Theoretical Economics 13 (2018), 1077–1110 1555-7561/20181077

Competition and networks of collaboration

Nikita Roketskiy
Department of Economics, University College London

I develop a model of collaboration between tournament participants in which agents collaborate in pairs, and an endogenous structure of collaboration is represented by a weighted network. The agents are forward-looking and capable of coordination; they value collaboration with others and higher tournament rankings. I use von Neumann–Morgenstern stable sets as a solution. I find stable networks in which agents collaborate only within exclusive groups. Both an absence of intergroup collaboration and excessive intragroup collaboration lead to inefficiency. I provide a necessary and sufficient condition for the stability of efficient outcomes in winner-takes-all tournaments. I show that the use of transfers does not repair efficiency.

Keywords. Network, collaboration, farsighted agent, stable set, tournament.

JEL classification. C71, D85.

1. Introduction

We often observe collaboration between direct competitors. For instance, firms that compete in the market for a final product frequently collaborate at the research and development (R&D) stage. Similarly, co-workers who compete for a promotion collaborate with their rivals. Agents in these environments face a dilemma: If they collaborate, they become stronger competitors, but they also strengthen their rivals' positions.

Under what conditions do competitors collaborate efficiently? And if those conditions do not hold, what are stable patterns of collaboration? Does competition suppress collaboration, and if it does, do agents use transfers to exchange utility for collaboration and restore efficiency?

I address these questions with a model in which an endogenous structure of collaboration is represented by a weighted network; i.e., I assume that a quantum of collaboration is a *bilateral* interaction. I restrict my attention to situations in which competition can be modeled as a *tournament*. In a tournament, a higher level of collaboration,

---

Nikita Roketskiy: n.roketskiy@ucl.ac.uk

This paper is a revised version of a chapter of my Ph.D. dissertation at NYU.
I am extremely grateful to my advisors, Alessandro Lizzeri and Debraj Ray, for their guidance, support, and encouragement. I thank (in alphabetical order) Heski Bar-Isaac, Joyee Deb, Guillaume Fréchette, Sergei Izmalkov, Philippe Jehiel, Melis Kartal, Sergei Kovbasyuk, Guy Laroque, John Leahy, Stefania Minardi, Leonardo Pejsachowicz, Michael Richter, Ariel Rubinstein, Karl Schlag, Ennio Stacchetti, Kevin Thom, Matan Tsur, and all seminar participants at NYU, UCL, HEC, UPF, CEU, UC3M, Cambridge-INET Workshop, Vienna Joint Economics seminar, Max Planck Institute for Research on Collective Goods, and UT Austin for their insightful comments and suggestions. I also thank a co-editor and three anonymous referees for their thoughtful reports.

 https://doi.org/10.3982/TE2492

measured by the number of collaborative partners and the intensity of the collaborative interaction, results in better performance and, therefore, a higher tournament ranking.

In my model, a finite population of identical agents participates in a tournament. Each agent may exert an effort to collaborate with any opponents of his choice. The collaboration is nonexclusive, and if the agent chooses a higher collaborative effort, his performance improve only if his collaborative partner reciprocates the effort. Once all of the collaboration has taken place, all agents are ranked according to their output, which is increasing in their reciprocated collaborative efforts. Agents value their output directly and indirectly through their preferences for higher tournament ranks.

I focus on a protocol-free formation of a collaborative network. To model this process, I take a cooperative route: I look at all possible suggestions that agents can *collectively* make and test them against the possible objections of other agents. This process results in stable sets of outcomes (networks of collaboration) that are immune to objections. Requiring that a stable outcome be immune to all objections is too strong, so I require only that a stable outcome be immune to objections that lead to other stable outcomes. Formally, I study von Neumann–Morgenstern stable sets of outcomes defined for a farsighted blocking relation.

My findings are threefold. First, I find stable networks of collaboration that have a group structure. When tournament prizes are large enough, agents are endogenously divided into several groups. Generally, agents collaborate at an excessively high level within each group, but collaboration across groups is absent. Put differently, these groups form complete components. Any complete component is strictly larger in size than a union of all complete components that are smaller in size. In particular, the largest complete component always contains a strict majority of all agents. The number of groups, their size, and the intensity of the within-group collaboration are determined by the intensity of the competition. For instance, when tournament prizes are small, the competition is mild and the efficient (complete) network of collaboration is stable.

The intuition behind this result builds on the observation that a large enough group can guarantee top tournament rankings for its members, irrespective of what the rest of the agents do. To achieve that, the group members must sacrifice collaboration with outsiders. Roughly speaking, a large enough group has a collective maxmin strategy that yields a high payoff for its members. Indeed, members of this group can refuse to collaborate with outsiders. If a group constitutes a majority, there are more collaborative opportunities within the group than outside of it; thus group members have a competitive advantage in the tournament.

The size of each group is endogenous. It can be found by maximizing an agent's payoffs across complete components of various sizes (assuming that the agent is part of these complete components). For example, the size of the largest group maximizes a participant's payoff across all possible groups that can be formed by a strict majority. One interesting interpretation of this criterion is the following: Imagine a by-invitation-only union in which all participants collaborate with each other. Start with a union that is formed by the smallest strict majority.[1] In my model, such a union stops inviting new members as soon as it reaches the size of the largest group.

[1]This is a necessary condition for union members to dominate the tournament.

My second finding is a necessary and sufficient condition for stability of efficient outcomes in winner-takes-all tournaments. I show that there exists a stable set that contains an efficient outcome if and only if a payoff of an agent in this outcome is weakly larger than a payoff of an agent in any complete component that constitutes a strict majority and guarantees its members top rankings in the tournament. Moreover, if such a stable set exists, it is a singleton. For winner-takes-all tournaments, this condition is equivalent to the prize in the tournament being sufficiently small. To the best of my knowledge, this result does not appear in the literature (with the notable exception of Dutta et al. (1998); however, a similar observation in their paper is derived only for a three-agent example, and it does not generalize).

The important driving force behind these two results is an externality caused by tournament competition. Consider a complete network of collaboration in which all agents tie for all rankings in the tournament. Reducing the intensity of a link between two agents moves both of them all the way down to the bottom two positions of the ranking or, equivalently, moves the rest of the agents away from the bottom two positions. In this case, the two agents who reduce the intensity of the link bear the opportunity cost, which equals the value of lost collaboration and the value of the top-ranking positions. At the same time, these agents impose a positive externality on the rest of the agents, since the rankings of the latter improve. Clearly, agents cannot exploit this positive externality to their benefit unilaterally, but collectively such an exploitation may be possible. For instance, consider all agents who sever a link with agent i. These agents internalize the effect of the positive externality they impose on each other.

I find that the requirement for the stability of efficient outcomes is very demanding. A natural question, then, is whether one can allow agents to buy missing collaboration from each other and restore efficiency. In particular, there are large gains from such a trade in stable outcomes, in which networks of collaboration feature group structure. In my most general version of the model, I allow agents to use monetary transfers to pay each other for collaboration. Transfers are modeled as voluntary bilateral agreements, in which a pair of agents jointly decide on the amount of money one agent pays the other.

I show that transfers do not resolve the tension between stability and efficiency. In particular, the opportunity to transfer money voluntarily does not affect the stability of outcomes in which agents in larger groups refuse to collaborate with agents in smaller groups. The absence of links between groups in these outcomes results in efficiency losses. I show that even if we allow agents to split the gains of restored links endogenously and without any restrictions, missing links are not restored. Intuitively, agents are substitutes for each other. When negotiating the price of a missing link, agents propose to implement new outcomes that generate larger welfare compared to the starting point of negotiation. However, these new outcomes are prone to collective deviations, and I show that the set of collective deviations is so rich that the long-term gain of implementing these outcomes is always zero. An important assumption in this part of the model is that the transfers are part of self-enforcing bilateral agreements and, therefore, are set in a decentralized manner.

These results are in line with the observation that the structure of collaboration between competitors is often asymmetric and inefficient. For example, Bekkers et al. (2002) show that the network of cross-licensing agreements between participants in the Global System for Mobile Communications (GSM) market has a tightly connected cluster of industry leaders. Some firms were left out of these agreements, despite having large portfolios of patents that were essential to GSM technology.

The rest of the paper is structured as follows. Related literature is discussed in Section 2. Section 3 contains a simple three-agent example that outlines the main findings of the paper. The setup of the general model in Section 4 is followed by the results in Section 5. Applications of the model are discussed in Section 6, and Section 7 concludes.

2. Related literature

This paper contributes to the literature on collaboration between rivals. Related models are studied by Bloch (1995), Yi (1998, 1997), and Yi and Shin (2000) in the context of coalition formation, and Joshi (2008), Goyal and Joshi (2003), Goyal and Moraga-Gonzalez (2001), Marinucci and Vergote (2011), Mauleon et al. (2014), and Grandjean and Vergote (2015) in the context of network formation. These studies focus on R&D collaboration among firms as the main application.

Other applications that are relevant for this paper are sabotage in tournaments and the interaction between discrimination and social status. Lazear (1989), Chen (2003), and Konrad (2000) study various aspects of sabotage in tournaments. McAdams (1995) studies racial discrimination that is fueled by a desire to obtain higher social status.

A stability concept used in this paper—the farsighted stable set—is closely related to various solutions used in the literature on coalition and network formation with farsighted agents. Several papers in this literature follow a cooperative approach and use farsighted stability concepts as solutions. This strand of the literature includes Greenberg (1990), Chwe (1994), Ray and Vohra (1997), Diamantoudi and Xue (2007), Herings et al. (2009), Page et al. (2005), Grandjean et al. (2010, 2011), and Mauleon et al. (2011). The version of the farsighted stable set used in this paper differs from the versions defined in the above papers in a few aspects. First, I allow for arbitrary acting coalitions (Herings et al. 2009 and Mauleon et al. 2011 restrict the acting coalition to be a singleton or a pair). Second, I allow agents to choose all of their actions (i.e., the intensity of collaboration and the sizes of transfers) in a cooperative manner (Herings et al. 2009, Page et al. 2005, Grandjean et al. 2010, 2011, and Mauleon et al. 2011 focus on pure network formation).

Another strand of the literature uses dynamic noncooperative models to describe the process of coalition or network formation. Among these are Aumann and Myerson (1988), Bloch (1996), Konishi and Ray (2003), and Dutta et al. (2005). Dynamic models can naturally accommodate the time preferences of agents involved in the network-formation process. However, this comes at the cost of less rich sets of coalitional deviations that agents are allowed to undertake. In most of these models, exogenously chosen proposers (or agenda setters) suggest the course of action.

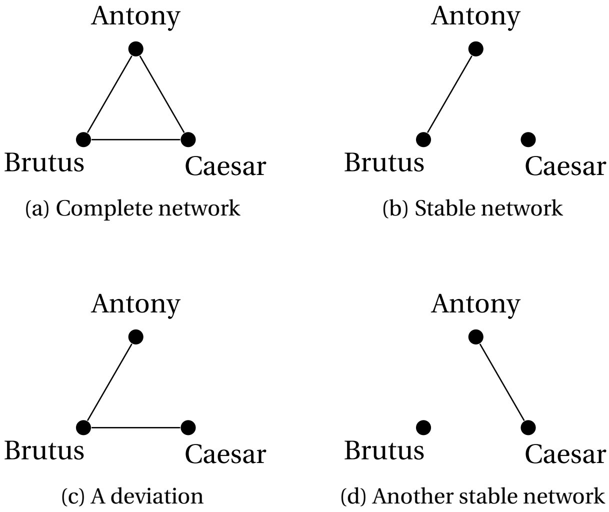

FIGURE 1. Various networks of collaboration.

