On Stars and Galaxies: Exploiting Social Influence in Networks

Vladyslav Nora* and Eyal Winter†

October 2020

Abstract

A principal wishes to induce an action from agents that belong to a social network. Agents' social benefit from taking the action increases with any additional friend who acts. On top of the social benefits, the principal offers external rewards in order to sustain a unique Nash equilibrium where everyone acts. We first show that in the influence mechanism that minimizes the principal's expenses, popular agents receive a preferential treatment from the principal. Using this observation, we identify networks that are most favorable for the principal to induce action. Such networks, "galaxies", partition nodes into stars and periphery, with every star being linked to all nodes, and every periphery node being linked only to stars. We discuss the relevance of this finding to social media platforms (such as Facebook and Twitter) in terms of manipulating the network, as well as to regulators who would attempt to prevent such manipulations.

Keywords: social networks, mechanism design, unique implementation, strategic complementarities, split graphs.

JEL Classification Numbers: C72, D82.

*Nazarbayev University. E-mail: vladyslav.nora@nu.edu.kz

†Lancaster University and The Hebrew University. Email: mseyal@mscc.huji.ac.il

1 Introduction

The economics literature on social influence in networks has thrived during the past few decades with both theoretical and empirical works. To an extent, the growing interest in these topics has been triggered by the proliferation of online social networks. Theory papers have provided important models for the process of network formation and the spread of influence within networks (Jackson and Wolinsky, 1996; Kempe et al., 2003; Ballester et al., 2006; Jackson, 2011). Empirical papers used existing data as well as field experiments to study how behavior diffuses within network (Ichino and Maggi, 2000; Sacerdote, 2001; Mas and Moretti, 2009; Banerjee et al., 2013).

However, less has been written on the mechanisms that enhance social influence and on how external forces to the network can utilize such mechanism to affect individuals' behavior. We also note that the literature is thin in addressing the question of which network structures are more susceptive to developing social influence – or put differently, which network structures are more suitable for external forces to utilize social influence in trying to affect individuals' behavior. In this paper we attempt to precisely answer these questions.

The role of external forces to the network in utilizing social influence has become salient following the Cambridge Analytica scandal that exposed ways in which Trump's 2016 election campaigners engineered their activities on Facebook to manipulate social influence using big data. It is also reasonable to assume that companies such as Facebook and LinkedIn gain by having more control over the flow of social influence through their platforms, as social influence can be utilized to make commercial ads more effective, and hence increase their market price and consequently also the revenues of the social network company. To the extent that the public and its regulators should be concerned about companies' excessive control over social influence, it is important to understand how such control can be optimally established and what sort of network structures are more vulnerable to such manipulation. Finally, our results are also relevant to benevolent authorities that seek to use social influence to increase social welfare without implementing legal measures. Social distancing and self-quarantine would be the relevant actions in the context of COVID-19 pandemic. Online public campaigns against smoking or alcohol consumption to improve public health also come to mind.

We use a simple model of social influence based on graphs, where nodes represent individuals/agents, and edges represent links/friendships. An external force or a principal desires to induce as many people as possible to take a certain action (i.e., vote for a certain candidate, buy a specific product, commit to a certain behavior that protects the environment, or self-isolate during an epidemic, etc.). Our main behavioral assumption is that by taking the action, each agent induces positive externality on each of his/her friends who took the action as well. We interpret this positive externality to be driven by psychological or social factors such as conformism but it can also take the form of consumption benefits when the underlying action involves purchasing a certain network product. In addition to this social payoff and the cost of taking the action, the principal contracts with each agent to promise a certain reward if the agent takes the action. These transfers should primarily be interpreted as resources invested by the principal to convince the agent to take the action or reduce the (mental) cost of taking the action. Each set of transfers induces a game between all members of the network with each player having two strategies: either taking the action or refraining from doing it. We will be interested in the scheme that minimizes the principal's total transfer to the agents subject to inducing all agents to take the action in a unique Nash equilibrium. The requirement of uniqueness is driven by the assumption that the principal cannot coordinate his/her agents to play a specific equilibrium and we shall discuss this assumption later.

Our first result characterizes the optimal mechanisms for utilizing social influence. In all mechanisms the equilibrium is obtained by iterative elimination of dominated strategies. But more importantly our result highlights the importance of degree centrality in networks, as it shows that those who enjoy higher centrality receive preferential treatment from the principal. Players of high-degree centrality receive higher-powered incentives than the rest and take the role of network leaders, allowing the principal to utilize his/her social influence on others. As periphery players expect leaders to take the action, a lower-powered incentive would suffice to induce them to act as well. Put differently, players with low-degree centrality face smaller (strategic) risk when taking the actions and hence receive lower premiums to compensate for this risk than high-centrality players.

Perhaps the most novel aspect of our paper is the objective of characterizing the network structures that are most vulnerable to the spread of social influence. Put differently, we are interested in those network architectures that allow the principal to induce all agents to take the action at the lowest possible overall cost. Our results here are informative not only for bodies that seek to utilize social influence online, but also those that seek to diminish it (such as internet regulators). We believe these results to be highly important because major internet players (such as Facebook, Google, and other social media platforms) have enormous power not only over the design of incentives to undertake specific actions for a given network, but also over the design of the network's structure itself. Facebook's algorithms that determine friend suggestions and whose post appears on whose feed are constantly restructuring the network. If these algorithms can be alerted when the network architecture approaches a certain critical level of vulnerability to the rapid spread of social influence, these companies may choose to take preventive measures or, if not, be forced to do so by regulators.

Our finding regarding the principal's optimal network structure is that under mild and intuitive conditions networks possessing the hub-and-spoke type of architecture allow for the most effective exploitation of social influence. These graphs, that we call "galaxies", have the following simple structure: the set of nodes (agents) is partitioned into two subsets, "stars" and "periphery". Every star is connected to all other nodes, while periphery nodes are connected only to stars.[1] A novel feature of our model that allows us to study optimal networks in a greater generality is the degree-dependent network effects. Specifically, we assume that while the influence of friends is linear in the number of friends taking the action, the degree of influence of a specific friend is declining with one's number of friends.

[1] The hub-and-spoke network structures related to our galaxies arise in several other contexts. König et al. (2014) derive nested split graphs as limits of a certain stochastic network formation process, whereas Goyal and Joshi (2003) and Goyal et al. (2006) uncover the role of interlinked stars in, respectively, oligopoly collaboration networks and co-authorship network among economists. While galaxy satisfies most of the properties of the structures described in this literature, it remains a special case of hub-and-spoke network.

To the best of our knowledge it is the first paper to explore this idea.

Finally, we consider two modifications of the benchmark model. First, we note that social influence is not always symmetric. This asymmetry is exemplified by online social networks such as Twitter and Instagram, where one can follow and be influenced by other users without necessarily exerting influence on them. We present a simple extension of our model, based on directed graphs, that accounts for asymmetric social influence. It turns out that the principal's optimal directed networks are acyclic tournaments: agents are arranged in an arbitrary sequence, and are influenced by all the agents in front of them. Second, at the end of the paper we investigate a model where agents who take the action actively influence their friends to follow suit. We show that many insights from the original model continue to hold, and, in particular, galaxy networks are still optimal.

1.1 Literature review

This paper contributes to a vast literature on peer effects and social networks. The literature studies how social interactions affect the adoption and spread of behaviors, technologies, and ideas. Methodologically, the existing papers follow one of the two avenues. The first strand of theoretical literature examines how locally interacting individuals coordinate their actions (Ellison, 1993; Morris, 2000; Jackson and Yariv, 2007; Jackson and Rogers, 2007). These papers consider games where players face a simple choice whether to adopt some behavior or not. Researchers study how adoption levels and dynamics of diffusion relate to characteristics of social interaction networks. Often in these papers a network is not perfectly known and is described solely by its degree distribution (Jackson and Rogers, 2007; Jackson and Yariv, 2007), whereas players repeatedly revise their actions according to a best response dynamics. The other strand of theoretical literature studies how the intensity of agents' actions, such as the effort they invest in production of a local public good, depends on their position in a network. This literature considers continuous-action models with linear best replies where equilibrium actions depend on a player's Bonacich centrality (Ballester et al., 2006; Bramoullé and Kranton, 2007; Ghiglino and Goyal, 2010; Bramoullé et al., 2014). We adopt the binary action framework of the former literature to raise and answer two novel questions: what are the optimal mechanisms for exploiting social influence and which networks are most susceptive to the external influence?

