

Near-Strong Equilibria in Network Creation Games

Ola Rozenfeld

Technion–Israel Institute of Technology
ola.rozenfeld@gmail.com,

Moshe Tennenholtz

Microsoft R&D Labs, Israel
moshet@ie.technion.ac.il

Abstract

We introduce a new solution concept for games, near-strong equilibrium, a variation of strong equilibrium. Previous work has shown the existence of 2-strong pure strategy equilibrium for network creation games with $1 < \alpha < 2$ and that k -strong equilibrium for $k \geq 3$ does not exist. In this paper we show that 3-near-strong equilibrium exists, and provide tight bounds on existence of k -near-strong equilibria for $k \geq 4$. Then we repeat our analysis for correlated mixed strategies, where we show that, surprisingly, 3-correlated-strong equilibrium exists, and also show bounds for existence of correlated k -strong equilibria. Moreover, the equilibrium profile can be arbitrarily close to the social optimum. For both pure and correlated settings, we show examples where no equilibrium exists. On the conceptual level, our work contributes to the recent literature of extensions of strong equilibrium, while providing positive results for stability against group deviations in one of the basic settings discussed in the algorithmic game theory literature.

1 Introduction

The Nash equilibrium is a solution concept most commonly used by game theorists. A known drawback of this concept is the assumption that agents do not cooperate in order to agree on a joint deviation. The strong equilibrium (SE), introduced in (Aumann, 1959), is an extension of NE that takes care of this problem: it is a strategy profile which is stable against joint deviations by coalitions of agents. Having an SE is highly desirable, but most games of interest do not possess such stability. Also, the concept might be too strong, in a sense: it does not take into account that deviations themselves might be unstable against sub-deviations. Indeed, if we begin by supposing that an agent will deviate from a proposed profile if a better choice is available, the same should hold if the profile is proposed to him by a deviating coalition of other agents. (Bernheim *et al.*, 1987) suggested the concept of *coalition-proof Nash equilibrium* (CPNE), which captures this idea: a profile is CPNE if it has no *self-enforcing* profitable group deviations, where self-enforcing means: has no self-enforcing profitable sub-deviations (recursive definition). A unilateral deviation is always self-enforcing. Although CPNE captures the above logic perfectly, it goes too far in a sense: people in real life are extremely unlikely to employ considerations of such complexity. In a general game, it will be probably computationally impossible to determine if a given group deviation is self-enforcing. Therefore, a simpler and stronger solution concept was suggested by (Milgrom and Roberts, 1994): to require self-enforcing group deviations to be stable only against *unilateral* sub-deviations¹.

¹Independently, the same concept was introduced in (Kaplan, 1992) as *semistrong equilibrium*.

We suggest to concentrate on the following (even stronger) solution concept, which is a variation of the above definition:

A profile of actions is a *near-strong equilibrium* (NSE) if for every beneficial joint deviation by a coalition of players, there exists a player in that coalition who, given that the rest of the coalition will stick to their deviation, would strictly prefer to betray them and return *to his original strategy*.

Our definition differs from (Milgrom and Roberts, 1994) only in that the betraying player's strategy is restricted to equal his original strategy. This difference is important, because it emphasizes the stability of the original profile from the agent-centric perspective: suppose a coalition K of players wishes to convince a member i of K to participate in a joint deviation. Their arguments are: all of them (the other players in K) agreed to play their corresponding new actions, and the new actions will result in a higher payoff for all $i \in K$, compared with the original profile. In a SE, that argumentation should suffice to convince each player $i \in K$ to join. However, player i might consider the following logic: either he trusts the rest of the coalition to behave according to the chosen deviation, or he does not. If he does not trust them to behave as agreed, there is certainly no incentive for him to deviate from the equilibrium; but if he does trust them, the equilibrium strategy still gives him a strictly higher payoff than the deviation strategy! So, in both cases, player i will actually be better off not joining the coalition than joining it. So, although NSE is formally weaker than strong equilibrium, it is conceptually similar: *an agent will prefer the original profile over all possible deviations, even when coalitions of agents can coordinate a joint deviation*.

The aim of this work is to explore how the concept of NSE can improve over SE: to show an interesting setting which does not possess SE, but in which NSE can be shown to exist. The setting we chose is the network creation game, introduced by (Fabrikant *et al.*, 2003) and since appearing in many works, including (Albers *et al.*, 2006; Andelman *et al.*, 2009). The basic version of the game is as follows: each player is associated with a vertex in a network, and the aim of all players is to be connected to all the other players with as small a distance as possible to each. However, each player controls his outgoing edges and must choose which edges to buy. Buying an edge has a fixed cost, α . In the simplest model, which we adopt, a player's cost is simply a sum of his distances to all other players in the resulting (undirected) graph, and the total price he paid for his chosen edges.

When $\alpha \leq 1$, the game is not very interesting, since the socially optimal profile (a clique) is a strong equilibrium. Similarly, when $\alpha \geq 2$, the out-star (the strategy profile in which one node buys edges to all other nodes, while they buy nothing) is both socially optimal and an SE (Andelman *et al.*, 2009). However, when $1 < \alpha < 2$ even a 3-SE was shown not to exist for $n > 5$ (Andelman *et al.*, 2009); therefore, this is the case we concentrate upon in this paper.

