{v \in \mathcal{V}} p_v(x) V_v \leq \mathbb{E}[D(\widetilde{x})] \ .$$

3.2 Chance constraints

Given a threshold value $\widetilde{V} \in \mathbb{R}$ and a tolerance $\epsilon \in [0, 1]$, a chance constraint [22, 9] on the disutility $D(x)$ bounds the likelihood that $D(x) \geq \widetilde{V}$ by ϵ , that means:

$$\mathbb{P}\left[D(x) \geq \widetilde{V}\right] \leq \epsilon \ . \tag{4}$$

This constraint over $x \in \mathcal{X}$ is equivalent to the generic constraint (3) taking $\Xi = \epsilon$, $\xi_v = 1$ for every $v \in \mathcal{V}$ and $\bar{v} := \arg\min_{v \in \mathcal{V}}\{V_v : V_v \geq \widetilde{V}\}$, i.e.

$$\sum_{v \geq \bar{v}} p_v(x) \leq \epsilon \ . \tag{5}$$

3.3 Bounded distortion risk

A distortion risk measure [3] is a real valued function ρ taking as argument a random variable Z that can be described as:

$$\rho : Z \to d^{-1}\left(\mathbb{E}\left[d(Z)\right]\right),$$

where $d : \mathbb{R} \to \mathbb{R}$ is an increasing bijective disutility function. The entropic risk measure $Z \to \alpha \ln \mathbb{E}[e^{Z/\alpha}]$ of parameter $\alpha > 0$ is a particular distortion risk measure. A constraint that bounds a distortion risk is a constraint over $x \in \mathcal{X}$ so that the distortion risk is less than a given threshold $\widetilde{\rho}$, i.e.

$$\rho(D(x)) \leq \widetilde{\rho} \ . \tag{6}$$

Constraint (6) is equivalent to $\mathbb{E}[d(D(x))] \leq d(\widetilde{\rho})$, i.e. $\sum_{v \in \mathcal{V}} p_v(x) d(V_v) \leq d(\widetilde{\rho})$, which is exactly the generic constraint (3) with $\Xi = d(\widetilde{\rho})$, $\xi_v = d(V_v)$ for every $v \in \mathcal{V}$, and $\bar{v} = 1$. Because d is increasing, we indeed have $\xi_1 \leq \xi_2 \leq ... \leq \xi_{|\mathcal{V}|}$.

3.4 First order stochastic dominance constraints

Let $F_Z : t \to \mathbb{P}[Z \leq t]$ denote the cumulative distribution of a random variable Z. Given two random variables Z and T, Z is said to stochastically dominate T in the first order, $Z \succeq_{(1)} T$, if $F_Z(t) \geq F_T(t)$ for all $t \in \mathbb{R}$ [14].

Given a reference random variable $D(\widetilde{x})$, we can write a constraint over $x \in \mathcal{X}$ such that $D(x)$ stochastically dominates $D(\widetilde{x})$ in the first order, i.e. $D(x) \succeq_{(1)} D(\widetilde{x})$. In our context where both random variables $D(x)$ and $D(\widetilde{x})$ have the same discrete support, this can be rewritten as follows: for every $\widetilde{v} \in \mathcal{V}$ we must have $F_{D(x)}(V_{\widetilde{v}}) \geq F_{D(\widetilde{x})}(V_{\widetilde{v}})$, i.e. $\sum_{v \leq \widetilde{v}} p_v(x) \geq \sum_{v \leq \widetilde{v}} p_v(\widetilde{x})$. In other words:

$$\sum_{v \geq \widetilde{v}+1} p_v(x) \leq 1 - \sum_{v \leq \widetilde{v}} p_v(\widetilde{x}) \qquad \forall \widetilde{v} \in \mathcal{V} \ . \tag{7}$$

The first order stochastic dominance constraint $D(x) \succeq_{(1)} D(\widetilde{x})$ can thus be represented by the $|\mathcal{V}|$ constraints in (7) which are of type (3) with $\Xi = 1 - \sum_{v \leq \widetilde{v}} p_v(\widetilde{x})$, $\xi_v = 1$ for every $v \in \mathcal{V}$, and $\bar{v} = \widetilde{v} + 1$.

3.5 Second order stochastic dominance constraints

The second order cumulative distribution function of a random variable Z is given by:

$$F_Z^{(2)}(\eta) := \int_{-\infty}^{\eta} F_Z(t) \mathrm{d}t \ .$$

Given two random variables Z and T, Z is said to stochastically dominate T in the second order, $Z \succeq_{(2)} T$, if $F_Z^{(2)}(\eta) \geq F_T^{(2)}(\eta)$ for all $\eta \in \mathbb{R}$, [14].

Given a reference random variable $D(\widetilde{x})$, we want to enforce the fact that $D(x)$ stochastically dominates $D(\widetilde{x})$ in the second order, i.e. $D(x) \succeq_{(2)} D(\widetilde{x})$. A result from [13] states that $D(x) \succeq_{(2)} D(\widetilde{x})$ is equivalent to

$$\mathbb{E}\left[(V_{\widetilde{v}} - D(x))_+\right] \geq \mathbb{E}\left[(V_{\widetilde{v}} - D(\widetilde{x}))_+\right] \qquad \forall \widetilde{v} \in \mathcal{V} \ .$$

We can rewrite this equivalently as:

$$-\sum_{v \in \mathcal{V}} p_v(x)(V_{\widetilde{v}} - V_v)_+ \leq -\sum_{v \in \mathcal{V}} p_v(\widetilde{x})(V_{\widetilde{v}} - V_v)_+ \qquad \forall \widetilde{v} \in \mathcal{V} \ . \tag{8}$$

In consequence, the second order stochastic dominance constraint $D(x) \succeq_{(2)} D(\widetilde{x})$ can be represented by the $|\mathcal{V}|$ constraints in (8) which are of type (3) with $\Xi = -\sum_{v \in \mathcal{V}} p_v(\widetilde{x})(V_{\widetilde{v}} - V_v)_+$, $\bar{v} = 1$ and $\xi_v = -(V_{\widetilde{v}} - V_v)_+$ for every $v \in \mathcal{V}$. Note that ξ_v are also in increasing order.

