

Pessimistic Leader-Follower Equilibria with Multiple Followers

Paper ID: 2378

Abstract

The problem of computing the strategy to commit to has been widely investigated in the scientific literature for the case where a single follower is present. In the multi-follower setting though, results are only sporadic. In this paper, we address the multi-follower case for normal-form games, assuming that, after observing the leader's commitment, the followers play pure strategies and reach a Nash equilibrium. We focus on the pessimistic case where, among many equilibria, one minimizing the leader's utility is chosen (the opposite case is computationally trivial). We show that the problem is NP-hard even with only two followers, and propose an exact exponential-time algorithm which, for any number of followers, finds either an equilibrium when it exists (the leader's utility admitting a maximum) or, if not, an α -approximation of the supremum, for any $\alpha > 0$.

1 Introduction

In recent years, *Leader-Follower* (or *Stackelberg*) *Games* (LFG) and their corresponding *Equilibria* (LFE) have attracted a growing interest in Artificial Intelligence. Considering, for simplicity, the two-player case, LFGs describe situations where one player (*leader*) commits to a *strategy* and the other player (*follower*) first observes the leader's commitment and, then, decides how to play. This is the case of, e.g., security games, where a defender (leader) is tasked to allocate scarce resources to protect valuable targets from an attacker (follower). See [Paruchuri *et al.*, 2008; Kiekintveld *et al.*, 2009; An *et al.*, 2011] for real-world applications of such games. Outside the security domain, applications can be found in, e.g., interdiction games [Caprara *et al.*, 2016; Matuschke *et al.*, 2017], toll-setting problems [Labbé and Violin, 2016], and network routing [Amaldi *et al.*, 2013].

While the majority of the game theoretical investigations on computing LFE have focused on the single-follower case, in this work we address the multi-follower scenario. When facing a multi-follower LFG, two aspects need to be considered: what game the followers play and, in it, how ties among multiple equilibria are broken. As to the type of game, the followers may play, e.g., hierarchically one at a time [Conitzer

and Sandholm, 2006], simultaneously in a correlated manner [Conitzer and Korzhyk, 2011], or simultaneously and noncooperatively [Basilico *et al.*, 2016]. As to breaking ties, two cases are possible: the *optimistic* one, where an equilibrium maximizing the leader's utility is chosen, and the *pessimistic* case, where the followers select an equilibrium which minimizes the leader's utility. While, in the optimistic case, a LFE always exists, in the pessimistic setting it may not exist, since the leader's utility function may admit a supremum but not a maximum [von Stengel and Zamir, 2010].

We focus on multi-follower LFGs in normal-form where the followers play *simultaneously* and *noncooperatively*, reaching a Nash equilibrium. We refer to the corresponding equilibrium as *Leader-Follower Nash Equilibrium* (LFNE).

Related Work. In single-follower games ($n = 2$ players), the follower's best-response to the leader's commitment is always (disregarding degenerate cases) a pure strategy. In the optimistic case, a LFE is found in polynomial time by solving a Linear Programming (LP) problem for each action of the follower [Conitzer and Sandholm, 2006]. Although the pessimistic case is more involved, it can still be solved with similar techniques [von Stengel and Zamir, 2010].

In the multi-follower case ($n \geq 3$), see [Basilico *et al.*, 2016], the problem is NP-hard and inapproximable, in polynomial time, to within any polynomial factor unless $P = NP$. While the optimistic version can be solved to global optimality with global optimization techniques (spatial branch-and-bound) up to a finite precision [Basilico *et al.*, 2016], no exact algorithms are known for the pessimistic case. When the followers are restricted to pure strategies though, the optimistic problem is in P , as discussed in the next section.

Original Contributions. We address the pessimistic multi-follower case, with followers playing pure strategies. This is the case of many applications such as potential games [Monderer and Shapley, 1996], congestion games [Rosenthal, 1973], and toll-setting problems [Labbé and Violin, 2016].

Differently from the optimistic/pessimistic case with a single follower and the optimistic case with many followers, which are all in P , we show that computing a LFNE in the pessimistic case (P-LFNE) with at least two followers playing pure strategies is NP-hard. This problem is crucial, as the difference in leader's utility between an optimistic LFE and a pessimistic one can be arbitrarily large, as we show in Section 2. We then propose an exact exponential-time algo-

rithm (a multi-lex-MILP) for computing a P-LFNE and, then, extend it to a branch-and-bound algorithm.

The status, in terms of complexity and known algorithms, of the problem of computing a LFNE, updated with the original results provided in this paper (reported in bold), is:

		Optimistic		Pessimistic	
		Pure	Mixed	Pure	Mixed
$n = 2$	Comp	P		P	
	Alg	multi-LP		multi-LP	
$n \geq 3$	Comp	P	NP-hard, \notin Poly-APX	NP-hard	NP-hard, \notin Poly-APX
	Alg	multi-LP	Global optimization	multi-lex-MILP	-

2 Problem Definition

Let $N = \{1, \dots, n\}$ be the set of players and, for each $p \in N$, let A_p be the corresponding set of actions, with $m_p = |A_p|$. For each player $p \in N$, let $x_p \in [0, 1]^{m_p}$, with $\sum_{a \in A_p} x_p^a = 1$, be her *strategy vector* (or strategy, for short), where each component x_p^a represents the probability that action $a \in A_p$ is played. For each $p \in N$, let also $\Delta_p = \{x_p \in [0, 1]^{m_p} : \sum_{a \in A_p} x_p^a = 1\}$ be the set of her strategies, called *strategy space*. A strategy is called *pure* when only one action is played with positive probability, and *mixed* otherwise. We denote the collection of strategies of the different players, or *strategy profile*, by $x = (x_1, \dots, x_n)$. When all strategies are pure, we denote the collection of actions played by the players, or *action profile*, by $a = (a_1, \dots, a_n)$.

