



Strong price of anarchy[☆]

Nir Andelman^a, Michal Feldman^{b,*}, Yishay Mansour^a

^a *School of Computer Science, Tel Aviv University, Israel*

^b *School of Business Administration and Center for the Study of Rationality, Hebrew University of Jerusalem, Israel*

Received 1 August 2006

Available online 28 March 2008

Abstract

A strong equilibrium is a pure Nash equilibrium which is resilient to deviations by coalitions. We define the strong price of anarchy (SPoA) to be the ratio of the worst strong equilibrium to the social optimum. Differently from the Price of Anarchy (defined as the ratio of the worst Nash Equilibrium to the social optimum), it quantifies the loss incurred from the lack of a central designer in settings that allow for coordination.

We study the SPoA in two settings, namely job scheduling and network creation. In the job scheduling game we show that for unrelated machines the SPoA can be bounded as a function of the number of machines and the size of the coalition. For the network creation game we show that the SPoA is at most 2. In both cases we show that a strong equilibrium always exists, except for a well defined subset of network creation games.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C62; C72

Keywords: Strong equilibrium; Price of anarchy; Strong price of anarchy; Coalitions; Congestion games; Network formation; Job scheduling

[☆] This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a Grant No. 1079/04 from the Israel Science Foundation, by a grant from BSF and an IBM faculty award. The second author is also supported by the Lady Davis Fellowship Trust. This publication only reflects the authors' views.

* Corresponding author.

E-mail addresses: andelman@cs.tau.ac.il (N. Andelman), mfeldman@cs.huji.ac.il (M. Feldman), mansour@tau.ac.il (Y. Mansour).

¹ The research was partially done when the author was at the School of Computer Science at the Hebrew University of Jerusalem.

1. Introduction

1.1. Problem statement

The theory of games is used to study situations that involve rational (selfish) agents who are motivated by optimizing their own utilities rather than reaching some social optimum. If there is no central authority that can enforce the desired behavior on the individual agents, the social optimum will typically not be reached, and some social loss will be incurred.

Quantifying the efficiency loss due to selfish behavior is a natural concern in such settings. Koutsoupias and Papadimitriou (1999) proposed to analyze this inefficiency from a worst-case perspective, quantifying the loss as the ratio between the worst possible solution reached with selfish agents and the social optimum. In order to formalize this metric, one needs to define what constitutes rational behavior, or what is an accepted *solution concept*. They chose Nash equilibrium, which is perhaps the most popular solution concept used in the literature. Papadimitriou (2001) later coined the term *Price of Anarchy* (PoA) to denote this ratio between the worst Nash equilibrium and the social optimum. This metric has been extensively studied in the contexts of selfish routing (Roughgarden and Tardos, 2002), job scheduling (Koutsoupias and Papadimitriou, 1999; Czumaj and Vöcking, 2002), network formation (Fabrikant et al., 2003; Albers et al., 2006), facility location (Vetta, 2002) and general congestion games (Christodoulou and Koutsoupias, 2005).²

In a Nash equilibrium (NE) no agent can improve its own utility by *unilaterally* changing its action. The price of anarchy, therefore, quantifies the efficiency loss that is incurred in settings where coordinated deviations cannot occur. Clearly, it is reasonable to assume that a state from which some agent can unilaterally improve its utility is not sustainable. However, when no unilateral deviations are profitable, it does not necessarily imply that the solution is sustainable, since a group of agents may *coordinate* a joint deviation if it is in the best interest of all of its members. In situations that allow for some level of coordination, the NE concept cannot be assumed to reflect rational behavior.

To address the issue of coordination, we adopt the solution concept of *strong equilibrium* (SE), proposed by Aumann (1959). In a strong equilibrium, no coalition (of any size) can deviate and improve the utility of *every* member of the coalition (while possibly lowering the utility of players outside the coalition). Clearly, every strong equilibrium is a Nash equilibrium, but the converse does not hold. In cases where a strong equilibrium exists, it seems to be a very robust notion. We define the *strong price of anarchy* (SPoA) to be the ratio of the worst strong equilibrium and the social optimum. The SPoA metric better suits situations that allow for some level of coordination among the participating agents, but it is well defined only when a strong equilibrium exists. Unfortunately, most games do not admit any strong equilibrium (even if mixed strategies are allowed).³ Thus, in order to analyze the SPoA, one must first prove that a strong equilibrium exists in the specific setting at hand.

Since the set of SE is contained in the set of NE, SPoA is always smaller or equal to the PoA. In cases where the SPoA yields substantially better results than the PoA, it suggests that coordination can significantly improve the efficiency loss. If, on the other hand, the SPoA and

² In (Christodoulou and Koutsoupias, 2005), the authors also study the correlated price of anarchy, which considers correlated equilibria rather than Nash equilibria.

³ This is in contrast to Nash equilibrium, which exists (perhaps in mixed strategies) in every finite game.

PoA yield similar results, the efficiency loss, as quantified by the PoA, cannot be improved by coordination. We will see examples of both cases in the games we study.

In our definition we also consider the size of the coalition as a parameter and define k -SPoA to be the ratio of the worst Nash equilibrium which is immune to coalitions of size up to k and the social optimum.⁴ This is a natural restriction in many settings, where the ability to coordinate may be limited.⁵

While the price of anarchy takes a worst-case approach, the *Price of Stability (PoS)* (Anshelevich et al., 2004) is defined as the ratio of the *best* Nash equilibrium to the social optimum. Nash equilibrium has a natural meaning of stability, since if any Nash equilibrium is proposed, no agent will have incentive to deviate from it unilaterally. Similarly, one can define the *Strong Price of Stability (SPoS)* as the ratio of the best strong equilibrium and the optimum. Since any SE is also a NE, the SPoS cannot be smaller than the PoS, but interestingly enough, it may be greater. As will be shown below, this cannot happen in the families of games studied in this paper, but such an example is presented in (Epstein et al., 2007) in a cost-sharing connection game.

1.2. Our games and results

1.2.1. Job scheduling

The first set of games is derived from job scheduling, where each player controls a single job and selects the machine on which the job is run. The cost to the player is the load on the machine it selected while the social cost is the *makespan*, i.e., the maximal load on any machine.

With unit-size jobs and identical machines, the job scheduling game is a congestion game (Rosenthal, 1973). In congestion games, each agent's strategy is a subset of resources, and the utility of an agent from each resource depends only on the number of agents choosing the same resource. The most prominent example of a congestion game is that of traffic routing, where the cost of a road signifies the delay experienced by traffic traversing that road, and is a function of the number of players choosing it.⁶ Holzman and Law-Yone (1997) have characterized the set of congestion games that admit strong equilibria.

In this work, however, we consider mostly the model of *unrelated machines*, which is not a congestion game per se. In unrelated machines the load of a job is a function of the machine it is scheduled on, and the load on each machine is the sum of the loads of the jobs scheduled on it. In computational settings, jobs may be, for example, CPU bounded or memory bounded, and thus can place different loads on different machines, depending on their characteristics. Similarly, in the context of traffic routing, different types of vehicles may place different loads on different roads depending on the vehicle's type (e.g., bicycle, motorcycle, car, truck) and the road's attributes (e.g., steepness, width, curvedness).

Even-Dar et al. (2003) have shown that job scheduling with unrelated machines are general ordinal potential games (Monderer and Shapley, 1996), and therefore always admit pure NE. Yet, it does not hold in general that any potential game admits a strong equilibrium (even for exact potential games, see Section 3.2). We show that in any job scheduling game with unrelated machines, a strong equilibrium always exists. Moreover, some optimal solution is also a strong

⁴ Namely, no coalition of at most k players can coordinate a deviation and improve the utility of every player in the coalition.

⁵ A similar restriction on the coalition size was considered in (Abraham et al., 2006) in the context of a stronger solution concept.

⁶ Job scheduling is equivalent to traffic routing on parallel routes.

equilibrium, and thus the strong price of stability is 1. As for the strong price of anarchy, while it is rather simple to show that for unrelated machines the PoA is unbounded (Awerbuch et al., 2006), we show that the SPoA is bounded as a function of the number of players and machines. More specifically, we show that⁷:

- For m machines the worst-case SPoA is at most $2m - 1$ and at least m (and for 2 machines it is 2).
- For m machines and n players the worst-case k -SPoA is at most $O(nm^2/k)$ and at least $\Omega(n/k)$.

The vast literature on strong equilibrium has focused on pure strategies and pure deviations, e.g., (Holzman and Law-Yone, 1997, 2003; Milchtaich, 1998; Bernheim et al., 1987). This has been mainly motivated by the fact that the strong equilibrium is already a solution concept that does not exist in many cases and allowing mixed deviations would only further reduce it. The only exception is Rozenfeld and Tennenholtz (2006) where correlated deviations are considered. We show that in the job scheduling setting, once we allow mixed deviations by coalitions, in many cases no strong equilibrium exists (in contrast to pure deviations, where always some strong equilibrium exists). More specifically, in the case of mixed strategies and mixed deviations, for $m \geq 5$ identical machines and $n > 3m$ identical jobs, there is no strong mixed equilibrium with respect to mixed deviations.

1.2.2. Network creation

The second game we consider is a network creation game. Network structures play an important role in a wide and varied set of economic situations, such as job opportunity dissemination,⁸ non-centralized market trade (Kranton and Minehart, 2001), and many others. Due to the large variety of situations in which network creation plays a role, there is no single best model that captures all of its aspects. Instead, many models have been proposed in the literature to describe network creation and analyze their stability and efficiency. We refer the reader to Jackson (2004) for a comprehensive survey on the topic. The common assumption in most of these models is the discretion of each player in forming his links.

We adopt the model proposed by Fabrikant et al. (2003) and later studied also by Albers et al. (2006). In this game, the players can be viewed as nodes in a graph. Each node buys links to other nodes at the cost of α per link. The set of edges in the resulting graph is the union of the links that the nodes bought.⁹ Only the node that bought the link pays for it, but once a link is bought, it can be used in both directions. The cost to each node is the cost of the links it bought plus the sum of the shortest distances to all the nodes in the resulting graph. The social cost is the sum of the players' (nodes') costs.

We show that for most values of α there is some strong equilibrium. Specifically, for $\alpha \in (0, 1]$ we show that the clique is a strong equilibrium and for $\alpha \geq 2$ a star is a strong equilibrium. For $\alpha \in (1, 2)$ we show that there is no strong equilibrium in general. More specifically, we show that

⁷ In what follows, as well as in the remainder of this paper, we use the Big-O and Big-Omega notations, which are used to describe the asymptotic behavior of functions. We denote $g(n) = O(f(n))$ if $\exists c$ such that $g(n) \leq c \cdot f(n) \forall n \geq 1$. Similarly, we denote $g(n) = \Omega(f(n))$ if $\exists c$ such that $g(n) \geq c \cdot f(n) \forall n \geq 1$.

⁸ Research in sociology and labor economics has indicated that social interactions are the leading source of job opportunities (Montgomery, 1991).

⁹ In this model, the existence of an edge between two nodes does not require the consent of both players.

there is no strong equilibrium when the coalition size is at least 3 and the number of players is at least 6. However, for either a smaller number of players (four or less) or smaller coalitions (size at most 2) there always exists a strong equilibrium.

Previous work has already provided an upper bound for the PoA of the network creation game (Fabrikant et al., 2003; Albers et al., 2006). Roughly, for $\alpha = O(\sqrt{n})$ and $\alpha = \Omega(n \log n)$ the PoA is constant. For $\alpha \in [\sqrt{n}, n]$ the PoA is $O(\alpha^{2/3}/n^{1/3})$ and for $\alpha \in [n, n \log n]$ the PoA is $O(n^{2/3}/\alpha^{1/3})$. We show that for any $\alpha \geq 2$ the SPoA is at most 2. This should be contrasted with the lower bound of 3 for the PoA (Fabrikant et al., 2003). In addition, our existence results show that the SPoS is 1, since there always exists an optimal network which is a strong equilibrium. A related notion to strong equilibrium in network creation games was studied in Dutta and Mutuswami (1997), which considered the notion of *strong stability* in a related, yet different, network creation game. Strong stability is stronger than SE considered here, since it does not require that all the deviating agents *strictly* improve their utilities. The main difference between their model and ours is that they consider non-directed networks, in which the existence a link between two players requires the consent of both players. They concentrate on the question of whether or not efficiency and stability can be reconciled (i.e., whether optimal networks can be formed by rational players), but do not attempt to quantify the efficiency loss that may arise in cases where inefficient networks are formed. Requiring the consent of both players to create a link between them was also considered in Corbo and Parkes (2005), who showed that the price of anarchy is greater than the price obtained in the unilateral case considered by Fabrikant et al. (2003).

