

Exploiting social influence in networks

VLADYSLAV NORA

Department of Economics, Nazarbayev University

EYAL WINTER

Department of Economics, Lancaster University and Department of Economics, and the Federmann
Center for the Study of Rationality, The Hebrew University

We study binary action network games with strategic complementarities. An agent acts if the aggregate social influence of her friends exceeds a transfer levied on the agent by a principal. The principal seeks to maximize her revenue while inducing everyone to act in a unique equilibrium. We characterize optimal transfers showing that agents who are more popular than their friends receive preferential treatment. Our main result is that under mild conditions complete core-periphery networks deliver the highest revenue to the principal. Furthermore, we show that the revenue is higher in networks where links are allocated unequally across agents. Hence, the principal benefits from creating “influentials” by linking well-connected hubs to less popular periphery.

KEYWORDS. Social networks, unique implementation, strategic complementarities, split graphs.

JEL CLASSIFICATION. C72, D82.

1. INTRODUCTION

The spread of behaviors on social networks is often mediated by external forces exemplified by firms or social media platforms. Using various incentive mechanisms, they are trying to get individuals to purchase products or adopt behaviors at maximal revenue or minimal cost. When decisions are strategic complements, such mechanisms typically induce multiple equilibria, and hence are prone to coordination failures. In this paper, we study how external forces can exploit social influence to overcome coordination failures and which networks are the most accommodating to such efforts.

We use a canonical model of a social network, where nodes represent individuals (agents), and edges represent links (friendships). An individual decides whether to take an action such as buying a product or voting for a candidate. Acting creates a positive

Vladyslav Nora: vladyslav.nora@nu.edu.kz

Eyal Winter: eyal.winter@mail.huji.ac.il

We are grateful to Francis Bloch, Andrea Galeotti, Sanjeev Goyal, Sergiu Hart, Matt Jackson, Marek Weretka, and seminar audience at OfCom UK, FCC, KSE, Hebrew University, University of Hamburg, VEST, and the participants of the conference on Game Theoretic and Behavioral Economic Insights on Social Media for their insightful comments that helped to improve the paper. Eyal Winter wishes to acknowledge financial support from the ESRC through a research grant ECA7674.

externality on her friends who took the action as well, encouraging them to act. A novel feature of our model is the degree-dependent network effects: individuals with many friends are less influenced by any one of them compared to their less popular peers. An agent acts if the aggregate influence of her friends exceeds a transfer levied on her by a principal. A set of transfers, or a mechanism, induces a binary action network game.

We study mechanisms that maximize the principal's revenue while inducing everyone to act in a unique equilibrium, i.e., optimal mechanisms. The requirement of uniqueness is a cornerstone of our analysis that captures the concern that the principal cannot coordinate agents to act in accordance with her preferred equilibrium. Indeed, experiments show that such an equilibrium tends to unravel (Devetag and Ortmann (2007)), which makes focusing on it unjustified in environments where the principal cannot shape individual beliefs, as is arguably the case in real world social networks.

Our model speaks to applications of strategic complementarities in networks. One example is a firm selling a network good. The principal is a firm who posts an individual price, and a transfer represents the price net of an intrinsic value of the good. Another example is when the principal is a social media platform willing to incentivize its users to adopt a new online application, such as a messenger, by providing a subsidy to each adopting user. A transfer is given by an adoption cost net of the subsidy. In this case, the platform minimizes the sum of subsidies, which is equivalent to maximizing the sum of the transfers. Such campaigns often start with significant persuasive efforts to recruit community leaders as early adopters with the intention that their public visibility will induce others to adopt as well. As we shall see, this feature of the incentive mechanism arises from our analysis of the model.

Our first result characterizes an optimal mechanism. Such a mechanism induces a cascade of iterative elimination of dominated strategies, leading to the outcome where everyone acts. But more importantly, the result highlights the role of relative rather than absolute degree centrality in networks. We show that being more popular than your friends guarantees preferential treatment from the principal. Such individuals pay lower transfers (or receive higher subsidies) than the rest and take the role of network leaders, allowing the principal to exploit their influence by raising transfers of their friends, who expect leaders to act and are willing to pay these higher transfers.

We apply our characterization of optimal mechanisms to study the properties of networks that make achieving coordination easier for the principal. Put differently, we are interested in networks that guarantee higher revenue. Our motivation is twofold. First, the principal might have some control over the network. For example, the newsfeed algorithm of Facebook partially determines how active a link between users is. If the majority of user i 's posts are hidden from user j and vice versa, then a nominal link would be rather inactive. Likewise, by means of friends suggestions, the platform influences the likelihood of new links emerging. Second, an interested party might want to compare exogenously given networks. For instance, to predict the commercial success of a new network product a firm must consider its associated network. Whether the product is a messenger or a specialized editing software determines the structural properties of a relevant network, and hence the product's profitability.

Our main result is that under mild conditions complete core-periphery networks guarantee the highest revenue, and hence allow for the most effective coordination. In these networks, the nodes are partitioned into two subsets, core and periphery. Every core node is connected to all nodes, and hence is a star, while every periphery node is connected only to stars. One implication of this result is that core-periphery structures ubiquitous in real life social networks are vulnerable to manipulation by external forces.

To gain further insights into the interplay between the revenue and a network structure, we first show that one network always delivers a higher revenue than another if it has more links with a small-degree agent as at least one of the endpoints. The reason is that small-degree agents are more susceptible to social influence and the principal can exploit it by charging them higher transfers. For networks with the same degree sequence, this implies that the revenue is always higher in a more disassortative network where connections tend to be between small- and high-degree agents. We also characterize networks that deliver the lowest revenue or, equivalently, require the most resources to induce action. We show that for a given number of links, these networks have as many isolated agents as possible, while other agents are tightly connected to each other. Finally, we discuss several extensions that include heterogeneous social influence and show that many insights from the original model continue to hold.

1.1 *Related literature*

The paper builds on a vast literature that studies how locally interacting individuals coordinate their actions (Contagion (2000), Jackson and Yariv (2007), Sadler (2020a)). These papers consider games where players face a simple choice of whether to adopt some behavior or not, and study how adoption levels and dynamics relate to characteristics of social interaction networks. Using the binary action framework of this literature, we explore two novel questions: what are the optimal mechanisms for solving coordination problems and which networks are more susceptible to manipulation by external forces?

Our analysis of influence mechanisms contributes to the literature on pricing and influence in networks.¹ Candogan, Bimpikis, and Ozdaglar (2012), Bloch and Querou (2013), and Fainmesser and Galeotti (2016) consider price-discriminating firms selling network goods and explore how prices and welfare depend on network characteristics. They assume a unique equilibrium, and hence no coordination problems among consumers. Belhaj and Deroïan (2019) study bilateral contracting in networks aimed at increasing the sum of agents' effort, and focus on equilibria that maximize the principal's objective. Computer science papers investigate the algorithmic aspects of pricing in networks (Hartline, Mirrokni, and Sundararajan (2008), Arthur, Motwani, Sharma, and

¹See Bloch (2016) for a comprehensive survey. An extensive literature on targeting and interventions in networks uses distinct modeling approaches. Ballester, Calvó-Armengol, and Zenou (2006) studies "key" players whose removal induces the greatest change in equilibrium aggregate action; Talamàs and Tamuz (2017) and Galeotti, Golub, and Goyal (2020) consider welfare-maximizing interventions; Bimpikis, Ozdaglar, and Yildiz (2016), Vohra (2020), and Sadler (2020b) consider influencing agents who update their beliefs in a DeGroot fashion.

Xu (2009)) and influence maximization (Kempe, Kleinberg, and Tardos (2003)).² Unlike the above papers, we explicitly address coordination problems, adopting the unique implementation approach currently unexplored in the network literature.³

Several network formation models speak to the empirical ubiquity of core-periphery networks. Hojman and Szeidl (2008) derive periphery-sponsored stars as a unique equilibrium of a network formation game where individuals benefit from indirect connections, while Galeotti and Goyal (2010) show that core-periphery networks arise as a consequence of strategic information acquisition and network formation. Belhaj, Bervoets, and Deroñan (2016) show that such networks maximize welfare when agents choose an effort level in a game of strategic complements. Our paper is complementary to this literature highlighting an unexplored property of core-periphery networks.⁴

The unique implementation approach was pioneered by Segal (1999, 2003), who develops a general contracting model, and Winter (2004), who explores incentives provision in organizations. Babaioff, Feldman, Nisan, and Winter (2012), Bernstein and Winter (2012), Halac, Kremer, and Winter (2020), and Halac, Lipnowski, and Rappoport (2021) are prominent papers in this vein. We contribute to this literature by incorporating local externalities captured by a social network.

The paper is organized as follows. We begin with the model and an example in Section 2. In Section 3, we characterize optimal mechanisms. In Section 4, we compare revenue across networks, while in Section 5 we characterize optimal networks. In Section 6, we discuss robustness of our results. All proofs are presented in the Appendix.

