

Leadership in Singleton Congestion Games: What is Hard and What is Easy

Matteo Castiglioni, Alberto Marchesi, Nicola Gatti

Politecnico di Milano, Italy

Stefano Coniglio

University of Southampton, UK

Abstract

We study the problem of computing *Stackelberg equilibria* in *Stackelberg games* whose underlying structure is a *congestion game*, focusing on *singleton congestion games*, i.e., on congestion games where each player can choose a single resource, and assuming that one of them acts as leader while the other ones act as followers. We provide a comprehensive picture of the computational complexity of finding equilibria in these games, investigating different forms of commitment (*pure-strategy* and *mixed-strategy*) and followers' tie-breaking rules (*optimistic* and *pessimistic*). We identify two features of such games, namely, *identical/different* action spaces and *monotonic/generic* cost functions, by which we provide a complete characterization of the cases in which the equilibrium-finding problem is either easy or hard. In particular, we show that, in the case where the action spaces are different, the cost the leader incurs in an optimistic or pessimistic Stackelberg equilibrium cannot be approximated in polynomial time up to any polynomial factor in the size of the game unless $P = NP$, independently of the cost functions being monotonic or generic. This result holds even when the commitment is restricted to pure strategies. For general mixed-strategy commitments, we show that a similar result also holds when the players have generic cost functions, even if their action spaces are identical. Furthermore, we prove that the case with identical action spaces and monotonic cost functions is easy. We also improve the efficiency of a polynomial-time algorithm available in the literature for the computation of a socially optimal Nash equilibrium in non-Stackelberg singleton congestion games with identical action spaces and generic cost functions, and we extend it to the computation of a Stackelberg equilibrium for the case where the leader is restricted to playing pure strategies. For the cases in which the problem of finding an equilibrium is hard, we show how, in

*A preliminary version of this work appeared in [1]. This extended version contains new computational complexity proofs for problems addressed only partially in the preliminary version, proofs of results concerning problems not addressed previously, new algorithms, and their experimental evaluation.

the optimistic setting where the followers break ties in favor of the leader, the problem can be formulated via mixed-integer linear programming techniques. We also provide an experimental evaluation of our algorithms both on random instances and on instances generated from our inapproximability reductions.

Keywords: Algorithmic Game Theory, Stackelberg Equilibria, Congestion Games, Computational Complexity, Bilevel Programming

1. Introduction

The Stackelberg paradigm was originally introduced by von Stackelberg in 1934 [2] to model economic situations where a firm moves first and, then, another firm moves second by reacting to the first firm’s move. Over the last years, *Stackelberg Games* (SGs) and the corresponding *Stackelberg Equilibria* (SEs) have received growing attention in the artificial intelligence literature, where the computational problem of finding SEs is often referred to as the problem of *computing optimal strategies to commit to* [3]. In their most classical form, SGs involve a player who, acting as leader, has the ability to commit to playing a strategy beforehand, while the other player, acting as follower, can play only after observing the leader’s commitment [4]. The leader’s commitment can be either pure (*à la* von Stackelberg [2]) if the follower can observe the action that the leader draws from her commitment or mixed (*à la* Conitzer and Sandholm [3] and von Stengel and Zamir [4]) if the follower cannot.

While, traditionally, the majority of the works on SGs have focused on the *single-leader single-follower* case, the variant with more than two players has recently started to receive more attention: See [5, 6, 7] for the *single-leader multi-follower* case, [8, 9, 10, 11, 12] for the *multi-leader single-follower* case, and [13, 14, 15, 16] for the *multi-leader multi-follower* case.

When multiple equilibria can arise in the follower’s game for a single a leader strategy, two main variants are typically considered: the *Optimistic Stackelberg Equilibrium* (OSE), achieved when the followers break ties in favor of the leader, and the *Pessimistic Stackelberg Equilibrium* (PSE), achieved when they break ties against her.¹

Applications of SGs can be found in, among others, security games [18, 19, 20, 21], toll-setting games [22, 23], interdiction games [24, 25], network routing [26], inspection games [27], and mechanism design [28].

In this work, we focus on the Stackelberg paradigm applied to multi-player SGs whose underlying structure is that of *Congestion Games* (CGs) (see [29, 30, 31, 32], and [33] for a recent survey), which model scenarios involving multiple players competing for the use of shared resources.

In a CG, the players’ actions are subsets of a given set of resources and the costs the players incur depend (monotonically or not) on the level of resource

¹In the literature, optimistic and pessimistic SEs are also known as *strong* and *weak* SEs, respectively, following the terminology of [17] which first introduced this distinction.

utilization, which is typically referred to as *resource congestion*. Despite their
35 simplicity, CGs model a wide spectrum of real-world problems including, among
others, job scheduling [34], facility location [35], and network routing [36]. Crucially, CGs always admit *Nash Equilibria* (NEs) in which the players play pure strategies [29]. Moreover, such equilibria are always achievable in a finite number of iterations by *best-response dynamics*, i.e., by applying an iterative procedure
40 by which, at each iteration, a player changes her action and switches to playing a best response to the actions currently played by the other players [37].

CG games can be characterized according to the combinatorial structure of the players' action spaces, *such as the case* where the players' actions are paths connecting a source to a destination [38] or spanning trees [39]. An interesting
45 case *on which we will focus in the paper* is that of *Singleton CGs* (SCGs) [31], i.e., CGs where each player selects a single resource.² Computing an NE in SCGs is easy [30] and, for the case in which all the players have the same action space (we will refer to these games as *symmetric*), finding an NE minimizing the social cost is also easy [31].

50 Except for some sporadic works which applied the Stackelberg paradigm to games with a CG-like structure where the leader is a welfare-maximizing central authority (see, e.g., [40] and its extensions), a systematic study of the Stackelberg paradigm applied to CGs is still missing.

In this paper, we focus on Stackelberg SCGs, i.e., on the fundamental case
55 of Stackelberg CGs in which each player selects a single resource, assuming that a single player takes the role of leader and commits to a strategy defined over her resources before the other players decide how to play. We *independently address the two cases where the leader commits to a pure or a mixed strategy, assuming that,* after observing the (mixed or pure) strategy the leader has
60 committed to, the followers always reach a pure-strategy NE. This is w.l.o.g. since, as the followers' game is a CG, it always admits at least a pure-strategy NE. For the sake of generality, we assume that the leader's costs may differ from the followers'.

Since in general SGs the leader may get a higher utility/incur a smaller cost
65 by committing to a mixed strategy [4], one may wonder whether this is also the case in Stackelberg SCGs. Among other findings, the results that we present in the paper show under which conditions this is the case.

1.1. Applications

Introduced in [34], one of the simplest problems which can be modeled as
70 Stackelberg SCGs is a job scheduling problem where the users (players) select which machines (resources) have to execute their jobs (such as in virtualized environments or data centers). The time needed to complete a job on a machine (the resource cost) depends on its workload (the resource congestion). Assuming a single job per player to be executed on a single machine without preemption,

²See [30] for a generalization of SCGs where each player's action space is expressed as a *matroid* defined over the set of resources.

75 the players' actions are singletons and the problem fits in the setting of SCGs. The case of a single-leader Stackelberg SCG arises when one of the players is the owner of the machines and is willing to share her resources with the others, but, being the owner, she decides which resource/machine to pick before the others do. Under the assumption that the players schedule their jobs in an open-loop
80 fashion, i.e., without any knowledge of the current congestion of the machines, it is plausible that the followers could not observe the machine on which the leader's job is running. It is therefore natural, for the leader, to try and achieve a smaller cost by committing to a mixed strategy.

Another application can be found in facility location problems [35, 41] where
85 the players are firms and they have to decide on which site to locate their infrastructures (which, depending on the specific application, may be, e.g., factories, shops, or mineral extraction plants). Each firm selects a location from a list of candidate sites (the resources) and the cost it incurs depends on the number of firms that made the same choice. In these scenarios, the single-leader
90 Stackelberg SCG case arises whenever one of the firms either has a competitive advantage over the others (due to, e.g., being a governmental agency) or, as, in the job scheduling problem, it owns the candidate sites and, thus, can decide before the other ones where to locate its infrastructures. Mixed-strategy commitments are plausible when the time between the choice of the location and the
95 beginning of the construction works of the facility is extremely long, due to, e.g., administrative issues and/or the time needed for obtaining the authorizations. In this case, the follower could prefer to choose her location before observing the beginning of the construction works of the leader's facility to avoid incurring an excessive delay w.r.t. the leader.

100 *It is worth pointing out that*, due to constituting the simplest scenario in which the Stackelberg paradigm is applied to a CG, *Stackelberg SCGs are* the natural starting point to tackle more general settings involving arbitrary actions and multiple leaders.

1.2. Original contributions

105 In this paper, we draw a complete picture of the computational complexity of the problem of finding an SE in SCGs (i.e., CGs where each player can select a single resource) with one leader and multiple followers and pure or mixed-strategy commitments, considering the cases of finding either an OSE or a PSE. Interestingly, we identify two features of SCGs 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
110 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. Moreover, for each of the four cases identified by these two independent features, we provide a concrete application of our model, motivating the study of the four scenarios beyond their purely theoretical interest.

Table 1: Summary of the original contributions on the problem of computing an O/PSE in Stackelberg SCGs with a single leader provided in this paper.

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

120 Table 1 summarizes the original results that we provide with this work.
 In particular, we show that, in SCGs where the players’ action spaces can be
 different, computing an O/PSE 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
 125 the leader to pure-strategy commitments, given that the leader has only one
 action available. Moreover, this also shows that, as we will better explain in
 the following, the same inapproximability result also holds for the problem of
 computing, in the same game setting, an NE which minimizes/maximizes the
 cost incurred by any given player, a result which may be of independent interest.

130 For the symmetric games where the players have identical action spaces, we
 show that the complexity of computing an O/PSE 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
 135 that, in symmetric games, the problem of computing an O/PSE can be solved
 in polynomial time when the cost functions are monotonic by proposing an
 algorithm for it. While proving the soundness of the algorithm is straightforward
 when the leader’s commitment is a pure strategy, the analysis is more involved
 with mixed-strategy commitments. As we will show, our result follows from
 140 the fact that mixed-strategy commitments do not allow the leader to incur a
 cost smaller than the one she incurs with a pure strategy. 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-
 145 time dynamic programming algorithm available in the literature for computing
 a socially optimal NE in non-Stackelberg SCGs with identical action spaces,
 which we improve and extend to solve our problem.

Finally, we provide two mathematical programming formulations to compute
 an OSE for games with different action spaces in, at most, exponential time, and
 150 a more compact one for the symmetric games. Experimentally, we also evaluate
 the scalability of the two formulations when fed to a state-of-the-art *Mixed-
 Integer Linear Programming* (MILP) solver and compare their performance—in
 terms of computing time and solution efficiency—to simple heuristic algorithms
 based on best-response dynamics.

155 1.3. Structure of the work

The remainder of the paper is organized as follows. Section 2 provides a
 review of the literature on the computation of SEs in single and multi-leader
 SGs. Section 3 introduces basic concepts and the notation we use, including the
 formal definitions of the game models we consider, also providing some concrete
 160 examples where they can be applied. Section 4 provides the main hardness
 results for the problem of computing an O/PSE in games where the players’
 action spaces are different. Section 5 does the same in games where the play-
 ers’ action spaces are identical, but the cost functions are generic. Section 6

establishes which problems can be solved efficiently, providing the corresponding polynomial-time algorithms. Section 7 proposes mathematical programming formulations for computing an OSE in the intractable cases. Section 8 provides experimental results assessing the scalability of the proposed formulations. Finally, Section 9 concludes the work summarizing the results and pointing out directions for future research. Full proofs are reported in Appendix A, while supplementary experimental results are provided in Appendix B.

2. State of the Art on Stackelberg Equilibrium computation

In this section, we provide a comprehensive review of the literature on SE computation. First, we focus on SGs with a single leader and one or more followers. Then, we conclude the section with a brief discussion on the works dealing with games involving multiple leaders and followers.

2.1. Stackelberg games with a single leader

When focusing on SGs involving one leader and one or more followers, we can classify most of the works in the literature according to two crucial aspects: the followers' tie-breaking rule (which distinguishes between OSEs and PSEs) and the structure of the underlying followers' game.

Table 2 summarizes most of the known computational results for the problem of computing SEs in single-leader SGs according to these two features.

The problem is known to be easy in 2-player normal-form games in both the optimistic and the pessimistic setting, as shown in, respectively, [3] and [4]. In particular, [5] shows that the problem of computing an OSE can be formulated as a single *Linear Program* (LP), while [4] illustrates that a PSE can be computed by solving a polynomial number of LPs.

When one considers the case of n -player normal-form games with $n \geq 3$, many cases are possible, depending on how the followers behave after observing the leader's commitment. A reasonable choice, which has been widely investigated in the literature, is to assume that they play simultaneously and noncooperatively, reaching an NE. In this case, it is known that finding an O/PSE is not in Poly-APX unless $P = NP$ even when there are only two followers (i.e., with $n = 3$) [42]. Computing an OSE becomes easy for any n if we restrict the followers to playing pure strategies only, as it requires the solution of an LP for each outcome of the followers' game, whose number is polynomial in the size of the game representation [7]. On the other hand, computing a PSE is still NP-hard even with only two followers playing pure strategies [6] and it is not in Poly-APX unless $P = NP$ with at least three followers (i.e., with $n \geq 4$) [7].

As for algorithms, the authors of [43, 44, 42] show how to formulate the problem of finding an OSE in n -player normal-form games as a *nonlinear* and *nonconvex* mathematical program, which they solve via spatial branch-and-bound techniques. As shown in [6], when the followers are restricted to play pure strategies a PSE can be found by employing an algorithm which solves multiple *lexicographic Mixed-Integer Linear Programs* (lex-MILPs), which, as

the authors show, can be further enhanced by embedding it in a branch-and-bound scheme.

As to works on n -player normal-form games where the followers do not play an NE, [5] shows that finding an OSE is easy when the followers can play correlated strategies, while [3] proves that the problem becomes NP-hard if the
210 followers play in a hierarchical fashion.

Besides normal-form games, the literature has devoted considerable attention to Bayesian 2-player normal-form games where the follower can be of different types, mainly due to their relevance in security games. In this setting, it is known
215 that finding an OSE is Poly-APX-complete [45] and that an equilibrium can be found by solving a MILP [18]. As recently shown in [46], the same hardness result also holds for the problem of computing a PSE. [46] also provides an algorithm for computing an equilibrium via the solution of exponentially-many LPs, without resorting to the normal-form representation.

Over the last years, the Stackelberg paradigm has also been applied to 2-player extensive-form games. In particular, the authors of [47] prove that finding an OSE is NP-hard even in games without nature. The results are extended by the authors of [48], who prove that computing a PSE is also NP-hard. Works
220 such as [49, 50] address the problem of computing an OSE in extensive-form games, providing worst-case exponential-time algorithms based on MILPs. In the context of extensive-form games, attempts have also been made towards the refinement of SEs. In particular, the authors of [51] introduce the idea of a *robust* SE, where an optimal commitment is found against a worst-case follower's utility model. Pursuing a different approach, the authors of [48] and [52] show
225 how to guarantee an optimal commitment off the equilibrium path by adopting the idea of *trembling-hand perfection* to the Stackelberg setting.

Other works attempted to relax the general structure of normal-form games, trying to identify games with many players where SEs can be efficiently computed. Along this line of research, the authors of [46] analyze n -player *polymatrix*
235 *games*, which are games where the players interact pairwise and each player takes part to a 2-player normal-form game with each of the other players. For these games, [42] shows that, when the followers play mixed strategies, finding an O/PSE is not in Poly-APX unless $P = NP$. The result is extended in [46], where the authors show that it holds even when the followers are restricted to
240 pure strategies. While, for fixed n , finding an OSE with the followers playing pure strategies is easy, the same does not hold for PSEs, as the problem is hard even with only three followers (i.e., with $n = 4$) [53].

For works studying other game models, such as stochastic games and Bayesian signaling games, we refer the reader to [54, 55, 56].

While the Stackelberg paradigm has already been applied to games with a structure resembling that of CGs, the only works which, to our knowledge, pursue this line of research are [40] and its extensions [57, 58, 59]. We remark,
245 though, that the author of [40] considers a different Stackelberg paradigm where the leader is an authority whose objective is to minimize the social cost of the NE reached by the followers. Differently, in this work, as also done in all the papers
250 we mentioned above, we adopt the classical Stackelberg paradigm introduced

by [4] and [3], assuming that the leader is a special player who has the ability to commit to a strategy beforehand with the aim of minimizing her own cost.

2.2. Stackelberg games with multiple leaders

255 SGs with multiple leaders are widely unexplored. In this setting, the key aspect by which the works in the literature can be categorized is how they model the leaders' commitments and, in particular, on what kind of equilibrium constraints they enforce on them.

260 The vast majority of the works on multi-leader SGs is restricted to the field of security games and address games with multiple uncoordinated defenders (the leaders) and a single attacker (the follower) [8, 9, 10, 11, 12]. All these works introduce models which enforce NE-like constraints on the leaders' commitments and, in general, they all suffer from the drawback that the resulting SEs may not exist.

265 Recently, the authors of [13] studied the Stackelberg paradigm with multiple leaders and followers in SGs with an arbitrary structure. Their framework generalizes the one introduced in [5] by considering the case in which the leaders' commitments are governed by a correlation device which must obey some particular constraints which, in turn, guarantee that an equilibrium always exists.

270 The operations research literature provides further works on multi-leader-follower settings—some of which can be found in the literature on *mathematical programs with equilibrium constraints* [60]. Such works assume that both leaders and followers be subject to NE-like constraints, with the latter playing in the game which results after observing the leaders' strategies [14, 15, 16]. Other works from the same field focus on oligopoly models where the leaders select the level of investment to maximize their profits [61, 62].

3. Preliminaries

280 First, we introduce our notation and provide the formal definitions of Stackelberg SCGs and their equilibria. We conclude the section with some application examples showing how the games we consider map to real-world problems.

3.1. Stackelberg singleton congestion games and their symmetric variant

Adopting the notation introduced in [63], we provide the following formal definition of the class of games we study:

285 **Definition 1** (Stackelberg SCG (SSCG)). *A Stackelberg SCG (SSCG) is a tuple (N, R, A, c_ℓ, c_f) , where:*

- $N = F \cup \{\ell\}$ is a finite set of players, ℓ being the leader and F the set of followers;
- R is a finite set of resources;
- $A = \{A_p\}_{p \in N}$ is the set of all players' actions, with $A_p \subseteq R$, for each $p \in N$, being the set of actions of player p ;

Table 2: Summary of the results on the computation of SEs in normal-form games, Bayesian normal-form games, extensive-form games, and polymatrix games. In all cases, we assume that the leader is allowed to commit to mixed strategies.

Followers' strategies		Optimistic		Pessimistic	
		Pure	Mixed	Pure	Mixed
Normal-form games					
$n = 2$	Complexity Algorithm	P [3] multi-LP [3]		P [4] multi-LP [4]	
$n = 3$	Complexity Algorithm	P [7] multi-LP [7]	NP-hard, \notin Poly-APX [42] spatial branch-and-bound [44]	NP-hard [6] multi-lex-MILP [6]	NP-hard, \notin Poly-APX [42] –
$n \geq 4$	Complexity Algorithm	P [7] multi-LP [7]	NP-hard, \notin Poly-APX [42] spatial branch-and-bound [44]	NP-hard [6], \notin Poly-APX [7] multi-lex-MILP [6]	NP-hard, \notin Poly-APX [42] –
Bayesian normal-form games					
$n = 2$	Complexity Algorithm	NP-hard [3], Poly-APX-complete [45] MILP [18]		NP-hard, Poly-APX-complete [46] multi-LP [46]	
Extensive-form games					
$n = 2$	Complexity Algorithm	NP-hard [54] MILP [49, 50]		NP-hard [48] multi-LP [4]	
Polymatrix games					
$n = 3$	Complexity Algorithm	P [7] multi-LP [7]	NP-hard, \notin Poly-APX [42] spatial branch-and-bound [44, 42]	NP-hard multi-lex-MILP [7]	NP-hard, \notin Poly-APX [42] –
$n \geq 4$ (fixed)	Complexity Algorithm	P [7] multi-LP [7]	NP-hard, \notin Poly-APX [42] spatial branch-and-bound [44, 42]	NP-hard, \notin Poly-APX [53] multi-lex-MILP [7]	NP-hard, \notin Poly-APX [42] –
$n \geq 4$ (free)	Complexity Algorithm	NP-hard, \notin Poly-APX [46] multi-LP [7]	NP-hard, \notin Poly-APX [42] spatial branch-and-bound [44]	NP-hard, \notin Poly-APX [53] multi-lex-MILP [7]	NP-hard, \notin Poly-APX [42] –

- $c_\ell = \{c_{i,\ell}\}_{i \in R}$ and $c_f = \{c_{i,f}\}_{i \in R}$ are, respectively, the leader's and followers' cost functions, with $c_{i,\ell}, c_{i,f} : \mathbb{N} \rightarrow \mathbb{Q}$ being the costs of resource i as a function of its congestion for, respectively, the leader and the followers.

We denote by n and r the number of players and of resources, i.e., $n := |N|$ and $r := |R|$. As usual, we assume $c_{i,\ell}(0) = c_{i,f}(0) = 0$ for every $i \in R$. We call the players' cost functions (*weakly*) *monotonic* if, for every resource $i \in R$, $c_{i,\ell}(x) \leq c_{i,\ell}(x+1)$ and $c_{i,f}(x) \leq c_{i,f}(x+1)$ for all $x \in \mathbb{N}$, and *strictly monotonic* if all the inequalities are strict. Whenever the inequalities are not satisfied, we say that the players' cost functions are *nonmonotonic*.

We call *strategy* of player $p \in N$ a probability distribution σ_p over A_p , where $\sigma_p(a_p)$ denotes the probability that $a_p \in A_p$ is played. Let Δ_p be the set of player p 's strategies. A strategy $\sigma_p \in \Delta_p$ is said *pure* if it prescribes player p to always play some action $a_p \in A_p$, i.e., if $\sigma_p(a_p) = 1$ and $\sigma_p(a'_p) = 0$ for all $a'_p \in A_p \setminus \{a_p\}$. Otherwise, σ_p is said *mixed*. A collection of strategies is called *strategy profile* in general and *action profile* if all the strategies it contains are pure. In this work, we collectively denote by $\sigma = (\sigma_\ell, a)$ a strategy profile in which the leader plays a (possibly) mixed strategy $\sigma_\ell \in \Delta_\ell$ and the followers play the pure strategies contained in the action profile $a = (a_p)_{p \in F} \in \times_{p \in F} A_p$.