We show:

1. The out-star is always a 3-NSE
2. An example where no 4-NSE exists
3. A bound of when the out-star is a k-NSE:
 - (a) The out-star is a k-NSE if $\alpha \geq 2(1 - \frac{1}{k'})$, where $k' = 2\lfloor \frac{k}{2} \rfloor$
 - (b) Otherwise, if $\alpha < 2(1 - \frac{1}{k'})$
 - i. For $k \leq \lfloor \frac{n+1}{2} \rfloor$, the out-star is not a k-NSE
 - ii. Otherwise, we show examples of n, α where the out-star still is a k-NSE, and other examples where it is not
4. Some empirical results for small n , e.g. the out-star is an NSE for $n \leq 6$

All the above analysis concerned only pure strategies. Next, we turn to the case when players can also use mixed strategies; since the solution concepts we deal with allow coordination by coalitions of players, it makes sense to allow the players to employ *correlated* mixed strategies. A natural extension of the strong equilibrium concept to the case of correlated mixed strategies is the correlated strong equilibrium (CSE), by (Rozenfeld and Tennenholtz, 2006).

Using correlated mixed strategies allows us to concentrate on symmetric strategy profiles – informally, profiles where all players are treated equally (and, in particular, incur the same costs). We concentrate on two intuitive symmetric strategy profiles: the ϵ -clique and the randomized out-star (which we abbreviate to "star"). For $0 < \epsilon \ll 1$ the ϵ -clique is a profile that can get arbitrarily close to a clique (with probability $1 - \epsilon$ the players form a fair clique), and the star is an out-star where the root is chosen with uniform distribution.

We show:

1. The ϵ -clique is a k -CSE for $k < \min\{n, 1 + \frac{\alpha}{\alpha-1}\}$. In particular, this implies:
 - The ϵ -clique is a 3-CSE for $n \geq 4$.
 - The ϵ -clique is a 4-CSE for $\alpha < \frac{3}{2}, n \geq 5$.
2. We derive a tight bound on when the star is a k -CSE
 - In particular, it implies that the star is a 4-CSE for $\alpha \geq \frac{3}{2}, n \geq 10$
3. We present a sound algorithm for proving that no symmetric k -CSE exists for a given instance
 - In particular, using the algorithm we show that no symmetric 4-CSE exists for $n = 5, \alpha = \frac{3}{2}$
4. Monotonicity:
 - (a) If the ϵ -clique is a k -CSE for n, k, α , then for all $\alpha' < \alpha$ it is still a k -CSE for n, α'
 - (b) If the star is a k -CSE for n, k, α , then for all $\alpha' > \alpha$ it is still a k -CSE for n, α'

Recall that the clique is the social optimum in our setting ($1 < \alpha < 2$). Using only pure strategies, there was no hope to implement any profile with sufficiently many edges as even a NE. Surprisingly, when we allow correlated mixed strategies, a strong positive result emerges: we can implement a fair profile arbitrarily close to the social optimum as a k -CSE. Note also that the ϵ -clique gets more stable for smaller values of α , which is very good, since it is for small α that the clique yields a much lower social cost than the star. For big values of α , the social cost of the star approaches that of the clique; and in these cases, too, the closer to optimal the star gets, the more stable it becomes.

We can now extend our near-strong equilibrium concept in a similar way to define a correlated near-strong equilibrium (CNSE). Recall that the aim of this work is to compare the existence of NSE to that of SE. However, in the correlated setting it turns out that the concept of CNSE is just as strong as CSE: we show that a symmetric strategy profile is a CNSE if and only if it is a CSE.

Note: due to lack of space, some of the proofs in the following sections were omitted.

2 Model and Preliminaries

The network creation game was introduced in (Fabrikant *et al.*, 2003). A player is associated with a vertex in a network, who wishes to connect to other players. The set of players is $V = \{1, \dots, n\}$, and a strategy of a player is to select the subset of other players to whom he buys an edge: $S_v = 2^{V \setminus \{v\}}$.

For a set of players $K \subseteq V$, let $S_K = \prod_{v \in K} S_v$, and let $S = S_V$. A strategy profile $s \in S$ induces a directed graph $G(s) = (V, E)$, where $E = \{(v, u) | u \in s_v\}$.

The cost that a player incurs consists of two parts: the price of the edges he bought (each edge has a fixed cost of α) and his distances to the other players in the resulting network. Formally, $c_v(s) = \alpha |s_v| + \text{Dist}(v)$, where $\text{Dist}(v) = \sum_{u \in V} \delta_s(v, u)$, and $\delta_s(v, u)$ is the length of the shortest path from v to u in the *undirected* graph induced by $G(s)$. So, the directions of the edges serve only to visualize which player is paying for them; network-wise, the network is treated as an undirected graph. A player has to be connected to all other players; otherwise, $c_v(s) = \infty$. *For the remainder of this work, we assume that $1 < \alpha < 2$.*

We recall that a profile of actions $s \in S$ is a *strong equilibrium* (SE) if for every coalition $K \subseteq V$ and every joint choice of actions $t_K \in S_K$ there exists $v \in K$ for whom $c_v(t_K, s_{-K}) \geq c_v(s)$ (Aumann, 1959). For $1 \leq k \leq n$ we say that $s \in S$ is a *k-strong equilibrium* (k-SE) if the above condition holds for all $K \subseteq V$ s.t. $|K| \leq k$.

