

# Dinkelbach-Type Algorithm for Computing Quantal Stackelberg Equilibrium

Jakub Černý<sup>1</sup>, Viliam Lisý<sup>2</sup>, Branislav Bošanský<sup>2</sup> and Bo An<sup>1</sup>

<sup>1</sup>School of Computer Science and Engineering, Nanyang Technological University, Singapore

<sup>2</sup>Dept. of Computer Science, Faculty of El. Engineering, Czech Technical University in Prague, Czechia

cerny@disroot.org, {viliam.lisy, branislav.bosansky}@fel.cvut.cz, boan@ntu.edu.sg

## Abstract

Stackelberg security games (SSGs) have been deployed in many real-world situations to optimally allocate scarce resource to protect targets against attackers. However, actual human attackers are not perfectly rational and there are several behavior models that attempt to predict subrational behavior. *Quantal response* is among the most commonly used such models and *Quantal Stackelberg Equilibrium* (QSE) describes the optimal strategy to commit to when facing a subrational opponent. Non-concavity makes computing QSE computationally challenging and while there exist algorithms for computing QSE for SSGs, they cannot be directly used for solving an arbitrary game in the normal form. We (1) present a transformation of the primal problem for computing QSE using a Dinkelbach's method for any general-sum normal-form game, (2) provide a gradient-based and a MILP-based algorithm, give the convergence criteria, and bound their error, and finally (3) we experimentally demonstrate that using our novel transformation, a QSE can be closely approximated several orders of magnitude faster.

## 1 Introduction

Game-theoretic algorithms have been used for improving physical security (see examples in [Tambe, 2011]), protecting wildlife in natural parks [Fang *et al.*, 2017], or beating human professionals in poker [Moravčík *et al.*, 2017; Brown and Sandholm, 2018]. Following the standard game-theoretic assumption of perfectly rational players is, however, not optimal against human opponents and there are more effective strategies against subrational opponents [Pita *et al.*, 2008; Delle Fave *et al.*, 2014; Nguyen *et al.*, 2016]. One of the most commonly used subrational models (a model of bounded rationality) is *Quantal Response* (QR) assuming that players choose better actions more often than worse actions instead of choosing only the best action [McKelvey and Palfrey, 1995]. In practical applications, the goal is to find the best strategy to deploy (i.e., to commit to) while assuming the other player plays the best response to this commitment – i.e., to compute

a *Stackelberg Equilibrium*. QR is used as the response function in case of human players and the desired solution concept is termed *Quantal Stackelberg Equilibrium* (QSE).

Stackelberg Security Games (SSGs) [Tambe, 2011] are particularly useful game model with many deployed real-world applications but with certain limitations. The problem must be formulated in terms of allocating limited resources to a set of targets. This is often impossible, e.g., in classical games from economics (i.e., prisoner's/traveler's dilemma, location game [Nudelman *et al.*, 2004]). In order to solve real-world problems beyond SSGs, we study optimal behavior against a quantal response opponent in more general models of normal-form games. We consider a very general class of response functions modelling both risk-averse and risk-seeking opponents. An optimal strategy of the rational player against such a subrational opponent is described by a leader-follower solution concept: the QSE. The leader computes and implements a fixed strategy and the follower responds to this strategy based on the quantal response function.

The natural formulation of the optimization problem computing QSE is non-concave. Moreover, we show that the number of local minima may grow linearly with the number of actions even in a zero-sum game. Hence, simple gradient ascent approaches are not likely to be effective. It is difficult to further analyze the problem of computing QSE directly, therefore, we derive a Dinkelbach-type formulation of the problem and use it to identify sufficient conditions for concavity of the problem. If the conditions are satisfied, the optimal solution can be found by gradient ascent. Furthermore, we generalize the linear relaxation developed for SSGs [Yang *et al.*, 2012] to general normal-form game and use it to formulate a mixed integer linear program (MILP) for approximately solving the problem. The approach of SSGs cannot be applied directly, because it uses specific properties of SSGs.

In the experiments, we compare the convergence speed of gradient ascent in direct formulation to approximated Dinkelbach-type MILP formulation of QSE with various quantal functions. We show that the Dinkelbach-type algorithm is up to 25.5-times faster than one restart of gradient ascent and it also provides solutions of higher quality. For the most common quantal function (i.e., logit function) and sufficiently large games, the gradient ascent couldn't find a solution of the same quality even when running 153-times longer than the Dinkelbach-type algorithm.

## 2 Background

Normal-form games model one-shot interactions between players and can be visually represented as game matrices. Formally, a two-player NFG is defined as a tuple  $G = (\mathcal{N}, \mathcal{A}, u)$ :  $\mathcal{N} = \{l, f\}$  is a set of players, the leader and the follower. We use  $i$  to refer to one of the players, and  $-i$  to refer to his opponent.  $\mathcal{A} = (\mathcal{A}_l, \mathcal{A}_f)$  denotes ordered sets of actions for each player. For each pair of actions  $(a_l, a_f)$ ,  $a_l \in \mathcal{A}_l, a_f \in \mathcal{A}_f$ , we define a *utility function* for each player  $u_i : \mathcal{A}_l \times \mathcal{A}_f \rightarrow \mathbb{R}$ . In zero-sum setting, we assume that  $u = u_f$ , setting  $u_l = -u_f$ .

*Pure strategies*  $\Pi_i$  assign one action to play in the game. A *mixed strategy*  $\delta_i \in \Delta_i$  is a probability distribution over  $\Pi_i$ . Because the actions are ordered,  $\delta_i^j$  denotes the probability of player  $i$  taking the  $j^{\text{th}}$  action. For any pair of strategies  $\delta_l \in \Delta_l, \delta_f \in \Delta_f$  we use  $u_i(\delta_l, \delta_f)$  for the expected outcome of the game for player  $i$  when players follow strategies  $\delta_l, \delta_f$ . A *best response* of player  $i$  to the opponent's strategy  $\delta_{-i}$  is a strategy  $\delta_i^{BR} \in BR_i(\delta_{-i})$ , where  $u_i(\delta_i^{BR}, \delta_{-i}) \geq u_i(\delta'_i, \delta_{-i})$  for all  $\delta'_i \in \Delta_i$ .

### Solution Concepts in NFGs

We provide the formal definitions of Nash Equilibrium (NE) and Quantal Stackelberg Equilibrium (QSE).