Let $a = (a_p)_{p \in F} \in \times_{p \in F} A_p$ be a followers' action profile. We denote by $\nu_i^a = |\{p \in F \mid a_p = i\}|$ the number of followers selecting resource $i \in R$ in a . This quantity is equal to the resource congestion caused by the followers' presence only. We call *followers' configuration* (induced by action profile a) the vector $\nu^a \in \mathbb{N}^r$ whose i -th component is ν_i^a for all $i \in R$.

For any $\sigma_\ell \in \Delta_\ell$, we define the followers' expected cost for resource $i \in R$ given the leader's strategy σ_ℓ as the function $c_{i,f}^{\sigma_\ell} : \mathbb{N} \rightarrow \mathbb{Q}$. Specifically, $c_{i,f}^{\sigma_\ell}$ is a function of the number $x \in \mathbb{N}$ of followers who select resource i , i.e.:

$$c_{i,f}^{\sigma_\ell}(x) = \sigma_\ell(i)c_{i,f}(x+1) + (1 - \sigma_\ell(i))c_{i,f}(x). \quad (1)$$

Note that, given a leader's strategy σ_ℓ and a followers' congestion x , all the followers who select resource $i \in R$ experience a congestion that may (with probability $\sigma_\ell(i)$) or may not (with probability $1 - \sigma_\ell(i)$) be incremented by one w.r.t. x , depending on whether the leader chooses resource i or not. Given the strategy profile $\sigma = (\sigma_\ell, a)$, the leader's cost is:

$$c_\ell^\sigma = \sum_{i \in A_\ell} \sigma_\ell(i)c_{i,\ell}(\nu_i^a + 1). \quad (2)$$

Note that, whenever the leader selects resource $i \in R$ (which happens with probability $\sigma_\ell(i)$), her costs depends on the followers' congestion ν_i^a plus an additional unit of congestion which due to her choosing that resource.

After observing a leader's commitment $\sigma_\ell \in \Delta_\ell$, the followers play an SCG where the resource costs are specified by the functions $c_{i,f}^{\sigma_\ell}$, for $i \in R$. We assume that, after witnessing the leader's commitment, the followers play a pure-strategy NE. This is always possible as, being a CG, the new SCG played by the followers always admits one [29].

Given a strategy profile $\sigma = (\sigma_\ell, a)$, the followers' action profile $a = (a_p)_{p \in F}$ is an NE for the leader's strategy σ_ℓ if the following holds:

$$c_{a_p, f}^{\sigma_\ell}(\nu_{a_p}^a) \leq c_{a'_p, f}^{\sigma_\ell}(\nu_{a'_p}^a + 1) \quad \forall p \in F, a'_p \in A_p.$$

In words, a is an NE for σ_ℓ if no follower has an incentive to unilaterally deviate from a_p by selecting another resource a'_p . In the following, for any given $\sigma_\ell \in \Delta_\ell$ we let E^{σ_ℓ} be the set of NEs in the followers' game resulting from σ_ℓ .

325 We also consider *symmetric SSCGs* (SSSCGs), a subclass of SSCGs in which every player can select every resource, i.e., where $A_p = R$ for all $p \in N$:

Definition 2 (Symmetric SSCG (SSSCG)). *We call an SSCG defined by a tuple (N, R, A, c_ℓ, c_f) symmetric if $A = \times_{p \in F} R$.*

Note that, in an SSSCG, all the followers are identical due to being allowed to choose the same resources. Thus, only the number of followers selecting each resource is significant and, therefore, the action profile a of any follower can be equivalently represented by the followers' configuration ν^a it induces. As a consequence, when studying SSSCGs we do not explicitly refer to the followers' action profiles but, rather, to their configurations $\nu \in \mathbb{N}^r$, with $\sum_{i \in R} \nu_i = n - 1$. When working with followers' configurations rather than with action profiles, with a little abuse of notation we adopt $\sigma = (\sigma_\ell, \nu)$ to denote a strategy profile. Adopting this notation, we have that ν is an NE for the leader's strategy $\sigma_\ell \in \Delta_\ell$ if the following holds:

$$c_{i, f}^{\sigma_\ell}(\nu_i) \leq c_{j, f}^{\sigma_\ell}(\nu_j + 1) \quad \forall i \in R : \nu_i > 0, j \in R.$$

3.2. Stackelberg equilibria in SSCGs

330 Given a leader's strategy, the followers' SCG may admit multiple NEs. As customary in the literature, we consider two extreme cases which lead to the definition of OSEs and PSEs [4]. In the first one, we assume that the followers act in favor of the leader, playing an NE minimizing her cost. In the second one, we assume that the followers act against the leader, always playing an
335 NE which results in the maximum leader's cost. Let us remark that we are not assuming that the followers can coordinate so as to reach a particular NE in the game resulting from the leader's commitment. The leader's cost at an OSE corresponds to the smaller cost the leader may incur under the optimistic assumption that the followers would (somehow) end up playing an NE which
340 minimizes the leader's cost. Differently, the leader's cost at a PSE corresponds to a cost value the leader could always incur independently of the followers' behavior. Formally:

Definition 3. *A strategy profile $\sigma = (\sigma_\ell, a)$ is an OSE if it solves the following bilevel programming problem:*

$$\min_{\sigma_\ell \in \Delta_\ell} \min_{a \in E^{\sigma_\ell}} c_\ell^{(\sigma_\ell, a)}.$$

As it is clear, an OSE always exists in SSCGs and, since the same objective function is minimized in both levels, the problem can be equivalently rewritten as the following single-level problem:

$$\min_{\sigma_\ell \in \Delta_\ell, a \in E^{\sigma_\ell}} c_\ell^{(\sigma_\ell, a)}.$$

Definition 4. A PSE, if it exists, is a strategy profile $\sigma = (\sigma_\ell, a)$ which solves the following bilevel problem:

$$\min_{\sigma_\ell \in \Delta_\ell} \max_{a \in E^{\sigma_\ell}} c_\ell^{(\sigma_\ell, a)}.$$

Let us recall that, in general, the problem in Definition 4 may not admit a minimum (but only an infimum) and, thus, a PSE may not exist [4].

345 **Proposition 1.** *There are SSSCGs in which a PSE does not exist.*

Proof. Consider the following instance of an SSSCGs (whose cost functions are reported in the table below), where $|F| = 1$ and $R = \{r_1, r_2\}$.

x	$c_{r_1, \ell}$	$c_{r_1, f}$	$c_{r_2, \ell}$	$c_{r_2, f}$
1	2	1	2	1
2	0	2	2	2

350 Clearly, the single follower selects r_1 if $\sigma_\ell(r_1) < \frac{1}{2}$, she chooses r_2 if $\sigma_\ell(r_1) > \frac{1}{2}$, and she is indifferent between r_1 and r_2 if $\sigma_\ell(r_1) = \frac{1}{2}$. Thus, the leader's cost is $2 - 2\sigma_\ell(r_1)$ if $\sigma_\ell(r_1) < \frac{1}{2}$, while it is 2 if $\sigma_\ell(r_1) \geq \frac{1}{2}$, since, in the pessimistic case, the follower selects r_2 rather than r_1 when $\sigma_\ell(r_1) = \frac{1}{2}$. As a result, the problem in Definition 4 achieves an infimum with value 1 at $\sigma_\ell(r_1) = \frac{1}{2}$, but it does not admit a minimum. Thus, the game does not admit a PSE. \square

355 3.3. Examples of applications

Our characterization of hard and easy SSCG instances is based on two independent features of these games, namely, identical/different action spaces and monotonic/generic cost functions.

360 We illustrate how the different scenarios identified by these two features arise in the applications we mentioned in Subsection 1.1.

- *Identical action spaces.* In job scheduling problems, the action spaces are identical whenever the jobs can be executed on any machine. Similarly, in facility location problems, this happens if the firms can build their facilities on any candidate site, without restrictions.
- 365 • *Different action spaces.* In job scheduling problems, the action spaces are different whenever each job can be executed only on the machines which meet the job's technical specifications (such as, e.g., running a certain operating system or being equipped with CPUs supporting a certain instruction set architecture). In facility location problems, the case of

370 different action spaces arises whenever each firm is willing to build its facility only on a subset of the candidate sites such as, e.g., those located in geographical areas which are more likely to meet the firm's projected demand.

- 375 • *Monotonic cost functions.* In both the job scheduling and the facility location problems, the case of monotonic costs is natural. In job scheduling, the time required to execute a job is likely to increase with the number of jobs simultaneously running on the machine (since, e.g., there is an overhead due to process switching). In facility location, the cost functions are monotonic if, e.g., the firms compete for different market segments (one per candidate site) and the larger is the number of firms selecting the same segment the smaller is the market share of each firm (and, thus, the larger the profit loss).
- 380 • *Generic cost functions.* In facility location problems, the case of generic cost functions arises if we consider costs which depend (additively) on both the profit loss and on the cost for the construction works, under the assumption that the latter be equal to a fixed cost (incurred when the construction works start) plus a variable cost decreasing with the congestion level due to economies of scale (since, e.g., building two facilities on the same site is likely to be cheaper than building them at different ones, as it allows for sharing workers between them).
- 385
- 390

4. SSCGs NP-hardness and inapproximability

Let us start our analysis with a negative result, showing that the problem of computing an O/PSE in SSCGs with different action spaces is computationally intractable even if the leader can only select a single resource and all the costs are monotonic functions of the resource congestion. This also shows that, in general non-Stackelberg SCGs with different action spaces, computing an NE which either maximizes or minimizes the usage of a resource (or the cost incurred by a player) is hard, which may be of independent interest. Moreover, given that our intractability results hold even when the leader has only one resource available, computing an O/PSE in SSCGs with different action spaces is intractable even if we restrict the leader to pure-strategy commitments.

First, we prove that finding an OSE is not in Poly-APX unless $P = NP$ using a reduction from 3SAT. As a result, the leader's cost in an OSE cannot be approximated, in polynomial time, up to any approximation factor which depends polynomially on the size of the game given as input, unless $P = NP$. Then, we show that the same intractability result holds for computing a PSE by means of a different reduction still based on 3SAT.

For the main results in this section (namely, Theorems 1 and 2) we only report proof sketches. The full proofs are provided in Appendix A.

410 *4.1. Computational complexity of finding an OSE in SSCGs*

We analyze the problem of computing an OSE in SSCGs in the general case with different action spaces. The hardness and inapproximability results that we present are based on a reduction from 3SAT (see [64] for its NP-completeness), which is defined as follows:

415 **Definition 5 (3SAT).** *Given a finite set C of 3-literal clauses defined over a finite set V of Boolean variables, is there a truth assignment to the variables which satisfies all the clauses?*

In the following, let $l \in \phi$ denote a literal (i.e., a variable or its negation) appearing in clause $\phi \in C$ and $v(l) \in V$ denote the variable corresponding to that literal. Moreover, given a 3SAT instance (C, V) , let m and s be, respectively, 420 the number of clauses and variables, i.e., $m := |C|$ and $s := |V|$.

We introduce our reduction in the proof of the following theorem.

Theorem 1. *Computing an OSE in SSCGs with different action spaces is NP-hard, even if the leader has only one action (i.e., she can only select a single 425 resource) and the cost functions are monotonic.*

Proof sketch. We provide a reduction from 3SAT. Specifically, given a 3SAT instance (C, V) and a real number $0 < \epsilon < 4$, we build an instance $\Gamma_\epsilon(C, V)$ of an SSCG admitting an OSE $\sigma = (\sigma_\ell, a)$ in which $c_\ell^\sigma = \epsilon$ if and only if (C, V) is satisfiable; if not, $c_\ell^\sigma = 4$ in any OSE. $\Gamma_\epsilon(C, V)$ has one resource r_ϕ for each 430 clause $\phi \in C$ and four resources, namely r_v , $r_{v,t}$, $r_{\bar{v}}$, and $r_{\bar{v},t}$, for each variable $v \in V$. It also has resources $r_{\phi,v}$ and $r_{\phi,\bar{v}}$ for every $\phi \in C$ and $v \in V$. Moreover, there is an additional resource r_t . Figure 1 shows an example of the instance $\Gamma_\epsilon(C, V)$. The following table specifies the cost functions, where, additionally, $c_{r_{\bar{v}},f} = c_{r_v,f}$, $c_{r_{\phi,\bar{v}},f} = c_{r_{\phi,v},f}$, $c_{r_{\bar{v},t},f} = c_{r_{v,t},f}$, and $c_{r_t,f} = c_{r_t,\ell}$ (let us remark 435 that, given $\epsilon < 4$, the costs are monotonic functions of the resource congestion):

x	$c_{r_\phi,f}$	$c_{r_v,f}$	$c_{r_{v,t},f}$	$c_{r_{\phi,v},f}$	$c_{r_t,f}$
1	2	0	0	1	ϵ
$[2, m]$	5	0	6	6	4
$[m+1, \infty]$	5	7	6	6	4

The leader can only select resource r_t (i.e., $\sigma_\ell(r_t) = 1$) and $c_\ell^\sigma = \epsilon$ if and only if $\nu_{r_t}^a = 0$, while $c_\ell^\sigma = 4$ otherwise. For each variable $v \in V$, the gadget composed of r_v , $r_{v,t}$, $r_{\bar{v}}$, and $r_{\bar{v},t}$ is built in such a way that, in order to have $\nu_{r_t}^a = 0$, either 440 all the m followers $p_{v,k}$ (with $k \in \{1, \dots, m\}$) select r_v (v is false) or all the m followers $p_{\bar{v},k}$ (with $k \in \{1, \dots, m\}$) select $r_{\bar{v}}$ (v is true). Due to how the costs are defined, in the first case each follower $p_{\phi,v}$ (one per clause $\phi \in C$) prefers $r_{\phi,v}$ over r_v while, in the other case, each follower $p_{\phi,\bar{v}}$ (for every $\phi \in C$) prefers $r_{\phi,\bar{v}}$ rather than $r_{\bar{v}}$. For each clause $\phi \in C$, there is a follower $p_{\phi,t}$ who has to choose between r_ϕ and r_t , preferring the latter unless she is the only follower 445 selecting r_ϕ . Moreover, there is a follower p_ϕ who can choose between r_ϕ and the resources $r_{\phi,l}$ for all literals $l \in \phi$, preferring the former only if she is the only one selecting it. In order to have $\nu_{r_t}^a = 0$, for every clause $\phi \in C$ $p_{\phi,t}$ must

therefore select r_ϕ and p_ϕ must choose one of the resources $r_{\phi,l}$. Due to how
 450 the costs are defined, this happens only if $p_{\phi,l}$ selects r_l , which is the case only
 if $\nu_{r_l}^a \leq m$. This implies that $v(l)$ is true if l is positive, while it is false if l
 is negative, thus satisfying clause ϕ . Thus, since the same argument holds for
 every $\phi \in C$, $c_\ell^\sigma = \epsilon$ if and only if the 3SAT instance is satisfiable. \square

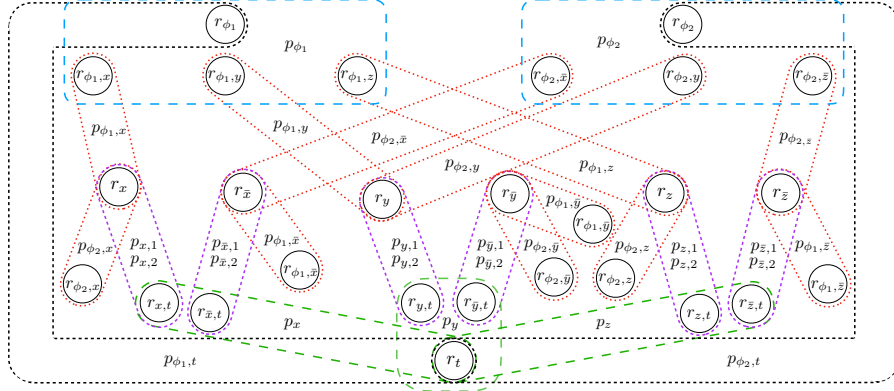


Figure 1: Example of a game instance $\Gamma_\epsilon(C, V)$ used in the reduction in the proof of Theorem 1, with $V = \{x, y, z\}$, $C = \{\phi_1, \phi_2\}$, $\phi_1 = x \vee y \vee z$, and $\phi_2 = \bar{x} \vee y \vee \bar{z}$.

The proof of Theorem 1 also shows the following:

455 **Corollary 1.** *In general non-Stackelberg SCGs with different action spaces, computing an NE minimizing the cost of a given player (or the usage of a given resource) is NP-hard even if the cost functions are monotonic.*

Proof. The result is easily proved by noticing that, in the $\Gamma_\epsilon(C, V)$ games defined in the proof of Theorem 1, since the leader can only use a single resource
 460 any OSE $\sigma = (\sigma_\ell, a)$ is also an NE. Thus, given that the followers behave optimistically, such games admit an NE with $c_\ell^\sigma = \epsilon$ if and only if the corresponding 3SAT instance is satisfiable; otherwise $c_\ell^\sigma = 4$ in any NE. As a result, due to 3SAT being NP-complete, computing an NE minimizing the cost of a given player (the leader) is NP-hard. Since $c_\ell^\sigma = \epsilon$ if and only if $\nu_{r_t}^a = 0$, the same
 465 holds for the problem of finding an NE which minimizes the usage of a given resource. \square

Theorem 1 also implies that the leader's cost in an OSE cannot be efficiently approximated up to any factor which depends polynomially on the size of the input:

470 **Corollary 2.** *The problem of computing an OSE in SSCGs with different action spaces is not in Poly-APX unless $P = NP$ even if the leader has only one action and the cost functions are monotonic.*

Proof. Given a 3SAT instance (C, V) , let us build an SSCG $\Gamma_\epsilon(C, V)$ as in the proof of Theorem 1. We have already proved that $\Gamma_\epsilon(C, V)$ admits an
475 OSE $\sigma = (\sigma_\ell, a)$ in which $c_\ell^\sigma = \epsilon$ if and only if (C, V) is satisfiable and that, otherwise, $c_\ell^\sigma = 4$. Let $\epsilon = \frac{4}{2^{n+r}}$. Assume that there exists a polynomial-time approximation algorithm \mathcal{A} with approximation factor $\text{poly}(n, r)$, i.e., a polynomial function of n and r . Assume (C, V) is satisfiable. \mathcal{A} applied to $\Gamma_\epsilon(C, V)$ would return a solution with $c_\ell^\sigma \leq \frac{4}{2^{n+r}} \text{poly}(n, r)$. Since, for n and r
480 large enough, $\frac{4}{2^{n+r}} \text{poly}(n, r) < 4$, \mathcal{A} would allow us to decide in polynomial time whether (C, V) is satisfiable, a contradiction unless $P = NP$. \square

Since that the intractability results, in Theorem 1 and Corollary 2 hold even when the leader can select only a single resource, we also obtain the following:

Corollary 3. *The problem of computing an OSE in SSCGs with different action
485 spaces is NP-hard and not in Poly-APX unless $P = NP$ even if we restrict the leader to pure-strategy commitments.*

Since the followers break ties in favor of the leader in the reduction, the results in Theorem 1 and Corollaries 2 and 3 do not apply to the problem of finding a PSE. We consider this case in the next subsection.

4.2. Computational complexity of finding a PSE in SSCGs 490

The hardness and inapproximability results that we are about to present for the problem of computing a PSE in SSCGs with different action spaces are still based on 3SAT but rely on a different reduction.

Theorem 2. *Computing a PSE in SSCGs with different action spaces is NP-
495 hard even if the leader has only one action and the cost functions are monotonic.*

Proof sketch. We provide a reduction from 3SAT with a structure similar to the one used in Theorem 1 for the optimistic setting. In this case, given a 3SAT instance (C, V) and a real number $0 < \epsilon < 4$, we build $\Gamma_\epsilon(C, V)$ so that it admits a PSE $\sigma = (\sigma_\ell, a)$ in which $c_\ell^\sigma = \epsilon$ if and only if (C, V) is not satisfiable;
500 otherwise, if (C, V) is satisfiable, $c_\ell^\sigma = 4$ holds in any PSE. Figure 2 shows an example of game $\Gamma_\epsilon(C, V)$. The following table specifies the cost functions, where, additionally, $c_{r_{\bar{v}}, f} = c_{r_v, f}$, $c_{r_{\phi, \bar{v}}, f} = c_{r_{\phi, v}, f}$, and $c_{r_t, f} = c_{r_t, \ell}$ (let us remark that, given $\epsilon < 4$, the costs are monotonic functions of the the resource congestion):

x	$c_{r_\phi, f}$	$c_{r_v, f}$	$c_{r_{v,t}, f}$	$c_{r_{\phi, v}, f}$	$c_{r_t, f}$
1	2	1	2	0	ϵ
$[2, m]$	5	1	5	7	ϵ
$m+1$	5	6	5	7	ϵ
$[m+s+1, \infty]$	5	6	5	7	4

As in the other reduction, the leader can only choose resource r_t , but, in this case, $c_\ell^\sigma = 4$ holds if and only if $\nu_{r_t}^a = m+s$ (i.e., all the followers $p_{\phi, t}$, for $\phi \in C$, and $p_{v, t}$, for $v \in V$, select resource r_t), while $c_\ell^\sigma = \epsilon$ otherwise. Differently from

the other reduction, for every variable $v \in V$ the gadget corresponding to v is composed of three resources r_v , $r_{\bar{v}}$, and $r_{v,t}$. It is built in such a way that, whenever follower $p_{v,t}$ selects resource r_t , at least one between p_v (if v is false) and $p_{\bar{v}}$ (if v is true) must choose $r_{v,t}$, preferring it over the other resource r_v (or $r_{\bar{v}}$). Given how the costs are defined, each follower $p_{\phi,v}$ prefers $r_{\phi,v}$ over r_v whenever r_v is chosen by p_v . Similarly, each follower $p_{\phi,\bar{v}}$ is better off selecting $r_{\phi,\bar{v}}$ rather than $r_{\bar{v}}$ if $p_{\bar{v}}$ chooses $r_{\bar{v}}$. For each clause $\phi \in C$, in addition to follower $p_{\phi,t}$ we have three other followers $p_{l,\phi}$, one per literal $l \in \phi$ (differently from the other reduction, where there is only p_ϕ). Whenever follower $p_{\phi,t}$ selects r_t , at least one follower $p_{l,\phi}$ must choose r_ϕ , otherwise $p_{\phi,t}$ would decrease her cost by switching to r_ϕ . Since, given how the costs are defined, $p_{l,\phi}$ prefers being the only follower using $r_{\phi,l}$ rather than selecting r_ϕ , follower $p_{\phi,l}$ must choose $r_{\phi,l}$. This implies that p_l would select resource r_l , and, thus, l evaluates to true, satisfying clause ϕ . Since the same argument holds for every $\phi \in C$, $c_\ell^\sigma = 4$ holds if and only if the 3SAT instance is satisfiable. \square




Figure 2: Example of a game instance $\Gamma_\epsilon(C, V)$ used in the reduction in the proof of Theorem 2 with $V = \{x, y, z\}$, $C = \{\phi_1, \phi_2\}$, $\phi_1 = x \vee y \vee z$, and $\phi_2 = \bar{x} \vee y \vee \bar{z}$.