Consider a *normal-form game* [Shoham and Leyton-Brown, 2008] where $U_p \in \mathbb{Q}^{m_1 \times \dots \times m_n}$ represents, for each player $p \in N$, her (multidimensional) utility (or payoff) matrix. Each of its components $U_p^{a_1, \dots, a_n}$ corresponds to the utility obtained by player p when all the players play actions a_1, \dots, a_n . For a strategy profile x , the expected utility of player $p \in N$ is the n th-degree polynomial $\sum_{a_1 \in A_1} \dots \sum_{a_n \in A_n} U_p^{a_1, \dots, a_n} x_1^{a_1} \dots x_n^{a_n}$.

An action profile $a = (a_1, \dots, a_n)$ is called a *pure Nash Equilibrium* (NE) if $U_p^{a_1, \dots, a_n} \geq U_p^{a'_1, \dots, a'_n}$ for each player $p \in N$ and action profile a' such that, for all $q \in N \setminus \{p\}$, $a'_q = a_q$ and, possibly, $a'_p \neq a_p$. This implies that, for each $p \in N$, if the players in $N \setminus \{p\}$ play as the equilibrium prescribes, player p cannot improve her utility by deviating from the equilibrium and playing some other action $a'_p \neq a_p$. More generally, a *mixed NE* is a strategy profile $x = (x_1, \dots, x_n)$ such that no player $p \in N$ can improve her utility playing a strategy $x'_p \neq x_p$, assuming the others play as the equilibrium prescribes. Observe that, in a normal-form game, a mixed NE always exists [Nash, 1951], while a pure NE may not.

In this paper, we are concerned with the computation of an equilibrium in a normal-form game where the followers, after observing the leader's commitment to a mixed strategy, play a pure NE in the resulting game, so as to minimize the leader's utility. We refer to it as *Pessimistic Leader-Follower Nash Equilibrium* (P-LFNE). For the sake of presentation, we consider the case of two followers (thus, with $n = 3$), where the n th player (player 3) takes on the role of leader. We also assume, w.l.o.g., $m_1 = m_2 = m_3 = m$. Our results can be easily generalized to the case of any $n > 3$.

Computing a P-LFNE amounts to solving the following

bilevel problem:

$$\sup_{x_3 \in \Delta_3} \min_{\substack{(i^*, j^*) \in \\ A_1 \times A_2}} \sum_{k \in A_3} U_3^{i^* j^* k} x_3^k : i^* \in \operatorname{argmax}_{i \in A_1} \left\{ \sum_{k \in A_3} U_1^{i j^* k} x_3^k \right\}$$

$$j^* \in \operatorname{argmax}_{j \in A_2} \left\{ \sum_{k \in A_3} U_2^{i^* j k} x_3^k \right\}.$$

For any $x_3 \in \Delta_3$ determined in the first level, the second level calls for a pair of followers' actions (i^*, j^*) forming a pure NE in the followers' game induced by x_3 which minimizes the leader's utility. Notice that, by definition, (i^*, j^*) is a NE for the given x_3 if and only if i^* maximizes the utility of follower 1 when assuming that follower 2 would play j^* , and j^* maximizes the utility of follower 2 when assuming that follower 1 would play i^* .

In the following, we write the above problem as $\sup_{x_3 \in \Delta_3} f(x_3)$, where f is the leader's utility in the pessimistic case. At any x_3 where the followers' game does not admit a pure NE we define $f(x_3) = -\infty$. Even if $\sup_{x_3 \in \Delta_3} f(x_3) < \infty$ due to the boundedness of f , f may not admit a maximum and, thus, the game may not admit a P-LFNE [von Stengel and Zamir, 2010].

The optimistic version of the problem (i.e., the O-LFNE finding problem), obtained by changing the second level objective from min to max, can be solved in polynomial time. This is because the formulation in [Basilico *et al.*, 2016] becomes an LP when the followers' pure strategies are given, and the set of pure strategies of a normal-form game can be enumerated in polynomial time. Note that a P-LFNE cannot be obtained by just perturbing the leader's strategy in an O-LFNE, as shown in Figure 1.




Figure 1: A game with $n = 3$, $A_1 = \{i_1, i_2\}$, $A_2 = \{j_1, j_2\}$, $A_3 = \{k_1, k_2\}$, parameterized by $\mu > 1$. The leader's utility is plotted as a function of $\rho \in [0, 1]$, with $x_3 = (1 - \rho, \rho)$. The unique O-LFNE is achieved at $\rho = \frac{1}{2}$ (utility 2μ) and the unique P-LFNE at $\rho = 1$ (utility μ), with an arbitrarily large (for $\mu \rightarrow \infty$) difference in utility. After a perturbation $\epsilon \rightarrow 0$ of x_3 in the O-LFNE, we obtain a utility of $\simeq (1 + \mu)/2$, whose difference w.r.t. the P-LFNE is arbitrarily large for $\mu \rightarrow \infty$.

3 Computational Complexity

In this section, we show that computing a P-LFNE is NP-hard. In decision form, the problem reads:

Definition 1 (P-LFNE-d). *Given a normal-form game with $n \geq 3$ players and a finite K , is there a P-LFNE where the leader achieves a utility $\geq K$?*

We show the hardness of the P-LFNE searching problem by a polynomial-time reduction of IND-SET, one of Karp's original 21 NP-complete problems [Karp, 1972], to P-LFNE-d. In decision form, IND-SET reads:

Definition 2 (IND-SET). Given an undirected graph $G = (V, E)$ and a finite J with $J \leq |V|$, does G contain an independent set (a subset of vertices $V' \subseteq V : \forall u, v \in V', \{u, v\} \notin E$) of size $\geq J$?