Our work is also related to a literature on pricing and targeting in social networks.[2] Like our paper, this literature considers an outside player who takes action to affect agents' behavior in a network. Typically these papers study a problem of price-discriminating firms in the presence of consumption externalities. The examples include Candogan et al. (2012), Bloch and Querou (2013), and Fainmesser and Galeotti (2016), who explore, albeit in different theoretical frameworks, how optimal prices and welfare depend on network characteristics. Two related problems are the analysis of "key" players whose removal induces the greatest change in equilibrium aggregate action (Ballester et al., 2006), and the optimal targeting of interventions in networks by a planner who seeks to maximize the welfare (Galeotti et al., 2017; Talamàs and Tamuz, 2017). On the other hand, the literature on the intersection of economics and computer science investigates the algorithmic aspects of optimal pricing strategies (Hartline et al., 2008; Arthur et al., 2009). It is motivated by so called "viral marketing", where a firm gives discounts or free samples to influential individuals, hoping to eventually reach more customers. Similarly, Domingos and Richardson (2001) and Kempe et al. (2003) consider a related algorithmic problem of maximizing a spread of influence over a social network.[3] In fact, we adopt a variant of a linear threshold model of social influence, a benchmark in many computer science papers, introduced in Granovetter (1978). In this model, each agent acts if and only if some linear function of the sum of actions of her friends exceeds this agent's individual threshold. Although our model can be interpreted in the context of price-discriminating monopolist selling an indivisible good, our approach is substantially different. First, our analysis of op-

[2]Bloch (2016) provides a comprehensive survey of this literature.

[3]However, the typical problem considered by this literature is different. In particular, for a parameter k one is asked to find k nodes, such that if these nodes act, eventually the maximal number of other nodes also choose to act. By contrast, we look for a profile of thresholds such that everyone acts in the unique equilibrium of the induced game on a network, and the sum of the thresholds is maximized. Whereas the former problem is NP hard, we provide a simple solution to the latter one.

timal influence mechanisms uses the unique implementation framework currently unexplored in the literature.[4] Second, we study a new intriguing question: which networks are more susceptive to the manipulation by the external forces?

Our unique implementation approach is motivated by the fact that the principal cannot coordinate a large group of agents into his most preferred equilibrium. Several experimental papers show that in the absence of such a coordination device, the principal's desired equilibrium tends to unravel (Devetag and Ortmann, 2007). The unique implementation approach has been used in many contracting environments in the literature. The study of unique implementation in contracting under multilateral externalities was initiated by Segal (1999) and Segal (2003), who develop a general contracting model, and Winter (2004), who explores incentives provision in organizations. Babaioff et al. (2012), Bernstein and Winter (2012), Halac et al. (2020a), and Halac et al. (2020b) are prominent papers that explore the implications of unique implementation in different environments.[5] Here we extend their analysis to include local externalities captured by a social network.

The paper is organized as follows. We introduce the model and discuss examples in Section 2. In Section 3 we characterize optimal influence mechanisms, while in Section 4 we present the result on networks most susceptive to the external influence. In Section 5 we extend the model to consider asymmetric social influence, and in Section 6 we consider the active influence model. We present the formal proofs in the Appendix.

[4]The existing models of pricing in networks typically assume unique equilibrium by restricting the value of network effect parameter to be sufficiently small.

[5]Park (2004), Weyl (2010), and Aoyagi (2013) also address the coordination problem among consumers in models of network goods. Yet, their solutions are based on adoption-contingent prices that depend on the total number of adopters. By contrast, we study the bilateral contracting where transfers of agents are not contingent on other parties' actions.

2 Model

2.1 Setup

A society consists of n individuals (agents) indexed by $i = 1, 2, \ldots, n$. 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 individual 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. The payoff from taking the action depends on the decisions of friends in the following way. Given network G and action profile $x = (x_1, \ldots, x_n)$, the payoff of individual i from taking the action is

$$U_i(x, G) = f(d_i) \sum_j g_{ij} x_j + t_i - c_i. \tag{1}$$

Here the term $\sum_j g_{ij} x_j$ is the number of friends of i who choose to act, the term $f(d_i)$ captures the social influence exerted on i by each such friend, t_i is an external reward for the action, and $c_i > 0$ is the individual cost of the action. Hence, a payoff from acting is comprised of a social benefit an individual derives from interacting with active friends and an external incentive. Social influence is always positive, $f(d_i) > 0$, meaning that the payoff from acting is linearly increasing in the number of active friends. Our central and novel assumption is that it is also a weakly decreasing function of the number of friends of individual i, i.e., $f(m) \geq f(m+1)$ for $m = 1, 2, \ldots, n$. The assumption reflects the idea that someone with more friends is swayed less by each one of them. We normalize the payoff from abstaining, $x_i = 0$, to zero. Hence, (1) can be viewed as the individual incentives to act given a social network and the behaviors of others.

Consider the resulting simultaneous move game with complete information. Because $f(d_i) > 0$, the game involves strategic complements. Therefore without the appropriate external rewards, typically, there are multiple equilibria, in some of which agents may fail to coordinate on acting.

2.2 The principal's problem

The principal wishes to induce all individuals to act by offering the rewards. An influence mechanism is a profile of rewards paid to agents if they choose to act, i.e., a vector $t = (t_1, \ldots, t_n)$. Because the game between the agents might have multiple equilibrium outcomes, we require that the rewards offered by the principal induce a unique equilibrium where all agents act. Formally, an influence mechanism t is incentive-inducing (INI) if $x = (1, \ldots, 1)$ 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 strictly dominant strategy for each agent i by offering $t_i > c_i$. Moreover, if t is INI, then so is each $t' > t$. However, the principal wants to induce action at the minimal total reward. Influence mechanism t is *optimal* if it has the lowest total reward among all INI mechanisms, i.e., $\sum t_i \leq \sum t'_i$ for each INI mechanism[6] t'.

The total reward required to induce all agents to act depends on the existing social network. Some networks require a lower total reward than others, and hence are more susceptible to external influence. A network is *optimal* if its corresponding optimal influence mechanism has the lowest total reward across all networks. In the remainder of this section we present two examples based on specific social influence functions f naturally arising in certain settings. We use the examples to illustrate the construction of optimal influence mechanisms and compare different networks in terms of their susceptibility to external influence.

Example 1. Consider the case where individuals directly care about a proportion and/or an absolute number of friends who take the action.[7] Then the social

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

[7] The examples of such situations studied in the literature include models of social comparison (Ghiglino and Goyal, 2010), conformity (Patacchini and Zenou, 2012; Liu et al., 2014), and general coordination games on network (Jackson and Zenou, 2015). For example, Ghiglino and

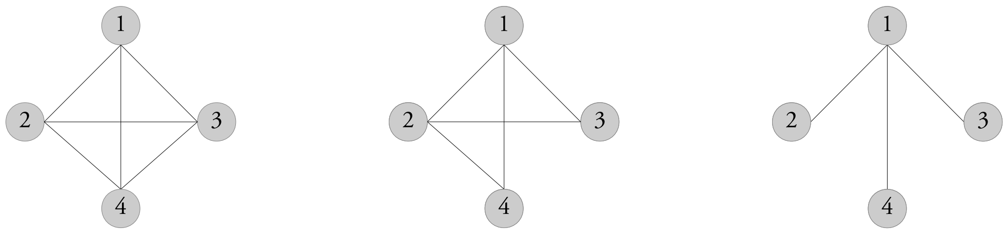

Figure 1: Comparing total rewards in social networks with four agents.

influence of each such friend on agent i is given by

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

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. Consider the networks in Figure 1 with four agents who have the same cost c. We begin by illustrating the construction of the optimal influence mechanisms for each network, and then compare the corresponding total rewards to see which network is easier for the principal to manipulate.