Theorem 2 also implies the following:

Corollary 4. *In general non-Stackelberg SCGs with different action spaces, computing an NE maximizing the cost of a given player (or the usage of a given resource) is NP-hard even if the cost functions are monotonic.*

Proof. In games $\Gamma_\epsilon(C, V)$ such as those used in the proof of Theorem 2, any PSE is also an NE (since the leader can choose a single action). Moreover, $\Gamma_\epsilon(C, V)$ admits a PSE $\sigma = (\sigma_\ell, a)$ in which $c_\ell^\sigma = 4$ if and only if the given 3SAT instance has answer *no*, otherwise $c_\ell^\sigma = \epsilon$. This proves the result for the problem of finding an NE maximizing the cost of a given player. Since $c_\ell^\sigma = 4$ if and only if $\nu_{r_t}^a = m + s$, the same holds for the problem of computing an NE maximizing the usage of a given resource. \square

535 Furthermore, from Theorem 2 it directly follows that the leader's cost in a PSE cannot be efficiently approximated up to any approximation factor which depends polynomially on the size of the input:

Corollary 5. *The problem of computing a PSE in SSCGs with different action spaces is not in Poly-APX unless $P = NP$, even if the leader has only one action and the cost functions are monotonic.*
540

Proof. Given a 3SAT instance (C, V) , let us build an instance $\Gamma_\epsilon(C, V)$ of SSCG as in the proof of Theorem 2. We have already proven that $\Gamma_\epsilon(C, V)$ admits a PSE $\sigma = (\sigma_\ell, a)$ in which $c_\ell^\sigma = \epsilon$ if and only if the 3SAT instance has answer *no*; otherwise, $c_\ell^\sigma = 4$ in any PSE. Let $\epsilon = \frac{4}{2^{n+r}}$. Assume that there
545 exists a polynomial-time approximation algorithm \mathcal{A} with approximation factor $\text{poly}(n, r)$, i.e., a polynomial function of n and r . Assume the answer to the 3SAT instance is *no*. \mathcal{A} applied to $\Gamma_\epsilon(C, V)$ would return a solution with $c_\ell^\sigma \leq \frac{4}{2^{n+r}} \text{poly}(n, r)$. Since, for n and r large enough, $\frac{4}{2^{n+r}} \text{poly}(n, r) < 4$, \mathcal{A} would allow us to decide in polynomial time whether the answer to the 3SAT
550 instance is *yes* or *no*, a contradiction unless $P = NP$. \square

Since in the reduction the leader only has one resource available we can conclude the following:

Corollary 6. *The problem of computing a PSE in SSCGs with different action spaces is NP-hard and not in Poly-APX unless $P = NP$ even if we restrict the leader to pure-strategy commitments.*
555

5. SSSCGs NP-hardness and inapproximability

In this section, we focus on SSSCGs (the subset of SSCGs in which the players have identical action spaces), showing that the problem of finding an O/PSE in such games is NP-hard and not in Poly-APX unless $P = NP$. This
560 result matches the other result that we have established for the problem of computing an O/PSE in general SSCGs with different action spaces. For SSSCGs, the inapproximability result relies on the nonmonotonicity of the players' cost functions and on the leader's ability to commit to mixed strategies. This must necessarily be the case since, as we will show in Section 6, the problem is easy
565 when the costs functions are monotonic and the players are symmetric (Theorem 9), and the same holds even with generic cost functions if we restrict the leader to pure-strategy commitments (Theorem 11 and its Corollary 7).

For the problem of computing an OSE, we rely on a reduction from K -PARTITION, a variant of PARTITION with an additional size constraint,
570 whereas we adopt a different reduction based on the classical version of PARTITION for the problem of computing a PSE. The two problems are defined as follows:

Definition 6 (PARTITION). *Given a finite set $S = \{x_1, \dots, x_{|S|}\}$ of positive integers $x_i \in \mathbb{Z}^+$ with $\sum_{i \in S} x_i$ even, is there a partition $(S', S \setminus S')$ of S , with
575 $S' \subseteq S$, such that $\sum_{x_i \in S'} x_i = \sum_{x_i \in S \setminus S'} x_i$?*

Definition 7 (*K*-PARTITION). Given a finite set $S = \{x_1, \dots, x_{|S|}\}$ of positive integers $x_i \in \mathbb{Z}^+$ with both $|S|$ and $\sum_{i \in S} x_i$ even and a positive integer $K \leq \frac{|S|}{2}$, is there a partition $(S', S \setminus S')$ of S , with $S' \subseteq S$ and $|S'| = K$ such that $\sum_{x_i \in S'} x_i = \sum_{x_i \in S \setminus S'} x_i$?

580 Letting $s = \frac{1}{2} \sum_{x_i \in S} x_i$, we assume for both problems that $x_i \leq s$ for all $x_i \in S$. Indeed, if some $x_i > s$ then $\sum_{x_i \in S'} x_i > s$ holds for every $S' \subseteq S$ and, thus, the answer to both PARTITION and *K*-PARTITION is trivially *no*.

PARTITION is well-known to be NP-complete [64]. To see that *K*-PARTITION is also NP-complete (its membership to NP is clear), it suffices to observe that
585 PARTITION has answer *yes* if and only if *K*-PARTITION has answer *yes* for some $K \in \{1, \dots, \frac{|S|}{2}\}$. This gives us a simple Cook reduction from PARTITION to *K*-PARTITION: after solving *K*-PARTITION $\frac{|S|}{2}$ times, once per value of $K \in \{1, \dots, \frac{|S|}{2}\}$, if answer *yes* is found for some K , PARTITION has answer *yes*; if, instead, answer *yes* is never found, PARTITION has answer *no*.

590 For the main results in this section (namely, Theorems 3 and 4 for the problem of finding an OSE, and Theorems 5 and 6 for the problem of finding a PSE) we only provide proof sketches. Full proofs are provided in Appendix A.

5.1. Computational complexity of finding an OSE in SSSCGs

We start our analysis with the problem of computing an OSE in SSSCGs.
595 We introduce our main reduction in the proof of the following theorem.

Theorem 3. *Computing an OSE in SSSCGs is NP-hard.*

Proof sketch. We provide a reduction from *K*-PARTITION. Given a finite set S of positive integers and a positive integer $K \leq \frac{|S|}{2}$, we build an SSSCG instance with $4|S| + 2$ followers, a resource r_i for each $x_i \in S$, and two additional resources r_{t_1} and r_{t_2} . Letting $w_i = \frac{x_i}{s}$ for all $x_i \in S$, the players' costs are the following ones:

x	$c_{r_i, f}$	$c_{r_i, \ell}$	$c_{r_{t_1}, f}$	$c_{r_t, \ell}$	$c_{r_{t_2}, f}$	$c_{r_{t_2}, \ell}$
1	$2s$	s	$3s^2$	s^4	1	s^4
2	0	s	$3s^2$	s^4	$4s^2$	s^4
3	$\frac{1}{w_i}$	ϵ	$3s^2$	s^4	$4s^2$	s^4
4	$\frac{2s - \frac{1}{w_i} + 1}{w_i}$	s	$3s^2$	s^4	$4s^2$	s^4
$[5, 4 S - 2K]$	$4s^2$	s	$3s^2$	s^4	$4s^2$	s^4
$4 S - 2K + 1$	$4s^2$	s	$2s$	s^4	$4s^2$	s^4
$4 S - 2K + 2$	$4s^2$	s	1	s^4	$4s^2$	s^4
$[4 S - 2K + 3, \infty]$	$4s^2$	s	0	s^4	$4s^2$	s^4

The costs are defined in such a way that there is an OSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$ if and only if the *K*-PARTITION instance has answer *yes*. Specifically,
605 in order to have $c_\ell^\sigma = \epsilon$ the followers' configuration must be such that $\nu_{r_{t_1}} =$

$4|S| - 2K + 1$, $\nu_{r_{t_2}} = 1$, and $\nu_{r_i} = 2$ for K different resources r_i , whereas the leader's strategy must satisfy $\sigma_\ell(r_i) = \frac{1}{w_i}$ for these resources. This is an OSE as c_ℓ^σ is minimized and the followers behave optimistically. Intuitively, if $\nu_{r_{t_1}}^a < 4|S| - 2K + 1$ then the followers selecting r_{t_1} would have an incentive to deviate to a resource r_i (incurring a smaller cost) while, if $\nu_{r_{t_1}}^a > 4|S| - 2K + 1$, at least one follower would switch to resource r_{t_1} , paying at most 1. For similar reasons, exactly one follower must choose resource r_{t_2} . Moreover, all the followers who do not choose r_{t_1} or r_{t_2} must select a resource r_i with $\nu_{r_i} = 2$, otherwise their cost would be greater than 1 and they would have an incentive to deviate to r_{t_1} . In order to have $c_\ell^\sigma = \epsilon$, the leader's strategy σ_ℓ must place positive probability only on resources r_i with $\nu_{r_i} = 2$, i.e., $\sigma_\ell(r_i) = w_i$ must hold for these resources. If $\sigma_\ell(r_i) < w_i$, a follower would deviate from r_{t_1} to r_i paying $(1 - \sigma_\ell(r_i))\frac{1}{w_i} + w_i\frac{2s - \frac{1}{w_i} + 1}{w_i} < 2s$ while, if $\sigma_\ell(r_i) > w_i$, the followers selecting r_i would pay $\frac{1}{w_i}\sigma_\ell(r_i) > 1$ and they would prefer to switch to r_{t_1} . Thus, letting S' be the set of $x_i \in S$ with $\sigma_\ell(r_i) = w_i$, $\sum_{x_i \in S'} \sigma_\ell(r_i) = 1$ implies $\sum_{x_i \in S'} w_i = 1$: thus, $\sum_{x_i \in S'} x_i = s$ and $(S', S \setminus S')$ is a partition of S . \square

Next, we show that even approximating the leader's cost in an OSE up to any polynomial factor of the input size is hard.

Theorem 4. *The problem of computing an OSE in SSCGs is not in Poly-APX unless $P = NP$.*

Proof sketch. The result is based on the same reduction from K -PARTITION used for Theorem 3. Indeed, it can be proved that, if the K -PARTITION instance has answer *no*, then $c_\ell^\sigma > 1$ holds for any OSE $\sigma = (\sigma_\ell, a)$ (see Appendix A for the technical details). The result is then proved by arguments similar to those employed in Corollaries 2 and 5. \square

5.2. Computational complexity of finding a PSE in SSSCGs

We focus now on the problem of computing a PSE in SSSCGs. The proof of the following theorem introduces our main reduction.

Theorem 5. *Computing a PSE in SSSCGs is NP-hard.*

Proof sketch. We provide a reduction from PARTITION. Given a finite set S of positive integers, we build an SSSCG instance with $3|S|$ followers, a resource r_i for each $x_i \in S$, and an additional resource r_t . Letting $w_i = \frac{x_i}{s}$ for all $x_i \in S$, the players' costs are the following ones:

x	$c_{r_i, f}$	$c_{r_i, \ell}$	$c_{r_t, f}$	$c_{r_t, \ell}$
1	0	ϵ	1	s^4
2	$\frac{1}{w_i - \frac{1}{s^4}}$	s^4	1	s^4
3	$\frac{1}{1 - w_i - \frac{1}{s^4}}$	ϵ	1	s^4
4	0	s^4	1	s^4
$[5, \infty]$	s	ϵ	1	s^4

640 The costs are defined in such a way that there is a PSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$ if and only if the PARTITION instance has answer *yes*. Given that the followers behave pessimistically, in the game resulting from σ_ℓ all the NEs must provide the leader with a utility $c_\ell^\sigma \leq \epsilon$. In order to have $c_\ell^\sigma = \epsilon$, the leader's strategy σ_ℓ must place positive probability only on resources r_i that

645 satisfy one between $\nu_{r_i} = 0$, $\nu_{r_i} = 2$, and $\nu_{r_i} \geq 4$. For each resource r_i , if $\sigma_\ell(r_i) < w_i - \frac{1}{s^4}$ then $\nu_{r_i} = 1$ holds in an NE of the resulting followers' game, since the follower selecting r_i would not have an incentive to deviate (as her cost would be $\frac{1}{w_i - \frac{1}{s^4}} \sigma_\ell(r_i) < 1$). Moreover, if $\sigma_\ell(r_i) > w_i + \frac{1}{s^4}$ then $\nu_{r_i} = 3$ holds in an NE of the resulting followers' game since the followers selecting r_i would

650 incur a cost of $\frac{1}{1 - w_i - \frac{1}{s^4}} (1 - \sigma_\ell(r_i)) < 1$ and no other follower would deviate to r_i as $s \sigma_\ell(r_i) > 1$. Thus, for every resource r_i it must be the case that either $w_i - \frac{1}{s^4} < \sigma_\ell(r_i) < w_i + \frac{1}{s^4}$ or $\sigma_\ell(r_i) = 0$. Since each $x_i \in S$ is integer and $\sum_{x_i \in S} \sigma_\ell(r_i) = 1$, we can conclude that the set S' of all the $x_i \in S$ such that $w_i - \frac{1}{s^4} < \sigma_\ell(r_i) < w_i + \frac{1}{s^4}$ defines a PARTITION. \square

655 Finally, we show that the same inapproximability result that we have established for OSEs also holds for PSEs.

Theorem 6. *The problem of computing a PSE in SSCGs is not in Poly-APX unless $P = NP$.*

Proof sketch. The result is based on the same reduction from PARTITION used for Theorem 5. Further analyses allow us to prove that, if the PARTITION

660 instance has answer *no*, then $c_\ell^\sigma > 1$ holds in any PSE $\sigma = (\sigma_\ell, \nu)$ (see Appendix A for the technical details). The claim is then proved via arguments similar to those we employed in Corollaries 2 and 5. \square

6. Polynomial-time algorithms for SSSCGs

665 In the previous sections, we have shown that the problem of computing an O/PSE in SSCGs is, both in the general case and when restricting ourselves to SSSCGs, computationally intractable. We provide, here, two positive results for SSSCGs, showing that, under certain conditions, the computation of an O/PSE in these games can be carried out in polynomial time.

670 First, we design a polynomial-time algorithm for finding an O/PSE in SSSCGs where the players' costs are monotonic functions of the resource congestion. The algorithm relies on the fact that, as we will show, in such games the leader cannot decrease her cost by playing mixed strategies and, thus, pure-strategy commitments are sufficient. We also exhibit a few examples showing that our

675 algorithm cannot be easily extended to more general settings as, if the players have either different action spaces or nonmonotonic cost functions, the leader could be better off playing mixed strategies, thus violating the fundamental assumption of our algorithm.

Finally, we show that, if we restrict our attention to pure-strategy commitments in SSSCGs, an O/PSE can be found in polynomial time by means of a

680

Dynamic Programming (DP) algorithm, even when the players' cost functions are generic.

6.1. Polynomial-time algorithms for computing an O/PSE in SSSCGs with monotonic cost functions

Let us recall that, in SSSCGs, an NE minimizing the social cost can be computed in polynomial time [31]. It is also easy to show that an NE minimizing/maximizing the cost incurred by one player can be found efficiently using an algorithm similar to that of [31] (see Section 6.3 for additional details). As a consequence, computing an O/PSE would also be easy if an equilibrium could only be induced by a leader's pure-strategy commitment. This is, unfortunately, not the case, as there are SSSCGs admitting O/PSEs in which the leader's commitment is a mixed strategy and the followers' configuration could only be induced by the leader committing to a mixed strategy.

Proposition 2. *There are SSSCGs with strictly monotonic cost functions which admit an O/PSE $\sigma = (\sigma_\ell, \nu)$ where the leader's strategy σ_ℓ is mixed and, additionally, the followers' configuration ν is an NE only for mixed-strategy commitments of the leader.*

Proof. Consider the following SSSCG with strictly monotonic cost functions where $|F| = 3$ and $R = \{r_1, r_2, r_3\}$.

x	$c_{r_1,\ell}$	$c_{r_1,f}$	$c_{r_2,\ell}$	$c_{r_2,f}$	$c_{r_3,\ell}$	$c_{r_3,f}$
1	1	1	3	4	1	1
2	2	3	4	5	2	3
3	3	6	5	6	3	6

The followers configuration $\nu = (1, 1, 1)^T$ in which each follower selects a different resource is not an NE if the leader commits to a pure strategy while, for instance, it is an NE for $\sigma_\ell(r_1) = \sigma_\ell(r_3) = \frac{1}{2}$ and $\sigma_\ell(r_2) = 0$. Moreover, notice that the game admits O/PSEs in which the leader's commitment is a mixed strategy. For instance, for $\sigma_\ell(r_1) = \sigma_\ell(r_3) = \frac{1}{2}$ and $\sigma_\ell(r_2) = 0$ the leader incurs a cost of 2 and there is no other strategy which allows her to pay less than 2. \square

Next, we focus on finding OSEs in SSSCGs with weakly monotonic cost functions. We show that, for every strategy profile $\sigma = (\sigma_\ell, \nu)$ where the leader's commitment σ_ℓ is a mixed strategy and ν is an NE in the followers' game, there is another strategy profile $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ where the leader's commitment $\hat{\sigma}_\ell$ is a pure strategy and for which her cost is no larger than the one for σ . In particular, this implies that, in order to achieve an OSE, the leader can always commit w.l.o.g. to a pure strategy. This is formalized in the following theorem, whose proof shows constructively how to build $\hat{\sigma}$ from σ .

The idea of the proof is the following. Let $\sigma = (\sigma_\ell, \nu)$ be a strategy profile where σ_ℓ is a mixed strategy and ν is an NE in the followers' game. Assume that the leader switched to selecting with probability one any of the resources for which she would incur the minimum cost when committing to σ_ℓ . As this would

increase the congestion of that resource, due to the cost functions being weakly monotonic the followers could only react by switching to another resource—this translates in the leader incurring a cost on that resource which is never larger than the one she would incur when committing to the mixed strategy σ_ℓ .

Theorem 7. *Every SSSCG with weakly monotonic cost functions admits an OSE $\sigma = (\sigma_\ell, \nu)$ in which σ_ℓ is pure.*

Proof. Given a strategy profile $\sigma = (\sigma_\ell, \nu)$ with σ_ℓ mixed and $\nu \in E^{\sigma_\ell}$, we show how to construct another strategy profile $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ with $\hat{\nu} \in E^{\hat{\sigma}_\ell}$ in which $\hat{\sigma}_\ell$ is pure and $c_\ell^{\hat{\sigma}} \leq c_\ell^\sigma$. Let $S = \{i \in R \mid \sigma_\ell(i) > 0\}$ be the set of resources played by the leader with positive probability in σ_ℓ and let $i^* \in \arg \min_{i \in S} c_{i,\ell}(\nu_i + 1)$. Clearly, since the leader's utility is a convex combination weighted by σ_ℓ of the costs she incurs in the resources chosen with positive probability, $c_\ell^\sigma = \sum_{i \in A_\ell} \sigma_\ell(i) c_{i,\ell}(\nu_i + 1) \geq c_{i^*,\ell}(\nu_{i^*} + 1)$. Moreover, since ν is an NE for σ_ℓ , the following holds by definition:

$$c_{i,f}^{\sigma_\ell}(\nu_i) \leq c_{j,f}^{\sigma_\ell}(\nu_j + 1) \quad \forall i \in R : \nu_i > 0, j \in R. \quad (3)$$

Let us define $\hat{\sigma}_\ell \in \Delta_\ell$ such that $\hat{\sigma}_\ell(i^*) = 1$. We now show that such $\hat{\sigma}_\ell$ is part of an OSE. Notice that $c_{i,f}^{\hat{\sigma}_\ell}(x) = c_{i,f}(x) \quad \forall x \in \mathbb{N}$ for every $i \in R \setminus \{i^*\}$ (as the leader does not select these resources), while $c_{i^*,f}^{\hat{\sigma}_\ell}(x) = c_{i^*,f}(x + 1) \quad \forall x \in \mathbb{N}$ (as the leader selects that resource). Since the followers behave optimistically, it is sufficient to exhibit a $\hat{\nu} \in E^{\hat{\sigma}_\ell}$ such that $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ satisfies $c_\ell^{\hat{\sigma}} \leq c_\ell^\sigma$. We construct a sequence of followers configurations which starts from ν and reaches such $\hat{\nu}$. Given $\hat{\sigma}_\ell$, let us consider the sequence $(\nu(0) = \nu, \nu(1), \dots, \nu(T) = \hat{\nu})$ such that each configuration differs from the previous one in that a single follower has changed resource, strictly decreasing her cost. Formally, this corresponds to showing that, for all $0 \leq t < T$, there is a pair $i, j \in R$ such that $\nu(t)_i > 0$, $\nu(t+1)_i = \nu(t)_i - 1$, $\nu(t+1)_j = \nu(t)_j + 1$, and $c_{i,f}^{\hat{\sigma}_\ell}(\nu(t)_i) > c_{j,f}^{\hat{\sigma}_\ell}(\nu(t+1)_j)$. Moreover, let us assume that a follower deviates to resource i^* , i.e., $\nu(t+1)_{i^*} > \nu(t)_{i^*}$, only if this is the only way of strictly decreasing some follower's cost. This is w.l.o.g., as it is consistent with the assumption of optimism. Let us now prove that the sequence of followers' configurations satisfies the following:

$$\nu(t+1)_{i^*} \leq \nu(t)_{i^*} \quad \forall 0 \leq t < T. \quad (4)$$

By contradiction, assume there exists $0 \leq t < T$ such that $\nu(t+1)_{i^*} > \nu(t)_{i^*}$. Then, there is a follower who can strictly decrease her cost in $\nu(t)$ by choosing i^* instead of some $j \neq i^* \in R : \nu(t)_j > 0$, i.e., $c_{i^*,f}^{\hat{\sigma}_\ell}(\nu(t)_{i^*} + 1) < c_{j,f}^{\hat{\sigma}_\ell}(\nu(t)_j)$ holds. Thus, given that $c_{i^*,f}^{\hat{\sigma}_\ell}(\nu(t)_{i^*} + 1) = c_{i^*,f}(\nu(t)_{i^*} + 2)$ and $c_{j,f}^{\hat{\sigma}_\ell}(\nu(t)_j) = c_{j,f}(\nu(t)_j)$, we conclude that:

$$c_{i^*,f}^{\sigma_\ell}(\nu_{i^*} + 1) \leq c_{i^*,f}(\nu(t)_{i^*} + 2) < c_{j,f}(\nu(t)_j), \quad (5)$$

where the first inequality holds since $\nu(t)_{i^*} = \nu_{i^*}$ (as Equation (4) holds for the elements of the sequence preceding $\nu(t)$ and the number of followers selecting i^*

cannot decrease with respect to ν_{i^*}). Two cases are possible. In the first one, $\nu(t)_j \leq \nu_j$, implying, by monotonicity, $c_{j,f}(\nu(t)_j) \leq c_{j,f}(\nu_j) \leq c_{j,f}^{\sigma_\ell}(\nu_j)$, which, together with Equations (3) and (5), leads to a contradiction. In the second case, $\nu(t)_j > \nu_j$ implies that there exists $k \neq i^* \in R$ such that $\nu(t)_k < \nu_k$ (and $\nu_k > 0$), otherwise $\sum_{i \in R} \nu(t)_i > n - 1$. It follows that $c_{j,f}(\nu(t)_j) \leq c_{k,f}(\nu(t)_k + 1) \leq c_{k,f}^{\sigma_\ell}(\nu_k)$, where the first inequality holds since, due to our assumptions on the sequence, it cannot be $c_{j,f}(\nu(t)_j) > c_{k,f}(\nu(t)_k + 1)$ as $\nu(t+1)_{i^*} > \nu(t)_{i^*}$, and the second inequality follows from $\nu(t)_k < \nu_k$. Thus, Equations (3) and (5) give a contradiction. As a result, Equation (4) holds, and, thus, $\hat{\nu}_{i^*} \leq \nu_{i^*}$. Given the monotonicity of the costs, we conclude that $c_\ell^{\hat{\sigma}} \leq c_\ell^\sigma$. To prove the claim, it now suffices to take as strategy profile $\sigma = (\sigma_\ell, \nu)$ an OSE in which σ_ℓ is mixed—since the leader’s cost at σ is the smallest possible, her cost at $\hat{\sigma}$ will be identical to it. \square

We prove, now, that a similar result holds for the pessimistic case, i.e., for computing a PSE. The result is weaker though, as it requires the stronger assumption that the followers’ cost functions be strictly monotonic.

