

Computing Correlated Strategies to Commit to with Multiple Leaders*

Matteo Castiglioni
Politecnico di Milano
Milan, Italy
matteo.castiglioni@polimi.it

Alberto Marchesi
Politecnico di Milano
Milan, Italy
alberto.marchesi@polimi.it

Nicola Gatti
Politecnico di Milano
Milan, Italy
nicola.gatti@polimi.it

ABSTRACT

We study the problem of computing *correlated strategies to commit to* in games with *multiple* leaders and followers. To the best of our knowledge, this problem is widely unexplored so far, as the majority of the works in the literature focus on games with a single leader and one or more followers. The fundamental ingredient of our model is that a leader can decide whether to participate in the commitment or to defect from it by taking on the role of follower. This introduces a preliminary stage where, before the underlying game is played, the leaders make their decisions to reach an *agreement* on the correlated strategy to commit to. We distinguish three solution concepts on the basis of the constraints that they enforce on the agreement reached by the leaders. Then, we provide a comprehensive study of the properties of our solution concepts, in terms of existence, relation with other solution concepts, and computational complexity.

1 INTRODUCTION

Over the last years, Stackelberg games are receiving an increasing attention from the algorithmic game theory community, thanks to their many applications in real-world scenarios, such as in security [39]. In the classical Stackelberg setting [40], there is a leader with the ability to play before the other player, who acts as follower by observing the realization of the leader’s strategy. In this work, we follow a different approach, where the leader looks for a *strategy to commit to* [12], and the follower observes the leader’s mixed strategy, without knowing its actual realization. An interpretation of this setting is provided by Von Stengel and Zamir [41]: any (*underlying*) game is extended as a sequential game in which the leader plays first, having a continuum of choices corresponding to mixed-strategy commitments.

The majority of the works in the literature focus on games with a single leader and a single follower [12, 41]. In this setting, the leader seeks for a utility-maximizing mixed strategy to commit to, while the follower plays a best response to the commitment. This model has been largely studied, especially for security problems [1, 21, 33].

Some works also study games with a single leader and multiple followers. Conitzer and Korzhyk [11] introduce a model where the leader commits to a *correlated strategy* and, accordingly, she draws recommendations for the followers, who must obey the incentive constraints of correlated equilibria [4]. The authors show that, in normal-form games, an optimal correlated strategy to

commit to can be computed in polynomial time. Other works study situations where the followers play a *Nash equilibrium* [30] in the game resulting from the leader’s mixed-strategy commitment [9, 10, 13, 28, 29, 41]. However, these models usually result in intractable computational problems even with a fixed number of followers.

Settings including multiple leaders are widely unexplored in the literature. In spite of this, many real-world applications naturally involve more than one player with competitive advantages, playing the role of leader. Some scenarios are, e.g., network platforms with premium (prioritized) users, markets where a group of firms forms a price-determining dominant cartel [15], and political elections in which some candidates choose policy positions in advance of challengers. [2].

Restricted to the security context, there are some works addressing games with multiple uncoordinated defenders (leaders) [17, 23, 25, 26, 38]. However, differently from our work, they all enforce Nash-like constraints on the leaders’ strategies. Moreover, their models suffer from two major drawbacks: (i) an exact equilibrium may not exist, and (ii) they strongly rely on problem-specific structures arising in security problems.

The operations research literature provides further works on multi-leader-follower settings, under the name of *mathematical programs with equilibrium constraints* [27]. They assume that both leaders and followers are subject to Nash constraints, with the latter playing in the game resulting from the leaders’ strategies [22, 24, 31]. Furthermore, other works from the same field focus on oligopoly models where the leaders select the level of investment to maximize profits [14, 35, 37]. All these works considerably depart from ours, as they use fundamentally different models and lack thorough game theoretic and computational studies.

1.1 Original Contributions

We introduce a new way to apply the Stackelberg paradigm to any finite (*underlying*) game. Our approach extends the idea of commitment to *correlated strategies* in settings involving multiple leaders and followers, generalizing the work of Conitzer and Korzhyk [11]. The crucial component of our framework is that a leader can decide whether to participate in the commitment or to defect from it by becoming a follower. This induces a preliminary *agreement stage* that takes place before the underlying game is played, where the leaders decide, in turn, whether to opt out from the commitment or not. We model this stage as a sequential game, whose size is factorial in the number of players. Our goal is to identify commitments guaranteeing some desirable properties that we define on the agreement stage. The first one requires that the leaders do not have any incentive to become followers. It comes in two

*This work has been published in the proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI 2019), see [7].

flavors, called *stability* and *perfect stability*, which are related to, respectively, Nash and subgame perfect equilibria of the sequential game representing the agreement stage. The second property is also defined in two flavors, namely *efficiency* and *perfect efficiency*, both enforcing Pareto optimality with respect to the leaders' utility functions, though at different levels of the agreement stage.

We introduce three solution concepts, which we generally call *Stackelberg correlated equilibria* (SCEs). They differ depending on the properties they call for. Specifically, SCEs, SCEs *with perfect agreement* (SCE-PAs), and SCE-PAs *with perfect efficiency* (SCE-PAPEs) require, respectively, stability and efficiency, perfect stability and efficiency, and both perfect stability and perfect efficiency.

First, we investigate the game theoretic properties of our solution concepts. We show that SCEs and SCE-PAs are guaranteed to exist in any game, while SCE-PAPEs may not. Moreover, we compare them with other solution concepts.

Then, we switch the attention to the computational complexity perspective. We show that, provided a suitably defined *stability oracle* is solvable in polynomial time, an SCE optimizing some linear function of leaders' utilities (such as the leaders' social welfare) can be computed in polynomial time, even in the number of players. The same holds for finding an SCE-PA, while we prove that computing an optimal SCE-PA is an intractable problem. Nevertheless, in the latter case, we provide an (exponential in the game size) upper bound on the necessary number of queries to the oracle.