**Definition 1.** Given a normal-form game  $G = (\mathcal{N}, \mathcal{A}, u)$ , a tuple of mixed strategies  $(\delta_l^N, \delta_f^N), \delta_l^N \in \Delta_l, \delta_f^N \in \Delta_f$  is a Nash Equilibrium if  $\delta_i^N$  is an optimal strategy of player  $i$  against strategy  $\delta_{-i}^N$ . Formally:

$$\delta_i^N \in BR(\delta_{-i}^N) \quad \forall i \in \{l, f\}. \quad (1)$$

In NE, both players make decisions fully rationally. In this work, we assume two possible diversions from this assumption. Both are often observed in laboratory experiments with human players [Camerer, 2011]. First, humans differ from rational players in their approaches towards risk. When exposed to uncertainty, humans might intentionally attempt to avoid the gamble (risk aversion) or, in contrast, they may seek higher risks (risk love). The approaches are modelled via strictly increasing expected-utility function  $e$ . While the function of risk-neutral players is the identity, risk-averse players behave according to a concave and risk-seeking players according to a convex function. Humans also tend to combine both approaches: loss-averse functions are risk-averse when the utility is non-positive and risk-seeking when the utility is positive [Yechiam and Hochman, 2013]. Vice versa, functions, which are risk-seeking when the utility is non-positive and risk-averse when the utility is positive are called loss-averse [Kahneman and Tversky, 2013].

Second, fully rational players always select the utility-maximizing option. Relaxing this assumption leads to a ‘‘statistical version’’ of best response, which takes into account the inevitable error-proneness of humans and allows the players to make systematic errors [McFadden, 1976].

**Definition 2.** Let  $e$  be an expected-utility function. Function  $QR : \Delta_l \rightarrow \Delta_f$  is a quantal response function if it is mono-

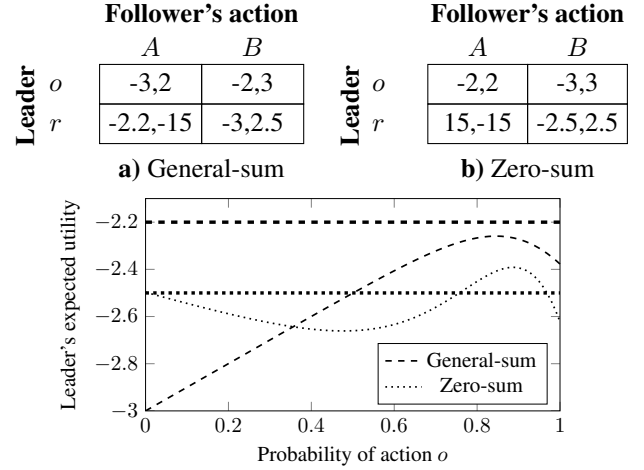


Figure 1: An example of a general-sum NFG (left matrix), the criterion functions of its QSEs (dashed curve) and its zero-sum variant (right matrix, dotted curve). Each row of the depicted matrices is labeled by a strategy of the first player (the leader), while every column is labeled by a strategy of the second player (the follower). The QSE maximizes the leader's expected utility. The values of Nash equilibria are depicted in corresponding bold straight lines.

tonically increasing in expected utility:

$$e(u_f(\delta, a_f^i)) \leq e(u_f(\delta, a_f^j)) \Rightarrow QR^i(\delta) \leq QR^j(\delta) \quad \forall \delta \in \Delta_l, a_f^i, a_f^j \in \mathcal{A}_f, \quad (2)$$

where  $QR^i(\delta)$  is the probability of playing action  $a^i$ . Quantal function  $QR$  is called canonical if

$$QR^i(\delta) = \frac{[e \circ q](u_f(\delta, a_f^i))}{\sum_{\pi \in \Pi_f} [e \circ q](u_f(\delta, \pi))} \quad (3)$$

for all  $\delta \in \Delta_l, a_f^i \in \mathcal{A}_f$  and some real-valued function  $q$ .

Note that whenever  $q$  is a strictly positive increasing function, the corresponding  $QR$  is a valid quantal response function. We call such functions  $q$  *generators* of canonical quantal functions. Because of the conceptual differences between the quantal and expected-utility functions (i.e., error proneness vs. risk assessment), we keep both functions separated.

**Definition 3.** Given a normal-form game  $G = (\mathcal{N}, \mathcal{A}, u)$  and a quantal response function  $QR$  of the follower, a mixed strategy  $\delta^{QS} \in \Delta_l$  describes a Quantal Stackelberg Equilibrium (QSE) if and only if

$$\delta^{QS} = \arg \max_{\delta \in \Delta_l} u_l(\delta, QR(\delta)). \quad (4)$$

In QSE, a leader is fully rational. He commits to an optimal strategy given that a follower quantal-responds according to given quantal and expected-utility functions. We refer to Equation (4) as to the criterion function of QSE.

**Example 1.** Consider the matrix game shown on the left in Figure 1. Let the follower behave according to a canonical quantal function with generator  $q(x) = e^{x/2}$  and utility function  $e(x) = x$  (identity). Below the matrix is depicted a criterion function of the game's QSE (dashed) and a criterion function of the game's zero-sum variant's QSE (dotted), in which we set  $u_l = -u_f$ . In the general-sum game, the optimal strategy of the leader is to commit to playing action  $o$

with probability 0.846. To achieve the QSE in the zero-sum variant, the leader plays  $o$  with probability 0.887.

The example shows that in general, finding the QSE is a non-linear non-concave problem. To design methods for efficient computation of QSE, we analyze the properties of QSE in more detail in the next section.

### 3 Properties of QSE in NFGs

In this section, we first study the relation of QSE to NE. Then we show why gradient ascent methods often fail to compute a global optimum. Note that because NFGs have a finite number of actions and all utilities are finite, there always exists a strategy maximizing the QSE criterion function with a finite value. Moreover, in zero-sum NFGs, the leader's utility in QSE is always greater or equal to his utility in NE.

**Observation 1.** Let  $G = (\mathcal{N}, \mathcal{A}, u)$  be a zero-sum NFG,  $QR$  be a quantal response function of the follower and  $\delta^N, \delta^{QS}$  be the mixed strategies of the leader in NE and QSE, respectively. Then  $u_l(\delta^N, BR(\delta^N)) \leq u_l(\delta^{QS}, QR(\delta^{QS}))$ .