A profile of actions $s \in S$ is a *near-strong equilibrium* (NSE) if for every coalition $K \subseteq V$ and every joint choice of actions $t_K \in S_K$ such that $\forall v \in K \ c_v(t_K, s_{-K}) < c_v(s)$, there exists $v \in K$ for whom $c_v(t_{K \setminus \{v\}}, s_{-K \cup \{v\}}) < c_v(t_K, s_{-K})$. For $1 \leq k \leq n$ we say that $s \in S$ is a *k-near-strong equilibrium* (k-NSE) if the above condition holds for all $K \subseteq V$ s.t. $|K| \leq k$.

For a set A , let $\Delta(A)$ denote the set of all probability distributions over A . For a correlated strategy profile $s \in \Delta(S)$, let $C_v(s)$ denote the expected cost of player v in s . We say that $s \in \Delta(S)$ is a *correlated strong equilibrium* (CSE) if for every coalition $K \subseteq V$ and a deviation $t_K \in \Delta(S_K)$ there exists $v \in K$ s.t. $C_v(t_K \times s_{[-K]}) \geq C_v(s)$ (Rozenfeld and Tennenholtz, 2006). Here, $s_{[-K]}$ means the marginal probability induced by s on $V \setminus K$. For $1 \leq k \leq n$ we say that $s \in S$ is a *correlated k-strong equilibrium* (k-CSE) if the above condition holds for all $K \subseteq V$ s.t. $|K| \leq k$. Note that this definition means that the deviating players should be able to gain in expectation, without knowing *anything* about the pure realization of s .

A justification of the above version of the definition of CSE is, perhaps, in order. The definition of a correlated Nash equilibrium (Aumann, 1974) presumed that before deciding on a deviation each player could observe his *signal* (his action in the pure realization of the correlated strategy). When defining a correlated strong equilibrium, the first instinct is to extend Aumann's definition of correlated Nash equilibrium (CNE) and require the profile to be stable against *ex-post* deviations by coalitions of players. However, there is a caveat in doing this. In Aumann's definition, it is vital that each player is informed only about *his own* signal. The underlying assumption of this model could be, for example, that a trusted authority rolls the dice, and sends to each player his selected strategy over a private channel. But how would this work for a coalition of players, who need to deviate jointly? Presumably, they would have to share their signals. But here is the problem: if they do share signals, then each one of them *by himself* possesses more information than he is allowed to by the CNE concept! We would have to consider situations such as this: a single agent cannot beneficially deviate knowing his own possible signals, but two agents can each beneficially deviate alone if each also gets to know the other's signal – even though the two of them together do not possess a joint deviation beneficial for both! No matter how we classify such cases, the justification is not at all intuitive. We are not saying that such definition is impossible or not interesting (on the contrary, we are working on it), but we have to be careful and visualize, in detail, the flow of information in the situation that we are trying to model.

Our definition, much like Aumann's, presumes the existence of a mediator – a trusted third party that can roll the dice to select a pure realization of a correlated strategy profile. In our model, the mediator does not output the signals at all – rather, it selects actions according to a preset correlated profile *on behalf* of the players who chose to use it (much like in (Rozenfeld and Tennenholtz, 2007), only without punishments). Alternatively, one could imagine some other means to enforce the players to follow through on the selected pure realization, once it is chosen

(e.g. a contract (Kalai *et al.*, 2007)). In the following, we assume that such means are available to any coalition of players wishing to implement a joint correlated strategy.

A profile $s \in \Delta(S)$ is a *correlated near-strong equilibrium* (CNSE) if for every coalition $K \subseteq V$ and a deviation $t_K \in \Delta(S_K)$ s.t. $\forall v \in K \ C_v(t_K \times s_{[-K]}) < C_v(s)$ there exists $v \in K$ for whom $C_v(t_{[K \setminus \{v\}] \times s_{[-K \cup \{v\}]}}) < C_v(t_K \times s_{[-K]})$. For $1 \leq k \leq n$ we say that $s \in S$ is a *correlated k-near-strong equilibrium* (k-CNSE) if the above condition holds for all $K \subseteq V$ s.t. $|K| \leq k$.

Now we want to define a symmetric profile. Intuitively, a strategy profile is symmetric if the names of the players don't matter. Formally, we will need some notation. Let π be a permutation of V . For a set $K \subseteq V$, we let $\pi(K) = \{\pi(v) | v \in K\}$. For a strategy profile $s \in S$, we denote by $\pi(s)$ the following strategy profile: for all $v \in V$, $\pi(s)_{\pi(v)} = s_v$. For a correlated strategy profile $s \in \Delta(S)$, we denote by $\pi(s)$ the following correlated strategy profile: for all $z \in S$ $\pi(s)(\pi(z)) = s(z)$. Let Π be the set of all possible permutations of V . For a correlated strategy profile $s \in \Delta(S)$, let $Sym(s) \in \Delta(S)$ be the following correlated strategy: first select $\pi \in \Pi$ with uniform probability, and then play $\pi(s)$. We say that a correlated strategy profile $s \in \Delta(S)$ is *symmetric* if $s = Sym(s)$.

