Chapter 10
Leadership Games: Multiple Followers, Multiple Leaders, and Perfection

Check for updates

Alberto Marchesi

10.1 Introduction

Over the last years, *algorithmic game theory* has received growing interest in AI, as it allows to tackle complex real-world scenarios involving multiple artificial agents engaged in a competitive interaction. These settings call for rational agents endowed with the capability of reasoning strategically, which is achieved by exploiting *equilibrium* concepts from game theory. The challenge is to design scalable computational tools that enable the adoption of such equilibrium notions in real-world problems.

The recent advances in the development of equilibrium-finding techniques have lead to the successful application of game-theoretic models in real-world settings. For instance, game theory has been extensively adopted in security domains, with the goal of devising protection strategies which are robust against strategic attackers [47]. Other application domains are found in the *Internet*, where interactions involving multiple strategic agents naturally arise, given the intrinsic distributed nature of the network. One examples is, among others, the problem of designing auction mechanisms for web advertising [23, 27]. Moreover, great achievements have been made towards the development of artificial agents capable of beating human professional in large two-player zero-sum recreational games like Chess [11], Go [46], and Poker [9, 10].

Despite the great attention devoted to algorithmic game theory in the last years, most of the works in the literature study (relatively) simple settings involving only two players with opposite objectives, i.e., two-player zero-sum games. In such models, there is a clear and well-established definition of solution, in which each player aims to maximize her utility given that the opponent acts so as to minimize it. In zero-sum games, this definition corresponds to that of *Nash equilibrium*. Thus, considerable efforts have been devoted to studying the problem of computing (possibly approximate) Nash equilibria in such settings. Instead, more complex games where

A. Marchesi (✉)
Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy
e-mail: alberto.marchesi@polimi.it

there are more than two players and/or arbitrary, i.e., general-sum, utilities are widely unexplored. In such scenarios, there is no clear definition of solution to a game, as this strongly depends on the specific application that one wish to represent. As a result, many solution concepts other than the Nash equilibrium have been introduced and studied. However, there is still a lot of work to be done on the computational side, as the algorithmic works on multi-player general-sum games are only few.

In this work, we study settings beyond two-player zero-sum games, focusing on a particular game paradigm which leads to the definition of what is known in the literature as the *Stackelberg equilibrium*.

10.2   The Stackelberg Paradigm

The Stackelberg paradigm was originally introduced by von Stackelberg in 1934 to model economic situations where a firm (the *leader*) moves first and, then, another firm (the *follower*) moves second by reacting to the first firm's move [50]. Recently, this paradigm was brought to new attention by the work of [7], who study a variant of the original Stackelberg paradigm in which the leader commits to a (possibly randomized, i.e., mixed) strategy beforehand, while the follower decides how to play after observing the leader's strategy. In general settings involving multiple players, a *Stackelberg game* is characterized by a group of players who act as leaders with the ability to commit to (possibly mixed) strategies beforehand, whereas the other players are followers who observe the commitment and decide how to play thereafter.

Over the last years, Stackelberg games and their corresponding Stackelberg equilibria have received growing attention in the AI literature, where the computational problem of finding such equilibria in often referred to as the problem of *computing optimal strategies to commit to* [20]. This surge of interest was motivated by the successful applications of Stackelberg games in many interesting real-world settings. In particular, among the others, the security domain is the most explored one, and, in it, different game models have been introduced, usually referred to as *security games* [2, 29, 42, 47]. In such models, there is a defender that has to protect some valuable targets from an attacker, who can wait while observing the defender's protection strategy before deciding where, when and how to attack. This scenario naturally fits into the Stackelberg model, where the defender is the leader and the attacker is the follower. Other interesting applications are found in *toll-setting games*, where the leader is a central authority which collects tolls from the users of a network who, acting as followers, decide on how to best travel through the network so as to minimize their cost after observing the pricing strategy chosen by the authority [30, 31]. Besides the security domain and toll-setting games, applications of Stackelberg games can be found in, among others, interdiction games [12, 39], network routing [1], inspection games [3], and mechanism design [44].