In conclusion, we study which classes of games admit a polynomial-time stability oracle, focusing on those with *polynomial type* [32]. We show that the problem solved by our oracle is strictly connected with the weighted deviation-adjusted social welfare problem introduced by Jiang and Leyton-Brown [18]. As a result, we get that our oracle is solvable in polynomial time in all the game classes where the same holds for the problem of finding an optimal correlated equilibrium.¹

2 PRELIMINARIES

In this section, we introduce some basic concepts about games and their equilibria used in the rest of the paper.

2.1 Finite Games and Their Equilibria

A (finite) game G is a tuple $(N, \{S_p\}_{p \in N}, \{u_p\}_{p \in N})$, where $N = \{1, \dots, n\}$ is a finite set of players, S_p is a finite set of player p 's strategies or actions, and $u_p : S \rightarrow \mathbb{R}$ is player p 's utility, defined over the set of strategy profiles $S = \times_{p \in N} S_p$. Given $s \in S$, let $s_{-p} \in S_{-p} = \times_{q \in N \setminus \{p\}} S_q$ be the partial profile obtained by dropping player p 's strategy s_p from s , so that $s = (s_p, s_{-p})$. We call $\mathcal{X} = \Delta(S)$ the set of *correlated distributions* defined over strategy profiles, i.e., each $x \in \mathcal{X}$ satisfies $\sum_{s \in S} x(s) = 1$ and $x(s) \geq 0$ for all $s \in S$. Moreover, overloading notation, $u_p(x) = \sum_{s \in S} x(s)u_p(s)$ is player p 's expected utility in $x \in \mathcal{X}$.

A correlated distribution $x \in \mathcal{X}$ is a *correlated equilibrium* (CE) [4] if, for every player $p \in N$ and strategies $s_p \neq s'_p \in S_p$, the following constraint holds:

$$\sum_{s_{-p} \in S_{-p}} x(s_p, s_{-p}) \left(u_p(s_p, s_{-p}) - u_p(s'_p, s_{-p}) \right) \geq 0. \quad (1)$$

¹Full proofs of all the results are in [8].

We can interpret a CE in terms of a mediator who draws some strategy profile $s \in S$ from a publicly known distribution x , and, then, it privately communicates each recommendation s_p to player p . The distribution is an equilibrium if no player has an incentive to deviate from the recommendation, as made formal by the *incentive constraints* of Eq. (1). Moreover, a *Nash equilibrium* (NE) [30] is a CE $x \in \mathcal{X}$ that can be written as the product distribution of players' *mixed strategies*, i.e., $x(s) = \prod_{p \in N} x_p(s_p)$ for all $s \in S$, where each $x_p \in \Delta(S_p)$ is a probability distribution over strategies S_p denoting a player p 's mixed strategy.

In the following, we denote with $\mathcal{X}_P^{\text{CE}} \subseteq \mathcal{X}$ the set of correlated distributions that satisfy the incentive constraints of Eq. (1) only for a subset of players $P \subseteq N$. Clearly, $\mathcal{X}^{\text{CE}} = \mathcal{X}_N^{\text{CE}}$ is the set of CEs of the game.

Different classes of games are employed depending on how strategies and utilities are represented. The most common representation is the *normal form*, which encodes each utility function u_p as an n -dimensional matrix indexed by $s \in S$. Thus, the size of a normal-form game is exponential in the number of players. Many other representations have been introduced in the literature. In Section 7, we are interested in those with *polynomial type* [32], where the number of players and the number of strategies are bounded by polynomials in the size of the game. Many important classes of games admit a polynomial-type representation, such as graphical games [20], polymatrix games [16], anonymous games [5], and congestion games [34].

2.2 Stackelberg Games and Equilibria

Any finite game has a Stackelberg counterpart where some of the players are *leaders* and the others are *followers*. The former have the ability to commit to a course of play beforehand, while the latter decide how to play after observing the commitment [41].

Definition 2.1. Given a finite game G , a *Stackelberg game* (SG) is a tuple (G, L, F) where L and F are the sets of *leaders* and *followers*, respectively, with $N = L \cup F$.

In single-leader single-follower SGs, the follower best responds to the leader's mixed-strategy commitment [12, 41].

Definition 2.2. Given an SG $(G, \{1\}, \{2\})$, a leader's mixed strategy $x_1 \in \Delta(S_1)$ defines a *Stackelberg equilibrium* (SE) if it maximizes u_1 given that, for each $x'_1 \in \Delta(S_1)$, the follower plays an $x_2(x'_1) \in \Delta(S_2)$ maximizing u_2 .²

The multi-follower case unfolds in different scenarios depending on how the followers are assumed to play. Conitzer and Korzhyk [11] study what they call *optimal correlated strategies to commit to*, where the leader commits to a utility-maximizing correlated distribution satisfying the incentive constraints (Eq. (1)) for the followers only. Formally:

Definition 2.3. Given an SG $(G, \{1\}, N \setminus \{1\})$, $x \in \mathcal{X}$ is an *optimal correlated strategy to commit to* if it maximizes the leader's utility $u_1(x)$ over the set $\mathcal{X}_{N \setminus \{1\}}^{\text{CE}}$.

²In the literature, different SE concepts are defined depending on how the follower is assumed to break ties. The *strong* and *weak* SEs are two notable cases [6], where the follower is assumed to break ties in favor and against the leader, respectively.

In our work, we pursue the approach of Conitzer and Korzhlyk [11], rather than letting the followers play an NE, as done, e.g., by Von Stengel and Zamir [41]. Indeed, while the two models provide the same leader’s utility in single-follower SGs (corresponding to that in an SE), the latter may be strictly better in SGs with two or more followers (see [11] for an example).