2. MODEL

2.1 Setup

There are n individuals (agents) indexed by $i = 1, 2, \dots, n$. Each individual i decides whether to act ($x_i = 1$), or not ($x_i = 0$). Individuals interact through a social network, represented by an undirected graph with a symmetric adjacency matrix G , where $g_{ij} = 1$ if and only if i and j are connected (friends) and $g_{ij} = 0$ otherwise; by convention $g_{ii} = 0$. We let d_i denote the number of friends of i , i.e., $d_i = \sum_j g_{ij}$. Individuals are prone to social influence that affects their incentives to take the action; they are encouraged to act when more of their friends do. Specifically, given a network G and an action profile $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$, we normalize the payoff of individual i from abstaining, $x_i = 0$, to zero, i.e., $U_i(0, x_{-i}, G) = 0$, and let the payoff from taking the action, $x_i = 1$, be

$$U_i(1, x_{-i}, G) = f(d_i) \sum_j g_{ij} x_j - t_i. \quad (1)$$

²A typical problem is to find k nodes such that if these nodes act, eventually the highest number of other nodes also choose to act. By contrast, we look for a profile of thresholds such that in the unique equilibrium everyone acts and the sum of the thresholds is maximal.

³The literature on network goods uses adoption-contingent prices (Weyl (2010), Masaki (2013)), whereas we study bilateral contracting where transfers are not contingent on actions of others.

⁴Other notable examples are Bala and Goyal (2000), Goyal and Joshi (2003), Goyal, Van Der Leij, and Moraga-González (2006), König, Testone, and Zenou (2014), Hiller (2017), and Herskovic and Ramos (2020).

The term $\sum_j g_{ij}x_j$ is the number of friends of i who choose to act, and $f(d_i) > 0$ captures a social influence exerted on i by each such friend. Hence, the payoff from acting is linearly increasing in the number of friends who act. We call f a *social influence function*, and assume that it is nonincreasing in the number of friends of an individual, i.e., $f(m) \geq f(m+1)$ for $m \geq 1$. The assumption reflects the idea that someone with more friends is swayed less by each one of them. The term $t_i \geq 0$ can be viewed as a threshold; an individual chooses to act when her aggregate social benefit from acting, $f(d_i) \sum_j g_{ij}x_j$, exceeds her threshold t_i . Because $f(d_i) > 0$, the resulting simultaneous move game with complete information has strategic complementarities, and, typically, there are multiple equilibria.

2.2 External influence

Consider the principal who influences individuals by choosing their thresholds. Specifically, we interpret a threshold t_i as a *transfer* from individual i to the principal. An *influence mechanism* is a profile of transfers, i.e., a vector $t = (t_1, \dots, t_n)$. Because the game between the agents might have multiple equilibrium outcomes, we require that the principal chooses the transfers to induce a unique equilibrium where all agents act. Formally, an influence mechanism t is *incentive-inducing* (INI) if $x = (1, \dots, 1)$ is a unique Nash equilibrium of the simultaneous move game induced by t . Clearly, such influence mechanisms exist because the principal can make acting a dominant strategy for each agent i by offering $t_i = 0$. Moreover, if t is INI, then so is each $t' < t$. However, the principal also wants to maximize the revenue while inducing action. Influence mechanism t is *optimal* if it has the highest revenue among all INI mechanisms, i.e., $\sum t_i \geq \sum t'_i$ for each INI mechanism⁵ t' . The maximal revenue that the principal can achieve while incentivizing all agents to act depends on a social network. Networks that deliver a higher revenue are more attractive to the principal. A network is *optimal* if its optimal influence mechanism has the highest revenue across all networks.

In the remainder of this section, we present an example based on a social influence function naturally arising in certain settings. We use the example to illustrate the construction of optimal influence mechanisms and compare the revenue across networks.

EXAMPLE 1. Consider the case where individuals directly care about a proportion and an absolute number of friends who take the action.⁶ Then the social influence of each such friend on agent i is given by

$$f(d_i) = \alpha + \frac{1}{d_i},$$

⁵Note, however, that a set of INI mechanisms is not closed, so an optimal INI mechanism may not exist. Let $I \in \mathbb{R}^n$ be a set of INI mechanisms and $\bar{I} \in \mathbb{R}^n$ be its closure. Formally, we say that influence mechanism t^* is optimal if $t^* \in \arg \max_{t \in \bar{I}} \sum t_i$. Hence, although our optimal mechanism t^* may admit multiple equilibria, for every $\varepsilon > 0$ there exists an INI mechanism t' whose revenue is only ε smaller, i.e., $\sum t_i^* - \varepsilon = \sum t'_i$.

⁶The examples of such situations studied in the literature include models of social comparison (Ghiglino and Goyal (2010)) and conformity (Liu, Patacchini, and Zenou (2014)). For example, Ghiglino and Goyal (2010) discuss two situations, when local aggregate action matters and when local average matters, and provide justification for each case. Both cases are subsumed by our model.




FIGURE 1. Comparing the principal's revenues across networks.

where $\alpha \geq 0$ is a constant part of the social influence from an acting friend. When α is small an individual cares mostly about the relative proportion of acting friends, whose number becomes important as α grows. We begin by illustrating the construction of optimal influence mechanisms in each of the networks in Figure 1 and then compare the corresponding revenues.

To begin, note that, in any network, the principal must induce at least one of the agents to act even when no one else does (otherwise there will be an equilibrium in which no one acts). Hence, the transfer of one of the agents must be at most zero. Now consider a complete network in panel (a). All agents are symmetric in the network and, therefore, we can let agent 1 pay $t_1 = 0$. Moreover, one of the remaining agents must pay at most $\alpha + 1/3$; otherwise there is an equilibrium where only agent 1 acts. Again, by symmetry we can let $t_2 = \alpha + 1/3$. Similarly, to induce one of the two remaining agents to act when both 1 and 2 act, the principal must ask for a transfer of at most $2\alpha + 2/3$. Let $t_3 = 2\alpha + 2/3$. Finally, agent 4 must pay at most $t_4 = 3\alpha + 1$. In fact, this is an optimal influence mechanism for the complete network with the revenue of $6\alpha + 2$. Next, consider a line network in panel (b). In an optimal influence mechanism, we let $t_1 = 0$ to induce agent 1 to act when no one else does. Now to induce 2 to act when only 1 does we let $t_2 = \alpha + 1/2$. Furthermore, to induce 3 and 4 to act when 1 and 2 do, we let $t_3 = t_4 = \alpha + 1$. The corresponding revenue is $3\alpha + 5/2$. Finally, a star network in panel (c) has an optimal influence mechanism where $t_1 = 0$ and $t_2 = t_3 = t_4 = \alpha + 1$. Note that agent 2 can pay more than in the line network because all her friends now act. The revenue is $3\alpha + 3$.

Two observations are in order. First, note that the principal achieves a higher revenue in the star network than in the line network, regardless of the influence function f . We revisit the example in Section 4, where we introduce and characterize a dominance order on networks with respect to the principal's revenue. Second, the principal obtains a higher revenue in the complete network if $\alpha > 1/3$, and in the star network if $\alpha < 1/3$. Both graphs belong to a general class of core-periphery networks. In Section 5, we prove our main result that an optimal network is indeed a core-periphery. We also show that the above social influence function results in a “bang bang” solution, where generically either a star or a complete network is optimal. $\diamond$

3. OPTIMAL INFLUENCE MECHANISMS

We begin our investigation by characterizing optimal influence mechanisms for any network. We show that individuals who have more connections than their friends have are asked to pay lower transfers. This allows the principal to exploit their influence by extracting surplus from their friends.

Fix a network G . We generalize the construction of the optimal influence mechanisms given in Example 1. We say, an influence mechanism t is *tight* if it is INI and there does not exist another INI influence mechanism t' such that for some i we have $t_i < t'_i$ and $t_j = t'_j$ for all $j \neq i$. In words, if t is tight, then we cannot increase the transfer of any single agent without violating the requirement of the uniqueness of equilibrium where all agents choose to act. Clearly, an optimal influence mechanism is tight. It turns out that there is surjection between a set of permutations of agents and a set of tight influence mechanisms.⁷

LEMMA 1. *An influence mechanism $t = (t_1, \dots, t_n)$ is tight if and only if there exists a permutation π such that for all i ,*

$$t_i = f(d_i) \sum_{j: \pi(j) < \pi(i)} g_{ij}, \quad (2)$$

where $\pi(i)$ denotes a place of i in the permutation π .

Given a permutation of agents, we can use (2) to construct a tight influence mechanism. In this mechanism, every agent transfers to the principal the social benefit obtained from her friends that are earlier in the corresponding permutation. Notice that despite the fact that in equilibrium everyone acts, in a tight influence mechanism the principal does not extract the entire social benefit from each agent. For instance, in Example 1 a permutation that corresponds to the constructed optimal influence mechanism in a complete network is $(1, 2, 3, 4)$. Here, the transfer of agent 2 is equal to her social benefit from agent 1 acting, i.e., $\alpha + 1/3$. However, her equilibrium social benefit from acting is $3\alpha + 1$, and so she is left with a surplus of $2\alpha + 2/3$.

Intuitively, a permutation represents an order in which agents iteratively eliminate dominated strategies. Indeed, in any incentive-inducing influence mechanism, there must exist an agent who acts regardless of other agents' decisions. That is, acting is her dominant strategy. This agent appears first in the permutation, and so her transfer is given by her social benefit when no one else acts, namely zero. Similarly, there must exist one agent for whom acting is a dominant strategy conditional on the first agent acting (otherwise we would have an equilibrium where all agents but the first one stays still). This agent is placed second in the permutation, and pays just as much as is her social benefit given that the first agent is acting, and so on. Hence, for each incentive-inducing mechanism we can inductively construct a corresponding permutation of agents.