3 Pure Strategies

In (Andelman *et al.*, 2009) it was shown that a 2-SE always exists, but even a 3-SE does not exist for $n > 5$. In particular, consider the following strategy profile o^* (the out-star):

Let one player (the root, denoted by r) purchase an edge to every other player, while all other players (which we call leaves) purchase no edges. In this profile, the cost of the root, $c_r(o^*) = \alpha(n-1) + n - 1 = (n-1)(\alpha+1)$ (he pays for $n-1$ edges, and his distance to any other node is 1). The cost of every other player is $c_v(o^*) = 1 + 2(n-2) = 2n-3$ (his distance to r is 1, and his distance to any other node is 2).

It is easy to see that o^* is a NE: $1 < \alpha < 2$ means that no player wishes to purchase an edge in order to decrease a distance to a single player from 2 to 1. By enumerating the few possible cases one can also verify that the out-star is a 2-SE. The reason the out-star is not a 3-SE is that any three leaves can deviate to the following strategy: they form a triangle in which each one of them purchases one edge. For each one of them, the deviation decreased his distances from the two other deviators by 1, decreasing his overall distance cost by 2; since his edge cost increased by $\alpha < 2$, the player strictly gains from the deviation.

But is this deviation stable? "Betraying" the deviators and returning to the original profile simply means buying no edges at all. Clearly, such strategy is more beneficial than the deviation – dropping the edge (v, u) decreases the edge cost of v by α while increasing his distance to u from 1 to 2, not affecting his distances to other players. So this deviation is not stable. As we will now show, this holds for every other beneficial deviation as well:

Lemma 1 *The out-star is a 3-NSE.*

In order to prove that, we first need the following:

Lemma 2 *Any stable deviation of k players from the out-star must include the root.*

Proof: Suppose for contradiction that r does not deviate, and let v be any node purchasing at least one edge in the deviation. Since r still buys an edge to every node, v has a path of length 2 to any other node, regardless of his chosen strategy. Therefore, each edge (v, u) decreases v 's total distance cost by exactly 1, while increasing his edge cost by $\alpha > 1$. Therefore, v will be strictly better off not joining the deviation. ■

Now we can prove lemma 1:

Proof: By lemma 2, the set of deviators is $K = \{r, v, u\}$. Suppose they have a beneficial deviation, $s_K \in S_K$. Now v has to purchase at least one edge – otherwise, o^* is not a 2-SE. Since $\alpha > 1$, v has to reduce his total distance cost by at least 2 in order to gain from the deviation. That means someone (other than r) has now to purchase an edge to v . That only leaves u . But then, by the same logic, u cannot gain from the deviation (there is no one left to purchase an additional edge to u). ■

Is the out-star a 4-NSE? In general, no. Let $n = 7$, $\alpha = 1.1$. Consider the deviation depicted in Fig.1:




Figure 1: Here, $n = 7$, $\alpha = 1.1$, and $K = \{0, 1, 2, 3\}$

Here, the root (1) benefits from the deviation, because he now has a distance of 2 to the nodes 5, 6 and 7, instead of buying a direct link, so he saves $3(\alpha - 1) > 0$. Obviously, he would not benefit from returning to the original profile, where he buys edges to all nodes. The other deviators purchase two edges each, while decreasing their distances by 3. Since $\alpha < 1.5$, they benefit from the deviation; since each one of them is responsible for connecting a node to the graph, he would suffer a cost of ∞ if returned to play his original strategy (not buying any edges). Therefore, the deviation is stable.

However, for $\alpha \geq 1.5$, the out-star is a 4-NSE. In general, we can show:

Theorem 3 *Let $2 < k \leq n$, and $k' = 2\lfloor \frac{k}{2} \rfloor$. If $\alpha \geq 2(1 - \frac{1}{k'})$, the out-star is a k -NSE.*

Thm. 3 raises two questions. Firstly, is the bound tight? Secondly, suppose the out-star is not a k -NSE. But maybe some other profile is? Are there examples where it can be proved that no k -NSE exists?

The following results address the first question:

Theorem 4 *Let $4 \leq k \leq \lfloor \frac{n+1}{2} \rfloor$, and let $\alpha < 2(1 - \frac{1}{k'})$, with $k' = 2\lfloor \frac{k}{2} \rfloor$. Then the out-star is not a k -NSE.*

Proof: It is sufficient to show a beneficial and stable deviation for a coalition of size k' . Let us denote the root by r and let $K' = K \setminus \{r\}$. The deviation is a generalized version of Fig. 1, and is defined as follows:

- The nodes in K' buy a fair clique between themselves; that is, each buys exactly $(k' - 2)/2$ edges to other deviators and they form a clique. It is possible since the clique size $(k' - 1)$ is odd.
- Each member of K' is assigned a distinct node in $N \setminus K$, and buys an edge to it. This is possible, because $2|K'| \leq n - 1$.
- The root buys edges to all members of K' , and buys edges to all the nodes in $N \setminus K$ that are not connected to a node in K' .