3 MULTI-LEADER-FOLLOWER STACKELBERG GAMES

We address SGs with multiple leaders and followers. The key components of our approach are the following. First, we allow the leaders to decide whether to participate in the commitment or to defect from it by taking on the role of followers. This is modeled by the *agreement stage* of the SG, whose result is the formation of an agreement involving a subset of the leaders. Second, in the spirit of CEs, we introduce a correlation device that, after the agreement, draws recommendations and privately communicates them to the players. Following Conitzer and Korzhlyk [11], we assume that the leaders involved in the agreement commit to play their recommendations, while the followers obey to the usual incentive constraints of CEs (see Eq. (1)). The correlation device may adopt different distributions depending on the sequence of defections that determined the agreement, and these distributions are publicly known. Our goal is to design the device, so as to achieve some desirable properties of the commitment, which we formally describe in the rest of the section.

Before going into our main definitions, we introduce some useful notation. Given a subset of players $P \subseteq N$, we denote with Π_P the collection of ordered subsets of P , including the empty set \emptyset . Given $\pi \in \Pi_P$ and $p \in P \setminus \pi$, we let πp be the ordered set obtained by appending p at the end of π . We use $\mathbf{x} = [x_\pi]$ to denote a vector of correlated distributions $x_\pi \in \mathcal{X}_{\pi \cup F}^{\text{CE}}$, one per ordered subset of leaders $\pi \in \Pi_L$, while $\mathbf{X} = \times_{\pi \in \Pi_L} \mathcal{X}_{\pi \cup F}^{\text{CE}}$ is the set of all such vectors. In words, $\pi \in \Pi_L$ represents a sequence of leaders’ defections in the agreement stage, while \mathbf{x} defines the publicly known correlated distributions adopted by the correlation device, with x_π being the one used when the sequence of defections is π .

Definition 3.1. Given a vector of distributions $\mathbf{x} = [x_\pi] \in \mathbf{X}$, an SG (G, L, F) is structured in the following two stages:

- **Agreement.** It goes on in rounds. In a given round, each leader, in turn, decides between OPT-IN and OPT-OUT.³ All the decisions are perfectly observable. If a player chooses OPT-OUT, then she leaves the set of leaders becoming a follower, and a new round starts. The stage ends when, during a round, all remaining leaders decide to OPT-IN. The result is the ordered subset $\pi \in \Pi_L$ of leaders who decided to OPT-OUT.^{4 5}
- **Play.** The correlation device draws some $s \in S$ according to the publicly known correlated distribution x_π . Then, each

player is privately told her recommendation and the underlying game G is played, with the leaders in $L \setminus \pi$ sticking to their recommendations.

Remark 1. *The agreement stage of an SG can be represented as a sequential (i.e., tree-form) game involving the leaders. In such game, the players play in turn, according to some fixed order, with only two actions available at each decision point: OPT-IN and OPT-OUT. When a player chooses OPT-OUT, then she never plays anymore. The game ends after a sequence of OPT-IN actions performed by all leaders who have not selected OPT-OUT yet. Thus, each leaf of the game corresponds to the ordered subset $\pi \in \Pi_L$ representing the sequence of leaders who performed OPT-OUT on the path to the leaf. Players’ payoffs are defined by $u_p(x_\pi)$ for $p \in L$. See Figure 1 (Right) for an example of sequential-game-representation of the agreement stage.*

Next, we introduce some desirable properties that the distributions of the correlation device should satisfy. In the following definitions, we assume that an SG (G, L, F) is given.

First, we introduce *stability*. In words, we require that the leaders in L do not have any incentive to become followers. We introduce two different notions of stability, as follows.

Definition 3.2. Given $\mathbf{x} = [x_\pi] \in \mathbf{X}$, for any $\pi \in \Pi_L$, x_π is *stable* if, for every $p \in L \setminus \pi$, $u_p(x_\pi) \geq u_p(x_{\pi p})$. Moreover:

- \mathbf{x} is *stable* if x_\emptyset is stable;
- \mathbf{x} is *perfectly stable* if x_π is stable for every $\pi \in \Pi_L$.

We denote with $\mathbf{X}^{\text{S}} \subseteq \mathbf{X}$ and $\mathbf{X}^{\text{PS}} \subseteq \mathbf{X}$ the sets of stable and perfectly stable distributions, respectively.

Remark 2. *The rationale behind stability is that of NE. Indeed, $\mathbf{x} \in \mathbf{X}$ is stable if and only if each leader playing OPT-IN is an NE of the sequential game representing the agreement stage. Intuitively, this is because, if $\mathbf{x} \in \mathbf{X}$ is stable, each leader must not have any incentive to play OPT-OUT given that the other leaders always play OPT-IN.*

Remark 3. *The rationale behind perfect stability is that of subgame perfection. Indeed, $\mathbf{x} \in \mathbf{X}$ is perfectly stable if and only if each leader playing OPT-IN is a subgame perfect equilibrium of the agreement stage. The reason is that perfect stability requires that playing OPT-IN is optimal at any decision point of the sequential game.*

The second property that we look for is *efficiency*. We require that the correlated distributions of the correlation device are *Pareto optimal* with respect to the utility functions of the leaders who decided to OPT-IN. Given $\mathbf{X}' \subseteq \mathbf{X}$, for $\pi \in \Pi_L$, we use $\mathcal{P}_{L \setminus \pi}(\mathbf{X}')$ to denote the set of Pareto optimal correlated distributions in the set $\{x'_\pi \mid \mathbf{x}' = [x'_\pi] \in \mathbf{X}'\}$, where the objectives are the functions u_p , for $p \in L \setminus \pi$. Formally:

Definition 3.3. Given $\mathbf{x} = [x_\pi] \in \mathbf{X}' \subseteq \mathbf{X}$, for any $\pi \in \Pi_L$, x_π is *efficient* on the set \mathbf{X}' if $x_\pi \in \mathcal{P}_{L \setminus \pi}(\mathbf{X}')$. Moreover:

- \mathbf{x} is *efficient* on \mathbf{X}' if x_\emptyset is efficient on \mathbf{X}' ;
- \mathbf{x} is *perfectly efficient* on \mathbf{X}' if x_π is efficient on \mathbf{X}' for every $\pi \in \Pi_L$.