First, note that, in any network, the principal must induce at least one of the agent to act even when no one else does (otherwise there will be an equilibrium in which no one acts). Hence one of the agents must be paid at least c. In a complete network on the left, all agents are symmetric in the network and therefore we can let agent 1 receive $t_1 = c$. Second, in a complete network one of the remaining agents must be paid at least $c - \alpha - 1/3$; otherwise there is an equilibrium where only agent 1 acts. Again, by symmetry we can let $t_2 = c - \alpha - 1/3$. Similarly, to induce one of the two remaining agents to act when both, 1 and 2 act, the principal must offer a reward of at least $c - 2\alpha - 2/3$. Let $t_3 = c - 2\alpha - 2/3$. Finally, agent 4 must be paid at least $t_4 = c - 3\alpha - 1$. In fact, this is an optimal influence mechanism for a complete network with the total reward of $4c - 6\alpha - 2$. Now

Goyal (2010) discuss two models, when local aggregate consumption of friends matters and when local average matters, and provide justification for each case. Both cases are subsumed by our model.

consider a network in the middle panel of Figure 1 with a severed link between 3 and 4. Here we can keep the rewards of 1 and 2 unchanged. Moreover, now to induce agent 3 to act when only 1 and 2 do, the principal can pay her less than before. Specifically, paying $t_3 = c - 2\alpha - 1$ is sufficient because all her friends are now active. However, in order for 4 to act when everyone else does, we must increase her reward by α because now she does not experience social influence from 3. Therefore the optimal reward mechanism in the middle network has the total reward of $4c - 5\alpha - 9/4$. Finally, we turn to the star network in the right panel of Figure 1. In the optimal reward mechanism 1 is paid $t_1 = c$ to act when no one else does. Furthermore, to induce every other agent to act we must pay at least $c - \alpha - 1$. The total reward is $4c - 3\alpha - 3$. We find that incentivizing agents in the middle network is always more expensive, while it is cheaper in the complete network if $\alpha > 1/3$, and cheaper in the star network if $\alpha < 1/3$. Otherwise optimal mechanisms in the two networks have the same total reward. In section 4 we derive a general result for this social influence function showing that it always has a "bang bang" solution. Namely the optimal network can be either a star or a complete graph. For other types of utility functions this may not be the case as example 2 demonstrates.

Next we provide an example of a social influence function arising from a limited cognitive capacity model.

Example 2. An individual has a a limited cognitive capacity to be influenced by a fixed number of $k \geq 1$ agents to whom she has links. Moreover she can be influenced only once by the same person. In particular, suppose that an individual is going to randomly sample with replacement $k \geq 1$ of her friends. Before knowing who is in her sample, her social incentives to act are proportional to the expected number of different active friends in the sample. For a sample of k taken with replacement out of d_i friends, the expected number of different active friends in the sample is[8]

$$\left(1-(1-\frac{1}{d_i})^k\right)\sum_j g_{ij}x_j.$$

Therefore we can model this situation by setting

$$f(d_i)=1-(1-\frac{1}{d_i})^k.$$

The case where $k=1$ corresponds to the average comparison model from the previous example where $\alpha=0$, and the case where $k\to\infty$ corresponds to the case where α is big and agents mostly care about the number of active friends. Clearly, f is decreasing in the number of friends. Suppose now that $k=2$ and agents have the same cost c as before. Using the argument mentioned above we can derive the corresponding optimal influence mechanisms for each network. For example, in the middle network an optimal mechanism is as follows. Agent 1 is paid to act independently of others, and hence $t_1=c$. Agent 2 is paid to act only when 1 acts, i.e., $t_2=c-(1-(1-1/3)^2)$. Agents 3 and 4 must be paid rewards that induce them to act when 1 and 2 act, i.e., $t_3=t_4=c-2\,(1-(1-1/2)^2)$. Hence, the total reward in the middle network is roughly $4c-3.56$. Similarly, we find that a total reward in the complete network is approximately $4c-2.63$, and that in the star network it is $4c-3$. Thus, it is now cheaper for the principal to incentivize agents in the middle network. In a later section we introduce a general network structure that unifies the above three networks, a galaxy, and show that under some natural assumptions an optimal network is, indeed, a galaxy.

2.3 Discussion

Our central requirement that the design of incentives generate a unique equilibrium guarantees that the mechanism screens out bad equilibria, including the one

[8] Let X be a random variable equal to the number of active agents in a sample of k. Then we have $X=X_1+\cdots+X_{\sum g_{ij}x_j}$, where each X_l is a random variable equal to 1 if active agent 1 is in the sample, and equal to 0 otherwise. Taking the expectation we have $\mathbb{E}(X)=\sum\mathbb{E}(X_l)$, where the summation is over the active agents in the friendhood of i. Then for each active friend l, $\mathbb{E}(X_l)=Prob(l\text{ is in the sample})=1-(\frac{d_i-1}{d_i})^k$. And hence $\mathbb{E}(X)=\left(1-(1-\frac{1}{d_i})^k\right)\sum_j g_{ij}x_j$.

where no player takes the desired action. Indeed, several experimental papers find that subjects are often trapped in bad equilibrium outcomes in environments with externalities (e.g., Devetag and Ortmann, 2007). Hence, unless the principal pays the extra premium to screen out the bad equilibria, such equilibria are likely to prevail.

We have assumed that the players take their actions simultaneously, i.e., without observing the action of anyone else. However, exactly the same results can be obtained if we impose an order of moves on the players, and assume that each player is informed about earlier decisions made by his friends. This alternative framework requires a stronger solution concept than subgame perfection (SPE), a subgame dominant equilibrium (SDE) (Halac et al., 2020a). Furthermore, it will require that the principal choose the order of moves. Unlike SPE, which requires that the strategy profile yield a Nash equilibrium in every subgame, SDE requires that this equilibrium also use a dominant strategy in every subgame. We elaborate on this result in the Appendix.

Finally, we focus on situations where social influence induces complementarities in agents' actions. In some cases, however, social influence can lower the incentives to act. In this case of strategic substitutability our problem is much simpler. Optimal influence mechanism provides rewards sufficient to induce all agents to act when everyone else does, thus making acting a dominant strategy. Furthermore, empty networks are trivially optimal, because every new link makes it more expensive for the principal to induce action.

3 Optimal influence mechanism

In this section we characterize optimal influence mechanisms for arbitrary networks. We shall show that these mechanisms crucially depend on agents' degree centrality. Central agents will receive favorable treatment to compensate them for the strategic risk that they face. This will allow the principal to spend fewer resources on more peripheral agents.

We now generalize the construction of the optimal influence mechanisms given in the examples from the previous section. We say, an influence mechanism t is *tight* if it is incentive-inducing (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 can not lower reward of any single agent without violating the requirement of the uniqueness of the "all act" equilibrium. 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 a tight influence mechanisms.[9]

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

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

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

In particular, given a permutation of agents, we can use (2) to construct a tight influence mechanism. In this mechanism every agent i is compensated for her individual cost of taking the action, c_i, net of the social benefit she receives from her friends that are earlier in the corresponding permutation. The term $f(d_i) \sum_{j:\pi(j)<\pi(i)} g_{ij}$ in (2) captures the extent to which the principal can exploit the social influence exerted on agent i by her friends in such a mechanism. Notice that despite the fact that in equilibrium every agent acts, in a tight influence mechanism the principal cannot extract the entire social benefit from each agent. For instance, in Example 1 a permutation that corresponds to the constructed optimal influence mechanism for a complete network is $(1, 2, 3, 4)$. Here agent 2 is compensated for her cost, c, net of the social influence exerted by agent 1, i.e., $\alpha + 1/3$. Note, however, that in equilibrium the social benefit that agent 2 derives 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 players' decisions. That is, acting is her dominant strategy. This agent appears first in the permutation, and she is paid just enough to act when no one else does. Similarly, there must

[9] All the proofs are deferred to the Appendix.

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 stay still). This agent is placed second in the permutation, and receives just as much as would induce her to act 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. Call a permutation of agents π *nonincreasing* if among any two connected agents the one with a strictly higher degree appears earlier in permutation π; i.e., for all agents i and j such that $g_{ij} = 1$ and $d_i > d_j$, we have $\pi(i) < \pi(j)$, where $\pi(i)$ is the place of agent i in permutation π.