Given the above lemma, finding an optimal mechanism reduces to a simpler problem of maximizing over permutations. For any permutation π , the revenue in the corre-

⁷All the proofs are deferred to the [Appendix](#).

sponding tight influence mechanism $t = (t_1, \dots, t_n)$ is

$$\begin{aligned} \sum_i t_i &= \sum_i f(d_i) \sum_{j: \pi(j) < \pi(i)} g_{ij} \\ &= \sum_{ij: \pi(j) < \pi(i)} g_{ij} f(d_i), \end{aligned} \quad (3)$$

where the last line follows from rearranging the summation. Each link in a network contributes a single term to (3). Specifically, consider a link between i and j . If $\pi(j) < \pi(i)$, then the link contributes $f(d_i)$ to (3), and if $\pi(j) > \pi(i)$, then the link contributes $f(d_j)$. Because f is nonincreasing, the contribution of a link between i and j such that $d_i > d_j$ is maximal when $\pi(i) < \pi(j)$. Moreover, the contribution depends only on the relative positions of i and j in the permutation, and not on the entire permutation. Call a permutation π of agents *nonincreasing* if among any two connected agents the one with a strictly higher degree appears earlier in the permutation, i.e., for all agents i and j such that $g_{ij} = 1$ and $d_i > d_j$, we have $\pi(i) < \pi(j)$. Clearly, the contribution of each link is maximal in every nonincreasing permutation and for a link between i and j it is given by $\max\{f(d_i), f(d_j)\}$. Hence, we obtain a characterization of an optimal influence mechanism.

PROPOSITION 1. *Fix a network G . An influence mechanism $t = (t_1, \dots, t_n)$ is optimal if it is induced by a nonincreasing permutation, i.e., there exists a nonincreasing permutation π of agents such that t is given by (2). Moreover, the revenue in an optimal influence mechanism is given by*

$$R(G) = \sum_{i < j} g_{ij} \max\{f(d_i), f(d_j)\}. \quad (4)$$

Intuitively, the principal arranges agents according to their resilience to social influence from others. Well-connected individuals are more resilient in a sense that their incentives to act increase less when a friend chooses to act compared to their less connected peers. Hence, it is better for the principal to persuade them directly by offering lower transfers, and instead exploit their social influence on their less connected, and hence more easily influenced friends, who are asked to pay higher transfers. This means placing high-degree agents earlier in a permutation. Notice that this is in contrast to arranging agents according to how influential they are. For example, consider a star network in panel (c) of Figure 1. In a case of a constant social influence function $f(d_i) = \alpha$, all agents are equally resilient (everyone is influenced in the same way when a friend chooses to act), but agent 1 is clearly the most influential. However, the revenue is the same in each permutation.

One interesting observation about the principal's revenue (4) is that it is decomposable as the sum of revenues generated by individual links determined only by the “local” properties of a network, namely agents' degrees. This means that changing a network structure in one place does not affect the values of remote links. Another observation is that whereas the equilibrium surplus generated by a link between i and j is $f(d_i) + f(d_j)$,

the principal extracts $\max\{f(d_i), f(d_j)\}$. Thus, only an agent with a smaller degree matters when evaluating a value of a link for the principal. Finally, note that a transfer itself is not monotone increasing in an agent's degree. For example, among the two agents who have the same degree, the one who has more friends with a lower degree than herself is offered to pay a lower transfer, and so gets to keep a higher share of the surplus. In other words, to pay a lower transfer it is neither necessary nor sufficient to be popular—instead one needs to be more popular than one's friends.

It is useful to contrast the results here with the revenue-maximizing influence mechanisms without the unique implementation requirement. If we assume that the agents will play an equilibrium that the principal wants, then the transfers must make $x = (1, \dots, 1)$ an equilibrium. So, each agent i must be willing to act when all her friends do, i.e., for each i we must have $d_i f(d_i) - t_i \geq 0$. Hence, it suffices to let

$$t_i = d_i f(d_i) \quad (5)$$

for each i . Note that these transfers are generally higher than those in (2), allowing the principal to extract the entire equilibrium surplus of every agent. However, a strategy profile where no one acts remains an equilibrium. A gap between this higher revenue and (4) is the principal's cost of coordinating individuals using transfers.⁸

4. REVENUE COMPARISON ACROSS NETWORKS

Equipped with the characterization of optimal influence mechanisms, we now develop an intuition about how a network structure affects the revenue. Note, however, that the revenue also depends on the shape of the social influence function. Thus, to focus solely on the role of a network structure, we ask what makes one network more attractive than another, regardless of the social influence function. While the results here are of interest in themselves, they also hint at the structure of optimal networks.

We say that a network G *dominates* a network G' if and only if the revenue in an optimal influence mechanism in G is weakly higher than in G' for each nonincreasing social influence function, i.e., $R(G) \geq R(G')$ for each nonincreasing f . A network is *undominated* if there does not exist another network which dominates it. This introduces a partial order on networks, which we now relate to a network structure.

Given a network G and a positive integer k , let $l_G(k)$ be a number of links in G such that the lowest degree among the two endpoints of a link is k , i.e., $l_G(k) = |\{\text{Link between } i \text{ and } j \mid \min\{d_i, d_j\} = k\}|$. A function l_G captures the distribution of links by the lowest degree of endpoints. The line network in Figure 1 has $l_G(1) = 2$, $l_G(2) = 1$, and $l_G(k) = 0$ for each $k \neq 1, 2$, and the star network has $l_G(1) = 3$ and $l_G(k) = 0$ for each $k \neq 1$. Hence, the distribution in the star network is shifted toward lower values. It turns out that such shifts in the distribution always increase the revenue.

⁸It is easy to see that the gap is, at most, $\frac{1}{2} \sum_i d_i f(d_i)$.

PROPOSITION 2. *A network G dominates a network G' if and only if*

$$\sum_{k=1}^h l_G(k) \geq \sum_{k=1}^h l_{G'}(k)$$

for each $h = 1, 2, \dots$

The above condition resembles the definition of the first-order stochastic dominance. Loosely speaking, G dominates G' if and only if $l_{G'}$ first order stochastically dominates l_G . The proof also parallels the standard argument that the first-order stochastic dominance is equivalent to a decision maker preferring one distribution to another regardless of what her utility function is.

To get basic intuition, consider again networks in panels (b) and (c) of Figure 1. In the star network, we have replaced the link between 2 and 4 by the link between 1 and 4. This lowers the degree of agent 2 and raises the degree of agent 1. However, the new link between 1 and 4 increases the principal's revenue by the same amount as the link between 2 and 4, because the lower degree among the two of its endpoints remains the same. On the other hand, the link between 1 and 2 now contributes more because the lower degree among the two of its endpoints has decreased: the degree of agent 2 is lower than before. This holds for all nonincreasing influence functions. Therefore, the star network is indeed always better for the principal than the line network.

Proposition 2 has especially clear implications for networks with the same degree sequence. In this case, we can relate the dominance order to the network disassortativity. Recall that a *degree sequence* of a network is a list of degrees of all the nodes in nonincreasing order. Given an integer k , let $H_k(G)$ be a number of links among agents with degrees strictly greater than k , i.e., $H_k(G) = |\{\text{Link between } i \text{ and } j \mid \min\{d_i, d_j\} > k\}|$; and $L_k(G)$ be a number of links between agents with degrees weakly lower than k , i.e., $L_k(G) = |\{\text{Link between } i \text{ and } j \mid \max\{d_i, d_j\} \leq k\}|$. We say a network G is *more disassortative* than a network G' if $L_k(G) \leq L_k(G')$ and $H_k(G) \leq H_k(G')$ for each $k \geq 1$. Roughly, a network is more disassortative if it has fewer links among similar-degree agents. For example, both networks in Figure 2 have degree sequence $(3, 3, 2, 2, 1, 1)$, but a network G is more disassortative than a network G' . Indeed, high-degree agents 1 and 2, as well as small-degree agents 3 and 4 are connected in G' , whereas 1 is connected to 4 and 2 is connected to 3 in G . One can easily show that G dominates G' . The following corollary of Proposition 2 generalizes this observation.




FIGURE 2. A network G is more dissortative than a network G' .

COROLLARY 1. *Suppose G and G' have the same degree sequence. Then G dominates G' if and only if G is more disassortative than G' .*

Observe that criss-crossing the two links as in panel (a) of Figure 2 does not change the degrees and benefits the principal. Indeed now less susceptible to social influence, high-degree agents 1 and 2 are connected to more susceptible small-degree agents 3 and 4. This allows the principal to more efficiently exploit the influence of high-degree agents by raising the transfers of their small-degree friends.

We close this section with an example of an undominated network with a given number of links, which also provides an intuition about the shape of optimal networks. Consider a greedy algorithm that aims to construct a network with the highest revenue given a number of links E . The algorithm initiates with an empty network and at each step connects the highest degree agent to the smallest degree agent, terminating when it runs out of links. Hence, at each step the algorithm maximizes the value of a newly created link. Formally, at step t , let E^t be the number of unused links, let G^t be the current network with the number of links equal to $E - E^t$, let $(d_i^t)_{i=1}^n$ be the degrees of agents in G^t , let $\hat{i}^t \in \arg \max_i d_i^t$ and is such that there exists j satisfying $g_{\hat{i}^t j}^t = 0$, and let $\check{i}^t \in \arg \min_i d_i^t$. Initialize at $E^0 = E$ and G^0 being an empty network, i.e., $g_{ij}^0 = 0$ for all i and j . At step $t \geq 1$:

1. If $E^{t-1} > 0$, connect $\hat{i}^{t-1}$ with $\check{i}^{t-1}$ with whom there is no link in G^{t-1} , i.e., $g_{\hat{i}^{t-1} \check{i}^{t-1}}^{t-1} = 0$. That is, let $g_{\hat{i}^{t-1} \check{i}^{t-1}}^t = 1$ and $E^t = E^{t-1} - 1$.
2. Terminate if $E^{t-1} = 0$ and let $G^t = G^{t-1}$.