We introduce three different solution concepts for our SGs, which we refer to as *Stackelberg correlated equilibria* (SCEs). They differ on the types of stability and efficiency that they prescribe. Formally:

Definition 3.4. Given an SG (G, L, F) , $\mathbf{x} \in \mathbf{X}$ is an:

³We assume that the leaders are asked to take a decision according to some ordering, e.g., $p \in L$ decides before $q \in L$ if $p < q$.

⁴The agreement stage is finite as there are at most $|L|$ rounds and each round involves at most $|L|$ decisions.

⁵Our results do not rely on the protocol implemented in the agreement stage. Others could be adopted, with the only requirement that they must record in which order the leaders do OPT-OUT.

- SCE if it is efficient on the set \mathbf{X}^S ;
- SCE with perfect agreement (SCE-PA) if it is efficient on the set \mathbf{X}^{PS} ;
- SCE with perfect agreement and perfect efficiency (SCE-PAPE) if it is perfectly efficient on the set \mathbf{X}^{PS} .

We denote with \mathbf{X}^{SCE} , \mathbf{X}^{SCE-PA} , and $\mathbf{X}^{SCE-PAPE}$ the sets of SCEs, SCE-PAs, and SCE-PAPEs, respectively.

Example 3.5. Consider the SG in Figure 1, where $L = \{1, 2\}$ and $F = \emptyset$. Let $\mathbf{x} = [x_\pi]$ be such that $x_{\emptyset}(s_{1,1}, s_{2,1}) = 1$, $x_{\{2\}}(s_{1,5}, s_{2,1}) = 1$, and $x_\pi(s_{1,1}, s_{2,2}) = 1$ for all the other $\pi \in \Pi_L$. Clearly, $x_\pi \in \mathcal{X}_\pi^{CE}$ for all $\pi \in \Pi_L$. Moreover, being x_\emptyset stable and Pareto optimal, \mathbf{x} is an SCE. Observe that, if player 2 performs OPT-OUT, \mathbf{x} prescribes an irrational behavior to player 1, as $u_1(x_{\{2\}}) = 0$, while she gets 1 by doing OPT-OUT. Thus, \mathbf{x} is not perfectly stable, as playing OPT-IN must be optimal at any decision point of the agreement stage. For instance, $\mathbf{x}' = [x'_\pi]$ with $x'_{\emptyset}(s_{1,2}, s_{2,1}) = 1$ and $x'_\pi(s_{1,3}, s_{2,2}) = 1$ for every other $\pi \in \Pi_L$ is an SCE-PA. However, notice that \mathbf{x}' is not an SCE-PAPE since $x'_{\{2\}}$ does not maximize player 1's utility. Instead, $\mathbf{x}'' = [x''_\pi]$ with $x''_{\emptyset}(s_{1,4}, s_{2,1}) = 1$, $x''_{\{2\}}(s_{1,3}, s_{2,1})$, and $x''_\pi(s_{1,4}, s_{2,2}) = 1$ for all the other $\pi \in \Pi_L$ is an SCE-PAPE.




Figure 1: Left: Example of two-player normal-form SG with $L = \{1, 2\}$. Right: Sequential game representing its agreement stage.

4 ON THE EXISTENCE OF SCE

We investigate the existence of our solution concepts in general SGs. We show that SCEs and SCE-PAs always exist, while we provide an SG where there is no SCE-PAPE.

The fundamental step for proving our existence results (Theorem 4.1) is to show that (i) \mathbf{X}^S and \mathbf{X}^{PS} are polytopes, and (ii) they are non-empty. The latter point is a direct consequence of the fact that all vectors $\mathbf{x} = [x_\pi] \in \mathbf{X}$ with $x_\pi = x$ for some CE $x \in \mathcal{X}^{CE}$ are perfectly stable.

THEOREM 4.1. *Every SG admits an SCE and an SCE-PA.*

PROPOSITION 4.2. *There are SGs with no SCE-PAPE.*

PROOF SKETCH. Consider the SG in Table 1, where $L = \{1, 2, 3\}$ and $F = \emptyset$. Any $\mathbf{x} = [x_\pi] \in \mathbf{X}^{SCE-PAPE}$ must be such that, for every x_π with player 3 in π , $u_3(x_\pi) = 1$ (as player 3 always gets 1 by deviating to $s_{3,3}$). Given the definition of stability and player 3's incentive constraints, $x_{\{1,2\}}$ and $x_{\{2,1\}}$ must always recommend $s_{3,3}$ to player 3. Moreover, by stability and efficiency, $x_{\{1\}}$ must always recommend $(s_{1,1}, s_{2,2}, s_{3,2})$, where player 1 gets a utility of

	$s_{2,1}$	$s_{2,2}$		$s_{2,1}$	$s_{2,2}$		$s_{2,1}$	$s_{2,2}$
$s_{1,1}$	0, 2, 0	2, 0, 0	$s_{3,1}$	0, 2, 0	2, 1, 1	$s_{3,2}$	0, 0, 1	0, 0, 1
$s_{1,2}$	0, 2, 0	1, 2, 1		0, 2, 0	0, 0, 0		0, 0, 1	0, 0, 1
							$s_{3,3}$	

Table 1: Three-player normal-form SG with no SCE-PAPE (players 1, 2, and 3 select rows, columns, and matrices, respectively).