3.6 Some difficult risk models

We say that constraint (3) is tractable if it describes a convex set on the decision variables or can be reasonably approximated with a handful of binary variables. While the previous examples show that the risk-aversion constraints can be expressed in a tractable form, there are some examples for which it is not clear whether there is a tractable transformation or not.

For example, constraining the variance of $D(x)$ to be under a given threshold σ^2, i.e. $\mathbb{V}[D(x)] \leq \sigma^2$, boils down to

$$\sum_{v \in \mathcal{V}} p_v(x) V_v^2 - \left(\sum_{v \in \mathcal{V}} p_v(x) V_v \right)^2 \leq \sigma^2,$$

which is a complicated constraint for general probability functions $p_v(x)$. The same can be said about the upper semideviation $USD : Z \to \mathbb{E}[(Z - \mathbb{E}[Z])_+]$ where enforcing $USD(D(x)) \leq \widetilde{U}$ is equivalent to:

$$\sum_{v \in \mathcal{V}} p_v(x) \left(V_v - \sum_{v' \in \mathcal{V}} p_{v'}(x) V_{v'} \right)_+ \leq \widetilde{U}$$

Both these constraints suggest non-convex constraints on the decision variables.

Interestingly, in a classic stochastic optimization setting where the probabilities are fixed and the uncertainty is affecting the payoffs alone, modelling chance constraints or using the value-at-risk turns the resulting problem NP-hard in general, whereas in our context, chance constraints are perfectly tractable computationally.

3.7 Minimizing a risk

We now make use of all the machinery available for risk-inducing constraints in the context of minimizing risk measures. First, the generic problem (2) can be equivalently recast as:

$$\begin{aligned} \min_{x \in \mathcal{X}, \eta} \quad & \eta \\ \text{s.t.} \quad & \eta \geq f_0(x) \\ & f(x) \leq 0 \end{aligned}$$

For several of the aforementioned risk measures such as distortions in Subsection 3.2 or the probability of having a poor outcome in Subsection 3.3, f_i for any i can be of the form

$$\begin{aligned} f_0 : \quad & x \to \sum_{v \geqslant \bar{v}_0} \xi_v^0 p_v(x) \\ f_i : \quad & x \to \sum_{v \geqslant \bar{v}_i} \xi_v^i p_v(x) - \Xi_i \\ \text{with} \quad & 0 \leq \xi_1 \leq \xi_2 \leq \ldots \leq \xi_{|\mathcal{V}|} \end{aligned}$$

Following the ideas in [30, 10], we can iteratively make guesses about the optimal value η^* with a binary search: when fixing $\eta = \widetilde{\eta}$, the latter problem reduces to

investigate whether there exists $x \in \mathcal{X}$ satisfying:

$$\sum_{v \geqslant \bar{v}_i} \xi_v^i p_v(x) \leq \Xi_i, \qquad \forall i \geqslant 1$$
$$\sum_{v \geqslant \bar{v}_0} \xi_v^0 p_v(x) \leq \widetilde{\eta}.$$

In the end, considering a risk measure in the objective is not harder - modulo the binary search - than considering a constraint equivalent.

4 VaR and CVaR minimization

In this section we consider two classic risk measures: the value at risk and the conditional value at risk. It remains open if it is possible to express these risk models in a tractable form. However, we see below that it is possible to minimize them in our context.

4.1 Value at Risk

The objective in this subsection is to minimize the Value-at-Risk of parameter $\epsilon \in]0,1[$ (VaR_ϵ) of the *disutility* of the defender. The Value-at-Risk-ϵ of a disutility random variable $D(x)$ is defined as $\mathrm{VaR}_\epsilon(Z) := \inf_{t\in\mathbb{R}} \{t : F_Z(t) \geq 1-\epsilon\}$. Because $D(x)$ has a discrete and finite probability distribution, the only values $\mathrm{VaR}_\epsilon(D(x))$ can possibly take are the payoffs $(V_v)_{v\in\mathcal{V}}$. In consequence, we have that

$$\mathrm{VaR}_\epsilon(D(x)) = \min_{\widetilde{v}\in\mathcal{V}} \left\{ V_{\widetilde{v}} : \sum_{v\leq\widetilde{v}} p_v(x) \geq 1-\epsilon \right\}.$$

The problem of finding a defence strategy $x \in \mathcal{X}$ that minimizes $\mathrm{VaR}_\epsilon(D(x))$ can then be cast as follows:

$$\min_{x\in\mathcal{X},\widetilde{v}\in\mathcal{V}} \left\{ V_{\widetilde{v}} : \sum_{v\leq\widetilde{v}} p_v(x) \geq 1-\epsilon \right\}. \tag{9}$$

After rearranging the minimizations in the latter problem (9), we obtain:

$$\min_{\widetilde{v}\in\mathcal{V}} \left\{ V_{\widetilde{v}} + \min_{x\in\mathcal{X}} \left\{ 0 : \sum_{v\leq\widetilde{v}} p_v(x) \geq 1-\epsilon \right\} \right\}. \tag{10}$$

Notice that the inner problem (10) in x given $\widetilde{v} \in \mathcal{V}$ is a feasibility problem that only requires to check if there exists some $x \in \mathcal{X}$ such that $\sum_{v\leq\widetilde{v}} p_v(x) \geq 1-\epsilon$, which is equivalent to

$$\sum_{v\geq\widetilde{v}+1} p_v(x) \leq \epsilon.$$

Proposition 1. *The feasibility of the inner problem in x from Problem* (10) *can be checked as follows: the inner problem in x is feasible iff the optimal objective value $u_{\widetilde{v}}$ of the following problem is lesser than or equal to ϵ:*

$$u_{\widetilde{v}} := \min_{x\in\mathcal{X}} \sum_{v\geq\widetilde{v}+1} p_v(x). \tag{11}$$

The last problem simulates the fact that if the chosen $\widetilde{v} \in \mathcal{V}$ is associated to a value $V_{\widetilde{v}}$ that is too low to guarantee that $F_{D(x)}(V_{\widetilde{v}}) \geq 1-\epsilon$ for at least one $x \in \mathcal{X}$, then it is an underestimator of the optimal objective value of the original problem (9). In consequence, the optimal objective value must lie strictly above $V_{\widetilde{v}}$, which allows us to eliminate from the candidates for the optimal objective value all the outcomes V_v such that $v \leq \widetilde{v}$.