1.3. Coalitions in games

The strong equilibrium solution concept is by no means the only solution concept that considers the formation of coalitions. Indeed, group incentives and coalition formation lie at the heart of *cooperative game theory*, and many solution concepts have been studied in this framework. The *core* is a central notion in cooperative game theory. Roughly speaking, it is an outcome in which all the agents form the grand coalition, and no coalition can deviate and improve the utility of each one of its members. In contrast to cooperative game theory, which focuses on coalitional moves, non-cooperative game theory concentrates on the strategic considerations of individual agents. Yet, many non-cooperative models have considered the issue of coalition formation. Some of these models assume transferable utilities while others assume non-transferable utilities. Some assume a static structure of the coalitions, while others assume a dynamic structure. Our work is within the framework of non-cooperative game theory, and assumes that coalitions are dynamically formed with non-transferable utilities.

The SE solution concept was refined by Bernheim et al. (1987) to the Coalition-Proof Nash Equilibrium (CPNE) solution concept, which appears to be well established. SE requires that there will be no beneficial joint deviation from a strategy profile by a coalition of any size. This is a very strong notion of stability, but may eliminate many profiles. In contrast, the CPNE solution concept has a weaker notion of stability. It considers only those deviations that are themselves resilient to further deviations by subsets of the original coalition. This implies that every SE is also a CPNE but not vice versa. Extending our analysis and results to the CPNE solution concept is an interesting future direction. In this work, we highlight one example where our analysis carry over to the CPNE case (see Footnote 13) and another example where our analysis does not extend to the case of CPNE (see Footnote 17 in the proof of Lemma B.1).

Recently, several works studied the effect of coalition formation in a somewhat different model (Hayrapetyan et al., 2006; Fotakis et al., 2006; Kuniavsky and Smorodinsky, 2007). Still within the framework of non-cooperative game theory, they considered the case in which a fixed set of coalitions is given (which is an exogenous partition over the set of agents), and the members of each coalition play as so to maximize their collective welfare (implying transferable utilities among the members of any given coalition but not between coalitions). The underlying congestion game coupled up with the coalition structure is not a congestion game any more, thus does not necessarily possess a pure Nash equilibrium. The above works study the existence of NE in these games, and also define the *price of collusion*, which specifies the factor by which the solution quality can deteriorate in the presence of coalitions. Interestingly enough, it is shown that the price of collusion can be arbitrarily high.

Coalitions have been also considered from a mechanism design perspective. A *group-strategyproof* mechanism is a mechanism that induces truth telling by individual agents and by coalitions of agents (i.e., where no coalition of agents can improve its utility by deviating from truth telling). Group-strategyproof mechanisms were studied, for example, in cost-sharing games (Moulin and Shenker, 2001) and in auctions (Goldberg and Hartline, 2005).

2. Model

In this section we provide general notations and definitions, while in Sections 3.1 and 4.1 we provide the notations and definitions for the specific games we study.

A game is denoted by a tuple $G = \langle N, (S_i), (c_i) \rangle$, where N is the set of players, S_i is the finite action space of player $i \in N$, and c_i is the cost function of player i .

We denote by $n = |N|$ the number of players. The joint action space of the players is $S = \times_{i=1}^n S_i$. For a joint action $s \in S$ we denote by s_{-i} the actions of players $j \neq i$, i.e., $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. Similarly, for a set of players Γ we denote by $s_{-\Gamma}$ the actions of players $j \notin \Gamma$. The cost function of player i maps a joint action $s \in S$ to a real number, i.e., $c_i : S \rightarrow \mathbb{R}$.

Nash Equilibrium (NE). A joint action $s \in S$ is a *pure Nash Equilibrium* if no player $i \in N$ can benefit from unilaterally deviating from his action to another action, i.e., $\forall i \in N \forall a \in S_i : c_i(s_{-i}, a) \geq c_i(s)$.

Resilience to coalitions. A *pure joint action* of a set of players $\Gamma \subset N$ (also called *coalition*) specifies an action for each player in the coalition, i.e., $\gamma \in \times_{i \in \Gamma} S_i$. A joint action $s \in S$ is not resilient to a *pure deviation* of a coalition Γ if there is a pure joint action γ of Γ such that $c_i(s_{-\Gamma}, \gamma) < c_i(s)$ for every $i \in \Gamma$ (i.e., the players in the coalition can deviate in such a way that *each* player reduces its cost). A pure Nash equilibrium $s \in S$ is *resilient to pure deviation of coalitions of size k* , if there is no coalition Γ of size at most k , such that s is not resilient to a pure deviation by Γ .¹⁰

Definition 2.1. A *k -strong equilibrium (k -SE)* is a pure Nash equilibrium that is resilient to pure deviation of coalitions of size at most k .

¹⁰ We define mixed Nash equilibrium and mixed deviations in Section 3.4.




Fig. 1. Illustration of the k -SE hierarchy (the set k -SE represents all the NE which are also k -SE).

Clearly, a k -SE is a refinement of NE. Let $\Phi(G, k)$ be the set of k -strong equilibria of the game G . By definition, for any k , $\Phi(G, k) \subseteq \Phi(G, k - 1)$ (see Fig. 1). Note that $\Phi(G, 1)$ coincides with the set of NE, and $\Phi(G, n)$ coincides with the classical notion of a *strong equilibrium* introduced by Aumann (1959).

Note that while in Nash equilibria we can restrict attention to pure deviations, this is not true for k -strong equilibrium, when $k \geq 2$. The conceptual reason is that we need to guarantee that *each* player in the coalition would benefit from the deviation. In Section 3.4 we show an example in which a coalition can benefit from a mixed deviation, yet in any pure deviation some player in the coalition does not benefit. (We defer the definition of a mixed deviation to the above section.)

In order to study the strong price of anarchy we need to define the *social cost* of a game G . Abstractly, there is a function f_G such that the social cost of $s \in S$ is $f_G(s)$. The optimal social cost is $OPT(G) = \min_{s \in S} f_G(s)$. In the cases discussed in this paper the social cost is a simple function of the costs of the players. More specifically, we discuss the linear case, i.e., $f_G(s) = \sum_{i=1}^n c_i(s)$, and the maximum, i.e., $f_G(s) = \max_{i=1}^n c_i(s)$. Next we define the strong price of anarchy (SPoA).

Definition 2.2. Let $\Phi(G, k)$ be the set of k -strong equilibria of the game G . If $\Phi(G, k) \neq \emptyset$ then the *k -strong price-of-anarchy* (k -SPoA) is the ratio between the *maximal* cost of a k -strong equilibrium and the social optimum, i.e., $\max_{s \in \Phi(G, k)} f_G(s) / OPT(G)$.

Similarly, we define the strong price of stability (SPoS).

Definition 2.3. Let $\Phi(G, k)$ be the set of k -strong equilibria of the game G . If $\Phi(G, k) \neq \emptyset$ then the *k -strong price-of-stability* (k -SPoS) is the ratio between the *minimal* cost of a k -strong equilibrium and the social optimum, i.e., $\min_{s \in \Phi(G, k)} f_G(s) / OPT(G)$.

We denote by SPoA the n -SPoA, and by SPoS the n -SPoS, allowing any size of a coalition. (Note that both SPoA and SPoS are defined only if some strong equilibrium exists.)

3. Job scheduling

In our job scheduling scenario there are m machines and n players (where each player controls a single job). In the job scheduling terminology, we will focus on unrelated machines, but also refer to identical machines. The missing proofs of this section appear in Appendix A.

3.1. Job scheduling model

A job scheduling setting is characterized by the tuple $\langle M, N, (w_i(J)) \rangle$, where $M = \{M_1, \dots, M_m\}$ is the set of machines, $N = \{1, \dots, n\}$ is the set of players (jobs) and $w_i(J) \in \mathbb{R}$ is the weight of player $J \in N$ on machine $M_i \in M$. A job scheduling setting has identical machines if for every $M_i, M_{i'} \in M$ and $J \in N$, we have $w_i(J) = w_{i'}(J)$. In identical machine settings we will use $w(J)$ to denote the weight of J (on any machine). If, in addition, the setting has unit jobs, then $w(J) = 1$ for every $J \in N$.

A job scheduling game has N as the set of players. The action space S_J of player $J \in N$ are all the individual machines, i.e., $S_J = M$. The joint action space is $S = \times_{J=1}^n S_J$. In a joint action $s \in S$ player J selects machine s_J as its action. We denote by B_i^s the set of players on machine M_i in the joint action $s \in S$, i.e., $B_i^s = \{J: s_J = M_i\}$. The load of a machine M_i , in the joint action $s \in S$, is the sum of the weights of the players that chose machine M_i , that is $L_i(s) = \sum_{J \in B_i^s} w_i(J)$. For a player $J \in N$, let $c_J(s)$ be the load that player J observes in the joint action s , i.e., $c_J(s) = L_i(s)$, where $s_J = M_i$. A job scheduling game is characterized by a tuple $\langle N, S, (c_J) \rangle$.

In job scheduling games the objective function (i.e., the social cost) is the *makespan*, which is the load on the most loaded machines (or equivalently, the highest load some player observes). Formally, $\text{makespan}(s) = \max_J c_J(s)$. A social optimum minimizes the makespan, i.e., $OPT = \min_s \text{makespan}(s)$. Thus, the strong price of anarchy (SPoA) in job scheduling games is the ratio between the makespan of the worst SE and the minimal makespan.

Notation: We define $w_{\min}(J) = \min_i w_i(J)$. We denote by $\min(J)$ the index of a machine on which player J has weight $w_{\min}(J)$, i.e., $\min(J) = \arg \min_i w_i(J)$ (if there is more than one such machine then select an arbitrary one). In addition, we denote by $OPT(J)$ the action of job J under a social optimum OPT .

3.2. Equilibrium existence

In this section we prove that in the job scheduling game, for any coalitions of size k , there is a k -SE, i.e., there exists a NE that is resilient to coalitions of size k (for any $k \leq n$). Job scheduling games with unrelated machines are not congestion games per se (Rosenthal, 1973). Therefore, even the existence of pure NE in these games requires a proof. Even-Dar et al. (2003) proved that by showing that any sequence of improvement steps, in a job scheduling game, converges to a NE. Our proof technique is similar. We first define a complete order on the joint actions.

Definition 3.1. A vector (l_1, l_2, \dots, l_m) is smaller than $(\hat{l}_1, \hat{l}_2, \dots, \hat{l}_m)$ lexicographically if for some i , $l_i < \hat{l}_i$ and $l_k = \hat{l}_k$ for all $k < i$. A joint action s is smaller than s' lexicographically if the vector of machine loads $L(s) = (L_1(s), \dots, L_m(s))$, sorted in non-increasing order, is smaller lexicographically than $L(s')$, sorted in non-increasing order. We denote this relationship by $s < s'$.

The following lemma would be helpful in establishing the lexicographic order of two joint actions.

Lemma 3.2. Consider two joint actions s and s' such that the load vectors $L(s)$ and $L(s')$ differ only in the loads of machines in a set $M' \subseteq M$. If for each $M_i \in M'$, $L_i(s) < \max_k \{L_k(s') \mid M_k \in M'\}$ then $s < s'$.

We now prove that the lexicographically minimal assignment is a k -SE.

Theorem 3.3. *In any job scheduling game, the lexicographically minimal joint action s is a k -SE equilibrium, for any k .*

Proof. Lemma A.1 in the Appendix shows that s is a NE. To show that s is a k -SE, assume by contradiction that there is a coalition Γ of size $k \leq n$ that can deviate such that each member of the coalition strictly decreases its observed load. Let Γ be such a coalition of the smallest size. Let the resulting joint action after the deviation be s' . Let $M(\Gamma, s) = \bigcup_{J \in \Gamma} \{s_J\}$ be the set of machines that the coalition Γ chooses in the joint action s .

