Beyond Homogeneous Adversaries: Stackelberg Security Games with Mixed Quantal Response

Hoang Giang Pham[1] , **Tien Mai**[1] , **Thuy Anh Ta**[2] and **Minh Hoàng Hà**[3]
[1]School of Computing and Information Systems, Singapore Management University, Singapore
[2]School of Computing, Phenikaa University, Hanoi, Vietnam
[3]SLSCM and CADA, Faculty of Data Science and Artificial Intelligence, College of Technology, National Economics University, Hanoi, Vietnam
hg.pham.2023@phdcs.smu.edu.sg, atmai@smu.edu.sg, anh.tathuy@phenikaa-uni.edu.vn, hoanghm@neu.edu.vn

Abstract

The quantal response (QR) model is widely used in Stackelberg security games (SSGs) to capture boundedly rational adversaries. Existing work on SSGs under QR, however, almost exclusively assumes a homogeneous attacker population, ignoring heterogeneity in attacker preferences and rationality. We study SSG with *mixed quantal response* attackers, where the follower population consists of multiple discrete attacker types, each following a type-specific QR model. The defender allocates limited resources across targets, while an attacker drawn from this heterogeneous population observes the defender's strategy and attacks a single target. This results in a highly non-convex equilibrium computation problem. We develop a *polynomial-time approximation scheme* for this setting when the number of attacker types is bounded, based on an *exponential cone programming* formulation combined with a carefully designed *Branch-and-Bound* procedure. Experiments demonstrate that our approach outperforms standard gradient-based methods and that explicitly modeling attacker heterogeneity yields significant gains over traditional SSG models with a single QR attacker.

1 Introduction

Stackelberg security games provide a powerful and widely adopted framework for modeling strategic interactions between a defender with limited resources and an adversary who observes and responds to the defender's strategy [Tambe, 2011; Sinha *et al.*, 2018]. SSGs have been successfully applied in a range of high-impact real-world domains, including airport and transportation security, wildlife protection, cybersecurity, and counterterrorism, where defenders must allocate scarce resources to protect multiple potential targets against intelligent attackers [Pita *et al.*, 2008; Fang *et al.*, 2017; An *et al.*, 2012; Sinha *et al.*, 2018]. The effectiveness of these deployments critically depends on accurately modeling attacker behavior and efficiently computing optimal Stackelberg strategies.

A key modeling challenge in SSGs is capturing the bounded rationality of human adversaries. While early work often assumed fully rational attackers who best respond to the defender's strategy, empirical evidence has repeatedly shown that such assumptions are unrealistic in practice. To address this, the QR model has been widely adopted as a more realistic alternative [Yang *et al.*, 2011; Yang *et al.*, 2012; Haghtalab *et al.*, 2016; Mai and Sinha, 2023]. Under QR, attackers select actions probabilistically, with higher-utility actions chosen more frequently, thereby capturing noise, bounded rationality, and decision-making errors. QR-based models have been extensively studied in SSGs and have also played a central role in related domains such as revenue management and discrete choice modeling, where they correspond to logit-based choice behavior.

Despite this progress, most existing work on SSGs under QR assumes a homogeneous attacker population governed by a single set of preference and rationality parameters. In many real-world settings, however, adversaries are heterogeneous, exhibiting substantial variation in objectives, sensitivities, and levels of rationality. Ignoring such heterogeneity can lead to systematically suboptimal defensive strategies. Attempts to extend QR models to accommodate multiple attacker types inherently introduce significant computational challenges; the resulting equilibrium computation problems are typically highly nonconvex, and current solution approaches either rely on local optimization heuristics or lack formal performance guarantees.

In this paper, we address this important gap by studying SSGs with a heterogeneous attacker population modeled via a mixed QR framework. Specifically, we consider a setting in which the attacker is drawn from a finite set of types, each following a type-specific QR model, yielding a defender objective that is a weighted sum of multiple logit-based fractional terms. We show that this heterogeneity fundamentally alters the computational structure of the problem, leading to *non-unimodal objectives* and rendering standard gradient-based and fractional programming methods ineffective. Our main contributions are threefold: (i) we develop a ***novel bilinear exponential-cone reformulation*** that isolates all sources of nonconvexity into a small number of bilinear terms; and (ii) leveraging this structure, we design a ***Branch-and-Bound (B&B) algorithm*** that yields a ***polynomial-time approxima-***

tion scheme (PTAS) when the number of attacker types is bounded; and (iii) through extensive experiments, we demonstrate that explicitly modeling attacker heterogeneity leads to ***significant performance improvements*** over traditional homogeneous QR models and that our method reliably outperforms gradient-based baselines. Our key innovation is a problem-specific reformulation of the original nonconvex optimization problem that isolates all sources of nonconvexity into exactly T dimensions, corresponding to the number of attacker types. This structural reduction enables a B&B algorithm whose complexity scales polynomially with the population size (i.e., the number of targets, which may be very large) and exponentially only in T.

To the best of our knowledge, this is the first work to establish a PTAS for a class of SSG under a mixed QR attacker model.

2 Literature Review

Our work most closely relates to SSGs with QR adversaries and to the broader literature on discrete choice modeling and optimization under heterogeneous agents.

SSG and QR. SSGs [Tambe, 2011] are a prominent class of security models with numerous real-world deployments, including airport security, wildlife protection, and counterterrorism [Pita *et al.*, 2008; Fang *et al.*, 2017; An *et al.*, 2012; Sinha *et al.*, 2018]. A central line of research in SSGs focuses on behavioral models of attackers, among which the QR model is the most extensively studied [Yang *et al.*, 2011; Yang *et al.*, 2012; Haghtalab *et al.*, 2016; Mai and Sinha, 2023]. The QR model captures bounded rationality by allowing attackers to choose actions probabilistically, with higher utility actions selected more frequently. Several variants of the QR model have also been explored to improve descriptive accuracy, including subjective utility QR [Nguyen *et al.*, 2013], prospect-theoretic models [Kar *et al.*, 2015], and nested logit-based QR models [Mai and Sinha, 2022].