The idea of the proof is similar to the previous one. Given a PSE $\sigma = (\sigma_\ell, \nu)$ in which σ_ℓ is a mixed strategy, we show that there exists another PSE $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ where the leader’s commitment $\hat{\sigma}_\ell$ is a pure strategy which selects with probability one any of the resources for which the leader incurs the minimum cost when committing to σ_ℓ . In order to show this, we prove, by contradiction, that any NE $\hat{\nu}$ in the followers’ game resulting from $\hat{\sigma}_\ell$ provides the leader with a cost smaller than or equal to the one for σ .

Theorem 8. *Every SSSCG in which the leader’s and followers’ cost functions are, respectively, weakly and strictly monotonic, admits a PSE $\sigma = (\sigma_\ell, \nu)$ in which σ_ℓ is pure.*

Proof. Assume there exists a PSE $\sigma = (\sigma_\ell, \nu)$ in which σ_ℓ is mixed. We show that there must be another PSE $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ such that $\hat{\sigma}_\ell$ is pure. Let us define $i^* \in R$ and $\hat{\sigma}_\ell \in \Delta_\ell$ as in the proof of Theorem 7, so that $c_\ell^\sigma \geq c_{i^*,\ell}(\nu_{i^*} + 1)$ and Equation (3) holds. Given that the followers behave pessimistically, we need to show that, for every $\hat{\nu} \in E^{\hat{\sigma}_\ell}$, $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ satisfies $c_\ell^{\hat{\sigma}} \leq c_\ell^\sigma$. By contradiction, assume $c_\ell^{\hat{\sigma}} > c_\ell^\sigma$, which implies $c_{i^*,\ell}(\hat{\nu}_{i^*} + 1) > c_{i^*,\ell}(\nu_{i^*} + 1)$. It easily follows from the monotonicity of the costs that $\hat{\nu}_{i^*} > \nu_{i^*}$. Thus, there must be a resource $j \in R$ such that $\hat{\nu}_j < \nu_j$ as, otherwise, $\sum_{i \in R} \hat{\nu}_i > n - 1$. Let us also remark that $\nu_j > 0$. Thus:

$$c_{i^*,f}^{\sigma_\ell}(\nu_{i^*} + 1) \leq c_{i^*,f}(\hat{\nu}_{i^*} + 1) \leq c_{j,f}(\hat{\nu}_j + 1) \leq c_{j,f}^{\sigma_\ell}(\nu_j), \quad (6)$$

where the first inequality follows from $\nu_{i^*} < \hat{\nu}_{i^*}$, the second one from the fact that $\hat{\nu}$ is an NE for $\hat{\sigma}_\ell$, and the third one from $\hat{\nu}_j < \nu_j$. Equation (3) implies $c_{j,f}^{\sigma_\ell}(\nu_j) \leq c_{i^*,f}^{\sigma_\ell}(\nu_{i^*} + 1)$. If $c_{j,f}^{\sigma_\ell}(\nu_j) < c_{i^*,f}^{\sigma_\ell}(\nu_{i^*} + 1)$, then Equation (6) leads to a contradiction. Otherwise, if $c_{j,f}^{\sigma_\ell}(\nu_j) = c_{i^*,f}^{\sigma_\ell}(\nu_{i^*} + 1)$ all the inequalities in Equation (6) hold as equalities. This, however, implies $c_{i^*,f}^{\sigma_\ell}(\nu_{i^*} + 1) = c_{i^*,f}(\hat{\nu}_{i^*} + 1)$ and $c_{j,f}(\hat{\nu}_j + 1) = c_{j,f}^{\sigma_\ell}(\nu_j)$, which is a contradiction since σ_ℓ is mixed and the followers’ cost functions are strictly monotonic. \square

760 Theorem 8 fails to hold if the followers' cost functions are weakly, rather than strictly, monotonic, as the following result shows:

Proposition 3. *There are SSSCGs with weakly monotonic cost functions where any PSE prescribes the leader to play a mixed strategy.*

765 *Proof.* Consider the following instance of SSSCG with weakly monotonic cost functions, where $|F| = 1$ and $R = \{r_1, r_2\}$.

x	$c_{r_1,\ell}$	$c_{r_1,f}$	$c_{r_2,\ell}$	$c_{r_2,f}$
1	1	1	1	1
2	2	1	2	1

Clearly, any followers' configuration is an NE in this game, independently of the leader's commitment. Whenever the leader commits to a pure strategy, be it the selection of r_1 or r_2 , the follower, due to the pessimistic assumption, chooses 770 the same resource, so to have the leader incur a cost as large as possible (of 2). By uniformly randomizing between the two resources, though, the leader can reduce her cost to $2\frac{1}{2} + \frac{1}{2} = 1.5$. \square

Relying on Theorems 7 and 8, we can compute an OSE (respectively, PSE) by enumerating the leader's pure strategies and, for each of them, computing a 775 followers' NE which results in the smallest (respectively, largest) leader's cost. Such NE can be computed by applying a simple greedy procedure which progressively assigns followers to resources. At each step, a single follower is assigned to the resource which is cheapest for her, given how the previously considered followers have been distributed over the resources. Moreover, at a given step, 780 among all the resources minimizing followers' cost the procedure selects one minimizing (respectively, maximizing) the leader's cost. An O/PSE is then obtained by picking any leader's pure strategy for which the leader's cost is the smallest.

785 The detailed procedure is described in Algorithm 1 where, for some $S \subseteq R$ and $i \in S$, the function **O-Pick**(S, i) (respectively, **P-Pick**(S, i)) returns some resource $j^* \in S$, giving precedence to resources $j^* \neq i$ (respectively, $j^* = i$).

Let us remark that, in Algorithm 1, $\sigma_\ell[\cdot]$, $\nu[\cdot, \cdot]$, and $c_\ell[\cdot]$ are the algorithm's variables and, for every $i \in R$, $\nu[i, j]$ denotes the number of followers selecting resource $j \in R$ in the NE which is reached when the leader's strategy is $\sigma_\ell[i]$.

790 **Theorem 9.** *Algorithm 1 is correct and it runs in time $O(nr \log r)$.*

Proof. We rely on the pseudocode reported in Algorithm 1 to show its correctness. Thanks to Theorems 7 and 8, we only need to prove that, for every $i \in R$ and after the execution of the while loop, the followers configuration ν is such that, for all $j \in R$, $\nu_j = \nu[i, j]$ is an NE for $\sigma_\ell[i]$ minimizing (or maximizing) 795 the leader's cost. First, let us show that ν is an NE. Suppose, by contradiction, that it is not. Then, there exists $j \in R : \nu_j > 0$ and $k \in R$ such that $c_{j,f}^{\sigma_\ell[i]}(\nu_j) > c_{k,f}^{\sigma_\ell[i]}(\nu_k + 1)$. Let $\bar{\nu}_k$ be the value of $\nu[i, k]$ during the step in which $\nu[i, j]$ is set to its final value ν_j . Clearly, $c_{j,f}^{\sigma_\ell[i]}(\nu_j) > c_{k,f}^{\sigma_\ell[i]}(\nu_k + 1) \geq c_{k,f}^{\sigma_\ell[i]}(\bar{\nu}_k + 1)$,

Algorithm 1: Algorithm computing an O/PSE in SSSCG with monotonic costs.

input : An SSSCG $\Gamma = (N, R, c_\ell, c_f)$

output: σ that is an O/P-LFE of Γ

Function Compute-O/P-LFE(Γ)

```

for  $i \in R$  do
   $\sigma_\ell[i] \leftarrow \sigma_\ell \in \Delta_\ell : \sigma_\ell(i) = 1;$ 
   $\nu[i, j] \leftarrow 0 \ \forall i, j \in R;$ 
  while  $\sum_{j \in R} \nu[i, j] < n$  do
     $S \leftarrow \arg \min_{j \in R} c_{j,f}^{\sigma_\ell[i]}(\nu[i, j] + 1);$ 
     $j^* \leftarrow \text{O/P-Pick}(S, i);$ 
     $\nu[i, j^*] \leftarrow \nu[i, j^*] + 1;$ 
   $c_\ell[i] \leftarrow c_{i,\ell}(\nu[i, i] + 1);$ 
 $i^* \leftarrow \arg \min_{i \in R} c_\ell[i];$ 
return  $\sigma = (\sigma_\ell[i^*], \nu[i^*, \cdot]);$ 

```

and the algorithm would have not incremented $\nu[i, j]$ during that step, a contradiction. Let us show now that $(\sigma_\ell[i], \nu)$ is an O/PSE. In the remainder of the proof, we focus on the optimistic case (the pessimistic one can be treated analogously). Suppose, by contradiction, that ν is not an NE minimizing the leader's cost for $\sigma_\ell[i]$ (i.e., not an OSE). Then, there exists another NE $\hat{\nu}$ for $\sigma_\ell[i]$ such that $c_{i,\ell}(\hat{\nu}_i + 1) < c_{i,\ell}(\nu_i + 1)$. Given the monotonicity of the costs, $\hat{\nu}_i < \nu_i$ must hold. Therefore, there must exist some $j \neq i \in R$ such that $\hat{\nu}_j > \nu_j$. Let us consider the step in which $\nu[i, i]$ is set to ν_i , and let $\bar{\nu}_j$ be the value of $\nu[i, j]$ during that step. Note that $c_{i,f}^{\sigma_\ell[i]}(\nu_i) < c_{j,f}^{\sigma_\ell[i]}(\bar{\nu}_j + 1)$ must hold as, otherwise, the algorithm would have incremented $\nu[i, j]$ instead of $\nu[i, i]$. But, then, $c_{j,f}^{\sigma_\ell[i]}(\bar{\nu}_j + 1) \leq c_{j,f}^{\sigma_\ell[i]}(\nu_j + 1) \leq c_{j,f}^{\sigma_\ell[i]}(\hat{\nu}_j)$, which implies $c_{i,f}^{\sigma_\ell[i]}(\hat{\nu}_i + 1) \leq c_{i,f}^{\sigma_\ell[i]}(\nu_i) < c_{j,f}^{\sigma_\ell[i]}(\bar{\nu}_j + 1) \leq c_{j,f}^{\sigma_\ell[i]}(\hat{\nu}_j)$, contradicting the fact that $\hat{\nu}$ is an NE for the given $\sigma_\ell[i]$.

Since the while loop is executed exactly r times, each execution carries out n steps. Using efficient data structures, each step takes time $O(\log r)$. Thus, the overall running time is $O(nr \log r)$. \square

Next, we provide a characterization of O/PSEs in SSSCGs with monotonic costs under the additional assumption that leader's and followers' costs be equal, which may be of independent interest besides the computation of O/PSEs.

Theorem 10. *Given an SSSCG with monotonic costs and $c_\ell = c_f = \{c_i\}_{i \in R}$, any O/PSE $\sigma = (\sigma_\ell, a)$ with σ_ℓ pure is an NE.*

Proof. Let $\sigma = (\sigma_\ell, \nu)$ be an O/PSE with $\sigma_\ell(i^*) = 1$ for some $i^* \in R$. Clearly, given that $\nu \in E^{\sigma_\ell}$, $c_i^{\sigma_\ell}(\nu_i) \leq c_j^{\sigma_\ell}(\nu_j + 1)$ holds for every $i \in R : \nu_i > 0$ and for every $j \in R$. Therefore, no follower has an incentive to change resource. Thus, it is sufficient to prove that the leader has no incentive to deviate from resource i^* unilaterally, i.e., without assuming that the followers would react

825 to her deviation (which is the case in the Stackelberg setting). If $\nu_{i^*} > 0$, we have $c_{i^*}(\nu_{i^*} + 1) = c_{i^*}^{\sigma_\ell}(\nu_{i^*}) \leq c_j^{\sigma_\ell}(\nu_j + 1) = c_j(\nu_j + 1)$ for every $j \neq i^* \in R$, and it immediately follows that the leader does not deviate and σ is an NE. The case in which $\nu_{i^*} = 0$ is more involved. By contradiction, assume that σ is not an NE. As a consequence, the leader must have an incentive to deviate to some resource $j \neq i^* \in R$, i.e., $c_{i^*}(\nu_{i^*} + 1) = c_{i^*}(1) > c_j(\nu_j + 1)$. Let $\hat{\sigma}_\ell$ with $\hat{\sigma}_\ell(j) = 1$ be the strategy the leader commits to. We prove (by contradiction) that, for every $\hat{\nu} \in E^{\hat{\sigma}_\ell}$, $\hat{\sigma} = (\hat{\sigma}_\ell, \hat{\nu})$ provides the leader with a cost strictly smaller than $c_{i^*}(1)$. Assume $c_j(\hat{\nu}_j + 1) \geq c_{i^*}(1)$. Three cases are possible. In the first one, $\hat{\nu}_j < \nu_j$ and $c_{i^*}(1) > c_j(\nu_j + 1) \geq c_j(\hat{\nu}_j + 1) \geq c_{i^*}(1)$. In the second one, $\hat{\nu}_j = \nu_j$ and $c_j(\hat{\nu}_j + 1) \geq c_{i^*}(1) > c_j(\nu_j + 1)$. In the third case, $\hat{\nu}_j > \nu_j$, which implies that there must be a resource $k \neq i^* \in R$ such that $\hat{\nu}_k < \nu_k$, and $c_{i^*}(1) > c_j(\nu_j + 1) \geq c_k(\nu_k) \geq c_k(\hat{\nu}_k + 1) \geq c_j(\hat{\nu}_j + 1) \geq c_{i^*}(1)$. As all the cases lead to a contradiction, it must be $c_j(\hat{\nu}_j + 1) < c_{i^*}(1)$. The proof is complete as, in $\hat{\sigma}$, the leader's cost is $c_j(\hat{\nu}_j + 1) < c_{i^*}(1)$, contradicting the fact that σ is an O/PSE. \square

6.2. On the necessity of the assumptions we made

We provide some examples showing why Algorithm 1 cannot be easily extended to more general settings—the reason being that Theorems 7 and 8 do not hold if the assumption of monotonicity is dropped.

845 First, let us analyze the general case of SSSCGs in which the costs need not be monotonic functions of the resource congestion:

Proposition 4. *There are SSSCGs in which, even if the cost functions of one player only are nonmonotonic, be it the leader or one of the followers, any O/PSE prescribes the leader to play a mixed strategy.*

850 *Proof.* Consider the following instance of SSSCG with nonmonotonic followers' cost functions where $R = \{r_1, r_2\}$ and there is a single follower ($|F| = 1$).

x	$c_{r_1,\ell}$	$c_{r_1,f}$	$c_{r_2,\ell}$	$c_{r_2,f}$
1	1	2	1	2
2	2	1	2	1

The follower selects r_2 whenever $\sigma_\ell(r_1) \leq \frac{1}{2}$, while, if $\sigma_\ell(r_1) \geq \frac{1}{2}$, she chooses r_1 . The leader's cost is $2 - \sigma_\ell(r_1)$ if $\sigma_\ell(r_1) \leq \frac{1}{2}$, and $1 + \sigma_\ell(r_1)$ if $\sigma_\ell(r_1) \geq \frac{1}{2}$. 855 There is, thus, a unique O/PSE that prescribes the leader to commit to σ_ℓ with $\sigma_\ell(r_1) = \sigma_\ell(r_2) = \frac{1}{2}$.

Consider now the following instance of SSSCG with nonmonotonic leader's cost functions, where $R = \{r_1, r_2\}$ and there is single follower ($|F| = 1$).

x	$c_{r_1,\ell}$	$c_{r_1,f}$	$c_{r_2,\ell}$	$c_{r_2,f}$
1	2	1	2	1
2	0	2	0	2

860 The follower selects r_2 if $\sigma_\ell(r_1) \geq \frac{1}{2}$ and r_1 if $\sigma_\ell(r_1) \leq \frac{1}{2}$. The leader's cost is thus $2\sigma_\ell(r_1)$ if $\sigma_\ell(r_1) \geq \frac{1}{2}$ and $2 - 2\sigma_\ell(r_1)$ if $\sigma_\ell(r_1) \leq \frac{1}{2}$. There is, thus, a unique O/PSE which prescribes the leader to commit to σ_ℓ with $\sigma_\ell(r_1) = \sigma_\ell(r_2) = \frac{1}{2}$. \square

865 Finally, we show that Theorems 7 and 8 do not hold for general SSCGs with different action spaces, even if all the cost functions are monotonic:

Proposition 5. *There are SSCGs with different action spaces and monotonic cost functions where any O/PSE prescribes the leader to play a mixed strategy.*

Proof. Consider the following SSCG with $R = \{r_1, r_2, r_3\}$, two followers $F = \{p_1, p_2\}$, and $A_{p1} = \{r_1, r_2\}$, $A_{p2} = \{r_2, r_3\}$, $A_\ell = \{r_1, r_2\}$:

x	$c_{r_1,f}$	$c_{r_1,\ell}$	$c_{r_2,f}$	$c_{r_2,\ell}$	$c_{r_3,f}$
1	1	0	0	1	3
2	1	1	2	1	3
3	1	1	4	1	3

870 If the leader plays $\sigma_\ell(r_1) = 1$, there is a unique NE where follower p_1 plays r_1 and follower p_2 plays r_2 . Indeed, p_2 incurs a cost of 0 and, thus, has no incentive to deviate, while p_1 would incur a cost of $2 > 1$ by deviating to r_2 . Thus, the leader's cost is 1. The leader's cost is also 1 if she played $\sigma_\ell(r_2) = 1$, as p_2 would also choose r_2 , while p_1 would choose r_1 . Let us show that the leader can commit to a mixed strategy and incur a cost smaller than 1. Indeed, with $\sigma_\ell(r_1) = \sigma_\ell(r_2) = \frac{1}{2}$, there is a followers' NE where p_1 chooses r_2 and p_2 chooses r_3 : p_1 , incurring a cost of 1 (smaller or equal than any other cost), has no incentive to deviate, while p_2 , currently incurring a cost of 3, by switching to r_2 would incur the same (expected) cost of 3 (i.e., a cost of 2 with probability $\frac{1}{2}$ and one of 4 with probability $\frac{1}{2}$), thus having no incentive to deviate. At that NE, the leader's cost is $0 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2} = \frac{1}{2}$. \square

6.3. Pure-strategy commitment in SSSCGs with generic costs

885 We propose, here, a simple polynomial-time algorithm for computing an O/PSE in SSSCGs with generic costs where the leader is restricted to pure-strategy commitments. It is based on a dynamic programming algorithm proposed in [31] for the computation of an optimal NE in symmetric non-Stackelberg SCGs. The original algorithm runs in $O(n^6 r^5)$. One can compute an O/PSE in r iterations, fixing, at each iteration, the action the leader would choose and calling the previous algorithm to compute an NE which either minimizes or maximizes the leader's cost. This takes, overall, $O(n^6 r^6)$.

890 We show, in the following, how to improve the complexity of the original algorithm to $O(n^4 r^3)$, thanks to which we can compute an O/PSE for the restricted case in $O(n^4 r^4)$. The algorithm is based on the same recursive formula shown in [31], which we reintroduce, here, in a different and, possibly, clearer way.

Let $A(h, B, M, V)$ be the cost of an optimal NE for a symmetric SCG without leadership restricted to h resources $\{1, 2, \dots, h\} \subseteq R$ and B players, where M is the largest cost incurred by a player and V is the smallest cost a player would incur if she were to switch to another resource.