*Proof.* From the definition of quantal response function  $u_f(\delta, BR(\delta)) \geq u_f(\delta, QR(\delta))$  for all  $\delta \in \Delta_l$ . Therefore, in zero-sum games it holds that  $u_l(\delta, BR(\delta)) \leq u_l(\delta, QR(\delta))$ . From the definition of QSE  $u_l(\delta, QR(\delta)) \leq u_l(\delta^{QS}, QR(\delta^{QS}))$ . Setting  $\delta = \delta^N$  yields the result.  $\square$

The same observation does not hold for general-sum games. As shown in Figure 1, the expected utility of the leader in the general-sum game is  $-2.26$ . In NE, the leader plays action  $o$  and the follower responds with  $B$ , resulting in the expected utility of the leader  $-2.2$ , strictly higher than in QSE. In other words, facing a subrational opponent does not have to be an advantage in a general-sum game. From now on, we will focus our analysis on canonical quantal functions.

**Observation 2.** Let  $G = (\mathcal{N}, \mathcal{A}, u)$  be a normal-form game and  $q$  be a generator of a quantal function with an expected-utility function  $e$ . Then the leader's strategy  $\delta^{QS}$  of QSE in  $G$  can be formulated as the following non-concave problem:

$$\delta^{QS} = \arg \max_{\delta \in \Delta_l} \frac{\sum_{\pi \in \Pi_f} u_l(\delta, \pi) [e \circ q](u_f(\delta, \pi))}{\sum_{\pi \in \Pi_f} [e \circ q](u_f(\delta, \pi))}. \quad (5)$$

Non-concave maximization problems can be solved by gradient ascent. The gradient ascents are an efficient class of algorithms with a guarantee to approach a local maximum in a polynomial number of steps in case the gradient of the criterion function is Lipschitz-continuous [Boyd and Vandenberghe, 2004]. Unfortunately, the criterion function of QSE can have a large number of such local optima.

**Conjecture 3.** For every  $n$ ,  $|\mathcal{A}_f| = n$ , there exists a zero-sum NFG  $G$  and a generator  $q$  of a canonical quantal function with an expected-utility function  $e$ , such that a criterion function of QSE in  $G$  has at least  $n - 1$  local maxima.

For every  $n$ , we define  $G_n = ((l, f), (\mathcal{A}_l, \mathcal{A}_f^n), u_f)$  and set  $\mathcal{A}_l = \{a^1, a^2\}$ ,  $\mathcal{A}_f^n = \{b^1, b^2, \dots, b^n\}$ . The utility function  $u_f$  is defined as  $u_f(a^1, b^1) = 0.5$ ,  $u_f(a^1, b^i) = -10^i$  for  $i \in [2, n]$  and  $u_f(a^2, b^i) = 0$  for all  $i \in [n]$ ,  $u_l = -u_f$ . Let  $q(x) = e^x$  and  $e(x) = x$  if  $x < 0$  and  $e(x) = 0$  for

#### Algorithm 1 Dinkelbach-Type Algorithm for QSE

---

```

1:  $UB \leftarrow \max_{\pi \in \Pi} u_l(\pi)$ ,  $LB \leftarrow \min_{\pi \in \Pi} u_l(\pi)$ 
2:  $\delta^* \leftarrow \arg \max_{\delta \in \Delta_l} \sum_{\pi \in \Pi_f} (u_l(\delta, \pi) - LB)[e \circ q](u_f(\delta, \pi))$ 
3: repeat
4:    $p \leftarrow (UB - LB)/2$ 
5:    $v \leftarrow \max_{\delta \in \Delta_l} \sum_{\pi \in \Pi_f} (u_l(\delta, \pi) - p)[e \circ q](u_f(\delta, \pi))$ 
6:    $\delta_p \leftarrow \arg \max_{\delta \in \Delta_l} \sum_{\pi \in \Pi_f} (u_l(\delta, \pi) - p)[e \circ q](u_f(\delta, \pi))$ 
7:   if  $v < 0$  then  $LB \leftarrow p$ ,  $\delta^* \leftarrow \delta_p$  else  $UB \leftarrow p$  end if
8: until  $UB - LB < \epsilon$ 
9: return  $\delta^*$ 
    
```

---

$x \geq 0$ . We were experimentally able to verify the conjecture up to games with  $n = 100$ , then the optima became too close to zero to verify their existence numerically. Applying restarted gradient ascent for solving QSE is thus impractical and can often lead to globally suboptimal local optima. Moreover, analysing the conditions when the criterion function of formulation (5) is concave is difficult, as concavity is not preserved under division<sup>1</sup>.

### 4 Dinkelbach-Type Formulation of QSE

This section introduces a transformation of the QSE formulation (5) into an equivalent optimization problem. This new formulation allows us to devise the conditions for approximating QSE in polynomial time.

The transformation is called the Dinkelbach's method for solving nonlinear fractional programming problems [Dinkelbach, 1967] and we can adapt it for finding the QSE because formulation (5) is a fractional program. The key idea of the Dinkelbach's algorithm is to express a problem in a form of  $\max_{x \in M} f(x)/g(x)$ ,  $g(x) > 0$  for some convex set  $M$  and continuous, real-valued functions  $f$  and  $g$  as an equivalent problem of finding a unique root of function  $F(p) = \max_{x \in M} \{f(x) - pg(x)\}$ ,  $p \in \mathbb{R}$ . It holds that  $\max_{x \in M} f(x)/g(x) = p^*$  if and only if  $F(p^*) = 0$ . Because  $F$  is a maximum of functions affine in  $p$ , it is convex. Finding the root is hence straightforward (e.g., a binary search method can be used for this purpose) and it can be done efficiently if and only if we are able to effectively determinate the value of function  $F$  for any  $p$ . Computing the maximum of the original formulation of QSE (5) is hence equivalent to solving a sequence of optimization subproblems

$$\max_{\delta \in \Delta_l} \sum_{\pi \in \Pi_f} (u_l(\delta, \pi) - p) [e \circ q](u_f(\delta, \pi)), \quad (6)$$

as described in Algorithm 1. We refer to Equation (6) as to the Dinkelbach subproblem of the Dinkelbach formulation of QSE. Algorithm 1 iteratively updates the upper bound ( $UB$ ) and lower bound ( $LB$ ) on the value of QSE according to a binary search method for finding a root of a function. Note that because of Observation 1, we are able to set the initial lower bound to the value of NE in case the game is zero-sum. This binary-search approach generalizes the same technique already used for computing logit QSE in security games [Yang et al., 2012]. The following proposition derives sufficient conditions for Algorithm 1 to run in polynomial time.

<sup>1</sup>Generalized concavity (e.g., pseudoconvavity or quasiconcavity) is preserved under division, but it imposes too strong assumptions on the game utilities to be used in practice.