Before presenting our reduction, we introduce the following class of normal-form games:

Definition 3. Given two reals b and c , with $1 > c > b > 0$, and an integer $m \geq 1$, let $\Gamma_b^c(m)$ be a class of normal-form games with three players ($n = 3$), the first two having $m + 1$ actions each, with action sets $A_1 = A_2 = A = \{1, \dots, m, f\}$, and the third one having m actions, with action set $A_3 = A \setminus \{f\}$. For every third player's action $k \in A \setminus \{f\}$, the other players play a game where:

- the payoffs on the main diagonal (where both players play the same action) satisfy $U_1^{kkk} = U_2^{kkk} = 1, U_1^{ffk} = c, U_2^{ffk} = b$ and, $\forall i \in A \setminus \{k, f\}, U_1^{iik} = U_2^{iik} = 0$;
- $\forall i, j \in A \setminus \{f\}$ with $i \neq j, U_1^{ijk} = U_2^{ijk} = b$;
- $\forall j \in A \setminus \{f\}, U_1^{fjk} = c$ and $U_2^{fjk} = 0$;
- $\forall i \in A \setminus \{f\}, U_1^{ifk} = 1$ and $U_2^{ifk} = 0$.

No restrictions are imposed on the third player's payoffs.

These games are of interest for the following reason:

Proposition 1. In a $\Gamma_b^c(m)$ game with $c \leq \frac{1}{m}$, for all $S \subseteq \{(i, i) : i \in A \setminus \{f\}\}$ with $S \neq \emptyset$ there is a leader's strategy $x_3 \in \Delta_3$ such that the outcomes $(i, i) \in S$ are the only pure NE in the resulting followers' game.

Proof. First, observe that the followers' payoffs not on the main diagonal do not depend on the leader's strategy x_3 . Thus, outcomes (i, j) , for any $i, j \in A \setminus \{f\}$ with $i \neq j$, cannot be NE, as the first follower would deviate by playing action f , obtaining a utility $c > b$. Analogously, outcomes (f, j) , with $j \in A \setminus \{f\}$, cannot be NE because the second follower would deviate by playing f (since $b > 0$). The same holds for outcomes (i, f) with $i \in A \setminus \{f\}$, since the second follower would be better off playing another action (as $b > 0$). The last outcome on the diagonal, (f, f) , cannot be a NE either, as the first follower would deviate from it (as she would get c in it, while she can obtain $1 > c$).

As a result, the only outcomes which can be pure NE are those in $\{(i, i) : i \in A \setminus \{f\}\}$. Clearly, when the leader plays a pure strategy k , the unique pure NE in the followers' game is (k, k) as, due to providing the followers with their maximum payoff, they would not deviate from it. Outcomes (i, i) with $i \in A \setminus \{f, k\}$ are not NE as, with them, the first follower would get $0 < c$. In general, if the leader plays an arbitrary mixed strategy $x_3 \in \Delta_3$, the resulting followers' game is such that the payoffs in (k, k) , with $k \in A \setminus \{f\}$, are (x_3^k, x_3^k) . Noticing that (k, k) is an equilibrium if and only if $x_3^k \geq c$ (as, otherwise, the first follower would deviate by playing action f), we conclude that the set of pure NE in the followers' game is $S = \{(k, k) : x_3^k \geq c\}$.

In order to guarantee that for every possible $S \subseteq \{(i, i) : i \in A \setminus \{f\}\}$ s.t. $S \neq \emptyset$ there is a leader's strategy such that S contains all the pure NE, we must allow the diagonal outcomes to be all (simultaneously) equilibria by properly choosing the value of c . This is done by imposing that,

when the leader plays $x_3 = (\frac{1}{m}, \frac{1}{m}, \dots, \frac{1}{m})$, all outcomes in $\{(i, i) : i \in A \setminus \{f\}\}$ are NE, which is obtained by selecting $c \leq \frac{1}{m}$. \square

Notice that, in a $\Gamma_b^c(m)$ game, the followers' game always admits a pure NE for any leader's commitment x_3 . See Figure 2 for an example.




Figure 2: A $\Gamma_b^c(m)$ game with $m = 3$ and $c \leq \frac{1}{m}$. The leader's strategy space Δ_3 is partitioned into $2^m - 1$ regions, one per subset of $\{(i, i) : i \in A \setminus \{f\}\}$ (the three NE in the followers' game, $(1, 1)$, $(2, 2)$, and $(3, 3)$, are labeled A, B, C).

Theorem 1. P-LFNE-d is strongly NP-complete even for $n = 3$.

Proof. Mapping. Given an instance of IND-SET, i.e., an undirected graph $G = (V, E)$ and a positive integer J , we construct $\Gamma(G)$, a special instance of P-LFNE-d of class $\Gamma_b^c(m)$, as follows. Assuming an arbitrary labeling of the vertices $\{v_1, v_2, \dots, v_m\}$, let $\Gamma(G)$ be an instance of $\Gamma_b^c(m)$ with $c < \frac{1}{(m+1)^2} \leq \frac{1}{m}$ and $0 < b < c < 1$, where each action $i \in A \setminus \{f\}$ is associated with vertex v_i . In compliance with Definition 3, in which no constraints are specified for the leader's payoffs, we define:

- for any pair of vertices $v_i, v_j \in V, U_3^{ijj} = U_3^{jjj} = -\frac{1}{c} - 1$ if $\{v_i, v_j\} \in E$, and $U_3^{ijj} = U_3^{jjj} = 1$ otherwise;
- $\forall k \in A \setminus \{f\}, U_3^{kkk} = 0$ and $U_3^{ffk} = 0$;
- $\forall i \neq j \in A$ and $k \in A \setminus \{f\}, U_3^{ijk} = U_3^{jik} = 0$.

Finally, let $K = \frac{J-1}{J}$. Clearly, this transformation can be performed in time polynomial in the number of vertices m .