On another hand, if the chosen $\widetilde{v} \in \mathcal{V}$ is associated to a value $V_{\widetilde{v}}$ that guarantees that $F_{D(x)}(V_{\widetilde{v}}) \geq 1-\epsilon$ for some $x \in \mathcal{X}$, then it is either an overestimator of the optimal objective value of the original problem (9), or the optimal objective value itself. In consequence, the optimal objective value must lie at $V_{\widetilde{v}}$ or under, which allows us to eliminate from the candidates for the optimal objective value all the outcomes V_v such that $v > \widetilde{v}$.

These observations suggest a binary search scheme iteratively looking for the index $v^* \in \mathcal{V}$ that corresponds to the true optimal value V_{v^*} of Problem (9). We summarize the procedure in Algorithm 1 where the routine `solve`$(\widetilde{v})$ takes as argument an index $\widetilde{v} \in \mathcal{V}$ and returns a tuple $\left(x^{\widetilde{v}}, u_{\widetilde{v}}\right)$ corresponding respectively to an optimal solution and the optimal value of Problem (11).

Algorithm 1:

Data: An instance of problem (9)
Result: An optimal solution x^* for (9)

```
(x*, u):=solve(1);
if u ≥ 1 − ϵ then
    return x*;
(x*, u):=solve(|V|);
U := |V|; L := 1;
while U > L + 1 do
    v := ⌈(L + U)/2⌉;
    (x, u):=solve(v);
    if u ≥ 1 − ϵ then
        U := v; x* := x;
    else
        L := v;
return x*;
```

Proposition 2. *Algorithm 1 returns an optimal solution for Problem* (9) *by solving* $O(\log_2 |\mathcal{V}|)$ *times a minimization problem* (11) *with different values of* $\bar{v}$.

4.2 Conditional Value At Risk

The objective in this subsection is to minimize the Conditional Value-at-Risk of parameter $\epsilon \in]0,1[$ (CVaR_ϵ) of the *disutility* of the defender, defined as

$$\mathrm{CVaR}_\epsilon(D(x)) := \inf_{t \in \mathbb{R}} \left\{ t + \epsilon^{-1} \mathbb{E}\left[(D(x) - t)_+ \right] \right\}.$$

Furthermore, as shown in [29], the minimum in ta is attained at $t^* = \mathrm{VaR}_\epsilon\left(D(x)\right)$ so that we also have the following alternative identity:

$$\mathrm{CVaR}_\epsilon(D(x)) := \mathrm{VaR}_\epsilon\left(D(x)\right) + \epsilon^{-1}\mathbb{E}\left[\left(D(x) - \mathrm{VaR}_\epsilon\left(D(x)\right)\right)_+\right].$$

We now want to determine an $x \in \mathcal{X}$ that minimizes the Conditional Value-at-Risk of $D(x)$, which is modeled by the following optimization problem

$$\omega^* := \min_{x\in\mathcal{X}, t\in\mathbb{R}} \left\{t + \epsilon^{-1}\mathbb{E}\left[\left(D(x) - t\right)_+\right]\right\}. \tag{12}$$

Recalling that $D(x)$ follows a discrete probability distribution, we also have:

$$\omega^* := \min_{x\in\mathcal{X}, t\in\mathbb{R}} \left\{t + \epsilon^{-1}\sum_{v\in\mathcal{V}} p_v(x)\left(V_v - t\right)_+\right\}. \tag{13}$$

In the previous section, we saw that $\mathrm{VaR}_\epsilon\left(D(x)\right) \in \mathrm{supp}\left(D(x)\right) = \left\{\left(V_v\right)_{v\in\mathcal{V}}\right\}$, meaning that (13) is equivalent to

$$\omega^* := \min_{x\in\mathcal{X}, \widetilde{v}\in\mathcal{V}} \left\{V_{\widetilde{v}} + \epsilon^{-1}\sum_{v\in\mathcal{V}} p_v(x)\left(V_v - V_{\widetilde{v}}\right)_+\right\}.$$

A basic algorithm First, notice that for any optimal solution (x^*, t^*) of (13), we have

$$\omega^* := \min_{x\in\mathcal{X}} \left\{t^* + \epsilon^{-1}\sum_{v\in\mathcal{V}} p_v(x)\left(V_v - t^*\right)_+\right\}.$$

In consequence, we can "guess" the optimal value of t by fixing it to $V_{\widetilde{v}}$ for every $\widetilde{v} \in \mathcal{V}$, then solve the corresponding problem in $x \in \mathcal{X}$:

$$\omega_{\widetilde{v}} := V_{\widetilde{v}} + \epsilon^{-1}\min_{x\in\mathcal{X}}\sum_{v\in\mathcal{V}} p_v(x)\left(V_v - V_{\widetilde{v}}\right)_+.$$

Because the outcomes are sorted in increasing order, the latter can be rewritten as

$$\omega_{\widetilde{v}} = V_{\widetilde{v}} + \epsilon^{-1}\underbrace{\min_{x\in\mathcal{X}}\sum_{v\geq\widetilde{v}+1} p_v(x)\left(V_v - V_{\widetilde{v}}\right)}_{u_{\widetilde{v}}}, \tag{14}$$

whose optimal solution is denoted $x^{\widetilde{v}}$. Keeping track of the values $w_{\widetilde{v}}$, we find $\omega^* := \arg\min_{v\in\mathcal{V}} w_v$ and return x^{v^*} as an optimal solution of the original problem (13). The procedure has to solve $|\mathcal{V}| = 2n$ times problem (14). We summarize the procedure in Algorithm 2 where the routine `solve`$(\widetilde{v})$ takes as argument an index $\widetilde{v} \in \mathcal{V}$ and returns a tuple $\left(x^{\widetilde{v}}, u_{\widetilde{v}}\right)$ corresponding respectively to an optimal solution and the optimal value of (14).

An improved algorithm Notice that Algorithm 2 requires to solve $2n$ optimization problems (14) with different values of $\widetilde{v}$, whereas minimizing VaR, only $O\left(\log_2 n\right)$ problems must be solved; as opposed to the classical optimization setting (where the uncertainty affects only the outcomes and the probabibilities are constant) where minimizing VaR is NP-hard whereas minimizing CVaR can be modelled via additional linear constraints and continuous variables. We now present a way to decrease the number of problems we need to solve.