**Proposition 4.** Let  $G = (\mathcal{N}, \mathcal{A}, u)$  be an NFG,  $q$  be a twice differentiable generator of a canonical quantal function and  $e$  be a twice differentiable expected-utility function, such that the gradient of the Dinkelbach formulation is Lipschitz-continuous. For each pure strategy  $\pi \in \Pi_f$ , denote  $\underline{u}_\pi^i = \min_{\pi_l \in \Pi_l} u_i(\pi, \pi_l)$  and  $\bar{u}_\pi^i = \max_{\pi_l \in \Pi_l} u_i(\pi, \pi_l)$ . Moreover, let  $\underline{u}_i = \min_{\pi \in \Pi_f} \underline{u}_\pi^i$  and  $\bar{u}_i = \max_{\pi \in \Pi_f} \bar{u}_\pi^i$ . Then QSE is polynomially approximable if for all  $\pi \in \Pi_f$  and any  $x \in [\underline{u}_\pi^l, \bar{u}_\pi^l]$ ,  $y \in [\underline{u}_\pi^f, \bar{u}_\pi^f]$  and  $p \in [\underline{u}_l, \bar{u}_l]$

$$(p - x) \left( q''(e(y))e'^2(y) + q'(e(y))e''(y) \right) \geq 0 \quad (7)$$

$$u_l(\pi)^T u_f(\pi) \leq 0,$$

where for any function  $g$ ,  $g'$  denotes a derivative of  $g$ .

*Proof.* If the Dinkelbach subproblem is concave, the gradient-ascent methods are guaranteed to reach a global optimum. We hence analyse the conditions when an equivalent minimization formulation is convex, i.e., its Hessian matrix is positive semidefinite (PSD). The Hessian matrix is of a form  $c \cdot u_f(\pi)u_f(\pi)^T - [[e \circ q](u_f(\pi)^T \delta)]' (u_l(\pi)u_f(\pi)^T + u_f(\pi)u_l(\pi)^T)$ , where  $c = \frac{(p - u_l(\pi)^T \delta) (q'(e(u_f(\pi)^T \delta))e''(u_f(\pi)^T \delta) + q''(e(u_f(\pi)^T \delta))e'^2(u_f(\pi)^T \delta))}{(p - u_l(\pi)^T \delta) (q'(e(u_f(\pi)^T \delta))e''(u_f(\pi)^T \delta) + q''(e(u_f(\pi)^T \delta))e'^2(u_f(\pi)^T \delta))}$  and  $u_l(\pi)^T \delta$  is the expected utility  $u_l(\delta, \pi)$  (and similarly for the follower). A well-known fact from linear algebra states that for non-zero  $n$ -dimensional real-valued vectors  $u, v$ , the matrix  $uv^T$  is positive semidefinite iff  $u^T v \geq 0$ . By this fact,  $u_f(\pi)u_f(\pi)^T$  is PSD. The matrix  $c \cdot u_f(\pi)u_f(\pi)^T$  is then PSD if  $c \geq 0$ . Because  $q, e$  are increasing functions, matrices  $-u_l(\pi)u_f(\pi)^T, -u_f(\pi)u_l(\pi)^T$  are PSD if  $u_l(\pi)^T u_f(\pi) \leq 0$ . Because the expected utility  $u_i(\pi)^T \delta$  always lies in the interval  $[\underline{u}_\pi^i, \bar{u}_\pi^i]$  and  $p$  is upper (lower) bounded by leader's maximal (minimal) utility in the game, by substituting for  $c$  we conclude that the Dinkelbach formulation of QSE is concave if conditions (7) holds.  $\square$

Note that with additional assumptions about functions  $q$  and  $e$  (e.g., concavity or convexity), the condition stated in Proposition 4 can be expressed just in the terms of maximal and minimal utility in the game and is hence easier to check.

## 5 Solving the Dinkelbach Subproblem

In case a game and the quantal and expected-utility functions do satisfy the condition stated in Proposition 4, the gradient-ascent methods will converge to a globally optimal strategy of the leader. In the opposite case, however, the guarantees are lost. In this section, we introduce efficient linearization methods with bounds on the solution quality to approximate the global optimum through a piecewise-linear approximation (PWLA) of the Dinkelbach subproblem in any game. The PWLA is subsequently turned into a mixed-integer linear program (MILP) and solved using standard methods.

### Limitations of PWLA in General NFGs

In SSGs, the formulation of the logit QSE has separable variables, i.e., the multivariate function can be written as a sum of univariate functions. Therefore, the authors were able to

linearize it efficiently [Yang *et al.*, 2012]. In contrast, the Dinkelbach formulation in NFGs is not separable in general and the probability simplex must be linearized into multivariate polytopes. The most straightforward method is to split the simplex into an exponential number of hypercubes of dimension  $|\mathcal{A}_l|$  using a fixed step size. To obtain a continuous PWLA, each hypercube must be split into an exponential number of triangulations. To escape the curse of dimensionality, instead of splitting a hypercube, we compute the PWLA in the whole hypercube by linear least squares. The resulting formulation is transformed to an MILP using the DLog construction [Ibaraki, 1976; Vielma *et al.*, 2010]. The DLog construction was shown to be the fastest by [Vielma *et al.*, 2010], because of introducing only a logarithmic number of binary variables in the number of polytopes. Since we construct one polytope for every hypercube and the number of hypercubes is exponential in  $|\mathcal{A}_l|$ , the final MILP is of a polynomial size in  $|\mathcal{A}_l|$ .

We implemented PWLA and experimentally verified that it won't scale to games with more than 5 actions of the leader. To improve its efficiency, it is necessary to find a precise enough separable approximation of the Dinkelbach subproblem:  $\max_{\delta \in \Delta_l} d_p^1(\delta) - p d_p^2(\delta)$ , where  $d_p^1(\delta) = \sum_{\pi \in \Pi_f} u_l(\delta, \pi) [e \circ q](u_f(\delta, \pi))$ ,  $d_p^2(\delta) = \sum_{\pi \in \Pi_f} [e \circ q](u_f(\delta, \pi))$ . Each of  $d_p^1, d_p^2$  is approximated as  $d_p^1(\delta) \approx \sum_{i=1}^{|\mathcal{A}_l|} \tilde{d}_{p,i}^1(\delta^i)$  and  $d_p^2(\delta) \approx \sum_{i=1}^{|\mathcal{A}_l|} \tilde{d}_{p,i}^2(\delta^i)$ , where  $\tilde{d}_{p,i}^1, \tilde{d}_{p,i}^2$  are (possibly) non-linear univariate functions.

### Separation via Substitution

To introduce a method for finding the separable approximation, we split the whole class of general-sum games into two disjoint classes: linearly dependent and linearly independent games.