Proposition 1. *An influence mechanism $t = (t_1, \ldots, t_n)$ is optimal if and only if it is induced by a nonincreasing permutation, i.e., there exists a nonincreasing permutation π such that t is given by (2).*

To understand why the mechanisms corresponding to nonincreasing permutations are optimal consider two agents i and j. We say i *influences* j in an incentive-inducing mechanism if i and j are connected and i precedes j in the corresponding permutation. When should i influence j, or vice versa? If i influences j, then the principal can extract $f(d_j)$ from j because when deciding to act j can count on i acting, as i is provided with sufficient incentives to do so. On the other hand, if j influences i, then the principal can extract $f(d_i)$ from player i. Since f is a nonincreasing function, in order to maximally lower the total reward the principal designs the mechanism so that an agent with a higher degree influences one with a lower degree, i.e., one who appears earlier in the permutation. Note, however, that reward is not monotone in an agent's degree: it is decreasing in the number of friends with a higher degree than hers. Agents with relatively more connections than their friends get higher rewards, whereas agents with relatively fewer connections get lower rewards.

4 Optimal networks

Which social networks are the easiest to manipulate? In this section we fully characterize such optimal networks. While the results here are of interest in themselves, they also help us to understand when an interested party will prefer one network to another. Indeed, in some circumstances the principal not only controls the external rewards, but also has a certain influence on the network architecture, which is arguably the case with online platforms such as Facebook and LinkedIn.

First, we introduce three assumptions about the function f, which captures how the social influence exerted on an agent by an active friend depends on the total number of friends the agent has. Acquiring a new active friend 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 his existing active friends. Specifically, because f is decreasing in an agent's degree, each active friend now exerts lower social influence on an agent. On the other hand, the agent obtains one more active friend who influences her to take the action. Our first assumption is that the overall effect of introducing a new active friend is nonnegative. Put differently, adding a link to an active agent can only increase the incentives of other agents to take the action. Formally, consider an agent with m friends and $k \leq m$ active friends. Suppose that this agent forms a friendship with another agent who takes the action. We assume that the corresponding change in social benefit she derives from acting is nonnegative:

$$(k+1)f(m+1) - kf(m) \geq 0, \tag{3}$$

for all $m \geq 1$ and $k \leq m$. Rewriting the above as

$$k\,(f(m+1) - f(m)) + f(m+1) \geq 0,$$

and noting that f is weakly decreasing in degree, we find that the left-hand side of (3) is also weakly decreasing in k. Therefore, if (3) holds for $k = m$, it must hold for all $k \leq m$. Let $f(m) - f(m+1) = \Delta(m)$. We obtain the following assumption.

Assumption B (Benefit). For each $m \geq 1$, we have $\Delta(m) \leq \frac{1}{m} f(m+1)$.

Intuitively, the assumption is satisfied when f falls slower than the reciprocal function $x \mapsto 1/x$. Hence, the extreme case where the above condition holds with equality is the case of the average comparison model from Example 1 where $f(m) = 1/m$. Then individuals care only about a fraction of their friends taking an action. So, adding a new friend who acts when all of the existing friends also act does not affect the incentives of an agent.

Since the influence of each existing friend declines with the addition of a new friend, it is natural to think that marginal dilution is smaller for someone with more connections. Our second assumption formalizes the idea, requiring f to be convex.

Assumption C (Convexity). For each $m \geq 1$, we have $\Delta(m) \geq \Delta(m+1)$.

Although convexity alone allows us to considerably narrow down a list of networks, we will use a stronger assumption in order to pin down a precise structure of optimal networks.

Assumption SC (Strong Convexity). For each $m \geq 1$, we have

$$\Delta(m) \geq \Delta(m+1)\frac{n+2}{n}.$$

Strong convexity requires that the absolute values of forward differences of f be decreasing at multiple $n/(n+2)$ or lower, where n is a total number of agents. One can check that the above examples satisfy both the B and SC assumptions. Furthermore, note that in a large network our strong convexity condition converges to a standard notion of convexity.

We call a network *galaxy* 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 in the graph, and nodes in P are connected only to the nodes in[10] S. A star

[10]These graphs are known in the literature as complete core-periphery graphs or complete split graphs or inter-linked stars. In a complete split graph on n vertices there is a clique of s vertices, $1 \leq s \leq n$, and and an independent set on the remaining $n - s$ vertices such that each vertex of the clique is adjacent to each vertex of the independent set.

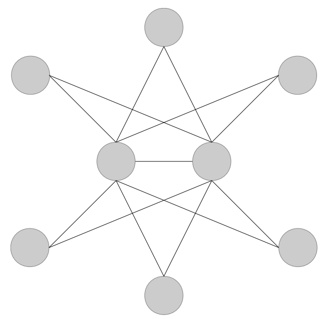

Figure 2: A galaxy with 2 stars and 6 periphery nodes.

network and a complete network are the special cases of a galaxy, and so is the network in the middle panel of Figure 1, as well as the network with two stars and six periphery nodes in Figure 2. In general, on n nodes there are only $n-1$ different galaxies (up to a permutation of nodes). Finally, we can state our first characterization of optimal networks.

Proposition 2. *Suppose assumptions B and SC hold. Then an optimal network is a galaxy.*

It appears that the easiest networks to manipulate are also the most unequal in terms of their degree distributions among all similarly dense networks.[11] The presence of highly connected stars allows for the most efficient exploitation of their social influence on periphery agents. However, the principal does not necessarily strive to achieve the maximal absolute inequality like in a simple star network: more connected networks may require lower total rewards because they allow for more opportunities for social influence. Specifically, an optimal number of stars in a galaxy depends on the shape of function f. For example, it is clear that if $f(d_i) = 1/d_i$, then a star network is optimal, and if $f(d_i) = k$ for some positive constant k, then a complete network is optimal. Similarly, in Example 2 we considered a function f which makes a galaxy with two stars and two periphery nodes optimal.

[11] We can make this statement precise using a notion of majorization. For any network, let $d = (d_1, \dots, d_n)$ denote a sequences of degrees of its nodes arranged in a non-increasing 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 galaxy majorizes a sequence of degrees of nodes in every network with the same number of edges.

We now briefly return to example 1 where agents care about some combination of local average and aggregate action. Our next result completely characterizes optimal networks for a specific social influence function in this example.

Corollary. *Suppose $f(d_i) = \alpha + 1/d_i$, where $\alpha \geq 0$. We have:*
(i) if $\alpha < \frac{1}{n-1}$, then an optimal network is a star,
(ii) if $\alpha > \frac{1}{n-1}$, then an optimal network is a complete network,
(iii) if $\alpha = \frac{1}{n-1}$, then each galaxy is optimal.

To better understand Proposition 2 consider an optimal influence mechanism for a galaxy. In any nonincreasing permutation stars must be followed by periphery nodes. Hence, the stars receive high differentiated rewards, whereas the periphery nodes receive the lower identical rewards (because periphery nodes are not connected between themselves). This situation corresponds to a dominance cascade, starting with the stars who act one after another and rely only on the other stars ahead of them, and finishing with the periphery agents relying only on the stars. That is, the principal provides high-powered incentives to the stars, and then exploits their social influence on the remaining agents. The formal proof of the proposition is rather long and we defer it to the Appendix. Instead, we now give some intuition for the result.

Consider an effect of a new link on a total reward in an optimal influence mechanism. From the principal's perspective connecting two agents has a cost and a benefit due to the induced change in an optimal influence mechanism. When the principal connects two agents, the one who appears earlier in the permutation will exert social influence on the other. Therefore the principal will have to increase the payoff of the influencing agent (due to the dilution), but can reduce the payoff of the influenced agent. Specifically, fix a network and an optimal influence mechanism, and consider connecting agents i and j, such that i precedes j in the corresponding nonincreasing permutation. Then i would suppose that she has one more inactive friend, and hence experiences the dilution of the social influence exerted by friends who appear earlier in the cascade (and who, she supposes, take the action). Thus in order to be induced to act, i must be paid more by the principal. On the other hand, now j gets to be influenced by i, and hence can be paid less (this follows from assumption B). Because convexity implies that, other things being equal, a cost of a link is lower for an agent with a higher degree, intuition suggests that an optimal distribution of links must be unequal: some individuals with many connections must influence less-connected friends. And, indeed, for a given a number of links a galaxy is the most unequal network in terms of its degree distribution.