2. Similarly, $x_{\{2\}}$ must always recommend $(s_{1,2}, s_{2,2}, s_{3,1})$ and, thus, player 2 receives a utility of 2. Thus, for stability, x_\emptyset must satisfy $u_1(x_\emptyset), u_2(x_\emptyset) \geq 2$, which is impossible. \square

As a result, in the rest of this work we focus on SCEs and SCE-PAs. We remark that the non-existence of SCE-PAPEs implies that, under the requirements of perfect stability and perfect efficiency, there cannot be an agreement involving all the leaders. This does not rule out the possibility that some subsets of leaders can still reach an agreement. However, these cases are much more involved, as the actual group of leaders reaching an agreement inevitably depends on the rules of the protocol implemented in the agreement stage.

5 SCES AND OTHER SOLUTION CONCEPTS

We show that the optimal correlated strategies to commit to introduced by Conitzer and Korzhyk [11] are a special case of SCEs. Intuitively, in single-leader SGs, efficiency is equivalent to the maximization of leader's utility, while stability does not enforce additional constraints on the commitment.

THEOREM 5.1. *Given an SG $(G, \{1\}, N \setminus \{1\})$, it holds $\mathbf{X}^{SCE} = \mathbf{X}^{SCE-PA} = \mathbf{X}^{SCE-PAPE}$ and, given some $\mathbf{x} = [x_\pi] \in \mathbf{X}^{SCE}$, x_\emptyset is an optimal correlated strategy to commit to.*

PROOF SKETCH. Since the SG has only one leader (player 1), $\mathbf{X}^S = \mathbf{X}^{PS}$, and, thus, $\mathbf{X}^{SCE} = \mathbf{X}^{SCE-PA}$. For the same reasons, $\mathbf{X}^{SCE-PA} = \mathbf{X}^{SCE-PAPE}$. Moreover, Pareto optimality is the same as maximizing the leader's utility function u_1 . Let $\mathbf{x} = [x_\pi] \in \mathbf{X}^{SCE}$ and assume, by contradiction, that x_\emptyset is not an optimal correlated strategy to commit to. Then, there would be another $\hat{x} \in \mathcal{X}_{N \setminus \{1\}}^{CE}$ such that $u_1(\hat{x}) \geq u_1(x_\emptyset)$. However, replacing x_\emptyset with \hat{x} in \mathbf{x} would give us another $\hat{\mathbf{x}} \in \mathbf{X}^S$, contradicting the efficiency of \mathbf{x} . \square

Given the relation between optimal correlated strategies to commit to and SEs in single-leader single-follower SGs:

COROLLARY 5.2. *Given an SG $(G, \{1\}, \{2\})$, any $\mathbf{x} = [x_\pi] \in \mathbf{X}^{SCE}$ is such that $u_1(x_\emptyset)$ is the leader's utility in an SE.*

For the relationships of SCEs with other non-Stackelberg solution concepts, see [8].

6 COMPUTATIONAL COMPLEXITY OF SCE

We study the computational complexity of SCEs and SCE-PAs in general SGs. We distinguish between the problem of finding an equilibrium and that of computing an *optimal* equilibrium, i.e., one maximizing a specific given linear function of leaders' utilities,

such as the leader’s social welfare. We introduce the following formal definitions (problems f-SCE-PA and o-SCE-PA(λ) are defined analogously for SCE-PAs).

Definition 6.1 (f-SCE). Given an SG (G, L, F) , find an SCE.

Definition 6.2 (o-SCE(λ)). Given an SG (G, L, F) and $\lambda = [\lambda_p] \in [0, 1]^{|L|}$, find an SCE $\mathbf{x} = [x_\pi] \in \mathbf{X}^{\text{SCE}}$ maximizing the objective function $f_\lambda = \sum_{p \in L} \sum_{s \in S} \lambda_p u_p(s) x_\emptyset(s)$.

Let us remark that, in general, the size of a vector $\mathbf{x} \in \mathbf{X}$ is factorial in the number of players. Thus, in the following, we assume that there is some compact representation for \mathbf{x} .⁶

We establish a tight connection between our problems and an auxiliary one, which is a generalization of the problem of finding an optimal CE. In the rest of the section, we assume to have access to an oracle solving this auxiliary problem, which we call *stability oracle*. In Section 7, we then investigate for which games the oracle can be efficiently implemented.

Definition 6.3. A *stability oracle* $\mathcal{O}(G, c, L, \{x_p\}_{p \in L' \subseteq L})$ is an algorithm that, given a finite game G , a coefficients vector $c = [c_p] \in [-1, 1]^N$, a set of leaders $L \subseteq N$, and a collection of correlated distributions $x_p \in \mathcal{X}$ for $p \in L' \subseteq L$, returns an $x \in \mathcal{X}_{N \setminus L}^{\text{CE}}$ maximizing $\sum_{p \in N} \sum_{s \in S} c_p u_p(s) x(s)$ subject to the stability constraints, i.e., $u_p(x) \geq u_p(x_p)$ for all $p \in L'$.⁷

In the following, we are interested in games where the stability oracle runs in polynomial time. Thus, we assume that \mathcal{O} always returns a correlated distribution with size polynomial in the size of the game.⁸ We also consider the decision form of the stability oracle, which reads as follows:

Definition 6.4. The *decision form* of a stability oracle \mathcal{O} is an algorithm $\mathcal{O}^D(x, L, \{x_p\}_{p \in L' \subseteq L})$ that, given $x \in \mathcal{X}$, $L \subseteq N$, and $x_p \in \mathcal{X}$ for $p \in L' \subseteq L$, answers YES if $x \in \mathcal{X}_{N \setminus L}^{\text{CE}}$ and x satisfies the stability constraints, and No otherwise.

In the following, given $L \subseteq N$ and $\lambda = [\lambda_p] \in [0, 1]^{|L|}$, we let $c_\lambda = [c_{\lambda, p}] \in [0, 1]^N$ be such that $c_{\lambda, p} = \lambda_p$ if $p \in L$, while $c_{\lambda, p} = 0$ if not. Moreover, given $p \in N$, we let $c_p = [c_{p, q}] \in [0, 1]^N$ be such that $c_{p, p} = -1$ and $c_{p, q} = 0$ for all $q \in N \setminus \{p\}$. Note that c_λ is the coefficients vector of the objective f_λ , while c_p corresponds to minimizing u_p .