3. Simple example

In this section, I present a simple three-agent example that illustrates my main findings. Consider three engineers—Antony, Brutus, and Caesar— who are participating in a winner-takes-all tournament. The objective of the tournament is to select the best design for a phone. Each engineer is an expert on a particular phone module: Antony's specialty is touch screens, Brutus's is batteries, and Caesar's is mobile processors and memory modules.

The engineers can ask each other to design high-quality proprietary modules for their phones or they can source low-quality generic modules from the market. When two engineers—say, Antony and Brutus—agree to collaborate, Antony can use a battery design developed by Brutus in exchange for his own touch screen design.[2] In this case, their products will have identical proprietary touch screens and batteries. It is convenient to represent a structure of bilateral collaboration by a network (see Figure 1) in which nodes correspond to agents and links correspond to collaborations.

For simplicity, assume that the quality of a final product is strictly increasing in the number of proprietary modules and does not depend on any other characteristics. Therefore, an engineer whose phone has the largest number of proprietary modules wins the tournament. Also, assume that even if an engineer does not win the tournament, he can use his prototype in the future. The latter means that developing a high-quality prototype is valuable: Let $f(k)$ be the value of a prototype with k proprietary components and let R be a prize in the tournament. An engineer with a prototype that has k proprietary components receives a payoff

$$f(k) + wR,$$

where $w \in [0, 1]$ is the engineer's chances of winning the tournament.

[2]Such a collaboration is essentially a cross-licensing agreement when engineers have patent protection for their proprietary components.

In this very stylized tournament, there is only one decision that each engineer has to make: with whom to collaborate. Consider Antony and Brutus. Collaboration between them does not change their relative positions in the tournament. Suppose that Antony has a better prototype than Brutus. I assume that if they collaborate, Antony's prototype will still be better than Brutus's. Moreover, collaboration contributes to the value of both prototypes and makes them more competitive than Caesar's prototype.

If the competitors are myopic and can only make one link change at a time, they will fully collaborate, and all three prototypes will be built with proprietary components. More formally, the unique pairwise stable network of collaboration is a complete one (see Figure 1(a)). This outcome is also the unique efficient outcome, since the tournament is a constant-sum game, and the value of prototypes is increasing with collaboration.

This paper focuses on a case in which agents are farsighted (i.e., they care about their long-run payoffs) and able to coordinate with each other. I show that the complete network of collaboration is no longer a plausible prediction. For instance, suppose that R is large and all three engineers are collaborating with each other. Any two engineers (e.g., Antony and Brutus) have a jointly profitable deviation. If they simultaneously refuse to share their modules with Caesar, the value of their prototypes drops from $f(2)$ to $f(1)$, but their individual chance of winning the tournament increases from $1/3$ to $1/2$, since Caesar's prototype becomes strictly worse than the other two prototypes (see Figure 1(b)). If $R/6 > f(2) - f(1)$, such a deviation is mutually beneficial for Antony and Brutus.

Naturally, one may cast doubt on the credibility of this deviation. For instance, both Brutus and Caesar prefer to restore their missing link so as to proceed from the outcome depicted in Figure 1(b) to the outcome depicted in Figure 1(c). Note that the credibility of the latter deviation is also not obvious, as both Antony and Caesar would like to seize their collaboration with Brutus and restore their missing link so as to proceed from the outcome depicted in Figure 1(c) to the one depicted in Figure 1(d). It is easy to see that there are no outcomes in this example that are immune to all coalitional deviations.

To resolve this problem, I relax the stability requirement. Suppose that stable outcomes are those that are immune only to *credible* coalitional deviations (i.e., to deviations to other stable outcomes).[3]

If $R/6 > f(2) - f(1)$, a set of all collaborative networks with exactly one link is stable. To show this, consider the following two arguments. First, there is no coalition of engineers who can and want to proceed from the outcome depicted in Figure 1(b) to the one depicted in Figure 1(d). Indeed, the only engineer who wants to follow this path is Caesar, and he cannot do anything to make this transition happen (he needs Antony's active participation, but Antony does not gain anything from this transition). Therefore, these three outcomes are immune to deviations to stable outcomes. Second, for any outcome with zero, two, or three links, there is a coalition of two engineers who want to proceed to an outcome in which they collaborate only with each other. Moreover, these two engineers can always implement this transition without relying on the third.

[3]This definition implies that a set of stable outcomes must be self-enforcing.

Therefore, outcomes with zero, two, or three links are not immune to deviations to stable outcomes.

To get a better intuition for the solution, consider two sets that are not stable: a singleton that contains a complete network and the set that contains two networks, as depicted in Figures 1(c) and 1(d). The first set does not satisfy the criteria for stability, because there are outcomes that are not included in it and that are immune to deviations to the (allegedly) stable outcomes. In particular, the networks in Figures 1(b) and 1(d) are immune to deviations to the complete network. The property that the complete network fails to satisfy is called external stability.

The second set, which consists of the two networks depicted in Figures 1(c) and 1(d), does not satisfy the criteria for stability because the network in Figure 1(c) is not immune to a deviation to the (allegedly) stable network in Figure 1(d). The property that this set fails to satisfy is called internal stability. I discuss internal and external stability in detail in Section 4.

When is the efficient level of collaboration stable in this example? All three engineers share their design when competition is not too fierce, compared to direct benefits from collaboration; i.e., when $R \leq 6(f(2) - f(1))$. This condition can be rewritten as

$$n \in \arg\max_{k>n/2}\{V_k\},$$

where $V_k = f(k) + R/k$ is the payoff of an agent participating in a large, fully collaborating group of size k, and n is the total number of players ($n = 3$ in this example). Intuitively, when a group is formed, its size is determined by the utility of its representative member. New members are added only if the current members benefit from the addition, and existing members are excluded if the remaining members benefit from the exclusion.

In the inefficient outcome in which, for instance, Antony and Brutus collaborate with each other and Caesar is on his own, there are gains from trade: Caesar could collaborate with the two other engineers and compensate them for their loss in the tournament. Despite the presence of gains from such a trade, voluntary transfers cannot destabilize the inefficient outcome mentioned above. The engineers are imperfect substitutes for each other. Therefore, in the situation in which Caesar pays for his collaboration with competitors, he can propose a new arrangement in which one of the competitors, e.g., Brutus, is dropped and the other, Antony, is compensated with a small amount for following this proposal.

The findings presented in this section do not depend on the simplifying assumptions about three-player winner-takes-all tournaments and the discreet and costless nature of collaboration. In the next section, I present a much richer model, followed by formal results that generalize the observations discussed here.

4. Model

Let $N = \{1, \dots, n\}$ be a set of identical agents competing in a tournament. Tournament participants engage in bilateral collaborations with each other. Agent $i \in N$ chooses a vector of efforts $\mathbf{x}_i = (X_{i,j})_{j\in N} \in \mathbb{R}_+^n$. A component $X_{i,j}$ of this vector is the amount of effort agent i contributes to the collaboration with agent j. A matrix of efforts is defined as $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$.

The structure of collaboration is described by a symmetric matrix G that is defined as

$$G_{i,j} = g\bigl(\min\{X_{i,j}, X_{j,i}\}\bigr).$$

It is useful to visualize this matrix as a network in which links between the agents represent bilateral collaboration. A link between agents i and j has an intensity $g(\min\{X_{i,j}, X_{j,i}\})$. I model the intensity as (a transformation of) the smaller of the two efforts to capture the idea that collaboration requires the consent and active participation of both collaborators.

The diagonal elements of matrix G play a special role in this model. For an agent i, $G_{i,i} = g(X_{i,i})$ is (a transformation of) the effort the agent spends working solo. Therefore, the diagonal elements of G capture the activities that do not require collaborative partners, but are beneficial for tournament participants. I assume no special order of choosing the two types of efforts. This assumption allows me to capture the idea that agents may respond to changes in their rivals' collaborative levels by adjusting their efforts toward working solo and vice versa. The diagonal of G can be used to compare the results of the main model to a scenario in which collaboration is infeasible.[4]

I assume that function $g : \mathbb{R}_+ \to \mathbb{R}_+$ is strictly increasing, concave, and bounded from above by $\overline{g} = \lim_{z\to\infty} g(z)$. The monotonicity property is self-explanatory. The concavity of g reflects the decreasing returns to the collaborative effort. The assumption that g is bounded means that the number of collaborative partners plays a crucial role in this model.[5] I use the normalization $g(0) = 0$.

The following notation is useful. For $M \subset N$, $I(M) \in \{0,1\}^{N\times N}$ is a matrix such that for all i, $[I(M)]_{i,j} = 1$ if $\{i,j\} \subset M$ and $[I(M)]_{i,j} = 0$ otherwise. In particular, matrix $I(\varnothing)$ describes the empty network and $I(N)$ describes the complete network in which every link has a unit intensity. For two matrices Y and Z, denote their Hadamard product by $Y \circ Z$: for all i, j, $[Y \circ Z]_{i,j} = Y_{i,j}Z_{i,j}$.

In the course of the tournament, agent i produces an output y_i that is determined by the total intensity of the agent's collaboration:

$$y_i(X) = \sum_{j=1}^{n} g\bigl(\min\{X_{i,j}, X_{j,i}\}\bigr) = \sum_{j=1}^{n} G_{i,j}.$$

To model the process of forming collaborative relationships, I follow a cooperative approach; i.e., I define a set of outcomes, agents' preferences, and a binary blocking relation on this set. Using these components, I study stable outcomes in the sense of von Neumann and Morgenstern (see von Neumann and Morgenstern 1944).

[4]The returns to working solo are assumed to be equal to the returns to working with a partner. This assumption does not play an important role in the analysis, and can be easily removed at the cost of introducing additional notation.

[5]I discuss this assumption in detail in Section 5.

An outcome in this model is a pair (X, T), where $X \in \mathbb{R}_+^{N\times N}$ is a matrix of efforts that define a structure of collaboration, and $T \in \mathbb{R}_+^{N\times N}$ is a matrix that describes a system of transfers between agents. I assume that $T_{i,j} \geq 0$ is the amount that agent i pays to agent j in the outcome (X, T). By $(X, 0_{n,n})$, I denote an outcome with zero transfers. Finally, by $\mathcal{U}$, I denote a set of all feasible outcomes.

The result of the tournament depends on the vector of the agents' outputs. In particular, given an outcome (X, T), the agents are ranked according to their outputs in descending order. Ties are resolved randomly using the uniform distribution. Let $R : N \to \mathbb{R}$ be a tournament prize schedule; i.e., $R(k)$ is the prize for an agent ranked kth in the tournament. I assume that R is decreasing and convex (the latter means that $R(k) - R(k+1)$ is decreasing in k), and I normalize the prize for the agent with the lowest ranking to be zero, i.e., $R(n) = 0$. For any $i, j : 1 \leq i \leq j \leq n$, let

$$r(i, j) = \frac{1}{j-i+1}\sum_{k=i}^{j} R(k)$$

be an expected prize for an agent who is randomly placed between rankings i and j in the tournament (by construction, this agent ties with $j - i$ other agents).

The agent's payoff is additive in his tournament prize, output, cost of effort, and transfers. The payoff of agent i in outcome (X, T) is

$$U_i(X, T) = r\big(p_i(X), q_i(X)\big) + f\big(y_i(X)\big) + \sum_{j=1}^{n}(T_{j,i} - T_{i,j} - cX_{i,j}),$$

where $c > 0$ is a constant marginal cost of effort, and p_i and q_i denote the lower and the upper bounds on possible rankings for agent i in the tournament. These bounds are defined as

$$p_i(X) = \big|\{k \in N : y_i(X) < y_k(X)\}\big| + 1$$

and

$$q_i(X) = n - \big|\{k \in N : y_i(X) > y_k(X)\}\big|.$$

I assume that function $f : \mathbb{R}_+ \to \mathbb{R}_+$ is increasing.