Our proof involves two stages. We first show that if in an optimal mechanism two agents i and j influence, respectively, x and y who follow them in a corresponding nonincreasing permutation, then by convexity at least one of them must influence both x and y; otherwise it would be profitable to let i influence y instead of j, or j influence x instead of i. We show that this implies that there is a nonincreasing permutation such that if i influences j, then it must be that each agent preceding i in the permutation also influences j. We use this observation along with some symmetry arguments to show that there in an independent set[12] of symmetric "periphery" agents who are influenced by everyone else.

In the second stage of the proof we show that each pair of non-periphery agents is connected. Consider the two cases. First, suppose that more than half of agents are periphery agents. Then we show that the benefit assumption implies that all non-periphery agents must be connected between themselves, and hence are stars. Second, suppose that less than half of the agents are periphery agents. Then using the strong convexity assumption we show that the benefit of influencing a periphery agent is less than the benefit of influencing a non-periphery agent, and hence again it is optimal to connect all non-periphery agents between themselves.[13] Therefore, the entire population is partitioned into stars and periphery as in a galaxy.

4.1 Optimal number of stars in a galaxy

Fix two functions f and h, such that Benefit and Strong convexity assumptions are satisfied. Moreover, without loss of generality let $f(1) = h(1) = 1$[14]. We say

[12] A 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.

[13] Hence, the term of $n/2$ in the convexity assumption.

[14] This is a normalization which is without loss of generality when we would like to compare optimal number of stars in galaxies. One can always appropriately scale the functions in order

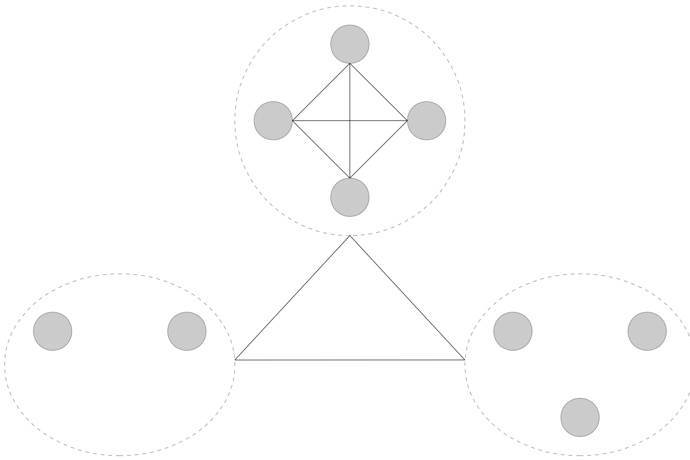

Figure 3: A 2-cluster galaxy with 4 stars.

that function f is *flatter* than function h if

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

for all $k = 1, 2, \ldots, n-1$. Intuitively, f admits a uniformly lower dilution effect than h. Given function f, let $s^*(f)$ denote the optimal number of stars in a galaxy.

Proposition. *If f is flatter than h, then $s^*(f) \geq s^*(h)$.*

Proof. For function f, let a total reward required in a galaxy with s stars be G_f^s, given by

$$G_f^s = \sum c_i - \frac{s(s-1)}{2} f(n-1) - (n-s)sf(s).$$

Suppose f is flatter than h. For $s = 1, 2, \ldots, n-1$, consider the difference between the corresponding rewards required in galaxies with s stars,

$$G_f^s - G_h^s = \frac{s(s-1)}{2} \left[h(n-1) - f(n-1)\right] + (n-s)s \left[h(s) - f(s)\right].$$

Note that the above must be weakly negative, meaning that under f total required reward is lower for all galaxies. Define $h(s) - f(s) = \delta_s$. For $k = 2, 3, \ldots, n-1$,

for this condition to hold.

we can rewrite $f(k)$ in terms of it's consecutive differences as

$$f(k) = 1 - \sum_{i=1}^{k-1} \Delta_f(i),$$

where $\Delta_f(i) = f(i) - f(i+1)$, for $i = 1, 2, \ldots, n-2$. Hence, we have

$$\begin{aligned} \delta_s &= \sum_{i=1}^{s-1} \Delta_f(i) - \sum_{i=1}^{s-1} \Delta_h(i), \\ &= \sum_{i=1}^{s-1} \left[\Delta_f(i) - \Delta_h(i)\right]. \end{aligned}$$

Because f is flatter than h, we have $\Delta_f(i) \leq \Delta_h(i)$ for each $i = 1, 2, \ldots, n-2$. Hence, $\delta_s \leq 0$ and is weakly decreasing in s, for $s = 1, 2, \ldots, n-1$. Next we show that

$$G_f^s - G_h^s \geq G_f^{s+1} - G_h^{s+1}, \tag{4}$$

for each $s = 1, 2, \ldots, n-1$. That is as one moves from h to f, the required rewards decrease more in galaxies with more stars. Consider the difference:

$$\begin{aligned} G_f^s - G_h^s - \left(G_f^{s+1} - G_h^{s+1}\right) &= \frac{s(s-1)}{2}\delta_{n-1} + (n-s)s\delta_s - \frac{(s+1)s}{2}\delta_{n-1} - (n-s-1)(s+1)\delta_{s+1}, \\ &= -s\delta_{n-1} + (n-s)s\delta_s - (n-s-1)(s+1)\delta_{s+1}, \\ &\geq -s\delta_{n-1} + (n-s)s\delta_s - (n-s-1)(s+1)\delta_s, \\ &= -s\delta_{n-1} + (2s-n+1)\delta_s, \\ &\geq -s\delta_s + (2s-n+1)\delta_s, \\ &= (s-n+1)\delta_s, \\ &\geq 0. \end{aligned}$$

Where the first two inequalities follow from $\delta_{n-1} \leq \delta_{s+1} \leq \delta_s \leq 0$.

Finally, fix $s^*(h)$, and consider moving from h to f. From (4) it follows that reward required in $s^*(h)$-galaxy decreases weakly more than in any galaxy with fewer stars. Hence, any galaxy with fewer stars cannot be optimal under f. □

4.2 k-cluster galaxies

Relaxing the assumption of strong convexity leads us to generalize the notion of a galaxy. We call a network a *k-cluster galaxy* if its nodes can be partitioned into a clique C_0 and k independent sets $C_1, \ldots, C_k$ such that each node in each of these sets is connected to every other node in a network. More formally, a network G is a k-cluster galaxy if there exists a partition of its nodes into $k+1$ subsets $C_0, \ldots, C_k$ satisfying the following three properties:
(i) if $i \in C_0$ and $j \in C_0$ then $g_{ij} = 1$;
(ii) if $i \in C_l$ and $j \in C_l$ then $g_{ij} = 0$ for each $l \geq 1$;
(iii) if $l < h$ and $i \in C_l$ and $j \in C_h$, then $g_{ij} = 1$.

Thus a galaxy is a 1-cluster galaxy. Figure 3 provides an example of a 2-cluster galaxy. Here each node inside any grey circle is connected to each node in the other two grey circles. We call nodes in the upper circle stars. Each star is connected to all 8 other nodes in the network. At the same time, nodes in the left circle are connected to 7 other nodes each, and nodes in the right circle are connected to 6 other nodes each.

Proposition 3. *Suppose assumptions B and C hold. Then an optimal network is a k-cluster galaxy.*

5 Directed networks

In some situations the social influence flows only in one direction. For instance, on Twitter or Instagram one can follow other users and be influenced by their posts without necessarily being followed by them in return. In this section we extend our model to capture this one-directional social influence.

We represent a directed social network by a directed graph with, possibly, an asymmetric adjacency matrix G, where $g_{ij} = 1$ if and only if i can influence j and $g_{ij} = 0$ otherwise; by convention $g_{ii} = 0$. We let d_i^{in} denote the in-degree of node i in G, i.e., $d_i^{in} = \sum_j g_{ji}$. It represents the number of individuals who can influence i. Given network G and action profile $x = (x_1, \ldots, x_n)$, the payoff of individual i

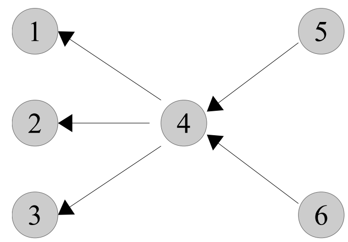

Figure 4: A directed social network.

from acting is given by

$$U_i(x, G) = f(d_i^{in}) \sum_j g_{ij} x_j + t_i - c_i.$$

We still assume that $f(d_i^{in}) > 0$ and that it is a weakly decreasing function of the number of individuals who can influence i, i.e., $f(m) \geq f(m+1)$ for $m = 1, 2, \ldots, n$. Note, however, that the social influence exerted on agent i does not depend on how many other agents i influences herself. As before, the payoff from abstaining is zero.