If. We show that, if the graph G contains an independent set of size $\geq J$, then $\Gamma(G)$ has a P-LFNE with leader's utility $\geq K$. Let V^* be an independent set with $|V^*| = J$. Consider the case in which outcomes (i, i) , with $v_i \in V^*$, are the only pure NE in the followers' game, and suppose the leader plays x_3 with $x_3^k = \frac{1}{|V^*|}$ if $v_k \in V^*$ and $x_3^k = 0$ otherwise. Since $U_3^{iik} = 1$ for all $k \in A \setminus \{f, i\}$, the leader's utility at any equilibrium (i, i) is:

$$\sum_{k \in A \setminus \{f\}} U_3^{iik} x_3^k = \sum_{k \in A \setminus \{f, i\}} U_3^{iik} x_3^k = \sum_{k \in A \setminus \{f, i\}} x_3^k = \frac{|V^*| - 1}{|V^*|} = \frac{J - 1}{J} = K.$$

Only if. We show that, if $\Gamma(G)$ has a P-LFNE with leader's utility $\geq K$, then G has an independent set of size $\geq J$. Due to Proposition 1, at any P-LFNE the leader plays a strategy \bar{x}_3 inducing a set of pure NE in the followers' game $S^* = \{(k, k) : \bar{x}_3^k \geq c\}$. We now show that, in a P-LFNE, the leader never plays two actions $i, j \in A \setminus \{f\}$ with probability $\geq c$ if $\{v_i, v_j\} \in E$. By contradiction, suppose the leader's equilibrium strategy \bar{x}_3 is such that $\bar{x}_3^i, \bar{x}_3^j \geq c$. The leader's

utility when the followers play the equilibrium (i, i) (the same holds for (j, j)) is:

$$\sum_{k \in A \setminus \{f\}} U_3^{iik} \bar{x}_3^k = \sum_{k \in A \setminus \{f, i, j\}} U_3^{iik} \bar{x}_3^k + \bar{x}_3^j \left(-\frac{1}{c} - 1\right).$$

In the right-hand side, the first term is < 1 (as the leader's payoffs are ≤ 1 and $\sum_{k \in A \setminus \{f, i, j\}} \bar{x}_3^k = 1 - \bar{x}_3^i - \bar{x}_3^j < 1$, since $\bar{x}_3^i, \bar{x}_3^j \geq c$). The second term is $\leq c(-\frac{1}{c} - 1) = -1 - c$ (as $\bar{x}_3^j \geq c$), which is < -1 . It follows that, since (i, i) (or, equivalently, (j, j)) always provides the leader with a < 0 utility, she would never play \bar{x}_3 in an equilibrium. This is because, by playing a pure strategy, she would obtain a utility of at least zero (as, when she plays a pure strategy, the followers' game admits a unique pure NE giving her a zero payoff). As a result, for any action k such that $\bar{x}_3^k \geq c$, we have $U_3^{kkk} = 0$, and $U_3^{iik} = 1$ for every i such that $\bar{x}_3^i \geq c$ (since v_i and v_k are not connected by an edge).

Assume the leader either plays an action with probability $\geq c$ or she does not play it at all. If this is the case, then the leader's utility at an equilibrium $(k, k) \in S^*$ is $1 - \bar{x}_3^k$. Since, due to the pessimistic assumption, the leader maximizes her utility in the worst NE, her best choice is to select an \bar{x}_3 such that all NE yield the same utility, that is: $\bar{x}_3^i = \bar{x}_3^j$ for every i, j such that $(i, i), (j, j) \in S^*$. This results in the leader playing all actions k such that $(k, k) \in S^*$ with the same probability $\bar{x}_3^k = \frac{1}{|S^*|}$, obtaining a utility of $\frac{|S^*| - 1}{|S^*|} = K$. Therefore, the vertices $\{v_k : (k, k) \in S^*\}$ must form an independent set of G , whose size is $|S^*| = J$.

We now show that the previous assumption always holds, i.e., that in a P-LFNE the leader is not better off playing any action with probability less than c . Indeed, this may not happen in presence of isolated vertices. Assume that v_i is isolated and that the leader plays actions $A \setminus \{i, f\}$ with probability $c^- = c - \epsilon$, for some $\epsilon > 0$, and action i with probability $1 - c^-(m - 1)$. Since the latter is greater than c by definition, the unique NE for the followers is (i, i) , providing the leader with a utility of $c^-(m - 1)$, which approaches $\frac{m-1}{m}$ for $c^- \rightarrow \frac{1}{m}$. In general, let us denote by ℓ the number of isolated vertices in G , and assume that the other $m - \ell \geq 2$ vertices form a complete graph. This represents the worst case possible as, for it, the leader cannot get a utility larger than $\frac{\ell}{\ell+1}$ without playing some actions with probability smaller than c , but, at the same time, she could get more by uniformly playing the ℓ actions associated with the isolated vertices, each with probability $\alpha = \frac{1 - (m - \ell)c^-}{\ell}$, while playing with probability c^- the other $m - \ell$ actions. If this is the case, the leader's utility is:

$$(\ell - 1)\alpha + (m - \ell)c^- = \frac{(\ell - 1) + (m - \ell)c^-}{\ell}.$$

Thus, we require $\frac{(\ell - 1) + (m - \ell)c^-}{\ell} < \frac{\ell}{\ell + 1}$ for every $\ell \in \{1, \dots, m - 2\}$, which implies that c must satisfy:

$$(\ell + 1)(\ell - 1) + (\ell + 1)(m - \ell)c < \ell^2 \Leftrightarrow c < \frac{1}{(\ell + 1)(m - \ell)},$$

in which we upper bounded c^- by c . Clearly, c satisfies the condition whenever it is less than the minimum value taken by $\frac{1}{(\ell + 1)(m - \ell)}$ for all ℓ . The first order derivative of the bound

is $\frac{(\ell + 1) - (m - \ell)}{(\ell + 1)^2(m - \ell)^2}$, which vanishes at $\ell = \frac{m - 1}{2}$. Given that, according to our definition, $c < \frac{1}{(m + 1)^2} < \frac{4}{(m + 1)^2}$, we obtain that the leader either plays an action with probability at least c or she never plays such action. The reduction is complete.

NP membership. Since we can verify in polynomial time whether a triple (i, j, x_3) is a P-LFNE, we deduce that P-LFNE-d is strongly NP-complete due to IND-SET being strongly NP-complete. \square

4 Enumerative Algorithm

We propose an exact exponential-time algorithm for the computation of $\sup_{x_3 \in \Delta_3} f(x_3)$. If there is no $x_3 \in \Delta_3$ where $f(x_3)$ achieves value $\sup_{x_3 \in \Delta_3} f(x_3)$ (as $f(x_3)$ does not admit a maximum), our algorithm also returns a strategy \hat{x}_3 providing the leader with a utility equal to an α -approximation (in additive sense) of $\sup_{x_3 \in \Delta_3} f(x_3)$, for any $\alpha > 0$.