We first note that if there is a job $J \in \Gamma$ that does not migrate, i.e. $s_J = s'_J$, then the set of jobs $\Gamma \setminus \{J\}$ also forms a coalition, contradicting the minimality of Γ . Therefore, for every job $J \in \Gamma$ we have $s_J \neq s'_J$.

We show that for every machine in $M_i \in M(\Gamma, s)$ that a job $J \in \Gamma$ wishes to leave there is a job $J' \in \Gamma$ that wishes to migrate to that machine, and vice versa, which implies that $M(\Gamma, s) = M(\Gamma, s')$. We have to consider two cases. In the first case there is a machine that some job $J \in \Gamma$ migrates to, but no job $J' \in \Gamma$ migrates from. Such a case would contradict the fact that s is a NE (established in Lemma A.1). The second case is that there is a machine that some job $J \in \Gamma$ migrates from, but no job $J' \in \Gamma$ migrates to. Such a case would contradict the minimality of Γ , since the set of jobs $\Gamma \setminus \{J\}$ forms a coalition as well.

We now have that only machines in $M' = M(\Gamma, s) = M(\Gamma, s')$ change their loads. For each machine in $M_i \in M'$ there is at least one job $J \in \Gamma$ that wishes to migrate to it, i.e., $s'_J = M_i$. Since each job $J \in \Gamma$ benefits from the coalition deviation, the new load on each machine $M_i \in M'$ must be strictly lower than $L_{s_J}(s)$, and therefore strictly lower than $\max_k \{L_k(s) \mid M_k \in M'\}$. By Lemma 3.2, this implies that $s' \prec s$, contradicting the minimality of s . \square

Since a lexicographically minimal joint strategy is a k -SE for any k , the following is an immediate corollary.

Corollary 3.4. *In any job scheduling game, for any k , the k -Strong Price of Stability (k -SPoS) is 1.*

It was shown in Even-Dar et al. (2003) that any job scheduling game is a general ordinal potential game (Monderer and Shapley, 1996). However, while Theorem 3.3 holds for any job scheduling game, it does not hold in general for any potential game (and not even for any exact potential game). For example, the prisoner’s dilemma is an exact potential game (Monderer and Shapley, 1996), but the only NE in this game (in which both players defect) is not resilient to a coalition of both players cooperating. Thus, the prisoner’s dilemma game has no SE.

The requirement that every member in a coalition strictly benefits from the deviation is a crucial assumption for the correctness of Theorem 3.3. If we relax the condition and require only that some member improves its cost and no other member of the coalition would lose from the deviation, there are job scheduling games that do not admit any SE.¹¹

¹¹ For example, consider the following setting: there are two identical machines, and three identical unit jobs. Clearly, in a NE, a pair of jobs is on one machine and the third job is on the other. However, under the relaxed improvement requirement, no equilibrium is 2-SE: The pair of jobs on the same machine can form a coalition where one job migrates




Fig. 2. An example of an assignment (a) that is a Nash equilibrium but not a strong equilibrium, since the jobs of size $\{5, 5, 3, 3\}$ all profit from the deviation demonstrated in (b).

While the proof technique for the existence of k -SE is similar to Even-Dar et al. (2003), it is important to note that there is no equivalence between deviation in coalitions and unilateral deviations. In particular, already with 3 identical machines,¹² there exist NE assignments that are not resilient to coalitional deviations, as the following example (illustrated in Fig. 2) demonstrates.¹³

Example 3.5. Consider $m = 3$ identical machines and 6 jobs with weights $w(J_1) = w(J_2) = 5$, $w(J_3) = w(J_4) = 4$ and $w(J_5) = w(J_6) = 3$, and consider the assignment s in which $s_{J_1} = s_{J_2} = M_1$, $s_{J_3} = s_{J_5} = M_2$ and $s_{J_4} = s_{J_6} = M_3$. One can verify that s is a NE. However, consider the coalition $\Gamma = \{J_1, J_2, J_5, J_6\}$ such that $s'_{J_1} = M_2$, $s'_{J_2} = M_3$ and $s'_{J_5} = s'_{J_6} = M_1$. Their costs in s are $c_{J_1}(s) = c_{J_2}(s) = 10$, $c_{J_5}(s) = c_{J_6}(s) = 7$, while their costs in s' are $c_{J_1}(s') = c_{J_2}(s') = 9$, $c_{J_5}(s') = c_{J_6}(s') = 6$. Thus, the deviation is beneficial to all of the coalition’s members.

3.3. Strong price of anarchy

In this section we study the SPoA in scenarios with identical and unrelated machines. For identical machines, it is known that $\text{PoA} \leq 2$ (Koutsoupias and Papadimitriou, 1999), while for unrelated machines, the PoA may be unbounded (Awerbuch et al., 2006). Consider the following motivating example for unrelated machines.

Example 3.6. Consider $m \geq 2$ machines and $n = m$ jobs, where $w_i(J_i) = \epsilon$ for all $1 \leq i \leq m$, and $w_i(J_j) = 1$ for all $i \neq j$. The joint action $(1, 2, \dots, m)$ has a minimal makespan of ϵ (and is also a NE). However, the joint action $(m, 1, 2, \dots, m - 1)$ is also a NE and has a makespan of 1. Therefore, the PoA is at least $1/\epsilon$, which can be arbitrarily large. However the only joint action that is resilient to a coalition of all the players is $(1, 2, \dots, m)$, and therefore in this example the SPoA is 1, which is significantly smaller than the PoA.

Example 3.6 motivates using the SPoA solution concept for unrelated machines. We now prove our main results for the job scheduling games, showing that the strong price of anarchy is

to the other machine, while the other job does not change machines. After the deviation, the migrating job remains with a load of 2, while the load observed by the idle job in the coalition decreases from 2 to 1.

¹² In a job scheduling game with 2 identical machines, one can verify that an assignment is NE if and only if it is SE.

¹³ The example above shows that a NE is not necessarily a SE, but it actually demonstrates a stronger result—that a NE is not necessarily a coalition-proof Nash equilibrium (see Section 1.3 for a definition), because the post-deviation assignment (in Fig. 2(b)) is robust against further deviations.

bounded in the unrelated machine setting. This shows that the unbounded PoA is not originated from selfishness, but rather from lack of coordination. Once coordination is allowed, even with selfish agents, the efficiency loss can be bounded.

We start with the following straightforward relationship between OPT and the weights.

Claim 3.7. *For any job scheduling game with unrelated machines, the following inequalities hold:*

$$OPT \geq \max_J w_{\min}(J), \tag{1}$$

$$OPT \geq \frac{1}{m} \sum_J w_{\min}(J) \tag{2}$$

where $OPT = \min_{s \in S} \max_i L_i(s)$. \square

We first bound the SPoA for games with two machines.

Theorem 3.8. *For any job scheduling game with 2 unrelated machines and n jobs, $SPoA \leq 2$.*

We next introduce some notations that will be useful. For simplicity, for the rest of this section we will assume WLOG that given a joint action s , the machine indices are sorted in a non-decreasing order of the loads under s , i.e., $L_1(s) \leq \dots \leq L_m(s)$.

Definition 3.9. We denote it by $M_i \mapsto_s M_j$, if there is a job J such that $M_j = \min(J)$, $s_J = M_i$ and $L_i(s) \geq L_j(s)$. Two machines M_i and M_j , such that $L_i(s) \geq L_j(s)$ are *connected under the joint action s* if $\exists i', j'$, such that $L_{i'}(s) \geq L_i(s)$, $L_j(s) \geq L_{j'}(s)$, and $M_{i'} \mapsto_s M_{j'}$. Let $C(s) = \{M_m, \dots, M_\ell\}$ denote the maximal suffix of machines, such that M_{i+1} is connected to M_i under joint action s .

To provide an intuition for the definition above, consider the assignment depicted in Fig. 3. Machine M_m is connected to machine M_{m-1} since there is some job J on M_m such that $\min(J) = m - 1$ (demonstrated by the arrow going from M_m to M_{m-1}). Similarly, machine M_{m-1} is connected to machine M_{m-2} since there is some job J on M_{m-1} such that $\min(J) = m - 2$. In a similar way, we show that M_{m-2} is connected to M_{m-3} , and M_{m-3} is connected to M_{m-4} . However, machine M_{m-4} is not connected to machine M_{m-5} since there is no “arrow” that goes from M_{m-4} or a machine to its left to M_{m-5} or a machine to its right. Therefore, $C(s) = \{M_m, M_{m-1}, M_{m-2}, M_{m-3}, M_{m-4}\}$.

By the definition of $C(s)$ and the relation $M_i \mapsto_s M_j$ we have,




Fig. 3. Illustration of $C(s)$.

Claim 3.10. *For every job J such that $s_J \in C(s)$ we have $\min(J) \in C(s)$.*

The next lemma bounds the difference between loads of machines in $C(s)$, under a NE s .

Lemma 3.11. *Let s be a NE. If $M_i \mapsto_s M_j$ then $L_i(s) \leq L_j(s) + OPT$. In addition, for any $i, j \in C(s)$ we have $L_i(s) \leq L_j(s) + (m - 1)OPT$*

Proof. Since s is a NE, for each $J \in B_i^s$ we have $L_i(s) \leq L_j(s) + w_j(J)$. From the definition of $M_i \mapsto_s M_j$, there exists $J \in B_i^s$ for which $M_j = \min(J)$. From inequality (1), $w_j(J) \leq OPT$, and we get: $L_i(s) \leq L_j(s) + OPT$.

By consecutive applications of this argument, the load of M_m and M_ℓ , the least loaded machine in C_m , cannot differ by more than $(m - 1)OPT$. Therefore, for any two machines M_i and M_j in C_m , $L_i(s) \leq L_j(s) + (m - 1)OPT$. \square

Theorem 3.12. *For any job scheduling game with m unrelated machines and n jobs, $SPoA \leq 2m - 1$.*

Proof. Let s be an arbitrary joint action that is a SE. Recall that we assume WLOG that the machines are sorted in a non-decreasing order of the loads.

If for some $M_i \in C(s)$ we have $L_i(s) \leq m \cdot OPT$ then by Lemma 3.11 $L_m(s) \leq (2m - 1) \cdot OPT$, and we are done. Otherwise, $\forall i \in C(s)$, $L_i(s) > m \cdot OPT$. We will show that such a joint action s is not resilient to a deviation of a coalition. Consider the joint action s' , where for J such that $s_J \in C(s)$ we have $s'_J = \min(J)$, and for J such that $s_J \notin C(s)$ we have $s'_J = s_J$. This implies that the coalition Γ includes all the jobs scheduled in s on machines in $C(s)$, i.e., $\Gamma = \bigcup_{M_i \in C(s)} B_i^s$.

Recall that by Claim 3.10 we have $\min(J) \in C(s)$. By inequality (2), $L_i(s') \leq m \cdot OPT < L_i(s)$, for any $M_i \in C(s)$. Therefore, each job $J \in C(s)$ is strictly better off under s' . \square

The following theorem shows that the SPoA might be linear in the number of machines m .

Theorem 3.13. *For every $m \geq 2$, there exists a job scheduling game with m unrelated machines for which $SPoA \geq m$.*

The proof appears in the Appendix, and is illustrated for $m = 4$ in Fig. 4.

Next, we derive bounds for coalitions whose size is smaller than n . We first present a lower bound for two machines. This lower bound shows that differently from the case of unrestricted

	M_1	M_2	M_3	M_4
J_1	①	∞	∞	①
J_2	①	②	∞	∞
J_3	∞	①	③	∞
J_4	∞	∞	①	④

Fig. 4. This example illustrates a job scheduling game in which $SPoA = m$ for $m = 4$. The assignment s such that $s_{J_i} = M_i$ (marked by the thick circles) is a SE with $c(s) = 4$. Yet, the optimal assignment is $OPT_{J_1} = M_m$ and $OPT_{J_i} = M_{i-1}$ for any $i \in \{2, \dots, m\}$ (marked by the thin circles), with $c(OPT) = 1$. Thus, $SPoA = 4$.

coalition sizes, where the SPoA is linear in m (see Theorem 3.12), if the coalition size is restricted by some $k < n$, the SPoA can be on the order of $\frac{n}{k}$. This is especially unfortunate in cases where the number of jobs is significantly greater than the number of machines.