**Definition 4.** A normal-form game  $G$  is called linearly dependent, if for each strategy  $\pi \in \Pi_f$  there exists a constant  $c_\pi \in \mathbb{R}$ , such that for each action  $a_l \in \mathcal{A}_l$  it holds that  $u_l(a_l, \pi) = c_\pi u_f(a_l, \pi)$ .

Note that all zero-sum games are linearly dependent since we can set  $c_\pi = -1$ . For linearly dependent games, we can introduce a new variable  $y_\pi$  for each pure strategy  $\pi \in \Pi_f$  with a linear substitution  $u_f(\delta, \pi) = y_\pi$ . The problem of solving the Dinkelbach subproblem is hence expressed as

$$\max_{\delta \in \Delta_l} \sum_{\pi \in \Pi_f} c_\pi y_\pi [e \circ q](y_\pi) - p [e \circ q](y_\pi) \quad (8)$$

$$y_\pi = u_f(\delta, \pi) = \sum_{i=1}^{|\mathcal{A}_l|} u_f(a_l^i, \pi) \delta^i \quad \forall \pi \in \Pi_f. \quad (9)$$

Both functions  $[e \circ q](y_\pi)$  and  $y_\pi [e \circ q](y_\pi)$  are easily linearizable univariate functions. In case each  $y_\pi$  is divided uniformly into  $K$  segments, the final standard MILP formulation has  $(K - 1) \times |\Pi_f|$  binary variables [Yang *et al.*, 2012]. The following proposition claims that the error goes to zero as  $\epsilon$  goes to zero and  $K$  approaches infinity.

**Proposition 5.** Let  $\delta^*$  be a strategy computed by Algorithm 1 with precision  $\epsilon$  and the subproblems solved via linearization of (8) and (9) with  $K$  segments. Let  $\delta^{QS}$  be a mixed strategy of the leader in QSE. Then for the utility difference

$d = |u_l(\delta^*, QR(\delta^*)) - u_l(\delta^{QS}, QR(\delta^{QS}))|$  it holds that

$$d \leq \epsilon + |A_f| \frac{C_1 + C_2}{4K^2[e \circ q](\underline{u}_f)},$$

where  $C_1 = \bar{u}_l(\bar{u}_f - \underline{u}_f)^2 \max_{x \in [\underline{u}_f, \bar{u}_f]} |[e \circ q]''(x)|$  and  $C_2 = (\bar{u}_f - \underline{u}_f)^2 \max_{x \in [\underline{u}_f, \bar{u}_f], \pi \in \Pi_f} |c_\pi| |[I \cdot (e \circ q)]''(x)|$ .

*Proof.* For any twice differentiable univariate function  $g$ , let  $\bar{g}$  be a uniform piece-wise linearization of  $g$  into  $K$  segments on interval  $[a, b]$ . The maximum difference in values of  $g$  and  $\bar{g}$  is then [Yano *et al.*, 2013]

$$|g(x) - \bar{g}(x)| \leq \bar{C} \frac{(b-a)^2}{8K^2}, \quad (10)$$

where  $\bar{C} = \max_{x \in [a, b]} |g''(x)|$ . The fractional criterion function (5) of the direct formulation of QSE can be expressed as  $QSE(\delta) = u_l(\delta, QR(\delta)) = N(\delta)/D(\delta)$ , where

$$N(\delta) = \sum_{\pi \in \Pi_f} c_\pi [I \cdot (e \circ q)](z_\pi)$$

$$D(\delta) = \sum_{\pi \in \Pi_f} [e \circ q](z_\pi)$$

$$z_\pi = u_f(\delta, \pi) \quad \forall \pi \in \Pi_f,$$

where  $I$  is the identity function. The linearized criterion function is then  $\overline{QSE}(\delta) = \bar{N}(\delta)/\bar{D}(\delta)$ , where

$$\bar{N}(\delta) = \sum_{\pi \in \Pi_f} c_\pi [I \cdot (e \circ q)](z_\pi)$$

$$\bar{D}(\delta) = \sum_{\pi \in \Pi_f} [e \circ q](z_\pi)$$

$$z_\pi = u_f(\delta, \pi) \quad \forall \pi \in \Pi_f.$$

Denote  $\delta^{\overline{QS}}$  the optimal solution of  $\max_{\delta \in \Delta_l} \overline{QSE}(\delta)$ . Then we have

$$d = \left| \frac{N(\delta^{QS})}{D(\delta^{QS})} - \frac{N(\delta^*)}{D(\delta^*)} \right| \leq \left| \frac{N(\delta^{QS})}{D(\delta^{QS})} - \frac{\bar{N}(\delta^{\overline{QS}})}{\bar{D}(\delta^{\overline{QS}})} \right| + \left| \frac{\bar{N}(\delta^{\overline{QS}})}{\bar{D}(\delta^{\overline{QS}})} - \frac{\bar{N}(\delta^*)}{\bar{D}(\delta^*)} \right| + \left| \frac{\bar{N}(\delta^*)}{\bar{D}(\delta^*)} - \frac{N(\delta^*)}{D(\delta^*)} \right|.$$

We bound each of these terms separately. For the first and the last term we use the first part of Lemma 7 of [Yang *et al.*, 2012]. By this Lemma, for each  $\delta \in \Delta_l$ , we have

$$\left| \frac{\bar{N}(\delta)}{\bar{D}(\delta)} - \frac{N(\delta)}{D(\delta)} \right| \leq \frac{1}{\bar{D}(\delta)} (d_D + d_N), \quad (11)$$

where  $d_D = |D(\delta) - \bar{D}(\delta)|$  and  $d_N = |N(\delta) - \bar{N}(\delta)|$ . From the linearization bound (10),

$$d_D \leq |A_f| \frac{(\bar{u}_f - \underline{u}_f)^2}{8K^2} \max_{x \in [\underline{u}_f, \bar{u}_f]} |[e \circ q]''(x)| = \tilde{C}_1$$

$$d_N \leq |A_f| \frac{(\bar{u}_f - \underline{u}_f)^2}{8K^2} \max_{\substack{x \in [\underline{u}_f, \bar{u}_f] \\ \pi \in \Pi_f}} |c_\pi| |[I \cdot (e \circ q)]''(x)| = \tilde{C}_2.$$

Because  $\bar{D}(\delta) \geq [e \circ q](\underline{u}_f)$  and  $QSE(\delta) \leq \bar{u}_l$ , together with (11), we conclude that

$$\left| \frac{\bar{N}(\delta)}{\bar{D}(\delta)} - \frac{N(\delta)}{D(\delta)} \right| \leq \frac{\bar{u}_l \tilde{C}_1 + \tilde{C}_2}{[e \circ q](\underline{u}_f)}. \quad (12)$$

Because  $\delta^{\overline{QS}}$  and  $\delta^{QS}$  are the maxima of  $\overline{QSE}$  and  $QSE$ ,

respectively, the difference in their values is also at most  $(\bar{u}_l \tilde{C}_2 + \tilde{C}_1)/[e \circ q](\underline{u}_f)$ . To see why this is true, for contradiction, assume that  $|\overline{QSE}(x) - QSE(x)| \leq c$ , but without a loss of generality  $\overline{QSE}(\delta^{\overline{QS}}) - QSE(\delta^{QS}) > c$ . By combining both inequalities we have  $QSE(\delta^{\overline{QS}}) > QSE(\delta^{QS})$ ,  $\delta^{QS}$  is hence not a maximum.