Despite the attention that Stackelberg games received from the AI literature, most of the works related to them focus, with some exceptions (see, e.g., [7, 19, 26]), on particular game settings that involve only two players (i.e., one leader and one follower) and enjoy specific structures, as it is the case in security games. It is worth

pointing out two works that study general Stackelberg games with a single leader and multiple follower; specifically, [7] study the case in which the followers play a Nash equilibrium given the leader's commitment, whereas [19] address the case where they play a *correlated equilibrium*.

Let us also notice that, while some works (see, e.g., [8, 16, 32]) address the computation of Stackelberg equilibria in games with a sequential (i.e., tree-form) structure, none of them investigates refinements of such equilibria. This is surprising as refinements have been extensively studied for the Nash equilibrium, since it is well-known that classical (unrefined) solution concepts may lead to a sub-optimal behavior off the equilibrium path in games with a sequential structure (see [24, 49] for some references on the topic).

10.3   Stackelberg Games with Multiple Followers

We address Stackelberg games with a *single leader* and *multiple followers*. Following [7], we study settings in which, after observing the leader's commitment, the followers play a Nash equilibrium in the resulting game. We refer to this solution as *Stackelberg-Nash equilibrium*. We focus on the case in which the followers are restricted to pure (i.e., non-mixed) strategies, as the general problem with followers playing mixed strategies is already known to be computationally intractable [4]. This restriction leads to interesting computational complexity results. Moreover, this is without loss of generality in games that always admit pure-strategy Nash equilibria, as it is the case for congestion games [43].

We study the problem of computing Stackelberg-Nash equilibria, focusing on two cases: the one in which the followers break ties in favor of the leader (what is usually referred to as a *strong* equilibrium), and the case where they break ties against the leader (leading to a *weak* equilibrium). Moreover, we analyze three different classes of games, namely, *normal-form* games, *polymatrix* games, and *congestion* games.

Table 10.1 shows our contributions related to Stackelberg games with a single leader, summarizing the computational complexity and the algorithmic aspects of the problems we study, with focus on normal-form, extensive-form, Bayesian, and polymatrix games. The table also shows, for comparison, other state-of-the-art results, including those about single-leader single-follower Stackelberg games (our original contributions are those without a reference). Our contributions on Stackelberg congestion games are instead detailed in Table 10.2.

10.3.1   Norma-Form Stackelberg Games

Since a strong Stackelberg-Nash equilibrium (with followers restricted to pure strategies) can be computed efficiently (in polynomial time) by solving multiple *linear programs* (LPs), we entirely devote our analysis to the weak case (with, again, followers restricted to pure strategies). In terms of computational complexity, we show that, differently from the strong case, in the weak one the equilibrium-finding prob-

Table 10.1 Summary of the results on the computation of Stackelberg equilibria in normal-form Stackelberg games, Bayesian Stackelberg games, extensive-form Stackelberg games, and Stackelberg polymatrix games. The state-of-the-art results are those with related references