In the remainder of the paper, we denote the closure of a set $X \subseteq \Delta_3$ relative to $\text{aff}(\Delta_3)$ by \bar{X} , its boundary relative to $\text{aff}(\Delta_3)$ by ∂X , and its complement relative to Δ_3 by X^c .¹

4.1 Computing $\sup_{x_3 \in \Delta_3} f(x_3)$

The key ingredient of our algorithm is what we call *outcome configurations*. We say that, for a given $x_3 \in \Delta_3$, a pair (S^+, S^-) with $S^+ \subseteq A_1 \times A_2$ and $S^- = A_1 \times A_2 \setminus S^+$ is an *outcome configuration* if, in the followers' game induced by x_3 , all pairs of actions $(i, j) \in S^+$ constitute a NE and all pairs $(i, j) \in S^-$ do not.

For any $(i, j) \in A_1 \times A_2$, we introduce $X(i, j)$ as the polytope of all $x_3 \in \Delta_3$ for which (i, j) is a NE, so defined:

$$X(i, j) := \left\{ x_3 \in \Delta_3 : \begin{array}{l} \sum_{k \in A_3} U_1^{ijk} x_3^k \geq \sum_{k \in A_3} U_1^{i'jk} x_3^k \quad \forall i' \in A_1 \\ \sum_{k \in A_3} U_2^{ijk} x_3^k \geq \sum_{k \in A_3} U_2^{ij'k} x_3^k \quad \forall j' \in A_2 \end{array} \right\}.$$

For all $(i, j) \in A_1 \times A_2$, we also introduce the set $X^c(i, j)$ of all $x_3 \in \Delta_3$ for which (i, j) is *not* a NE.

Let $D_1(i, j, i') := \left\{ x_3 \in \Delta_3 : \sum_{k \in A_3} U_1^{ijk} x_3^k < \sum_{k \in A_3} U_1^{i'jk} x_3^k \right\}$ and

$D_2(i, j, j') := \left\{ x_3 \in \Delta_3 : \sum_{k \in A_3} U_2^{ijk} x_3^k < \sum_{k \in A_3} U_2^{ij'k} x_3^k \right\}$. They

correspond to the values of x_3 for which player 1 (resp., 2) would deviate from (i, j) by playing action i' (resp., j'). $X^c(i, j)$ is thus the set:

$$X^c(i, j) := \left(\bigcup_{i' \in A_1 \setminus \{i\}} D_1(i, j, i') \right) \cup \left(\bigcup_{j' \in A_2 \setminus \{j\}} D_2(i, j, j') \right).$$

Notice that $\partial X^c(i, j) \subseteq \partial \Delta_3$, as any other accumulation point of $X^c(i, j)$ satisfies one of the inequalities in $D_1(i, j, i')$ or $D_2(i, j, j')$ as an equation and, thus, it is not in $X^c(i, j)$.

Let now $X(S^+) := \bigcap_{(i, j) \in S^+} X(i, j)$ (a closed polytope) and $X(S^-) := \bigcap_{(i, j) \in S^-} X^c(i, j)$ (the union of non open not closed polytopes). As for $X^c(i, j)$, $X(S^-)$ satisfies $\partial X(S^-) \subseteq \partial \Delta_3$. Its closure $\bar{X}(S^-)$ is obtained by turning the $<$ constraints into \leq in the definition of $D_1(i, j, i')$ and $D_2(i, j, j')$. Finally, let us define $P := \{(S^+, S^-) : S^+ \subseteq 2^{A_1 \times A_2}, S^- = 2^{A_1 \times A_2} \setminus S^+\}$.

¹Here, $\text{aff}(\Delta_3)$ denotes the *affine hull* of Δ_3 , which is, in our case, the hyperplane in \mathbb{R}^m containing Δ_3 .

Theorem 2. Let $\psi(x_3; S^+) := \min_{(i,j) \in S^+} \sum_{k \in A_3} U^{ijk} x_3$. The following holds:

$$\sup_{x_3 \in \Delta_3} f(x_3) = \max_{(S^+, S^-) \in P} \max_{\substack{x_3 \in X(S^+) \\ X(S^+) \cap X(S^-) \neq \emptyset}} \psi(x_3; S^+).$$

Proof. Let $\Delta'_3 := \{x_3 \in \Delta_3 : \text{there exists a pure NE in the followers' game induced by } x_3\}$. Clearly, since $f(x_3) > -\infty$ only on this set, $\sup_{x_3 \in \Delta_3} f(x_3) = \sup_{x_3 \in \Delta'_3} f(x_3)$. Thus, we can focus w.l.o.g. on Δ'_3 .

The sets $X(S^+) \cap X(S^-) \neq \emptyset$, obtained for all $(S^+, S^-) \in P$, form a partition of Δ'_3 . Since, at any $x_3 \in X(S^+) \cap X(S^-)$, the only pure NE induced by x_3 in the followers' game are those in S^+ , we have $f(x_3) = \psi(x_3; S^+)$. Since $\sup_{x_3 \in \Delta'_3} f(x_3)$ is the largest value attained by the supremum of $f(x_3)$ over all sets forming any partition of Δ'_3 , we deduce:

$$\sup_{x_3 \in \Delta_3} f(x_3) = \max_{(S^+, S^-) \in P} \sup_{\substack{x_3 \in X(S^+) \\ X(S^+) \cap X(S^-) \neq \emptyset}} \psi(x_3; S^+).$$

It remains to show that, for all $X(S^+) \cap X(S^-) \neq \emptyset$, $\sup_{x_3 \in X(S^+) \cap X(S^-)} \psi(x_3; S^+) = \max_{x_3 \in X(S^+) \cap X(S^-)} \psi(x_3; S^+)$.