For bounding the second term we use the properties of the Dinkelbach formulation. From Algorithm 1, for each  $\delta \in \Delta_l$ , we have  $\bar{N}(\delta) - UB\bar{D}(\delta) \leq 0$  and  $\bar{N}(\delta^*) - LB\bar{D}(\delta^*) \geq 0$ . Therefore, for the optimal strategy  $\delta^{\overline{QS}}$  it holds that  $LB \leq \bar{N}(\delta^{\overline{QS}})/\bar{D}(\delta^{\overline{QS}}) \leq UB$  and hence

$$\left| \frac{\bar{N}(\delta^{\overline{QS}})}{\bar{D}(\delta^{\overline{QS}})} - \frac{\bar{N}(\delta^*)}{\bar{D}(\delta^*)} \right| \leq UB - LB \leq \epsilon. \quad (13)$$

By combining the bounds (12) and (13) we conclude that

$$d \leq \epsilon + 2 \frac{\bar{u}_l \tilde{C}_1 + \tilde{C}_2}{[e \circ q](\underline{u}_f)}. \quad \square$$

In NFGs which are linearly independent, we can take advantage of the fact that all quantal and expected-utility functions are invertible because they are strictly monotone. Let  $g^{-1}$  be an inverse of a function  $g$ . A substitution  $[e \circ q](u_f(\delta, \pi_l)) = y_\pi$  leads to an easily linearizable constraint

$$u_f(\delta, \pi) = \sum_{i=1}^{|\mathcal{A}_l|} u_f(\delta, \pi^i) \delta^i = [e \circ q]^{-1}(y_\pi). \quad (14)$$

With an additional substitution  $u_l(\delta, \pi) = z_\pi$ , the problem of solving the Dinkelbach subproblem is expressed as

$$\max_{\delta \in \Delta_l} \sum_{\pi \in \Pi_f} z_\pi y_\pi - p y_\pi \quad (15)$$

$$[e \circ q]^{-1}(y_\pi) = \sum_{i=1}^{|\mathcal{A}_l|} u_f(a_l^i, \pi) \delta^i \quad \forall \pi \in \Pi_f \quad (16)$$

$$z_\pi = \sum_{i=1}^{|\mathcal{A}_l|} u_l(a_l^i, \pi) \delta^i \quad \forall \pi \in \Pi_f, \quad (17)$$

where  $z_\pi y_\pi$  are  $|\Pi_f|$  bilinear terms<sup>2</sup>. For linearizing the bilinear terms, we can use either McCormick's envelopes [McCormick, 1976], which are fully linear, or the MDT method [Kolodziej *et al.*, 2013]. The size of the MILP formulation for the McCormick's envelopes is  $|\Pi_f|$  real variables and  $(K-1) \times |\Pi_f|$  binary variables. Experimental evaluation of the MDT method showed that while it provides close-to-optimal solutions, the additional binary variables required by the method significantly slow down the computation. We hence focus more on the McCormick's envelopes.

**Proposition 6.** *Let  $\delta^*$  be a strategy computed by Algorithm 1 with precision  $\epsilon$ , the subproblems solved via linearization of (15) – (17) with  $K$  segments and bilinear terms relaxed using McCormick's envelopes. Let  $\delta^{QS}$  be a mixed strategy of the leader in QSE. Then for the utility difference  $d = |u_l(\delta^*, QR(\delta^*)) - u_l(\delta^{QS}, QR(\delta^{QS}))|$  it holds that*

$$d \leq \epsilon + |A_f| \frac{C + M}{4[e \circ q](\underline{u}_f)},$$

<sup>2</sup>Note that if we would use the same substitution as in the case of linearly dependent games, the number of bilinear terms would be  $K \times |\Pi_f|$ , i.e.,  $K$ -times larger.

where  $C = (\bar{u}_l + 1)h^2 \max_{x \in [\bar{u}_f, \underline{u}_f]} |[e \circ q]''(x)|$ ,  $M = (\bar{u}_l - \underline{u}_l)(\bar{u}_f - \underline{u}_f)/4$  and  $h = ([e \circ q](\bar{u}_f) - [e \circ q](\underline{u}_f))/K \max_{x \in [e \circ q](\underline{u}_f), [e \circ q](\bar{u}_f)} |[e \circ q]^{-1}]'(x)|$ .

*Proof.* We follow a reasoning of proof of Proposition 5 up to Equation (11). In linearly independent games, we linearize the composition of inverse functions  $[e \circ q]^{-1}$  in constraint (16), instead of function  $[e \circ q]$ . Therefore, to estimate the differences  $|N(\delta) - \bar{N}(\delta)|$  and  $|D(\delta) - \bar{D}(\delta)|$ , we seek to bound the linearization error  $|[e \circ q](x) - \overline{[e \circ q]^{-1}}^{-1}(x)|$ .