By $\mathbf{U}_M(X, T)$, I denote a vector of utilities for the set of agents M in outcome (X, T). Also, for two vectors $\mathbf{U}_M$ and $\mathbf{V}_M$, I say that $\mathbf{U}_M \gg \mathbf{V}_M$ if $\forall i \in M$, $U_i > V_i$.

In this specification, the agents may derive a positive net value of collaboration without taking into consideration a tournament outcome. In the vast majority of the literature (see Goyal and Joshi (2003), Goyal and Moraga-Gonzalez (2001), Marinucci and Vergote (2011), and others), collaboration is assumed to be costly. I consider the possibility of both costly and costless links.[6] The latter is interesting for two reasons. First, this assumption relates better to some of the applications that I discuss in Sections 1

[6] The case of costless links models a situation in which any two agents are always better off collaborating with each other if the actions taken by all other agents remain unchanged.

and 6; second, it allows me to highlight a novel interaction between collaboration and competition.

In the main specification of the model, I assume that the agents derive a value only from their direct connections. One could get similar results if indirect connections were assumed to be valuable for the agents. I consider such an extension in Section 5.3.

Since agents' utilities are linear in transfers, f is increasing, and g is strictly increasing and concave, the set of efficient outcomes consists of all outcomes in which all agents collaborate at the optimal level with all available partners. In the efficient outcomes, networks of collaboration are complete.

Remark 4.1. An outcome (X, T) is efficient if and only if, for all i, j, $X_{i,j} = x^*$, where

$$x^* = \arg\max_{x \geq 0} \{f(ng(x)) - cnx\}.$$

A corresponding network of collaboration for an efficient outcome is always complete.

Proof. Start with the observation that the Pareto frontier is a flat surface with a slope of 45 degrees. Therefore, one can use a utilitarian welfare criterion. Consider an outcome (X, T). Observe that $\sum_{i=1}^n \sum_{j=1}^n T_{i,j} = 0$. The social welfare in this outcome is

$$W = \sum_{i=1}^n U_i(X, T) = \sum_{i=1}^n \left(f\left(\sum_{j=1}^n G_{i,j} \right) - c \sum_{j=1}^n X_{i,j} \right) + \sum_{i=1}^n R(i).$$

This expression achieves the maximum if and only if $X = x^* I(N)$. □

The main result of this paper involves the stability of efficient outcomes in this model. As shown in Remark 4.1, besides efficiency, these outcomes have another potentially desirable property—completeness of the network of collaboration.

4.1 *Network formation and stability*

When modeling the formation of collaborative relationships, I follow the usual practice in cooperative games: I define a notion of stability using a binary blocking relation on the set of feasible outcomes.

To understand the idea behind the blocking relation, consider a group (or a coalition) of players carrying out a transition from one outcome to another. Once the transition takes place, farsighted agents expect additional transitions. Eventually, as the result of a sequence of such transitions, the agents arrive at the "terminal" outcome from which no further transitions are attempted. A necessary condition for rational agents to engage in such a sequence of transitions is that ultimately, in the terminal outcome, they are better off. I implicitly assume that the agents do not derive the utility from transitory outcomes along a transition. More precisely, if there are two different transitions between outcomes (X, T) and (X', T'), agents do not distinguish between these two transitions, because the final destination is the same. One way to justify this assumption is to interpret the transitions as proposals and counterproposals (or objections) that agents make to each other without engaging in the actual modification of physical outcomes. These proposals are meant to convince everyone to proceed to a stable outcome right away.

The following definition formalizes the idea of a feasible transition, i.e., what each coalition can do in terms of shaping outcomes. Note that the feasibility of a transition does not depend on agents' preferences.

Definition 4.2. A coalition M can enforce a transition from outcome (X, T) to outcome (X', T'), i.e., $(X, T) \overset{M}{\to} (X', T')$ if, for all $i, j \in N$,

(i) $X'_{i,j} \neq X_{i,j}$ implies $i \in M$,

(ii) $T'_{i,j} > T_{i,j}$ implies $i, j \in M$,

(iii) $T'_{i,j} < T_{i,j}$ implies $i \in M$ or $j \in M$.

According to this definition, all agents who are active during a transition from one outcome to the other must be contained in the coalition that enforces the transition. In this definition, it is postulated that players can unilaterally choose collaborative efforts. Recall, however, that an increase in this effort does not necessarily translate into an increase in a collaborative intensity, because $G_{i,j} = g(\min\{X_{i,j}, X_{j,i}\})$. For example, if $X_{i,j} = X_{j,i}$, both agent i and agent j must increase their efforts to increase the intensity of their collaboration. Any agent can always unilaterally decrease the intensity of the collaboration with any of his partners. This dichotomy reflects the fact that collaboration is achieved through a bilateral agreement and is a standard assumption in the literature on the formation of undirected networks.

A reduction in the amount of money transferred can be achieved unilaterally, either by refusing to pay (on the side of the sender) or refusing to accept (on the side of the receiver).

The next definition introduces a blocking relation that formalizes, among other things, the assumption that agents are rational and farsighted.

Definition 4.3. An outcome $(X, T) \in \mathcal{U}$ setwise farsightedly blocks $(X', T') \in \mathcal{U}$ or

$$(X, T) \rhd \left(X', T'\right)$$

if there exists a finite sequence $\{(S_k, X^k, T^k)\}_{k=1}^{K}$ $\forall k = 1, \dots, K$, $S_k \subset N$, and $(X^k, T^k) \in \mathcal{U}$ such that

(i) $(X', T') = (X^1, T^1) \overset{S_1}{\to} (X^2, T^2) \overset{S_2}{\to} \cdots \overset{S_K}{\to} (X, T)$,

(ii) $\mathbf{U}_{S_k}(X, T) \gg \mathbf{U}_{S_k}(X^k, T^k)$ for all $k \leq K$.

To establish the intuition for this definition, assume that all agents view outcome (X, T) as stable (this assumption is confirmed in the definition of stable sets of outcomes, Definition 4.4). This outcome blocks the other outcome (X', T') if the following conditions hold:

(i) There exists a sequence of transitions that starts at (X', T') and arrives at (X, T); every outcome of this sequence, except for (X, T), is assigned an active coalition that enforces a corresponding step in the transition.

(ii) Every member of an active coalition strictly prefers the final destination of the transition (X, T) to the outcome in which the coalition becomes active. In other words, every agent who has to modify his choice of efforts and transfers so that the transition proceeds benefits from the transition once it is complete. This condition mimics the transition process in which coalition members are asked whether they wish to proceed with the transition or stay in the current outcome.

The blocking relation makes little sense on its own, and its adequacy should not be judged in the absence of the stability concept. It is implicitly assumed that all agents who participate in a sequence of transitions from (X', T') to (X, T) believe that the latter outcome is final or, in other words, stable. When this definition is used to check for stability, a blocking outcome is always stable and a blocked outcome is arbitrary. A stability notion that I use in conjunction with this blocking relation is the von Neumann–Morgenstern stable set defined for an abstract problem $(\mathcal{U}, \rhd)$.[7]

Definition 4.4. A set of outcomes $\mathcal{R} \subset \mathcal{U}$ is farsighted stable[8] if it satisfies internal stability (IS) and external stability (ES) conditions:

(IS) For any $(X, T), (X', T') \in \mathcal{R}$, $(X, T) \not\rhd (X', T')$.

(ES) For any $(X', T') \notin \mathcal{R}$, there exist $(X, T) \in \mathcal{R}$, $(X, T) \rhd (X', T')$.

Internal stability requires that stable outcomes do not block other stable outcomes, while external stability requires that all outcomes that are not part of a stable set are blocked by stable outcomes.

A stable set of outcomes is a collection of all outcomes that are unblocked by elements of this stable set. Let $Y : 2^{\mathcal{U}} \to 2^{\mathcal{U}}$ be a function that, for a set of outcomes $\mathcal{X}$, returns a set $Y(\mathcal{X})$ of all outcomes that are unblocked by any outcome in $\mathcal{X}$: $Y(\mathcal{X}) = \{(X, T) \in \mathcal{U} : (X', T') \not\rhd (X, T)\ \forall (X', T') \in \mathcal{X}\}$. Then $\mathcal{R}$ is farsighted stable if and only if

$$\mathcal{R} = Y(\mathcal{R}).$$

Farsighted stability is a set-valued solution. An element of a stable set is not considered stable in isolation (unless the stable set is a singleton). The stability of a single element hinges on the stability of all other elements in the stable set. This means, for instance, that there can be more than one stable set.

Note that in both the internal and external stability conditions, the outcomes that are blocking or not blocking other outcomes come from the conjectured stable set. Put

[7]One can also define an abstract core for $(\mathcal{U}, \rhd)$. However, in my model, for the most interesting values of parameters, this abstract core is empty.

[8]There is little agreement on naming various stability concepts in the recent literature on cooperative games; this name is chosen following Ray and Vohra (2015).

differently, this definition ignores the instances in which an unstable outcome blocks a stable outcome, because the transition that implements this blocking is not credible. Indeed, the agents who participate in this transition should not be concerned about their well-being in the unstable outcome. This connection between Definitions 4.3 and 4.4 is crucial for understanding the meaning of a farsighted blocking relation.

As Chwe (1994) shows, internal and external stability, together, imply that stable sets possess a consistency property. Any collective short-run profitable deviation to an unstable outcome is punished by a low long-run payoff for at least one of the agents who participated in the deviation. I discuss this property in detail in the Supplemental Appendix.[9]

Farsighted stability implicitly requires all agents to agree on the set of outcomes that, once reached, are not followed by any deviations. It also requires that all agents involved in a sequential deviation agree on the exact path of this deviation. Put differently, I do not allow a situation in which an agent initiates a certain transition that ends up different from his original plan due to the actions of the other agents involved.

An alternative way to model agents' beliefs is to define an expectation function with a Markov property: a function that returns a stable outcome for any given outcome if the latter is treated as an initial condition for the transition process. This function is similar in spirit to a subgame perfect Nash equilibrium (SPNE) strategy profile in a dynamic game. The expectation function was proposed by Jordan (2006) and was later used by Acemoglu et al. (2012), Acemoglu et al. (2015), Dutta and Vohra (2017), and others. The advantage of the latter approach over sequential blocking is that the agent's expectations do not depend on whether the agent is currently involved in a transition between outcomes. Also, the expectation function allows the acting agents to maximize their preferences rather than just improve their well-being. However, the drawback of this approach is that the expectation function may not exist if the set of feasible transitions is large.

5. Stable sets

The main result of this paper is twofold. First, Theorem 5.3 shows that there always exist stable sets in which each outcome has a group structure. These stable sets exist both when transfers are allowed and when they are not. When the group structure in these stable sets is nondegenerate (i.e., there is more than one group), all of the outcomes in these sets are inefficient. Second, I examine the hypothesis that there exist other stable sets that contain efficient outcomes. For winner-takes-all tournaments without transfers, Theorem 5.6 provides necessary and sufficient conditions for efficient outcomes to be included in a stable set. If these conditions are satisfied, the stable set is unique and does not contain inefficient outcomes.

5.1 *Groups*

In this section, I characterize a special class of stable sets that always exist in this model. They have the following properties. First, agents are partitioned into groups of a certain

[9] The Supplemental Appendix is available at https://ssrn.com/abstract=3220682.

size. Each group must be larger than the union of all of the smaller groups, i.e., the largest group must contain a strict majority, the second largest group must contain a strict majority once the largest group is removed, and so on.

Second, agents collaborate with all members of their own group and not with anyone else. All agents within a group exert the same effort, but these efforts may differ across groups. This property implies that all members of a group will tie in the tournament.