Theorem 3.14. *For every $n, k \geq 2, k < n$, there exists a job scheduling game with 2 unrelated machines and n jobs, s.t. k -SPoA $\geq \frac{n}{2k}$.*

Moreover, for $k < m$, the k -SPoA can be unbounded as demonstrated in Example 3.6. The configuration presented in the example is resilient to any coalition of size at most $m - 1$. Therefore, the k -SPoA is unbounded for $k \leq m - 1$.

For $k \geq m$, we already know by Theorem 3.13 that m -SPoA $\geq m$. On the positive side, the following theorem shows that the k -SPoA can be bounded from above¹⁴ for any $m \leq k < n$.

Theorem 3.15. *For any job scheduling game with m unrelated machines and n jobs, for any $k \geq m, k$ -SPoA $\leq \frac{2nm}{z} + 3m$, where $z = \lfloor k/m \rfloor$.*

Proof. We first present and prove the following lemma:

Lemma 3.16. *Fix a joint action s and a machine M_i . If for every subset $\Gamma \subseteq B_i^s, |\Gamma| \leq z$, the following inequality holds:*

$$\sum_{J \in \Gamma} w_i(J) - \sum_{J \in \Gamma} w_{\min}(J) \leq \beta,$$

then $L_i(s) \leq m \cdot OPT + \lceil \frac{n}{z} \rceil \beta$.

Proof. Let $\Gamma_1 = \{J \in B_i^s \mid M_i = \min(J)\}$ and $\Gamma_2 = \{J \in B_i^s \mid M_i \neq \min(J)\}$. Partition Γ_2 into $\ell = \lceil \frac{|\Gamma_2|}{z} \rceil$ subsets $\Gamma_{2,l}$ of size at most z . By the assumption in the lemma, for every subset $\Gamma_{2,l}$:

$$\sum_{J \in \Gamma_{2,l}} w_i(J) \leq \sum_{J \in \Gamma_{2,l}} w_{\min}(J) + \beta.$$

Summing over all l , we get:

$$\sum_{l=1}^{\ell} \sum_{J \in \Gamma_{2,l}} w_i(J) \leq \sum_{l=1}^{\ell} \sum_{J \in \Gamma_{2,l}} w_{\min}(J) + \ell \beta.$$

Therefore,

$$L_i(s) = \sum_{J \in B_i^s} w_i(J) = \sum_{J \in \Gamma_1} w_i(J) + \sum_{J \in \Gamma_2} w_i(J) \leq \sum_{J \in \Gamma_1} w_i(J) + \sum_{l=1}^{\ell} \sum_{J \in \Gamma_{2,l}} w_{\min}(J) + \ell \beta.$$

However, by inequality (2), $\sum_{J \in \Gamma_1} w_i(J) + \sum_l \sum_{J \in \Gamma_{2,l}} w_{\min}(J) \leq m \cdot OPT$.

Therefore, $\sum_{J \in B_i^s} w_i(J) \leq m \cdot OPT + \ell \beta$. \square

¹⁴ A tight bound for the k -SPoA is given in Fiat et al. (2007), where it is shown to be $\theta(m(n - m + 1)/(k - m + 1))$. (It is said that $g(n) = \theta(f(n))$ if $\exists c_1, c_2$ such that $c_1 \cdot f(n) \leq g(n) \leq c_2 \cdot f(n) \forall n$.)

We continue with the proof of the theorem. Consider the set $C(s)$. If there exists a machine $M_i \in C(s)$ such that for every subset $\Gamma \subset B_i^s, |\Gamma| \leq z$,

$$\sum_{J \in \Gamma} w_i(J) - \sum_{J \in \Gamma} w_{min}(J) \leq (2m - 1)OPT,$$

then, by Lemma 3.16, $L_i(s) \leq m \cdot OPT + (\frac{n}{z} + 1)(2m - 1)OPT$, and by Lemma 3.11, $L_m(s) \leq L_i(s) + (m - 1)OPT$, and we get

$$L_m(s) \leq (4m - 2)OPT + \frac{n}{z}(2m - 1)OPT \leq \left(\frac{2nm}{z} + 3m\right)OPT. \tag{3}$$

Otherwise, for every machine $M_i \in C(s)$, there exists a subset $\Gamma_i \subset B_i^s$, where $|\Gamma_i| \leq z$, for which

$$\sum_{J \in \Gamma_i} w_i(J) - \sum_{J \in \Gamma_i} w_{min}(J) > (2m - 1)OPT. \tag{4}$$

We show that in this case the joint action s is not a k -SE. Consider the following joint action s' : for $J \notin \bigcup_i \Gamma_i, s_J = s'_J$, and for $J \in \bigcup_i \Gamma_i, s'_J = \min(J)$ (i.e., in joint action s' each job from the Γ_i sets chooses its minimal work machine). Let $L_i(s \setminus \Gamma_i)$ denote the load of machine i excluding the jobs in Γ_i . That is, $L_i(s \setminus \Gamma_i) = L_i(s) - \sum_{J \in \Gamma_i} w_i(J)$. By inequality (4), $L_i(s \setminus \Gamma_i) < L_i(s) - (2m - 1)OPT$, and by inequality (2), $L_i(s') \leq L_i(s \setminus \Gamma_i) + m \cdot OPT$. Thus, for every $M_i \in C(s)$, $L_i(s') < L_i(s) - (m - 1)OPT$. By Lemma 3.11, for every $M_i, M_j \in C(s)$, $L_i(s) \leq L_j(s) + (m - 1)OPT$, therefore, $L_i(s') < L_j(s)$. This implies that s is not resilient to deviation of the coalition $\Gamma = \bigcup_i \Gamma_i$, where $|\Gamma| \leq zm \leq k$. \square

The above result should be contrasted both with the observation that the PoA can be unbounded and with Theorem 3.12. It shows that coordination among the players can significantly improve the PoA, and that its effect becomes more significant as the size of the coalition increases. When the coalition size is unrestricted, $SPoA \leq 2m - 1$, regardless of the number of jobs. However, when the coalition size is restricted, the efficiency loss may increase. Intuitively, when more coordination is possible, the players may coordinate their actions more efficiently and reduce the efficiency loss incurred due to strategic considerations. This intuition, however, should be taken with some caution, as in some cases, perhaps counter-intuitively, a larger degree of coordination can decrease the efficiency level (see Epstein et al., 2007 and Hayrapetyan et al., 2006).

The theorems above show that with unrelated machines the SPoA can be significantly improved over the PoA. For identical machines, it is known that $PoA \leq 2$. We next show that the SPoA does not improve on the PoA.

Theorem 3.17. *For every $m \geq 2$, there exists a job scheduling game with m identical machines and n jobs, s.t. $SPoA \geq \frac{2}{1+\frac{1}{m}}$.*

Remark. In this section we discussed the quality of SE, as measured by the SPoA. One can ask a similar question regarding the quality of Coalition-Proof Nash Equilibrium (CPNE). Regarding the case of two unrelated machines, our proof for SPoA seems to carry over to CPNE, since for two machines every optimal assignment is both a SE and a CPNE. For the general case, of multiple machines, we believe that our proof methodology can be extended to handle CPNE. Technically, we need to show that given a CPNE s , the maximum load of machines in $C(s)$ is

not much higher than OPT (say, at most $(3m - 1)OPT$). This can be done by arguing that if the load in $C(s)$ is higher, then every machine in $C(s)$ has a load higher than $(2m - 1)OPT$, and in such a case there is a stable deviation for jobs in $C(s)$. One such stable deviation is a deviation to SE for the jobs in $C(s)$ while using only machines in $C(s)$ (clearly, jobs in $C(s)$, either before or after the deviation, would not be able to profit from a deviation to machines not in $C(s)$). Since we showed that a SE is at most $(2m - 1)OPT$, every job in $C(s)$ has gained by the deviation.

3.4. Mixed deviations and mixed equilibrium

A natural extension of the SE solution concept would be to consider mixed strategies and deviations. A mixed strategy is a distribution over the action space, and similarly, a mixed coalition deviation assigns a new mixed strategy to every player in the coalition.

If players are only allowed to deviate unilaterally (as in NE), it is known that it is sufficient to consider only pure deviations. In contrast to NE, a pure SE might not be preserved when mixed deviations are allowed.¹⁵ We will show that when mixed deviations are allowed, many job scheduling games do not have a SE.

We will use the notation $\pi_J(i)$ to denote the probability that player J chooses machine M_i and let the joint strategy be $\pi = (\pi_1, \dots, \pi_n)$. The following example shows a pure SE, in a job scheduling game, which is not preserved when mixed deviations are allowed:

Example 3.18. Consider 2 identical machines and 3 unit jobs, J_1, J_2 and J_3 . In any NE with pure strategies, two jobs are assigned to one machine, while the third is assigned to the other machine. Clearly, this is also a SE. WLOG, we assume J_1 and J_2 are assigned to M_1 , and J_3 to M_2 in s .

Consider a coalition Γ consisting of J_1 and J_2 , where the mixed deviations are $\pi_1 = \pi_2 = (\frac{3}{4}, \frac{1}{4})$. The original load on M_1 in s is 2. After the deviation, J_1 and J_2 observe an expected load of $1\frac{7}{8}$. Since both players improve their costs, there is no pure NE that is a 2-SE.

Although Example 3.18 shows that the existence of pure SE cannot be guaranteed when mixed deviations are allowed, in the above example there is a mixed SE.¹⁶ However, in many cases allowing mixed deviations by a coalition eliminates all SE. The following theorem proves that this occurs even for identical machines and unit size jobs.

Theorem 3.19. For $m \geq 5$ identical machines and $n > 3m$ unit jobs, there is no 4-SE when mixed deviations are allowed.

Theorem 3.19 required that the coalitions would be of size 4, in order to demonstrate deviations with unit size jobs. The following theorem shows that with weighted jobs, there are settings where even coalitions of size as small as 2 eliminate all SE.

Theorem 3.20. There exists a job scheduling game with 2 identical machines and 3 jobs, where no joint mixed strategy is a 2-SE, when mixed deviations are allowed.

¹⁵ Rozenfeld and Tennenholtz (2006) consider an even stronger solution concept of correlated equilibria, and have shown that in a congestion game, it is possible that there is no strong correlated equilibrium in mixed strategies.

¹⁶ The SE has $\pi_1 = (1, 0)$, $\pi_2 = (0, 1)$ and $\pi_3 = (1/2, 1/2)$.

4. Network creation

In this section we study a network creation game which was introduced by Fabrikant et al. (2003). The game models the tradeoff of the agents (nodes) between buying links (edges) and reducing the distances to other nodes. In this section we discuss both the existence of a SE and the SPoA. The missing proofs of this section appear in Appendix B.

4.1. Network creation model

In the network creation game, there are n players, each of which is associated with a separate network vertex. The players buy edges to other nodes and the resulting network is an undirected graph. The cost of each player consists of two components. First, a player pays a cost of $\alpha > 0$ per edge it buys. Second, a player incurs a distance cost equal to the sum of the distances to the other nodes.

Formally, we represent the set of players by a vertex set $V = \{1, \dots, n\}$. For a player $v \in V$, an action $s_v \in S_v$ is a subset of the edges that include v , i.e., $s_v \subset \{(v, u) \mid u \in V \setminus \{v\}\}$. The action set of player v is S_v , which is the union of all the possible actions s_v .

Given a joint action $s = (s_1, \dots, s_n)$, let the resulting graph $G(s) = (V, E)$ consists of the edge set $E = \bigcup_{v \in V} s_v$. Let $\delta_s(v, w)$ be the length of the shortest path between v and w in $G(s)$.