COROLLARY 2. *If $E \geq n - 1$, then the algorithm generates an undominated network.*

Note that an undominated network must be connected because an additional link to an isolated agent always increases the revenue. Thus, the condition $E \geq n - 1$. A generated network is not a unique undominated network with a given number of links. Rather the algorithm demonstrates the role of disassortativity in the emergence of undominated networks: they have a special structure where small-degree agents are never connected.

We call a network *complete core-periphery network* if its nodes can be partitioned into two subsets, S (stars) and P (periphery), such that nodes in S are connected to all the nodes, and nodes in P are connected only to the nodes in S .⁹ A star network and a complete network are the special cases, and so is the network with two stars and six periphery nodes in Figure 3. There are $n - 1$ complete core-periphery networks with n nodes (up to a permutation of nodes). A network generated by the algorithm is a (possibly incomplete) core-periphery.¹⁰ Note, however, that the greedy algorithm might fail to generate a network with the highest revenue given a number of links. The reason is that the algorithm does not account for the fact that the values of the existing links are affected by the addition of new links. In the following section, we show that under appropriate assumptions on the social influence function an optimal network is a complete core-periphery.

⁹These graphs are also known in the literature as complete split graphs or interlinked stars.

¹⁰The network may have a single periphery agent that is connected only to a subset of stars.




FIGURE 3. A complete core-periphery with 2 stars and 6 periphery nodes.

5. OPTIMAL NETWORKS

In this section, we characterize optimal networks. To pin down their precise structure, we begin by introducing two assumptions about a social influence function. First, we assume that $mf(m)$ is nondecreasing.

ASSUMPTION 1. *For each $m \geq 1$, we have $mf(m) \leq (m+1)f(m+1)$.*

The assumption is satisfied whenever f falls slower than the reciprocal function $x \mapsto 1/x$. Hence, in the case of the average comparison model from Example 1 where $f(m) = 1/m$, the above condition holds with equality.

The influence of each existing friend is diluted with an addition of a new friend. However, it is natural to think that a marginal dilution is smaller for someone who has more connections. Our second assumption formalizes the idea, requiring f to be strongly convex in the following sense.

ASSUMPTION 2. *For each $m \geq 1$, we have*

$$f(m) - f(m+1) \geq [f(m+1) - f(m+2)] \frac{n+2}{n}.$$

Strong convexity means that the absolute values of forward differences of f are decreasing at multiple $n/(n+2)$ or lower, where n is a total number of agents. Note that in a large network our strong convexity condition converges to a standard notion of convexity of f . The social influence function in Example 1 satisfies both assumptions.

PROPOSITION 3. *Suppose that Assumptions 1 and 2 hold. Then an optimal network is a complete core periphery.*

An optimal influence mechanism for a complete core periphery induces a dominance cascade that starts with the stars, each relying on the stars ahead, and finishes with the periphery agents relying on all the stars. The principal collects lower transfers from the stars, and exploits the stars' influence on periphery agents by raising their transfers. Note that complete core-periphery networks are the most unequal in terms of

their degree distributions among all similarly dense networks.¹¹ However, the principal does not necessarily strive to achieve the maximal inequality like in a star network: more links means more opportunities for the social influence, which may be exploited by the principal. To explain the role of our assumptions, we next examine an effect of a new link on the principal's revenue.

First, consider the case when we do not need to uniquely implement $x = (1, \dots, 1)$ and are content with it simply being an equilibrium. Acquiring a new friend who acts has two countervailing effects on an agent's incentives to act. On the one hand, an agent experiences the dilution of the social influence from her existing friends who act. Specifically, because f is nonincreasing each friend now exerts weakly lower social influence on an agent. On the other hand, the agent obtains one more friend who acts, and hence influences her to act as well. Specifically, according to (5), adding a link between i and j changes the revenue by

$$\underbrace{(d_i + 1)f(d_i + 1) - d_i f(d_i)}_{\text{Change in } t_i} + \underbrace{(d_j + 1)f(d_j + 1) - d_j f(d_j)}_{\text{Change in } t_j} \geq 0.$$

Assumption 1 guarantees that the change in both transfers is nonnegative. Thus, a new link always increases the revenue and a complete network is optimal.

In the unique implementation case, the principal extracts a value of a link only from a smaller degree agent, and so a new link can increase as well as decrease the revenue. To see why, consider adding a link between 2 and 3 in the star network from Example 1. The transfers of 1 and 4 are unchanged, but now 2 must pay a lower transfer $t_2 = \alpha + 1/2$ and 3 can pay a higher transfer $t_3 = 2\alpha + 1$. Generally, connecting i and j such that $d_i < d_j$ has a cost and a benefit. By Assumption 1, a transfer from a smaller degree agent rises, hence the benefit

$$(d_i^+ + 1)f(d_i + 1) - d_i^+ f(d_i) \geq 0,$$

where d_i^+ is the number of i 's friends with a higher degree than i .¹² A transfer from a higher degree agent falls, hence the cost

$$d_j^+ [f(d_j) - f(d_j + 1)] \geq 0.$$

While it is clear that convexity of f makes the cost decreasing in a degree, suggesting that having well-connected hubs is optimal, a more subtle observation is that a stronger assumption 2 is required to guarantee that the hubs are interlinked.

Now we sketch the proof of Proposition 3. Suppose that in an optimal network agent i is a friend of a smaller degree agent x . If $d_i < d_j$, then convexity of f suggests that j

¹¹We can make this statement precise using a notion of majorization. Let $d = (d_1, \dots, d_n)$ denote a sequence of degrees of agents arranged in a nonincreasing order. Then d majorizes d' if for each $k = 1, 2, \dots, n$ we have $\sum_{i=1}^k d_i \geq \sum_{i=1}^k d'_i$, with equality if $k = n$. Then a sequence of degrees of nodes in a core periphery majorizes a sequence of degrees of nodes in every network with the same number of edges.

¹²For clarity, we assume that neither i nor j have friends with the same initial degrees as themselves, and hence whose relative position in a nonincreasing permutation would change after an addition of a new link. This is without loss of generality because swapping the positions of the same-degree agents keeps a permutation nonincreasing.

would have a lower cost of a link to x , and thus must also be x 's friend. In other words, the “smaller-degree neighborhoods” of agents are nested. This observation along with some symmetry arguments implies that there is an independent set¹³ of periphery agents, and every nonperiphery agent is a hub connected to every periphery agent. Finally, we show that hubs must form a clique. Assumption 1 guarantees that a benefit of a link between two hubs outweighs its cost when hubs are in minority. The role of Assumption 2 is to make benefit increasing in d_i whenever $d_i^+ \geq n/2$. Then when a majority of agents are hubs, a benefit of a link between them is higher than a benefit of a link between a hub and a periphery. Thus, it is optimal to connect all hubs.

To illustrate Proposition 3, we briefly return to Example 1. It is clear that if $f(d_i) = \alpha$, then a complete network is optimal.¹⁴ Indeed, there is no dilution and each link has the same value for the principal regardless of the agents' degrees. Hence, an optimal network simply maximizes the number of links. In contrast, when $f(d_i) = 1/d_i$, the dilution is maximal, and a star network is optimal. We characterize optimal networks for the intermediate cases below.

COROLLARY 3. *Suppose $f(d_i) = \alpha + 1/d_i$, where $\alpha \geq 0$. We have:*

- (i) *if $\alpha < 1/(n-1)$, then a unique optimal network is a star,*
- (ii) *if $\alpha > 1/(n-1)$, then a unique optimal network is a complete network,*
- (iii) *if $\alpha = 1/(n-1)$, then a network is optimal if and only if it is a complete core periphery.*

How does the shape of the social influence function affect a number of stars in an optimal core periphery? Fix two functions f and h such that assumptions hold. Without loss of generality, let $f(1) = h(1) = 1$.¹⁵ We say that a function f is *flatter* than a function h if

$$h(k) - h(k+1) \geq f(k) - f(k+1),$$

for all $k \geq 1$. Given a function f , let $s^*(f)$ denote the smallest number of stars in a corresponding optimal core periphery.

PROPOSITION 4. *Suppose that Assumptions 1 and 2 hold. If f is flatter than h , then $s^*(f) \geq s^*(h)$.*


Intuitively, if f is flatter than h , then it admits a weaker dilution effect, and hence a new link has a lower cost and a higher benefit. Therefore, an optimal network under f has more links, while remaining a core periphery. So, it must have more stars.

Thus far, we have focused on what makes networks more attractive for the principal. To better understand how the network structure determines the revenue, we close this

¹³A set of nodes N is called *independent* if there is no link between any two nodes in N ; a set of nodes N is called a *clique* if every two nodes in N are connected.

¹⁴Although Assumption 2 is violated, an optimal network is still a complete core periphery.

¹⁵This is without loss since we can always scale the functions in order for this condition to hold.

FIGURE 4. Networks G and G' are resilient, while network G'' is not.

section with a glimpse of an opposite question—what are the networks that deliver the lowest revenue? Strategic complementarity implies that an empty network is trivially such a network, and hence we shall fix the number of links. We say a network with E links is *resilient* if its optimal influence mechanism has the lowest revenue across all networks with E links. Note that if $f(d_i) = \alpha$, then every network is resilient.