Third, the effort exerted by an agent in a group k must be large enough to guarantee that his output is weakly larger than the largest output that agents in the smaller groups can produce. To calculate this bound on effort, one should look at the counterfactual outcome in which all members of all groups that are smaller than group k collaborate with each other at the maximum (unachievable) level $\overline{g}$. This condition ensures that agents' tournament rankings are increasing in the size of their group and that agents with low rankings cannot overthrow agents with high rankings.

Finally, the size of each group is chosen to maximize the payoff of its representative member, taking the sizes of all larger groups as given.

To formalize this construction, consider a group of r agents collectively trying to outperform other $q < r$ agents who are collaborating with each other at a very high (unfeasible) level of effort, but not collaborating with the rest of the agents. The group is guaranteed to succeed in this task if each member exerts an effort x toward every available partner, where x satisfies $rg(x) \geq q\overline{g}$ or

$$x \geq g^{-1}\left(\frac{q\overline{g}}{r}\right). \tag{1}$$

If the members of the group want to maximize their output net of cost of effort, under the requirement that they outperform other $q < r$ agents for all possible levels of collaboration between the latter, they must solve the problem

$$v(r, q) = \max_{x \geq g^{-1}(\frac{q\overline{g}}{r})} \bigl\{f\bigl(rg(x)\bigr) - crx\bigr\}. \tag{2}$$

The effort level that solves this problem is an analog of a maxmin strategy. The sizes of the groups can now be defined.

Definition 5.1. Consider a sequence $\{m_k\}_{k=1}^K$. Let $M_0 = 0$ and, for $k \geq 1$, let $M_k = \sum_{i=1}^k m_i$. The sequence $\{m_k\}_{k=1}^K$ is group-optimal if $M_K = n$ and for all $k \geq 1$,

$$m_k \in \operatorname*{arg\,max}_{\frac{n-M_{k-1}}{2} < m \leq n - M_{k-1}} \bigl\{r(1 + M_{k-1}, m + M_{k-1}) + v(m, n - m - M_{k-1})\bigr\}.$$

For clarity of exposition, I assume that the group-optimal sequence is unique. All of the results easily generalize to multiple group-optimal sequences by taking a union across these sequences. Given a group-optimal sequence $\{m_k\}_{k=1}^K$, let

$$V_k = r(1 + M_{k-1}, M_k) + v(m_k, n - M_k)$$

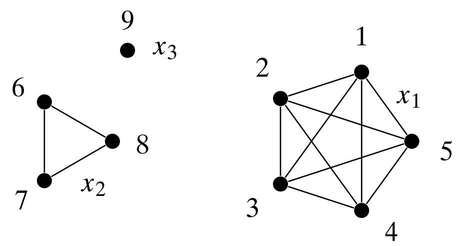

Figure 2. An outcome that has a group structure induced by a sequence $\{5, 3, 1\}$.

be a payoff of a representative member of group k and let

$$x_k = \operatorname*{argmax}_{x \geq g^{-1}(\frac{q\overline{g}}{r})} \{f(m_k g(x)) - cm_k x\}$$

be an effort exerted by this member toward a collaboration with another member of the same group.

The definition of m_1 and V_1 considers a set of all outcomes in which a majority group of size m forms a complete component in which all members of the group collaborate at the payoff-maximizing level, subject to the constraint that their effort must be sufficiently high to dominate all outsiders in the tournament independent of the efforts of the outsiders. The size of the majority group m_1 is chosen to maximize the payoff of a single member. The criterion for m_k is identical to the criterion for m_1, formulated with respect to a "residual" problem in which the sizes and structure of all of the larger groups are fixed.

Definition 5.2. An outcome (X, T) has a group structure induced by a sequence $\{m_k\}_{k=1}^K$ if there exists a partition $\mathcal{N} = \{N_1, \ldots, N_K\}$ of the set N such that

(i) $\forall k, |N_k| = m_k$,

(ii) $X = \sum_{k=1}^K x_k I(N_k)$.

An example of an outcome that satisfies Definition 5.2 is given in Figure 2. This outcome is induced by a sequence $\{5, 3, 1\}$ and effort levels x_1, x_2, and x_3. There are three complete components or groups of size 5, 3, and 1. A member of group k exerts efforts x_k along every link that is present in Figure 2.

Stable sets of outcomes always exist in this model, and at least one consists of outcomes that have a group structure. In these outcomes, the agents may use transfers within a group, but these transfers do not affect the distribution of payoffs, i.e., for each agent, the sums of outgoing and incoming transfers are equal.

Theorem 5.3. *Set of outcomes $\mathcal{R}$ is stable if every outcome $(X, T) \in \mathcal{R}$ has a group structure induced by a group-optimal sequence and satisfies*

(i) $X_{i,j} = 0$ *implies* $T_{i,j} = 0$,

(ii) *for any* $i \in N$, $\sum_{j\in N} T_{i,j} = \sum_{j\in N} T_{j,i}$.

Proof. Denote a set that satisfies the conditions of the theorem by $\mathcal{R}$. One has to show that set $\mathcal{R}$ is internally and externally stable.

I start with internal stability. I show that for any $(X', T'), (X, T) \in \mathcal{R}$, $(X, T) \not\triangleright (X', T')$. Let $\mathcal{H} = \{H_1, \dots\}$ be a partition that induces (a network of collaboration in) (X, T) and let $\mathcal{F} = \{F_1, \dots\}$ be a partition that induces (X', T'). Also, let $B = \{i \in N : U_i(X, T) > U_i(X', T')\}$.

The following argument formalizes the idea that agents in set F_1 do not participate in the transition from (X', T') to (X, T), because their utility cannot be further increased. Agents from set F_2 do not participate in this transition, because for any of them to increase their utility, they must get a spot in set H_1. However, for that to happen, at least one agent from F_1 must participate in the transition. A similar argument applies to sets F_2, F_3, etc.

Formally, denote an index of a largest set populated by agents from B in (X, T) by k, i.e., for all $j < k$, $B \cap H_j = \varnothing$ and $B \cap H_k \neq \varnothing$. Let $M = \bigcup_{j\leq k} F_j$ and note that $|M| > N/2$. For any $S \subset N \setminus M$ and for any $(\hat{X}, \hat{T})$, $(X', T') \overset{S}{\to} (\hat{X}, \hat{T})$, I have $U_M(X', T') = U_M(\hat{X}, \hat{T})$. Hence, if $(X, T) \triangleright (X', T')$, it must be that $U_M(X', T') = U_M(X, T)$, which contradicts $B \cap M \neq \varnothing$ (by construction of set B, if $i \in B \cap H_k$, it must be the case that $i \in F_j$ for some $j < k$).

To show that $\mathcal{R}$ satisfies external stability, for any $(X', T') \notin \mathcal{R}$, I construct $(X, T) \in \mathcal{R} : (X, T) \triangleright (X', T')$. By the definition of set $\mathcal{R}$, every element of this set has a group structure. I partition the transition from (X', T') to (X, T) into K stages in such a way that in the course of stage k, only agents who form a group of size m_k are active and, at the end of the stage, this group is formed.

The following result is used to complete the proof.

Definition 5.4. An outcome $\gamma = (X, T)$ contains a top component if $\exists M \subset N$, $|M| = m_1$, $\forall i \in M$, $X_{i,j} = x_1 \mathbb{I}\{j \in M\}$, $\sum_{j\in M} T_{i,j} = \sum_{j\in M} T_{j,i}$, and $\sum_{j\notin M} T_{i,j} = 0$.

Lemma 5.5. *Denote a set of agents whose payoff is below* V_1 *by* $A(X, T) = \{i : U_i(X, T) < V_1\}$. *For any outcome* (X, T), *either* (X, T) *contains a top component or one can always find* (X', T') *such that*

(i) $(X, T) \overset{A(X,T)}{\to} (X', T')$

(ii) $A(X, T) \subsetneq A(X', T')$

(iii) (X', T') *does not contain a top component.*

Proof. Suppose that (X, T) does not contain a top component. Consider an outcome $(\tilde{X}, \tilde{T})$ such that $\tilde{X}_{i,j} = X_{i,j}\mathbb{I}\{i, j \notin A(X, T)\}$ and $\tilde{T}_{i,j} = T_{i,j}\mathbb{I}\{i \notin A(X, T)\}$. There are two cases to consider: either (i) $(\tilde{X}, \tilde{T})$ contains a top component or (ii) the opposite.

If it is case (ii), then $\exists\chi > 0$, for $\hat{X}_{i,j} = X_{i,j}\mathbb{I}\{i, j \notin A(X,T)\} + \chi\mathbb{I}\{i, j \in A(X,T)\}$ and $\hat{T}_{i,j} = T_{i,j}\mathbb{I}\{i \notin A(X,T)\}$, we have $A(X,T) \subsetneq A(\hat{X},\hat{T})$. Also, $(\tilde{X},\tilde{T})$ not containing a top component implies that $(\hat{X},\hat{T})$ also does not contain a top component. Therefore, $(\hat{X},\hat{T})$ satisfies all three conditions of the lemma.

Consider case (i), in which $(\tilde{X},\tilde{T})$ contains a top component. Since (X,T) does not contain a top component and $(\tilde{X},\tilde{T})$ does, there exists a player $k \notin A(X,T)$ such that $\sum_{i\in A(X,T)} G_{i,k} > 0$. Consider an outcome $(X',\hat{T})$, such that $X'_{i,j} = X_{i,j}(\mathbb{I}\{i, j \notin A(X,T)\} + \mathbb{I}\{i = k \text{ and } j \in A(X,T)\} + \mathbb{I}\{j = k \text{ and } i \in A(X,T)\})$. The outcome $(X',\hat{T})$ does not contain a top component. Moreover, by convexity of R, $A(F',\hat{T}) = N \setminus \{k\} \supset A(F,T)$; hence, $(F',\hat{T})$ satisfies the conditions of the lemma. □

Consider an outcome $(\hat{X},\hat{T})$ that emerges at the end of stage $k-1$. There is either a set of agents N_k such that $\forall i \in N_k$, $X_{i,j} = X_{j,i} = x_k\mathbb{I}\{j \in N_k\}$ or the opposite. In the former case, stage k is degenerate. In the latter case, let $A(\hat{X},\hat{T}) = \{i : U_i(\hat{X},\hat{T}) < V_k\}$. If $|A(\hat{X},\hat{T})| < (n - M_{k-1})/2$, applying Lemma 5.5 repeatedly obtains a sequence of outcomes such that the last element of the sequence, $(\tilde{X},\tilde{T})$, satisfies $|A(\tilde{X},\tilde{T})| \geq m_k$. Moreover, the transition between the elements of the sequence can be enforced by agents in corresponding sets $A(\cdot,\cdot)$, and these sets are nested. In the final step of the transition, select m_k agents from $A(\tilde{X},\tilde{T})$ (including all agents who were active in all of the previous steps), and call this set N_k. An outcome (X^*,T^*), such that $X^*_{ij} = x_k\mathbb{I}\{i, j \in N_k\} + \tilde{X}_{ij}\mathbb{I}\{i, j \notin N_k\}$ and $T^*_{ij} = \tilde{T}_{ij}\mathbb{I}\{i, j \notin N_k\}$, finalizes stage k.

The same result holds if transfers are not allowed, i.e., if the set of feasible outcomes is

$$\mathcal{U}_0 = \{(X,T) \in \mathbb{R}_+^{N\times N} \times 0_{n,n}\}.$$

Indeed, the transitions are constructed in such a way that transfers are reduced in the course of a transition. If, in the origin of a transition, all transfers are equal to zero, the whole transition sequence is contained in $\mathcal{U}_0$. □

From an efficiency perspective, the outcomes presented in Theorem 5.3 have too much intragroup collaboration and too little intergroup collaboration.