The cost for a player v under joint action s is $c_v(s)$, and is composed from two parts. The *buying cost* is $B_s(v) = \alpha |s_v|$, which charges α for each edge v buys. The *distance cost* is $Dist_s(v) = \sum_{w \in V} \delta_s(v, w)$. The cost for player v is $c_v(s) = B_s(v) + Dist_s(v)$. When clear from the context we will omit the subscript s and use $\delta(v, w)$, $B(v)$, and $Dist(v)$ rather than $\delta_s(v, w)$, $B_s(v)$, and $Dist_s(v)$, respectively.

For a joint action $s \in S$, let the social cost be the total cost of all players, i.e., $cost(s) = \sum_{v \in V} c_v(s)$, and the optimal social cost is $OPT = \min_{s \in S} cost(s)$.

Remark. In our analysis it will sometimes be convenient to assume that the edges have a direction. A directed edge (v, w) indicates that the player v buys an edge to w .

4.2. Equilibrium existence

It was shown in Fabrikant et al. (2003) that for $\alpha < 1$ the clique is the social optimum and also the unique NE. For $1 < \alpha < 2$, the clique is the social optimum, but it is no longer a NE, and the star is the worst NE. Finally, for $\alpha \geq 2$, the star is the social optimum, and also a NE, but not a unique one. In this section we analyze the existence of SE for the different values of α . Our main positive result is that for any $\alpha \geq 2$ there is a SE.

Our first finding is that for a sufficiently large α (specifically, $\alpha \geq 2$), a star is a SE. The intuition behind the stability of the star structure is that it has a nice balance of short distances (a distance of 2 per node, where the minimum possible is 1) and low buying costs to all the nodes. Since a node cannot decrease its distance cost by more than $n - 2$, a star is obviously a SE for $\alpha \geq n - 2$. In what follows we show that already for $\alpha \geq 2$, a star is a SE.

Theorem 4.1. *Let s^* be a joint action where $s_r^* = \emptyset$ and $s_v^* = \{(v, r)\}$, for $v \neq r$ (i.e., $G(s^*)$ is a star in which all the nodes buy edges to the root r). For $\alpha \geq 2$, the joint action s^* is a SE.*

Proof. For contradiction, assume there exists a coalition Γ and a deviation s' , in which all nodes in Γ strictly gain from a deviation to s' . Clearly, $r \notin \Gamma$, since in s^* the root r has the lowest possible cost (it does not buy any edges and enjoys the minimum possible distance cost, i.e., distance of 1 to all nodes). For any node $v \in \Gamma$, let x_v denote the number of its *new outgoing* edges, and y_v denote the number of its *new incoming* edges. Obviously, all the new edges originate from nodes in the coalition. Thus, it must hold that $\sum_{v \in \Gamma} x_v \geq \sum_{v \in \Gamma} y_v$. We separate the analysis to two cases:

Case (a): For every $v \in \Gamma$, $x_v = y_v$. If v does not remove its original edge to r , $B(v)$ increases by αx_v , and $Dist(v)$ decreases by $x_v + y_v = 2x_v$. Therefore, v 's cost change is $\alpha x_v - (x_v + y_v) = (\alpha - 2)x_v \geq 0$, since $\alpha \geq 2$. Thus, if $x_v = y_v$, v may improve its cost only if it removes the edge to r . However, if all the nodes in Γ remove their edges to r , the only way for v to remain connected to r (to prevent a distance cost of ∞) is to buy an edge to a node $u \notin \Gamma$. In such a case, $\sum_{v \in \Gamma} x_v > \sum_{v \in \Gamma} y_v$, hence, there exists a node $v \in \Gamma$ for which $x_v > y_v$, which is a contradiction.

Case (b): There exists a node v for which $x_v > y_v$. If v does not remove its original edge to r , the change in v 's cost is $\alpha x_v - (x_v + y_v) \geq \alpha x_v - 2x_v + 1$ which is positive for $\alpha \geq 2$ (which implies that the cost of v increased). If v removes its edge to r , the change in v 's cost is $\alpha x_v - (x_v + y_v) - \alpha + 1 \geq \alpha x_v - 2x_v + 2 - \alpha = (x_v - 1)(\alpha - 2) \geq 0$, since $x_v \geq 1$ and $\alpha \geq 2$. In each case, some $v \in \Gamma$ does not strictly gain from joining the coalition, and therefore s^* is a SE. \square

Theorem 4.1 shows that for $\alpha \geq 2$, there exists a star that is a SE. Similarly, we can show that a star in which the root buys edges to all the nodes is also a SE (proof omitted). We conjecture that for $\alpha \geq 2$, any star is a SE, regardless of how the edges are bought (we can prove this conjecture only for $\alpha \geq (n - 2)$, see Theorem B.9 in the Appendix).

The next theorem establishes the existence of a SE for $\alpha \leq 1$.

Theorem 4.2. For $\alpha < 1$, s is a SE iff $G(s)$ is a clique. For $\alpha = 1$, if $G(s)$ is a clique, then s is a SE.

Proof. For $\alpha < 1$ every NE is a clique (Fabrikant et al., 2003), which implies that if s is a SE then $G(s)$ is a clique. For the other direction (which applies to $\alpha \leq 1$), consider a joint action s such that $G = G(s)$ is a clique. Suppose that there exists a coalition Γ that deviates to s' , such that the obtained graph is $G' = G(s')$, which is possibly not a clique. Let x denote the number of edges that are “missing” from the clique, i.e., $x = |E_G| - |E_{G'}|$. (If G' is a clique then $x = 0$.) For each missing edge, there exists a node $v \in \Gamma$ whose buying cost, $B(v)$, decreased by $\alpha \leq 1$. Thus $\sum_{v \in \Gamma} B(v)$ decreased by exactly $\alpha x \leq x$. However, for each missing edge, there exists at least one node in Γ whose distance cost increased by 1. Thus, $\sum_{v \in \Gamma} Dist(v)$ increased by at least x . Therefore, the sum of the costs for nodes in the coalition has not decreased. Therefore, there exists a node $u \in \Gamma$ such that $B(u) + Dist(u)$ has not decreased. In contradiction to the assumption that every $v \in \Gamma$ gains from the deviation to s' . \square

The following corollary follows immediately from Theorems 4.1 and 4.2 above, and indicates that for any $\alpha \leq 1$ or $\alpha \geq 2$, there exists a SE which is optimal.

Corollary 4.3. For any $\alpha \leq 1$ or $\alpha \geq 2$ and any n , we have $SPoS = 1$.

We next show that for $\alpha \in (1, 2)$ there is no SE (even if we limit the coalition size to 3).

Theorem 4.4. *For any $\alpha \in (1, 2)$, and any $n \geq 7$, there does not exist any 3-SE.*

The proof of Theorem 4.4 is quite involved and appears in Appendix B. In the following, we will attempt to give a very high level view of the proof. Consider a graph $G(s)$ that has an independent set of size at least 3. We can build a coalition composed of three nodes from the independent set, each buying one edge (and thus forming a triangle). Each node paid $\alpha < 2$ and its distance to the other two nodes is reduced by at least 2. Therefore, all the three nodes gain from this deviation. So our first observation is that in any 3-SE there cannot exist an independent set of size 3 (Lemma B.1). Next we show that there cannot exist any triangle in $G(s)$ (Lemma B.3). Based on those two lemmas, we show that the degree of each node must be at least $n - 3$ (Lemma B.4). Finally, we show that in such a graph, the removal of any edge is beneficial to its buyer. (The complete proof appears in Appendix B, where we also show that the theorem holds for $n = 6$ as well.)

To complete the analysis for $\alpha \in (1, 2)$, it is easy to see that for $n = 2$ any single edge is a SE, and for $n = 3$ any tree is a SE. In addition, one can verify that for $n = 4$, any ring in which each node buys a single edge is a SE. For $n = 5$, we show in the Appendix (Theorem B.5) that there does not exist any 5-SE, and for $n = 6$, we show in the Appendix (Theorem B.7) that there does not exist any 3-SE.

Interestingly enough, while coalitions of size 3 or more exclude any SE, the following theorem shows that 2-SE do exist for any number of players.

Theorem 4.5. *Let s^* be a joint action where $s_r^* = \emptyset$ and $s_v^* = \{(v, r)\}$, for $v \neq r$ (i.e., $G(s^*)$ is a star in which all the nodes buy edges to the root r). For $\alpha \in (1, 2)$ and any n , the joint action s^* is a 2-SE.*

Proof. Obviously, s^* is resilient to coalitions of size 1 since it is a NE (Fabrikant et al., 2003). We next show that it is resilient to any coalition of size 2. First note that since the root (r) does not buy any edges and enjoys the minimum possible distance cost, it will not belong to any coalition. For any other two nodes u_1, u_2 , we show that there does not exist a graph $G' = G(s')$ they can form in which both nodes gain. In any such coalition, there must exist an edge between u_1 and u_2 ; otherwise, each one of them can deviate unilaterally, in contradiction to the fact that s^* is a NE. Suppose WLOG that u_1 bought the edge (u_1, u_2) . For any node v , other than u_2 and r , the only way node u_1 can decrease the distance to it to 1 is by buying the edge (u_1, v) . But since $\alpha > 1$ this will result in a net loss. Therefore, u_1 does not buy any edges to nodes other than r or u_2 . We have two cases involving nodes u_2 and r : Case (a): u_1 does not remove the edge (u_1, r) . Then, the buying cost increases by $\alpha > 1$, while the distance cost decreases only by 1. Such a deviation results in a net loss for u_1 . Case (b): u_1 removes the edge (u_1, r) . Then, the buying cost does not change, but the total distance does not decrease (and may increase). Thus, u_1 did not gain from the deviation. Therefore, there is no coalition of size 2 where both players gain from the deviation. \square




Fig. 5. This network is a SE for $\alpha \geq 3$, and the *OPT* is a star. In this example, $SPoA = 29/27 > 1$ for $\alpha = 3$.

4.3. Strong price of anarchy

Recall that for $\alpha < 1$ the clique is the only SE. For $\alpha = 1$, it is easy to see that $PoA < 2$ (since in any NE the distance between any two nodes cannot exceed 2). For $\alpha \in (1, 2)$ we do not have any SE for $n \geq 5$. In this section we bound the SPoA for $\alpha \geq 2$.

We start with an example demonstrating that there are cases in which the SPoA is strictly greater than 1.

Example 4.6. Consider the network $G(s)$ presented in Fig. 5, and suppose $\alpha \geq 3$. We claim that s is a SE. To see this, consider first nodes v_1 and v_4 . $B(v_1) = B(v_4) = 0$, and $Dist(v_1) = Dist(v_4) = 6$. The minimum possible distance cost of a node for $n = 4$ is 3. Therefore, for $\alpha \geq 3$, nodes v_1 and v_4 will not join any coalition. Thus, a coalition can be formed only by v_2 and v_3 , who buy a total of 3 links in s . Since the graph must stay connected in any profitable deviation, they must keep buying at least 3 links in any deviation. By deviating, $v \in \{v_2, v_3\}$ can at most reduce its distance costs from 4 to 3, which is the minimum possible distance cost for $n = 4$. Since $\alpha \geq 3$, nodes v_2 and v_3 cannot reduce their total costs if they increase the number of links they buy. Therefore, the total number of links bought must remain 3 in any deviation. However, in this case, the only way for $v \in \{v_2, v_3\}$ to reduce its cost is by reducing the distance cost to 3, which is impossible to accomplish for both nodes with only 3 links. We conclude that s is a SE for $\alpha \geq 3$. Summing up the social cost in s for $\alpha = 3$, we get: $c(s) = 6 + 6 + 4 + 4 + 3\alpha = 29$. In contrast, the cost of *OPT*, which is a star, for $\alpha = 3$ is: $c(OPT) = 3 + 5 + 5 + 5 + 3\alpha = 27$. We get that $SPoA = 29/27 > 1$.

Similarly to Albers et al. (2006), we make use of the following lemma, which bounds the PoA as a function of n, α and the distance cost of any node.

Lemma 4.7. *Let s be a NE. For any node v we have $cost(s) \leq (n - 1)(2\alpha + n - 1 + Dist(v))$.*

Our main result is the following.