Proposition 6. $A(h, B, M, V)$ satisfies the following recursive equation:

$$A(h, B, M, V) = \min_{\substack{p \in \{0, \dots, B\} \\ m \in \mathbb{Z}^+, v \in \mathbb{Z}^+}} A(h-1, p, m, v) + (B-p) c_h(B-p) \quad (7)$$

$$\text{s.t. } m \leq M \quad (8)$$

$$v \geq V \quad (9)$$

$$c_h(B-p) \leq M \quad (10)$$

$$c_h(B-p+1) \geq V \quad (11)$$

$$c_h(B-p) \leq v \quad (12)$$

$$c_h(B-p+1) \geq m. \quad (13)$$

Proof. We show that all the constraints are necessary for the definition of $A(h, B, M, V)$ to be respected. If Constraint (8) were not satisfied, $m > M$ would imply that there is at least a resource among those in $\{1, \dots, h-1\}$ costing strictly more than M . If Constraint (9) were not satisfied, $v < V$ would imply that the cost to deviate to a resource among those in $\{1, \dots, h-1\}$ is strictly smaller than V . If Constraint (10) were not satisfied, $c_h(B-p) > M$ would imply that M is smaller than the cost of the most expensive chosen resource. If Constraint (11) were not satisfied, $c_h(B-p+1) < V$ would imply that V is larger than the cheapest cost a player would incur upon deviating to another resource. If Constraint (12) were not satisfied, $c_h(B-p) > v$ would imply that each of the $B-p$ players who chose resource h would have an incentive to deviate to any of the resources in $\{1, \dots, h-1\}$. If Constraint (13) were not satisfied, $c_h(B-p+1) < m$ would imply that at least one of the p players who selected a resource in $\{1, \dots, h-1\}$ (i.e., all those incurring a cost of m) would have an incentive to deviate to resource h . \square

We now show how to simplify the recursive formula for $A(h, B, M, V)$:

Theorem 11. $A(h, B, M, V)$ satisfies the following recursive equation:

$$A(h, B, M, V) = \min_{p \in \{0, \dots, B\}} A(h-1, p, m(p)^*, v(p)^*) + (B-p) c_h(B-p) \quad (14)$$

$$\text{s.t. } c_h(B-p) \leq M \quad (15)$$

$$c_h(B-p+1) \geq V, \quad (16)$$

where $m(p)^* = \min\{M, c_h(B-p+1)\}$ and $v(p)^* = \max\{V, c_h(B-p)\}$.

Proof. Constraints (8)–(13) and (9)–(12) imply, respectively, $m \leq \min\{M, c_h(B-p+1)\}$ and $v \geq \max\{V, c_h(B-p)\}$. Hence, $m(p)^*$ and $v(p)^*$ are feasible for Problem (7)–(13). Note that, if $m' > m$ and $v' < v$, the feasible region underlying $A(h, p, m', v')$ contains the one underlying $A(h, p, m, v)$, which implies

$A(h, p, m', v') \leq A(h, p, m, v)$. The claim follows since $m(p)^*$ and $v(p)^*$ are, respectively, the largest and smallest values m and v can take. \square

Corollary 7. *In symmetric non-Stackelberg SCGs, an optimal NE can be found in $O(n^4 r^3)$. In SSSCGs with the leader restricted to pure strategies, an O/PSE can be found in $O(n^4 r^4)$.*

Proof. Since there are at most nr different values of $c_j(x)$, for all $j \in R$ and $x \in N$, there are at most nr values of M and at most nr values of V . There are also exactly r values of h and exactly n of B . Hence, the dynamic programming table of $A(h, B, M, V)$ contains $O(n^3 r^3)$ entries. Due to Theorem (11), computing an entry of the table requires $O(n)$. Overall, an optimal NE is computed in $O(n^4 r^3)$. For the case with leadership restricted to pure strategies, it suffices to run the algorithm for each resource the leader may choose, i.e., $O(r)$ times, obtaining a complexity of $O(n^4 r^4)$. \square

7. Mixed-Integer Linear Programming Formulations for Computing OSEs in Intractable SSCGs and SSSCGs

In this section, we provide two MILP formulations for the problem of computing an OSE (in exponential time in the worst case) in SSCGs and SSSCGs for which the problem is intractable (see Sections 4 and 5). Our goal is to provide methods which work suitably well in practice, even though their worst-case running time is exponential.³

We start from SSSCGs, for which the MILP formulation is simpler, and then extend the result to the more general case of SSCGs.

7.1. Computing an OSE in SSSCGs (with generic costs)

For the ease of notation, let $V = \{1, \dots, n-1\}$ be the set of possible congestion levels induced by the followers on a resource. Let, for every resource $i \in R$ and value $v \in V$, the binary variable y_{iv} be equal to 1 if and only if $\nu_i = v$, i.e., if and only if v followers select resource $i \in R$. We use these variables to achieve a binarized representation of the followers' configuration $\nu \in \mathbb{N}^r$, namely, $\nu_i = \sum_{v \in V} v y_{iv}$ for all $i \in R$. Let, for each $i \in R$, $\alpha_i \in [0, 1]$ be equal to $\sigma_\ell(i)$. Let also, for each $i \in R$ and $v \in V$, the auxiliary variable z_{iv} be equal to the bilinear term $y_{iv}\alpha_i$.

The complete MILP formulation reads:

$$\min \sum_{i \in R} \sum_{v \in V} c_{i,\ell}(v+1) z_{iv} \quad (17a)$$

$$\text{s.t. } \sum_{v \in V} y_{iv} \leq 1 \quad \forall i \in R \quad (17b)$$

³We recall that, while we do not directly propose algorithms for the computation of PSEs for these intractable cases, their computation can be carried out with the general method proposed in [6, 7] for general SGs in normal form.

$$\sum_{i \in R} \sum_{v \in V} v y_{iv} = n - 1 \quad (17c)$$

$$\sum_{v \in V} \left(y_{jv} c_{j,f}(v+1) + z_{jv} (c_{j,f}(v+2) - c_{j,f}(v+1)) \right) \geq \sum_{v \in V} \left(y_{iv} c_{i,f}(v) + z_{iv} (c_{i,f}(v+1) - c_{i,f}(v)) \right) \quad \forall i \neq j \in R \quad (17d)$$

$$z_{iv} \leq \alpha_i \quad \forall i \in R, \forall v \in V \quad (17e)$$

$$z_{iv} \leq y_{iv} \quad \forall i \in R, \forall v \in V \quad (17f)$$

$$z_{iv} \geq \alpha_i + y_{iv} - 1 \quad \forall i \in R, \forall v \in V \quad (17g)$$

$$z_{iv} \geq 0 \quad \forall i \in R, \forall v \in V \quad (17h)$$

$$\sum_{i \in R} \alpha_i = 1 \quad (17i)$$

$$\alpha_i \geq 0 \quad i \in R \quad (17j)$$

$$y_{iv} \in \{0, 1\} \quad \forall i \in R, \forall v \in V. \quad (17k)$$

Function (17a) represents the leader's expected cost (to be minimized). Constraints (17b) ensure that at most one variable y_{iv} be equal to 1 for each resource $i \in R$, thus guaranteeing that the congestion level of each resource be uniquely determined (note that $\sum_{v \in V} y_{iv} = 0$ if no followers select resource $i \in R$). Constraints (17c) guarantee that the followers' configuration be well-defined, i.e., that $\sum_{i \in R} \nu_i$ be equal to $n - 1$ (the number of followers). Constraints (17d) force the followers' configuration defined by the y_{iv} variables to be an NE for the leader's strategy identified by the α_i variables. This follows from the fact that $\sum_{v \in V} \left(y_{iv} c_{i,f}(v) + z_{iv} (c_{i,f}(v+1) - c_{i,f}(v)) \right)$ (recall that $z_{iv} = y_{iv} \alpha_i$) is equal to the cost incurred by the followers who select resource $i \in R$, while $\sum_{v \in V} \left(y_{jv} c_{j,f}(v+1) + z_{jv} (c_{j,f}(v+2) - c_{j,f}(v+1)) \right)$ (recall that $z_{jv} = y_{jv} \alpha_j$) is equal to the cost they would incur after deviating to resource $j \in R$. Let us remark that Constraints (17d) are trivially satisfied if $y_{iv} = 0$ for all $v \in V$. This is correct as, if no followers choose resource $i \in R$, no equilibrium conditions need to be enforced. Constraints (17e)–(17h) are McCormick envelope constraints [65] which guarantee $z_{iv} = y_{iv} \alpha_i$ whenever $y_{iv} \in \{0, 1\}$.

We remark that Formulation (17) features $r(2n - 1)$ variables, nr of which binary, and $r(r - 1) + r(3n - 2) + 2$ constraints.

7.2. Computing an OSE in SSCGs

We now extend Formulation (17) to the case where the followers may have different action spaces, i.e., to general SSCGs.

For the ease of notation, let, for every $i \in R$, $\bar{v}_i = |\{p \in F \mid i \in A_p\}|$ be the maximum number of followers who can select resource i , and let $V(i) = \{1, \dots, \bar{v}_i\}$ be the set of possible congestion levels for resource i . For every follower $p \in F$ and resource $i \in A_p$, let the binary variable x_{pi} be equal to 1 if

and only if player p selects resource i , i.e., if and only if $a_p = i$. All the variables in Formulation (17) are used with the same meaning.

The complete MILP formulation reads:

$$\min \sum_{i \in R} \sum_{v \in V(i)} c_{i,\ell}(v+1) z_{iv} \quad (18a)$$

$$\text{s.t. } \sum_{i \in A_p} x_{pi} = 1 \quad \forall p \in F \quad (18b)$$

$$\sum_{v \in V(i)} y_{iv} \leq 1 \quad \forall i \in R \quad (18c)$$

$$\sum_{v \in V(i)} v y_{iv} = \sum_{p \in F: i \in A_p} x_{pi} \quad \forall i \in R \quad (18d)$$

$$\begin{aligned} & \sum_{v \in V(i)} \left(y_{jv} c_{j,f}(v+1) + z_{jv} (c_{j,f}(v+2) - c_{j,f}(v+1)) \right) \geq \\ & \geq \sum_{v \in V(i)} \left(y_{iv} c_{i,f}(v) + z_{iv} (c_{i,f}(v+1) - c_{i,f}(v)) \right) \quad \forall p \in F, i \neq j \in A_p \end{aligned} \quad (18e)$$

$$z_{iv} \leq \alpha_i \quad \forall i \in R, \forall v \in V(i) \quad (18f)$$

$$z_{iv} \leq y_{iv} \quad \forall i \in R, \forall v \in V(i) \quad (18g)$$

$$z_{iv} \geq \alpha_i + y_{iv} - 1 \quad \forall i \in R, \forall v \in V(i) \quad (18h)$$

$$z_{iv} \geq 0 \quad \forall i \in R, \forall v \in V(i) \quad (18i)$$

$$\sum_{i \in R} \alpha_i = 1 \quad (18j)$$

$$\alpha_i \geq 0 \quad \forall i \in R \quad (18k)$$

$$\alpha_i = 0 \quad \forall i \in R \setminus A_\ell \quad (18l)$$

$$x_{pi} \in \{0, 1\} \quad \forall p \in F, \forall i \in A_p \quad (18m)$$

$$y_{iv} \in \{0, 1\} \quad \forall i \in R, \forall v \in V(i). \quad (18n)$$

980 Objective Function (18a), Constraints (18c), and Constraints (18e)–(18k) have the same meaning as their counterparts in Formulation (17). Constraints (18b) ensure that each follower selects exactly one resource. Constraints (18d) guarantee that the followers' configuration be well-defined, i.e., that, for each $i \in R$, $\nu_i = \sum_{v \in V} v y_{iv}$ be equal to $\sum_{p \in F} x_{pi}$, which is the number of followers who
985 select resource i . Notice that, differently from the previous formulation, Constraints (18e) are enforced for each follower $p \in F$ here, and only for pairs of resources $i, j \in R$ follower p has access to. Note also that, via Constraints (18l), α_i is forced to be 0 for all the resources $i \in R$ the leader has no access to.

990 We observe that Formulation (18) features $\sum_{p \in F} |A_p| + 2 \sum_{i \in R} \bar{v}_i + r = O(r(3n+1))$ variables, $\sum_{p \in F} |A_p| + \sum_{i \in R} \bar{v}_i = O(2rn)$ of which binary, and $n+2r+3 \sum_{i \in R} \bar{v}_i + \sum_{p \in F} |A_p| (|A_p| - 1) = O(n+2r+3nr+nr(r-1))$ constraints.

7.3. Computing a pure-strategy OSE in SSSCGs and SSCGs with generic costs

Formulations (17) and (18) can be easily modified for the case in which the leader’s commitment is forced to be in pure strategies. This can be achieved by imposing the variables α_i to be binary. Notice that, when both α_i and y_{iv} are binary variables, z_{iv} becomes binary as well due to the McCormick constraints. The resulting formulations are ILPs.

One may wonder on the practical advantages of introducing an ILP formulation for computing a pure-strategy OSE in SSSCGs since, in Section 6.3, we have shown that this can be done in $O(n^4 r^4)$ by dynamic programming. As we will see comment on in Section 8, preliminary experiments show that, due to the high order of complexity of the dynamic programming algorithm, it is, in practice, more efficient to solve this formulation with state-of-the-art ILP algorithms than running the dynamic programming algorithm, even if solving Formulation (17) with branch-and-bound (and its variants, such as branch-and-cut) may take an exponential amount of computing time in the worst case.

8. Experimental evaluation

Since Algorithm 1 has a very low complexity— $O(nr \log r)$ —its efficiency is clear and it does not need to be established via computational experiments.

As to the dynamic programming algorithm proposed in Section 6.3 for SSSCGs with generic costs when the leader is restricted to pure-strategy commitments, preliminary tests have shown that this algorithm takes several hours to solve instances which are solved only in a matter of seconds with a state-of-the-art ILP algorithm applied to Formulation (17). This happens in, e.g., instances with 10 resources and 25 followers, on which the dynamic programming algorithm takes more than 11,000 seconds while with the ILP formulation we can solve them in less than a second. For this reason, in the remainder of the section we solely focus on the mathematical programming formulations.

Notice that games with monotonic cost functions and identical action spaces can be solved efficiently using Algorithm 1 (which we proposed in Section 6). For this reason, we focus on games with generic cost functions and/or different action spaces, assessing how state-of-the-art branch-and-bound methods behave when solving our formulations on instances of increasing size.

For the purpose, we experiment with Formulations (17) and (18) on a testbed of randomly generated game instances of two classes:

- *SSSCG instances*: we assume a number of followers in $\{20, 40, 60, 80, 100\}$, with r resources in the range $\{10, 20, 30, 40, 50\}$ and players’ costs randomly generated by sampling from $\{1, \dots, (n-1)r\}$ with a uniform probability.⁴

⁴The value $(n-1)r$ is chosen as, when looking for pure-strategy NEs, cost functions taking $(n-1)r$ different values are sufficient to represent every possible SCG.

1030 • *SSCG instances*: we assume a number of followers in $\{20, 40, 60, 80, 100\}$, with $r = 30$ resources and a number of actions $|A_p|$ per player in the range $\{7, 15, 22\}$, generated by sampling uniformly at random without replacement; the players' costs are sampled from $\{1, \dots, (n-1)r\}$ with uniform probability.

1035 We also test our MILP formulations on the worst-case game instances generated by following the reductions of Theorems 1 and 3:

- *SSSCG instances*: instances built following the reduction of Theorem 3 starting from K -PARTITION instances with $|S| \in \{50, 100, 150, 250, 300\}$ integer numbers with values sampled from $\{1, \dots, 100\}$.
- 1040 • *SSCG instances*: instances built following the reduction of Theorem 1 using random 3SAT instances with $|V| \in \{3, 5, 7, 9, 11, 13\}$ variables and $|C| = k|V|$ clauses, where $k \approx 4.26$ is the phase-transition parameter which typically characterizes hard-to-solve 3SAT instances [66].

1045 We generate 15 instances per combination of the parameters. All the experiments are run on a UNIX machine with a total of 32 cores working at 2.3 GHz, equipped with 128 GB of RAM. Each game instance is solved on a single core within a time limit of 7200 seconds. We use Python 2.7, solving the MILP formulations with GUROBI 7.0.

1050 We use, as baseline for the comparisons, a simple algorithm which, starting from a randomly generated assignment of the players to the resources, simulates best-response dynamics halting after a time limit of 10 minutes. When ties arise, i.e., whenever there are two or more players who are not playing their best response, we select a player lexicographically and make her switch to playing her (currently) best response. We refer to this algorithm as a *best response dynamics heuristic* since it is not exact when applied to the intractable cases of SSCGs and SSSCGs.⁵ On average, within the time limit of 10 minutes we observe a number of deviations to a best response of the order of 10^5 . Let us recall that the method always produces, by design, pure-strategy NEs.⁶

1060 Figure 3 and Figure 5 (a) report the results for SSSCGs with generic costs. Figure 3 (a) displays the average computing time required by Formulation (17), as a function of the number of followers and for different numbers of resources. One can see that, with Formulation (17), an optimal solution is always found

⁵We also evaluated different heuristic algorithms combining Algorithm 1 together with best response dynamics, e.g., using the solution returned by Algorithm 1 as a starting point for best-response dynamics instead of using randomly generated starting points. However, these approaches exhibited worse empirical performances than the best-response dynamics heuristic for all the settings which we have considered, including the more symmetric ones, showing that the algorithm does not benefit from the degree of symmetry of the instance. For this reason, we defer their analysis to Appendix B.

⁶Notice that, for games with identical action spaces, one could think of using the dynamic programming algorithm presented in Section 6.3 to find a pure-strategy OSE as heuristic approximation of a mixed-strategy OSE. However, as we mentioned above, the dynamic programming algorithm does not scale well enough in practice.

within the time limit of 7200 seconds in all the instances. This suggests that, even if the problem is hard in the worst case, an optimal solution can be found in a reasonable amount of time on randomly generated instances. Figure 5 (a) displays the results for worst-case instances. Surprisingly, within the time limit of 7200 seconds we are able to solve games with up to 302 resources and 1202 followers. Thus, while the instances generated by our reduction from K -PARTITION are the hardest ones asymptotically, they are solved more easily than randomly generated instances of the same size. Figure 3 (b) reports, as a function of the number of followers, the average leaders' cost of the solutions obtained with Formulation (17), compared to the average cost obtained with the best-response dynamics heuristic. As the figure shows, the difference in leader's utility between solutions found with the two methods can be quite large as the number of followers increases, up to a factor of 6 with $n = 100$, showing a clearly growing trend.

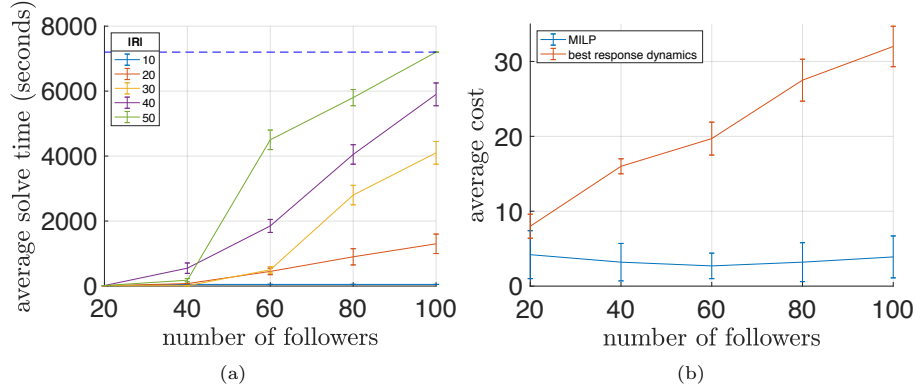


Figure 3: Results for the computation of an OSE in SSSCGs with generic costs. (a) Average computing time required by Formulation (17), as a function of the number of followers and for different numbers of actions available to each player. (b) Average leader's cost of the solutions obtained with Formulation (17) and with the best-response dynamics heuristic as a function of the number of followers, with 30 resources.

Figure 4 and Figure 5 (b) report the results for SSCGs with generic costs and 30 resources. Figure 4 (a) reports the average computing time required by Formulation (18) to find an OSE, as a function of the number of followers and for a different number of actions available to each player. Similarly to the case of SSSCGs, the chart shows that with Formulation (18) we can find an optimal solution within the time limit of 7200 seconds in all the instances. This suggests that, even if the problem is hard in the worst case, also for SSCGs one can find an optimal solution in a reasonable amount of computing time on randomly generated instances. The chart also shows, though, that the time required to solve this class of problems is much larger than the time required to solve their SSSCGs counterparts. Figure 5 (b) displays results for worst-case instances generated using our reduction from 3SAT. As for SSSCGs, these instances are not harder than random ones for the instance sizes used in our experimental

1090 setting. In particular, within the time limit of 7200 seconds, we are able to
 solve instances with up to 1538 resources and 2983 followers. Figure 4 (b)
 reports, for games with 15 actions per player, the average leader’s cost of the
 solutions obtained with the MILP Formulation (18) and with the best-response
 dynamics heuristic, as a function of the number of followers. Differently from
 1095 the case of SSSCGs, we observe that for SSCGs the heuristic returns solutions
 which, empirically, appear to be within a constant approximation factor of the
 optimal ones which is never larger than 5.

Overall, the results suggest the practical viability of our MILP formula-
 tions for finding provably optimal solutions also for games where a simple best-
 1100 response heuristic provides poor-quality solutions.

Surprisingly, the results that we have obtained for the worst-case instances
 are comparable to those for random games, empirically showing that, for the
 games that we study, random instances are not easier to solve than structured
 ones, differently from what is often observed in other cases (see, for instance,
 1105 [67] for the case of normal-form games).

9. Conclusions and future works

We have analyzed Stackelberg games where the underlying structure is a
 congestion game, focusing on the case in which the players’ actions are single-
 tons. We have shown that the problem of computing a *Stackelberg Equilibrium*
 1110 (SE) in such games is hard, except for the case in which all the players share
 the same resources and the cost functions are monotonically increasing in the
 congestion level. More precisely, we have shown that, for games where either the
 players have different action spaces and their cost functions are monotonic or
 their action spaces are the same but their cost functions are generic (monotonic

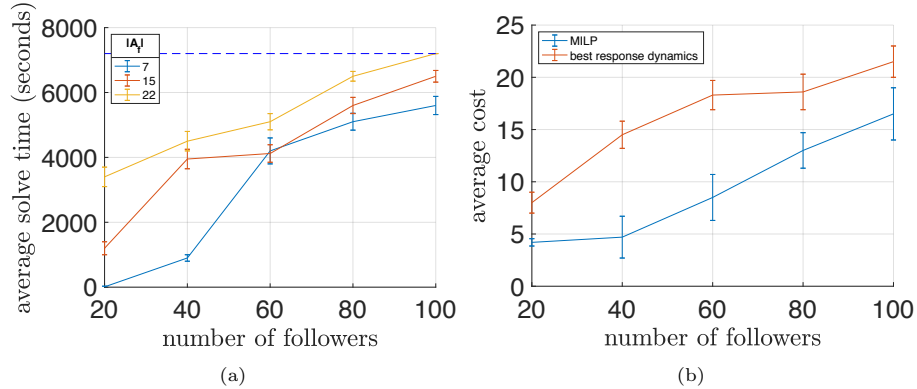


Figure 4: Results for the computation of an OSE in SSCGs with generic costs and 30 resources. (a) Average computing time required by Formulation (18), as a function of the number of followers and for different numbers of actions available to each player. (b) Average leader’s cost of the solutions obtained with the Formulation (18) and with the best-response dynamics heuristics, as a function of the number of followers and with 15 actions per player.