Followers' strategies		Strong Stackelberg(-Nash) equilibrium		Weak Stackelberg(-Nash) equilibrium	
		Pure	Mixed	Pure	Mixed
Normal-form Stackelberg games					
$n = 2$	Complexity	P [20]		P [7]	
	Algorithm	Multi-LP [20]		Multi-LP [7]	
$n = 3$	Complexity	P	NP-hard, $\notin$ Poly-APX [6]	NP-hard	NP-hard, $\notin$ Poly-APX [6]
	Algorithm	Multi-LP	Spatial branch-and-bound [5]	Multi-lex-MILP	–
$n \geq 4$	Complexity	P	NP-hard, $\notin$ Poly-APX [6]	NP-hard, $\notin$ Poly-APX	NP-hard, $\notin$ Poly-APX [6]
	Algorithm	Multi-LP	Spatial branch-and-bound [5]	Multi-lex-MILP	–
Bayesian Stackelberg games					
$n = 2$	Complexity	NP-hard [20], Poly-APX-complete [33]		NP-hard, Poly-APX-complete	
	Algorithm	MILP [42]		Multi-LP	
Extensive-form Stackelberg games					
$n = 2$	Complexity	NP-hard [32]		NP-hard	
	Algorithm	MILP [8, 16]		Multi-LP [7]	
Stackelberg polymatrix games					
$n = 3$	Complexity	P	NP-hard, $\notin$ Poly-APX [6]	NP-hard	NP-hard, $\notin$ Poly-APX [6]
	Algorithm	Multi-LP	Spatial branch-and-bound [5, 6]	Multi-lex-MILP	–
$n \geq 4$ (fixed)	Complexity	P	NP-hard, $\notin$ Poly-APX [6]	NP-hard, $\notin$ Poly-APX	NP-hard, $\notin$ Poly-APX [6]
	Algorithm	Multi-LP	Spatial branch-and-bound [5, 6]	Multi-lex-MILP	–
$n \geq 4$ (free)	Complexity	NP-hard, $\notin$ Poly-APX	NP-hard, $\notin$ Poly-APX [6]	NP-hard, $\notin$ Poly-APX	NP-hard, $\notin$ Poly-APX [6]
	Algorithm	Multi-LP	Spatial branch-and-bound [5]	Multi-lex-MILP	–

lem is NP-hard with two or more followers, while, when the number of followers is three or more, the problem cannot be approximated in polynomial time to within any polynomial multiplicative factor unless P = NP (i.e., in formal terms, it is *not* in the class Poly-APX unless P = NP). To establish these two results, we introduce two reductions, one from Independent Set and the other one from 3-SAT.

After analyzing the complexity of the problem, we focus on its algorithmic aspects. First, we formulate the problem as a *bilevel programming problem*. We then show how to recast it as a single-level *quadratically constrained quadratic program* (QCQP),

Table 10.2 Summary of the results on the computation of Stackelberg equilibria in Stackelberg singleton congestion games with a single leader

Strong Stackelberg-Nash equilibrium

Leader's commitment			Pure	Mixed
Identical action spaces (symmetric games)	Monotonic costs	Complexity	P	P
		Algorithm	Greedy	Greedy
	Generic costs	Complexity	P	NP-hard, $\notin$ Poly-**APX**
		Algorithm	Dynamic programming	MILP
Different action spaces	Monotonic costs	Complexity	NP-hard, $\notin$ Poly-**APX**	NP-hard, $\notin$ Poly-**APX**
		Algorithm	MILP	MILP
	Generic costs	Complexity	NP-hard, $\notin$ Poly-**APX**	NP-hard, $\notin$ Poly-**APX**
		Algorithm	MILP	MILP

Weak Stackelberg-Nash equilibrium

Leader's commitment			Pure	Mixed
Identical action spaces (symmetric games)	Monotonic costs	Complexity	P	P
		Algorithm	Greedy	Greedy
	Generic costs	Complexity	P	NP-hard, $\notin$ Poly-**APX**
		Algorithm	Dynamic programming	Multi-lex-MILP
Different action spaces	Monotonic costs	Complexity	NP-hard, $\notin$ Poly-**APX**	NP-hard, $\notin$ Poly-**APX**
		Algorithm	Multi-lex-MILP	Multi-lex-MILP
	Generic costs	Complexity	NP-hard, $\notin$ Poly-**APX**	NP-hard, $\notin$ Poly-**APX**
		Algorithm	Multi-lex-MILP	Multi-lex-MILP