Theorem 4.8. *For any $\alpha \geq 2$ and any n , we have $SPoA \leq 2$.*

The proof of Theorem 4.8 follows directly from the next two lemmas.

Lemma 4.9. *Let s be a NE. Assume that for every node v , such that $s_v \neq \emptyset$, we have that $Dist(v) > 3n - 5$. Then s is not a SE.*

Proof. Let Γ be the set of all nodes v that bought some edge in s , i.e., $\Gamma = \{v \mid s_v \neq \emptyset\}$. We will show that Γ can deviate, such that all its members would benefit from a deviation. In the deviation we build a tree T in which each node in Γ buys at most the same number of edges as in s and it strictly reduces the distances to other nodes, i.e., every $v \in \Gamma$ lowers its cost in the deviation.

Assume that there is some node $r \notin \Gamma$. Let T be the following tree. The root of the tree is r . The nodes in the first level are the nodes in Γ . The nodes in the second level are the remaining

$n - |\Gamma| - 1$ nodes. Each node in Γ buys an edge to the root r and at most $|s_v| - 1$ edges to nodes in the second level (the leaves). Clearly the number of edges that each node in Γ bought can only decrease. To see that we have enough edges to connect all the $n - |\Gamma| - 1$ leaves, note that in s at least $n - 1$ edges are bought (otherwise some node is disconnected, and all the nodes have infinite cost). We need only $n - 1$ edges to connect all the nodes in T , so we have a sufficient number of edges.

Fix a node $v \in \Gamma$. The distances $Dist(v)$ in T is at most $1 + 2(|\Gamma| - 1) + 3(n - |\Gamma| - 1) \leq 3n - 5$, since $|\Gamma| \geq 1$. Hence, node v improved on its distance cost in s and did not increase its buying cost. Therefore, in this case, s is not a SE.

In the case in which there is no $r \notin \Gamma$ we can select any node to be the root and the remaining nodes will buy an edge to it. Since all the nodes bought at least one edge, the cost of buying edges can only decrease per node. The distances of a node v is now at most $2(n - 2) + 1 \leq 3n - 5$ for $n \geq 2$, hence v improved on its distance cost in s . \square

Lemma 4.10. *Let s be a NE. Assume that for some node v , such that $s_v \neq \emptyset$, we have that $Dist(v) \leq 3n - 5$. Then $\frac{cost(s)}{cost(OPT)} \leq 2$.*

Proof. By Lemma 4.7 we have that

$$\begin{aligned} cost(s) &\leq (n - 1)(2\alpha + n - 1 + Dist(v)) \leq (n - 1)(2\alpha + n - 1 + 3n - 5) \\ &= 2(n - 1)(\alpha + 2n - 3). \end{aligned}$$

For OPT (which is a star) we have:

$$cost(OPT) = \alpha(n - 1) + (n - 1)(2(n - 2) + 1) + (n - 1) = (n - 1)(\alpha + 2n - 2).$$

Therefore:

$$\frac{cost(s)}{cost(OPT)} \leq \frac{2(n - 1)(\alpha + 2n - 3)}{(n - 1)(\alpha + 2n - 2)} \leq \frac{2(n - 1)(\alpha + 2n - 2)}{(n - 1)(\alpha + 2n - 2)} = 2. \quad \square$$

While the theorem above shows that \leq -SPoA2, the analysis requires that in some cases the coalition may include all n agents. Providing an upper bound for k -SPoA for $k < n$ is left as an open question.

Acknowledgments

We thank Noam Nisan for helpful comments on an earlier draft of the paper.

Appendix A. Job scheduling

Lemma 3.2. *Consider two joint actions s and s' such that the load vectors $L(s)$ and $L(s')$ differ only in the loads of machines in a set $M' \subseteq M$. If for each $M_i \in M'$, $L_i(s) < \max_k \{L_k(s') \mid M_k \in M'\}$ then $s < s'$.*

Proof. Let $SL(s)$ and $SL(s')$ denote the vectors $L(s)$ and $L(s')$ sorted in a non-increasing order, respectively. Let $k' = \arg \max_k \{L_k(s') \mid M_k \in M'\}$ and $k = \arg \max_k \{L_k(s) \mid M_k \in M'\}$ be the index of the most loaded machine from M' in $L(s')$ and $L(s)$, respectively. In the sorted load vectors $SL(s)$ and $SL(s')$, the smallest index that differs between the vectors is of a machine that

has a load of $L_{k'}(s')$ in $SL(s')$, and $L_k(s)$ in $SL(s)$. Since $L_{k'}(s') < L_k(s)$, $SL(s)$ is lexicographically smaller than $SL(s')$, and $s < s'$. \square

Lemma A.1. *The lexicographically minimal joint action s is a NE.*

Proof. For contradiction, assume that there exists a job J , where $s_J = M_i$, that can benefit from deviating to $s'_J = M_k$. Let s' denote the joint action after J deviates. Since $w_i(J) > 0$, we have $L_i(s') < L_i(s)$. Since J benefits from the deviation, we have $L_k(s') < L_i(s)$. Therefore, $L_i(s) > \max\{L_i(s'), L_k(s')\}$. Since $L(s)$ and $L(s')$ differ only in the loads on M_i and M_k , by Lemma 3.2 we have $s' < s$, which contradicts the minimality of s . \square

Theorem 3.8. *For any job scheduling game with 2 unrelated machines and n jobs, $SPoA \leq 2$.*

Proof. Let s be a SE and, WLOG, $L_2(s) \geq L_1(s)$. In the case that for every $J \in B_2^s$ we have $w_2(J) \leq w_1(J)$, by inequality (2), $L_2(s) \leq 2 \cdot OPT$, and we are done. Otherwise, there exists some $J \in B_2^s$ such that $w_2(J) > w_1(J)$. Since s is a SE, it is in particular a NE, which means that no job on M_2 can gain by unilaterally migrating to M_1 . Therefore, $L_2(s) \leq L_1(s) + w_1(J)$. By inequality (1), we get:

$$L_2(s) \leq L_1(s) + OPT. \tag{5}$$

The following are the possible cases relating OPT , $L_1(s)$ and $L_2(s)$:

1. If $L_1(s) \leq L_2(s) < OPT$, this is impossible (a contradiction to the minimality of OPT).
2. If $OPT < L_1(s) \leq L_2(s)$, then s is not resilient to a coalition of size n (since by deviating to OPT all the players strictly gain).
3. If $L_1(s) \leq OPT \leq L_2(s)$, then from inequality (5), we get: $L_2(s) \leq L_1(s) + OPT \leq 2OPT$.

Taking the maximum over all cases, we get: $SPoA \leq 2$. \square

Theorem 3.13. *For every $m \geq 2$, there exists a job scheduling game with m unrelated machines for which $SPoA \geq m$.*

Proof. Consider a job scheduling game with m jobs and m unrelated machines, where for each job J_ℓ , $\ell = 2, \dots, m$: $w_\ell(J_\ell) = \ell$, $w_{\ell-1}(J_\ell) = 1$, and $w_i(J) = \infty$ for $i \neq J, J - 1$. For job J_1 , $w_1(J_1) = w_m(J_1) = 1$, and $w_i(1) = \infty$, for $i \neq 1, m$. The joint action that achieves social optimum is: $OPT(J_\ell) = M_{\ell-1}$ for $\ell = 2, \dots, m$, and $OPT(J_1) = M_m$, which yields a makespan of 1. However, the following joint action s has a cost of m : for $\ell = 1, \dots, m$, $s_{J_\ell} = M_\ell$. (To see that s is a SE note that for any coalition Γ the player with the lowest index can not lower its cost from a deviation of players in Γ .) Since the makespan of s is m we have that $SPoA \geq m$. \square

Theorem 3.14. *For every $n, k \geq 2$, $k < n$, there exists a job scheduling game with 2 unrelated machines and n jobs, s.t. $k\text{-SPoA} \geq \frac{n}{2k}$.*

Proof. Consider the following job scheduling game. Let $w_1(J_i) = 1$ and $w_2(J_i) = 1/(n - 1)$, for $2 \leq i \leq n$, and let $w_1(J_1) = 2k$ and $w_2(J_1) = n + 1 + k + \epsilon$. In this game $OPT(J_1) = M_1$ and $OPT(J_2) = \dots = OPT(J_n) = M_2$, which yields a cost of $2k$. The joint action $s_1 = M_2$ and $s_2 = \dots = s_n = M_1$ is a k -SE. (To see that it is a k -SE note that if J_1 migrates to M_1 the new

load is $n + 2k$. This implies that at least $k + 1$ jobs have to migrate from M_1 in order that it will be beneficial for J_1 to migrate to M_1 .) Therefore, $k\text{-SPoA} \geq \frac{n-1+k+\epsilon}{2k} \geq \frac{n}{2k}$. \square

Theorem 3.17. *For every $m \geq 2$, there exists a job scheduling game with m identical machines and n jobs, s.t. $\text{SPoA} \geq \frac{2}{1+\frac{1}{m}}$.*

Proof. Consider the following game of $n = 2m$ jobs running on m machines:

$$w(J_1) = \dots = w(J_m) = 1 \quad \text{and} \quad w(J_{m+1}) = \dots = w(J_{2m}) = \frac{1}{m}.$$

The optimum is to have one job of weight 1 and one job of weight $\frac{1}{m}$ on each machine, which yields a makespan of $1 + \frac{1}{m}$. Consider the joint action s as follows:

$$\forall i \in 1, \dots, m - 2, \quad s_i = M_i, \quad s_{m-1} = s_m = M_{m-1}, \quad \text{and} \quad s_{m+1} = \dots = s_{2m} = M_m.$$

In this case, we have $L_i(s) = 1$ for $i \neq m - 1$, and $L_{m-1}(s) = 2$. Clearly, the first $m - 2$ jobs can not gain from any deviation, since each is alone on a machine. Since the load on the first $m - 2$ machines is 1, no other job can gain from migrating to these machines. Therefore we can concentrate on the last two machines.

For one of the two jobs on M_{m-1} to improve its load, it needs to migrate alone to M_m and have at least one job from M_m migrate back to M_{m-1} . However, a job in M_m would gain from migrating to M_{m-1} only if both jobs on M_{m-1} migrate to M_m . This implies that the joint action s is resilient to deviation of coalitions of any size. Since the makespan of s is 2, we have that $\text{SPoA} \geq \frac{2}{1+\frac{1}{m}}$. \square

A.1. Mixed deviations

In the remainder of this appendix we discuss the case of mixed deviations. We start with the following lemma which greatly limits the structure of a mixed SE.

Lemma A.2. *Given a k -SE with mixed strategies π , for some $k \geq 2$, let J_1 and J_2 be two jobs with strictly mixed strategies. The supports of π_1 and π_2 must be disjoint.*

Proof. For contradiction, assume the jobs J_1 and J_2 have strictly mixed strategies and their supports intersect. We will show that π is not resilient to a coalition of J_1 and J_2 .

Suppose that the supports of J_1 and J_2 intersect in at least two machines. WLOG, let M_1 and M_2 denote these machines. Consider the following mixed deviation π' for a coalition of J_1 and J_2 :

$$\begin{aligned} \pi'_1 &= (\pi_1(1) + \pi_1(2), \quad 0, & \pi_1(3), \quad \pi_1(4), \quad \dots, \quad \pi_1(m)), \\ \pi'_2 &= (0, & \pi_2(1) + \pi_2(2), \quad \pi_2(3), \quad \pi_2(4), \quad \dots, \quad \pi_2(m)). \end{aligned}$$

The strategies of other jobs in π' are the same as in π . In π , J_1 is indifferent between machines M_1 and M_2 , since π is a NE. Therefore, by having machine M_1 removed from the support of J_2 , the expected load observed by J_1 in π' is reduced by $(\pi_1(1) + \pi_1(2))\pi_2(1)w_1(J_2)$. Similarly, the expected load observed by J_2 in π' is reduced by $(\pi_2(1) + \pi_2(2))\pi_1(2)w_2(J_1)$. Therefore, both jobs benefit from the deviation, and hence π is not resilient to coalitions of size 2.