Optimal influence mechanisms for directed networks are more involved. In particular, ordering agents according to their in-degrees no longer induces optimal influence mechanisms.

Example 3. Consider a directed network in Figure 4. Here agents 1, 2, and 3 are all influenced by agent 4, who is in turn is influenced only by agents 5 and 6. Although agent 4 has the highest in-degree, in a permutation corresponding to an optimal influence mechanism she must be placed after agents 5 and 6. In the corresponding reward profile $t_i = c_i - f(1)$ for $i = 1, 2, 3$, $t_4 = c_4 - 2f(2)$, and $t_i = c_i$ for $i = 5, 6$.

Although there is no simple structure for optimal influence mechanisms in directed networks, it is clear that optimal directed networks must be easier to manipulate than optimal undirected ones. Indeed, in an optimal influence mechanism for an undirected network, the principal exploits only the social influence of agents who are earlier in a nonincreasing permutation on agents who are later in the permutation. Now, for a given undirected network we can construct a directed network by directing links toward agents who are later in the nonincreasing permutation. This, clearly, can only decrease the total reward in an optimal influence mechanism because each directed link now has positive benefit at no cost. In particular, we do not need to give higher-powered incentives to earlier players in the permutation to compensate them for the strategic uncertainty from the possibility that subsequent agents in the order may avoid taking the action. Consider the example below.

Example 4. There are four agents who have the same cost c and $f(d_i) = \alpha + 1/d_i$. Suppose first that the social influence is undirected. As we have seen in Example 1, if $\alpha > 5/12$, then the complete network on the left in Figure 1 is more susceptible to external influence than the other two networks. Because these are all the galaxies with four agents, Proposition 2 implies the complete network is also optimal with the total reward $4c - 6\alpha - 2$. Now consider a directed network obtained from the complete one by directing each link toward a later agent in permutation $(1, 2, 3, 4)$. In an optimal influence mechanism, $t_1 = c$, $t_2 = c - \alpha - 1$, $t_3 = c - 2\alpha - 1$, $t_4 = c - 3\alpha - 1$. Hence, the total reward is $4c - 6\alpha - 3$, which is strictly lower than in an undirected network.

The above example suggests the following intuition. Consider adding a directed link from one agent to another. By the benefit assumption the incentives of an influenced agent (the one to whom the link is directed) to take the action increase if an influencer (the one from whom the link is directed) takes the action. On the other hand, the incentives of the influencer are not affected, which is in contrast to the situation with undirected links where both agents experience the dilution of influence from others. Hence, adding a directed link from the perspective of the principal has only the benefit without the cost. Thus in an optimal directed network each pair of agents must be connected by a one-directional link.

We can now describe the architecture of optimal directed networks. A *tournament* is a directed network obtained by assigning a direction for each edge in an undirected complete network. A tournament is *acyclic* if there are no cycles.

Proposition 4. *Suppose that assumption B holds. Then an optimal directed net-*

work is an acyclic tournament.

Proof. Fix an optimal directed network G and let t be the corresponding optimal influence mechanism. From the proof of Proposition 1 we know that there is a permutation of agents π, such that $t = t(\pi)$. We shall show that network G must be an acyclic tournament such that each node points only to nodes after it in π; i.e., for each i and j, $g_{ij} = 1$ if and only if $\pi(i) < \pi(j)$. For the sake of contradiction suppose that there exist i and j such that $\pi(i) < \pi(j)$ and $g_{ij} = 0$. Consider introducing a link from i to j. Then, keeping the permutation of agents unchanged, the corresponding change in the total reward is negative: t_i is unchanged and t_j decreases by the Benefit assumption (j is influenced by one more active agent). Therefore the total reward in an optimal permutation after the introduction of a link from i to j must also decrease, a contradiction. Assume now by way of contradiction that there are i and j such that $g_{ij} = 1$ and $\pi(i) > \pi(j)$. Fix permutation π, and note that removing the link from i to j does not affect t_i, and can only decrease t_j because the in-degree of j decreases by one. Therefore there must exist an optimal directed network that is an acyclic tournament. □

We conclude with a couple of simple observations about distinctions between optimal directed and undirected networks and corresponding influence mechanisms. First, when social influence is asymmetric, dense tournament networks, where every agent either influences or is influenced by every other agent, are optimal. By contrast, galaxies with only a few links can be optimal if the influence is symmetric. Second, in the model of undirected social influence the optimal number of stars in a galaxy is tied to the shape of function f, whereas in the directed model there is a unique optimal network (up to a permutation of agents). Hence, an optimal topology of a directed network is more "robust" to the specification of the model: the principal does not need to know neither the exact function f, nor whether it is convex or not.

6 Active social influence

In our benchmark model a social influence exerted by an agent varies across her friends: more-connected friends are influenced less than the less-connected ones. One way to think about this model is that social influence is passive and hence its strength depends only on the characteristics of the influenced agents. It is not difficult to think of the opposite situation where social influence is active. A familiar case is when someone who recently became a vegan is actively persuading her friends to adopt this new lifestyle. In a reduced form model that describes this situation a strength of social influence naturally depends on the characteristics of the influencing agent, and hence the same agent will be influenced differently by every friend. Here we investigate this alternative view of the mechanics of social influence, and show that many insights from the original model continue to hold.[15]

Given network G and action profile $x = (x_1, \ldots, x_n)$, we let the payoff of individual i from taking an action be given by

$$U_i(x, G) = \sum_j g_{ij} x_j f(d_j) + t_i - c_i. \tag{5}$$

The difference between the above expression and (1) in the original model is the term $\sum_j g_{ij} x_j f(d_j)$. Here $f(d_j)$ captures the social influence that an agent j exerts on each of her friends, and in particular on agent i, whereas everything else is as before. Hence, the social benefit of i from taking the action is comprised of the sum of heterogeneous influences from her friends. As before, social influence is positive, $f(m) > 0$, and is a weakly decreasing function of the number of friends of individual i, i.e., $f(m) \geq f(m+1)$ for $m = 1, 2, \ldots, n$. However, now the interpretation is that an agent splits her attention between influencing each of her friends, and hence someone with more friends will influence each of them less. We call models with payoff functions given by (1) and (5) models of *passive* and *active* social influence.

[15] In fact, one can consider a more general model that incorporates both, passive and active, influence. In such a model a payoff function of an individual i will be given by $U_i(x, G) = f(d_i) \sum_j g_{ij} x_j h(d_j) + t_i - c_i$, where f and h are two nonincreasing functions. However, for the sake of clarity we do not introduce such a general model in this paper.

It is straightforward to confirm that the result analogous to Lemma 1 carries over to the model of active social influence, respectively.

Lemma 2. *In a model of active social influence an influence mechanism $t = (t_1, \ldots, t_n)$ is tight if and only if there exists permutation π such that for all i,*

$$t_i = c_i - \sum_{j:\pi(j)<\pi(i)} g_{ij} f(d_j), \tag{6}$$

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

It appears that there is a simple relation between tight mechanisms in both models. Consider a value that a link between agents i and j contributes to the total reward in a tight mechanism in a passive influence model. As pointed out in Section 3, if i precedes j in the corresponding permutation, then this link adds $f(d_j)$ to the total reward. Otherwise it adds $f(d_i)$. By contrast, from (6) it follows that in the active influence model the same link instead contributes $f(d_i)$ and $f(d_j)$, respectively. Therefore, given a tight influence mechanism t in the passive influence model, we can obtain a tight influence mechanism $\hat{t}$ in the active influence model with the same total reward, i.e., $\sum t_i = \sum \hat{t}_i$. Specifically, let π be a permutation corresponding to t, and $\hat{\pi}$ be a permutation with the reverse order, i.e. $\pi(i) < \pi(j)$ if and only if $\hat{\pi}(i) > \hat{\pi}(j)$. Then $\hat{t}$ is the influence mechanism obtained from $\hat{\pi}$ by (6).