Let us show that the deviation is profitable. For r , the difference is that he dropped the edges to some of the nodes in $N \setminus K$ (saving α from each), and as result, his distance to these nodes increased to 2. Since $\alpha > 1$, he gained from the deviation. Each $v \in K'$ gained $k' - 2 + 1 = k' - 1$ in distances (from forming a clique and having a direct connection to a node in $N \setminus K$). He paid $\alpha \left(\frac{k'-2}{2} + 1 \right) = \alpha \frac{k'}{2} < k' - 1$ for his edges. His distance to r stayed 1 and his distances to $N \setminus K$ stayed 2, therefore he gains by deviating. Now, let us show that no one would benefit from going back to the original strategy. For r , we have already seen that he is worse off if he buys edges to all the nodes. For $v \in K'$, dropping his edges leaves him disconnected from his appropriate node in $N \setminus K$, so he cannot deviate. ■

So for $k \leq \lfloor \frac{n+1}{2} \rfloor$, the bound of Thm. 3 is tight. Otherwise, it is tight for some cases and not tight for others:

Example 1 Let $n = 15, k = 11, \alpha < \frac{9}{5}$. Suppose that $r = 1$. Consider the following deviation s by $\{1, \dots, 11\}$: $s_1 = \{2, \dots, 11\}$

The 9 players $\{2, \dots, 10\}$ form a fair clique, with each player purchasing 4 edges, and in addition they all point to 11. 11 points to $V \setminus K$:

$$\begin{array}{ll} s_2 = \{3, 4, 5, 6, 11\} & s_6 = \{7, 8, 9, 10, 11\} \\ s_3 = \{4, 5, 6, 7, 11\} & s_7 = \{8, 9, 10, 2, 11\} \\ s_4 = \{5, 6, 7, 8, 11\} & s_8 = \{9, 10, 2, 3, 11\} \\ s_5 = \{6, 7, 8, 9, 11\} & s_9 = \{10, 2, 3, 4, 11\} \\ s_{10} = \{2, 3, 4, 5, 11\} & s_{11} = \{12, 13, 14, 15\} \end{array}$$

One can easily verify that the deviation is beneficial and stable. Here, unlike in the proof of Thm. 4, the deviation is stable not because each deviator is responsible for connecting a node to the graph (that holds only for node 11), but because dropping the edge to 11 will increase the player's total distances by 5, making betrayal not beneficial for him for $\alpha < \frac{9}{5}$.

So in this case, the bound of Thm. 3 is tight.

Example 2 Let $n = 7, k = 7, \frac{3}{2} \leq \alpha < \frac{5}{3}$. Here, the out-star is an NSE.

Proof: By checking all possible deviations with NSESAT (see below). ■

So in this case the bound of Thm. 3 is not tight.

NSESAT: During the course of this work we often ran into the problem of classifying small instances. Even for number of agents as small as 5, it was often not easy to answer questions such as "Is a given profile a k-NSE (or k-SE)?" The question of existence of k-NSE is even more difficult, since even for very small instances there is a huge number of possible profiles to check. Therefore we developed a computer program for these tasks, called NSESAT. The main idea was to reduce the problem of finding a stable beneficial deviation of at most k players to SAT, and then use a known SAT solver to solve it (we used MiniSat, introduced in (Eén and Sörensson, 2003)). We also added some simple optimizations to break symmetry in symmetric strategy profiles. Despite the fact that the total number of deviations of k players in an n vertex graph is order of $2^{(n-1)k}$, MiniSat handled the decision problem extremely well, allowing us to solve instances for n as high as 11! Using NSESAT, we quickly established the following fact:

Fact 5 The out-star is an NSE for $n < 7$.

But this fact raised a problem – it meant that in order to find an example where 4-NSE does not exist, we had to start with n of at least 7. The number of non-isomorphic directed graphs on 7

nodes is 882033440 (Sloane, N. J. A. Sequence A000273 in "The On-Line Encyclopedia of Integer Sequences"). Even with MiniSat deciding on each graph in under 1 second, the straightforward approach of checking all possible graphs was infeasible. Fortunately, we could use the following two necessary conditions for Nash equilibria to narrow the search:

1. The diameter of the undirected graph has to equal 2 (otherwise, if it equals 1, an agent will want to drop an edge, and if there exists a shortest path of at least 3, an agent will want to add an edge).
2. Lemma B.2 in (Andelman *et al.*, 2009) provided a structural property of NE for $1 < \alpha < 2$, which could be efficiently verified on the underlying undirected graph.

Using the above conditions, we were able to start with all the simple undirected graphs of 7 nodes (1044) and reduce the number of potentials for 4-NSE to only 46. All the possible ways to direct these graphs gave us an order of 120000 strategy profiles, which we were able to check in under 10 hours. Our conclusion:

Fact 6 *For $n = 7, \alpha = 1.1$ a 4-NSE does not exist.*

Some additional facts that we were able to establish using NSESAT are: for $\alpha \geq 1.5$ the out-star is an NSE for $n \leq 10$. For $n = 11$, the out-star is an NSE for $\alpha \geq \frac{5}{3}$. There are examples where the out-star is not an NSE, but another profile is (such as the in-star, where all leaves buy edges to the root).

Overall, it appears impossible to achieve a better bound on existence of NSE than the one we have with the out-star. Which is unfortunate; not because the bound is not good enough, but rather because the out-star has several serious drawbacks, which would probably make it impractical as recommended strategy profile. Firstly, the profile is extremely unfair – one agent has to incur all the costs of the network. Secondly, its social cost is very high – with $1 < \alpha < 2$, the clique is the social optimum, and a star is the worst NE possible.

In order to address the first issue, it is intuitive to consider correlated mixed strategies (if the players can roll a dice to choose the root, then, at least in terms of expected cost, the profile will be fair). As we will show in the next section, surprisingly, allowing for correlated mixed strategies addresses the second problem as well – we will often be able to implement nearly optimal fair strategy profiles as correlated strong equilibrium.