which we show to be impractical to solve due to admitting a supremum, but not a maximum. We then introduce a restriction based on a *mixed-integer linear program* (MILP) which, while forsaking optimality, always admits an optimal (restricted) solution. Next, we propose an exact algorithm to compute the value of the supremum of the problem based on an enumeration scheme which, at each iteration, solves a *lexicographic* MILP (lex-MILP) where the two objective functions are optimized in sequence. Subsequently, we embed the enumerative algorithm within a branch-and-bound scheme, obtaining an algorithm which is, in practice, much faster. We also extend the algorithms so that, for cases where the supremum is not a maximum, they

return a strategy by which the leader can obtain a utility within an additive loss α with respect to the supremum, for any $\alpha > 0$. To conclude, we experimentally evaluate the scalability of our methods over a testbed of randomly generated instances.

A preliminary version of our results on normal-form Stackelberg games appeared in [17], while a complete version is [18].

10.3.2   Stackelberg Polymatrix Games

We identify two classes of Stackelberg polymatrix games that allow to characterize the complexity of computing Stackelberg-Nash equilibria (with followers restricted to pure strategies). The key property of these games is that, once fixed the number of players, computing a strong or weak equilibrium presents the same complexity, namely polynomial (again assuming that the followers play pure strategies). These games are of practical interest in security problems. Moreover, they are equivalent to *Bayesian* Stackelberg games with one leader and one follower, where the latter may be of different types. Our first class is equivalent to games with *interdependent types*, while the second one is equivalent to games with *independent types* (i.e., the leader's utility is independent of the follower's type). Thus, every result that holds for a game class also holds for its equivalent class.

We investigate whether the problem keeps being easy when the number of players is *not* fixed. We show that it is NP-hard to compute a weak Stackelberg-Nash equilibrium, and we provide an exact (exponential-time) algorithm (conversely, to compute a strong equilibrium, one can adapt the algorithm provided in [20] for Bayesian games, by means of our mapping). We also prove that, in all the instances where the weak Stackelberg-Nash equilibrium is a supremum but not a maximum, an α-approximation of the supremum can be found in polynomial time (also in the number of players) for any given additive loss $\alpha > 0$. As for approximation complexity, we show that the problem is Poly-APX-complete. This also shows that, in Bayesian Stackelberg games with uncertainty over the follower, computing a weak Stackelberg-Nash equilibrium is as hard as finding a strong one [33].

Next, we investigate whether, in general polymatrix games with followers restricted to play pure strategies, the problem admits polynomial-time approximation algorithms. We provide a negative answer, showing that in the strong case the problem is *not* in Poly-APX if the number of players is non-fixed, unless $P = NP$. We also prove that the same inapproximability result holds for the weak case, even with a fixed number of players.

Our detailed results on Stackelberg polymatrix games appeared in [21] (see [22] for an extended version).

10.3.3   Stackelberg Congestion Games

We provide a comprehensive study of the computational complexity of finding Stackelberg-Nash equilibria in congestion games. These are games with a large

number of players that compete for the use of some shared resources, where the cost of each resource is a function of the number of players using that resource, i.e., its *congestion*. Notice that, in such setting, assuming that the followers play a pure-strategy Nash equilibrium is without loss of generality, as congestion games always admit one [43].

First, we focus on games with *singleton* actions, i.e., where each player selects only one resource at a time. We draw a complete picture of the computational complexity of the problem of finding equilibria in Stackelberg singleton congestion games, with pure or mixed-strategy commitments, and considering the cases of finding either a strong equilibrium or a weak one. Interestingly, we identify two features which allow for thoroughly characterizing hard and easy game instances. The first one concerns the relationship among the action spaces of the players, with two possibilities: the one where the players are *symmetric* as they have identical action spaces and therefore they share the same set of resources, and the one where their action spaces may differ. The second feature is related to the shape of the players' cost functions. Two cases are possible: the one where these functions are *monotonically increasing* in the resource congestion and the one in which they may be not.