PROPOSITION 5. *Suppose that f is strictly decreasing and Assumption 2 holds. If there exists an integer s such that $E = \binom{s}{2}$, then a resilient network consists of a clique of s agents and $n - s$ isolated agents.¹⁶ If there exists an integer s such that $\binom{s}{2} < E < \binom{s+1}{2}$, then a resilient network consists of a connected component of $s + 1$ agents and $n - s - 1$ isolated agents.*

Panels (a) and (b) in Figure 4 illustrate resilient networks with 6 and 8 links and 6 nodes. When $E = \binom{s}{2}$ for some s , the result pins down a unique network, whereas it provides only a partial characterization when $\binom{s}{2} < E < \binom{s+1}{2}$ for some s . For example, one can show that a network in panel (c) is not resilient under our assumptions.

In a resilient network, all the social influence is concentrated inside a tightly interconnected group, and it has as many as possible isolated agents who pay zero transfers. We prove the result by showing that if we could reallocate all the links from the smallest degree agent to others, thus isolating this agent, then we would decrease the revenue. One way to grasp the intuition is to consider a greedy algorithm to construct a network with E links, analogous to the algorithm in Section 4. Begin with an empty network. At each step, connect two highest degree agents and terminate when out of links. Thus, at each step the algorithm minimizes the value of a newly created link. When $E = \binom{s}{2}$ for some s , the algorithm generates a clique of s agents as in the proposition above.¹⁷

6. CONCLUDING REMARKS

In this section, we briefly discuss the robustness of our results. So far, we assumed that an influence exerted by an agent varies across her friends—popular individuals are influenced less than others. Alternatively, a strength of influence may depend on the characteristics of the influencer and the same individual may be influenced differently by every friend. More generally, agents can be heterogeneous with respect to both, their

¹⁶These networks are called the dominant group architecture by Goyal and Joshi (2003).

¹⁷Note that the algorithm fails to produce a resilient network if $\binom{s}{2} < E < \binom{s+1}{2}$.

susceptibility to the influence from others, as well the influence they have on others. Some of our main insights extend to these environments.

Given a network G and an action profile $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$, we let the payoff of individual i from taking the action be given by

$$U_i(1, x_{-i}, G) = \sum_j g_{ij} w_{ij} x_j - t_i, \quad (6)$$

and the payoff from abstaining be $U_i(0, x_{-i}, G) = 0$. Thus, an agent i is influenced by an agent j according to a weight $w_{ij} > 0$. Apart from the heterogeneous influences represented by a matrix $(w_{ij})_{i,j}$, the model is as in Section 2. Whereas it is straightforward to confirm that the result analogous to Lemma 1 holds, the general model is no longer tractable.¹⁸ Instead we discuss several interesting special cases.

- *Degree-dependent susceptibility*: $w_{ij} = f(d_i)$ for each i and j and f is nonincreasing. This is our benchmark model from Section 2.
- *Degree-dependent influence*: $w_{ij} = f(d_j)$ for each i and j and f is nonincreasing. The interpretation is that an agent splits her effort between influencing each of her friends, and hence someone with more friends will influence each of them less. Optimal mechanisms are obtained from *nondecreasing* permutations of agents, i.e., permutations where among the two connected agents the one with a strictly higher degree appears later in a permutation. Moreover, the revenue is the same as in the benchmark model (provided the same network and f). Hence, our optimal network result holds.¹⁹
- *Exogenous influence*: $w_{ij} = \omega_j$ for each i and j . Individuals are heterogeneous with respect to their influence on others, but the level of influence is exogenous. We can interpret it as stemming from public credentials, such as those of political or religious leaders. Because agents can be ordered with respect to their level of influence, an optimal mechanism is characterized by permutations where the influential agents appear earlier. It is clear that in a model where social influence is independent of degree, complete networks are trivially optimal.

APPENDIX

PROOF OF LEMMA 1. Given a permutation π , let t be a corresponding influence mechanism defined by (2). We shall show that t is tight. First, we show that t is INI. Note

¹⁸The general model does not admit a characterization of optimal mechanisms analogous to Proposition 1. For example, a natural conjecture would be to order agents with respect to the “influence index” $W_i = \sum_j g_{ij} w_{ji}$. One can check that the conjecture fails in an example with 3 agents where $g_{12} = g_{23} = 1$, $g_{13} = 0$ and $w_{12} = w_{31} = 1$, $w_{21} = w_{13} = 3$. Whereas 1 is the most influential, it is optimal to put 3 at the first place in the permutation.

¹⁹If f is nondecreasing, then an optimal influence mechanism in a model with degree-dependent susceptibility (influence) is obtained from a nondecreasing (nonincreasing) permutation of agents. Moreover, in both cases an optimal network is complete because there is no trade-off between introducing new connections and diluting the influence of the existing friends.

that t_i is sufficient to induce i to act given that all agents preceding i in π act, no matter what the other agents do. Thus, an agent on the first place in the permutation, $\pi^{-1}(1)$, acts no matter what the others do. By induction suppose that for $k = 1, \dots, n-1$ agents $\pi^{-1}(1), \dots, \pi^{-1}(k)$ act. Then an agent $\pi^{-1}(k+1)$ also acts. Hence, each agent acts and t is INI.

Second, we show that increasing a transfer of any single agent creates an equilibrium where some of the agents do not act. For agent j , let

$$F_j = \{i | \exists j_1, \dots, j_m \text{ such that } g_{jj_1} = g_{j_1 j_2} = \dots = g_{j_m i} = 1 \\ \text{and } \pi(j) < \pi(j_1) < \dots < \pi(j_m) < \pi(i)\}.$$

If we increase the transfer of agent $\pi^{-1}(n)$, then she would strictly prefer not to act. Moreover, she strictly prefers not to act if any of her friends do not act. By induction suppose that for $k = 1, \dots, n-1$, each agent $\pi^{-1}(k+1), \dots, \pi^{-1}(n)$ strictly prefers not to act if all her friends following her in π and at least one of her friends preceding her in π does not act. If we increase the transfer of agent $\pi^{-1}(k)$, then there would exist an equilibrium where agent $\pi^{-1}(k)$ and all agents in $F_{\pi^{-1}(k)}$ do not act, while everyone else does. Therefore, t is tight.

Given a tight t , we construct a permutation π such that t is given by (2). First, there must exist an agent a_1 such that $t_{a_1} = 0$, otherwise there would exist an equilibrium where no one acts. Let $\pi(a_1) = 1$ and proceed to inductively define π . Suppose that for $k = 1, \dots, n-1$, each of the agents $a_1, \dots, a_k$ acts if agents with a lower index than theirs act, and correspondingly $\pi(a_i) = i$ and $t_{a_i} \leq f(d_{a_i}) \sum_{j=1}^{i-1} g_{a_i j}$ for $i = 1, \dots, k$. Then there exists an agent a_{k+1} who weakly prefers to act when $a_1, \dots, a_k$ act and others do not. Otherwise, there would exist an equilibrium where $a_1, \dots, a_k$ act and others do not, contradicting that t is INI. Hence, we must have

$$t_{a_{k+1}} \leq f(d_{a_{k+1}}) \sum_{i=1}^k g_{a_{k+1} a_i}. \quad (7)$$

Let $\pi(a_{k+1}) = k+1$. If at any step of the induction argument there are several such agents, then pick the one with the lowest index. Moreover, suppose that for $k = 1, \dots, n$ and some agent a_k we have that (7) holds with a strict inequality. But then t cannot be tight because by slightly increasing t_{a_k} , we can increase the revenue while keeping the mechanism INI. Thus, we have established a surjection from a set of permutations to a set of tight influence mechanisms. $\square$

PROOF OF PROPOSITION 2. First, we will need the following standard result (Abel's lemma). Let $a_1, \dots, a_n, b_1, \dots, b_n$ be real numbers. Set $A_k = \sum_{j=1}^k a_j$. Then

$$\sum_{k=1}^n a_k b_k = \sum_{k=1}^{n-1} A_k (b_k - b_{k+1}) + A_n b_n. \quad (8)$$

Also, note that we can rewrite (4) as

$$R(G) = \sum_k l_G(k) f(k). \quad (9)$$

Now suppose that $\sum_{j=1}^k l_G(j) \geq \sum_{j=1}^k l_{G'}(j)$ for each $k = 1, 2, \dots$. We shall show that G dominates G' , i.e., $R(G) \geq R(G')$ for each nonincreasing f . Using (9), this is equivalent to $\sum_k [l_G(k) - l_{G'}(k)] f(k) \geq 0$ for each nonincreasing f . By (8), we get

$$\sum_{k=1}^{\hat{d}} [l_G(k) - l_{G'}(k)] f(k) = \sum_{k=1}^{\hat{d}-1} A_k [f(k) - f(k+1)] + A_{\hat{d}} f(\hat{d}),$$

where $A_k = \sum_{i=1}^k [l_G(i) - l_{G'}(i+1)]$, and $\hat{d}$ is the highest degree among the nodes in G and G' . The above expression is nonnegative for each nonincreasing f , because by assumption, $A_k \geq 0$ for each $k = 1, 2, \dots$

Suppose that G dominates G' . For the sake of contradiction, suppose that there exists a positive integer $\bar{k}$ such that $\sum_{j=1}^{\bar{k}} l_G(j) < \sum_{j=1}^{\bar{k}} l_{G'}(j)$. Take f such that $f(k) = f(k+1)$ for each $k \neq \bar{k}$, and $f(\bar{k}) > f(\bar{k}+1)$. Then we have

$$\sum_{k=1}^{\hat{d}} [l_G(k) - l_{G'}(k)] f(k) = A_{\bar{k}} [f(\bar{k}) - f(\bar{k}+1)].$$