Most prior work on QR-based SSGs, however, assumes a homogeneous attacker population governed by a single set of QR parameters. This assumption significantly simplifies equilibrium computation and often leads to objectives with favorable unimodality or quasi-concavity properties, enabling the use of gradient-based or fractional programming techniques [Tambe, 2011; Mai and Sinha, 2023; Mai *et al.*, 2024]. In contrast, relatively little is known about SSGs with heterogeneous attackers following different QR behaviors. While some works consider multiple attacker types [Sinha *et al.*, 2018], existing solution approaches typically rely on local optimization methods and do not provide global optimality or approximation guarantees.

From an algorithmic perspective, a closely related line of work studies gradient-based methods for learning and solving games with QR responses [Mertikopoulos and Sandholm, 2016; Li *et al.*, 2020; Sinha *et al.*, 2018]. In particular, [Perrault *et al.*, 2020] applies differentiating-through-optimization techniques to SSGs with QR adversaries. As with other gradient-based approaches, these methods guarantee only convergence to stationary points and may fail in the presence of highly nonconvex objectives. Our work differs fundamentally by focusing on ***global optimization*** and ***polynomial-time approximation guarantees*** for SSG with heterogeneous QR attackers.

Beyond QR, several studies analyze alternative behavioral models in Stackelberg settings, including complexity and learnability aspects [Jiang *et al.*, 2013; Cerny *et al.*, 2020; Milec *et al.*, 2021]. These works primarily address single-type or structurally restricted models and do not consider the computational challenges induced by attacker heterogeneity under mixed QR responses.

Discrete Choice Modeling and Heterogeneity. Our study of SSGs under a mixed QR model is closely related to the literature on discrete choice modeling and mixed logit models [Train, 2009; McFadden and Train, 2000]. Discrete choice models have a long history in modeling human decision-making under uncertainty and bounded rationality. While the multinomial logit (MNL) and QR models are widely used due to their analytical tractability, they implicitly assume homogeneous agents, which can lead to suboptimal solutions in practice. To capture heterogeneity, mixture models and mixed logit formulations have been extensively studied in economics, transportation, and revenue management [McFadden and Train, 2000; Hess and Polak, 2005]. In these models, agents are drawn from a finite or continuous distribution over preference parameters, resulting in choice probabilities that are mixtures of logit responses.

From an optimization standpoint, incorporating heterogeneity via mixed MNL or mixed QR models substantially complicates decision-making problems. In particular, optimization under mixed logit or mixture-of-logit models is known to be highly nonconvex and computationally challenging, and existing work often focuses on heuristics or restricted settings with limited constraints [Li *et al.*, 2019; Marandi and Lurkin, 2023; van de Geer and den Boer, 2022]. These challenges are closely mirrored in SSGs under mixed QR attackers, where the defender's objective becomes a weighted sum of multiple fractional logit terms with distinct normalizations.

Our work bridges these two literatures by studying SSGs under a mixed QR model that explicitly captures attacker heterogeneity while retaining a finite and interpretable set of attacker types. Unlike prior approaches, we exploit the specific algebraic structure of the mixed QR objective to develop a bilinear exponential-cone reformulation and a B&B algorithm with provable approximation guarantees. To the best of our knowledge, this is the first work to establish a *PTAS* for a class of Stackelberg security games under a mixed QR attacker model, and correspondingly the first such result for constrained pricing problems under mixed logit demand.

3 SSGs with Mixed QR Followers

In this section, we introduce an SSG model in which the follower (attacker) population exhibits heterogeneous, boundedly rational behavior. This heterogeneity is captured via a *discrete mixed QR* model, which is closely related to the mixed logit model in discrete choice theory [McFadden and Train, 2000].

Game Model and Defender Strategy. We consider a standard SSG with a single defender and a single attacker. The defender aims to protect a finite set of targets $\mathcal{M} = \{1, \dots, m\}$ by allocating limited security resources. The defender commits to a mixed strategy represented by marginal protection probabilities $x_j \in [0, 1], \quad \forall j \in \mathcal{M}$, subject to a global resource constraint $\sum_{j=1}^{m} x_j \leq C$. Here, $C > 0$ denotes the total defensive resource capacity available to the defender. More generally, x_j can be interpreted as the amount of defensive resources allocated to target j, with $x_j \geq 0$. The constraint $\sum_{j=1}^{m} x_j \leq C$ captures the defender's overall resource budget. Note that, for simplicity, we assume a single capacity constraint; however, our approach readily extends to more general settings with multiple constraints of the form $\sum_{j \in \mathcal{M}_k} x_j \leq C_k$ for $k = 1, \dots, K$. If a target j is attacked, the defender receives a reward r_j^d if the target is protected and a penalty l_j^d otherwise, with $r_j^d > l_j^d$. The defender's expected utility for the target j is therefore $u_j^d(x_j) = x_j r_j^d + (1 - x_j) l_j^d = \delta_j^d x_j + l_j^d$, where $\delta_j^d = r_j^d - l_j^d \geq 0$.