4 Correlated Mixed Strategies

Let us again recall the problem with implementing good social outcomes as strong equilibrium. Suppose we recommend the agents to form a clique. What is the best response of a single agent? It is to buy no edges at all. Since some of the other agents still connect to him, dropping his edges only increases his corresponding distances from 1 to 2, but saves him α . But what if the agents form a *randomized clique*? A randomized clique (denoted by r-#) is a symmetric correlated strategy profile where all (undirected) edges are bought with probability 1, and for each edge (u, v) , the buyer of the edge (u or v) is selected independently and uniformly. Recall that in our model, a deviating agent does not possess the "signal" to his selected strategy; he only knows what randomized strategy the other players will play.

Not buying any edges is no longer a best response of a single agent u to a randomized clique. Why not? Buying an edge (u, v) is now even less beneficial for u than in a pure clique – now with probability 0.5 v will buy the edge himself, so the *expected distance* to v if no edge is bought is 1.5 instead of 2. It is clearly better to lose 0.5 (increasing the distance from 1 to 1.5) than to spend

α . However, there is one problem. Since the buyer of each edge is selected independently, there is a small chance that u is selected to buy all his edges. Since the cost of not being connected to another agent is ∞ , then, no matter how small the chance, u cannot afford to take it! He must purchase at least one edge (however, one edge will suffice – the randomized clique is not a 1-CSE).

Let us now define a profile we call the randomized out-star (or, simply, the star, denoted by r^*): the agents select one agent uniformly to be the root. The root purchases edges to all other agents, while the other agents purchase nothing.

Let $0 < \epsilon \ll 1$. The ϵ -clique is a symmetric correlated strategy profile (denoted by $\epsilon\text{-}\#$) defined as follows: with probability $1 - \epsilon$, the randomized clique is played, and with probability ϵ , the star is played.

Theorem 7 *For $k < \min\{n, 1 + \frac{\alpha}{\alpha-1}\}$, there exists $\epsilon > 0$ s.t. $\epsilon\text{-}\#$ is a k -CSE.*

Proof: Firstly, note that $\epsilon\text{-}\#$, as well as any other symmetric profile with less than optimal social cost, can never be an n-CSE, because the n players could always deviate to the optimum, $r\text{-}\#$. (But, in a sense, this observation does not take away from the result – the whole point is to implement a socially good outcome, so if the only problem with it is that the agents can deviate to another outcome that is even better for everyone, then it is not really a problem).

Let $k < n$, and K be the set of the deviating agents. Since with positive probability, $\epsilon(1 - \frac{k}{n}) > 0$, the agents in $V \setminus K$ do not buy any edges, any profitable deviation of K must include links to all the agents in $V \setminus K$ (since the deviators do not possess the signals to the chosen pure realization of the correlated strategy, they must consider the worst case). This is the key idea that allows us to show that the agents will not be able to reduce their total cost.

Let us denote by C_{OPT} the summary cost of k agents in a fair clique:

$$C_{OPT} = \alpha \frac{k(n-1)}{2} + k(n-1)$$

The deviation that minimizes the total cost of the agents (their collective best response strategy) is to form a clique among themselves and to purchase an edge to all of $V \setminus K$. Since we are interested in their summary costs, assume w.l.o.g. that a single agent buys all the edges to $V \setminus K$ (obviously, the agents can easily share these costs by selecting this agent uniformly). Let us denote this deviation by g^k , and let us denote by C_{DEV} the total cost the deviating agents incur when playing g^k when the other agents play $r\text{-}\#$:

$$C_{DEV} = \alpha \left(\frac{k(k-1)}{2} + n - k \right) + n - 1 + (k-1)(k-1 + 1.5(n-k))$$

Here, $n-1$ are the total distances of the agent who bought the additional edges, and $k-1 + 1.5(n-k)$ are the total distances of every other deviator. Recall that 1.5 is the expected distance to any node in $V \setminus K$, since they play $r\text{-}\#$. The summary cost of K using the above deviation against the $\epsilon\text{-}\#$ will be slightly higher than C_{DEV} . But, as we will now show, $C_{DEV} > C_{OPT}$:

$$\begin{aligned} C_{DEV} - C_{OPT} &= \alpha \left(\frac{k(k-1)}{2} + n - k \right) + n - 1 + (k-1)(k-1 + 1.5(n-k)) - \alpha \frac{k(n-1)}{2} + k(n-1) = \\ &= \frac{\alpha}{2}(n-k)(2-k) + \frac{1}{2}(n-k)(k-1) = \frac{1}{2}(n-k)(\alpha(2-k) + k-1) > 0 \quad \Leftrightarrow \\ &\quad \alpha(2-k) + k-1 > 0 \quad \Leftrightarrow \quad k < 1 + \frac{\alpha}{\alpha-1} \end{aligned}$$

Hence, there exists $\epsilon > 0$ s.t. the payoff of k players in $\epsilon\text{-}\#$ will be below C_{DEV} , and therefore, the players will not have a beneficial joint deviation. ■

In particular, $\epsilon\text{-}\#$ is a 3-CSE for all $n \geq 4$, and 4-CSE for $n \geq 5, \alpha < 1.5$. Note that the bound on k increases as α approaches 1, meaning that $\epsilon\text{-}\#$ can be implemented as a more stable equilibrium. Formally,

Corollary 8 *If the ϵ -clique is a k -CSE for n, k, α , then for all $\alpha' < \alpha$ it is still a k -CSE for n, α' .*

But what about when α is big? Turns out that in these cases r^* (the randomized star) becomes a k -CSE.