These outcomes have many missing links in networks of collaboration. A large group of agents isolates itself from others to ensure top tournament ranking for its members. An absence of collaboration between groups is an extreme measure. Indeed, there are other outcomes that induce the same distribution of tournament ranks and feature strictly more direct net benefits from collaboration. In other words, there are outcomes that Pareto-dominate the stable outcomes found in Theorem 5.3. However, deviations to Pareto-improving outcomes are not credible, because collaboration between agents who are ranked differently in the tournament opens a door for further modification of a collaborative network. In particular, agents who are ranked low may threaten others with dropping the existing links. This may lead to losses for agents who are ranked high in the tournaments because they may lose both the value of deleted links and their high ranking. To neutralize threats of this kind, the dominant majority severs all links to all other agents in the stable outcomes found in Theorem 5.3.

Also, for some parameters of the model, there is excessive within-group collaboration. Collaboration within large groups may have an inefficiently high intensity because the members of these groups are threatened by competition from lower-ranked agents. This concern is formalized in inequality (1) when the group-optimal sequence is defined. This competition may not materialize in stable outcomes, but the agents must still take it into account because it can be part of a credible blocking transition.

When the size of group k is chosen, group members face the following trade-off: Making the group smaller leads to a higher expected prize in the tournament, but expanding the group results in more opportunities for collaboration and makes it less costly to compete with the remaining agents. Formally, $r(1 + M_{k-1}, m + M_{k-1})$ is decreasing in m and $v(m, n - m - M_{k-1})$ is increasing in m (because $f(m(g(x)))$ is increasing in m and the set of the permissible effort levels expands with m).

The equation for m_1 is related to union mentality; to see this, consider a problem of a homogeneous union inviting new members. The optimal size of the union, from the point of view of its existing members, is m_1. Each member of such a union evaluates new members based on their potential contribution to existing members' well-being. This decision rule leads to inefficient allocation of membership, because the well-being of outsiders (i.e., potential members) is ignored.

It is well known that in group-formation (or coalition-formation) models, a union mentality results in inefficient outcomes (for a summary of these results, see Ray 2007). Note, however, that in those models, each member of a group has veto power over the inclusion of new members. This veto power reflects the assumption that group membership is exclusive. This is not the case in my model, in which it is feasible for any member of a group to collaborate with outsiders. Nevertheless, there is a stable set of outcomes with a full separation of groups. This means that the notion of a group arises endogenously.

Theorem 5.3 makes the connection between the results obtained in the literature on coalition-formation and network-formation models.[10] Theorem 5.3 justifies the notion of a coalition—or simply a group of agents—that is characterized by exclusive membership and a lack of connections with outsiders. The vast majority of the literature assumes that a coalitional structure is a partition of the set of agents: An agent cannot be a member of more than one coalition at any given moment in time (see Ray 2007, Section 14.4). In my model, this property is endogenous and can be derived from stability conditions. Moreover, the stable set of coalitions in a coalition-formation model (either a canonical cooperative model or a model with sequential proposals, as in Bloch 1995, 1996), in which agents are endowed with the same preferences as in the current model, is the same as the set of groups (network components) in Theorem 5.3. Of course, this result must be taken with a grain of salt, because there may exist other stable sets of networks in which the structure of connections between agents cannot be reduced to groups.

The rules of the tournament—or, more precisely, the feature by which tournament participants are awarded prizes based on a ranking of their outputs—allow me to construct outcomes in which smaller groups of agents cannot change the payoffs of members of larger groups without their consent. This feature underlies the internal stability

[10] A similar result in a different setting appears in Erol and Vohra (2014).

of the stable sets characterized in Theorem 5.3. That is why the largest group should constitute a strict majority; otherwise, the remaining agents may overrule the outcome. This feature is also important for the existence of stable sets: It limits the externalities sufficiently for a stable set to exist for any vector of parameters.

Definition 5.1 pins down the size of each group in the stable networks characterized in Theorem 5.3 uniquely (up to indifference). The literature on pairwise stable networks of collaboration, such as Goyal and Joshi (2003, Proposition 3.5) and Marinucci and Vergote (2011, Propositions 1 and 2), puts bounds on sizes of interconnected groups. The multiplicity of pairwise stable networks in those models follows from the inability to rule out failures of coordination. In my model, group members collectively decide on the group's composition by maximizing participants' payoffs; hence, there is no scope for miscoordination.

Apart from constructing interesting stable outcomes, Theorem 5.3 also solves an important technical problem: It establishes the existence of farsighted stable sets in the model. Indeed, the set of outcomes found in Theorem 5.3 is well defined for any vector of the parameter values. There is no general existence theorem for farsighted stable sets in hedonic games.[11] By proving that any element in the core is a singleton stable set, Mauleon et al. (2011) show that farsighted stable sets always exist in a one-to-one two-sided matching framework. Several papers (see Ray and Vohra 1997, Levy 2004, and Acemoglu et al. 2012) use acyclicity conditions imposed on the superposition of feasible transitions and individual preferences over outcomes to enable the use of backward induction in constructing farsighted stable sets. I show the existence of farsighted stable sets in my model without relying on these commonly used assumptions.

The previous result suggests that competitive forces may lead to inefficient outcomes. In the case of winner-takes-all tournaments, this observation generalizes to any farsighted stable set. The efficient outcome is stable: It belongs to some stable set if and only if the stakes in the competition are low. This is the second main result of this paper.

Theorem 5.6. *Suppose the following conditions:*

(i) Transfers are not allowed, i.e., the set of feasible outcomes is

$$\mathcal{U}_0 = \{(X, T) \in \mathbb{R}_+^{N\times N} \times 0_{n,n}\}.$$

(ii) The tournament is winner-takes-all, i.e., $R(k) = R(n)$ for all $k > 1$.[12]

Then there exists a stable set $\mathcal{R}$ that contains an efficient outcome if and only if a group-optimal sequence is $\{n\}$ or, equivalently,

$$n = \underset{\frac{n}{2}<m\leq n}{\operatorname{argmax}}\{r(1, m) + v(m, n - m)\}. \tag{3}$$

Proof. If (3) holds, Theorem 5.3 implies that there exists a singleton stable set that contains the efficient outcome with the complete network of collaboration.

[11] See the discussion of known existence results in Ray and Vohra (2015).

[12] This condition can be relaxed to $R(2) < r(1, n)$.

Suppose that (3) does not hold, and let $(X, 0_{n,n})$ be an outcome in which X has a group structure induced by some group-optimal sequence $\{m_k\}_{k=1}^K$. Note that $m_1 \neq n$. Let $(X^*, 0_{n,n})$ be the efficient outcome, i.e., $\forall i, j, X^*_{i,j} = x^*$, where

$$x^* = \operatorname{argmax}\{f\big(ng(x)\big) - cnx\},$$

and let $\mathcal{R}$ be a farsighted stable set.

Assume to the contrary that $(X^*, 0_{n,n}) \in \mathcal{R}$. Since $(X^*, 0_{n,n}) \not\triangleright (X, 0_{n,n})$ and $(X, 0_{n,n}) \triangleright (X^*, 0_{n,n})$, it must be that $(X, 0_{n,n}) \notin \mathcal{R}$, and there must exist $(X', 0_{n,n}) \in \mathcal{R}$ such that $(X', 0_{n,n}) \triangleright (X, 0_{n,n})$. Then $(X^*, 0_{n,n}) \triangleright (X', 0_{n,n})$, which is a contradiction.

To show this, define a set of winners in the tournament,

$$H(Z) = \left\{i \in N : \forall k \in N : \sum_{j\in N} g\big(\min\{Z_{i,j}, Z_{j,i}\}\big) \geq \sum_{j \in N} g\big(\min\{Z_{k,j}, Z_{j,k}\}\big)\right\},$$

and a set of agents who are immediately willing to make a transition into $(X^*, 0_{n,n})$,

$$B(Z) = \{i \in N : U_i(Z, 0_{n,n}) < f\big(ng\big(x^*\big)\big) - cnx^* + r(1, n)\}.$$

Since $r(1, n) > R(2) = 0$, for all $i \notin H(Z)$,

$$f\left(\sum_{j=1}^{n} g\big(\min\{Z_{i,j}, Z_{j,i}\}\big)\right) - \sum_{j=1}^{n} cX_{i,j} < f\big(ng\big(x^*\big)\big) - cnx^* + r(1, n)$$

or

$$U_i\big(Z', 0_{n,n}\big) < U_i\big(X^*, 0_{n,n}\big).$$

Therefore, $\forall Z$, $N \setminus H(Z) \subset B(Z)$.

There are two cases to consider: either (i) $|B(X')| < n/2$ or (ii) $|B(X')| \geq n/2$ (if $n = |H(X')|$, it is impossible that $(X', 0_{n,n}) \triangleright (X, 0_{n,n})$).

In case (i), since $(X', 0_{n,n}) \triangleright (X, 0_{n,n})$ for any $i \in H(X')$, $U_i(X', 0_{n,n}) > V_1$. Therefore, either $\forall i \in H(X')$, $j \notin H(X')$, $\min\{X'_{i,j}, X'_{j,i}\} = 0$ or the opposite. In the latter case, select a pair $a \in H(X')$, $b \notin H(X')$, $\min\{X'_{a,b}, X'_{a,b}\} > 0$ and consider an outcome X^1:

$$X^1_{i,j} = X'_{i,j}\mathbb{I}\{i \text{ or } j \notin B\big(X'\big)\} - X'_{b,a}\mathbb{I}\{i = b \text{ and } j = a\}.$$

Clearly, $H(X^1) \subsetneq H(X')$. Repeat this procedure iteratively until either $|B(X^k)| \geq n/2$ or $\forall i \in H(X^k)$, $j \notin H(X^k)$, $\min\{X^k_{i,j}, X^k_{j,i}\} = 0$.

If $\forall i \in H(X^k)$, $j \notin H(X^k)$, $\min\{X^k_{i,j}, X^k_{j,i}\} = 0$ or $|B(X^k)| \geq n/2$, there exists χ such that an outcome X'' satisfying

$$\forall i, j \in N : X''_{i,j} = X^k_{i,j}\mathbb{I}\{i, j \notin B\big(X^k\big)\} + \chi\mathbb{I}\{i, j \in B\big(X^k\big)\}$$

results in a low payoff for all agents:

$$\forall i \in N : U_i\big(X'', 0_{n,n}\big) < f\big(ng\big(x^*\big)\big) - cnx^* + r(1, n).$$

The sequence of outcomes that results from this construction enforces $(X^*, 0_{n,n}) \triangleright (X', 0_{n,n})$. □

Theorem 5.6 is qualitatively different from the results obtained in the literature. The model in which collaboration is costless (recall that, in this model, the moderate amount of collaboration is beneficial for its participants even if they are ranked very low in the tournament) is considered a simple case in the literature: either the efficient outcome is guaranteed to be pairwise stable or all agents exert an inefficiently large collaborative effort. In both cases, the pairwise stable outcomes are symmetric. If links are moderately costly, the efficient outcome is stable, but there may be inefficient outcomes that are also stable. Theorem 5.6 says that even in a simple case with costless links, efficiency is incompatible with stability if potential gains from competition are high. More precisely, the efficient outcome can neither be singleton stable nor can it coexist with any other outcomes in any stable set.

The condition (3) is equivalent to $R(1) \le R^*$, where

$$R^* = \min_{\frac{n}{2}<m<n} \left\{ \frac{mn}{n-m} \big(v(n,0) - v(m,n-m) \big) \right\} \ge 0$$

for winner-takes-all tournaments. It requires that the tournament prize is low compared to the direct value of collaboration. Theorem 5.6 suggests that if this condition is not satisfied, a strict subset of agents would be willing to sacrifice some collaboration in exchange for the top tournament ranking. A collective tactic that achieves top tournament rankings for a large group of agents has a maxmin property. By following this tactic, the agents obtain top rankings no matter what outsiders do.