The conclusion is that, for any network, optimal mechanisms in the active influence model are obtained from nondecreasing permutations of agents, and have exactly the same total reward as the optimal mechanisms in the passive influence model (given the same function f). Hence, our optimal network result carries over from the benchmark model without modification. We summarize these observations in the next proposition.

Proposition 5. *Consider the active influence model. Then:*

(i) An influence mechanism $t = (t_1, \ldots, t_n)$ is optimal if and only if it is induced by a nondecreasing permutation, i.e., there exists a nondecreasing permutation π such that t is given by (6).

(ii) Suppose that assumptions B and SC (C) hold. Then an optimal network is a galaxy (k-cluster galaxy).

In the active influence model the role of the degree centrality in the determination of agents' rewards is flipped. Now agents with a lower number of friends tend to be favored by the principal. It may help to spell out why this result is not that surprising. The principal wishes to exploit the social influence exerted by agents. Between the two friends one must influence the other. In the passive influence model the principal lets an agent with fewer friends be influenced by the other because her attention is less diluted and she is swayed more by each friend. By contrast, in the active influence model the principal lets an agent with fewer friends actively influence the other because her influence is less diluted.

References

Masaki Aoyagi. Coordinating adoption decisions under externalities and incomplete information. *Games and Economic Behavior*, 77(1):77 – 89, 2013.

David Arthur, Rajeev Motwani, Aneesh Sharma, and Ying Xu. Pricing strategies for viral marketing on social networks. *ArXiv e-prints*, 2009.

Moshe Babaioff, Michal Feldman, Noam Nisan, and Eyal Winter. Combinatorial agency. *Journal of Economic Theory*, 147(3):999–1034, 2012.

Coralio Ballester, Antoni Calvó-Armengol, and Yves Zenou. Who's who in networks. Wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013.

Shai Bernstein and Eyal Winter. Contracting with heterogeneous externalities. *American Economic Journal: Microeconomics*, 4(2):50–76, 2012.

Francis Bloch. Targeting and pricing in social networks. In *The Oxford Handbook of the Economics of Networks*. 2016.

Francis Bloch and Nicolas Querou. Pricing in social networks. *Games and Economic Behavior*, 80(0):243 – 261, 2013.

Yann Bramoullé and Rachel Kranton. Public goods in networks. *Journal of Economic Theory*, 135(1):478–494, 2007.

Yann Bramoullé, Rachel Kranton, and Martin D'amours. Strategic interaction and networks. *American Economic Review*, 104(3):898–930, 2014.

Ozan Candogan, Kostas Bimpikis, and Asuman Ozdaglar. Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905, 2012.

Giovanna Devetag and Andreas Ortmann. When and why? A critical survey on coordination failure in the laboratory. *Experimental Economics*, 10(3):331–344, 2007.

Pedro Domingos and Matt Richardson. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 57–66, 2001.

Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5): 1047–1071, 1993.

Itay P Fainmesser and Andrea Galeotti. Pricing network effects. *The Review of Economic Studies*, 83(1):165–198, 2016.

Andrea Galeotti, Benjamin Golub, and Sanjeev Goyal. Targeting interventions in networks. *arXiv preprint arXiv:1710.06026*, 2017.

Christian Ghiglino and Sanjeev Goyal. Keeping up with the neighbors: Social interaction in a market economy. *Journal of the European Economic Association*, 8(1):90–119, 2010.

Sanjeev Goyal and Sumit Joshi. Networks of collaboration in oligopoly. *Games and Economic behavior*, 43(1):57–85, 2003.

Sanjeev Goyal, Marco J Van Der Leij, and José Luis Moraga-González. Economics: An emerging small world. *Journal of political economy*, 114(2):403–412, 2006.

Mark Granovetter. Threshold models of collective behavior. *American journal of sociology*, 83(6):1420–1443, 1978.

Marina Halac, Ilan Kremer, and Eyal Winter. Raising capital from heterogeneous investors. *American Economic Review*, 110(3):889–921, 2020a.

Marina Halac, Elliot Lipnowski, and Daniel Rappoport. Rank uncertainty in organizations. *Working paper*, 2020b.

Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, pages 189–198, 2008.

Andrea Ichino and Giovanni Maggi. Work environment and individual background: Explaining regional shirking differentials in a large Italian firm. *The Quarterly Journal of Economics*, 115(3):1057–1090, 2000.

Matthew O. Jackson. An overview of social networks and economic applications. In *Handbook of Social Economics*, volume 1, pages 511–585. Elsevier, 2011.

Matthew O. Jackson and Brian W. Rogers. Relating network structure to diffusion properties through stochastic dominance. *The BE Journal of Theoretical Economics*, 7(1):1–16, 2007.

Matthew O. Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.

Matthew O. Jackson and Leeat Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, 2007.

Matthew O Jackson and Yves Zenou. Games on networks. In *Handbook of Game Theory with Economic Applications*, volume 4, pages 95–163. Elsevier, 2015.

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

Michael D König, Claudio J Tessone, and Yves Zenou. Nestedness in networks: A theoretical model and some applications. *Theoretical Economics*, 9(3):695–752, 2014.

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

Alexandre Mas and Enrico Moretti. Peers at work. *American Economic Review*, 99(1):112–45, 2009.

Stephen Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, 2000.

In-Uck Park. A simple inducement scheme to overcome adoption externalities. *B.E. Journal of Theoretical Economics*, 4(1), 2004.

Eleonora Patacchini and Yves Zenou. Juvenile delinquency and conformism. *Journal of Law, Economics, and Organization*, 28(1):1–31, 2012.

Bruce Sacerdote. Peer effects with random assignment: Results for Dartmouth roommates. *The Quarterly Journal of Economics*, 116(2):681–704, 2001.

Ilya Segal. Contracting with externalities. *The Quarterly Journal of Economics*, 114(2):pp. 337–388, 1999.

Ilya Segal. Coordination and discrimination in contracting with externalities: Divide and conquer? *Journal of Economic Theory*, 113(2):147 – 181, 2003.

Eduard Talamàs and Omer Tamuz. Network cycles and welfare. *Available at SSRN 3059604*, 2017.

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

Eyal Winter. Incentives and discrimination. *The American Economic Review*, 94 (3):764–773, 2004.

7 Appendix

Proof of Lemma 1

Proof. For a permutation π, let $t(\pi) = (t_1, \ldots, t_n)$ be an influence mechanism assigned to π by (2). First, we show that $t(\pi)$ is INI. If not, then there must exist an equilibrium where there is a nonempty subset of agents S such that agents in S do not act and agents outside of S do, i.e., $x_i = 0$ if and only if $i \in S$. Note that t_i specified by (2) is sufficient to induce i to act, given that all agents preceding i in π act, no matter what the other agents do. Thus $\pi(1) \notin S$ because agent $\pi(1)$ weakly prefers to act no matter what the other agents do. By induction suppose that for $k = 1, \ldots, n-1$ it must be that agents $\pi(1), \ldots, \pi(k)$ are not in S. Then $\pi(k+1)$ weakly prefers to act, and hence is also not in S. It follows that S must be empty. Hence, $t(\pi)$ is INI.

Now we show that we cannot lower a reward of any single agent without violating the uniqueness condition. Note, that we cannot lower the reward of $\pi(n)$ given by (2) because otherwise agent $\pi(n)$ would strictly prefer not to act. Moreover given $t_{\pi(n)}$, agent $\pi(n)$ strictly prefers not to act if any of the agents connected to $\pi(n)$ and preceding her in π do not act. For agent j, let

$$
\begin{aligned}
F_j = \{i | \pi(i) > \pi(j), \exists j_1, \ldots, j_m \\
\text{s.t. } g_{jj_1} = g_{j_1 j_2} = \cdots = g_{j_m i} = 1 \text{ and } \pi(j) < \pi(j_k) < \pi(i) \text{ for all } k\}.
\end{aligned}
$$

Now suppose that for $k = 1, \ldots, n-1$, it must be that each agent from $\pi(k+1), \ldots, \pi(n)$ strictly prefers not to act if all the agents connected and following her in π and at least one of the agents connected and preceding her in π does not act.