1115 or not), it is not possible to approximate in polynomial time the leader’s cost
 at an either *optimistic* SE (OSE) or *pessimistic* SE (PSE) to within any factor
 polynomial in the size of the game unless $P = NP$.

We have proposed a polynomial-time algorithm for finding an O/PSE for
 the case where the players have identical action spaces and their cost functions
 1120 are monotonic, and we have shown that games in this class always admit a
 pure-strategy SE. We have also shown how to improve the complexity of the
 state-of-the-art algorithm for the computation of an optimal NE in singleton
 congestion games, which has allowed us to compute an O/PSE in polynomial
 time for the case where the leader is restricted to pure strategies. For the
 1125 intractable cases with different action spaces and generic cost functions, we have
 proposed a mixed-integer linear programming formulation for finding an OSE,
 and a more compact one for the case in which the action spaces are identical. We
 have shown that state-of-the-art integer linear programming solvers scale well
 in practice when solving our formulations on random game instances, allowing
 1130 for tackling games with up to 40 resources and 100 followers.

In the future, we will investigate whether congestion games with a special
 structure allow for efficient solution algorithms, *e.g.*, when users can be of a fixed
 number of types [68]. We will also investigate more dominance relations for our
 dynamic programming algorithms in order to reduce the number of states they
 1135 explore and, ultimately, improving their efficiency. Another natural research
 direction is addressing the multi-leader case and, in particular, investigating
 whether, for the cases in which the problem we studied admits a polynomial-
 time algorithm, the introduction of more than a single leader makes the problem
 harder.




Figure 5: Results for the computation of an OSE in worst-case instances built on the base of our inapproximability reductions. (a) Average computing time required by Formulation (17) for instances based on K -PARTITION (Theorem 3), as a function of the number of resources. (b) Average computing time required by Formulation (18) for instances based on 3SAT (Theorem 1), as a function of the number of resources.

1140 References

References

- [1] A. Marchesi, S. Coniglio, N. Gatti, Leadership in singleton congestion games, in: IJCAI, 2018.
- [2] H. von Stackelberg, Marktform und Gleichgewicht, 1934.
- 1145 [3] V. Conitzer, T. Sandholm, Computing the optimal strategy to commit to, in: Proceedings of the 7th ACM conference on Electronic commerce, 2006, pp. 82–90.
- [4] B. von Stengel, S. Zamir, Leadership games with convex strategy sets, Games and Economic Behavior 69 (2) (2010) 446–457.
- 1150 [5] V. Conitzer, D. Korzhyk, Commitment to correlated strategies, in: AAAI, 2011, pp. 632–637.
- [6] S. Coniglio, N. Gatti, A. Marchesi, Pessimistic leader-follower equilibria with multiple followers, in: IJCAI, 2017, pp. 171–177.
- 1155 [7] S. Coniglio, N. Gatti, A. Marchesi, Computing a pessimistic leader-follower equilibrium with multiple followers: the mixed-pure case, CoRR abs/1808.01438. [arXiv:1808.01438](#).
- [8] A. Smith, Y. Vorobeychik, J. Letchford, Multidefender security games on networks, PERF E R 41 (4) (2014) 4–7.
- [9] J. Lou, Y. Vorobeychik, Equilibrium analysis of multi-defender security games, in: IJCAI, 2015.
- 1160 [10] A. Laszka, J. Lou, Y. Vorobeychik, Multi-defender strategic filtering against spear-phishing attacks, in: AAAI, 2016.
- [11] J. Lou, A. Smith, Y. Vorobeychik, Multidefender security games, IEEE INTELL SYST 32 (1) (2017) 50–60.
- 1165 [12] J. Gan, E. Elkind, M. Wooldridge, Stackelberg security games with multiple uncoordinated defenders, in: AAMAS, 2018.
- [13] M. Castiglioni, A. Marchesi, N. Gatti, 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, 2019, pp. 123–129.
- 1170 [14] J.-S. Pang, M. Fukushima, Quasi-variational inequalities, generalized nash equilibria, and multi-leader-follower games, COMP MAN SC 2 (1) (2005) 21–56.
- [15] S. Leyffer, T. Munson, Solving multi-leader-common-follower games, OPT MET SO 25 (4) (2010) 601–623.
- 1175

- [16] A. Kulkarni, U. Shanbhag, A shared-constraint approach to multi-leader multi-follower games, *SET-VALUED VAR ANAL* 22 (4) (2014) 691–720.
- [17] M. Breton, A. Alj, A. Haurie, Sequential stackelberg equilibria in two-person games, *J OPTIMIZ THEORY APP* 59 (1) (1988) 71–97.
- 1180 [18] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, S. Kraus, Playing games for security: an efficient exact algorithm for solving bayesian stackelberg games, in: *AAMAS*, 2008, pp. 895–902.
- 1185 [19] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordóñez, M. Tambe, Computing optimal randomized resource allocations for massive security games, in: *AAMAS*, 2009, pp. 689–696.
- [20] B. An, J. Pita, E. Shieh, M. Tambe, C. Kiekintveld, J. Marecki, Guards and Protect: Next generation applications of security games, *ACM SIGecom Exchanges* 10 (1) (2011) 31–34.
- 1190 [21] M. Tambe, *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*, Cambridge University Press, 2011.
- [22] M. Labbé, P. Marcotte, G. Savard, A bilevel model of taxation and its application to optimal highway pricing, *Management science* 44 (12-part-1) (1998) 1608–1622.
- 1195 [23] M. Labbé, A. Violin, Bilevel programming and price setting problems, *ANN OPER RES* 240 (1) (2016) 141–169.
- [24] A. Caprara, M. Carvalho, A. Lodi, G. J. Woeginger, Bilevel knapsack with interdiction constraints, *INFORMS J COMPUT* 28 (2) (2016) 319–333.
- [25] J. Matuschke, S. T. McCormick, G. Oriolo, B. Peis, M. Skutella, Protection of flows under targeted attacks, *OPER RES LETT* 45 (1) (2017) 53–59.
- 1200 [26] E. Amaldi, A. Capone, S. Coniglio, L. G. Gianoli, Network optimization problems subject to max-min fair flow allocation, *IEEE COMMUN LETT* 17 (7) (2013) 1463–1466.
- 1205 [27] R. Avenhaus, A. Okada, S. Zamir, Inspector leadership with incomplete information, in: *Game equilibrium models IV*, Springer, 1991, pp. 319–361.
- [28] W. H. Sandholm, Evolutionary implementation and congestion pricing, *The Review of Economic Studies* 69 (3) (2002) 667–689.
- [29] R. W. Rosenthal, A class of games possessing pure-strategy nash equilibria, *International Journal of Game Theory* 2 (1) (1973) 65–67.
- 1210 [30] H. Ackermann, H. Röglin, B. Vöcking, On the impact of combinatorial structure on congestion games, *Journal of the ACM (JACM)* 55 (6) (2008) 25.

- [31] S. Jeong, R. McGrew, E. Nudelman, Y. Shoham, Q. Sun, Fast and compact: A simple class of congestion games, in: AAAI, 2005, pp. 489–494.
- 1215 [32] P.-A. Chen, C.-J. Lu, Generalized mirror descents in congestion games, Artificial Intelligence 241 (2016) 217–243.
- [33] D. Fotakis, A selective tour through congestion games, in: Algorithms, Probability, Networks, and Games, Springer, 2015, pp. 223–241.
- [34] S. Suri, C. D. Tóth, Y. Zhou, Selfish load balancing and atomic congestion
1220 games, Algorithmica 47 (1) (2007) 79–96.
- [35] D. Konur, J. Geunes, Competitive multi-facility location games with non-identical firms and convex traffic congestion costs, Transportation Research Part E: Logistics and Transportation Review 48 (1) (2012) 373–385.
- [36] T. Roughgarden, Selfish routing and the price of anarchy, Vol. 174, MIT
1225 press Cambridge, 2005.
- [37] D. Monderer, L. S. Shapley, Potential games, Games and economic behavior 14 (1) (1996) 124–143.
- [38] A. Fabrikant, C. Papadimitriou, K. Talwar, The complexity of pure nash
1230 equilibria, in: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, ACM, 2004, pp. 604–612.
- [39] R. Werneck, J. Setubal, A. da Conceicao, Finding minimum congestion spanning trees, Journal of Experimental Algorithmics (JEA) 5 (2000) 11.
- [40] T. Roughgarden, Stackelberg scheduling strategies, SIAM Journal on Computing 33 (2) (2004) 332–350.
- 1235 [41] H. Küçükaydın, N. Aras, İ. K. Altınel, A leader–follower game in competitive facility location, Computers & Operations Research 39 (2) (2012) 437–448.
- [42] N. Basilico, S. Coniglio, N. Gatti, A. Marchesi, Bilevel programming methods for computing single-leader-multi-follower equilibria in normal-form and polymatrix games, EURO Journal on Computational Optimization
1240 [doi:10.1007/s13675-019-00114-8](https://doi.org/10.1007/s13675-019-00114-8).
- [43] N. Basilico, S. Coniglio, N. Gatti, Methods for Finding Leader-Follower Equilibria with Multiple Followers: (Extended Abstract), in: Proc. of AAMAS, International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 1363–1364.
1245
- [44] N. Basilico, S. Coniglio, N. Gatti, A. Marchesi, Bilevel programming approaches to the computation of optimistic and pessimistic single-leader-multi-follower equilibria, LEIBNIZ INTERNATIONAL PROCEEDINGS IN INFORMATICS 75 (2017) 1–14.

- 1250 [45] J. Letchford, V. Conitzer, K. Munagala, Learning and approximating the optimal strategy to commit to, in: *International Symposium on Algorithmic Game Theory*, Springer, 2009, pp. 250–262.
- [46] G. De Nittis, A. Marchesi, N. Gatti, Computing the optimal strategy to commit to in polymatrix games, in: *AAAI*, 2018, pp. 82–90.
- 1255 [47] J. Letchford, V. Conitzer, Computing optimal strategies to commit to in extensive-form games, in: *EC*, 2010.
- [48] G. Farina, A. Marchesi, C. Kroer, N. Gatti, T. Sandholm, Trembling-hand perfection in extensive-form games with commitment, in: *IJCAI*, 2018.
- [49] B. Bošanský, J. Cermak, Sequence-form algorithm for computing stackelberg equilibria in extensive-form games, in: *AAAI*, 2015.
- 1260 [50] J. Cermak, B. Bošanský, K. Durkota, V. Lisy, C. Kiekintveld, Using correlated strategies for computing Stackelberg equilibria in extensive-form games, in: *AAAI*, 2016.
- [51] C. Kroer, G. Farina, T. Sandholm, Robust Stackelberg equilibria in extensive-form games and extension to limited lookahead, in: *AAAI*, 2018.
- 1265 [52] A. Marchesi, G. Farina, C. Kroer, N. Gatti, T. Sandholm, Quasi-perfect stackelberg equilibrium, in: *AAAI*, 2019.
- [53] G. De Nittis, A. Marchesi, N. Gatti, Computing the strategy to commit to in polymatrix games (Extended Version), *CoRR* abs/1807.11914. **arXiv: 1807.11914**.
- 1270 [54] J. Letchford, L. MacDermed, V. Conitzer, R. Parr, C. L. Isbell, Computing optimal strategies to commit to in stochastic games., in: *AAAI*, 2012.
- [55] Y. Vorobeychik, S. P. Singh, Computing stackelberg equilibria in discounted stochastic games., in: *AAAI*, 2012.
- 1275 [56] H. Xu, R. Freeman, V. Conitzer, S. Dughmi, M. Tambe, Signaling in bayesian stackelberg games, in: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 150–158.
- [57] D. Fotakis, Stackelberg strategies for atomic congestion games, *Theory of Computing Systems* 47 (1) (2010) 218–249.
- 1280 [58] V. Bonifaci, T. Harks, G. Schäfer, Stackelberg routing in arbitrary networks, *Mathematics of Operations Research* 35 (2) (2010) 330–346.
- [59] V. Bilò, C. Vinci, On stackelberg strategies in affine congestion games, in: *International Conference on Web and Internet Economics*, Springer, 2015, pp. 132–145.
- 1285

- [60] Z.-Q. Luo, J.-S. Pang, D. Ralph, Mathematical programs with equilibrium constraints, 1996.
- [61] H. Serali, A multiple leader stackelberg model and analysis, OPER RES 32 (2) (1984) 390–404.
- 1290 [62] V. DeMiguel, H. Xu, A stochastic multiple-leader stackelberg model: analysis, computation, and application, OPER RES 57 (5) (2009) 1220–1235.
- [63] Y. Shoham, K. Leyton-Brown, Multiagent systems: Algorithmic, game-theoretic, and logical foundations, Cambridge University Press, 2008.
- 1295 [64] M. R. Garey, D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-completeness, WH Freeman and Company, 1979.
- [65] G. McCormick, Computability of global solutions to factorable nonconvex programs: Part I – Convex underestimating problems, Math. Program. 10 (1) (1976) 147–175.
- 1300 [66] P. Cheeseman, B. Kanefsky, W. Taylor, Where the really hard problems are, in: IJCAI, 1991, pp. 331–337.
- [67] T. Sandholm, A. Gilpin, V. Conitzer, Mixed-integer programming methods for finding nash equilibria, in: AAAI, 2005, pp. 495–501.
- 1305 [68] A. Marchesi, M. Castiglioni, N. Gatti, 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, 2019, pp. 485–491.

Appendix A. Omitted proofs

Theorem 1. *Computing an OSE in SSCGs with different action spaces is NP-hard, even if the leader has only one action (i.e., she can only select a single resource) and the cost functions are monotonic.*

Proof. We provide a reduction from 3SAT showing that the existence of a polynomial-time algorithm for computing an OSE in SSCGs would allow us to solve any 3SAT instance in polynomial time. Specifically, given a 3SAT instance (C, V) and a real number $0 < \epsilon < 4$, we build an instance $\Gamma_\epsilon(C, V)$ of an SSCG admitting an OSE in which the leader's cost is ϵ if and only if (C, V) is satisfiable; if not, the leader's cost is 4 in any OSE.

Mapping. $\Gamma_\epsilon(C, V)$ is defined as follows:

- $N = F \cup \{\ell\}$, with $F = \{p_\phi, p_{\phi,t} \mid \phi \in C\} \cup \{p_v \mid v \in V\} \cup \{p_{v,k}, p_{\bar{v},k} \mid v \in V, k \in \{1, \dots, m\}\} \cup \{p_{\phi,v}, p_{\phi,\bar{v}} \mid \phi \in C, v \in V\}$;
- $R = \{r_t\} \cup \{r_\phi \mid \phi \in C\} \cup \{r_v, r_{v,t}, r_{\bar{v},t}, r_{\bar{v},t} \mid v \in V\} \cup \{r_{\phi,v}, r_{\phi,\bar{v}} \mid \phi \in C, v \in V\}$;
- $A_{p_\phi} = \{r_\phi\} \cup \{r_{\phi,l} \mid l \in \phi\}$, $A_{p_{\phi,t}} = \{r_\phi, r_t\} \quad \forall \phi \in C$;
- $A_{p_{v,k}} = \{r_{v,t}, r_v\}$, $A_{p_{\bar{v},k}} = \{r_{\bar{v},t}, r_{\bar{v}}\} \quad \forall v \in V, k \in \{1, \dots, m\}$;
- $A_{p_v} = \{r_t, r_{v,t}, r_{\bar{v},t}\} \quad \forall v \in V$;
- $A_{p_{\phi,v}} = \{r_v, r_{\phi,v}\}$, $A_{p_{\phi,\bar{v}}} = \{r_{\bar{v}}, r_{\phi,\bar{v}}\} \quad \forall \phi \in C, v \in V$;
- $A_\ell = \{r_t\}$.

The cost functions take values according to the following table, and satisfy $c_{r_{\bar{v}},f} = c_{r_v,f}$, $c_{r_{\phi,\bar{v}},f} = c_{r_{\phi,v},f}$, $c_{r_{\bar{v},t},f} = c_{r_{v,t},f}$, and $c_{r_t,f} = c_{r_t,\ell}$ (let us remark that, given $\epsilon < 4$, they are all monotonic functions of the resource congestion):

x	$c_{r_{\phi},f}$	$c_{r_v,f}$	$c_{r_{v,t},f}$	$c_{r_{\phi,v},f}$	$c_{r_t,f}$
1	2	0	0	1	ϵ
$[2, m]$	5	0	6	6	4
$[m+1, \infty]$	5	7	6	6	4

Figure 1 shows an example of the game instance $\Gamma_\epsilon(C, V)$.

Given a 3SAT instance (C, V) , $\Gamma_\epsilon(C, V)$ can be constructed in polynomial time, as it features $n = 2m + s + 4ms + 1$ players and $r = m + 4s + 2ms + 1$ resources. Since, in $\Gamma_\epsilon(C, V)$, the leader can only select a single resource, r_t , the only leader's commitment is $\sigma_\ell \in \Delta_\ell : \sigma_\ell(r_t) = 1$. As a result, the leader's cost is ϵ if and only if no follower selects resource r_t ; otherwise, it is 4.

If. Assume that (C, V) is satisfiable, and let $\tau : V \rightarrow \{T, F\}$ be a truth assignment satisfying all the clauses in C . Using τ , we show how to recover

a followers' action profile $a = (a_p)_{p \in F} \in \times_{p \in F} A_p$ such that $a \in E^{\sigma_\ell}$, with
1340 $\sigma = (\sigma_\ell, a)$ providing $c_\ell^\sigma = \epsilon$. Note that, since ϵ is the minimum cost the leader
can achieve and the followers behave optimistically, σ is an OSE. In particular,
let $a_{p_{\phi,t}} = r_\phi$, for all $\phi \in C$. Moreover, if $\tau(v) = T$, let $a_{p_v} = r_{\bar{v},t}$ and $a_{p_{\phi,v}} = r_v$,
 $a_{p_{\phi,\bar{v}}} = r_{\phi,\bar{v}}$ for all $\phi \in C$, while, for all $k \in \{1, \dots, m\}$, let $a_{p_{v,k}} = r_{v,t}$ and
 $a_{p_{\bar{v},k}} = r_{\bar{v}}$. Instead, if $\tau(v) = F$, let $a_{p_{\bar{v}}} = r_{v,t}$ and $a_{p_{\phi,\bar{v}}} = r_{\bar{v}}$, $a_{p_{\phi,v}} = r_{\phi,v}$ for
1345 all $\phi \in C$, while, for all $k \in \{1, \dots, m\}$, let $a_{p_{\bar{v},k}} = r_{\bar{v},t}$ and $a_{p_{v,k}} = r_v$. Notice
that, since either $\tau(v) = T$ or $\tau(v) = F$, two cases are possible. If $\tau(v) = T$,
we have $\nu_{r_v}^a = m$ (followers $p_{\phi,v}$), $\nu_{r_{\bar{v}}}^a = m$ (followers $p_{\bar{v},k}$), $\nu_{r_{v,t}}^a = m$ (followers
 $p_{v,k}$), and $\nu_{r_{\bar{v},t}}^a = 1$ (follower $p_{\bar{v}}$). If $\tau(v) = F$, we have $\nu_{r_{\bar{v}}}^a = m$ (followers $p_{\phi,\bar{v}}$),
 $\nu_{r_v}^a = m$ (followers $p_{v,k}$), $\nu_{r_{\bar{v},t}}^a = m$ (followers $p_{\bar{v},k}$), and $\nu_{r_{v,t}}^a = 1$ (follower p_v).
1350 Assume, w.l.o.g., $\tau(v) = T$, as the other case is analogous. First, no follower
 $p_{\phi,v}$ would deviate from r_v to $r_{\phi,v}$, as, otherwise, she would incur a cost of at
least 1, rather than 0. The same holds for followers $p_{\phi,\bar{v}}$, as their cost is at most
6 while, if any of them switched to $r_{\bar{v}}$, she would incur a cost of 7. Similarly,
followers $p_{v,k}$ would not deviate from $r_{v,t}$ (as $6 < 7$) and followers $p_{\bar{v},k}$ would
1355 not deviate from $r_{\bar{v}}$ (as $0 < 6$). Since $\nu_{r_{\bar{v},t}}^a = 1$, follower $p_{\bar{v}}$ would not deviate
from $r_{\bar{v},t}$ (as $0 < 6$ and $0 < 4$). Furthermore, since τ is a truth assignment
satisfying (C, V) , at least one literal $l \in \phi$ evaluates to true under τ for every
 $\phi \in C$. Let $a_{p_\phi} = r_{\phi,l}$ for every $\phi \in C$. Since l evaluates to true, it must be
 $a_{p_{\phi,l}} = r_l$, thus p_ϕ is the only follower who selects $r_{\phi,l}$. As a result, p_ϕ incurs
1360 a cost of 1, and she has no incentive to deviate. Finally, $p_{\phi,t}$ does not deviate
from r_ϕ to r_t as $2 < 4$. Thus, we can conclude that a is an NE and that, since
no follower chose r_t , the leader's cost is ϵ .

Only if. Suppose there exists an OSE $\sigma = (\sigma_\ell, a)$ in which the leader's cost
is ϵ . We show that, in polynomial time, one can recover a truth assignment
1365 τ that satisfies all the clauses in C from $a = (a_p)_{p \in F}$. First, let us note that
no follower selects r_t in a as, otherwise, the leader's cost would be $4 > \epsilon$. As
a consequence, all followers $p_{\phi,t}$ and p_v must select one of the other resources
available to them, i.e., $a_{p_{\phi,t}} = r_\phi$ and $a_{p_v} \in \{r_{v,t}, r_{\bar{v},t}\}$. Moreover, there cannot
be two followers using resource r_ϕ for every $\phi \in C$ as, otherwise, $p_{\phi,t}$ would
1370 have an incentive to deviate from r_ϕ to r_t (as $5 > 4$). Thus, $a_{p_\phi} \neq r_\phi$, and,
for all $\phi \in C$, there must be a literal $l \in \phi$ such that $a_{p_\phi} = r_{\phi,l}$. In addition,
there cannot be two followers selecting $r_{\phi,l}$ as, otherwise, p_ϕ would have an
incentive to deviate to r_ϕ (as $5 < 6$). Thus, it must be the case that $a_{p_{\phi,l}} = r_l$.
This implies that $\nu_{r_l}^a \leq m$ as, otherwise, the cost of $p_{\phi,l}$ would be $7 > 6$, and
1375 that follower would change resource, switching to $r_{\phi,l}$. Thus, at least one of the
followers $p_{l,k}$ must select $r_{l,t}$ as, otherwise, $\nu_{r_l}^a > m$. As a consequence, if l is
positive and $v(l) = v$, p_v selects $r_{\bar{v},t}$ as, if she selected $r_{v,t}$, she would have an
incentive to deviate (as $6 > 4$). Moreover, no other follower would select $r_{\bar{v},t}$
as, otherwise, p_v would deviate to r_t (as $6 > 4$). This implies that $\nu_{r_{\bar{v},t}}^a = 1$
1380 (follower p_v) and all the followers $p_{\bar{v},k}$ select resource $r_{\bar{v}}$, while the followers
 $p_{\phi,\bar{v}}$ choose resources $r_{\phi,\bar{v}}$. On the other hand, if l is negative and $v(l) = v$,
similar arguments allow us to conclude that $\nu_{r_{v,t}}^a = 1$ (follower p_v) and all the
followers $p_{v,k}$ select resource r_v , while the followers $p_{\phi,v}$ choose resources $r_{\phi,v}$.