The set of outcomes described in Theorem 5.3 plays an important role in Theorem 5.6, as suggested by condition (3). Suppose that (3) does not hold. If agents are in a stable set $\mathcal{P}$, either $\mathcal{P} = \mathcal{R}$ (and an efficient outcome is not inside the set) or the outcomes in $\mathcal{P}$ block the outcomes in $\mathcal{R}$. In the latter case, these outcomes either block or are blocked by an efficient outcome.

In mainstream models of tournaments with costly effort and no possibility of collaboration, equilibrium outcomes are usually inefficient. In those models, every agent in the efficient outcome has an individual incentive to raise his effort and collect a higher tournament prize. My model rules out this source of inefficiency: Agents are capable of coordination and individual incentives yield to collective interests. Indeed, if one removes the possibility of collaboration from the current model, the outcome in which every agent exerts an effort,

$$x^a = \arg\max_{x \ge 0} \big\{ f\big(g(x)\big) - cx \big\},$$

is a singleton stable set. Therefore, the inefficiency highlighted in Theorems 5.3 and 5.6 is caused by agents' cooperative behavior in the presence of competition.

5.2 *The role of transfers*

A common intuition suggests that the ability to use transfers should allow agents to reach an efficient outcome and stay in it. As shown in Theorem 5.3, when tournament prizes are large, agents create gaps in stable collaborative networks to sustain the difference in rankings between the fully connected majority and the rest of the population. This difference in ranking results in an extra payoff. One may argue that the minority can offer transfers to the majority in exchange for missing links; it is possible for agents to emulate an unequal division of tournament prizes through a system of transfers, while enjoying the maximum value of collaboration. However, this does not always happen in stable outcomes. More precisely, the outcomes that I find in Theorem 5.3 are stable independent of whether agents can use voluntary bilateral transfers.

Remark 5.7. Suppose that the set of feasible outcomes is

$$\mathcal{U}_0 = \{(X, T) \in \mathbb{R}_+^{N\times N} \times 0_{n,n}\}.$$

Outcome $\mathcal{R}$ is a stable set if every outcome $(X, T) \in \mathcal{R}$ has a group structure induced by a group-optimal sequence.

This remark follows directly from the proof of Theorem 5.3.

Note that a central planner can easily implement an efficient outcome by collecting the total surplus in every outcome and redistributing it uniformly across agents. However, if the process of setting up transfers is decentralized, the efficiency of a stable outcome is guaranteed only when the prizes in the tournament are small. In that case, transfers play no role, as stated in Remark 5.7.

Transfers do not necessarily help with efficiency and do not realize potential gains from trade, because they lack endogenous credibility. A minority may pay a majority to restore missing links, but there exists a similar outcome in which members of the minority swap roles with some members of the majority: the latter should pay the former. Note that this argument does not rely on symmetry; even in a model with moderate heterogeneity, agents are imperfect substitutes for each other and the same argument applies. Alternatively, one may think of this situation as a competition a là Bertrand, in which every agent is both a buyer and a seller of missing links.

The lack of credibility is neither a general property of transfers nor an artifact of the solution concept. It is tournament-induced externalities that make transfers endogenously noncredible. To see that transfers may be endogenously credible in similar environments without externalities, consider the following modification of the model. For simplicity, suppose that there are $n = 2$ agents who have an opportunity to collaborate with each other. Let the prize in the tournament be zero. Also, suppose that the agents differ in terms of their cost of effort. In particular, suppose that $c_1 > c_2 > 0$. For simplicity, assume that for all x, $f(x) = x$. If transfers are allowed, an efficient outcome (X^*, T) must satisfy $X^*_{1,2} = X^*_{2,1}$ and

$$X^*_{1,2} = \arg\max_{x\geq 0}\left\{g(x) - \frac{c_1 + c_2}{2}x\right\}.$$

It is very costly for agent 1 to collaborate at this level. For agent 1, the optimal choice of $X_{1,2}$ is

$$X^1_{1,2} = \arg\max_{x \geq 0}\{g(x) - c_1 x\} < X^*_{1,2}.$$

Similarly, for agent 2, the optimal choice of $X_{2,1}$ (conditional on agent 1 fully reciprocating) is

$$X^2_{2,1} = \arg\max_{x \geq 0}\{g(x) - c_2 x\} > X^*_{1,2}.$$

If transfers are not allowed, for any $x \in [X^1_{1,2}, X^2_{2,1}]$ that is individually rational for agent 1, i.e., that satisfies $g(x) - c_1 x \geq 0$, an outcome $(X, 0_{2,2})$ such that

$$X = \begin{pmatrix} X^*_{1,1} & x \\ x & X^*_{2,2} \end{pmatrix}$$

is a singleton stable set. All of these outcomes, except for at most one, are inefficient. However, if transfers are allowed, none of these networks of collaboration remains stable except for the efficient ones (X^*, T). In this case, agents use transfers to exploit gains from trade (i.e., to compensate agent 1 for the extra effort he is exerting in the efficient outcome) and to depart from inefficient outcomes. Transfers do restore efficiency in this simple example,[13] but do not necessarily do it in the main model, because a tournament introduces an externality that large groups of agents may exploit to divert surplus from the remaining agents.

It is useful to look at the main model from a coalition-formation perspective. The conditions for a nontrivial group-optimal sequence in Theorem 5.3 have the same flavor as the condition of unbalancedness in the Bondareva–Shapley theorem for total utility (TU) games (with the difference that once the group is formed in my model, group members choose their collaborative effort levels). Indeed, define a coalition as a set of all agents who belong to the same component and define the value of that coalition as a maximum sum of agents' utilities. In this case, for a coalition S of size $m > n/2$, the value is $V(S) = m[r(1, m) + v(m, n-m)]$. The condition for the nontrivial group-optimal sequence $\{m_k\}_{k=1}^K$ boils down to the absence of balance, $V(S)/|S| > V(N)/n$. When considering small coalitions, one must recall that the value of a coalition depends, in general, on the whole coalition structure; therefore, when computing the value, one must assume that other coalitions are structurally sound, i.e., they do not want to merge or split.

The model is not neutral to the introduction of transfers. Transfers may allow agents to exploit some gains from trade if these gains are not associated with externalities. Also, the presence of transfers imposes a restriction on the payoffs in various outcomes inside a stable set. In particular, in the presence of transfers, outcomes in any stable set must induce at least two distinct vectors of payoffs.

Theorem 5.8. *If the group-optimal sequence is such that $m_1 < n$, there exists no stable set $\mathcal{R}$ such that for any $(X, T), (X', T') \in \mathcal{R}$ and for all $i \in N$, $U_i(X, T) = U_i(X', T')$.*

[13] This result holds for an arbitrary number of agents and an arbitrary increasing function f.

Proof. I show that any set of outcomes characterized by a single payoff vector necessarily violates external stability.

Take a set $\mathcal{R}$ such that for any $(X, T), (X', T') \in \mathcal{R}$ and for all $i \in N$, $U_i(X, T) = U_i(X', T')$. Without loss of generality, assume that agents are enumerated in such a way that $i > j$ implies that $U_i(X, T) \geq U_j(X, T)$. Note that $nV_1 > \sum_{i\in N} U_i(X, T)$.

I construct an outcome $(\hat{X}, \hat{T})$ such that it is not blocked by any outcome in $\mathcal{R}$. Partition a set $\{1, \dots, n\}$ into two sets, $N_1 = \{1, \dots, m_1\}$ and $N_2 = \{m_1+1, \dots, N\}$, and consider an outcome $(\hat{F}, \hat{T})$ such that

(i) $\hat{X}_{i,j} = x_1 \mathbb{I}\{\{i, j\} \subset N_1\}$

(ii) $\hat{X}_{i,j} = 0$ implies $T_{i,j} = 0$

(iii) for all $i \in N_1$, $U_i(\hat{X}, \hat{T}) > U_i(F, T)$.

There always exists a system of transfers that satisfies condition (iii), because

$$\frac{1}{m_1}\sum_{i=1}^{m_1} U_i(X, T) \leq \frac{1}{n}\sum_{i=1}^{n} U_i(X, T) < V_1 = \frac{1}{m_1}\sum_{i=1}^{m_1} U_i(\hat{X}, \hat{T}).$$

By construction, for any $S \subset N_2$ and for all (X', T'), $(\hat{X}, \hat{T}) \xrightarrow{S} (X', T')$, $\mathbf{U}_{N_1}(X', T') = \mathbf{U}_{N_1}(\hat{X}, \hat{T}) > \mathbf{U}_{N_1}(X, T)$. Therefore, $(\hat{X}, \hat{T})$ is not blocked by any outcome that induces the payoff vector $\mathbf{U}(X, T)$. □

This theorem, when applied to efficient outcomes, dictates that when the optimal group sequence is nontrivial, an efficient outcome cannot constitute a singleton stable set.

5.3 *Extensions and special cases of the model*

One important feature of the current model is that the existence of stable outcomes with group structure does not rely on the assumption of costly collaboration. Formally, the current model does not cover the case in which $c = 0$, because the agent's maximization problem (2) does not have a solution. However, one can extend the model to accommodate this case. This extension of the model generalizes the example presented in Section 3.

Suppose that $c = 0$. Allow the agents to choose the infinite effort, set $g(\infty) = \overline{g}$, and normalize $\overline{g}$ to be 1. For simplicity, assume that agents cannot exert any intermediate effort level. i.e., for all i, j, $X_{i,j} \in \{0, \infty\}$. In this case, a collaboration can be fully described by an undirected graph $G \in \{0, 1\}^{n\times n}$.

The payoff of agent i in outcome (G, T) is

$$U_i(G, T) = r\big(p_i(G), q_i(G)\big) + f\big(y(G, i)\big) + \sum_{j\in N}(T_{j,i} - T_{i,j}),$$

where $y(G, i) = \sum_{j\in N} G_{i,j}$ is the output of agent i in outcome (G, T). All of the other definitions carry over to this extension without modification.

A group-optimal sequence $\{m_k\}_{k=1}^K$ solves

$$m_k \in \operatorname*{arg\,max}_{\frac{n-M_{k-1}}{2}<m\leq n-M_{k-1}} \big\{r(1+M_{k-1}, m+M_{k-1})+f(m)\big\}.$$

In this extension, the choice of collaborative intensity is limited; therefore, when defining the group optimal sequence, I can omit the first maximization problem that defines function v. Similar to the main model, the set of all networks that have a group structure induced by a group-optimal sequence is a stable set, independent of whether transfers are allowed. Also, in the winner-takes-all tournaments, the efficient outcome, which is a complete network of collaboration, belongs to a stable set if and only if the group optimal sequence is $\{n\}$,

$$n \in \operatorname*{arg\,max}_{\frac{n}{2}<m\leq n}\big\{r(1,m)+f(m)\big\}$$

or, equivalently,

$$R(1) \leq \min_{m>\frac{n}{2}}\left\{\frac{mn}{n-m}\big(f(n)-f(m)\big)\right\}.$$

These results extend Theorems 5.3 and 5.6 and Remark 5.7 to the case of costless collaboration.

This version of the model can also be extended to allow for the agent's output to depend on his indirect connections. Given a network of collaboration G, let each agent i produce output $y(G,i)$, which depends on the amount of indirect collaboration in which the agent is involved, i.e.,

$$y(G,i)=\sum_{k=0}^{\infty}\sum_{j=1}^{n}\alpha_k\big(G^k\big)_{ji},$$

where α_k represents the weight that is assigned to an indirect collaboration with agents who are k connections away from agent i. I normalize $\alpha_0=\alpha_1=1$ and I assume that α_k is decreasing in k. There are two special commonly used cases for this formulation: (i) when $\alpha_k=0$ for all $k>1$, the output is equal to the degree of the agent in G, and (ii) when $\alpha_k=\alpha^k$, the output is equal to the Katz centrality measure of node i in network G.