In particular, we show that, in games where the players' action spaces can be different, computing a (strong or weak) Stackelberg-Nash equilibrium is *not* in Poly-APX unless P = NP even when the players' cost functions are monotonic, the leader has only one action available, and her costs are equal to the followers'. This result also holds if we restrict the leader to pure-strategy commitments, given that the leader has only one action available. For symmetric games where the players have identical action spaces, we show that the complexity of computing an equilibrium depends on the nature of the players' cost functions. For the case where the players' costs are generic (monotonic or not) functions of the resource congestion, we prove that the problem is *not* in Poly-APX unless P = NP. On the other hand, we show that, in symmetric games, the problem of computing a strong or weak Stackelberg-Nash equilibrium can be solved in polynomial time when the cost functions are monotonic by proposing an algorithm for it. We also consider the case where the leader is restricted to pure-strategy commitments, providing a polynomial-time algorithm for its solution which applies even to symmetric games with generic cost functions. This algorithm is based on a polynomial-time dynamic programming algorithm available in the literature for computing a socially optimal Nash equilibria in non-Stackelberg singleton congestion games with identical action spaces, which we improve and extend to solve our problem.

Then, we switch the attention to games beyond singleton ones. We show that having actions made of only one resource is necessary to have efficient (polynomial-time) algorithms. Indeed, we prove that finding a strong Stackelberg-Nash equilibrium is NP-hard and *not* in Poly-APX unless P = NP, even if players' actions contain only two resources, costs are monotonic, and players are symmetric. We also introduce and study singleton congestion games in which the players are partitioned into *classes*, with followers of the same class sharing the same set of actions. These are a generalization of singleton games with symmetric players, capturing the common case in which users can be split into (usually few) different classes, such as, e.g., users with

different priorities. For these games, we provide a dynamic programming algorithm that computes a strong Stackelberg-Nash equilibrium in polynomial time, when the number of classes is fixed and the leader is restricted to play pure strategies. On the other hand, we prove that, if the leader is allowed to play mixed strategies, then the problem becomes NP-hard even with only four classes and monotonic costs.

Finally, for all the settings we study, we design MILP formulations for computing a strong Stackelberg-Nash equilibrium, and we experimentally evaluate them on a testbed containing both randomly generated game instances and worst-case instances based on our hardness reductions.

The results related to singleton games appeared in [36] and its extended version [15]. Instead, all the other results are provided by [34] (see [35] for an extended version).

10.4   Stackelberg Games with Multiple Leaders

We study games with multiple leaders, providing 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 [19]. 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 on the agreement stage. The first one requires that the leaders do not have any incentive to become followers. It comes in two flavors, called *stability* and *perfect stability*, which are related to, respectively, Nash and subgame perfect equilibria of the sequential game of 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*. They differ depending on the properties they call for. Specifically, (simple) Stackelberg correlated equilibria, Stackelberg correlated equilibria *with perfect agreement*, and Stackelberg correlated equilibria *with perfect agreement and perfect efficiency* 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 Stackelberg correlated equilibria with or without perfect agreement are guaranteed to exist in any game, while Stackelberg correlated equilibria with perfect agreement and perfect efficiency may not. Moreover, we compare the former with other solution concepts, both Stackelberg and non-Stackelberg ones.

Then, we switch the attention to the computational complexity perspective. We show that, provided a suitably defined *stability oracle* is solvable in polynomial time,

a Stackelberg correlated equilibrium 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 *a* Stackelberg correlated equilibrium with perfect agreement, while we prove that computing an optimal one 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 *succinct games of polynomial type* [41]. The problem solved by our oracle is strictly connected with the *weighted deviation-adjusted social welfare problem* introduced by [28]. As a result, we get that our oracle is solvable in polynomial time in all game classes where the same holds for finding an optimal correlated equilibrium.

Our results on Stackelberg games with multiple leaders appeared in [13] (see [14] for an extended version).