As a result, either $\nu_{r_{v,t}}^a = 1$ or $\nu_{r_{\bar{v},t}}^a = 1$. In conclusion, we can define a truth assignment τ such that $\tau(v) = \text{T}$ if $a_{p_v} = r_{\bar{v},t}$ and $\tau(v) = \text{F}$ if $a_{p_v} = r_{v,t}$. Clearly, τ is well-defined. Moreover, as previously shown, for every $\phi \in C$ there exists a literal $l \in \phi$ such that $a_{p_{\phi,l}} = r_l$, which, letting $v = v(l)$, implies that $\nu_{r_{v,t}}^a = 1$. Thus, $\tau(v(l)) = \text{T}$ if l is positive, while $\nu_{r_{v,t}}^a = 1$ and $\tau(v(l)) = \text{F}$ if l is negative. Hence, τ satisfies all the clauses. \square

Theorem 2. *Computing a PSE in SSCGs with different action spaces is NP-hard even if the leader has only one action and the cost functions are monotonic.*

Proof. We provide a reduction from 3SAT showing that the existence of a polynomial-time algorithm for computing a PSE in SSCGs would allow us to solve any 3SAT instance in polynomial time. Specifically, given a 3SAT instance (C, V) and a real number $0 < \epsilon < 4$, we build an SSCG instance $\Gamma_\epsilon(C, V)$ such that it admits a PSE where the leader's cost is ϵ if and only if the 3SAT instance admits a *no* answer, i.e., if and only if (C, V) is not satisfiable. Instead, if the 3SAT instance has answer *yes*, i.e., if (C, V) is satisfiable, then the leader's cost is 4 in any PSE.

Mapping. $\Gamma_\epsilon(C, V)$ is defined as follows:

- $N = F \cup \{\ell\}$, where $F = \{p_{\phi,t} \mid \phi \in C\} \cup \{p_{v,t}, p_v, p_{\bar{v}} \mid v \in V\} \cup \{p_{l,\phi} \mid \phi \in C, l \in \phi\} \cup \{p_{\phi,v}, p_{\phi,\bar{v}} \mid \phi \in C, v \in V\}$;
- $R = \{r_t\} \cup \{r_\phi \mid \phi \in C\} \cup \{r_{v,t}, r_v, r_{\bar{v}} \mid v \in V\} \cup \{r_{\phi,v}, r_{\phi,\bar{v}} \mid \phi \in C, v \in V\}$;
- $A_{p_{\phi,t}} = \{r_\phi, r_t\} \quad \forall \phi \in C$;
- $A_{p_v} = \{r_{v,t}, r_v\}, A_{p_{\bar{v}}} = \{r_{v,t}, r_{\bar{v}}\}, A_{p_{v,t}} = \{r_{v,t}, r_t\} \quad \forall v \in V$;
- $A_{p_{l,\phi}} = \{r_\phi\} \cup \{r_{\phi,l}\} \quad \forall \phi \in C, \forall l \in \phi$;
- $A_{p_{\phi,v}} = \{r_v, r_{\phi,v}\}, A_{p_{\phi,\bar{v}}} = \{r_{\bar{v}}, r_{\phi,\bar{v}}\} \quad \forall \phi \in C, v \in V$;
- $A_\ell = \{r_t\}$.

The cost functions take values according to the following table, and satisfy $c_{r_{\bar{v}},f} = c_{r_v,f}$, $c_{r_{\phi,\bar{v}},f} = c_{r_{\phi,v},f}$, and $c_{r_t,f} = c_{r_t,\ell}$ (let us remark that, given $\epsilon < 4$, they are all monotonic functions of the resource congestion):

x	$c_{r_{\phi},f}$	$c_{r_v,f}$	$c_{r_{v,t},f}$	$c_{r_{\phi,v},f}$	$c_{r_t,f}$
1	2	1	2	0	ϵ
$[2, m]$	5	1	5	7	ϵ
$m+1$	5	6	5	7	ϵ
$[m+s+1, \infty]$	5	6	5	7	4

Figure 2 shows an example of the game instance $\Gamma_\epsilon(C, V)$.

Given (C, V) , $\Gamma_\epsilon(C, V)$ can be constructed in polynomial time, as it features $n = 3m + m + 3s + 2ms + 1$ players and $r = m + 3s + 2ms + 1$ resources. Observe

that, in $\Gamma_\epsilon(C, V)$, the leader can only select a single resource r_t and, hence, the only leader's commitment is $\sigma_\ell \in \Delta_\ell : \sigma_\ell(r_t) = 1$. As a result, the leader's cost is 4 if and only if all followers $p_{\phi,t}$ and $p_{v,t}$ select resource r_t ; otherwise, it is ϵ .

If. Suppose that the 3SAT instance has answer *no*, i.e., there is no truth assignment to the variables in V that satisfies all the clauses in C . We prove that, in that case, $\Gamma_\epsilon(C, V)$ admits a PSE with leader's cost equal to ϵ . By contradiction, let us assume there exists a PSE $\sigma = (\sigma_\ell, a)$ in which the leader's cost c_ℓ^σ is $4 > \epsilon$. We show that $a = (a_p)_{p \in F}$ can be employed to recover, in polynomial time, a truth assignment τ that satisfies all the clauses in C , which is a contradiction. First, let us notice that all the followers $p_{\phi,t}$ and $p_{v,t}$ select r_t in a as, otherwise, the leader's cost would be $\epsilon < 4$. As a result, $a_{p_{\phi,t}} = r_t$ for all $\phi \in C$ and $a_{p_{v,t}} = r_t$ for all $v \in V$. Thus, for every $v \in V$, at least one between p_v and $p_{\bar{v}}$ must select $r_{v,t}$ as, otherwise, player $p_{v,t}$ would deviate from r_t (as $2 < 4$). If $a_{p_v} = r_{v,t}$, then all the followers $p_{\phi,v}$ select r_v as, otherwise, p_v would have an incentive to deviate from $r_{v,t}$ (since $\nu_{r_v}^a < m$ and p_v would incur a cost of $1 < 2$ by switching to r_v). Similarly, if $a_{p_{\bar{v}}} = r_{v,t}$, then all the followers $p_{\phi,\bar{v}}$ select $r_{\bar{v}}$. Let us define a truth assignment τ such that $\tau(v) = \text{T}$ if $a_{p_v} = r_v$, $\tau(v) = \text{F}$ if $a_{p_{\bar{v}}} = r_{\bar{v}}$ and $\tau(v)$ is either T or F whenever $a_{p_v} = a_{p_{\bar{v}}} = r_{v,t}$. Clearly, τ is well-defined. Since $a_{p_{\phi,t}} = r_t$ for all $\phi \in C$, there must be at least one follower using resource r_ϕ for every $\phi \in C$ as, otherwise, $p_{\phi,t}$ would have an incentive to deviate from r_t to r_ϕ (as $4 > 2$). Thus, for each $\phi \in C$ there must be a literal $l \in \phi$ such that $a_{p_{l,\phi}} = r_\phi$. This implies that $a_{p_{\phi,l}} = r_{\phi,l}$ as, otherwise, follower $p_{l,\phi}$ would deviate from r_ϕ to $r_{\phi,l}$ (as $2 > 0$). As a result, it must be the case that $a_{p_l} = r_l$, since, if $a_{p_l} = r_{v(l),t}$, then $p_{\phi,l}$ would select r_l instead of $r_{\phi,l}$. Thus, $\tau(v(l)) = \text{T}$ if l is positive or $\tau(v(l)) = \text{F}$ if it is negative. Therefore, τ satisfies all the clauses, which is a contradiction.

Only if. Suppose that the 3SAT instance admits answer *yes*, i.e., there exists a truth assignment to the variables which satisfies all the clauses in C . We prove that in any PSE of $\Gamma_\epsilon(C, V)$ the leader's cost is $4 > \epsilon$. Let $\tau : V \rightarrow \{\text{T}, \text{F}\}$ be one such truth assignment. We show how to recover from τ a followers' action profile $a = (a_p)_{p \in F} \in \prod_{p \in F} A_p$ such that $a \in E^{\sigma_\ell}$, with $\sigma = (\sigma_\ell, a)$ providing $c_\ell^\sigma = 4$. Since 4 is the maximum cost the leader can achieve and the followers behave pessimistically, σ is clearly a PSE. In particular, let $a_{p_{\phi,t}} = r_t$ for all $\phi \in C$ and $a_{p_{v,t}} = r_t$ for all $v \in V$. Moreover, if $\tau(v) = \text{T}$, let $a_{p_v} = r_v$, $a_{p_{\bar{v}}} = r_{v,t}$, and, for all $\phi \in C$, $a_{p_{\phi,v}} = r_{\phi,v}$ and $a_{p_{\phi,\bar{v}}} = r_{\bar{v}}$. Additionally, for every clause $\phi \in C$ and $l \in \phi$ such that $v(l) = v$, let $a_{p_{l,\phi}} = r_\phi$ if l is positive, while $a_{p_{l,\phi}} = r_{\phi,l}$ if it is negative. Conversely, if $\tau(v) = \text{F}$, let $a_{p_{\bar{v}}} = r_{\bar{v}}$, $a_{p_v} = r_{v,t}$, and, for all $\phi \in C$, $a_{p_{\phi,v}} = r_{\phi,\bar{v}}$ and $a_{p_{\phi,v}} = r_v$. Furthermore, for every clause $\phi \in C$ and $l \in \phi$ such that $v(l) = v$, let $a_{p_{l,\phi}} = r_\phi$ if l is negative, and $a_{p_{l,\phi}} = r_{\phi,l}$ if it is positive. Notice that, since either $\tau(v) = \text{T}$ or $\tau(v) = \text{F}$, one between p_v and $p_{\bar{v}}$ selects $r_{v,t}$. Assume, w.l.o.g., $\tau(v) = \text{T}$ (as the other case is analogous). First, no follower $p_{\phi,v}$ would deviate from $r_{\phi,v}$ to r_v , as, otherwise, she would incur a cost of 1, rather than 0. The same holds for followers $p_{\phi,\bar{v}}$, as their cost is 1 while, if any of them switched to $r_{\phi,\bar{v}}$, she would incur a cost of 7, because $a_{p_{\bar{v},\phi}} = r_{\phi,\bar{v}}$. Similarly, since there is one follower selecting $r_{v,t}$, follower $p_{v,t}$

would not deviate from r_t (as $4 < 5$), while follower p_v would not deviate from r_v because her cost is $1 < 5$ and $p_{\bar{v}}$ would not switch from $r_{v,t}$ (as she would get 6 rather than 1). Furthermore, since τ is a truth assignment satisfying (C, V) , for each clause $\phi \in C$ there exists at least one literal $l \in \phi$ that evaluates to true under τ . Thus, $p_{l,\phi}$ would not deviate from r_ϕ (as she pays either 2 or 5 instead of 7). Thus all the followers $p_{\phi,t}$ would not deviate from r_t (as $4 < 5$) and we can conclude that a is an NE. Since $m + s$ follower use r_t , $c_\ell^\sigma = 4$. \square

Theorem 3. *Computing an OSE in SSSCGs is NP-hard.*

Proof. We prove the theorem using a reduction from K -PARTITION, showing that the existence of a polynomial-time algorithm for computing an OSE in SSSCGs would allow us to solve K -PARTITION in polynomial time. Let (S, K) be a given K -PARTITION instance, and let us recall that we assumed $x_i \leq s$ for all $x_i \in S$, where $s = \frac{1}{2} \sum_{x_i \in S} x_i$. Clearly, any valid partition $(S', S \setminus S')$ is uniquely defined by a subset $S' \subseteq S$ such that $\sum_{x_i \in S'} x_i = s$ and $|S'| = K$. Let $w_i = \frac{x_i}{s}$ for all $x_i \in S$. Due to having $x_i \leq s$ for all $x_i \in S$, we also have $w_i \leq 1$. Given (S, K) , we build an instance $\Gamma_\epsilon(S, K)$ of SSSCG with $0 < \epsilon < 1$ such that there exists an OSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$ if and only if the K -PARTITION instance (S, K) admits answer *yes*.

Mapping. $\Gamma_\epsilon(S, K)$ is defined as follows:

- $N = F \cup \{\ell\}$, with $|F| = 4|S| + 2$;
- $R = \{r_{t_1}\} \cup \{r_{t_2}\} \cup \{r_i \mid x_i \in S\}$;

The players' cost functions are specified in the following table:

x	$c_{r_i,f}$	$c_{r_i,\ell}$	$c_{r_{t_1},f}$	$c_{r_t,\ell}$	$c_{r_{t_2},f}$	$c_{r_{t_2},\ell}$
1	$2s$	s	$3s^2$	s^4	1	s^4
2	0	s	$3s^2$	s^4	$4s^2$	s^4
3	$\frac{1}{w_i}$	ϵ	$3s^2$	s^4	$4s^2$	s^4
4	$\frac{2s - \frac{w_i}{s} + 1}{w_i}$	s	$3s^2$	s^4	$4s^2$	s^4
$[5, 4 S - 2K]$	$4s^2$	s	$3s^2$	s^4	$4s^2$	s^4
$4 S - 2K + 1$	$4s^2$	s	$2s$	s^4	$4s^2$	s^4
$4 S - 2K + 2$	$4s^2$	s	1	s^4	$4s^2$	s^4
$[4 S - 2K + 3, \infty]$	$4s^2$	s	0	s^4	$4s^2$	s^4

Clearly, $\Gamma_\epsilon(S, K)$ can be built in polynomial time, as it features $n = 4|S| + 3$ players and $r = |S| + 2$ resources.

If. Suppose that the K -PARTITION instance (S, K) admits a *yes* answer. Let $S' \subseteq S$ be a set of integers with $|S'| = K$ and $\sum_{x_i \in S'} x_i = s$. We prove that $\Gamma_\epsilon(S, K)$ admits an OSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$. Given S' , we build the followers' configuration $\nu \in \mathbb{R}^r$ and the leader's strategy $\sigma_\ell \in \Delta_\ell$. Let $\nu_{r_i} = 2$ and $\sigma_\ell(r_i) = w_i$ for all $x_i \in S'$, while, for every $x_i \notin S'$, let $\nu_{r_i} = 0$

and $\sigma_\ell(r_i) = 0$. Moreover, let $\nu_{r_{t_1}} = 4|S| - 2K + 1$, $\sigma_\ell(r_{t_1}) = 0$, $\nu_{r_{t_2}} = 1$, and $\sigma_\ell(r_{t_2}) = 0$. First, let us observe that the leader's strategy σ_ℓ is well-defined, as

$$\sum_{x_i \in S} \sigma_\ell(r_i) + \sigma_\ell(r_{t_1}) + \sigma_\ell(r_{t_2}) = \sum_{x_i \in S'} \sigma_\ell(r_i) = \sum_{x_i \in S'} w_i = \sum_{x_i \in S'} \frac{x_i}{s} = 1,$$

where the last equality follows from the fact that S' defines a partition of S . Next, we show that ν is an NE for σ_ℓ with the following argument.

- All the followers who selected resource r_i , with $x_i \in S'$, do not have any incentive to change resource, as their cost is $w_i \cdot \frac{1}{w_i} = 1$ and they cannot improve it by switching to another resource. Indeed, if they selected a resource r_j with $x_j \in S'$, they would incur a cost of $\frac{1}{w_j} \cdot (1 - w_j) + \frac{2s - \frac{1}{w_j} + 1}{w_j} \cdot w_j = 2s > 1$. Similarly, their cost would be $2s$ if they choose r_j with $x_j \notin S'$. They would not benefit from choosing resource r_{t_1} , as they would incur a cost of 1, which is the same as their current cost, and they would not switch to resource r_{t_2} , as their cost would become $4s^2 > 1$.
- All the followers who selected resource r_{t_1} incur a cost of $2s$. Thus, they do not have an incentive to deviate to a resource r_i with $x_i \in S'$, as they would still incur a cost of $2s$. The same holds for resources r_i with $x_i \notin S'$. Similarly, if they chose to play r_{t_2} , they would incur a cost of $4s^2 > 2s$.
- The follower who chose resource r_{t_2} does not deviate, as her cost is 1 and she would incur a cost of $2s$ and 1 if she switched to resource r_i (for some $x_i \in S$) and r_{t_1} , respectively.

Overall, the leader's cost is:

$$\begin{aligned} c_\ell^\sigma &= \sum_{x_i \in S} \sigma_\ell(r_i) c_{r_i, \ell}(\nu_{r_i} + 1) + \sigma_\ell(r_{t_1}) c_{r_{t_1}, \ell}(\nu_{r_{t_1}} + 1) + \sigma_\ell(r_{t_2}) c_{r_{t_2}, \ell}(\nu_{r_{t_2}} + 1) = \\ &= \sum_{x_i \in S'} \sigma_\ell(r_i) c_{r_i, \ell}(\nu_{r_i} + 1) = \sum_{x_i \in S'} \epsilon w_i = \epsilon. \end{aligned}$$

Only if. Suppose that $\Gamma_\epsilon(S, K)$ has an OSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$. Then, $\sigma_\ell(r_{t_1}) = \sigma_\ell(r_{t_2}) = 0$ must hold. Moreover, the leader must place positive probability only on resources r_i with $\nu_{r_i} = 2$. Clearly, there is always a resource r_i with $\nu_{r_i} = 2$ and $\sigma_\ell(r_i) > 0$. Next, we prove that $\nu_{r_{t_1}} = 4|S| - 2K + 1$. By contradiction, assume that $\nu_{r_{t_1}} \neq 4|S| - 2K + 1$. Three cases are possible.

- $\nu_{r_{t_1}} = 0$ implies that either there exists at least one resource r_i with $\nu_{r_i} \geq 5$ or $\nu_{r_{t_2}} = 2$, but, then, the followers who chose r_i or, respectively, r_{t_2} , would deviate by choosing r_{t_1} , decreasing their cost from $4s^2$ to $3s^2$.
- $1 \leq \nu_{r_{t_1}} \leq 4|S| - 2K$ implies that the followers who selected r_{t_1} incur a cost of $3s^2$. Thus, they would deviate to some resource r_i with $\nu_{r_i} = 2$, since their cost would be at most $\frac{2s - \frac{1}{w_i} + 1}{w_i} < 3s^2$.

- $\nu_{r_{t_1}} \geq 4|S| - 2K + 2$ implies that the followers' cost when they deviate by playing resource r_{t_1} is 0. Thus, the followers who selected a resource r_i with $\nu_{r_i} = 2$ and $\sigma_\ell(r_i) > 0$ would change resource, since their current cost is strictly greater than 0.

The only remaining option for $\nu \in E^{\sigma_\ell}$ is $\nu_{r_{t_1}} = 4|S| - 2K + 1$. Then, $\nu_{r_{t_2}} = 1$ must hold as, if $\nu_{r_{t_2}} = 0$, a follower would switch from resource r_{t_1} to resource r_{t_2} (incurring a cost of 1 instead of $2s > 1$), while, if $\nu_{r_{t_2}} \geq 2$, the followers who selected resource r_{t_2} would deviate to resource r_{t_1} (incurring a cost of 1 instead of $4s^2 > 1$). Let us now consider a resource r_i with $\nu_{r_i} = 2$. We prove that $\sigma_\ell(r_i) = w_i$ by contradiction. Two cases are possible.

- If $\sigma_\ell(r_i) < w_i$, the followers' cost by switching to resource r_i satisfies

$$\frac{1}{w_i}(1 - \sigma_\ell(r_i)) + \frac{2s - \frac{1}{w_i} + 1}{w_i}\sigma_\ell(r_i) < \frac{1}{w_i}(1 - w_i) + \frac{2s - \frac{1}{w_i} + 1}{w_i}w_i = 2s,$$

where the inequality holds since the left-most quantity is a convex combination of $\frac{1}{w_i}$ and $\frac{2s - \frac{1}{w_i} + 1}{w_i}$ with weights $(1 - \sigma_\ell(r_i))$ and $\sigma_\ell(r_i)$, and, since $\frac{1}{w_i} < \frac{2s - \frac{1}{w_i} + 1}{w_i}$, its maximum for $\sigma_\ell(r_i) \leq w_i$ is attained at $\sigma_\ell(r_i) = w_i$. Thus, we deduce that a follower would deviate from resource r_{t_1} to resource r_i (as her current cost is $2s$), contradicting the fact that $\nu \in E^{\sigma_\ell}$.

- If $\sigma_\ell(r_i) > w_i$, we reach a contradiction since the cost incurred by the followers who are using resource r_i would be $\frac{1}{w_i}\sigma_\ell(r_i) > 1$ and they would deviate playing resource r_{t_1} , decreasing their cost to 1.