Mixed QR Attacker Model. We now consider an attacker model under mixed QR, in which the attacker population consists of multiple heterogeneous types. Specifically, the attacker population is represented by a finite set of types $\Theta = \{1, \dots, T\}$, where type $t \in \Theta$ occurs with probability $\pi_t > 0$, and $\sum_{t=1}^{T} \pi_t = 1$. Each attacker type t is characterized by a QR model with type-specific parameters. Given a defender strategy x, the utility of attacking a target j for a type-t attacker is $u_{jt}^a(x_j) = -\delta_j^a x_j + r_{jt}^a$, where $\delta_j^a \geq 0$ represents the attacker's sensitivity to the defender's protection effort and r_{jt}^a captures other target-specific factors that affect the utility of attacker type t. Here, to ensure the existence of a PTAS, we assume that the sensitivity parameters δ_j^a are identical across all attacker types, while all other utility components may remain attacker-type dependent.

Conditional on type t, the attacker follows a QR (logit) choice model. The probability that a type-t attacker chooses to attack a target j is given by

$$q_{jt}(x) = \frac{\exp\left(u_{jt}^a(x_j)\right)}{U_{0t} + \sum_{k=1}^{m} \exp\left(u_{kt}^a(x_k)\right)},$$

where $U_{0t} \geq 0$ represents the utility of the outside option corresponding to not attacking any target. The defender does not observe the realized attacker type. Consequently, the marginal probability that target j is attacked is given by the mixture $\bar{q}_j(x) = \sum_{t=1}^{T} \pi_t \, q_{jt}(x)$. This mixed QR model generalizes the standard QR model by allowing heterogeneity in attacker preferences and sensitivities and corresponds to a discrete mixed logit model in the discrete choice literature.

Equilibrium Problem Formulation. Under the mixed QR attacker model with an outside option, the defender's expected utility is given by

$$f(x) = \sum_{j=1}^{m} \bar{q}_j(x) \, u_j^d(x_j) = \sum_{t=1}^{T} \pi_t \sum_{j=1}^{m} q_{jt}(x) \, u_j^d(x_j),$$

where $q_{jt}(x)$ denotes the probability that a type-t attacker attacks target j. Substituting the logit choice probabilities with an outside option yields

$$f(x) = \sum_{t=1}^{T} \pi_t \frac{\sum_{j=1}^{m} \exp\left(u_{jt}^a(x_j)\right) u_j^d(x_j)}{U_{0t} + \sum_{k=1}^{m} \exp\left(u_{kt}^a(x_k)\right)},$$

where $U_{0t} \geq 0$ represents the utility of the outside option corresponding to not attacking any target for a type-t attacker.

The defender seeks to compute a Stackelberg strategy that maximizes expected utility against the mixed QR attacker population. This leads to the following optimization problem:

$$\max_{x \in \mathcal{X}} \quad f(x), \tag{1}$$

where $\mathcal{X}$ denotes the feasible set of defender strategies, incorporating both the bound constraints $x_j \in [0, 1]$ and the total resource constraint $\sum_{j=1}^{m} x_j \leq C$. Problem (1) is generally *non-convex* and *non-separable* due to the ratio-of-sums structure induced by the mixed QR model and the presence of the outside option. As a result, exact equilibrium computation is intractable in general, motivating the development of efficient approximation algorithms, which we study in subsequent sections.

We note that, when $T = 1$ (typical setting in prior SSG under QR works), the defender's objective reduces to a single fractional function, and existing work typically applies Dinkelbach-type transformations [Tambe, 2011; Yang *et al.*, 2011; Mai and Sinha, 2023; Mai *et al.*, 2024] to convert the problem into a sequence of convex programs with guaranteed global convergence. This approach crucially relies on the presence of a single denominator. In contrast, when $T > 1$, the objective becomes a weighted sum of multiple fractional terms with distinct denominators, for which no single Dinkelbach parameter can simultaneously linearize all ratios. As a result, standard fractional programming techniques no longer apply, and the objective is generally non-concave with multiple local optima, motivating the need for the global optimization framework developed in this work.

Non-unimodality. In the following, we show that when $T > 1$, the objective function is generally non-unimodal, implying that gradient-based methods can easily become trapped in suboptimal local optima. We state the result in the following proposition.

Proposition 1 (Unimodality vs. Multimodality of the Objective)**.** *If the number of attacker types is $T = 1$, then the defender's objective function $f(x)$ is unimodal over the feasible region, and hence any local maximizer is globally optimal. In contrast, when $T > 1$, the objective function is generally non-concave and may admit multiple local optima.*

While the proof of unimodality is deferred to the appendix, the lack of unimodality can be demonstrated by the example in Fig. 1, which shows that multiple local optima may arise when $T = 2$ or $T = 3$.

The above proposition (see also the proof in the appendix) shows that when $T = 1$, the problem admits a concave reformulation and can therefore be solved efficiently in polynomial time. In contrast, when $T > 1$, the defender's objective becomes a weighted sum of T fractional logit terms with distinct denominators. Such sum-of-ratios structures are known

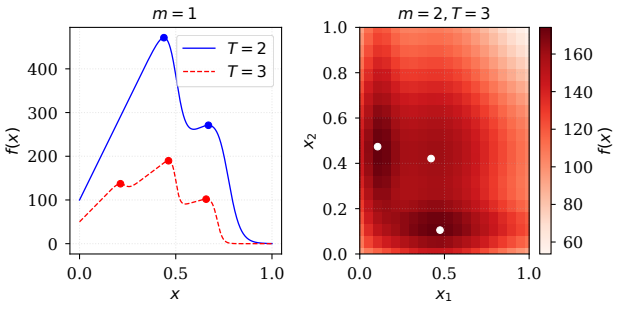

Figure 1: Visualization of the multi-modal objective function $f(x)$: single target ($m = 1$) with $T \in \{2, 3\}$ attackers and two targets ($m = 2$) with $T = 3$ attackers. Local maxima are highlighted.

to be computationally challenging to solve to global optimality in general [Gilbert *et al.*, 2013; Li *et al.*, 2019]. In the following, we present a positive result showing that, when $T > 1$ but bounded, the problem admits an efficient **PTAS** for computing an ε-optimal solution.