However, by assumption $A_{\bar{k}} < 0$, and we obtain the desired contradiction. $\square$

PROOF OF COROLLARY 1. Suppose G dominates G' . Fix $h = 1, 2, \dots$. Note that

$$\sum_{k=1}^h l_G(k) = E - H_h(G), \quad (10)$$

where E is the total number of links, which is the same in both networks. Thus, by dominance we have $E - H_h(G) \geq E - H_h(G')$, implying that $H_h(G) \leq H_h(G')$. Let D_h be the sum of degrees of agents with a degree weakly lower than h , which is also the same in both networks. Note that

$$\begin{aligned} D_h &= 2L_h(G) + E - L_h(G) - H_h(G) \\ &= 2L_h(G') + E - L_h(G') - H_h(G'). \end{aligned}$$

Combining it with the above, we get $L_h(G) \leq L_h(G')$. Finally, if G is more disassortative than G' , then from (10) it immediately follows that G dominates G' . $\square$

PROOF OF COROLLARY 2. For brevity, we consider only the case where there exists an integer s such that $E = \frac{1}{2}s(s-1) + s(n-s)$, i.e., the algorithm generates a complete core periphery with s stars. Note that a complete core periphery with s stars maximizes the number of links in a network given that there are $n-s$ agents with degree less than or

equal to s . Moreover, the maximal number of links is increasing in s . Now for the sake of contradiction suppose that a core periphery with s stars is dominated by G , and hence the number of links in G must be greater or equal to E . Then by Proposition 2 there exists $d < n - 1$ such that $L_d(G) + M_d(G) > s(n - s)$ and $L_s(G) + M_s(G) \geq s(n - s)$. Let $n_d(G)$ denote the number of nodes in G with degree higher than d . From the above, we get $n_d(G) < s$. Hence, the number of links in G must be less than E , a contradiction. $\square$

Proof of Proposition 3

We prove the result with help of four lemmas. Fix an optimal network G . Without loss of generality, assume that if $g_{ij} = 1$ and $d_i > d_j$, then $i < j$, and hence the identity permutation, id , is nonincreasing and induces an optimal influence mechanism. Let N_i denote a set of friends of agent i , i.e., $N_i = \{j | g_{ij} = 1\}$. For a permutation π and an agent i , let $N_i^{\pi,-} \subseteq N_i$ denote a subset of i 's friends who follow i in permutation π , i.e., $N_i^{\pi,-} = \{j | g_{ij} = 1 \text{ and } \pi(j) > \pi(i)\}$. We say that agents in $N_i^{\pi,-}$ are *influenced* by i . Let $d_i^{\pi,-} = |N_i^{\pi,-}|$ and $d_i^{\pi,+} = |N_i \setminus N_i^{\pi,-}|$. We call agent i a *sink* if $N_i^{\pi,-} = \emptyset$.

LEMMA 2. *Fix four different agents i, j, x , and y such that $\max\{d_x, d_y\} < \min\{d_i, d_j\}$. If $g_{ix} = g_{jy} = 1$, then either $g_{iy} = 1$, or $g_{jx} = 1$, or both.*

PROOF. For the sake of contradiction, suppose that $g_{iy} = g_{jx} = 0$. Consider replacing a link between j and y by a link between i and y . Since the benefits of the two links are the same, the corresponding change in the revenue is given by the change in the cost of a link, given by

$$d_i^{id,+} f(d_i) + d_j^{id,+} f(d_j) - d_i^{id,+} f(d_i + 1) - d_j^{id,+} f(d_j - 1).$$

Similarly, the change in the revenue from replacing a link between i and x by a link between j and x is

$$d_i^{id,+} f(d_i) + d_j^{id,+} f(d_j) - d_i^{id,+} f(d_i - 1) - d_j^{id,+} f(d_j + 1).$$

Because G is optimal, each replacement must weakly decrease the revenue:

$$d_i^{id,+} (f(d_i) - f(d_i + 1)) - d_j^{id,+} (f(d_j - 1) - f(d_j)) \leq 0,$$

$$d_j^{id,+} (f(d_j) - f(d_j + 1)) - d_i^{id,+} (f(d_i - 1) - f(d_i)) \leq 0.$$

Combining the inequalities, we get

$$\frac{f(d_j) - f(d_j + 1)}{f(d_i - 1) - f(d_i)} \leq \frac{d_i^{id,+}}{d_j^{id,+}} \leq \frac{f(d_j - 1) - f(d_j)}{f(d_i) - f(d_i + 1)}. \quad (11)$$

By convexity, we have

$$f(d_j) - f(d_j + 1) \leq f(d_j - 1) - f(d_j),$$

$$f(d_i) - f(d_i + 1) \leq f(d_i - 1) - f(d_i),$$

with equality only when the right-hand sides are zero. Clearly, if at least one right-hand side is not zero, then (11) is inconsistent. Hence, one of the two replacements must strictly decrease total revenue. Alternatively, if both RHSs are zero, then the cost of adding a link between i and y and a link between j and x on top of the existing links is zero, and hence it strictly increases the revenue. $\square$

LEMMA 3. *There exists a nonincreasing permutation π of agents such that*

$$N_{\pi^{-1}(1)}^{\pi, -} \supseteq N_{\pi^{-1}(2)}^{\pi, -} \supseteq \cdots \supseteq N_{\pi^{-1}(n)}^{\pi, -}. \quad (12)$$

PROOF. We shall construct a permutation π by inductively defining a permutation π_k for $1 \leq k \leq n$ and letting $\pi = \pi_n$. A permutation π_k will satisfy three properties: (i) π_k is nonincreasing, (ii) $N_{\pi_k^{-1}(1)}^{\pi_k, -} \supseteq N_{\pi_k^{-1}(2)}^{\pi_k, -} \supseteq \cdots \supseteq N_{\pi_k^{-1}(k)}^{\pi_k, -}$, and (iii) $\pi_k^{-1}(l) = l$ for $l > k$. Begin with an identity permutation id , letting $\pi_1 = \pi_2 = id$. Indeed, because agent 1 is the highest degree agent she has zero cost of a link, and thus must be connected to each node, implying that $N_1^{id, -} \supseteq N_2^{id, -}$. For the induction step, suppose that for $1 \leq k < n$ there is a permutation π_k satisfying the three properties above. We construct a permutation π_{k+1} . Suppose that $\pi_k^{-1}(k) = x$. First, we show that either $N_x^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$ or $N_x^{\pi_k, -} \subseteq N_{k+1}^{\pi_k, -}$. For the sake of contradiction, suppose there exist i and j such that $i \in N_x^{\pi_k, -}$, $i \notin N_{k+1}^{\pi_k, -}$ and $j \in N_{k+1}^{\pi_k, -}$, $j \notin N_x^{\pi_k, -}$. If $i \neq k+1$, then by Lemma 2 we have either $g_{xj} = 1$, or $g_{(k+1)i} = 1$, or both, a contradiction. So, suppose that $i = k+1$. By the induction assumption, $N_x^{\pi_k, -} \subseteq N_{\pi_k^{-1}(l)}^{\pi_k, -}$ for each $l < k$, and hence each node that follows and is connected to x is also connected to each node before x in π_k . Hence, $k+1$ must have strictly more friends preceding it in π_k than x , i.e., $d_{k+1}^{\pi_k, +} > d_x^{\pi_k, +}$. Moreover, it has a weakly lower degree than x because π_k is nonincreasing. Now consider replacing a link between $k+1$ and j by a link between x and j . It follows that the revenue must increase because the benefit accrued to j is the same but the cost of a link is lower for x than for $k+1$. Therefore, if $i = k+1$, then $N_x^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$, a contradiction. Thus, we have established that either $N_x^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$ or $N_x^{\pi_k, -} \subseteq N_{k+1}^{\pi_k, -}$. Now if $N_x^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$, then let $\pi_{k+1} = \pi_k$. On the other hand, if $N_x^{\pi_k, -} \subseteq N_{k+1}^{\pi_k, -}$, then define π_{k+1} in the following way. Move x one position further in the permutation, i.e., let $\pi_{k+1}^{-1}(k+1) = x$. Then, by the same argument as above either $N_{\pi_k^{-1}(k-1)}^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$ or $N_{\pi_k^{-1}(k-1)}^{\pi_k, -} \subseteq N_{k+1}^{\pi_k, -}$. If $N_{\pi_k^{-1}(k-1)}^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$, then let $\pi_{k+1}^{-1}(k) = k+1$, and $\pi_{k+1}^{-1}(l) = \pi_k^{-1}(l)$ for $l \neq k, k+1$. On the other hand, suppose that $N_{\pi_k^{-1}(k-1)}^{\pi_k, -} \subseteq N_{k+1}^{\pi_k, -}$ and $\pi_k^{-1}(k-1) = z$. Then, by the same argument as above, z is not connected to $k+1$. Now let $\pi_{k+1}^{-1}(k) = z$, and if $N_{\pi_k^{-1}(k-2)}^{\pi_k, -} \supseteq N_{k+1}^{\pi_k, -}$, then let $\pi_{k+1}^{-1}(k-1) = k+1$, and $\pi_{k+1}^{-1}(l) = \pi_k^{-1}(l)$ for $l \neq k-1, k, k+1$. Continue moving $k+1$ to the top of the permutation in this way until a set of agents influenced by it is nested in the set of agents influenced by an agent preceding it in π_k . Each step of the above procedure is well-defined, and in the end, it produces a permutation π_{k+1} satisfying the properties. Iterating the procedure yields π_n , and finally letting $\pi = \pi_n$ we obtain the required permutation. $\square$

LEMMA 4. If $k \leq m$, then $(k+1)f(m+1) - kf(m)$ is:

- (i) nonincreasing in k , given m , and
- (ii) nondecreasing in m , given $k \geq n/2$.