We have shown that $\sigma_\ell(r_i) = w_i$ for every resource r_i with $\nu_{r_i} = 2$. Finally, let r_i be a resource with $\nu_{r_i} \neq 2$. Clearly, it must be the case that $\sigma_\ell(r_i) = 0$ since the leader's cost is ϵ . Moreover, it cannot be the case that $\nu_{r_i} = 1$, as, if it were the case, the follower would deviate to resource r_{t_1} with a cost of 1, instead of $2s$. Similarly, $\nu_{r_i} \geq 3$ cannot hold, as one of the followers who are selecting resource r_i would deviate playing r_{t_1} , since her current cost is greater than 1. Thus, either $\nu_{r_i} = 2$ or $\nu_{r_i} = 0$. As a consequence, there are K resources r_i with $\nu_{r_i} = 2$ and $\sigma_\ell(r_i) = w_i$, and $|S| - K$ resources r_i with $\nu_{r_i} = 0$ and $\sigma_\ell(r_i) = 0$. Let us define S' as the set of integers $x_i \in S$ such that the corresponding resources r_i satisfy $\nu_{r_i} = 2$. Since $\sum_{x_i \in S'} \sigma_\ell(r_i) = 1$ and $\sigma_\ell(r_i) = w_i$ for all such resources r_i , we can conclude that $\sum_{x_i \in S'} w_i = \sum_{x_i \in S'} \frac{x_i}{s} = 1$, and, thus, $\sum_{x_i \in S'} x_i = s$. As a result, $(S', S \setminus S')$ is solution to K -PARTITION. \square

Theorem 4. *The problem of computing an OSE in SSCGs is not in Poly-APX unless $P = NP$.*

Proof. In order to prove the result, we rely on the reduction introduced in the proof of Theorem 3. We have already shown that in an OSE $\sigma = (\sigma_\ell, \nu)$ of $\Gamma_\epsilon(S, K)$ it holds $c_\ell^\sigma = \epsilon$ if and only if the corresponding instance of K -PARTITION (S, K) admits a *yes* answer. Now, we prove that, when the K -PARTITION instance admits a *no* answer, $c_\ell^\sigma \geq 1$ in any OSE. By contradiction,

assume that there exists an OSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma < 1$. Let $S' \subseteq S$ be the set of integers corresponding to a group of resources r_i with $\nu_{r_i} = 2$ (at least one must exist since the leader's cost is smaller than 1). Then, $\sum_{x_i \in S'} \sigma_\ell(r_i) > \frac{s-1}{s}$ since $\sum_{x_i \in S \setminus S'} \sigma_\ell(r_i) + \sigma_\ell(r_{t_1}) + \sigma_\ell(r_{t_2})$ must be smaller than $\frac{1}{s}$ in order to have $c_\ell^\sigma < 1$. Moreover, $\sigma_\ell(r_{t_1}) \leq \frac{1}{s^4}$ and $\sigma_\ell(r_{t_2}) \leq \frac{1}{s^4}$ must both hold as, if not, we would get $c_\ell^\sigma \geq 1$. We prove, now, that $\nu_{r_{t_1}} = 4|S| - 2K + 2$ by contradiction. We identify three cases:

- $\nu_{r_{t_1}} = 0$ implies that either there exists at least one resource r_i with $\nu_{r_i} \geq 5$ or $\nu_{r_{t_2}} = 2$, and, thus, either a follower who selected resource r_i or one who selected resource r_{t_2} would have an incentive to deviate to resource r_{t_1} (as $4s^2 > 3s^2$).
- $1 \leq \nu_{r_{t_1}} \leq 4|S| - 2K - 1$ implies that one of the followers who selected r_{t_1} would have an incentive to deviate to resource r_i with $\nu_{r_i} = 2$, as she would incur a cost smaller than or equal to $\frac{2s - \frac{1}{w_i} + 1}{w_i} < 3s^2$.
- $\nu_{r_{t_1}} = 4|S| - 2K$ implies that the cost incurred by the followers who selected resource r_{t_1} is greater or equal than $3s^2(1 - \frac{1}{s^4}) + \frac{2}{s^3}$, as $\sigma_\ell(r_{t_1}) \leq \frac{1}{s^4}$. Thus, since $\frac{2s - \frac{1}{w_i} + 1}{w_i} < 2s^2 < 3s^2 - \frac{3}{s^2} + \frac{2}{s^2}$, these followers would deviate from r_{t_1} to a resource r_i with $\nu_{r_i} = 2$.
- $\nu_{r_{t_1}} \geq 4|S| - 2K + 2$ implies that the followers' cost after deviating to resource r_{t_1} would be 0 and, since there exists at least one resource r_i with $\nu(r_i) = 2$ and $\sigma(r_i) > 0$, one of the followers who selected such resource would switch from it in favor of r_{t_1} .

Thus, $\nu_{r_{t_1}} = 4|S| - 2K + 1$. Let us consider resource r_{t_2} . If $\nu_{r_{t_2}} = 0$, the followers' cost incurred when deviating to resource r_{t_2} would be smaller than or equal to $(1 - \frac{1}{s^4}) + \frac{4}{s^2}$ (as $\sigma_\ell(r_{t_2}) \leq \frac{1}{s^4}$), while the cost incurred by choosing resource r_{t_1} is at least $2s(1 - \frac{1}{s^4}) + \frac{1}{s^4} > (1 - \frac{1}{s^4}) + \frac{4}{s^4}$. Instead, if $\nu_{r_{t_2}} \geq 2$, the followers' cost for resource r_{t_2} is $4s^2$ and they would have an incentive to deviate to r_{t_1} to decrease their cost to 1 or less. Thus, $\nu_{r_{t_2}} = 1$. We deduce $\sigma_\ell(r_{t_1}) = 0$ as, otherwise (i.e., with $\sigma_\ell(r_{t_1}) > 0$), a follower would deviate from resource r_{t_2} to r_{t_1} , decreasing her cost to 1 or less. Let us focus on resources r_i with $\nu_{r_i} = 2$. If $\sigma_\ell(r_i) < w_i$, the followers' cost of deviating to r_i is

$$\frac{1}{w_i}(1 - \sigma_\ell(r_i)) + \frac{2s - \frac{1}{w_i} + 1}{w_i}\sigma_\ell(r_i) < \frac{1}{w_i}(1 - w_i) + \frac{2s - \frac{1}{w_i} + 1}{w_i}w_i = 2s,$$

and they would deviate from r_{t_1} to r_i , as their current cost is $2s$. Instead, if $\sigma_\ell(r_i) > w_i$ the cost of any follower who selected r_i is greater than 1 and she would deviate to resource r_{t_1} to decrease her cost to 1. Thus, $\sigma_\ell(r_i) = w_i$ for all resources r_i with $\nu_{r_i} = 2$. Now, let us consider a resource r_i with $\nu_{r_i} \neq 2$. Clearly, $\sigma_\ell(r_i) \leq \frac{1}{s}$ must hold since $c_\ell^\sigma \leq 1$. If $\nu_{r_i} = 1$, the followers' cost for resource r_i is at least $2s\frac{1}{s} > 1$ while, if $\nu_{r_i} \geq 3$, the followers' cost for resource r_i is at least $\frac{1}{w_i} > 1$. In both cases, the followers who selected resource

1580 r_i would have an incentive to deviate to r_{t_1} (as they would pay 1). Thus, either $\nu_{r_i} = 2$ or $\nu_{r_i} = 0$. As a consequence, there are K resources r_i with $\nu_{r_i} = 2$ and $\sigma_\ell(r_i) = w_i$ and $|S| - K$ resources r_i with $\nu_{r_i} = 0$. If $c_\ell^\sigma \leq 1$, there must be a subset $S' \subseteq S$ with $\sum_{x_i \in S'} \sigma_\ell(r_i) = \sum_{x_i \in S'} w_i > \frac{s-1}{s}$, which implies that $\sum_{x_i \in S'} \frac{x_i}{s} > \frac{s-1}{s}$ and $\sum_{x_i \in S'} x_i > s - 1$. Note that $x_i \in \mathbb{N}$ and
1585 $\sum_{x_i \in S'} x_i = s \sum_{x_i \in S'} w_i = s \sum_{x_i \in S'} \sigma_\ell(r_i) \leq s$. Thus, $\sum_{x_i \in S'} x_i = s$ and $(S', S \setminus S')$ is solution to K -PARTITION. So far, we have proven that $\Gamma_\epsilon(S, K)$ admits an OSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$ if and only if (S, K) has answer *yes* and that, otherwise, $c_\ell^\sigma \geq 1$ in any OSE. Let $\epsilon = \frac{1}{2^{n+r}}$. Assume that there exists a polynomial-time approximation algorithm \mathcal{A} with approximation factor
1590 $\text{poly}(n, r)$, i.e., a polynomial function of n and r . Assume (S, K) has answer *yes*. \mathcal{A} applied to $\Gamma_\epsilon(S, K)$ would return a solution with $c_\ell^\sigma \leq \frac{1}{2^{n+r}} \text{poly}(n, r)$. Since, for n and r large enough, $\frac{1}{2^{n+r}} \text{poly}(n, r) < 1$, \mathcal{A} would allow us to decide in polynomial time whether (S, K) has a *yes* or *no* answer, a contradiction unless $P = \text{NP}$. \square

1595 **Theorem 5.** *Computing a PSE in SSSCGs is NP-hard.*

Proof. We provide a reduction from PARTITION showing that the existence of a polynomial-time algorithm for computing a PSE in SSSCGs would allow us to solve PARTITION in polynomial time. Given a PARTITION instance with a set S of positive integers, let, as in the previous proof, $s = \frac{1}{2} \sum_{x_i \in S} x_i$ and
1600 $w_i = \frac{x_i}{s}$ for all $x_i \in S$. Let us also recall that we assumed, w.l.o.g., $x_i \leq s$ for all $x_i \in S$, and, thus, $w_i \leq 1$. Given S , we build an instance $\Gamma_\epsilon(S)$ of SSSCG with $0 < \epsilon < 1$ such that $c_\ell^\sigma = \epsilon$ in a PSE $\sigma = (\sigma_\ell, \nu)$ if and only if the PARTITION instance admits answer *yes*.

Mapping. $\Gamma_\epsilon(S)$ is defined as follows:

- 1605
- $N = F \cup \{\ell\}$, with $|F| = 3|S|$;
 - $R = \{r_t\} \cup \{r_i \mid i \in S\}$;

with the following cost functions:

x	$c_{r_i, f}$	$c_{r_i, \ell}$	$c_{r_t, f}$	$c_{r_t, \ell}$
1	0	ϵ	1	s^4
2	$\frac{1}{w_i - \frac{1}{s^4}}$	s^4	1	s^4
3	$\frac{1}{1 - w_i - \frac{1}{s^4}}$	ϵ	1	s^4
4	0	s^4	1	s^4
$[5, \infty]$	s	ϵ	1	s^4

Clearly, $\Gamma_\epsilon(S)$ can be built in polynomial time, as it features $n = 3|S|$ players, and $r = |S| + 1$ resources.
1610

If. Suppose that the PARTITION instance admits a *yes* answer, and let $S' \subseteq S$ be such that $\sum_{x_i \in S'} x_i = s$. We show that there exists a PSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$. Let $\sigma_\ell(r_i) = w_i$ for all $x_i \in S'$, $\sigma_\ell(r_i) = 0$ for all $x_i \notin S'$, and

1615 $\sigma_\ell(r_t) = 0$. We prove that $c_\ell^\sigma = \epsilon$ for any $\sigma = (\sigma_\ell, \nu)$ with $\nu \in E^{\sigma_\ell}$. Assume, by contradiction, that there exists an NE $\nu \in E^{\sigma_\ell}$ such that $c_\ell^\sigma = \epsilon$. This implies that there exists a resource r_i with $x_i \in S'$ and either $\nu_{r_i} = 1$ or $\nu_{r_i} = 3$. If $\nu_{r_i} = 1$, the cost incurred by the followers who select r_i is $\frac{1}{w_i - \frac{1}{s^4}} w_i > 1$ and any of them would deviate to resource r_t to decrease her cost to 1. If $\nu_{r_i} = 3$, the followers' cost is $\frac{1}{1 - w_i - \frac{1}{s^4}} (1 - w_i) > 1$ and any of them would deviate to resource r_t . In both cases, this contradicts the fact that ν is an NE, and, thus, it must be that $\nu(r_i) \neq 1$ and $\nu(r_i) \neq 3$ for all $x_i \in S'$. As a result, $c_\ell^\sigma = \epsilon$ for any $\sigma = (\sigma_\ell, \nu)$ with $\nu \in E^{\sigma_\ell}$.

Only if. Suppose that $\Gamma_\epsilon(S)$ admits a PSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$. Then, $\sigma_\ell(r_t) = 0$ and $\sigma_\ell(r_i) > 0$ only if resource r_i is such that $\nu_{r_i} \neq 1$ and $\nu_{r_i} \neq 3$. Let us define $R' \subseteq R$ as the set of resources r_i with $\sigma_\ell(r_i) \leq w_i - \frac{1}{s^4}$, R'' as the set of resources r_i with $w_i - \frac{1}{s^4} < \sigma_\ell(r_i) < w_i + \frac{1}{s^4}$, and R''' as the set of resources r_i with $\sigma_\ell(r_i) \geq w_i + \frac{1}{s^4}$. Let $\nu \in \mathbb{R}^r$ be a followers' configuration such that $\nu_{r_i} = 1$ for all $r_i \in R'$, $\nu_{r_i} = 0$ for all $r_i \in R''$, $\nu_{r_i} = 3$ for all $r_i \in R'''$, and $\nu_{r_t} = 3|S| - \sum_{r_i \in R \setminus \{r_t\}} \nu_{r_i}$. First, we show that $\nu \in E^{\sigma_\ell}$. Indeed, all the followers who selected resource r_t incurs a cost of 1, all those who selected a resource $r_i \in R'$ incur a cost of $\frac{1}{w_i - \frac{1}{s^4}} \sigma_\ell(r_i) < 1$, and all those who selected resource $r_i \in R'''$ incur a cost of $\frac{1}{1 - w_i - \frac{1}{s^4}} (1 - \sigma_\ell(r_i)) < 1$. If any follower deviated, she would incur a cost greater than or equal than 1. In particular, no follower would deviate to a resource $r_i \in R'$, as she would incur a cost that is a convex combination of values greater than 1. Similarly, no follower would deviate to a resource $r_i \in R''$ or $r_i \in R'''$, as she would incur a cost of, respectively, $\frac{1}{w_i - \frac{1}{s^4}} \sigma(r_i) > 1$ or $s\sigma_\ell(r_i) > 1$. Finally, no follower has an incentive to switch to resource r_t , as her cost would not decrease. This shows that, in the followers' game resulting from σ_ℓ , there exists an NE such that, whenever the leader selects a resource r_i in $R' \cup R''$, she incurs a cost of s^4 . Thus, given that $c_\ell^\sigma = \epsilon$, $R' = R''' = \emptyset$ must hold. Let us define $S' \subseteq S$ as the set of integers $x_i \in S$ whose corresponding resource r_i is such that $w_i - \frac{1}{s^4} < \sigma_\ell(r_i) < w_i + \frac{1}{s^4}$. For all the other resources r_i , it must be $\sigma_\ell(r_i) = 0$. Since $\sum_{x_i \in S'} \sigma_\ell(r_i) = 1$, we have $\sum_{x_i \in S'} (w_i - \frac{1}{s^4}) < 1 < \sum_{x_i \in S'} (w_i + \frac{1}{s^4})$, and, therefore,

$$\frac{s-1}{s} < 1 - \sum_{x_i \in S'} \frac{1}{s^4} < \sum_{x_i \in S'} w_i < 1 + \sum_{x_i \in S'} \frac{1}{s^4} < \frac{s+1}{s},$$

which implies $s-1 < \sum_{x_i \in S'} x_i < s+1$. Since $\sum_{x_i \in S'} x_i$ is an integer quantity, we deduce $\sum_{x_i \in S'} x_i = s$, implying that S' is a solution to PARTITION. \square

1625 **Theorem 6.** *The problem of computing a PSE in SSCGs is not in Poly-APX unless $P = NP$.*

Proof. In order to prove the result, we rely on the reduction introduced in the proof of Theorem 5. We have already shown that in a PSE $\sigma = (\sigma_\ell, \nu)$ of $\Gamma_\epsilon(S)$ it holds $c_\ell^\sigma = \epsilon$ if and only if the corresponding instance of PARTITION admits a yes answer. Now, we show that, if the partition problem has no answer, 1630

then $c_\ell^\sigma \geq 1$ in any PSE. Suppose, by contradiction, that there is a leader's strategy σ_ℓ such that all NEs of the resulting followers' game provide the leader with a cost smaller than 1. Then, $\sigma_\ell(r_i) < \frac{1}{s^4}$ for all resources r_i such that $\nu_{r_i} = 3$, $\sigma_\ell(r_i) < \frac{1}{s^4}$ for all resources r_i such that $\nu_{r_i} = 1$, and $\sigma_\ell(r_t) < \frac{1}{s^4}$.
1635 If there is a resource r_i with $\sigma_\ell(r_i) > w_i + \frac{1}{s^4}$, we have already proven that there is an NE with $\nu_{r_i} = 3$ providing the leader with a cost greater than $s^4 \sigma_\ell(r_i) > 1$. Consider the set $S'' \subseteq S$ of integers x_i corresponding to resources r_i with $\sigma_\ell(r_i) \leq w_i - \frac{1}{s^4}$. We have already shown that there is an NE with $\nu_{r_i} = 1$ for all $x_i \in S''$. Since the leader can select these resources with, at most,
1640 probability $\frac{1}{s^4}$ (as $\sum_{x_i \in S''} \sigma_\ell(r_i) \leq \frac{1}{s^4}$), there is a set S' of resources r_i with $w_i - \frac{1}{s^4} < \sigma_\ell(r_i) < w_i + \frac{1}{s^4}$ and $\sum_{x_i \in S'} \sigma_\ell(r_i) \geq 1 - \frac{1}{s^4}$. From $\sum_{x_i \in S'} (w_i + \frac{1}{s^4}) > \sum_{x_i \in S'} \sigma_\ell(r_i) \geq 1 - \frac{1}{s^4}$, we obtain $\sum_{x_i \in S'} w_i > 1 - \frac{1}{s^4} - \frac{|S'|}{s^4} > \frac{s-1}{s}$. From $\sum_{x_i \in S'} (w_i - \frac{1}{s^4}) < \sum_{x_i \in S'} \sigma_\ell(r_i) \leq 1$, we deduce $\sum_{x_i \in S'} w_i < 1 + \frac{|S'|}{s^4} < \frac{s+1}{s}$. Thus, $s - 1 < \sum_{x_i \in S'} x_i < s + 1$ and, since $\sum_{x_i \in S'} x_i$ is an integer quantity,
1645 we have that $\sum_{x_i \in S'} x_i = s$, showing that S' is a solution to PARTITION. We have proven that $\Gamma_\epsilon(S)$ admits a PSE $\sigma = (\sigma_\ell, \nu)$ in which $c_\ell^\sigma = \epsilon$ if and only if the PARTITION instance has a *yes* answer, while, otherwise, $c_\ell^\sigma \geq 1$ in any PSE. Let $\epsilon = \frac{1}{2^{n+r}}$. Assume that there exists a polynomial-time approximation algorithm \mathcal{A} with approximation factor $\text{poly}(n, r)$, i.e., a polynomial function
1650 of n and r . Assume the PARTITION instance has answer *yes*. \mathcal{A} applied to $\Gamma_\epsilon(S)$ would return a solution with $c_\ell^\sigma \leq \frac{1}{2^{n+r}} \text{poly}(n, r)$. Since, for n and r large enough, $\frac{1}{2^{n+r}} \text{poly}(n, r) < 1$, \mathcal{A} would allow us to decide in polynomial time whether the PARTITION instance has a *yes* or *no* answer, a contradiction unless $P = NP$. \square

1655 Appendix B. Additional experimental results

In this section, we present a heuristic algorithm which combines Algorithm 1 with best-response dynamics using the result produced by the former as starting solution for the latter.

We remark that when Algorithm 1 runs on games with different action spaces
1660 and/or generic cost functions we have no guarantees on the returned solution. In particular, such solution may not be an OSE, neither a feasible strategy profile, i.e., one in which the followers' configuration is an NE given the leader's strategy. However, using best-response dynamics starting from the solution returned by Algorithm 1 guarantees to find a strategy profile which is feasible.
1665 The reason why we study this heuristic is that, in principle, one could expect that, in SSCGs with almost identical action spaces and almost monotonic cost functions, Algorithm 1 may return a strategy profile from which best-response dynamics could lead to solutions better than those obtained by starting from randomly generated followers' configurations (as in the plain best-response dynamics heuristic we considered in Section 8).
1670

In more detail, our heuristic algorithm works as follows. We first randomly order the followers and assign each of them, one at a time, to a resource according to the procedure in Algorithm 1, i.e., each follower is assigned to the

resource with the current smallest cost, breaking the ties in favor of the leader. Then, we run best-response dynamics starting from the obtained followers' configuration. We compare this heuristic algorithm to the best-response dynamics heuristic with randomly generated starting points and our MILP formulations, using three classes of SSCGs instances:

- *Almost monotonic cost functions and identical action spaces.* We assume $r = 20$ resources, $|F| \in \{20, 40, 60, 80, 100\}$ followers, and so-called *single-peak* cost functions, which are built by increasing the resource cost until a given level of congestion, or *peak* (sampled uniformly at random in $\{1, \dots, |N|\}$) and decreasing it after the peak, sampling each increment/decrement uniformly at random in $\{1, \dots, (n-1)r\}$.
- *Monotonic cost functions and almost identical action spaces.* We assume $r = 20$ resources, $|F| \in \{20, 40, 60, 80, 100\}$ followers, cost functions increasing by a value in $\{1, \dots, (n-1)r\}$ for each additional unit of congestion, and 15 actions per follower.
- *Generic cost functions and almost identical action spaces.* We assume $r = 20$ resources, $|F| \in \{20, 40, 60, 80, 100\}$ followers, 15 actions per follower, and cost functions randomly generated by sampling from $\{1, \dots, (n-1)r\}$ with uniform probability.

We generate 15 instances for each combination of the parameters.

In games with almost monotonic (i.e., single-peak) cost functions and identical action spaces, our heuristic combining Algorithm 1 with best-response dynamics performs quite well, even if its performance is worse than the one achieved with the plain best-response dynamics heuristic we introduced in Section 8. In games with monotonic cost functions and almost identical action spaces, the performance of the two heuristics is comparable. This is reasonable, since in these game there are few pure NEs and both algorithms achieve the same solutions. The results on these two classes of instances are reported in Figure B.6 (a).

When we completely drop the monotonicity assumption of the cost functions (see Figure B.6 (b)), the heuristic based on Algorithm 1 performs much worse than the plain best-response dynamics heuristic.




Figure B.6: Comparison of the average leader's cost of the solutions returned by our MILP formulations, the best-response dynamics heuristic (brd), and best-response dynamics initialized with Algorithm 1 (greedy+brd). (a) Average leader's cost as a function of the number of followers with 20 resources, 15 actions per player, and monotonic and almost monotonic (i.e., single-peak) cost functions. (b) Average leader's cost as a function of the number of followers, with 20 resources, 15 actions per player, and random generic cost functions.