10.5   Trembling-Hand Perfection in Stackelberg Games

We study Stackelberg games with a sequential structure, usually referred to as *extensive-form* Stackelberg games. In particular, we show that classical Stackelberg equilibria may prescribe the players to play sub-optimally off the equilibrium path, as it is the case for the Nash equilibrium. Thus, in order to amend these weaknesses, we propose a way to refine Stackelberg equilibria thorough trembling-hand perfection, which is based on the idea that each player might play each action with low-but-non-zero probabilities, usually called *trembles* [45].

We show that for every perturbation scheme (i.e., any possible way of introducing trembles), the set of limit points of Stackelberg equilibria for perturbed games with vanishing perturbations is always a nonempty subset of the Stackelberg equilibria of the non-perturbed game. This does not hold when focusing only on strong (or weak) equilibria: for a given game, the set of strong Stackelberg equilibria (or weak Stackelberg equilibria) in the non-perturbed game may be disjoint from the set of limit points of strong Stackelberg equilibria (or weak Stackelberg equilibria) in the perturbed game. We resort to the perturbation schemes used for *quasi-perfect* equilibria [48] and *extensive-form* perfect equilibria [45] to define their Stackelberg counterpart—and their strong and weak versions—as refinements of the Stackelberg equilibrium.

Next, we focus on quasi-perfection. We formally define the *quasi-perfect Stackelberg equilibrium* refinement game theoretically in the same axiomatic fashion as the quasi-perfect equilibrium was defined for non-Stackelberg games [48]. Thus, our definition is based on a set of properties of the players' strategies, and it cannot be directly used to search for a quasi-perfect Stackelberg equilibrium. Subsequently, we define a class of perturbation schemes for the sequence form such that any limit point of a sequence of Stackelberg equilibria in perturbed games with vanishing perturbation is a quasi-perfect Stackelberg equilibrium. This class of perturbation schemes strictly includes those used to find a quasi-perfect equilibrium by [40]. Then, we

extend the algorithm by [16] to the case of quasi-perfect Stackelberg equilibrium computation. We derive the corresponding mathematical program for computing a *Stackelberg extensive-form correlated equilibrium* when a perturbation scheme is introduced and we discuss how the individual steps of the algorithm change. In particular, the implementation of our algorithm is much more involved, requiring the combination of branch-and-bound techniques with arbitrary-precision arithmetic to deal with small perturbations. This does not allow a direct application of off-the-shelf solvers. Finally, we experimentally evaluate the scalability of our algorithm.

In conclusion, we also study the computational complexity of finding Stackelberg equilibrium refinements, showing that the problem of deciding the existence of a Stackelberg equilibrium—refined or not—that gives the leader expected value at least ν is NP-hard.

Our results appeared in [25] and [38] (see [37] for an extended version).