Assume that  $[e \circ q]^{-1}$  is uniformly linearized in points  $x_0, x_1, \dots, x_{K-1}$ , i.e.,  $[e \circ q]^{-1}(x_i) = \overline{[e \circ q]^{-1}}(x_i) = y_i, \forall i \in [K]$ . Every piece-wise linear function is invertible. Since  $[e \circ q](y_i) = x_i = \overline{[e \circ q]^{-1}}^{-1}(y_i)$ , function  $\overline{[e \circ q]^{-1}}^{-1}$  is a linearization of function  $[e \circ q]$ . Function  $[e \circ q]^{-1}$  is not identity in general. Hence, the linearization is non-uniform. Since the distance between any  $x_i, x_{i+1}$  is  $([e \circ q](\bar{u}_f) - [e \circ q](\underline{u}_f))/K$ , the maximum difference in values of any consecutive  $y_i, y_{i+1}$  is bounded by

$$h = \frac{(\bar{e}q - eq)}{K} \max_{x \in [eq, \bar{e}q]} |[e \circ q]^{-1}]'(x)|,$$

where  $eq = [e \circ q](\underline{u}_f)$  and  $\bar{e}q = [e \circ q](\bar{u}_f)$ . Using the bound on a maximum difference between  $[e \circ q](x)$  and  $\overline{[e \circ q]^{-1}}^{-1}(x)$  by [Yano *et al.*, 2013], we have

$$d_D \leq |A_f| \frac{h^2}{8} \max_{x \in [\underline{u}_f, \bar{u}_f]} |[e \circ q]''(x)|. \quad (18)$$

For bounding the difference  $|N(\delta) - \bar{N}(\delta)|$ , we need to take into account also the error arising from relaxing the bilinear terms using McCormick's envelopes. The McCormick's envelope is a convex polytope defined by four linear inequalities [McCormick, 1976]. Using simple algebra, it can be shown that the difference between a bilinear term  $z_\pi y_\pi$  in Equation (15) and its approximation by McCormick's envelope is at most  $M$ , where

$$M = \frac{(\bar{u}_l - \underline{u}_l)(\bar{u}_f - \underline{u}_f)}{4}. \quad (19)$$

Using the bound (19), the difference between  $N(\delta)$  and  $\bar{N}(\delta)$  can be bounded as

$$d_N \leq |A_f| \left( \frac{h^2}{8} \max_{x \in [\underline{u}_f, \bar{u}_f]} |[e \circ q]''(x)| + M \right). \quad (20)$$

We proceed as in Proposition 5, using Equations (11), (13), and a modified Equation (12), which uses bounds (18) and (20). The result follows.  $\square$

Note that because of the relaxation, the bound does not go to zero with increasing  $K$  and decreasing  $\epsilon$ . However, it can be fairly small for games with a higher value of  $[e \circ q](\underline{u}_f)$ .

### Separation via Additive Approximations

Besides substitution, we considered computing separable approximations also through generalized additive models (GAMs) [Hastie, 2017]. GAMs can be used to approximately solve the Dinkelbach subproblem in case the link function of

the model is an identity. We implemented two methods for fitting GAMs in the subproblem: a spline method ( $\tilde{d}_p^1$  and  $\tilde{d}_p^2$  are smoothing splines [Wahba, 1990]) and a deep neural network method ( $\tilde{d}_p^1$  and  $\tilde{d}_p^2$  are constructed from activation functions [Potts, 1999]). Our experimental evaluation showed that piecewise-linear separable approximations of the Dinkelbach subproblem lead to fast computation of a solution of the approximation (especially for linearly independent games), but neither the spline method nor the neural networks provide approximations precise enough for the Dinkelbach method to converge to an optimum.

## 6 Experiments

Finally, we demonstrate practical aspects of proposed algorithms for computing QSE in NFGs. As a benchmark, we use the original formulation solved by gradient ascent (GA). We compare it to the Dinkelbach-type algorithm (DTA) – Algorithm 1 with subproblems solved via substitutional PWLA. All implementations were done in C++17. We used an implementation of the SLSQP GA algorithm in the NLOPT 2.6.1 library for non-linear optimization. A single-threaded IBM CPLEX 12.8 solver carried all MILP computations. The experiments were performed on a 3.2GHz CPU with 16GB RAM. Because the algorithm is domain independent, we used Randomly Generated Games (RGGs) for evaluation.

**Experimental Setting.** We consider six different generators of canonical quantal functions (four convex functions:  $e^{0.4x}$ ,  $e^{0.8x}$ ,  $e^{1.2x}$  and  $-1050/(x+20)$ , and two concave functions:  $\ln x + 12$  and  $\sqrt{x+11}$ ) and three expected-utility functions (one risk-neutral:  $x$ , one loss-attentive:  $x^3/100$  and one loss-averse:  $20/(1+e^{-x})-10$ ). We construct games with action spaces of sizes from  $3,000 \times 50$  up to  $7,500 \times 50$ . The tolerance parameter for the GA algorithm in NLOPT was set to  $10^{-6}$  (equivalent to the default CPLEX tolerance parameter) and  $\epsilon = 2\%$  of the leader's utility range for the DTA's binary search. The linearization uses  $K = 5$ . For each combination of game size  $\times$  generator function  $\times$  expected-utility function, we constructed and solved 5 instances. The utilities of the follower were generated uniformly randomly from interval  $[-10, 10]$ , while the constants  $c_\pi$  for linearly dependent RGGs were generated from interval  $[-3, 3]$ .

**Runtimes.** In Figure 2 we present the speedup results achieved with the DTA in linearly dependent RGGs. The x-axis shows the game size, while the y-axis varies the mean speedup in the game. We measure the speedup as a ratio between runtime of the DTA and runtime of one restart of the GA. Every point in the graph corresponds to the mean over the sampled instances and also shows the achieved standard error. The results show that with the increasing size of the game, the speedup of the DTA also increases, being up to 25.5-times faster than one restart of GA for games with 7500 leader's actions. It suggests that for even larger games, the DTA should perform significantly better than the GA.

**Solution Quality.** The relative errors of computed solutions are presented in Figure 3. The x-axis varies the game size. The y-axis shows the mean ratio of the difference in the defender's expected utility computed using one restart of the




Figure 2: The mean risk-neutral speedups (left), mean loss-attentive speedups (middle) and mean loss-averse speedups (right) in linearly dependent Randomly Generated Games for different sizes of games. Every point shows also a standard error.




Figure 3: The mean risk-neutral deviations (left), mean loss-attentive deviations (middle) and mean loss-averse deviations (right) in linearly dependent Randomly Generated Games for different sizes of games. Every point shows also a standard error.

size	Risk-neutral restarts			Loss-attentive restarts			Loss-averse restarts		
	ln	sqrt	hyp	ln	sqrt	hyp	ln	sqrt	hyp
3000	7.60 ± 3.68	5.40 ± 3.30	1.80 ± 0.71	1.00 ± 0.00	4.80 ± 3.39	1.00 ± 0.00	19.00 ± 1.15	6.25 ± 4.58	1.60 ± 0.35
4000	1.00 ± 0.00	6.25 ± 4.60	1.40 ± 0.35	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	5.75 ± 4.75	13.60 ± 7.31	1.00 ± 0.00
5000	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	4.00 ± 2.38	2.30 ± 1.53	1.20 ± 0.17

Table 1: The expected number of GA restarts needed to reach 1% deviation from the DTA solution. All *exp* functions required > 20 restarts.