Algorithm 2: Minimize CVaR

Data: An instance of problem (13)
Result: An optimal solution x^* for (13)

$v = 1$;
$w^* = +\infty$;
while $v \neq |\mathcal{V}|$ **do**
  (x, u):=`solve`(v);
  if $w^* > V_v + \epsilon^{-1}u$ **then**
    $x^* := x$;
    $w^* := V_v + \epsilon^{-1}u$;
  $v++$;
return x^*;

Proposition 3. *The function $t \to t + \epsilon^{-1} \min_{x\in\mathcal{X}} \sum_{v\in\mathcal{V}} p_v(x)[V_v - t]_+$ is continuous and piecewise concave with breakpoints $(V_v)_{v\in\mathcal{V}}$. Unfortunately, there is no guarantee that even the same function in its discrete form - i.e. restricting its domain to $(V_v)_{v\in\mathcal{V}}$ - is convex. However, we can find a locally optimal solution for the original problem solving $O(\log_2 n)$ problems in $x \in \mathcal{X}$ with t fixed to some V_v.*

The last proposition allows us to return an upper bound that is hopefully better than just solving the problems in a sequential order. Together with the next proposition, we show how to prune values V_v without solving the problem they are associated with.

Proposition 4. *Recalling that $VaR_\epsilon(D(x)) \leq CVaR_\epsilon(D(x))$, each time we solve a problem with fixed $t = V_{\bar{v}}$, we can eliminate from the list of candidates all the V_v's lying over $w_{\bar{v}}$ as they cannot possibly produce a solution improving the current best objective value. Marking each $v \in \mathcal{V}$ when we solve its corresponding problem in x during the local minimization via binary search in v, or when we eliminate it by bounds, we can accelerate the practical convergence of the first algorithm.*

Notice that in the worst case we will solve at most $|\mathcal{V}| = 2n$ problems, which is no worse than using Algorithm 2. We summarize the procedure in Algorithm 3 where the routine `binary_search`$(\mathcal{V}^+)$ takes as argument a subset $\mathcal{V}^+ \subseteq \mathcal{V}$ of marked outcomes and returns a locally optimal solution x found by binary search with its objective value u and the outcome number v it is associated with. The routine also updates the set $\mathcal{V}^+$ with the previously nonmarked outcomes it visited during the binary search.

5 Quantal response (QR)

5.1 Defining the response probabilities $p_v(x)$

Recalling that the expected utility of the attacker when the target i is attacked is $U_i(x_i) = x_i P_i + (1 - x_i)R_i$, if the attacker is not perfectly rational and follows a QR

Algorithm 3: Improved CVaR Algorithm

Data: An instance of problem (13)
Result: An optimal solution x^* for (13)

$\mathcal{V}^+ := \emptyset$;
$w^* = +\infty$;
while $\mathcal{V}^+ \neq \mathcal{V}$ **do**
  (x, u, v):=`binary_search`$(\mathcal{V}^+)$;
  if $w^* > V_v + \epsilon^{-1}u$ **then**
    $x^* := x$;
    $w^* := V_v + \epsilon^{-1}u$;
    $\mathcal{V}^+ := \mathcal{V}^+ \cup \{v' \in \mathcal{V} : w^* \leq V_{v'}\}$
return x^*;

of rationality factor $\lambda > 0$ [24], the probability that target i is attacked is given by:

$$y_i(x) = \frac{e^{\lambda U_i(x_i)}}{\sum_{j=1}^{n} e^{\lambda U_j(x_j)}}. \tag{15}$$

Defining $R := \max_{i\in\{1,\dots,n\}} R_i$, for theoretical complexity and computational tractability purposes, it is better [10] to divide by $e^{\lambda R}$ both the numerator and denominator in (15): $y_i(x) = e^{\lambda(U_i(x_i)-R)} / \sum_{j=1}^{n} e^{\lambda\left(U_j(x_j)-R\right)}$. Defining $\beta_i := e^{\lambda(R_i-R)} \geqslant 0$, $\gamma_i := \lambda(R_i - P_i) \geqslant 0$ and $\delta_i := \bar{R}_i - \bar{P}_i \geqslant 0$ we obtain that

$$y_i(x) = \frac{\beta_i e^{-\gamma_i x_i}}{\sum_{j=1}^{n} \beta_j e^{-\gamma_j x_j}}.$$

We link the QR Stackelberg security game with the generic notation as follows: The set of payoffs is $\{(V_v)_{v\in\mathcal{V}:=\{1,\dots,2n\}}\} := \{(-\bar{P}_i)_{i\in\{1,\dots,n\}}, (-\bar{R}_i)_{i\in\{1,\dots,n\}}\}$ i.e. the set of all possible disutilities sorted in increasing order. Letting $i(v)$ being the target associated to outcome V_v - be it a penalty or a reward - the probabilities of having each outcome are as follows:

$$p_v(x) := \begin{cases} x_{i(v)}\beta_{i(v)} e^{-\gamma_{i(v)}x_{i(v)}} / \sum_{j=1}^{n} \beta_j e^{-\gamma_j x_j} & \text{If outcome } V_v \text{ is a reward} \\ \left(1 - x_{i(v)}\right)\beta_{i(v)} e^{-\gamma_{i(v)}x_{i(v)}} / \sum_{j=1}^{n} \beta_j e^{-\gamma_j x_j} & \text{If outcome } V_v \text{ is a penalty} \end{cases}$$

For convenience, let define for each target $i \in \{1, ..., n\}$ the index $v^P(i)$ (respectively $v^R(i)$) corresponding to the payoff of its penalty (respectively reward).