4 Bilinear Exponential-Cone Reformulation

We present the first step toward deriving a PTAS for the challenging SSG problem described above. The key idea is to introduce an innovative change of variables that reformulates the problem as an exponential-cone program, in which all sources of nonconvexity are consolidated into a small set of bilinear terms. This structural property enables the design of a polynomial-time B&B algorithm and facilitates a rigorous analysis of the overall computational complexity using tools from interior-point methods [Nesterov and Nemirovskii, 1994; MOSEK ApS, 2024].

We begin by recalling the definition of the *exponential cone* $\mathcal{K}_{\exp} \subset \mathbb{R}^3$ [Chares, 2009], which is given by

$$\mathcal{K}_{\exp} = \left\{(u, v, w) \in \mathbb{R}^3 \;\middle|\; v > 0,\ ve^{u/v} \le w\right\} \cup \left\{(u, 0, w) \in \mathbb{R}^3 \;\middle|\; u \le 0,\ w \ge 0\right\}.$$

The exponential cone provides a convex representation of exponential and logarithmic relations, such as $w \ge e^u$ or equivalently $u \le \ln w$, and is supported by modern conic optimization solvers, including MOSEK.

We summarize our key structural result below. It shows that the equilibrium optimization problem $\max_x f(x)$ can be reformulated as a bilinear exponential-cone program, in which all sources of nonconvexity are confined to a small number of bilinear constraints—exactly T, corresponding to the number of attacker types.

Theorem 1 (Bilinear Exponential-Cone Reformulation)**.** *The Stackelberg equilibrium optimization problem* $\max_{x \in \mathcal{X}} f(x)$ *under the mixed QR attacker model is equivalent to the following bilinear exponential-cone program:*

$$\max_{w,\,\theta,\,y,\,z,\,s,\,v} \quad \sum_{t=1}^{T} \pi_t \theta_t \qquad \text{(SSG-BiConv)}$$

$$\text{s.t.} \quad y_t = \theta_t z_t, \quad \forall t \in [T], \qquad (2)$$

$$y_t \le \sum_{j=1}^{m} e^{r^a_{jt}} \left(-\tfrac{\delta^d_j}{\delta^a_j} s_j + l^d_j w_j\right), \quad \forall t \in [T],$$

$$z_t \ge U_{0t} + \sum_{j=1}^{m} e^{r^a_{jt}} w_j, \quad \forall t \in [T],$$

$$\sum_{j=1}^{m} \frac{1}{\delta^a_j}(-v_j) \le C,$$

$$-\delta^a_j \le v_j \le 0, \quad \forall j \in [m],$$

$$(-s_j, w_j, 1) \in \mathcal{K}_{\exp}, \quad \forall j \in [m],$$

$$(v_j, 1, w_j) \in \mathcal{K}_{\exp}, \quad \forall j \in [m],$$

$$w_j > 0, \quad \forall j \in [m].$$

Moreover, all constraints in (SSG-BiConv) *are convex and admit exact exponential-cone representations, except for the T bilinear equalities $y_t = \theta_t z_t$, which constitute the only source of nonconvexity.*

While the detailed proof of Theorem 1 is provided in the appendix, we summarize the main reformulation steps here. The key idea is to apply the exponential change of variables $w_j := e^{-\delta^a_j x_j}$, $x_j = -\frac{1}{\delta^a_j} \ln w_j$, which transforms the exponentiated attacker utilities in the QR model as $\exp(u^a_{jt}(x_j)) = e^{r^a_{jt}} w_j$. This change converts each attacker type's expected utility into a fractional form

$$\theta_t = \frac{\sum_{j=1}^{m} u^d_j\left(-\frac{1}{\delta^a_j} \ln w_j\right) e^{r^a_{jt}} w_j}{U_{0t} + \sum_{j=1}^{m} e^{r^a_{jt}} w_j},$$

which is captured by introducing auxiliary variables (y_t, z_t) satisfying $y_t = \theta_t z_t$. Substituting the affine defender utility $u^d_j(x_j) = \delta^d_j x_j + l^d_j$ yields terms of the form $-(\delta^d_j/\delta^a_j) w_j \ln w_j + l^d_j w_j$. The nonlinear expressions $w_j \ln w_j$ and $\ln w_j$ are convex for $w_j > 0$ and admit exact exponential-cone representations via auxiliary variables s_j and v_j: $(-s_j, w_j, 1) \in \mathcal{K}_{\exp}$, $(v_j, 1, w_j) \in \mathcal{K}_{\exp}$. After these transformations, all constraints in the reformulated problem are convex and conic-representable, except for the bilinear equalities $y_t = \theta_t z_t$, $\forall t \in [T]$, which constitute the only remaining source of nonconvexity. Importantly, the number of such bilinear constraints is exactly T, the number of attacker types, which is the structural property exploited by the approximation algorithm developed later.

The formulation (SSG-BiConv) isolates all sources of nonconvexity into a small number of bilinear constraints of the form $y_t = \theta_t z_t$, one for each attacker type $t \in [T]$, while all remaining constraints are convex and admit exact exponential-cone representations. This structural property is critical from an algorithmic standpoint. In particular,

when the bilinear equalities are relaxed using standard convex envelopes, such as McCormick relaxations over bounded domains, the resulting problem reduces to a purely convex exponential-cone program. Such programs can be solved to arbitrary ε-accuracy in time polynomial in the problem size and $\log(1/\varepsilon)$ using modern conic solvers. This observation can be leveraged to design a provably efficient B&B algorithm: the exponential-cone relaxation yields tight and efficiently computable upper bounds at each node, while branching is required only on the T bilinear terms. Consequently, when the number of attacker types T is treated as a constant or bounded from above, this approach yields a PTAS for solving (SSG-BiConv).