Suppose the supports of J_1 and J_2 intersect in exactly one machine. WLOG, let M_1 denote the machine only in the support of J_1 , M_2 denote the machine only in the support of J_2 and M_3 denote the machine in the intersection of the supports. (M_1 and M_2 exist since both π_1 and π_2 are strictly mixing.) Let $\rho = \frac{1}{2} \min\{\pi_1(3), \pi_2(3)\}$. Consider the following mixed deviation π' for a coalition of J_1 and J_2 :

$$\begin{aligned} \pi'_1 &= (\pi_1(1) + \rho, 0, \pi_1(3) - \rho, \pi_1(4), \dots, \pi_1(m)), \\ \pi'_2 &= (0, \pi_2(2) + \rho, \pi_2(3) - \rho, \pi_2(4), \dots, \pi_2(m)). \end{aligned}$$

The strategies of other jobs in π' are the same as in π . Let $\beta(q) = (q - \rho)\rho$. The expected load observed by J_1 in π' is reduced by $\beta(\pi_1(3))w_3(J_2)$ and the expected load of J_2 in π' is reduced by $\beta(\pi_2(3))w_3(J_1)$. Again, both jobs benefit from the deviation, contradicting the assumption that π is a 2-SE. \square

Theorem 3.19. *For $m \geq 5$ identical machines and $n > 3m$ unit jobs, there is no 4-SE, if mixed deviations are allowed.*

Proof. We first consider equilibria with pure strategies. Since all jobs are unit sized, the only equilibrium with pure strategies is when the load on each machines is either $\lfloor \frac{n}{m} \rfloor$ or $\lceil \frac{n}{m} \rceil$.

Let $k = \lceil \frac{n}{m} \rceil$. Since $n > 3m$, there exists a machine with at least 4 jobs assigned to it. WLOG, assume M_1 is one of these machines, and J_1, J_2, J_3, J_4 are four of the k jobs that chose it.

Consider the following mixed deviation of these jobs:

$$\begin{aligned} \pi_1 &= \left(\frac{1}{2}, \frac{1}{2}, 0, 0, 0, 0, \dots, 0\right), & \pi_2 &= \left(\frac{1}{2}, 0, \frac{1}{2}, 0, 0, 0, \dots, 0\right), \\ \pi_3 &= \left(\frac{1}{2}, 0, 0, \frac{1}{2}, 0, 0, \dots, 0\right), & \pi_4 &= \left(\frac{1}{2}, 0, 0, 0, \frac{1}{2}, 0, \dots, 0\right). \end{aligned}$$

The strategies of the remaining jobs are unchanged. In the original joint strategy, the load observed by each of these jobs is k . In the new strategy π , if any of these jobs migrates, it observes a load of at most $k + 1$. If a job remains on M_1 , it observes the load incurred by itself, $k - 4$ jobs that did not change their strategy, and an expectation half of the remaining three jobs. In total, the expected load observed by each of the first four jobs in π is at most $1 + \frac{(k-2.5)+k}{2} = k - \frac{1}{4}$. Since all jobs in the coalition benefit from the deviation, no pure NE in this setting is a 4-SE.

We now consider equilibria with mixed strategies. Clearly, the expected load on each machine has to be between $k - 1$ and k . By Lemma A.2, on each machine there is at most one job that has a mixed strategy.

WLOG, assume M_1 is the most loaded machine. If there are 4 jobs that purely choose M_1 as their strategy, then the same deviation described for the pure case holds for these jobs. Otherwise, $k = 4$ and there are 3 jobs that purely choose M_1 , and another job that has a mixed strategy and M_1 is in its support vector. WLOG, assume J_1 is the job on M_1 that has a mixed strategy and that M_2 is one of the other machines in its support. Let p denote the probability that J_1 chooses M_1 . We also assume that the other jobs that choose M_1 are J_2, J_3 and J_4 (the expected load on M_1 is therefore $3 + p$).

Consider the following mixed deviation π of these jobs:

$$\begin{aligned} \pi_1 &= \left(\frac{p}{2}, 1 - \frac{p}{2}, 0, 0, 0, 0, \dots, 0\right), & \pi_2 &= \left(\frac{1}{2}, 0, \frac{1}{2}, 0, 0, 0, \dots, 0\right), \\ \pi_3 &= \left(\frac{1}{2}, 0, 0, \frac{1}{2}, 0, 0, \dots, 0\right), & \pi_4 &= \left(\frac{1}{2}, 0, 0, 0, \frac{1}{2}, 0, \dots, 0\right). \end{aligned}$$

The strategies of the remaining jobs are unchanged. In the original joint strategy, the load observed by J_1 is 4 on both M_1 and M_2 , since it observes itself and three other jobs. The deviation decreases the expected load J_1 observes on M_1 to 2.5, while the observed load on M_2 remains 4. In expectation, J_1 observes a decreased load of $(\frac{p}{2} \cdot \frac{5}{2}) + (1 - \frac{p}{2})4 = 4 - \frac{3p}{4}$. As for the other jobs in the coalition, in the original joint strategy, the expected load observed by each job is $3 + p$. In π , the expected load observed by each job if it remains on M_1 is $2 + \frac{p}{2}$, and at most $4 + p$ if it migrates. In total the expected load is at most $\frac{(2+p/2)+(4+p)}{2} = 3 + \frac{3p}{4}$. Since all jobs in the coalition benefit from the deviation, no mixed NE in this setting is a 4-SE. \square

Theorem 3.20. *There exists a job scheduling game on 2 identical machines and 3 jobs, where no joint mixed strategy is a 2-SE, when mixed deviations are allowed.*

Proof. Consider 2 identical machines and 3 jobs with weights $w(J_1) = 1 - \epsilon$, $w(J_2) = 1$, $w(J_3) = 1 + \epsilon$, where ϵ is a small value that will be determined later. In any pure NE, J_1 and J_2 are assigned to the same machine, while J_3 is assigned alone. WLOG, we assume J_1 and J_2 are assigned to M_1 , and J_3 to M_2 .

Consider the following mixed deviation π of J_1 and J_2 :

$$\pi_1 = \left(\frac{3}{4}, \frac{1}{4}\right), \quad \pi_2 = \left(\frac{3}{4}, \frac{1}{4}\right).$$

The load on M_1 in the original assignment is $2 - \epsilon$. After the deviation, J_1 observes an expected load of $1 - \epsilon + \frac{3}{4} \cdot \frac{3}{4} + \frac{1}{4}(1 + \epsilon + \frac{1}{4}) = 1\frac{7}{8} - \frac{3}{4}\epsilon$, while J_2 observes an expected load of $1 + \frac{3}{4} \cdot \frac{3}{4}(1 - \epsilon) + \frac{1}{4}(1 + \epsilon + \frac{1}{4}(1 - \epsilon)) = 1\frac{7}{8} - \frac{3}{8}\epsilon$. For any $0 < \epsilon < \frac{1}{5}$ both jobs improve their observed load, and therefore there is no pure NE that is a 2-SE.

It remains to show that no mixed NE is resilient to coalitions of size 2. By Lemma A.2 it is sufficient to consider only NE where each machine is included in at most the support of one job that plays a strictly mixed strategy. Since there are only 2 machines, there can be only one job J that is using a strictly mixed strategy (otherwise there will be intersecting supports). In any mixed NE, J needs to be indifferent between M_1 and M_2 , which is impossible, since the other two jobs have different weights. \square

Appendix B. Network creation

Theorem 4.4. *For any $\alpha \in (1, 2)$, and any $n \geq 7$, there does not exist any 3-SE.*

Proof. We first establish the following sequence of lemmas.

Lemma B.1. *For $\alpha \in (1, 2)$, in any 3-SE, s , there does not exist any independent set of size 3 in $G(s)$.*

Proof. For contradiction, assume that there exists an independent set I such that $|I| \geq 3$. Let $\Gamma \subset I$ be any subset of three nodes. For a deviation, let the nodes in Γ form a triangle in which

each node buys a single edge. In the original graph, $G(s)$, $\forall v, u \in \Gamma, v \neq u$, we had $\delta(v, u) \geq 2$. Therefore, for each $v \in \Gamma$, the distance cost, $Dist(v)$, decreased by at least 2 after the deviation. Since $B(v)$ increased only by $\alpha < 2$, each $v \in \Gamma$ lowered its cost by deviating.¹⁷ \square

Lemma B.2. For $\alpha \in (1, 2)$, in any NE s , if there exists a set of nodes U that form a clique in $G(s)$, then if $u_1 \in U$ buys the edge to $u_2 \in U$, there must exist a node w_2 that is directly connected to u_2 but not to any other node $u \in U \setminus \{u_2\}$.

Proof. Suppose there does not exist such w_2 , we will show that u_1 strictly gains by removing (u_1, u_2) , contradicting the assumption that s is a NE. For any $v \in V \setminus U$, if the shortest path from u_1 to v does not go through any node in $U \setminus \{u_1\}$, then $\delta(u_1, v)$ is not affected by the removal. Otherwise, since the diameter of the graph cannot exceed 2 (for $\alpha \in (1, 2)$), v must be directly connected to some $x_v \in U \setminus \{u_1\}$. By our assumption, v must be directly connected to some $u \in U \setminus \{u_1, u_2\}$, but then, $\delta(u_1, v)$ is not affected by the removal either. We conclude that if u_1 unilaterally deviates, and does not buy the edge (u_1, u_2) , then its distance cost, $Dist(u_1)$, increases by only 1, while its buying cost, $B(u_1)$, decreases by $\alpha > 1$, thus its cost strictly decreases. \square

We use the above lemma to prove that a 3-SE cannot include triangles.

Lemma B.3. For $\alpha \in (1, 2)$, in any 3-SE s , there does not exist any triangle in $G(s)$.

Proof. Suppose that the set of nodes $U = \{u_1, u_2, u_3\}$ forms a triangle. It is easy to see that in any triangle, there must exist a node $u \in U$ that buys exactly one triangle edge. Assume WLOG that u_2 buys the edge (u_2, u_3) , and u_1 buys the edge (u_1, u_2) (see Fig. 6(a)). We will show that u_2 strictly gains by removing the edge (u_2, u_3) .

For any $v \in V \setminus U$, if the shortest path from u_2 to v does not go through u_1 or u_3 , then removing (u_2, u_3) does not affect $\delta(u_2, v)$. Otherwise, v must be directly connected to either u_1 or u_3 (for any two nodes $v_1, v_2 \in V$, it must hold that $\delta(v_1, v_2) \leq 2$, otherwise, since $\alpha < 2$, each




Fig. 6. $\alpha \in (1, 2)$. (a) Illustration of a triangle and connected nodes. (b) The nodes of a pentagon can improve by buying a pentagram.

¹⁷ The proof of this lemma can serve as a simple example to the difference between the strong equilibrium and the coalition-proof Nash equilibrium solution concepts. An independent set of size 3 cannot exist in any SE since deviating by forming a triangle will improve the utility of each one of the players. Yet, as demonstrated by Lemma B.3, a triangle is not resilient to a deviation by one of the players forming the triangle. Therefore, this proof does not eliminate the existence of an independent set of size 3 in a coalition-proof Nash equilibrium.

one of these nodes gain by buying the edge between them). If it is directly connected to u_1 , then again, removing (u_2, u_3) does not affect $\delta(u_2, v)$. Otherwise, v is directly connected to u_3 , but not to u_1 or u_2 . By Lemma B.2, there exists a node w_2 that is directly connected to u_2 but not to u_1 or u_3 . Thus, by Lemma B.1, v must be directly connected to node w_2 , otherwise $\{v, u_1, w_2\}$ form an independent set of size 3. So $\delta(u_2, v)$ remains 2 and is not affected by the deviation. Therefore, $Dist(u_2)$ increases only by 1 (since $\delta(u_2, u_3)$ increases by 1), and $B(u_2)$ decreases by $\alpha > 1$, so it implies that u_2 strictly gains. \square

Using the above lemmas, we derive a lower bound on the degree of each node in any 3-SE. Let $deg(v, G)$ be the degree of node v in the graph G .

Lemma B.4. For $\alpha \in (1, 2)$, in any 3-SE s , for every v , we have $deg(v, G(s)) \geq n - 3$.