5.2 Efficient solution

We now see that any constraint of type (3) can be put in a tractable way in any optimization framework: In fact, they can be either piecewise linearly approximated

or in some reasonable cases, be equivalent to convex constraints. Because the adversary follows a QR, any type (3) constraint becomes

$$\sum_{i:v^P(i)\geq\bar{v}} \xi_{v^P(i)} \frac{\beta_i e^{-\gamma_i x_i}}{\sum_{j=1}^{n} \beta_j e^{-\gamma_j x_j}} (1-x_i) + \sum_{i:v^R(i)\geq\bar{v}} \xi_{v^R(i)} \frac{\beta_i e^{-\gamma_i x_i}}{\sum_{j=1}^{n} \beta_j e^{-\gamma_j x_j}} x_i \leq \Xi$$

$$\text{i.e.} \sum_{i:v^P(i)\geq\bar{v}} \xi_{v^P(i)} \beta_i e^{-\gamma_i x_i} (1-x_i) + \sum_{i:v^R(i)\geq\bar{v}} \xi_{v^R(i)} \beta_i e^{-\gamma_i x_i} x_i \leq \Xi \sum_{i=1}^{n} \beta_i e^{-\gamma_i x_i}. \tag{16}$$

Proposition 5. *The following statements hold:*

1. *The left-hand side of* (16) *is separable in the variables x_i and can be consequently piecewise linearly approximated via the use of integer variables [33].*
2. *If the vector ξ is such that $\xi_1 \leq \xi_2 \leq ... \leq \xi_{|\mathcal{V}|}$,* (16) *can be cast as a the following convex constraint:*

$$\sum_{i:v^P(i)\geq\bar{v}} \xi_{v^P(i)} \beta_i z_i - \Xi \sum_{i=1}^{n} \beta_i z_i + \sum_{i:v^P(i)\geq\bar{v}} \xi_{v^P(i)} \frac{\beta_i}{\gamma_i} z_i \ln z_i - \sum_{i:v^R(i)\geq\bar{v}} \xi_{v^R(i)} \frac{\beta_i}{\gamma_i} z_i \ln z_i \leq 0$$

after using the change of variables $x_i := -\ln(z_i)/\gamma_i$.

Proof. The first part is immediate. For the second part, after using the change of variables $x_i := -\ln(z_i)/\gamma_i$, we obtain

$$\sum_{i:v^P(i)\geq\bar{v}} \xi_{v^P(i)} \beta_i z_i - \Xi \sum_{i=1}^{n} \beta_i z_i + \sum_{i:v^P(i)\geq\bar{v}} \xi_{v^P(i)} \frac{\beta_i}{\gamma_i} z_i \ln z_i - \sum_{i:v^R(i)\geq\bar{v}} \xi_{v^R(i)} \frac{\beta_i}{\gamma_i} z_i \ln z_i \leq 0. \tag{17}$$

Because of the last term in the left hand side, it is not obvious that constraints (17) define a convex set. However, if a term appears in the last sum of the left-hand side, it also appears in the penultimate term given that 1) $v^P(i) > v^R(i)$ and 2) the components ξ_v are in increasing order. Let us consider a single target i that appears in the complicating last term: its "contribution" wrt each x_i in constraint (17) is:

$$\xi_{v^P(i)} \beta_i z_i - \Xi \beta_i z_i + \xi_{v^P(i)} \frac{\beta_i}{\gamma_i} z_i \ln z_i - \xi_{v^R(i)} \frac{\beta_i}{\gamma_i} z_i \ln z_i. \tag{18}$$

The first two terms in (18) do not cause any harm to the overall convexity because of their linearity whereas the two last terms can be factored into:

$$\left(\xi_{v^P(i)} - \xi_{v^R(i)}\right) \frac{\beta_i}{\gamma_i} z_i \ln z_i. \tag{19}$$

By hypothesis we have $\xi_{v^P(i)} \geq \xi_{v^R(i)}$, making the term (19) convex. □

The last Proposition tells us that whenever the adversary follows a QR, using risk-aversion inducing constraints or objective functions is tractable in practice. More precisely, if $\mathcal{X}$ is defined by linear constraints,

1. the first result of Proposition 5 tells us that the full optimization problem can be cast as a mixed-integer linear optimization problem, and
2. the second part tells us that if $\mathcal{X}$ is defined by r linear inequalities with non-negative coefficients $(a^j)\top x \leq b_j, \forall j \in \{1, ..., r\}$, then the constraints defining $\mathcal{X}$ after the change of variables $x_i := -\ln(z_i)/\gamma_i$ translate into the r following convex constraints:

$$-\sum_{i=1}^{n} \frac{a_j^i}{\gamma_i} \ln(z_i) \leq b_i, \qquad \forall j \in \{1, ..., r\}$$

which is readily solvable by off the shelf interior point algorithms (e.g. [27, 35])

6 Prospect Theory

Another case in which problem (1) can lead to a tractable solution problem has to do with approximate solutions of the SSG when both leader and follower consider a Prospect Theory decision model. As mentioned in Section 1 prospect theory assumes players deviate from the expected objective through distortion functions that modify the valuations and the probability of occurrence. Under this model the disutility of the leader is

$$PT(D(x)) = \sum_{v \in \mathcal{V}} \pi(p_v)\nu(V_v) \ ,$$

where, given parameters $\lambda, \alpha, \beta \geq 0$ and $\delta \in [0, 1]$, the distortion functions are

$$\nu(z) = \begin{cases} (z - C)^\alpha & \text{if } z \geq C \\ -\lambda(-z + C)^\beta & \text{if } z < C \end{cases} \qquad \text{and} \qquad \pi(x) = \frac{x^\delta}{(x^\delta + (1 - x)^\delta)^{\frac{1}{\delta}}} \ .$$

Where C represents the reference point for the valuation function. Similar expressions are used for the follower for its own distortion functions π' and ν'.

To propose a tractable model we assume that the follower selects an optimal pure strategy (a target to attack), i.e. $y \in \{0, 1\}^n$. An assumption that holds for a linear objective of the subproblem. Since p_v is either $y_{i(v)}x_{i(v)}$ or $y_{i(v)}(1 - x_{i(v)})$, we have that $\pi(p_v)$ is either $y_{i(v)}\pi(x_{i(v)})$ or $y_{i(v)}\pi(1 - x_{i(v)})$. Therefore we express (1) as:

$$\begin{aligned} \max \quad & \sum_{i=1}^{n} y_i \left[\pi(x_i)\nu(\bar{R}_i) + \pi(1 - x_i)\nu(\bar{P}_i)\right] \\ \text{s.t.} \quad & x \in \mathcal{X} \\ & y = \operatorname{argmax} \sum_{i=1}^{n} y_i \left[\pi'(x_i)\nu'(P_i) + \pi'(1 - x_i)\nu'(R_i)\right] \\ & \qquad \text{s.t.} \quad \sum_{i=1}^{n} y_i = 1, \quad y \in \{0, 1\}^n \ . \end{aligned}$$