5 B&B with Provable Polynomial-Time Performance Guarantees

We now describe a B&B approach for solving the bilinear exponential-cone formulation (SSG-BiConv). In the following, we outline the key components of the proposed algorithm; the full algorithmic details are deferred to the appendix. We then provide a rigorous analysis of its computational complexity.

Convex Relaxation via McCormick Envelopes. At each node of the B&B tree, we construct a convex relaxation of (SSG-BiConv) by replacing each bilinear equality $y_t = \theta_t z_t$ with its McCormick relaxation over bounded domains. Specifically, given bounds $\theta_t \in [\underline{\theta}_t, \overline{\theta}_t]$ and $z_t \in [\underline{z}_t, \overline{z}_t]$, the bilinear term is relaxed by the standard McCormick inequalities [McCormick, 1976]:

$$y_t \geq \underline{\theta}_t z_t + \theta_t \underline{z}_t - \underline{\theta}_t \underline{z}_t,\ y_t \geq \overline{\theta}_t z_t + \theta_t \overline{z}_t - \overline{\theta}_t \overline{z}_t,$$
$$y_t \leq \underline{\theta}_t z_t + \theta_t \overline{z}_t - \underline{\theta}_t \overline{z}_t,\ y_t \leq \overline{\theta}_t z_t + \theta_t \underline{z}_t - \overline{\theta}_t \underline{z}_t.$$

After this replacement, the resulting problem is a purely convex exponential-cone program.

Node Evaluation. The convex relaxation at each node is solved using a standard exponential-cone solver (e.g., MOSEK) [MOSEK ApS, 2024], yielding an upper bound on the optimal objective value of the original nonconvex problem over the corresponding subregion. Since exponential-cone programs can be solved to ε-accuracy in polynomial time, node evaluation is computationally efficient.

Branching Strategy. Branching is performed exclusively on the bilinear variables by partitioning the domain of either θ_t or z_t for some attacker type $t \in [T]$. For example, a common strategy is to bisect the interval of θ_t with the largest McCormick relaxation gap. This generates two child nodes with tightened bounds, leading to progressively tighter convex relaxations.

Termination and Approximation Guarantee. Under the Best-First Search policy, the algorithm terminates immediately upon retrieving a node that satisfies the bilinear feasibility tolerance ε, as this node is guaranteed to be globally optimal. Since the number of bilinear terms T is fixed, the spatial branching tree size scales polynomially in $1/\varepsilon$, as discussed below.

5.1 Complexity Analysis

In the following, we analyze the computational complexity of the B&B above. The next proposition provides an explicit upper bound on the number of nodes explored by the algorithm.

Proposition 2 (Explicit Node Bound)**.** *Fix $\varepsilon > 0$, define $\underline{U}_d := \min_j l_j^d$, $\overline{U}_d := \max_j (l_j^d + \delta_j^d)$, and $\overline{Z} := \max_{t\in[T]} \left(U_{0t} + \sum_{j=1}^m e^{r_{jt}^a}\right)$. Then the total number of nodes explored by the B&B algorithm is bounded by*

$$N_{\text{nodes}} \leq 2 \left\lceil \frac{T}{2\varepsilon} (\overline{U}_d - \underline{U}_d)\, \overline{Z} \right\rceil^T - 1.$$

In particular, for fixed T, the number of nodes is polynomial in $1/\varepsilon$ and the problem size.

While the detailed proof is provided in the appendix, we briefly explain the main ideas underlying the bound. The analysis exploits the fact that all nonconvexity in (SSG-BiConv) is confined to the T bilinear equalities $y_t = \theta_t z_t$. At each node of the B&B tree, these bilinear terms are relaxed using McCormick envelopes over a bounded rectangular domain for (θ_t, z_t). The relaxation gap of each bilinear term decreases quadratically with the diameter of its domain. The algorithm recursively partitions the feasible domain of each bilinear pair (θ_t, z_t) until the maximum possible McCormick relaxation error is at most ε/T, ensuring that the overall optimality gap is bounded by ε. Since the initial ranges of θ_t and z_t are explicitly bounded by problem-dependent constants $(\overline{U}_d - \underline{U}_d)$ and $\overline{Z}$, respectively, the number of subdivisions required along each dimension scales on the order of $\sqrt{(\overline{U}_d - \underline{U}_d)\overline{Z}/\varepsilon}$. Taking the product over the T independent bilinear terms yields the stated node bound.

This bound on the number of B&B nodes can be combined with the computational complexity of solving the exponential-cone program at each node—using interior-point methods [Nesterov and Nemirovskii, 1994]—to obtain the overall complexity of the algorithm, as stated in the following theorem.

Theorem 2 (Overall Complexity)**.** *For any $\varepsilon > 0$, the proposed B&B algorithm computes an ε-optimal solution to* (SSG-BiConv) *in time*

$$O\left(\left(\frac{T(\overline{U}_d - \underline{U}_d)\,\overline{Z}}{2\varepsilon}\right)^T \cdot m^{3.5} \cdot \log\frac{1}{\varepsilon}\right),$$

In particular, for fixed T, the algorithm runs in time polynomial in m and $1/\varepsilon$ and thus constitutes a PTAS.