Let us denote by C_* the total cost of any k players in r^* :

$$C_* = \frac{k}{n}(\alpha(n-1) + n - 1 + (n-1)(1 + 2(n-2))) = \frac{k(n-1)}{n}(\alpha + 2n - 2)$$

Let us denote by C_{DEV^*} the total cost of k players in their joint best response deviation (which is g^k , the same as against $r\text{-}\#$):

$$\begin{aligned} C_{DEV^*} &= \alpha\left(\frac{k(k-1)}{2} + n - k\right) + \frac{k}{n}(n-1 + (k-1)(k-1 + 2(n-k))) + \frac{n-k}{n}(n-1 + (k-1)(k-1 + 2(n-k)) - (k-1)) = \\ &= \frac{\alpha}{2}(k^2 - 3k + 2n) - k^2 - n + 2nk - k + 1 + \frac{k(k-1)}{n} \end{aligned}$$

Theorem 9 *The star, r^* , is a k -CSE if and only if $C_{DEV^*} \geq C_*$.*

Unfortunately, unlike in Thm. 7, the bound cannot be simplified to eliminate one of the variables. However, it can be used to derive various bounds for fixed values of one or two variables. For example, it is simple to see that for $\alpha \geq 1.5$ and $n \geq 10$ the star is a 4-CSE. Also, it is easy to derive the following result:

Corollary 10 *If the star is a k -CSE for n, k, α , then for all $\alpha' > \alpha$ it is still a k -CSE for n, α' .*

Note that when α approaches 2, the social cost of star approaches the social optimum; so Thms. 7 and 9 together imply that as α approaches its bounds, it becomes easier to implement a near-optimal outcome; the intermediate values of α are the most problematic. For example, for $\alpha = 1.5, 4 < n < 10$ neither $\epsilon\text{-}\#$ nor r^* are a 4-CSE. But what about other profiles? Can we prove that no 4-CSE exists? We will now show a way to do this for symmetric profiles.

For $s \in \Delta(S)$, let $edges(s)$ denote the expected number of bought edges in s : $edges(s) = \sum_{z \in S} s(z) \sum_v |z(v)|$. Let $maxEdges = \frac{n(n-1)}{2}$. Let $C_G(s)$ denote the total cost of k agents when they deviate to g^k from a symmetric correlated strategy profile s .

Lemma 11 $C_G(s) = \alpha\left(\frac{k(k-1)}{2} + n - k\right) + n - 1 + (k-1)(k-1 + (2 - 0.5p)(n-k))$, where $p = \frac{edges(s)}{maxEdges}$.

Proof: Consider the expected distance from any node $v \in K$ who does not buy the $n - k$ additional edges to any node $u \in V \setminus K$. We need to show that $E_{s_{[-K]}} \delta(v, u) = 2 - 0.5p$, and the result will follow. The profile is symmetric, therefore the edge (v, u) is bought by u with probability $0.5p$, in which case the distance is 1, and is not bought with probability $1 - 0.5p$, in which case the distance is 2. On average, it is $0.5p + 2(1 - 0.5p) = 2 - 0.5p$. ■

Since $C_G(s)$ depends only on $edges(s)$, we will abuse notation and denote it by $C_G(t)$, with $t = edges(s)$.

Lemma 12 *For $t \in [n - 1, maxEdges]$, let $v(t) = \frac{n}{k}(C_G(t) - C_{OPT})$, and $t' = maxEdges - \frac{v(t)}{2 - \alpha}$. If $s \in \Delta(S)$ is a symmetric k -CSE then:*

1. $v(edges(s)) > 0$
2. $edges(s) \geq t'$

Proof:

1. Follows from the fact that the total cost of k players in s is strictly above C_{OPT} (since, as we know, the clique is not a k -CSE), and s is a k -CSE (therefore k players cannot beneficially deviate to $Sym(g^k)$).
2. Let $t = edges(s)$. Since the deviation to $Sym(g^k)$ is not beneficial to K , $C_v(s) \leq \frac{C_G(t)}{k} = \frac{C_{OPT}}{k} - \frac{v(t)}{n}$, and therefore $\sum_{v \in V} C_v(s) \leq OPT - v(t)$ (where OPT is the optimal social cost). Each edge removed from the optimal profile (the clique) increases the social cost by at least $2 - \alpha$. Therefore, in order to satisfy the cost bound, s must remove at most $\frac{v(t)}{2-\alpha}$ edges. ■

Now, suppose we want to prove that no symmetric k -CSE exists for given n, k, α . We know that any CSE has to have at least $t_0 = n - 1$ edges, because the graph has to be connected; therefore, we can apply Lemma 12 and derive a new lower bound, t_1 , for the expected number of edges in a symmetric k -CSE. If $t_1 \leq t_0$, we don't have any additional information, and the process stops (we cannot prove anything). But if $t_1 > t_0$, we can apply Lemma 12 again! This gives us a new bound, t_2 . So, we can continue the process and keep deriving lower bounds t_3, \dots, t_j, \dots for $edges(s)$ in any symmetric k -CSE s . If, at any point in this process, it holds that $C_G(t_j) \leq C_{OPT}$, we have our proof – a symmetric k -CSE does not exist (by Lemma 12, part 1). Similarly, if $t_j \leq t_{j+1}$, the process stops and we cannot prove anything. The only other option is that the series t_j converge. In this case, if $\lim_{j \rightarrow \infty} t_j = maxEdges$, we have our proof – no symmetric CSE exists (because this means that no profile s with $edges(s) < maxEdges$ can be a k -CSE; and we already know that the clique is not a k -CSE).