References

1. Amaldi E, Capone A, Coniglio S, Gianoli LG (2013) Network optimization problems subject to max-min fair flow allocation. IEEE Commun Lett 17(7):1463–1466
2. An B, Pita J, Shieh E, Tambe M, Kiekintveld C, Marecki J (2011) Guards and protect: next generation applications of security games. ACM SIGecom Exchanges 10(1):31–34
3. Avenhaus R, Okada A, Zamir S (1991) Inspector leadership with incomplete information. In: Game equilibrium models IV, pp 319–361. Springer, Berlin
4. Basilico N, Coniglio S, Gatti N (2017) Methods for finding leader–follower equilibria with multiple followers. arXiv:1707.02174
5. Basilico N, Coniglio S, Gatti N, Marchesi A (2017) Bilevel programming approaches to the computation of optimistic and pessimistic single-leader-multi-follower equilibria. In: SEA, vol 75, pp 1–14. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH. Dagstuhl Publishing, Wadern
6. Basilico N, Coniglio S, Gatti N, Marchesi A (2019) Bilevel programming methods for computing single-leader-multi-follower equilibria in normal-form and polymatrix games. EURO J Comput Optim 8:1–29
7. von Stengel B, Zamir S (2010) Leadership games with convex strategy sets. Games Econ Behav 69(2):446–457
8. Bošanský B, Cermak J (2015) Sequence-form algorithm for computing stackelberg equilibria in extensive-form games. In: Proceedings of the twenty-ninth AAAI conference on artificial intelligence (AAAI 2015), pp 805–811 (2015)
9. Brown N, Sandholm T (2018) Superhuman ai for heads-up no-limit poker: libratus beats top professionals. Science 359(6374):418–424
10. Brown N, Sandholm T (2019) Superhuman ai for multiplayer poker. Science 365(6456):885–890
11. Murray C, Joseph Hoane A Jr, Hsu F-H (2002) Deep blue. Artif Intell 134(1–2):57–83
12. Caprara A, Carvalho M, Lodi A, Woeginger GJ (2016) Bilevel knapsack with interdiction constraints. INFORMS J Comput 28(2):319–333
13. Castiglioni M, Marchesi A, Gatti N (2019) Be a leader or become a follower: the strategy to commit to with multiple leaders. In: Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI 2019, Macao, China, August 10–16, 2019, pp 123–129
14. Castiglioni M, Marchesi A, Gatti N (2019) Be a leader or become a follower: the strategy to commit to with multiple leaders (extended version). CoRR, arxiv:abs/1905.13106

15. Castiglioni M, Marchesi A, Gatti N, Coniglio S (2019) Leadership in singleton congestion games: what is hard and what is easy. Artif Intell 277:103177
16. Cermak J, Bošanský B, Durkota K, Lisý V, Kiekintveld C (2016) Using correlated strategies for computing stackelberg equilibria in extensive-form games. In: Proceedings of the thirtieth AAAI conference on artificial intelligence (AAAI 2016)
17. Coniglio S, Gatti N, Marchesi A (2017) Pessimistic leader-follower equilibria with multiple followers. In: Proceedings of the twenty-sixth international joint conference on artificial intelligence, IJCAI-17, pp 171–177
18. Stefano C, Nicola G, Alberto M (2020) Computing a pessimistic stackelberg equilibrium with multiple followers: the mixed-pure case. Algorithmica 82:1189–1238
19. Conitzer V, Korzhyk D (2011) Commitment to correlated strategies. In: Proceedings of the twenty-fifth AAAI conference on artificial intelligence (AAAI 2011), pp 632–637
20. Conitzer V, Sandholm T (2006) Computing the optimal strategy to commit to. In: Proceedings of the 7th ACM conference on electronic commerce, pp 82–90
21. De Nittis G, Marchesi A, Gatti N (2018) Computing the strategy to commit to in polymatrix games. In: Proceedings of the thirty-second AAAI conference on artificial intelligence (AAAI-18), pp 989–996
22. De Nittis G, Marchesi A, Gatti N (2018) Computing the strategy to commit to in polymatrix games (extended version). CoRR, arxiv:abs/1807.11914
23. Farina G, Gatti N (2017) Adopting the cascade model in ad auctions: efficiency bounds and truthful algorithmic mechanisms. J Artif Intell Res 59:265–310
24. Farina G, Gatti N, Sandholm T (2018) Practical exact algorithm for trembling-hand equilibrium refinements in games. In: Advances in neural information processing systems, pp 5039–5049
25. Farina G, Marchesi A, Kroer C, Gatti N, Sandholm T (2018) Trembling-hand perfection in extensive-form games with commitment. In: Proceedings of the twenty-seventh international joint conference on artificial intelligence, IJCAI 2018, July 13–19, 2018, Stockholm, Sweden, pp 233–239
26. Gan J, Elkind E, Wooldridge M (2018) Stackelberg security games with multiple uncoordinated defenders. In: Proceedings of the 17th international conference on autonomous agents and multiagent systems, pp 703–711. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC
27. Gatti N, Lazaric A, Rocco M, Trovò F (2015) Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. Artif Intell 227:93–139
28. Xin Jiang A, Leyton-Brown K (2011) A general framework for computing optimal correlated equilibria in compact games. In: International workshop on internet and network economics, pp 218–229. Springer, Berlin
29. Kiekintveld C, Jain M, Tsai J, Pita J, Ordóñez F, Tambe M (2009) Computing optimal randomized resource allocations for massive security games. In: AAMAS, pp 689–696
30. Labbé M, Marcotte P, Savard G (1998) A bilevel model of taxation and its application to optimal highway pricing. Manag Sci 44(12-part-1):1608–1622
31. Labbé M, Violin A (2016) Bilevel programming and price setting problems. Ann Operat Res 240(1):141–169
32. Letchford J, Conitzer V (2010) Computing optimal strategies to commit to in extensive-form games. In: Proceedings of the 11th ACM conference on electronic commerce, pp 83–92. ACM, New Year
33. Letchford J, Conitzer V, Munagala K (2009) Learning and approximating the optimal strategy to commit to. In: International symposium on algorithmic game theory, pp 250–262. Springer, Berlin
34. Marchesi A, Castiglioni M, Gatti N (2019) Leadership in congestion games: multiple user classes and non-singleton actions. In: Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI 2019, Macao, China, August 10–16, 2019, pp 485–491
35. Marchesi A, Castiglioni M, Gatti N (2019) Leadership in congestion games: multiple user classes and non-singleton actions (extended version). CoRR, arxiv:abs/1905.13108