The complexity bound in Theorem 2 immediately implies the existence of a PTAS for (SSG-BiConv) when the number of attacker types T is bounded. Specifically, the overall runtime grows exponentially only in T, while remaining polynomial in the number of targets m and in $1/\varepsilon$, the desired approximation accuracy. Since T captures the degree of attacker heterogeneity and is typically small in practical security applications, this result establishes the tractability of the mixed QR Stackelberg security game in realistic regimes. From a

complexity-theoretic perspective, the result shows that the intrinsic difficulty of the problem is driven entirely by the dimensionality of heterogeneity: once T is fixed, the combination of exponential-cone relaxations and B&B yields globally optimal solutions up to arbitrary precision in polynomial time. This sharply contrasts with general nonconvex optimization, where approximation guarantees of this form are rarely available, and highlights the importance of exploiting the problem's conic and low-rank bilinear structure.

6 Experimental Results

We evaluate our approach on synthetic and benchmark security game instances. The experiments focus on comparing our B&B method with state-of-the-art gradient-based algorithms. We also compare the mixed QR model with the standard single-type QR model to demonstrate the benefits of explicitly modeling heterogeneous attackers. Additional experimental results are reported in the appendix.

Data Generation and Experimental Setting. We compare our B&B method against a gradient-based baseline on synthetic datasets with the number of targets $m \in \{20, 50, 100, 200\}$ and the number of attackers $T \in \{1, 2, 3, 4\}$. For each target $j \in \mathcal{M}$, the defender's penalty l_j^d and reward r_j^d are independently sampled from the intervals [1, 400] and [600, 1000], respectively. The term r_{jt}^a of the attacker t for attacking the target j is sampled uniformly from [-5, 5]. We enforce five capacity constraints. For each $k \in \{1, \ldots, 5\}$, we set the capacity $C_k \in [0.2m, 0.5m]$ and define the constraint over a random selection of half the total targets. For each pair (m, T), we generate 10 independent instances and report aggregated results after running each instance for at most 1 hour with an optimal tolerance of $\epsilon = 0.001$. Our main experiments consider instances with a small number of attacker types ($T \leq 4$), for which the proposed method reliably computes globally optimal solutions. Additional experiments with larger T are reported in the appendix, demonstrating that our approach remains competitive and consistently outperforms gradient-based methods even when exact optimality can no longer be guaranteed.

Both the B&B and gradient-based algorithms are implemented in Python. All experiments are conducted on a personal computer running Windows 10, equipped with an Intel(R) Core(TM) i7-1255U CPU @ 1.70GHz and 16 GB of RAM. We use MOSEK Optimization Tools version 11 to solve the exponential-cone subproblems and the SciPy library (version 1.17.0) for the gradient-based method.

Single-type QR vs. Mixed QR. As previously noted, for $T = 1$, the objective function constitutes a single fractional function that is unimodal, thereby ensuring efficient convergence to the global optimum using established methods. For scenarios with $T > 1$, we examine the efficacy of approximating the multi-attacker SSG as a single-attacker problem to leverage these existing techniques. Specifically, we construct a surrogate single-attacker instance by aggregating the parameters of multiple attackers via their mean values to evaluate the resulting objective function. Figure 2 demonstrates the inadequacy of the typical setting in prior SSG under QR works ($T = 1$) in capturing the complex dynamics of the

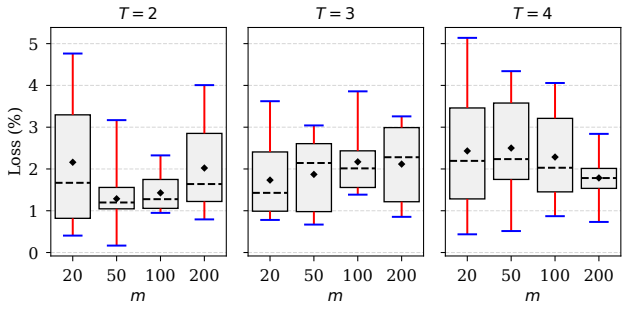

Figure 2: The loss of objective function under the mixed MNL model compared to MNL model.

mixed QR attacker model ($T \geq 2$). The results reveal a persistent objective loss incurred by this approximation across all target sizes (m) and attacker configurations. Rather than converging to the true optimal values, the simplified MNL formulation consistently yields suboptimal solutions, with approximation errors reaching up to 5% in the worst-case scenarios (indicated by the upper whiskers). Even on average, the method suffers from a systematic loss typically ranging between 1% and 2.5%, underscoring that the structural intricacies of the multi-attacker setting cannot be fully represented by the single-attacker surrogate.

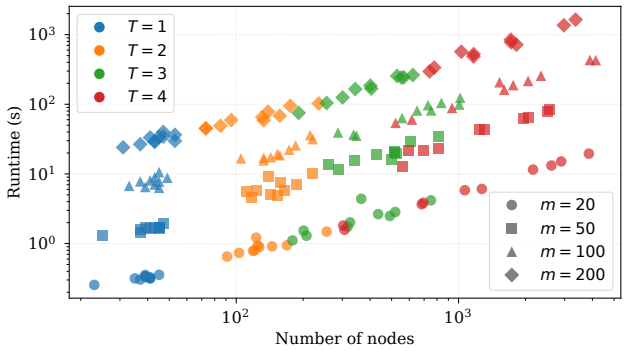

Figure 3: The relation between the number of nodes explored and the runtime in the B&B method.

Node exploration in B&B. Figure 3 illustrates the relationship between the runtime and the number of nodes explored in the B&B tree for varying T and m. A strong positive linear correlation is observed between the number of nodes and runtime when plotted on a log-log scale, indicating a consistent computational cost per node across different problem configurations. As the number of attackers T increases from 1 to 4, both the number of nodes and the total runtime increase exponentially, evidenced by the distinct clusters of data points shifting upwards and to the right. Specifically, problems with $T = 4$ require exploring significantly more nodes (often exceeding 10^3) and result in runtimes several orders of magnitude higher compared to $T = \{1, 2, 3\}$. While the number of targets m also contributes to increased runtime (indicated

by different markers within each cluster), its impact is less dominant than the exponential growth driven by the number of attackers T. This behavior underscores that the primary driver of complexity in this SSG setting is the number of attackers, which expands the search space exponentially.