6.1 Computing SCEs

We show that, in games admitting a polynomial-time stability oracle, an optimal SCE can be computed in polynomial time. Intuitively, o-SCE(λ) is solved by $\mathbf{x} = [x_\pi]$ computed as: $x_{\{p\}} = \mathcal{O}(G, c_p, L \setminus \{p\}, \emptyset)$ for $p \in L$, $x_\emptyset = \mathcal{O}(G, c_\lambda, L, \{x_{\{p\}}\}_{p \in L})$, and $x_\pi = \mathcal{O}(G, c_\lambda, \emptyset, \emptyset)$ for every other ordered subset $\pi \in \Pi_L$. Formally:

THEOREM 6.5. *Given an SG (G, L, F) and $\lambda \in [0, 1]^{|L|}$, o-SCE(λ) can be solved with $|L| + 2$ queries to an oracle \mathcal{O} .*

⁶As we see next, for all our positive results we can safely assume that there is a compact representation for $\mathbf{x} \in \mathbf{X}$ (e.g., \mathbf{x} only requires a polynomial number of polynomially-sized distributions).

⁷Note that, given a finite game G , $\mathcal{O}(G, c, \emptyset, \emptyset)$ returns an optimal CE $x \in \mathcal{X}^{\text{CE}}$ for the objective function defined by $c \in [0, 1]^N$.

⁸Indeed, this assumption is not restrictive, as all the games we study in Section 7 admit a poly-time oracle \mathcal{O} with this property.

COROLLARY 6.6. *Given an SG (G, L, F) , if there is a poly-time oracle \mathcal{O} , then o-SCE(λ) can be solved in polynomial time.*

6.2 Computing SCE-PAs

First, we provide a positive result: one can find an SCE-PA with polynomially many invocations to a stability oracle. It is sufficient to compute $\mathbf{x} = [x_\pi]$ where $x_{\{p\}} = \mathcal{O}(G, c_p, \emptyset, \emptyset)$ for $p \in L$ and $x_\emptyset = \mathcal{O}(G, c_\lambda, L, \{x_{\{p\}}\}_{p \in L})$. Thus:

THEOREM 6.7. *Given an SG (G, L, F) , f-SCE-PA can be solved with $|L| + 1$ queries to an oracle \mathcal{O} .*

COROLLARY 6.8. *Given an SG (G, L, F) , if there is a poly-time oracle \mathcal{O} , then f-SCE-PA can be solved in polynomial time.*

Now, we switch to the problem of computing an optimal SCE-PA, showing that it cannot be solved efficiently, even with access to a polynomial-time stability oracle. Specifically, we prove a stronger negative result: even the easier problem of verifying the perfect stability of a given $\mathbf{x} \in \mathbf{X}$ is computationally intractable. Our statement is based on a reduction from the coNP-complete problem of deciding whether a given formula in *disjunctive normal form* (DNF) is a tautology or not [3].

THEOREM 6.9. *Given an SG (G, L, F) and $\mathbf{x} \in \mathbf{X}$, verifying whether $\mathbf{x} \in \text{or } \notin \mathbf{X}^{\text{PS}}$ is not in P unless NP = coNP, even with access to a polynomial-time decision-form oracle \mathcal{O}^D .*

PROOF SKETCH. Given a formula Φ in DNF, we construct an SG (G, L, F) involving a leader p_v for each variable $v \in V$ and a single follower p_f . Each p_v has two actions, s_T and s_F , which define the truth value of v . As a result, any $s \in S$ corresponds to a truth assignment τ^s defined by leaders’ strategies. The follower has a strategy s_v for each variable $v \in V$. Table 2 reports the leaders’ utilities (the follower always gets 0). We build $\mathbf{x} = [x_\pi] \in \mathbf{X}$ with $x_\emptyset(s) = 1$ for some $s \in S$ such that $s_{p_v} = s_T$ for all $v \in V$. Furthermore, for every $v \in V$ and $\pi \in \Pi_{L \setminus \{p_v\}}$, we let $x_{\pi p_v}(s) = 1$ for $s \in S$ with $s_p = s_F$ for all $p \in \pi p_v$, $s_p = s_T$ for all $p \in L \setminus \pi p_v$, and $s_{p_f} = s_v$.

If Φ is a tautology. For every $\pi \in \Pi_L$, x_π recommends all the leaders in π to play s_F . Note that, for every $v \in V$ and $\pi \in \Pi_{L \setminus \{p_v\}}$, $u_{p_v}(x_\pi) = \#\mathbf{F}(\tau^s) = |\pi|$, while, if p_v decides to OPT-OUT, she is recommended to play s_F (one more variable is set to false) and, being $s_{p_f} = s_v$, she gets the same utility. As a result, all distributions x_π are stable.

Only if. Suppose Φ is not a tautology. Let $s \in S$ be such that $\Phi(\tau^s) = \mathbf{F}$. If $s_{p_v} = s_T$ for every $v \in V$, then x_\emptyset is not stable as the leaders would OPT-OUT (getting at least $0 > -1$). Otherwise, there exists $s' \in S$ such that $\Phi(\tau^{s'}) = \mathbf{T}$ and $x_\pi(s') = 1$, $x_{\pi p_v}(s) = 1$ for some $v \in V$, $\pi \in \Pi_{L \setminus \{p_v\}}$. Then, $u_{p_v}(x_\pi) = \#\mathbf{F}(\tau^{s'}) = |\pi|$ and $u_{p_v}(x_{\pi p_v}) = |V| > |\pi|$. Thus, x_π is not stable, as leader p_v would OPT-OUT. \square