As shown in the Supplemental Appendix, the results remain qualitatively the same compared to the case in which only direct connections contribute to the agents' output. This happens for two reasons. First, indirect connections are assumed to contribute less than direct ones (α_k is decreasing in k). Second, an agent's output does not depend on the connections of agents who do not belong to the same component ($(G^k)_{ji}=0$ for all k if agents i and j belong to different components of network G). Intuitively, Theorem 5.3 characterizes the stable set of outcomes in which agents are connected if and only if they have the same payoff and ranking. From this perspective, indirect connections are not different from direct ones: If there are two indirectly connected agents who have different payoffs, there must exist two directly connected agents who have different payoffs. The formal statements and proofs of these results are relegated to the Supplemental Appendix.

Another interesting special case of the model is when $f(z) = 0$ for all z. This is the opposite of the case with costless links, and it corresponds to the situation in which collaboration is costly and has no direct benefits to the participants. Therefore, the only reason that agents would want to collaborate is to gain an advantage over their competitors in the tournament. When collaboration provides no direct benefit to the participants, there is no welfare loss from the fact that agents are not collaborating between the groups in the stable set characterized in Theorem 5.3. However, there are welfare losses from the excessively intensive collaboration within the larger groups. In this case, the role of inequality (1) is particularly stark: It puts a lower bound on the amount of inefficiency in this stable set. The trade-off between making a dominant group smaller or larger becomes a trade-off between a higher expected ranking and a smaller cost of dominating the remaining agents in the tournament. This case is extensively studied in the literature on R&D collaboration. As in this literature, the smallest group of agents in this stable set does not collaborate at all ($x_K = 0$).

5.4 *Other forms of competition*

Unfortunately, studying a model that nests several forms of competition is quite difficult, for both theoretical and expositional reasons. Here I present a very simple example that points in the direction of a condition on payoffs that one can use to extend the results of this paper to other forms of competition, such as Tullock contests or monopolistic, Cournot, or Bertrand competition. Using the same example, I argue that stable sets may not exist if this condition is not satisfied.

Consider three agents who are competing with each other in a Tullock contest for a prize R. Each agent chooses a vector of collaborative efforts. For simplicity, I assume that any effort level is restricted to either 0 or 1, and therefore collaboration between the agents can be fully described by a graph $G \in \{0, 1\}^{3\times 3}$. As in the main model, $G_{i,i}$ represents agent i's effort toward working solo. The per-unit cost of effort is c and agent i's output is

$$y_i(G) = \Big(\sum_j G_{i,j}\Big)^\alpha,$$

where $\alpha > 0$, and the direct value of the agent's output is zero. To further simplify this example, I assume that the agents cannot use transfers.

The probability of agent i winning the contest (and getting the prize R) is proportional to his output; therefore, the payoff of agent i is

$$U_i(G) = \begin{cases} \dfrac{\Big(\sum\limits_j G_{i,j}\Big)^\alpha}{\sum\limits_k \Big(\sum\limits_j G_{k,j}\Big)^\alpha} R - c\sum\limits_j G_{i,j} & \text{if } \sum\limits_{k,j} G_{k,j} > 0, \\ \dfrac{1}{3}R & \text{otherwise.} \end{cases}$$

Agent i's expected share of the prize increases in his own degree and decreases in the average degree of all agents. In contrast to the main model, the agent's payoff is sensitive to the efforts of *all* other agents in *any* outcome.

For the purpose of this example, I consider two cases: A set of outcomes, each of which features a group structure (similar to the one in Theorem 5.3), is stable if $\alpha = 1$; the same set is not stable if $\alpha = 3$. The difference between the two cases is due to a condition that enables the internal stability of this set in the case of $\alpha = 1$. Intuitively, this condition ensures that the agents in the smaller group can maximize their own payoffs by minimizing the payoffs of the members of the larger group. A similar condition is satisfied in the main model.

If $\alpha = 1$, there exists a stable set that is similar to the one characterized in Theorem 5.3. For instance, if $R > 30c$, the set of all outcomes in which only two out of three agents collaborate with each other is stable.[14] This example satisfies the condition used in Grandjean and Vergote (2015): The agent's payoff is increasing in his degree and decreasing in the degree of others. In pure network-formation models, this condition is sufficient for the existence of stable sets of outcomes with dominant-group architecture. Note, however, that this condition is not necessary and is not satisfied in the main model of this paper.

If $\alpha = 3$, the results obtained in the main model no longer hold. Let $R \in (12c, 17c)$. In this case, the group-optimal sequence is $\{2, 1\}$, but any set of outcomes in which only two out of three agents collaborate is unstable. To see this, consider the three outcomes

$$G = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad G' = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \text{and} \quad G'' = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Agent 3 strictly prefers outcome G over G'', and, therefore, G blocks G'' and G'' does not block G. Thus, an externally stable set that consists of outcomes with only two agents collaborating must include at least two outcomes that are obtained by a permutation of G or G'. However, any such set does not satisfy internal stability. For example, agent 3 can initiate a blocking transition from G to a permutation of G': At G, he may unilaterally lower the payoffs of agents 1 and 2 by increasing his solo effort.

Moreover, if $\alpha = 3$, stable sets may fail to exist altogether.[15] The difference between this paper's results and those of Grandjean and Vergote (2015) suggests that the existence of stable sets can be obtained by restricting either a payoff structure (i.e., as in this paper, by specifying a particular form of competition) or a set of feasible outcomes (i.e., as in Grandjean and Vergote 2015, by focusing on a pure network-formation model).

If $\alpha = 3$, the agent's solo effort level that maximizes his payoff is different from the effort level that minimizes the payoffs of his rivals. Because of this difference, large externally stable sets tend to be internally unstable. If $\alpha = 1$ (or if agents participate in a tournament rather than a Tullock contest), these two effort levels are the same.

[14] In every outcome of this stable set, all agents exert the maximum effort toward working solo.

[15] For instance, if $R \in (12c, 17c)$, stable sets do not exist in this example.

To fully generalize the results of Theorem 5.3, one must first provide a general sufficient condition for the existence of farsighted stable sets. This important criterion is currently unresolved, even in a narrower class of models such as characteristic function games (see Ray and Vohra 2015), and is beyond the scope of this paper.

6. Discussion of the results

This paper contributes to the literature on network formation and its applications to R&D collaboration, discrimination, and tournaments.

The paper is closely related to Goyal and Joshi (2003), Goyal and Moraga-Gonzalez (2001), and Marinucci and Vergote (2011). These papers develop models of R&D collaboration between market competitors. In these models, firms can resort to joint research to save on R&D costs. The common finding in this literature is that networks of collaboration that consist of several components are possible in equilibrium.

My model produces several important results that do not appear in the literature on R&D collaboration. First, I argue that under certain conditions, efficient outcomes may be unstable. More precisely, I provide a necessary and sufficient condition for the existence of a farsighted stable set that contains an efficient outcome. If this condition is not satisfied, efficient outcomes cannot be stable. Results in the previous literature often do not rule out efficient networks as equilibrium outcomes under similar conditions.

Second, in my model, the sizes of the complete components in stable networks are uniquely determined by the shape of payoff functions, whereas in Goyal and Joshi (2003) and Marinucci and Vergote (2011), the local incentives of individual agents put bounds on the sizes of the components. Also, the mechanics of my model are different from those in prior papers on R&D collaboration. In Goyal and Joshi (2003), a link is missing from a stable outcome because forming it is individually costly for at least one of the two nodes.[16] Decreasing the cost of the link leads to larger stable components. In particular, if one assumes that links are beneficial rather than costly, the unique stable outcome is a complete network. In my paper, the links are missing because of the positive externality on the rest of the agents. Therefore, even when links are beneficial, complete networks may not be stable.

In addition, it is worth pointing out that Goyal and Joshi (2003), Goyal and Moraga-Gonzalez (2001), and Marinucci and Vergote (2011) model competition differently from this paper (the closest being Marinucci and Vergote 2011, who model competition as a winner-takes-all tournament with stochastic outcomes). Finally, the literature focuses mainly on the case of pure network formation, whereas my model allows for a richer description of collaborative relationships between agents.

This paper helps to explain the difference in results between R&D models that use a coalition- or network-formation approach. An extensive literature on collaboration between firms looks at coalitions of firms rather than at bilateral agreements between them (e.g., Bloch 1995, 1996, Yi 1998, 1997, Yi and Shin 2000, and Joshi 2008). Surveys of the literature can be found in Bloch (2002) and Ray (2007). The predictions obtained in

[16]Other papers on networks of R&D collaboration, such as Goyal and Moraga-Gonzalez (2001) and Marinucci and Vergote (2011), share this feature with Goyal and Joshi (2003).

this literature are different from the findings obtained in the network-formation models discussed above. In particular, the grand coalition (which is the analog of the complete network) is usually not stable, because there exists a smaller coalition that prefers to reduce the amount of collaboration in exchange for greater market power. For example, Bloch (1995) employs a dynamic game in which firms sequentially propose to form alliances to reduce the marginal cost of production. Once the alliances are formed, firms engage in Cournot competition. Bloch (1995) shows that the alliance structure in the market is usually asymmetric and inefficient. These results are obtained under the assumption that participation in a coalition is exclusive. I obtain similar results, but I do not use the exclusivity assumption: In my paper, groups are *endogenously* exclusive. Therefore, my model is useful for understanding the relationship between coalition- or alliance-formation and network-formation models.

Grandjean and Vergote (2015) consider a network-formation model in which the agent's payoff is increasing in his own degree and decreasing in the degree of his competitors. They show that if the payoff of any two agents with the same degree always increases when they are connected by a link and if the payoff of agents in a small clique increases in the size of the clique, there exists a stable set of networks. These networks are either two-clique networks or dominant-group networks. In contrast to Grandjean and Vergote (2015), this paper looks at a particular form of competition—tournaments—but allows for a richer set of actions available to agents. It also provides necessary and sufficient conditions for the stability of efficient outcomes in winner-takes-all tournaments.

My theoretical findings successfully capture some properties of collaborative networks that are observed in practice. One salient illustration that supports my theoretical results is a study of the early GSM market by Bekkers et al. (2002), who examine the emergence of GSM technology in the 1990s. They document that large portfolios of standard-essential patents for GSM technology were owned by several companies: Nokia, Motorola, Alcatel, Phillips, Bull, Telia, and others. Five of these companies—Ericsson, Nokia, Siemens, Motorola, and Alcatel—signed numerous cross-licensing agreements that allowed them to use each other's patents without paying royalties. This network of cross-licensing agreements provided its participants with a market advantage over firms that were not included. Not surprisingly, the same five companies later dominated the market for GSM infrastructure and terminals, having a total market share of 85% in 1996. At the same time, three other companies—Phillips, Bull, and Telia—held roughly as many patents as Alcatel, but were not able to convert them into a significant market share. Moreover, they performed worse than Ericsson and Siemens, which had considerably smaller patent portfolios, yet were ranked the largest and third-largest GSM companies, respectively, in 1996.

My model suggests that if the stakes in the winner-takes-all competition are high enough, the efficient network of collaboration, in which agents sign all available collaborative agreements, is not stable. Moreover, there are stable networks in which a group of firms that dominates the market (let us call them insiders) does not collaborate with other, outsider firms. Despite the fact that this tactic destroys the value of collaboration between insiders and outsiders, it is profitable for the insiders because it allows them to maintain their dominant position in the market. Indeed, Bekkers et al. (2002) claim that the structure of cross-licensing agreements in the GSM industry in the 1990s, directed by Motorola, was instrumental in crowding out potential rivals such as Phillips. This story is not unique; for instance, the 2009–2013 smart phone patent war had similar features.