PROOF. Part (i) follows from f being a nonincreasing function. To prove (ii) note that by strong convexity, for $n/2 \leq k \leq m$ we have

$$f(m) - f(m+1) \geq [f(m+1) - f(m+2)] \left(1 + \frac{1}{k}\right).$$

Rewriting, we get

$$\begin{aligned} (f(m+1) - f(m+2))(k+1) &\leq kf(m) - kf(m+1), \\ (k+1)f(m+1) - kf(m) &\leq (k+1)f(m+2) - kf(m+1). \end{aligned} \quad \square$$

LEMMA 5. Fix a nonincreasing π and agents i and j such that $d_i = d_j = m$, $d_i^{\pi,+} = d_j^{\pi,+} = k$. If $2k < m$, then $g_{ij} = 1$.

PROOF. For the sake of contradiction, suppose that $g_{ij} = 0$. Consider the change in the revenue due to adding a link between i and j :

$$\underbrace{kf(m+1) + (k+1)f(m+1)}_{\text{After adding a link}} - \underbrace{2kf(m)}_{\text{Before adding a link}}.$$

Rewriting, we find that a new link increases the revenue if

$$f(m+1) - 2k(f(m) - f(m+1)) > 0.$$

By Assumption 1, we have

$$f(m+1) - 2k(f(m) - f(m+1)) \geq f(m+1) - \frac{2k}{m}f(m+1).$$

Hence, the revenue increases when $2k < m$, contradicting the optimality of G . $\square$

Now we are ready to prove Proposition 3.

PROOF OF PROPOSITION 3. Take a nonincreasing permutation π satisfying (12), and let $\pi^{-1}(k) = s_k$ for $k = 1, \dots, n$. We show that each nonsink is connected to each sink. First, we show that each sink is connected to the same set of nonsinks. For the sake of contradiction, suppose that x and y are two sinks and nonsink s_j connects to x but not to y . Then by (12) each s_k such that $k \leq j$ connects to x and each s_l such that $l \geq j$ does not connect to y . Hence, there are strictly fewer agents connected to y than to x , i.e., $d_x > d_y$. Moreover, by (12) it follows that each agent connected to y also connects to x . Then take all agents connected to x and not to y . Removing the links from these agents to x must decrease the revenue. But then adding the links from these agents to y creates

the same benefit as adding them to x , but has a lower cost by convexity. Hence, each agent connected to x must also connect to y , a contradiction. Second, notice that there cannot exist a nonempty set of nonsinks not connected to sinks because the last such agent in π must be herself a sink.

It remains to show that all nonsinks are connected. Let s_k be the last nonsink in sequence $(s_1, s_2, \dots, s_n)$. First, we show that s_{k-1} connects to s_k . For the sake of contradiction, suppose not. Then $N_{s_{k-1}}^{\pi,-} = N_{s_k}^{\pi,-}$ and $x \in N_{s_{k-1}}^{\pi,-}$ if and only if x is a sink. Suppose, first, that $d_{s_{k-1}} < d_{s_k}$. Then by an argument similar to the above, there exists s_j , $j < k-1$, such that $s_{k-1} \notin N_{s_j}^{\pi,-}$ and $s_k \in N_{s_j}^{\pi,-}$. Take all such agents. By symmetry, adding links between these agents and s_{k-1} reduces the total revenue, because the costs are lower and the benefit is the same as when adding links between these nodes and s_k . It follows that s_{k-1} and s_k are symmetric. Now by Lemma 5, s_{k-1} must be connected to s_k if $d_{s_k}^{+,+} < d_{s_k}^{\pi,-}$, where $d_{s_k}^{\pi,-}$ is also the number of sinks. Instead suppose that $d_{s_{k-1}}^{+,+} = d_{s_k}^{+,+} \geq d_{s_{k-1}}^{\pi,-} = d_{s_k}^{\pi,-}$. Then

$$\begin{aligned} d_{s_k}^{\pi,+} + d_{s_k}^{\pi,-} &\leq n-2, \\ 2d_{s_k}^{\pi,-} &\leq n-2, \\ d_{s_k}^{\pi,-} &\leq n/2-1, \end{aligned}$$

where the second inequality follows from $d_{s_k}^{\pi,+} \geq d_{s_k}^{\pi,-}$. Hence, there are weakly fewer sinks than $n/2-1$ and, therefore, the in-degree of each sink must be strictly greater than $n/2$ because it is connected to each nonsink. Take any sink x , and consider the benefit created by a link between s_{k-1} and x . It is given by $d_x^+ f(d_x) - (d_x^+ - 1)f(d_x - 1)$. We compare this benefit to the one created by instead connecting s_{k-1} to s_k , given by $(d_{s_k}^+ + 1)f(d_{s_k} + 1) - d_{s_k}^+ f(d_{s_k})$. We have

$$\begin{aligned} (d_{s_k}^+ + 1)f(d_{s_k} + 1) - d_{s_k}^+ f(d_{s_k}) &> (d_x^+ + 1)f(d_{s_k} + 1) - d_x^+ f(d_{s_k}), \\ &> (d_x^+ + 1)f(d_x + 1) - d_x^+ f(d_x), \end{aligned}$$

where the first inequality follows because x connects to each nonsink, and s_k is at least not connected with s_{k-1} , and so we have $d_{s_k}^+ < d_x^+$, and the second inequality follows from Lemma 4 because $d_x^+ > n/2$ and $d_{s_k} \geq d_x$. Therefore, it is profitable to add a link between s_{k-1} and s_k instead of a link between s_{k-1} and x , and thus s_{k-1} and s_k must be connected. Finally, suppose that nonsink s_j and s_{j+1} are not connected, and all nonsinks after $j+1$ connect to the following nonsinks. The argument above applies, and hence the two nonsinks must be connected. $\square$

Remaining proofs

PROOF OF COROLLARY 3. By Proposition 3, an optimal network is a complete core periphery. The revenue in such a network with s stars is

$$\frac{s(s-1)}{2} f(n-1) + (n-s)sf(s). \quad (13)$$

Substituting the expression for f , we get a quadratic $As^2 + Bs + n$, where $A = \frac{1}{2(n-1)} - \frac{\alpha}{2}$ and $B = \alpha(n - \frac{1}{2}) - \frac{2n-1}{2(n-1)}$. If $\alpha < \frac{1}{n-1}$, then the function is convex and the maximum is achieved either when $s = 1$ or $s = n - 1$. Substituting the values, we find that $s = 1$, in other words a star, is optimal. If $\alpha > \frac{1}{n-1}$, then the function is concave and is maximized at $s^* = -\frac{B}{2A}$. Some algebra reveals that s^* is constant in α and is equal to $n - \frac{1}{2}$. Thus, the maximum in integer values is achieved when $s = n$ or $s = n - 1$, both cases corresponding to a complete network. Finally, when $\alpha = \frac{1}{n-1}$, then the revenue is independent of s . $\square$

PROOF OF PROPOSITION 4. Let f be flatter than h . For $s = 2, 3, \dots, n - 1$, we have

$$f(s) = 1 - \sum_{i=1}^{s-1} [f(i) - f(i+1)].$$

Let $\delta_s = f(s) - h(s)$ for $s = 1, 2, \dots, n - 1$. From the above, we get

$$\delta_s = \sum_{i=1}^{s-1} ([h(i) - h(i+1)] - [f(i) - f(i+1)]),$$

where each term is positive by assumption. Thus, $\delta_s \geq 0$ and is nondecreasing.

For f , let G_f^s denote the revenue in a complete core periphery with s stars given by (13). We show that $G_f^s - G_h^s$ is nondecreasing in s . We have

$$\begin{aligned} G_f^s - G_h^s - (G_f^{s+1} - G_h^{s+1}) &= \frac{s(s-1)}{2} \delta_{n-1} + (n-s)s\delta_s \\ &\quad - \dots - \frac{(s+1)s}{2} \delta_{n-1} - (n-s-1)(s+1)\delta_{s+1} \\ &= -s\delta_{n-1} + (n-s)s\delta_s - (n-s-1)(s+1)\delta_{s+1} \\ &\leq -s\delta_{n-1} + (n-s)s\delta_s - (n-s-1)(s+1)\delta_s \\ &= -s\delta_{n-1} + (2s-n+1)\delta_s \\ &\leq -s\delta_s + (2s-n+1)\delta_s \\ &= (s-n+1)\delta_s \\ &\leq 0. \end{aligned}$$

The first two inequalities follow from δ_s being nondecreasing. Finally, fix $s^*(h)$ and consider moving from h to f . From the above, it follows that the revenue in a complete core periphery with $s^*(h)$ stars increases weakly more than in any complete core periphery with fewer stars. Hence, a complete core periphery with fewer stars cannot be optimal under f . $\square$

PROOF OF PROPOSITION 5. Suppose that G is a resilient network with E links. Let i be an agent with the smallest positive degree and C be a subset of agents with positive degrees, excluding i . It suffices to show that $\binom{|C|}{2} < E$. For the sake of contradiction, suppose that $E \leq \binom{|C|}{2}$, and so it is possible to isolate i by sequentially deleting an existing

link between each pair i and j , and replacing it by a link between some pair in C . Consider a link reallocation procedure. At step t , let G^t be the current network and $(d_i^t)_{i=1}^n$ be the degrees of agents in G^t . At each step $t \geq 1$, we conduct one of the two types of link reallocation:

1. If there exists $j, k \in C$ such that $g_{ij}^{t-1} = 1$ and $g_{jk}^{t-1} = 0$, then let $g_{ij}^t = 0$, $g_{jk}^t = 1$, and $g_{xy}^t = g_{xy}^{t-1}$ for each pair x, y different from i, j or j, k .
2. If there exists $j \in C$ such that $g_{ij}^{t-1} = 1$, but does not exist $k \in C$ such that $g_{jk}^{t-1} = 0$, then take any pair $m, l \in C$ such that $g_{ml}^{t-1} = 0$ and let $g_{ij}^t = 0$, $g_{ml}^t = 1$, and $g_{xy}^t = g_{xy}^{t-1}$ for each pair x, y different from i, j or m, l .