This problem with non-linear objectives in both problems can be solved approximatedly with piecewise linear approximations and integer variables. For this consider that every x_i variable is partitioned into K segments with breakpoints $c_0, c_1, \ldots, c_K$,

with $c_0 = 0$ and $c_K = 1$. The perturbation functions π and π' take values b_k and b'_k at breakpoints c_k for $k \in K$. To simplify the constraints we will use a variable $z_{iK+1} = 0$.

$$
\begin{array}{ll}
\max & \gamma \\
\text{s.t.} & x \in \mathcal{X} \\
& \sum_{i=1}^{n} y_i = 1, \quad y \in \{0,1\}^n \\
& \sum_{k \in K} z_{ik} = 1, \quad z_i \in \{0,1\}^{|K|} \qquad i \in I \\
& \sum_{k \in K} \hat{z}_{ik} = 1, \quad \hat{z}_i \in \{0,1\}^{|K|} \qquad i \in I \\
& \sum_{k \in K} w_{ik} = 1, \quad w_i \in [0,1]^{|K|} \qquad i \in I \\
& \sum_{k \in K} \hat{w}_{ik} = 1, \quad \hat{w}_i \in [0,1]^{|K|} \qquad i \in I \\
& w_{ik} \leq z_{ik} + z_{ik+1}, \quad \hat{w}_{ik} \leq \hat{z}_{ik} + \hat{z}_{ik+1} \qquad i \in I, k \in K \\
& x_i = \sum_{k \in K} c_k w_{ik}, \quad 1 - x_i = \sum_{k \in K} c_k \hat{w}_{ik} \qquad i \in I \\
& q_i = \sum_{k \in K} b_k w_{ik}, \quad \hat{q}_i = \sum_{k \in K} b_k \hat{w}_{ik} \qquad i \in I \\
& q'_i = \sum_{k \in K} b'_k w_{ik}, \quad \hat{q}'_i = \sum_{k \in K} b'_k \hat{w}_{ik} \qquad i \in I \\
& 0 \leq a - [q'_i \nu'(P_i) + \hat{q}'_i \nu'(R_i)] \leq M(1 - y_i) \qquad i \in I \\
& M(1 - y_i) + \left[q_i \nu(\bar{R}_i) + \hat{q}_i \nu(\bar{P}_i) \right] \geq \gamma \qquad i \in I
\end{array}
$$

Here variables z_{ik}, $\hat{z}_{ik}$ indicate which interval of the of the piecewise approximation is used for x_i and for $1 - x_i$ respectively. The value of the convex combination is given by variables w_{ik} and $\hat{w}_{ik}$, respectively. The values q_i, $\hat{q}_i$, q'_i, and $\hat{q}'_i$ give the expresions of $\pi(x_i)$, $\pi(1 - x_i)$, $\pi'(x_i)$, and $\pi'(1 - x_i)$, respectively. We consider that this approximate mixed integer optimization problem is a tractable model for the prospect theory approach.

7 Computational results

7.1 Expected value and Entropy minimization with QR adversaries

In [10], we studied risk-neutral and risk-averse objective models that minimize either the expected value $\mathbb{E}[D(x)]$ or an entropic risk measure $\alpha \ln \mathbb{E}[\exp(D(x)/\alpha)]$. The resulting models were able to solve instances within an hour with up to $n = 10.000$ targets for 1) a basic model where

$$\mathcal{X}_0 := \left\{ x \in [0,1]^n : \sum_{i=1}^{n} x_i \leq m \right\}$$

and 2) a more concrete model with disjunctive and precedence constraints

$$\mathcal{X}_1 := \left\{ x \in \mathcal{X}_0 : \sum_{i \in \mathcal{D}_d} x_i \leq 1, \forall d \in \{1, ..., D\}, x_i \leq x_j, \forall (i,j) \in \mathcal{E} \right\}.$$

statistic	Objective minimized $\mathcal{E}_{\alpha=1}$	$\mathcal{E}_{\alpha=2}$	$\mathcal{E}_{\alpha=5}$	$\mathcal{E}_{\alpha=7}$
$\mathbb{V}$	-32	-26	-17	-12
$\mathbb{E}$	-15	-5	-2	-1
Worst case $\mathbb{P}$	-68	-44	-14	-11
$\text{VaR}_{\epsilon=10\%}$	-9	-10	-6	-5
Exec. time (s)	6.324	5.131	4.085	3.862

Table 1: Difference of the optimal strategies in function of α as a % of the $\mathbb{E}$ solution's statistics ($\mathcal{X}_1$,$n = 1000$,$m = 100$).

All payoffs $R_i, \bar{R}_i, P_i$ and $\bar{P}_i$ belong to $[-10, 10]$ in this subsection and the next.

As we can see in Figure 1,the cumulative distributions corresponding to the risk averse strategies (i.e. minimizing entropic risk measures of parameters $\alpha = 5$ and $\alpha = 10$) are stochastically dominating the risk neutral strategies (i.e. minimizing the expected loss) in the tail of the distribution.

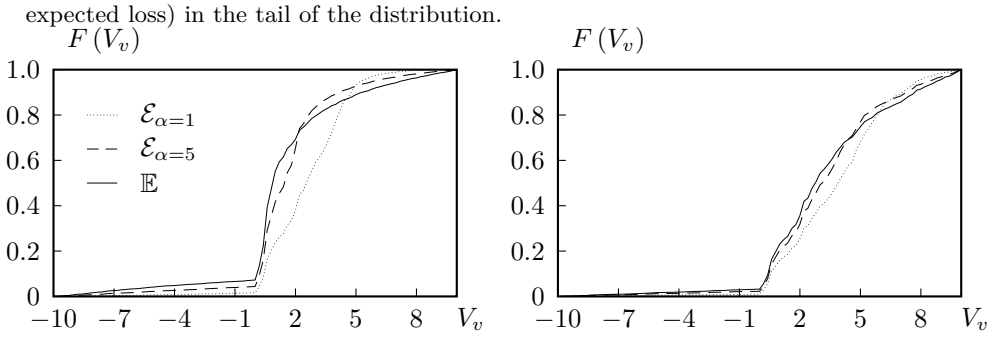

Figure 1: Loss CDFs of the minimizers of $\mathcal{E}_\alpha$ and maximizers of $\mathbb{E}$ with $\mathcal{X}_0$ (left) and $\mathcal{X}_1$ (right)

Further, we studied the influence of the entropic risk parameter α: in Table 1 we can see that as α increases - i.e. the defender becomes less risk-averse - the benefit in terms of variance, $\text{VaR}_{\epsilon=10\%}$ and the worst case probability reduction becomes less important but is significant for lower values of α; On another hand, these benefits come at the moderate cost of having an increased expected loss by about 1-15%.