Note that, since $X(S^+) \cap X(S^-)$ is, in general, not a closed set, $\psi(x_3; S^+)$ admits a supremum over it but not necessarily a maximum. Since $\psi(x_3; S^+)$ is continuous (it is the point-wise minimum of finitely many continuous functions), its supremum over $X(S^+) \cap X(S^-)$ equals its maximum over the closure $\overline{X(S^+) \cap X(S^-)}$ of $X(S^+) \cap X(S^-)$. The claim follows since $\overline{X(S^+) \cap X(S^-)} = X(S^+) \cap \overline{X(S^-)}$. \square

Corollary 1. If, according to the formula in Theorem 2, the value of $\sup_{x_3 \in \Delta_3} f(x_3)$ is attained at some $(S^+, S^-) \in P$ for which there exists an $x_3^* \in X(S^+) \cap X(S^-)$ such that $\psi(x_3^*; S^+) = \max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$, then $\sup_{x_3 \in \Delta_3} f(x_3) = \max_{x_3 \in \Delta_3} f(x_3) = f(x_3^*)$.

Proof. Since $x_3^* \in X(S^+) \cap X(S^-)$, the only pure NE in the followers' game induced by x_3^* are those in S^+ . Hence, $f(x_3^*) = \psi(x_3^*; S^+)$ and the result follows. \square

Theorem 3. There is a finite exponential-time algorithm which computes $\sup_{x_3 \in \Delta_3} f(x_3)$ and, whenever $\sup_{x_3 \in \Delta_3} f(x_3) = \max_{x_3 \in \Delta_3} f(x_3)$, also returns a strategy x_3^* with $f(x_3^*) = \max_{x_3 \in \Delta_3} f(x_3)$.

Proof. The algorithm relies on the expression given in Theorem 2. All pairs $(S^+, S^-) \in P$ can be constructed by enumeration in time exponential in m . Given $\epsilon \geq 0$, let:

$$D_1(i, j, i'; \epsilon) := \left\{ x_3 \in \Delta_3 : \sum_{k \in A_3} U_1^{ijk} x_3^k + \epsilon \leq \sum_{k \in A_3} U_1^{i'jk} x_3^k \right\} \text{ and}$$

$$D_2(i, j, j'; \epsilon) := \left\{ x_3 \in \Delta_3 : \sum_{k \in A_3} U_2^{ijk} x_3^k + \epsilon \leq \sum_{k \in A_3} U_2^{ij'k} x_3^k \right\}. \text{ Ac-}$$

cordingly, let:

$$X^c(i, j; \epsilon) := \left(\bigcup_{i' \in A_1 \setminus \{i\}} D_1(i, j, i'; \epsilon) \right) \cup \left(\bigcup_{j' \in A_2 \setminus \{j\}} D_2(i, j, j'; \epsilon) \right)$$

and define $X(S^-; \epsilon) := \bigcap_{(i,j) \in S^-} X^c(i, j; \epsilon)$.

Since, for $\epsilon = 0$, $X(S^-; 0) = \overline{X(S^-)}$, checking whether $X(S^+) \cap X(S^-) \neq \emptyset$ is equivalent to verifying whether there

exists some $\epsilon > 0$ such that $X(S^+) \cap X(S^-; \epsilon) \neq \emptyset$. This can be done by solving:

$$\max_{\epsilon \geq 0, x_3 \in \Delta_3} \{\epsilon : x_3 \in X(S^+) \cap X(S^-; \epsilon)\}.$$

The condition $X(S^+) \cap X(S^-) \neq \emptyset$ is verified if and only if the problem admits a solution (ϵ, x_3) with $\epsilon > 0$. The problem is a *Mixed-Integer Linear Program* (MILP). To see this, it suffices to express each $X^c(i, j; \epsilon)$ as the set of $x_3 \in \Delta_3$ such that there are binary variables $z_{ijj'}$ $\in \{0, 1\}$, for $i' \in A_1 \setminus \{i\}$, and $z_{ijj'}$ $\in \{0, 1\}$, for $j' \in A_2 \setminus \{j\}$, satisfying:

$$\sum_{k \in A_3} U_1^{ijk} x_3^k + \epsilon \leq \sum_{k \in A_3} U_1^{i'jk} x_3^k + M^{ijj'}(1 - z_{ijj'}) \quad \forall i' \in A_1 \setminus \{i\}$$

$$\sum_{k \in A_3} U_1^{ijk} x_3^k + \epsilon \leq \sum_{k \in A_3} U_2^{ij'k} x_3^k + M^{ijj'}(1 - z_{ijj'}) \quad \forall j' \in A_2 \setminus \{j\}$$

$$\sum_{i' \in A_1 \setminus \{i\}} z_{ijj'} + \sum_{j' \in A_2 \setminus \{j\}} z_{ijj'} = 1,$$

with $M^{ijj'} := \max_{k \in A_3} \{U_1^{ijk} - U_1^{i'jk}\}$ and $M^{ijj'} := \max_{k \in A_3} \{U_2^{ijk} - U_2^{ij'k}\}$. $X(S^-; \epsilon)$ is the intersection of such sets, for all $(i, j) \in S^-$.

Finally, observe that $\max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$ is equal to the value that $\psi(x_3; S^+)$ achieves at an optimal solution of a lexicographical maximization problem where, among all solutions which maximize $\psi(x_3; S^+)$, one with the largest ϵ is sought. It reads:

$$\text{lex-max}_{\epsilon \geq 0, x_3 \in X(S^+) \cap X(S^-; \epsilon)} [\psi(x_3; S^+); \epsilon].$$

Due to the original problem being a restriction of the lex one obtained for $\epsilon = 0$, we deduce that optimal solutions to the latter are also optimal for the former. Therefore, $\max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$ can be computed in, at most, exponential time by solving the following lex-MILP:

$$\max_{\substack{\eta \in \mathbb{R}, \epsilon \geq 0 \\ x_3 \in \Delta_3}} \left\{ \begin{array}{l} [\eta; \epsilon] : \eta \leq \sum_{k \in A_3} U_3^{ijk} x_3^k \quad \forall (i, j) \in S^+ \\ x_3 \in X(S^+) \cap X(S^-; \epsilon) \end{array} \right\}.$$