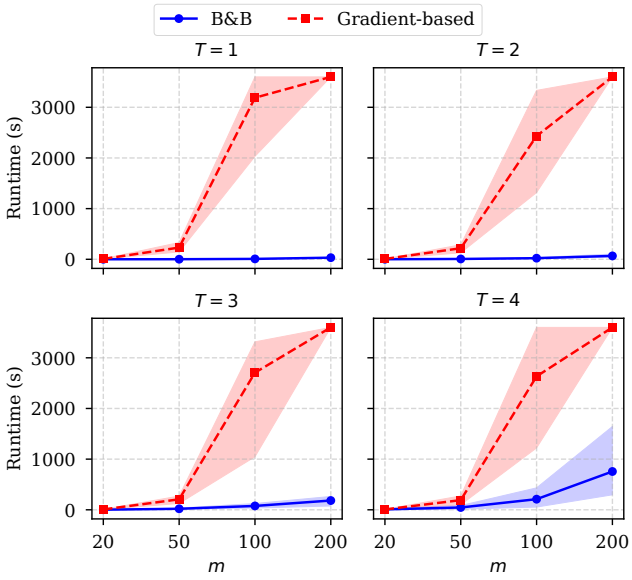

Figure 4: Runtime comparison of B&B and gradient-based methods.

B&B vs. Gradient-based Method. Figure 4 compares the runtime scalability of the B&B and gradient-based algorithms. The red dashed line with square markers denotes the gradient-based method, and the blue solid line with circular markers represents our B&B algorithm. The shaded regions indicate the variance in solving time across instances. The first plot highlights the effectiveness of the proposed B&B algorithm in standard SSG-QR scenarios ($T = 1$). It achieves global optimality in less than 40 seconds, even when handling the maximum instance size of $m = 200$. Figure 4 indicates that the B&B method consistently outperforms the gradient-based baseline across all tested parameter settings, successfully solving all instances optimally with drastically lower runtime. In contrast, the gradient-based method faces major scalability issues, hitting the 3600-second ceiling as m approaches 100 or 200, whereas the B&B algorithm remains computationally efficient. Even as problem complexity increases to the largest datasets ($m = 200, T = 4$), the average runtime of B&B is approximately 5 times lower than that of the gradient-based method, which consistently saturates the time limit in high-complexity regimes. Figure 4 also proves that gradient-based methods are highly sensitive to increases in both T and m. This is because gradient-based approaches operate on the continuous relaxation of the problem, where the computational cost per iteration is governed by the dimensionality of the constraints and variables. Since the number of collision avoidance constraints scales with the product of attackers and targets ($T \times m$), an increase in m directly inflates the size of the Jacobian and Hessian matrices, significantly slowing down convergence, whereas for B&B, the combinatorial explosion caused by T overshadows the linear complexity added by m.

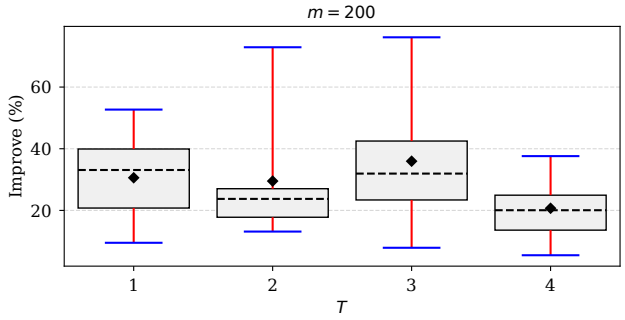

Figure 5: The gap between B&B and gradient-based methods on datasets with a large number of targets.

Although the gradient-based method can solve instances to near-optimal solutions (gap $< 0.01\%$) when $m \leq 100$, the objective gap compared to B&B becomes significant when $m = 200$, as detailed in Figure 5. The figure highlights the substantial improvement in solution quality achieved by the B&B algorithm compared to the gradient-based baseline for datasets with the largest number of targets. The results demonstrate a significant performance gap, as B&B consistently yields higher objective values across all attacker configurations. In this visualization, the diamond markers denote the mean percentage improvement, while the dashed horizontal lines within the boxes indicate the median values. The advantage is particularly pronounced for $T = 2$ and $T = 3$, where the maximum improvement exceeds 76%, and mean improvements hover near 30% and 35%, respectively. Even at $T = 4$, the method maintains a median improvement of approximately 20%, underscoring the inability of the gradient-based approach to escape local optima in complex, high-dimensional target spaces.

7 Conclusion

We studied SSGs with heterogeneous attackers modeled via a mixed QR framework. By developing a bilinear exponential-cone reformulation that confines all nonconvexity to the number of attacker types, we designed a B&B algorithm that yields a PTAS when the number of attacker types is bounded. Experimental results demonstrate that explicitly modeling attacker heterogeneity significantly improves solution quality and that our approach consistently outperforms gradient-based methods in nonconvex settings.

Our PTAS guarantee relies on treating the number of attacker types as a constant, and the algorithm's complexity grows exponentially with this dimension, reflecting an inherent trade-off between modeling expressiveness and computational tractability. Important directions for future work include extending the framework to richer or continuous heterogeneity models, relaxing structural assumptions such as homogeneous sensitivity parameters, and more tightly integrating learning and optimization in data-driven security games.