7.2 VaR_ϵ and $\mathbb{P}[D(x) \geqslant \widetilde{V}]$ minimization

Some preliminar experiments were conducted on minimizing VaR_ϵ and $\mathbb{P}[D(x) \geqslant \widetilde{V}]$ with mid-sized instances of $\mathcal{X}_0$ with $n = 400$, $m = 60$. The thresholds $\widetilde{V}$ used when minimizing $\mathbb{P}[D(x) \geqslant \widetilde{V}]$ where chosen to be 100%, 50% and 20% of the worst case disutility $-V_1$, noted respectively $\mathbb{P}_{100}$, $\mathbb{P}_{50}$ and $\mathbb{P}_{20}$.

The results are summarized in Table 2 where we can see that the variance is consistently decreased by using risk measures instead of the expected value, and the more risk averse the defender is, the greater the loss in expected outcome. Because minimizing VaR involves the solution of $O(\log_2 n)$ subproblems, it is significantly slower than minimizing the probability of being over a threshold.

statistic	Objective minimized						
	$\text{VaR}_{20\%}$	$\text{VaR}_{10\%}$	$\text{VaR}_{5\%}$	$\text{VaR}_{1\%}$	$\mathbb{P}_{100}$	$\mathbb{P}_{50}$	$\mathbb{P}_{20}$
$\mathbb{V}$	-9	-18	-26	-36	+9	+1	-34
$\mathbb{E}$	+6	+17	+29	+49	+10	+2	+45
$\text{VaR}_{\epsilon=2\%}$	-10	-6	-4	+5	+18	+5	+4
$\text{CVaR}_{\epsilon=2\%}$	-4	-10	-10	-8	+8	+3	-9
Exec. time (%)	+576	+595	+606	+600	-25	-36	-31

Table 2: Statistics of the optimal strategies for different objectives as a % of the $\mathbb{E}$ solution's ($\mathcal{X}_0$,$n = 400$,$m = 60$).

7.3 Prospect Theory

Here we present computational results evaluating the change in the solution of using and not using a prospect theory model over a small random instance with $n = 8$ targets. Payoffs are generated from $[-10, 10]$. We consider seven instances with this data, changing the number of security resources that the leader uses, with $m = \{1, 2, \ldots, 7\}$. We consider a piecewise linear approximation of the probability distortion function by partitioning $[0, 1]$ in five uniformly spaced break points $K = 5$. We consider three different models, depending on which player considers a prospect theory or an expected utility objective. In particular, model *Neither* assumes both the leader and follower minimize the expected utility; model *Only Follower* has a follower with prospect theory and the leader with expected utility; and model *Both* assumes both players use a prospect theory objective.

In Table 3 we present the leader utility objective (expected utility for Neither and Only Follower) and a prospect theory objective in Both over the different instances. We observe as instance number increases (and more security resources are used) the disutility decreases for all models. In addition, notice that changing the follower utility function does not cause significant change on the leader utility. Finally, the decrease in leader utility when the leader uses prospect theory is related to the diminishing returns of the utility perturbation because $0 \leq \alpha < 1$.

Model	Instances						
	1	2	3	4	5	6	7
Neither	4.4	5.8	3.8	5.5	1.1	1.6	0.5
Only Follower	4.2	5.9	3.7	5.7	0.9	1.2	0.6
Both	1.1	1.9	1.1	1.8	0.2	0.1	0

Table 3: Leader utility objective function. *Model* identifies if objective is prospect theory or expected utility.

In Table 4 we present the change in leader expected utility as we modify the reference point C. The change is given as the difference between the leader expected utility of the Only Follower model minus the Neither model. As we change the follower reference point from -10 to 10 for all instances the leader expected utility difference is U-shaped. This difference decreases and then increases. An explanation for this is because for a reference point close to 0, the distortion of the utility value of the follower is not so large and thus does not change much from the expected utility

behavior. Largest changes are for extreme reference values in instances 3, 4 and 5. Because in this situation leader policy can be more different. In instance 1 most targets are not protected while in instance 7 most targets are protected.

Reference Point	Instances						
	1	2	3	4	5	6	7
-10	0	-0.1	1.5	4.3	3	1.8	0.7
-8	0	-0.1	1.5	3.3	2.9	1.8	0.7
-6	-0.1	-0.1	1.5	2.2	2.8	1.7	0.7
-4	-0.1	-0.1	1.5	2	2.1	1.6	0.7
-2	-0.1	0	1.9	2.9	2.7	1.8	0.7
0	-0.2	0.1	-0.1	0.2	-0.2	-0.4	0.1
2	0	0	0	0.4	0.5	0.3	0.1
4	-3	-0.7	1.4	1.8	1.3	0.8	0.2
6	-2.3	1.3	3	2.7	1.9	1.2	0.2
8	-1.9	2	3.5	3.2	2.1	1.2	0.2
10	0.4	3.2	3.8	3.4	2.1	1.2	0.2

Table 4: Expected leader utility difference (*Only Follower* − *Neither*) for different follower reference points.

8 Conclusions

The Stackelberg Security Game considered has a leader utility that results in a discrete random variable where the probability of the events depend on the players' decisions. The more common situation in optimization under uncertainty is that the decision variables influence the utility values, not the probability distribution.

We present several formulations for risk models of uncertainty, that for the leaders utility, provide convex constraints or that can be approximated efficiently with a few integer variables. These are referred to as tractable models. We show that the difficulty of computing certain statistics changes depending on whether the decision variables determine the probability or the utility. In particular VaR becomes tractable while variance seems intractable for the leader utility, a situation that is reversed when the probabilities are given and utility values depend on decision variables.