Assume that $\sup_{x_3 \in \Delta_3} f(x_3)$ is attained at some (S^+, S^-) and let (x_3^*, ϵ^*) be the optimal solution of the corresponding lex-MILP. Clearly, $\epsilon^* > 0$ if and only if x_3^* satisfies the assumptions of Corollary 1. Therefore, $f(x_3^*) = \max_{x_3 \in \Delta_3} f(x_3)$. \square

4.2 Finding an α -Approximate Strategy

Assume that, for the strategy x_3^* where $s := \sup_{x_3 \in \Delta_3} f(x_3)$ is attained, $f(x_3^*) < s$ and, thus, the problem does not admit a maximum. Given any $\alpha > 0$, we look for a strategy \hat{x}_3 such that $s - f(\hat{x}_3) \leq \alpha$, i.e., for an α -approximate strategy \hat{x}_3 . Its existence is guaranteed by the following lemma:

Lemma 1. Let $X \subseteq \mathbb{R}^n$, for some $n \in \mathbb{N}$, $Y \subseteq \mathbb{R}$, $f : X \rightarrow Y$, and $s := \sup_{x \in X} f(x)$. Assume $s < \infty$. For any $\alpha > 0$, there exists $x \in X : s - f(x) \leq \alpha$.

Proof. By contradiction, assume there is some $\alpha > 0$ such that, for every $x \in X$, $s - f(x) > \alpha$. Then, for all $x \in X$, $f(x) < s - \alpha$, implying $\sup_{x \in X} f(x) \leq s - \alpha < s$. \square

Theorem 4. Assume that $f(x_3)$ does not admit a maximum over Δ_3 . Then, for any $\alpha > 0$, there is an algorithm which computes an α -approximate strategy \hat{x}_3 in, at most, exponential time.

Proof. According to the formula in Theorem 2, assume $\sup_{x_3 \in \Delta_3}$ is achieved at some (S^+, S^-) and x_3^* , with $\psi(x_3^*, S^+) = \max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$. We deduce $x_3^* \in \overline{X(S^-)} \setminus X(S^-)$ as, otherwise, f would achieve a maximum at x_3^* , in contrast with the assumption in the statement of the theorem.

Aiming at an α -approximate strategy \hat{x}_3 , we can, w.l.o.g., look for a point $\hat{x}_3 \in X(S^+) \cap X(S^-)$ where $f(\hat{x}_3) \geq s - \alpha$, whose existence is guaranteed by Lemma 1. To find it, it suffices to solve the following MILP (which takes, at most, exponential time):

$$\max_{\epsilon \geq 0, x_3 \in \Delta_3} \left\{ \begin{array}{l} \epsilon : \sum_{k \in A_3} U_3^{ijk} x_3^k \geq s - \alpha \quad \forall (i, j) \in S^+ \\ x_3 \in X(S^+) \cap X(S^-, \epsilon) \end{array} \right\},$$

Notice that the problem achieves a solution with $\epsilon > 0$ such that there exists $x_3 \in X(S^+) \cap X(S^-)$, whose existence is guaranteed by the fact that the formula in Theorem 2 only considers pairs $(S^+, S^-) \in P : X(S^+) \cap X(S^-) \neq \emptyset$. \square

5 Branch-and-Bound Algorithm

As it is clear, the computation of $\sup_{x_3 \in \Delta_3} f(x_3)$ as in Theorem 3 is impractical, as it requires the explicit enumeration of all the outcome configurations of a game, many of which, most likely, yield empty regions $X(S^+) \cap X(S^-)$. A more efficient algorithm, albeit still exponential, can be obtained by means of a branch-and-bound technique.

Rather than defining $S^- = (A_1 \times A_2) \setminus S^+$, assume now $S^- \subseteq (A_1 \times A_2) \setminus S^+$. In such case, we call the corresponding pair (S^+, S^-) a *relaxed outcome configuration*. Starting from some pair of followers' actions $(i, j) \in A_1 \times A_2$, the algorithm, through a sequence of branching operations, constructs and explores two search trees whose nodes correspond to relaxed outcome configurations. One tree, which considers the case where (i, j) is a NE, has the pair $(\{(i, j)\}, \emptyset)$ as root, while the other, which considers the case where (i, j) is not a NE, has root $(\emptyset, \{(i, j)\})$.

Notice that, when $S^- \subset (A_1 \times A_2) \setminus S^+$, solving $\max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$ might not give a strategy x_3 such that, if the leader plays it, the only pure NE in the followers' game are those in S^+ , even if $x_3 \in X(S^+) \cap X(S^-)$. This is because, due to $S^+ \cup S^- \subset A_1 \times A_2$, there might be another NE $(i', j') \in A_1 \times A_2 \setminus (S^+ \cup S^-)$ providing the leader with a strictly smaller utility than all the pairs in S^+ . Since, if that is the case, the followers would play (i', j') for that x_3 , $\psi(x_3; S^+)$ is not a valid candidate for being $\sup_{x_3 \in \Delta_3} f(x_3)$.

In order to detect whether one such (i', j') exists, it suffices to carry out a *feasibility check* for x_3 by looking for, in the followers' game, a pure NE different from those in S^- which minimizes the leader's utility. This can be done by inspection in $O(m^2)$. If the feasibility check returns some $(i', j') \notin S^+$, the tree is expanded by performing a *branching* operation. Two nodes are introduced: a *left node* (S_L^+, S_L^-) , where $S_L^+ = S^+ \cup \{(i', j')\}$ and $S_L^- = S^-$ (accounting for the case where (i', j') is a NE), and a *right node* (S_R^+, S_R^-) , where $S_R^+ = S^+$ and $S_R^- = S^- \cup \{(i', j')\}$ (accounting for the case where (i', j') is not a NE). Instead, if $(i', j') \in S^+$, then $\psi(x_3; S^+)$

represents a valid candidate solution, and, thus, (S^+, S^-) is a leaf node. Observe that the feasibility check can safely ignore all $(i, j) \in S^-$ since, whenever (i, j) is a NE for x_3 , due to how the tree is built, there is another node where $(i, j) \in S^+$ (which accounts for the case where (i, j) is a NE).