The procedure first exhausts all reallocations of type (1) and then of type (2). Initialize at $G = G^0$ and terminate if $d_i^{t-1} = 0$.

Consider type (1) reallocation at step t . A value of a link between j and k in G^t cannot exceed the value of a link between j and i in G , because type (1) reallocations do not decrease the degrees of agents in C , and thus $\min\{d_j^{t-1}, d_k^{t-1}\} \geq d_i$. Moreover, the revenue generated by each other link in G^{t-1} does not increase in G^t because we have increased k 's degree. Consider type (2) reallocation at step t . A value of a link between m and l in G^t cannot exceed the value of a link between j and i in G . Indeed, as mentioned before, type (1) reallocations do not decrease the degrees of agents in C , and type (2) reallocations can decrease a degree only of agents adjacent to i . However, type (2) reallocation decreases the degree of j and increases the degrees of m and l . First, we evaluate the increase in the revenue due to the decrease in j 's degree. Note that agents connected to i in G^{t-1} have the highest degree in G^{t-1} equal to $|C|$. Let S^{t-1} be a set of such agents. Because the degree of j decreases by one, the revenue generated by each of $|S^{t-1}| - 1$ links between agents in $S^{t-1} \setminus j$ and j increases by $f(|C| - 1) - f(|C|)$, while the revenue from each other link adjacent to j does not change because j remains the highest degree agent in G^t . Hence, we have a total increase in the revenue of $(|S| - 1)[f(|C| - 1) - f(|C|)]$. Second, we evaluate the decrease in the revenue due to the increases in the degrees of m and l . Consider the revenue generated by each of $2|S^{t-1}|$ links between m and l and each $h \in S^{t-1}$. It must decrease by at least $f(|C| - 2) - f(|C| - 1)$ by convexity, because the highest degree of m and l in G^{t-1} is $|C| - 2$. Hence, the revenue decreases by at least $2|S|[f(|C| - 2) - f(|C| - 1)]$. Thus, the net change in the revenue after reallocation (2) at step t is negative because

$$2|S|[f(|C| - 2) - f(|C| - 1)] > (|S| - 1)[f(|C| - 1) - f(|C|)],$$

by convexity. Therefore, it follows that after each reallocation the total value that links between agents in C contribute to the revenue decreases. When the procedure terminates at step T , the value of each reallocated link is not higher in G^T than in G . Therefore, the revenue in G^T is not higher than in G . Thus, $\binom{|C|}{2} < E$. $\square$

REFERENCES

Arthur, David, Rajeev Motwani, Aneesh Sharma, and Ying Xu (2009), "Pricing strategies for viral marketing on social networks. ArXiv e-prints." [4]

Babaioff, Moshe, Michal Feldman, Noam Nisan, and Eyal Winter (2012) *Combinatorial agency. Journal of Economic Theory*, 147, 999–1034. [4]

Bala, Venkatesh and Sanjeev Goyal (2000), “A noncooperative model of network formation.” *Econometrica*, 68, 1181–1229. [4]

Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou (2006), “Who’s who in networks. Wanted: The key player.” *Econometrica*, 74, 1403–1417. [3]

Belhaj, Mohamed, Sebastian Bervoets, and Frédéric Deroïan (2016), “Efficient networks in games with local complementarities.” *Theoretical Economics*, 11, 357–380. [4]

Belhaj, Mohamed and Frédéric Deroïan (2019), “Group targeting under networked synergies.” *Games and Economic Behavior*, 118, 29–46. [3]

Bernstein, Shai and Eyal Winter (2012), “Contracting with heterogeneous externalities.” *American Economic Journal: Microeconomics*, 4, 50–76. [4]

Bimpikis, Kostas, Asuman Ozdaglar, and Ercan Yildiz (2016), “Competitive targeted advertising over networks.” *Operations Research*, 64, 705–720. [3]

Bloch, Francis (2016), “Targeting and pricing in social networks.” In *The Oxford Handbook of the Economics of Networks*. [3]

Bloch, Francis and Nicolas Querou (2013), “Pricing in social networks.” *Games and Economic Behavior*, 80, 243–261. [3]

Candogan, Ozan, Kostas Bimpikis, and Asuman Ozdaglar (2012), “Optimal pricing in networks with externalities.” *Operations Research*, 60, 883–905. [3]

Contagion, Stephen Morris (2000) *The Review of Economic Studies*, 67, 57–78. [3]

Devetag, Giovanna and Andreas Ortmann (2007), “When and why? A critical survey on coordination failure in the laboratory.” *Experimental Economics*, 10, 331–344. [2]

Fainmesser, Itay P. and Andrea Galeotti (2016), “Pricing network effects.” *The Review of Economic Studies*, 83, 165–198. [3]

Galeotti, Andrea, Benjamin Golub, and Sanjeev Goyal (2020), “Targeting interventions in networks.” *Econometrica*, 88, 2445–2471. [3]

Galeotti, Andrea and Sanjeev Goyal (2010), “The law of the few.” *American Economic Review*, 100, 1468–1492. [4]

Ghiglino, Christian and Sanjeev Goyal (2010), “Keeping up with the neighbors: Social interaction in a market economy.” *Journal of the European Economic Association*, 8, 90–119. [5]

Goyal, Sanjeev and Sumit Joshi (2003), “Networks of collaboration in oligopoly.” *Games and Economic behavior*, 43, 57–85. [4, 15]

Goyal, Sanjeev, Marco J. Van Der Leij, and José Luis Moraga-González (2006), “Economics: An emerging small world.” *Journal of Political Economy*, 114, 403–412. [4]

Halac, Marina, Ilan Kremer, and Eyal Winter (2020), “Raising capital from heterogeneous investors.” *American Economic Review*, 110, 889–921. [4]

Halac, Marina, Elliot Lipnowski, and Daniel Rappoport (2021), “Rank uncertainty in organizations.” *American Economic Review*, 111, 757–786. [4]

Hartline, Jason, Vahab Mirrokni, and Mukund Sundararajan (2008), “Optimal marketing strategies over social networks.” In *Proceedings of the 17th International Conference on World Wide Web, WWW’08*, 189–198. [3]

Herskovic, Bernard and Joao Ramos (2020), “Acquiring information through peers.” *American Economic Review*, 110, 2128–2152. [4]

Hiller, Timo (2017), “Peer effects in endogenous networks.” *Games and Economic Behavior*, 105, 349–367. ISSN 0899-8256. [4]

Hojman, Daniel A. and Adam Szeidl (2008), “Core and periphery in networks.” *Journal of Economic Theory*, 139, 295–309. [4]

Jackson, Matthew O. and Leeat Yariv (2007), “Diffusion of behavior and equilibrium properties in network games.” *American Economic Review*, 97, 92–98. [3]

Kempe, David, Jon Kleinberg, and Éva Tardos (2003), “Maximizing the spread of influence through a social network.” In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146, ACM. [4]

König, Michael D., Claudio J. Tessone, and Yves Zenou (2014), “Nestedness in networks: A theoretical model and some applications.” *Theoretical Economics*, 9, 695–752. [4]

Liu, Xiaodong, Eleonora Patacchini, and Yves Zenou (2014), “Endogenous peer effects: Local aggregate or local average?” *Journal of Economic Behavior and Organization*, 103, 39–59. [5]

Masaki, Aoyagi (2013), “Coordinating adoption decisions under externalities and incomplete information.” *Games and Economic Behavior*, 77, 77–89. [4]

Sadler, Evan (2020a), “Diffusion games.” *American Economic Review*, 110, 225–270. [3]

Sadler, Evan (2020b), “Influence campaigns.” Available at SSRN 3371835. [3]

Segal, Ilya (1999), “Contracting with externalities.” *The Quarterly Journal of Economics*, 114(2), 337–388. [4]

Segal, Ilya (2003), “Coordination and discrimination in contracting with externalities: Divide and conquer?” *Journal of Economic Theory*, 113, 147–181. [4]

Talamàs, Eduard and Omer Tamuz (2017), “Network cycles and welfare.” Available at SSRN 3059604. [3]

Vohra, Akhil (2020), “Strategic influencers and the shaping of beliefs.” [3]

Weyl, E. Glen (2010), “A price theory of multi-sided platforms.” *American Economic Review*, 100, 1642–1672. [4]

Winter, Eyal (2004), “Incentives and discrimination.” *The American Economic Review*, 94, 764–773. [4]

Co-editor Simon Board handled this manuscript.

Manuscript received 13 October, 2021; final version accepted 3 March, 2023; available online 13 March, 2023.