36. Marchesi A, Coniglio S, Gatti N (2018) Leadership in singleton congestion games. In: Proceedings of the twenty-seventh international joint conference on artificial intelligence, IJCAI 2018, July 13–19, 2018, Stockholm, Sweden, pp 447–453
37. Marchesi A, Farina G, Kroer C, Gatti N, Sandholm T (2018) Quasi-perfect stackelberg equilibrium. CoRR, arxiv:abs/1811.03871
38. Marchesi A, Farina G, Kroer C, Gatti N, Sandholm T (2019) Quasi-perfect stackelberg equilibrium. In: The thirty-third AAAI conference on artificial intelligence, AAAI 2019, Honolulu, Hawaii, USA, January 27–February 1, 2019, pp 2117–2124
39. Matuschke J, McCormick ST, Oriolo G, Peis B, Skutella M (2017) Protection of flows under targeted attacks. Oper Res Lett 45(1):53–59
40. Miltersen PB, Sørensen TB (2010) Computing a quasi-perfect equilibrium of a two-player game. Econ Theory 42(1):175–192
41. Papadimitriou CH, Roughgarden T (2008) Computing correlated equilibria in multi-player games. J ACM (JACM) 55(3):14
42. Paruchuri P, Pearce JP, Marecki J, Tambe M, Ordonez F, Kraus S (2008) Playing games for security: an efficient exact algorithm for solving bayesian stackelberg games. In: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems, pp 895–902
43. Rosenthal RW (1973) A class of games possessing pure-strategy nash equilibria. Int J Game Theory 2(1):65–67
44. Sandholm WH (2002) Evolutionary implementation and congestion pricing. Rev Econ Stud 69(3):667–689
45. Selten R (1975) Reexamination of the perfectness concept for equilibrium points in extensive games. Int J Game Theory 4(1):25–55
46. Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M, et al (2016) Mastering the game of go with deep neural networks and tree search. Nature 529(7587):484
47. Milind T (2011) Security and game theory: algorithms, deployed systems, lessons learned. Cambridge University Press, Cambridge
48. Van Damme E (1984) A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. Int J Game Theory 13(1):1–13
49. Van Damme E (1987) Stability and perfection of nash equilibria. Springer, Berlin, Heidelberg
50. Von Stackelberg H (1934) Marktform und gleichgewicht. Verlag von Julius Springer, Berlin