GA and the DTA to the absolute difference between the maximal and minimal utility in the game. While the DTA is guaranteed to approximate a global solution, there is no guarantee of a solution quality for the GA. The results show that the GA never finds a better solution than the DTA (otherwise the ratio would be negative). However, the GA’s solution quality is comparable to the DTA for concave functions and the hyperbole. The quality degrades rapidly for the exponential functions for all three expected-utility functions, especially for higher coefficients. For the loss-averse function, the quality of the GA’s solution with exponential quantal function with coefficient 1.2 is on average at least 42.5-times worse than the solution of the DTA.

**GA Restarts.** The number of restarts of the GA required to reach a deviation from the DTA’s solution less than 1% is shown in Table 1. We restarted the GA 20 times and counted the first restart when the deviation was sufficiently small. None of the exponential functions was able to find a good-enough solution within the 20 restarts. This suggests that to obtain comparable results with the most commonly used quantal function – the exponential function with a higher coefficient – in games with at least 7000 actions of the leader, the GA has to run at least 153× longer than the DTA. In contrast, the concave functions and the hyperbole computed nearly-optimal solutions within the first few restarts. This result is consistent with the deviations reported in Figure 3.

Due to space constraints, we do not report the results achieved with linearly independent games and the GA solving the Dinkelbach subproblem. Because of McCormick’s envelopes, the DTA achieved even faster speedups: up to 3030-

times faster than one restart of GA for exponential functions and games with at least 7000 actions. However, due to the relaxation, the quality of the solutions was often worse (but no more than 12% for the exponential functions and 6% for other quantal functions) than GA’s solutions. The GA running on the Dinkelbach subproblem has the theoretical guarantees (Proposition 4), but its runtimes are comparable to GA on the direct formulation.

### 7 Conclusion

We introduced a Dinkelbach-type formulation of computing a boundedly-rational quantal Stackelberg equilibrium in normal-form games. In contrast to the direct formulation, the Dinkelbach formulation has both the theoretical advantages (i.e., we can explicitly check how closely we approximate the global optimum) as well as positive computational consequences – the formulation offers up to 25.5-times speedup when compared with the original formulation.

### Acknowledgments

This research is supported by the SIMTech-NTU Joint Laboratory on Complex Systems, the Czech Science Foundation (grant no. 18-27483Y and 19-24384Y) and by the OP VVV MEYS funded project CZ.02.1.01/0.0/0.0/16 019/0000765 “Research Center for Informatics”. Bo An is partially supported by Singtel Cognitive and Artificial Intelligence Lab for Enterprises (SCALE@NTU), which is a collaboration between Singapore Telecommunications Limited (Singtel) and Nanyang Technological University (NTU) that is funded by the Singapore Government through the Industry Alignment Fund – Industry Collaboration Projects Grant.

## References

- [Boyd and Vandenberghe, 2004] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [Brown and Sandholm, 2018] Noam Brown and Tuomas Sandholm. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- [Camerer, 2011] Colin F Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2011.
- [Delle Fave *et al.*, 2014] Francesco Maria Delle Fave, Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe, Sarit Kraus, and John P Sullivan. Game-theoretic patrolling with dynamic execution uncertainty and a case study on a real transit system. *Journal of Artificial Intelligence Research*, 50:321–367, 2014.
- [Dinkelbach, 1967] Werner Dinkelbach. On nonlinear fractional programming. *Management Science*, 13(7):492–498, 1967.
- [Fang *et al.*, 2017] Fei Fang, Thanh H Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Brian C Schwedock, Milind Tambe, and Andrew Lemieux. PAWS - a deployed game-theoretic application to combat poaching. *AI Magazine*, 2017.
- [Hastie, 2017] Trevor J Hastie. Generalized additive models. In *Statistical models in S*, pages 249–307. Routledge, 2017.
- [Ibaraki, 1976] Toshimide Ibaraki. Integer programming formulation of combinatorial optimization problems. *Discrete Mathematics*, 16(1):39–52, 1976.
- [Kahneman and Tversky, 2013] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, pages 99–127. World Scientific, 2013.
- [Kolodziej *et al.*, 2013] Scott Kolodziej, Pedro M Castro, and Ignacio E Grossmann. Global optimization of bilinear programs with a multiparametric disaggregation technique. *Journal of Global Optimization*, 57(4):1039–1063, 2013.
- [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, 1976] Daniel L. McFadden. Quantal choice analysis: A survey. In *Annals of Economic and Social Measurement, Volume 5, number 4*, pages 363–390. NBER, 1976.
- [McKelvey and Palfrey, 1995] Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.
- [Moravčík *et al.*, 2017] Matěj Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. DeepStack: Expert-level artificial intelligence in no-limit poker. *Science*, 2017.
- [Nguyen *et al.*, 2016] Thanh H Nguyen, Arunesh Sinha, Shahrzad Gholami, Andrew Plumtre, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, et al. Capture: A new predictive anti-poaching tool for wildlife protection. In *Proceedings of the 15th International Conference on Autonomous Agents & Multiagent Systems*, pages 767–775, 2016.
- [Nudelman *et al.*, 2004] Eugene Nudelman, Jennifer Wortman, Yoav Shoham, and Kevin Leyton-Brown. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 4, pages 880–887, 2004.
- [Pita *et al.*, 2008] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 125–132, 2008.
- [Potts, 1999] William JE Potts. Generalized additive neural networks. In *Proceedings of the 5th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 194–200. ACM, 1999.
- [Tambe, 2011] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, New York, NY, USA, 2011.
- [Vielma *et al.*, 2010] Juan Pablo Vielma, Shabbir Ahmed, and George Nemhauser. Mixed-integer models for non-separable piecewise-linear optimization: Unifying framework and extensions. *Operations Research*, 58(2):303–315, 2010.
- [Wahba, 1990] Grace Wahba. *Spline Models for Observational Data*, volume 59. SIAM, 1990.
- [Yang *et al.*, 2012] Rong Yang, Fernando Ordonez, and Milind Tambe. Computing optimal strategy against quantal response in security games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 847–854, 2012.
- [Yano *et al.*, 2013] Masayuki Yano, James Douglass Penn, George Konidaris, and Anthony T Patera. *Math, Numerics & Programming (for Mechanical Engineers)*. MIT, 2013.
- [Yechiam and Hochman, 2013] Eldad Yechiam and Guy Hochman. Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, 139(2):497, 2013.