As a byproduct of Theorem 6.5 we have that, when looking for optimal SCEs, one can restrict the attention to those $\mathbf{x} \in \mathbf{X}$ admitting a representation whose size is polynomial in the size of the game. For Theorem 6.7, the same holds when searching for an SCE-PA. However, Theorem 6.9 implies that optimal SCE-PAs require an exponential number of different distributions. Moreover, even when $\mathbf{x} \in \mathbf{X}$ can be easily represented in a compact form (as

		$\Phi(\tau^S) = T$		$\Phi(\tau^S) = F$	
		$s_{p_f} = s_v$	$s_{p_f} \neq s_v$	$\exists v : s_{p_v} = s_F$	$\forall v : s_{p_v} = s_T$
ST		0	$\#F(\tau^S)$	$ V $	-1
SF		$\#F(\tau^S) - 1$	$ V $	$ V $	0

Table 2: Leader p_v 's ($v \in V$) utilities in the SG for the reduction of Theorem 6.9. On rows, there are p_v 's strategies s_T and s_F , whereas, on columns, we report the four possible cases for $s \in S$. $\#F(\tau^S)$ denotes the number of variables set to false by τ^S .

in the proof of Theorem 6.9), we cannot check in polynomial time whether $\mathbf{x} \in X^{PS}$ or not.

This poses a new intriguing question: can we restrict the attention to $\mathbf{x} \in X$ whose size is less than factorial in the number of players? We show that the answer is positive. It is sufficient to consider $\mathbf{x} \in X$ whose size is exponential in the number of players, as only the unordered set of defecting leaders and the last of them who decided to OPT-OUT matter.

THEOREM 6.10. *Given an SG (G, L, F) and $\mathbf{x} = [x_\pi] \in X^{PS}$, there is an $\mathbf{x}' = [x'_\pi] \in X^{PS}$ s.t. $x'_\emptyset = x_\emptyset$ and $x'_{\pi p} = x'_{\pi' p}$ for every $p \in L$ and $\pi, \pi' \in \Pi_{L \setminus \{p\}}$ defining the same set.*

Theorem 6.10 allows us to reduce the number of queries to a stability oracle that are necessary to find an optimal SCE-PA.

THEOREM 6.11. *Given an SG (G, L, F) and $\lambda \in [0, 1]^{|L|}$, o-SCE-PA(λ) can be solved with $|L|2^{|L|-1} + 1$ queries to \mathcal{O} .*

Finally, we can provide an example showing that Theorem 6.11 is tight, which leads to the following proposition.

Solving o-SCE-PA(λ) requires to take into account the last player who performed OPT-OUT, while focusing only on the set of defecting leaders is not sufficient.

7 STABILITY ORACLE FOR COMPACT GAMES

We study which classes of games admit a polynomial-time stability oracle \mathcal{O} , focusing on those with polynomial type.⁹ We only provide our main final result; a detailed description of all the ancillary results is in [8].

Inspired by the classical approaches for finding CEs in games with polynomial type [18, 19, 32], we solve $\mathcal{O}(G, c, L, \{x_p\}_{p \in L' \subseteq L})$ in polynomial time using the ellipsoid method. This requires that a suitably defined separation problem ($\text{Sep}(z, t)$) can be computed in polynomial time. Our main result is that $\text{Sep}(z, t)$ can be reduced to the *weighted deviation-adjusted social welfare problem* ($w\text{-DaSW}(y, v, t)$) introduced by Jiang and Leyton-Brown [18] for finding an optimal (according to some linear function of players' utilities) CE. This establishes a strict connection between the problem solved by our stability oracle and that of computing optimal CEs. As a consequence, given the results of Jiang and Leyton-Brown [18], \mathcal{O} can be computed in polynomial time for all the compact

⁹We remark that, for normal-form games, a polynomial-time stability oracle \mathcal{O} can be implemented by using a variation of the LP for finding optimal CEs [36].

games where finding an optimal CE is computationally tractable. Thus:

THEOREM 7.1. *The following games admit a polynomial-time stability oracle \mathcal{O} : anonymous games, symmetric games, and bounded-treewidth graphical and polymatrix games.*

Finally, our results also imply that the polynomial-time stability oracle \mathcal{O} always outputs a polynomially-sized distribution (see [8, Corollary 11.1]).

8 DISCUSSION

This paper introduces a new way to apply the Stackelberg paradigm to any (underlying) finite game. Differently from previous works, our approach deals with scenarios involving multiple leaders by introducing a preliminary agreement stage in which each leader can decide whether to be a leader or become a follower. We introduce and study three natural solution concepts that differ depending on the properties that they require on the agreement stage (others, e.g., requiring stability and perfect efficiency, will be explored in future).

Our equilibria generalize the optimal correlated strategies to commit to introduced by Conitzer and Korczyk [11] for single-leader multi-follower Stackelberg games. At the same time, they also provide a significant advancement over the multi-leader solution concepts introduced in the security context (see, e.g., [17]). First, correlated-strategy commitments are more natural than leaders' strategies satisfying some Nash-like constraints. Second, our equilibria are founded on strong game-theoretic groundings, as they are guaranteed to exist independently of the game structure. Last but not least our solutions apply to general games.

Finally, our computational findings exploit a general framework relying on a game-independent stability oracle. Thus, our positive results can be extended to other game classes by simply designing polynomial-time oracles.

ACKNOWLEDGMENTS

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Market".