Proof. For contradiction suppose that there exists a node v such that $deg(v, G(s)) \leq n - 4$. Then, there are at least 3 nodes that are not directly connected to v . If any 2 of these nodes are not directly connected, then these 2 nodes together with v form an independent set of size 3, contradicting Lemma B.1. Therefore, these nodes must form a clique. However, this is a contradiction to Lemma B.3. \square

We now complete the proof of the theorem. By Lemma B.4, the degree of each node in any 3-SE must be at least $n - 3$. Then, for $n \geq 7$, any edge removal can strictly decrease the cost of the node that bought it. Consider the edge (w, u) . If w removes the edge, $B(w)$ decreases by $\alpha > 1$. We claim that $Dist(w)$ increases only by 1 (i.e., the only effect is that $\delta(w, u)$ increases from 1 to 2). To see this, note that for $n \geq 7$, if the degree of any node is at least $n - 3$, then after removing (w, u) , their degrees are at least $n - 4$, and since for any $n \geq 7$, it holds that $n - 4 + n - 4 > n - 2$, they must have a common neighbor. In addition, for any node $u' \neq w, u$, both w and u' must have a common neighbor, since $n - 4 + n - 3 > n - 2$ (where $n - 3$ and $n - 4$ are the minimal respective degrees of u' and w). Therefore, by removing the edge (w, u) , $Dist(w)$ increases by 1, while $B(w)$ decreases by $\alpha > 1$, so w strictly gains from the removal. This completes the proof of the theorem. \square

Theorem B.5. For $\alpha \in (1, 2)$, $n = 5$, there does not exist any 5-SE.

Proof. In our proof we use the following lemma.

Lemma B.6. For $\alpha \in (1, 2)$, in any n -SE s , there does not exist any cycle of length 5 in $G(s)$.

Proof. Suppose that there exist a cycle of length 5 $(u_1, u_2, u_3, u_4, u_5, u_1)$ (the directions of the edges do not matter). First note that by Lemma B.3, there cannot exist any other edge between these nodes (otherwise, it forms a triangle). But then, the nodes of the cycle can strictly gain by buying a pentagram $(u_1, u_3), (u_3, u_5), (u_5, u_2), (u_2, u_4), (u_4, u_1)$ (see Fig. 6(b)), since for each node u , $B(u)$ increases by $\alpha < 2$, but $Dist(u)$ decreases by 2. \square

Since we consider $n = 5$, by Lemmas B.3 and B.6, there do not exist odd cycles (i.e., cycles of size 3 or 5). Thus, the graph must be a bipartite graph. But then, there must exist an independent set of size 3, in contradiction to Lemma B.1. This completes the proof of the theorem. \square

Theorem B.7. *For $\alpha \in (1, 2)$, $n = 6$, there does not exist any 3-SE.*

Proof. In order to prove the theorem, we first show that there is no cycle of length 5 in any 3-SE.

Lemma B.8. *For $\alpha \in (1, 2)$, $n = 6$, in any 3-SE s , there does not exist any cycle of length 5 in $G(s)$.*

Proof. Let v_1, \dots, v_6 denote the nodes of $G(s)$, which contains a cycle of length 5, and assume without loss of generality that v_1, \dots, v_5 form a cycle. If there exist any other edges between the nodes of the cycle, then $G(s)$ contains a triangle, and by Lemma B.3 s is not a 3-SE. Clearly, v_6 should be connected to the cycle. However, if v_6 is connected to at least three other nodes, then v_6 and two of its neighbors form a triangle, and s is not a 3-SE by Lemma B.3. Otherwise, v_6 and two of the nodes that are not connected to v_6 form an independent set of size 3, and by Lemma B.1 s is not a 3-SE. \square

Based on Lemmas B.3 and B.8, for $n = 6$, in any 3-SE, there do not exist any cycles of odd length, and therefore, the network must be a bipartite graph. But any bipartite graph for $n = 6$ contains an independent set of size 3, and by Lemma B.1 cannot be a 3-SE. This completes the proof of the theorem. \square

Theorem B.9. *Let s be a joint action such that $G(s)$ is a star. For $\alpha \geq n - 2$, s is an n -SE.*

Proof. As in the proof of Theorem 4.1, for any node $v \in \Gamma$, let x_v and y_v denote the respective numbers of its *new outgoing* and *new incoming* edges. Obviously, all the new edges originate from nodes in the coalition. Thus, it must hold that $\sum_{v \in \Gamma} x_v \geq \sum_{v \in \Gamma} y_v$. If $r \notin \Gamma$, we show that the star is a SE for $\alpha \geq 2$, and if $r \in \Gamma$, we show that the star is a SE for $\alpha \geq n - 2$.

Case (a): The root, r , does not belong to the coalition Γ . In this case we claim that Γ would be also a deviation for s^* , where $s_r^* = \emptyset$ and $s_v^* = \{(v, r)\}$, for $v \neq r$ (i.e., $G(s^*)$ is a star in which all the nodes buy edges to the root r). This is since every node in Γ has a higher cost in s^* compared to s . Therefore, the fact that s is an n -SE follows from Theorem 4.1.

Case (b): The root, r , belongs to the coalition Γ . Let X denote the set of nodes to which r buys new edges, and let Y denote the set of nodes from which r removes edges. Since in the original star r has the minimum possible distance cost, r will join the coalition only if $\alpha|X| < \alpha|Y|$. That is, $|X| < |Y|$. For any coalition, each node $v \in Y$ must be connected to the graph after the deviation. We will show that there does not exist a coalition in which all the nodes that belong to Y stay connected, and thus reach a contradiction. For a node $v \neq r$, buying a new edge costs $\alpha \geq n - 2$, while it can gain no more than $n - 2$ in distance cost (since in the original graph it had a distance cost of $1 + 2(n - 2)$). Thus, the only nodes that might gain from buying new edges are nodes in X (since they removed their edge from r). However, using the same reasoning, they will not buy more than a single edge. Therefore, in order for the set of nodes Y to be connected to the graph, it must hold that $|X| \geq |Y|$, contradicting the assumption that since $r \in \Gamma$ we have that $|X| < |Y|$. \square

Lemma 4.7. *Let s be a NE. For any node v we have $cost(s) \leq (n - 1)(2\alpha + n - 1 + Dist(v))$.*

Proof. The proof follows a similar proof in Albers et al. (2006), with minor modifications. Fix v and consider the shortest path tree $T(v)$. For any vertex $u \in V$, let E_u be the number of tree

edges bought by u in $T(v)$. Clearly, v bought only tree edges while other vertices may have bought non-tree edges. We now prove that for every vertex $u \neq v$,

$$c_u(s) \leq \alpha(E_u + 1) + \text{Dist}_s(v) + n - 1 - \delta_s(v, u). \quad (6)$$

Since s is a NE, $c_u(s)$ is lower bounded by the following alternative action: Vertex u discards all non-tree edges, and buys an additional edge to v . The new cost for buying edges is $\alpha(E_u + 1)$. Since only non-tree edges were deleted, the distance between u and any other vertex $w \neq u$ is at most $1 + \delta_s(v, w)$. Summing over all vertices except for u , the new distance cost for u yields the bound in Eq. (6). Since the number of tree edges is $n - 1$, summing over all $n - 1$ vertices $u \neq v$, adding $c_v(s) = \alpha E_v + \text{Dist}_s(v)$, and noting that $\sum_{u \neq v} \delta_s(v, u) = \text{Dist}(v)$ completes the proof. \square

References

- Abraham I., Dolev, D., Gonen, R., Halpern, J., 2006. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In: Symposium on Principles of Distributed Computing (PODC), pp. 53–62.
- Albers, S., Elits, S., Even-Dar, E., Mansour, Y., Roditty, L., 2006. On Nash equilibria for a network creation game. In: Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 89–98.
- Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T., 2004. The price of stability for network design with fair cost allocation. In: Annual Symposium on Foundations of Computer Science (FOCS), pp. 295–304.
- Aumann, R., 1959. Acceptable points in general cooperative n-person games. In: Contributions to the Theory of Games, vol. 4.
- Awerbuch, B., Azar, Y., Richter, Y., Tsur, D., 2006. Tradeoffs in worst-case equilibria. Theoret. Comput. Sci. 361 (2–3), 200–209. (A preliminary version in the 1st International Workshop on Approximation and Online Algorithms, 2003.)
- Bernheim, D.B., Peleg, B., Whinston, M.D., 1987. Coalition-proof Nash equilibria: I. Concepts. J. Econ. Theory 42, 1–12.
- Christodoulou, G., Koutsoupias, E., 2005. On the price of anarchy and stability of correlated equilibria of linear congestion games. In: Annual European Symposium on Algorithms (ESA), pp. 59–70.
- Corbo, J., Parkes, D., 2005. The price of selfish behavior in bilateral network formation. In: ACM Symposium on Principles of Distributed Computing (PODC), pp. 99–107.
- Czumaj, A., Vöcking, B., 2002. Tight bounds for worst-case equilibria. In: ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 413–420.
- Dutta, B., Mutuswami, S., 1997. Stable networks. J. Econ. Theory 76, 322–344.
- Epstein, A., Feldman, M., Mansour, Y., 2007. Strong equilibrium in cost sharing connection games. In: ACM Conference on Electronic Commerce (EC), pp. 84–92.
- Even-Dar, E., Kesselman, A., Mansour, Y., 2003. Convergence time to Nash equilibria. In: International Colloquium on Automata, Languages and Programming (ICALP), pp. 502–513.
- Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C., Shenker, S., 2003. On a network creation game. In: ACM Symposium on Principles of Distributed Computing (PODC), pp. 347–351.
- Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., 2007. Strong price of anarchy for machine load balancing. In: International Colloquium on Automata, Languages and Programming (ICALP), pp. 583–594.
- Fotakis, D., Kontogiannis, S., Spiraklis, P., 2006. Atomic congestion games among coalitions. In: International Colloquium on Automata, Languages and Programming (ICALP), pp. 573–584.
- Goldberg, A.V., Hartline, J., 2005. Collusion-resistant mechanisms for single-parameter agents. In: ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 620–629.
- Hayrapetyan, A., Tardos, E., Wexler, T., 2006. The effect of collusion in congestion games. In: 38th ACM Symposium on Theory of Computing (STOC), pp. 89–98.
- Holzman, R., Law-Yone, N., 1997. Strong equilibrium in congestion games. Games Econ. Behav. 21, 85–101.
- Holzman, R., Law-Yone, N., 2003. Network structure and strong equilibrium in route selection games. Math. Soc. Sci. 46, 193–205.
- Jackson, M., 2004. Survey of models of network formation: Stability and efficiency. Chapter 1 in: Demange, G., Wooders, M. (Eds.), Group Formation in Economics: Networks: Clubs and Coalitions. Cambridge Univ. Press.

- Koutsoupias, E., Papadimitriou, C.H., 1999. Worst-case equilibria. In: *Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 404–413.
- Kranton, R., Minehart, D., 2001. A Theory of Buyer-Seller Networks. *Amer. Econ. Rev.* 91, 485–508.
- Kuniavsky, S., Smorodinsky, R., 2007. Coalitional congestion games. MSc thesis. Technion, Israel.
- Milchtaich, I., 1998. Crowding games are sequentially solvable. *Int. J. Game Theory* 27, 501–509.
- Monderer, D., Shapley, L.S., 1996. Potential games. *Games Econ. Behav.* 14, 124–143.
- Montgomery, J., 1991. Social networks and labor market outcomes. *Amer. Econ. Rev.* 81, 1408–1418.
- Moulin, H., Shenker, S., 2001. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *J. Econ. Theory* 18, 511–533.
- Papadimitriou, C.H., 2001. Algorithms, games, and the Internet. In: *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 749–753.
- Rosenthal, R.W., 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* 2, 65–67.
- Roughgarden, T., Tardos, E., 2002. How bad is selfish routing? *J. ACM* 49 (2), 236–259.
- Rozenfeld, O., Tennenholtz, M., 2006. Strong and correlated strong equilibria in monotone congestion games. Working paper. Technion, Israel.
- Vetta, A.R., 2002. Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. In: *Symposium on the Foundations of Computer Science (FOCS)*, pp. 416–425.