Some empirical results: for $5 \leq n \leq 20, 4 \leq k \leq n$, and $\alpha \in 1, 1.05, \dots, 1.95$, we tested which of the following holds in each case: $\epsilon\#$ is a k -CSE, r^* is a k -CSE, both of these profiles are k -CSE, or no symmetric k -CSE exists. In all these runs, we have not encountered a case where neither $\epsilon\#$ nor r^* were a k -CSE, but the iterative proof that no k -CSE exists failed. These empirical results strongly suggest the following:

Conjecture: For every n, k, α , if a symmetric k -CSE exists, then at least one of $\epsilon\#, r^*$ is a k -CSE.

Also, although the assumption of symmetry was crucial to our negative result, we conjecture that in this setting, the existence of k -CSE implies the existence of a symmetric k -CSE. Whether this conjecture is correct is an interesting open question for future work.

Now that we have explored the existence of (symmetric) k -CSE, we would like to find out whether k -CNSE might exist in cases where k -CSE does not. Unfortunately, the following result implies otherwise:

Lemma 13 *Let $s \in \Delta(S)$ be a symmetric strategy profile and let $t_K \in \Delta(S_K)$ be a profitable deviation by players $K \subseteq V$. Then there exists a deviation $q_K \in \Delta(S_K)$ which is both profitable and stable.*

Proof: Let $q_K \in \Delta(S_K)$ be defined as follows: with probability $1 - \epsilon$ the players play t_K , and with probability ϵ they play the following profile: choose a root $v \in K$ uniformly; the root buys edges to all the other nodes; nodes in $K \setminus \{v\}$ buy no edges. $0 < \epsilon < 1$ is selected so that $C_v(q_K \times s_{[-K]}) < C_v(s)$ for any $v \in K$ (this is possible, since $C_v(t_K \times s_{[-K]}) < C_v(s)$). Since every node in K is now responsible for connecting himself to other nodes with positive probability, the deviation is stable (the only case where $v \in K$ can betray without incurring a cost of ∞ is if he purchases all edges to $K \setminus \{v\}$ in s with probability 1; due to symmetry of s , this is impossible). ■

The proof uses the same idea that allowed us to implement good outcomes as k -CSE – using a small probability of having a player disconnected from the graph. An interesting idea for future

work is to try and work with a more realistic model, where a player does not incur a cost of ∞ for being disconnected.

References

- [Albers *et al.*, 2006] Susanne Albers, Stefan Eilts, Eyal Even-dar, Yishay Mansour, and Liam Roditty. On nash equilibria for a network creation game. In *In Proc. of SODA*, pages 89–98, 2006.
- [Andelman *et al.*, 2009] Nir Andelman, Michal Feldman, and Yishay Mansour. Strong price of anarchy. *Games and Economic Behavior*, 65(2):289–317, 2009.
- [Aumann, 1959] R.J. Aumann. Acceptable points in general cooperative n-person games. In A.W. Tucker and R.D. Luce, editors, *Contribution to the Theory of Games, Vol. IV, Annals of Mathematics Studies, 40*, pages 287–324. 1959.
- [Aumann, 1974] R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- [Bernheim *et al.*, 1987] B. D. Bernheim, B. Peleg, and M. Whinston. Coalition proof nash equilibrium i: Concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.
- [Eén and Sörensson, 2003] Niklas Eén and Niklas Sörensson. An extensible sat-solver. In Enrico Giunchiglia and Armando Tacchella, editors, *SAT*, volume 2919 of *Lecture Notes in Computer Science*, pages 502–518. Springer, 2003.
- [Fabrikant *et al.*, 2003] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker. On a network creation game. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 347–351, New York, NY, USA, 2003. ACM.
- [Kalai *et al.*, 2007] Adam Tauman Kalai, Ehud Kalai, Ehud Lehrer, and Dov Samet. Voluntary commitments lead to efficiency. Discussion Papers 1444, Northwestern University, Center for Mathematical Studies in Economics and Management Science, April 2007. available at <http://ideas.repec.org/p/nwu/cmsems/1444.html>.
- [Kaplan, 1992] G Kaplan. Sophisticated outcomes and coalitional stability. M.sc. thesis, Dept. of Statistics, Tel Aviv University, 1992.
- [Milgrom and Roberts, 1994] Paul Milgrom and John Roberts. Strongly coalition-proof equilibria in games with strategic complementarities. Working Papers 95002, Stanford University, Department of Economics, 1994.
- [Rozenfeld and Tennenholtz, 2006] Ola Rozenfeld and Moshe Tennenholtz. Strong and Correlated Strong Equilibria in Monotone Congestion Games. In *Proc. of 2nd international Workshop on Internet & Network Economics*, 2006.
- [Rozenfeld and Tennenholtz, 2007] Ola Rozenfeld and Moshe Tennenholtz. Routing Mediators. In *Proc. 20th International Joint Conference on Artificial Intelligence*, 2007.