REFERENCES

- [1] B. An, J. Pita, E. Shieh, M. Tambe, C. Kiekintveld, and J. Marecki. 2011. Guards and Protect: Next generation applications of security games. *ACM SIGEC EX 10*, 1 (2011), 31–34.
- [2] Simon P Anderson and Gerhard Glomm. 1992. Incumbency effects in political campaigns. *Public Choice* 74, 2 (1992), 207–219.
- [3] S. Arora and B. Barak. 2009. *Computational complexity: a modern approach*.
- [4] R.J. Aumann. 1974. Subjectivity and correlation in randomized strategies. *J MATH ECON* 1, 1 (1974), 67–96.
- [5] M. Blonski. 2000. Characterization of pure strategy equilibria in finite anonymous games. *J MATH ECON* 34, 2 (2000), 225–233.
- [6] M. Breton, A. Alj, and A. Haurie. 1988. Sequential Stackelberg equilibria in two-person games. *J OPTIMIZ THEORY APP* 59, 1 (1988), 71–97.
- [7] Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. 2019. Be a Leader or Become a Follower: The Strategy to Commit to with Multiple Leaders. In *IJCAL* 123–129.
- [8] M. Castiglioni, A. Marchesi, and N. Gatti. 2019. Be a Leader or Become a Follower: The Strategy to Commit to with Multiple Leaders (Extended Version). *CoRR abs/1905.13106* (2019). arXiv:1905.13106
- [9] S. Coniglio, N. Gatti, and A. Marchesi. 2017. Pessimistic Leader-Follower Equilibria with Multiple Followers. In *IJCAL*.
- [10] Stefano Coniglio, Nicola Gatti, and Alberto Marchesi. 2018. Computing a Pessimistic Leader-Follower Equilibrium with Multiple Followers: the Mixed-Pure Case. *CoRR abs/1808.01438* (2018). arXiv:1808.01438

- [11] V. Conitzer and D. Korzhyk. 2011. Commitment to Correlated Strategies.. In *AAAI*.
- [12] V. Conitzer and T. Sandholm. 2006. Computing the optimal strategy to commit to. In *ACM EC*.
- [13] G. De Nittis, A. Marchesi, and N. Gatti. 2018. Computing the strategy to commit to in polymatrix games. In *AAAI*.
- [14] V. DeMiguel and H. Xu. 2009. A stochastic multiple-leader Stackelberg model: analysis, computation, and application. *OPER RES* 57, 5 (2009), 1220–1235.
- [15] Effrosyni Diamantoudi. 2005. Stable cartels revisited. *Economic Theory* 26, 4 (2005), 907–921.
- [16] B.C. Eaves. 1973. Polymatrix games with joint constraints. *SIAM J APPL MATH* 24, 3 (1973), 418–423.
- [17] J. Gan, E. Elkind, and M. Wooldridge. 2018. Stackelberg Security Games with Multiple Uncoordinated Defenders. In *AAMAS*.
- [18] A.X. Jiang and K. Leyton-Brown. 2011. A General Framework for Computing Optimal Correlated Equilibria in Compact Games. In *WINE*.
- [19] A.X. Jiang and K. Leyton-Brown. 2015. Polynomial-time computation of exact correlated equilibrium in compact games. *GAME ECON BEHAV* 91 (2015), 347–359.
- [20] M. Kearns, M. Littman, and S. Singh. 2013. Graphical models for game theory. *arXiv:1301.2281* (2013).
- [21] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordóñez, and M. Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *AAMAS*.
- [22] A.A. Kulkarni and U.V. Shanbhag. 2014. A shared-constraint approach to multi-leader multi-follower games. *SET-VALUED VAR ANAL* 22, 4 (2014), 691–720.
- [23] A. Laszka, J. Lou, and Y. Vorobeychik. 2016. Multi-defender strategic filtering against spear-phishing attacks. In *AAAI*.
- [24] S. Leyffer and T. Munson. 2010. Solving multi-leader-common-follower games. *OPT MET SO* 25, 4 (2010), 601–623.
- [25] J. Lou, A.M. Smith, and Y. Vorobeychik. 2017. Multidefender security games. *IEEE INTELL SYST* 32, 1 (2017), 50–60.
- [26] J. Lou and Y. Vorobeychik. 2015. Equilibrium analysis of multi-defender security games. In *IJCAI*.
- [27] Z-Q. Luo, J-S. Pang, and D. Ralph. 1996. *Mathematical programs with equilibrium constraints*.
- [28] A. Marchesi, M. Castiglioni, and N. Gatti. 2019. Leadership in Congestion Games: Multiple User Classes and Non-Singleton Actions. In *IJCAI*.
- [29] A. Marchesi, S. Coniglio, and N. Gatti. 2018. Leadership in Singleton Congestion Games. In *IJCAI*.
- [30] J. Nash. 1951. Non-cooperative games. *ANN MATH* (1951), 286–295.
- [31] J-S. Pang and M. Fukushima. 2005. Quasi-variational inequalities, generalized Nash equilibria, and multi-leader-follower games. *COMP MAN SC* 2, 1 (2005), 21–56.
- [32] C.H. Papadimitriou and T. Roughgarden. 2008. Computing correlated equilibria in multi-player games. *J ACM* 55, 3 (2008), 14.
- [33] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. 2008. Playing games for security: an efficient exact algorithm for solving Bayesian Stackelberg games. In *AAMAS*.
- [34] R.W. Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *INT J GAME THEORY* 2, 1 (1973), 65–67.
- [35] H.D. Sherali. 1984. A multiple leader Stackelberg model and analysis. *OPER RES* 32, 2 (1984), 390–404.
- [36] Y. Shoham and K. Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*.
- [37] A. Sinha, P. Malo, A. Frantsev, and K. Deb. 2014. Finding optimal strategies in a multi-period multi-leader-follower Stackelberg game using an evolutionary algorithm. *COMPUT OPER RES* 41 (2014), 374–385.
- [38] A. Smith, Y. Vorobeychik, and J. Letchford. 2014. MultiDefender security games on networks. *PERF E R* 41, 4 (2014), 4–7.
- [39] M. Tambe. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*.
- [40] H. Von Stackelberg. 1934. *Marktform und Gleichgewicht*.
- [41] B. Von Stengel and S. Zamir. 2010. Leadership games with convex strategy sets. *GAME ECON BEHAV* 69, 2 (2010), 446–457.