More generally, my model provides several important insights into such phenomena as patent wars and other types of market competition outside of the price domain. First, bilateral agreements such as cross-licensing are a powerful instrument in shaping a landscape for future market competition. For instance, they can be used to create persistent asymmetric market outcomes in symmetric environments. Second, if the stakes in the competition are high, asymmetric inefficient outcomes (e.g., an inefficient level of cross-licensing) are inevitable. Finally, the prospect of these outcomes forces firms to join exclusive alliances in which bilateral agreements play the role of a skeleton that holds alliances together.

My results relate to the program proposed by Salop and Scheffman (1983), who state that firms can capture the market by increasing the costs of production for their rivals. In another paper, Salop and Scheffman (1987) describe various strategies that firms can use to raise their competitors' costs. They find that some of these strategies can be more effective than predatory pricing. For instance, a coalition of firms can use the mechanism described in my model to gain control over the market. This coalition does not need to engage in predatory pricing to raise the joint share of the market; instead, it can limit access to its intellectual property and, hence, create a competitive advantage for its members.

The findings in my paper complement the results in the literature on sabotage in tournaments. Lazear (1989), Chen (2003), and Konrad (2000) suggest that agents may sabotage their rivals if the cost of sabotage is low. I argue that if costs are large, agents still can sabotage their rivals, but they must coordinate their actions to save on costs. This gives rise to a collective sabotage. I show that when the competition is for a large prize, collective sabotage is self-enforcing and often unavoidable, i.e., it takes place in every stable outcome.

Another application of my model is related to the theory proposed by McAdams (1995), who suggests that racial discrimination in the United States is fueled by the desire to maintain the gap in social status between the white majority and the ethnic and racial minorities. According to McAdams, if people value high social status, they may sacrifice mutually beneficial interracial interactions so as to gain higher status. Note that in this theory, race is a marker that is irrelevant for the fundamental economic characteristics of agents. However, since it is easily observable, it is convenient to use it for specifying social norms that support the difference in social status. In other countries, in which the population is more racially homogeneous, other markers, such as nationality, ethnicity, or religion, are used for discrimination. Sometimes the markers are almost artificial and are not derived from any observable characteristics of an individual. Examples of such markers are the castes in India, Pakistan, Nepal, and Sri Lanka.

McAdams (1995) provides evidence that discrimination is often sustained through threats of exile. If a member of a discriminating majority interacts with members of a discriminated minority, he or she risks being ostracized. My paper provides a mechanism for sustaining such social norms when agents are allowed to undertake collective deviations from the social norm. Despite the fact that my theoretical findings are qualitatively the same for any number of agents, the model is better suited to small communities in which coordination between individuals is easier to implement.

7. Conclusion

This paper proposes a model of bilateral collaboration between farsighted agents in tournaments. The model sheds light on the tension between agents' objectives to outperform their rivals and to obtain as much help from their rivals as possible. When tournament rewards are large, this trade-off is resolved in favor of the former objective: In stable outcomes, agents engage in fewer collaborative relationships than required by efficiency. A refusal to engage in efficiency-improving collaboration serves an important purpose: It allows some agents to secure high rankings in the tournament. In the stable outcomes I find, missing collaboration is not arbitrary. Agents endogenously sort into several groups of different sizes and refuse to collaborate with anyone who belongs to smaller groups. As a result, the network of collaboration consists of multiple complete components. I characterize the size of each group and the intensity of within-group collaboration in these outcomes.

The other main contribution of this paper is a necessary and sufficient condition for the stability of efficient outcomes in winner-takes-all tournaments. I find that the unique efficient outcome is not stable whenever the tournament prize is large enough. This result supports the observation that agents may collectively sacrifice collaboration to obtain higher rankings in tournaments. In fact, this result suggests that if agents sufficiently value high tournament rankings, such destructive behavior is unavoidable.

I also find that the ability to use transfers to compensate for missing collaboration does not necessarily restore efficiency. More precisely, there are stable outcomes in which there are gains from trade (i.e., in which restoring a missing collaborative link generates a surplus), but agents cannot agree on a self-enforcing system of transfers that is compatible with efficiency.

The setup of my model is close to that of existing models of network formation, but the results I obtain are more in agreement with results in coalition-formation models. Therefore, my paper contributes to settling the differences between conflicting results in these two strands of the literature.

The results in this paper can provide insights into many seemingly unrelated phenomena, ranging from R&D collaboration to discrimination and promotion tournaments. Even though the model is relatively stylized, I believe that it pins down a common feature that unites the aforementioned applications. In situations in which individual incentives unambiguously point to an efficient outcome, there is still scope for inefficiency. In the environments described above, economic agents can make proposals to many participants simultaneously—proposals that open doors for coalitional deviations. My findings suggest that when this happens, efficient outcomes may be unachievable, as there may exist a coalition that benefits from a deviation to a stable inefficient outcome.

References

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin (2012), "Dynamics and stability of constitutions, coalitions, and clubs." *American Economic Review*, 102, 1446–1476. [1089, 1095]

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin (2015), "Political economy in a changing world." *Journal of Political Economy*, 123, 1038–1086. [1089]

Aumann, Robert and Roger Myerson (1988), "Endogenous formation of links between players and coalitions: An application of the shapley value." In *The Shapley Value* (A. E. Roth, ed.). Cambridge University Press. [1080]

Bekkers, Rudi, Geert Duysters, and Bart Verspagen (2002), "Intellectual property rights, strategic technology agreements and market structure: The case of GSM." *Research Policy*, 31, 1141–1161. [1080, 1105]

Bloch, Francis (1995), "Endogenous structures of association in oligopolies." *Rand Journal of Economics*, 26, 537–556. [1080, 1094, 1104, 1105]

Bloch, Francis (1996), "Sequential formation of coalitions in games with externalities and fixed payoff division." *Games and Economic Behavior*, 14, 90–123. [1080, 1094, 1104]

Bloch, Francis (2002), "Coalitions and networks in industrial organization." *Manchester School*, 70, 36–55. [1104]

Chen, Kong-Pin (2003), "Sabotage in promotion tournaments." *Journal of Law, Economics, and Organization*, 19, 119–140. [1080, 1106]

Chwe, Michael Suk-Young (1994), "Farsighted coalitional stability." *Journal of Economic Theory*, 63, 299–325. [1080, 1089]

Diamantoudi, Effrosyni and Licun Xue (2007), "Coalitions, agreements and efficiency." *Journal of Economic Theory*, 136, 105–125. [1080]

Dutta, Bhaskar and Rajiv Vohra (2017), "Rational expectations and farsighted stability." *Theoretical Economics*, 12, 1191–1227. [1089]

Dutta, Bhaskar, Anne van den Nouweland, and Stef Tijs (1998), "Link formation in cooperative situations." *International Journal of Game Theory*, 27, 245–256. [1079]

Dutta, Bhaskar, Sayantan Ghosal, and Debraj Ray (2005), "Farsighted network formation." *Journal of Economic Theory*, 122, 143–164. [1080]

Erol, Selman and Rakesh Vohra (2014), *Network Formation and Systemic Risk.* Report, University of Pennsylvania. [1094]

Goyal, Sanjeev and Jose Luis Moraga-Gonzalez (2001), "R&D networks." *RAND Journal of Economics*, 32, 686–707. [1080, 1085, 1104]

Goyal, Sanjeev and Sumit Joshi (2003), "Networks of collaboration in oligopoly." *Games and Economic Behavior*, 43, 57–85. [1080, 1085, 1095, 1104]

Grandjean, Gilles, Ana Mauleon, and Vincent Vannetelbosch (2010), "A characterization of farsightedly stable networks." *Games*, 1, 226–241. [1080]

Grandjean, Gilles, Ana Mauleon, and Vincent Vannetelbosch (2011), "Connections among farsighted agents." *Journal of Public Economic Theory*, 13, 935–955. [1080]

Grandjean, Gilles and Wouter Vergote (2015), *Network Formation Among Rivals*. Report, CEREC. [1080, 1103, 1105]

Greenberg, Joseph (1990), *The Theory of Social Situations: An Alternative Game-Theoretic Approach*. Cambridge University Press. [1080]

Herings, P. Jean-Jacques, Ana Mauleon, and Vincent Vannetelbosch (2009), "Farsightedly stable networks." *Games and Economic Behavior*, 67, 526–541. [1080]

Jordan, James S. (2006), "Pillage and property." *Journal of Economic Theory*, 131, 26–44. [1089]

Joshi, Sumit (2008), "Endogenous formation of coalitions in a model of a race." *Journal of Economic Behavior and Organization*, 65, 62–85. [1080, 1104]

Konishi, Hideo and Debraj Ray (2003), "Coalition formation as a dynamic process." *Journal of Economic Theory*, 110, 1–41. [1080]

Konrad, Kai A. (2000), "Sabotage in rent-seeking contests." *Journal of Law, Economics, and Organization*, 16, 155–165. [1080, 1106]

Lazear, Edward P. (1989), "Pay equality and industrial politics." *The Journal of Political Economy*, 97, 561–580. [1080, 1106]

Levy, Gilat (2004), "A model of political parties." *Journal of Economic Theory*, 115, 250–277. [1095]

Marinucci, Marco and Wouter Vergote (2011), "Endogenous network formation in patent contests and its role as a barrier to entry." *The Journal of Industrial Economics*, 59, 529–551. [1080, 1085, 1095, 1104]

Mauleon, Ana, Jose J. Sempere-Monerris, and Vincent Vannetelbosch (2014), "Farsighted R&D networks." *Economics Letters*, 125, 340–342. ISSN 0165-1765. [1080]

Mauleon, Ana, Vincent J. Vannetelbosch, and Wouter Vergote (2011), "Von Neumann–Morgenstern farsightedly stable sets in two-sided matching." *Theoretical Economics*, 6, 499–521. [1080, 1095]

McAdams, Richard H. (1995), "Cooperation and conflict: The economics of group status production and race discrimination." *Harvard Law Review*, 108, 1003–1084. [1080, 1106]

von Neumann, John and Oskar Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton University Press. [1084]

Page, Frank H., Jr., Myrna H. Wooders, and Samir Kamat (2005), "Networks and farsighted stability." *Journal of Economic Theory*, 120, 257–269. ISSN 0022-0531. [1080]

Ray, Debraj (2007), *A Game-Theoretic Perspective on Coalition Formation*. Oxford University Press. [1094, 1104]

Ray, Debraj and Rajiv Vohra (1997), "Equilibrium binding agreements." *Journal of Economic Theory*, 73, 30–78. [1080, 1095]

Ray, Debraj and Rajiv Vohra (2015), "The farsighted stable set." *Econometrica*, 83, 977–1011. ISSN 1468-0262. [1088, 1095, 1104]

Salop, Steven C. and David T. Scheffman (1983), "Raising rivals' costs." *The American Economic Review*, 73, 267–271. [1106]

Salop, Steven C. and David T. Scheffman (1987), "Cost-raising strategies." *The Journal of Industrial Economics*, 36, 19–34. [1106]

Yi, Sang-Seung (1997), "Stable coalition structures with externalities." *Games and Economic Behavior*, 20, 201–237. [1080, 1104]

Yi, Sang-Seung (1998), "Endogenous formation of joint ventures with efficiency gains." *The RAND Journal of Economics*, 29, 610–631. [1080, 1104]

Yi, Sang-Seung and Hyukseung Shin (2000), "Endogenous formation of research coalitions with spillovers." *International Journal of Industrial Organization*, 18, 229–256. [1080, 1104]

Co-editor Dilip Mookherjee handled this manuscript.

Manuscript received 1 April, 2016; final version accepted 25 November, 2017; available online 28 November, 2017.