Due to $S^- \subseteq (A_1 \times A_2) \setminus S^+$, solving $\max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$ for some (S^+, S^-) gives an *upper bound* on the leader's utility under the condition that all the pairs in S^+ constitute a NE and those in S^- do not. Thus, the values obtained by solving $\max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$ can be used as bounds.

However, since $\max_{x_3 \in X(S^+) \cap \overline{X(S^-)}} \psi(x_3; S^+)$ is not well-defined for nodes where $S^+ = \emptyset$, for such nodes we solve, instead, the optimistic problem (see Section 2) with additional MILP constraints (see Section 4) which impose that all the pairs in S^- are not NE.

Finally, suppose that $\sup_{x_3 \in \Delta_3} f(x_3)$ is attained at (S^+, S^-) , and let (x_3^*, ϵ^*) be the optimal solution of the associated lex-MILP problem. If $\epsilon^* > 0$, then $f(x_3^*) = \max_{x_3 \in \Delta_3} f(x_3)$ due to Corollary 1. Otherwise, the problem admits a supremum but not a maximum. In this case, we look for an α -approximate strategy \hat{x}_3 using the algorithm of Theorem 4. However, since (S^+, S^-) is a relaxed outcome configuration here, there might be a pair $(i', j') \in A_1 \times A_2 \setminus (S^+ \cup S^-)$ which is a NE for \hat{x}_3 providing the leader with a utility strictly smaller than $\psi(\hat{x}_3; S^+)$. In order to establish whether such (i', j') exists, we perform a feasibility check for \hat{x}_3 . If we obtain $(i', j') \notin S^+$, we perform a branching operation (adding (i', j') either to S^+ or S^-), and repeat the procedure of Theorem 4 on each of the resulting relaxed outcome configurations, iterating until we find an α -approximate strategy for which the feasibility check returns $(i', j') \in S^+$. Observe that, due to the correctness of the algorithm for the computation of the supremum, at x_3^* there cannot be a NE (i', j') worse than those in S^+ . Therefore, there must be a strategy which is a convex combination of x_3^* and \hat{x}_3 where either (i', j') is not a NE or it yields a leader's utility not worse than that obtained with the NE in S^+ . Thus, the algorithm is guaranteed to converge.

Preliminary results (omitted for reasons of space) on three-player GAMUT instances [Nudelman *et al.*, 2004] show that, with our branch-and-bound algorithm, we can solve games with up to 30 actions within 1000 seconds.

6 Conclusions

In this paper, we have shown that the problem of computing a pessimistic leader-follower equilibrium with multiple followers playing pure strategies simultaneously and noncooperatively is NP-hard even when the number of followers is two. Moreover, we have described an exact exponential-time algorithm for finding an equilibrium strategy (or, if it does not exist, an α -approximate one), which we have then enhanced with a branch-and-bound scheme. Future developments include applications to structured games (e.g., congestion games) and the generalization to the mixed case—even though we conjecture that this problem is much harder, probably Σ_2^P -complete, given its bilevel structure.

References

- [Amaldi *et al.*, 2013] E. Amaldi, A. Capone, S. Coniglio, and L. G. Gianoli. Network optimization problems subject to max-min fair flow allocation. *IEEE COMMUN LETT*, 17(7):1463–1466, 2013.
- [An *et al.*, 2011] B. An, J. Pita, E. Shieh, M. Tambe, C. Kiekintveld, and J. Marecki. Guards and Protect: Next generation applications of security games. *ACM SIGecom Exchanges*, 10(1):31–34, 2011.
- [Basilico *et al.*, 2016] N. Basilico, S. Coniglio, and N. Gatti. Methods for finding leader-follower equilibria with multiple followers: (extended abstract). In *AAMAS*, pages 1363–1364, 2016.
- [Caprara *et al.*, 2016] A. Caprara, M. Carvalho, A. Lodi, and G. J. Woeginger. Bilevel knapsack with interdiction constraints. *INFORMS J COMPUT*, 28(2):319–333, 2016.
- [Conitzer and Korzhyk, 2011] V. Conitzer and D. Korzhyk. Commitment to correlated strategies. In *AAAI*, pages 632–637, 2011.
- [Conitzer and Sandholm, 2006] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *ACM EC*, pages 82–90, 2006.
- [Karp, 1972] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [Kiekintveld *et al.*, 2009] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordóñez, and M. Tambe. Computing optimal randomized resource allocations for massive security games. In *AAMAS*, pages 689–696, 2009.
- [Labbé and Violin, 2016] M. Labbé and A. Violin. Bilevel programming and price setting problems. *ANN OPER RES*, 240(1):141–169, 2016.
- [Matuschke *et al.*, 2017] J. Matuschke, S. T. McCormick, G. Oriolo, B. Peis, and M. Skutella. Protection of flows under targeted attacks. *OPER RES LETT*, 45(1):53–59, 2017.
- [Monderer and Shapley, 1996] D. Monderer and L. S. Shapley. Potential games. *GAME ECON BEHAV*, 14(1):124–143, 1996.
- [Nash, 1951] J. F. Nash. Non-cooperative games. *ANN MAT*, 54(2):286–295, 1951.
- [Nudelman *et al.*, 2004] E. Nudelman, J. Wortman, Y. Shoham, and K. Leyton-Brown. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *AAMAS*, pages 880–887, 2004.
- [Paruchuri *et al.*, 2008] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Playing games for security: an efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS*, pages 895–902, 2008.
- [Rosenthal, 1973] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *INT J GAME THEORY*, 2(1):65–67, 1973.
- [Shoham and Leyton-Brown, 2008] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations*. Cambridge University Press, 2008.
- [von Stengel and Zamir, 2010] B. von Stengel and S. Zamir. Leadership games with convex strategy sets. *GAME ECON BEHAV*, 69:446–457, 2010.