Our computational results illustrate the tractability of the approach in two situations, when the follower uses a quantal response model and when the follower responds with pure strategies which enables the use of distortion functions (prospect theory) for the leader and follower. In the former we show that we can compute different risk measures (entropic risk, VaR, and chance constraint), in the later we compare the use or not of prospect theory for different players and the effect of the reference point. In [11], some particular cases of this model are studied. Specifically, that work considers explicitly the possibility of multiple adversaries. Extending this work for the multiple adversaries setting presented in [11] is straight forward, since the derivations in Section 3 can be used in every utility function and the leader utility is the weighted sum of the interaction between the leader with each follower. Further work is necessary to evaluate the tractability of a multiple follower SSG if the utilities depend in a more complicated non-linear way of the multiple players decisions. Another line of future

research is exploring the use of these formulations in other stochastic optimization problems where the decision variables influence the probability distribution.

References

[1] B. An, F. Ordóñez, M. Tambe, E. Shieh, R. Yang, C. Baldwin, J. DiRenzo III, K. Moretti, B. Maule, and G. Meyer. A deployed quantal response-based patrol planning system for the U.S. Coast Guard. *Interfaces*, 43(5):400–420, 2013.

[2] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.

[3] A. Balbás, J. Garrido, and S. Mayoral. Properties of distortion risk measures. *Methodology and Computing in Applied Probability*, 11(3):385–399, 2009.

[4] M. Ben-Akiva and S. R. Lerman. *Discrete choice analysis: theory and application to travel demand*. Transportation Studies, 2018.

[5] V. M. Bier. Choosing what to protect. *Risk Analysis*, 27(3):607–620, 2007.

[6] G. Brown, M. Carlyle, J. Salmerón, and K. Wood. Defending critical infrastructure. *Interfaces*, 36(6):530–544, 2006.

[7] T. Brünner, J. Reiner, M. Natter, and B. Skiera. Prospect theory in a dynamic game: Theory and evidence from online pay-per-bid auctions. *Journal of Economic Behavior & Organization*, 164:215–234, 2019.

[8] C. Camerer. Behavioral economics: Reunifying psychology and economics. *Proceedings of the National Academy of Science*, 96:10575–10577, 1999.

[9] A. Charnes and W. W. Cooper. Chance-constrained programming. *Management science*, 6(1):73–79, 1959.

[10] R. Chicoisne and F. Ordóñez. Risk averse stackelberg security games with quantal response. In *Proceedings of GameSec 2016, New York*, volume LNCS 9996, pages 83–100, 2016.

[11] R. Chicoisne and F. Ordóñez. Algorithms for a risk-averse stackelberg game with multiple adversaries. *HAL archives ouvertes*, 2021.

[12] F. M. Delle Fave, A. X. Jiang, Z. Yin, C. Zhang, M. Tambe, S. Kraus, and J. P. Sullivan. Game-theoretic security patrolling with dynamic execution uncertainty and a case study on a real transit system. *Journal of Artificial Intelligence Research*, 50:321–367, 2014.

[13] D. Dentcheva and A. Ruszczynski. Optimization with stochastic dominance constraints. *SIAM J. Optim.*, 14(2):548–566, 2003.

[14] D. Dentcheva and A. Ruszczyński. Semi-infinite probabilistic optimization: first-order stochastic dominance constrain. *Optimization*, 53(5-6):583–601, 2004.

[15] D. H. Gensch and W. W. Recker. The multinomial, multiattribute logit choice model. *Journal of Marketing Research*, 16(1):124–132, 1979.

[16] P. Haile, A. Hortaçsu, and G. Kosenok. On the empirical content of quantal response equilibrium. *The American Economic Review*, 98(1):180–200, 2008.

[17] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordóñez. Software assistants for randomized patrol planning for the LAX airport police and the federal air marshal service. *Interfaces*, 40(4):276–290, 2010.

[18] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263—291, 1979.

[19] D. Kar, T. H. Nguyen, F. Fang, M. Brown, A. Sinha, M. Tambe, and A. X. Jiang. Trends and applications in stackelberg security games. *Handbook of Dynamic Game Theory*, pages 1–47, 2017.

[20] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordóñez, and M. Tambe. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the 8th AAMAS Conference, Budapest, Hungary*, pages 689–696. International Foundation for AAMAS, 2009.

[21] H. Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.

[22] J. Mayer. Computational techniques for probabilistic constrained optimization problems. In *Stochastic Optimization*, pages 141–164. Springer, 1992.

[23] R. McDermott. Prospect theory in political science: Gains and losses from the first decade. *Political Psychology*, 25(2):289—312, 2004.

[24] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

[25] R. B. Myerson. *Game theory*. Harvard university press, 2013.

[26] P. Paruchuri, J. Pearce, J. Marecki, M. Tambe, F. Ordóñez, and S. Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *Proceedings of the 7th AAMAS Conference, Estoril, Portugal*. International Foundation for AAMAS, 2008.

[27] H. Pirnay, R. Lopez-Negrete, and L. Biegler. Optimal sensitivity based on ipopt. *Mathematical Programming Computations*, 4(4):307–331, 2012.

[28] J. W. Pratt. Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society*, 32(1/2):122–136, 1964.

[29] R. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

[30] E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, and G. Meyer. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *Proc. of The 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.

[31] D. Stahl II and P. Wilson. Experimental evidence on players' models of other players. *Journal of economic behavior & organization*, 25(3):309–327, 1994.

[32] A. Tversky and D. Kahneman. Rational choice and the framing of decisions. *The Journal of Business*, 59(4):251—278, 1986.

[33] J. Vielma. Mixed integer linear programming formulation techniques. *SIAM Review*, 57:3–57, 2015.

[34] H. Von Stackelberg. *The theory of the market economy*. William Hodge, 1952.

[35] A. Wachter and L. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.

[36] J. Wright and K. Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. In *Proceedings of the 24th AAAI conference on artificial intelligence, Atlanta, GA*, 2010.

[37] R. Yang, C. Kiekintveld, F. Ordóñez, M. Tambe, and R. John. Improving resource allocation strategy against human adversaries in security games. In *22th IJCAI Proceedings, Barcelona, Spain*, volume 22, pages 458–464. AAAI Press, 2011.