Acknowledgments

This research was supported by the Vingroup Innovation Foundation (VINIF), Vietnam, under project code VINIF.2024.DA072.

References

[An *et al.*, 2012] Bo An, Eric Shieh, Milind Tambe, Rong Yang, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. Protect–a deployed game theoretic system for strategic security allocation for the united states coast guard. *Ai Magazine*, 33(4):96–96, 2012.

[Cerny *et al.*, 2020] Jakub Cerny, Viliam Lisý, Branislav Bošanský, and Bo An. Dinkelbach-type algorithm for computing quantal stackelberg equilibrium. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 246–253, 2020.

[Chares, 2009] Robert Chares. *Cones and interior-point algorithms for combinatorial optimization*. PhD thesis, Université catholique de Louvain, 2009.

[Fang *et al.*, 2017] Fei Fang, Thanh H Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Brian C Schwedock, Milin Tambe, and Andrew Lemieux. PAWS—A deployed game-theoretic application to combat poaching. *AI Magazine*, 38(1), 2017.

[Gilbert *et al.*, 2013] François Gilbert, Patrice Marcotte, and Gilles Savard. Mixed-logit network pricing. *Computational Optimization and Applications*, 57:105–127, 2013.

[Haghtalab *et al.*, 2016] Nika Haghtalab, Fei Fang, Thanh H. Nguyen, Arunesh Sinha, Ariel D. Procaccia, and Milind Tambe. Three strategies to success: Learning adversary models in security games. In *25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[Hess and Polak, 2005] Stephane Hess and John W. Polak. Mixed logit modelling of airport choice in multi-airport regions. *Journal of Air Transport Management*, 11(2):59–68, 2005.

[Jiang *et al.*, 2013] Albert Xin Jiang, Thanh H Nguyen, Milind Tambe, and Ariel D Procaccia. Monotonic maximin: A robust stackelberg solution against boundedly rational followers. In *International Conference on Decision and Game Theory for Security*, pages 119–139. Springer, 2013.

[Kar *et al.*, 2015] Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe. "a game of thrones" when human behavior models compete in repeated stackelberg security games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1381–1390, 2015.

[Li *et al.*, 2019] Hongmin Li, Scott Webster, Nicholas Mason, and Karl Kempf. Product-line pricing under discrete mixed multinomial logit demand. *Manufacturing & Service Operations Management*, 21:14–28, 2019.

[Li *et al.*, 2020] Jiayang Li, Jing Yu, Yu Nie, and Zhaoran Wang. End-to-end learning and intervention in games. *Advances in Neural Information Processing Systems*, 33, 2020.

[Mai and Sinha, 2022] Tien Mai and Ankur Sinha. Choices are not independent: Stackelberg security games with nested quantal response models. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 2022.

[Mai and Sinha, 2023] Tien Mai and Ankur Sinha. Safe delivery of critical services in areas with volatile security situation via a stackelberg game approach. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023.

[Mai *et al.*, 2024] Tien Mai, Anirban Bose, Ankur Sinha, and Thanh H. Nguyen. Stackelberg network interdiction under boundedly rational adversary. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2024.

[Marandi and Lurkin, 2023] Ahmadreza Marandi and Virginie Lurkin. An exact algorithm for the static pricing problem under discrete mixed logit demand. *EURO Journal on Computational Optimization*, 11:100073, 2023.

[McCormick, 1976] Garth P. McCormick. Computability of global solutions to factorable nonconvex programs: Part i — convex underestimating problems. *Mathematical Programming*, 10(1):147–175, 1976.

[McFadden and Train, 2000] Daniel McFadden and Kenneth Train. Mixed MNL models for discrete response. *Journal of applied Econometrics*, 15(5):447–470, 2000.

[Mertikopoulos and Sandholm, 2016] Panayotis Mertikopoulos and William H Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, 2016.

[Milec *et al.*, 2021] David Milec, Jakub Černỳ, Viliam Lisỳ, and Bo An. Complexity and algorithms for exploiting quantal opponents in large two-player games. In *AAAI Conference on Artifical Intelligence*, 2021.

[MOSEK ApS, 2024] MOSEK ApS. Mosek modeling cookbook. Technical report, 2024.

[Nesterov and Nemirovskii, 1994] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*, volume 13. Siam, 1994.

[Nguyen *et al.*, 2013] Thanh Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. Analyzing the effectiveness of adversary modeling in security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 2013.

[Perrault *et al.*, 2020] Andrew Perrault, Bryan Wilder, Eric Ewing, Aditya Mate, Bistra Dilkina, and Milind Tambe. End-to-end game-focused learning of adversary behavior in security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1378–1386, 2020.

[Pita *et al.*, 2008] James Pita, Manish Jain, Fernando Ordónez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Armor security for los angeles international airport. In *AAAI*, pages 1884–1885, 2008.

[Sinha *et al.*, 2018] Arunesh Sinha, Fei Fang, Bo An, Christopher Kiekintveld, and Milind Tambe. Stackelberg security games: Looking beyond a decade of success. In *27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

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

[Train, 2009] Kenneth E Train. *Discrete choice methods with simulation*. Cambridge university press, 2009. https://eml.berkeley.edu/books/choice2.html.

[van de Geer and den Boer, 2022] Ruben van de Geer and Arnoud V. den Boer. Price optimization under the finite-mixture logit model. *Management Science*, 68:7480–7496, 2022.

[Yang *et al.*, 2011] Rong Yang, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, and Richard John. Improving resource allocation strategy against human adversaries in security games. In *International Joint Conference on Artificial Intelligence (IJCAI)*, page 458, 2011.

[Yang *et al.*, 2012] Rong Yang, Fernando Ordonez, and Milind Tambe. Computing optimal strategy against quantal response in security games. In *11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.