Then we cannot lower the reward of agent $\pi(k)$ given by (2). Indeed, if we lower the reward then there exists an equilibrium where agent $\pi(k)$ and all agents in $F_{\pi(k)}$ do not act, while everyone else does. Hence, this establishes that we cannot lower a reward of any single agent. Therefore $t(\pi)$ is tight.

For each tight t, we show that there exists a permutation π such that $t = t(\pi)$. Note, that there must exist an agent a_1 who weakly prefers to act even when no one else does. If not, then there would exist an equilibrium where no one acts, contradicting that t is INI. Hence we must have

$$t_{a_1} \geq c_{a_1}. \tag{7}$$

Clearly, (7) implies that a_1 weakly prefers to act no matter what the other agents do. We let $\pi(1) = a_1$. Now we proceed to define a permutation corresponding to t by induction. Suppose that for $k = 1, \ldots, n-1$, there is an agent a_k who weakly prefers to act if agents before her, $a_1, \ldots, a_{k-1}$, act regardless of what others do. Then there exists an agent a_{k+1} who weakly prefers to act when agents $a_1, \ldots, a_k$ act and others do not. Otherwise there would exist an equilibrium where $a_1, \ldots, a_k$ act and others do not, contradicting that t is INI. Hence, we must have

$$t_{a_{k+1}} \geq c_{a_{k+1}} - f(d_{a_{k+1}}) \sum_{i=1}^{k} g_{a_{k+1}a_i}. \tag{8}$$

Clearly, (8) implies that a_{k+1} weakly prefers to act when $a_1, \ldots, a_k$ also act, no matter what the other agents do. Let $\pi(k+1) = a_{k+1}$. 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 $k = 1, \ldots, n$ and some agent a_k we have that (8) holds with a strict inequality. But then t cannot be tight because by slightly lowering t_{a_k}, we can reduce the total reward while keeping t INI. Thus we have established a surjection from a set of permutations to a set of tight influence mechanisms. $\square$

Proof of Proposition 1

Proof. Given any permutation π, the total reward in the corresponding tight influence mechanism is

$$\sum_i t_i = \sum_i c_i - \sum_i \sum_{j:\pi(j)<\pi(i)} g_{ij} f(d_i). \tag{9}$$

Hence, each pair i and j, such that $g_{ij} = 1$, contributes to (9) either $-f(d_i)$ if $\pi(j) < \pi(i)$, and $-f(d_j)$ otherwise. Therefore, because f is weakly decreasing, to maximize (9) for each pair i and j such that $g_{ij} = 1$ and $d_i > d_j$ we must have $\pi(i) < \pi(j)$, and hence π corresponding to an optimal influence mechanism must be nonincreasing. □

Proof of Proposition 2

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 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 agent i, let $N_i^{\pi,-} \subseteq N_i$ denote a subset of i's friends who follow i in permutation π, i.e., $N_i^{\pi,-} = \{j | \pi(j) > \pi(i)\}$. We say that agents in $N_i^{\pi,-}$ are *influenced* by i. Let $d_i^{\pi,-} = |N_i^{\pi,-}|$ and $d_i^{\pi,+} = |N_i \setminus N_i^{\pi,-}|$.

Lemma 3. *Fix four different agents* i, j, x*, and* y *such that* $\max\{i, j\} < \min\{x, y\}$. *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. Then a total reward in an optimal influence mechanism changes by

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

Similarly, consider replacing a link between i and x by a link between j and x.

The corresponding change in total reward is

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

Because G is optimal, each of the above replacements must weakly increase a total reward:

$$\begin{aligned}d_i^{id,+}\left(f(d_i)-f(d_i+1)\right)-d_j^{id,+}\left(f(d_j-1)-f(d_j)\right)\geq 0,\\ d_j^{id,+}\left(f(d_j)-f(d_j+1)\right)-d_i^{id,+}\left(f(d_i-1)-f(d_i)\right)\geq 0.\end{aligned}$$

Combining the inequalities we get

$$\frac{f(d_j-1)-f(d_j)}{f(d_i)-f(d_i+1)}\leq\frac{d_i^{id,+}}{d_j^{id,+}}\leq\frac{f(d_j)-f(d_j+1)}{f(d_i-1)-f(d_i)}. \tag{10}$$

By convexity we have

$$\begin{aligned}f(d_j)-f(d_j+1)\leq f(d_j-1)-f(d_j),\\ f(d_i)-f(d_i+1)\leq f(d_i-1)-f(d_i),\end{aligned}$$

with equality only when the RHSs are zero. Clearly, if at least one RHS is not zero, then (10) is inconsistent. Hence, one of the two replacements must strictly decrease total reward. On the other hand, 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 decreases total reward. □

Lemma 4. *There exists a nonincreasing permutation π of agents such that*

$$N^{\pi,-}_{\pi^{-1}(1)}\supseteq N^{\pi,-}_{\pi^{-1}(2)}\supseteq\cdots\supseteq N^{\pi,-}_{\pi^{-1}(n)}. \tag{11}$$

Proof. Let $\pi^{-1}(k)=s_k$ for $k=1,\ldots,n$. To construct $(s_1,s_2,\ldots,s_k)$ we begin from a nonincreasing identity permutation id. Clearly, $N_1^{id,-}\supseteq N_2^{id,-}$ because 1 has zero cost of a link and thus must be connected to each node. Let, $s_1=1$ and $s_2=2$. For the induction argument suppose that for $k\leq n$ there is a nonincreasing permutation π_k such that (i) $N^{\pi_k,-}_{\pi_k^{-1}(1)} \supseteq N^{\pi_k,-}_{\pi_k^{-1}(2)} \supseteq \cdots \supseteq N^{\pi_k,-}_{\pi_k^{-1}(k)}$, (ii) $\pi_k^{-1}(l) = l$ for $l > k$. Now we construct nonincreasing permutation π_{k+1} satisfying the two conditions above. Suppose that $\pi_k^{-1}(k) = x$. We show that either $N^{\pi_k,-}_x \supseteq N^{\pi_k,-}_{k+1}$ or $N^{\pi_k,-}_x \subseteq N^{\pi_k,-}_{k+1}$. For the sake of contradiction suppose there exist i and j such that $i \in N^{\pi_k,-}_x$, $i \notin N^{\pi_k,-}_{k+1}$ and $j \in N^{\pi_k,-}_{k+1}$, $j \notin N^{\pi_k,-}_x$. First, if $i \neq k+1$, then by Lemma 1 we have either $g_{xj} = 1$, or $g_{(k+1)i} = 1$, or both, a contradiction. Second, suppose that $i = k+1$. By the induction assumption, $N^{\pi_k,out}_x \subseteq N^{\pi_k,out}_{\pi_k^{-1}(l)}$ 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., $d^{\pi_k,+}_{k+1} > d^{\pi_k,+}_x$. 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 corresponding change in total reward must be strictly negative 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 $N^{\pi_k,-}_x \supseteq N^{\pi_k,-}_{k+1}$, a contradiction. Thus we have established that either $N^{\pi_k,-}_x \supseteq N^{\pi_k,-}_{k+1}$ or $N^{\pi_k,-}_x \subseteq N^{\pi_k,-}_{k+1}$. Now if $N^{\pi_k,-}_x \supseteq N^{\pi_k,-}_{k+1}$, then let $\pi_{k+1} = \pi_k$. Clearly, such π_{k+1} satisfies (i) and (ii). On the other hand, if $N^{\pi_k,-}_x \subset N^{\pi_k,-}_{k+1}$, define π_{k+1} in the following way. Move x one position up in the permutation; i.e., let $\pi_{k+1}^{-1}(k+1) = x$. Then, by the same argument as above either $N^{\pi_k,-}_{\pi_k^{-1}(k-1)} \supseteq N^{\pi_k,-}_{k+1}$ or $N^{\pi_k,-}_{\pi_k^{-1}(k-1)} \subseteq N^{\pi_k,-}_{k+1}$. If $N^{\pi_k,-}_{\pi_k^{-1}(k-1)} \supseteq N^{\pi_k,-}_{k+1}$, 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 $N^{\pi_k,-}_{\pi_k^{-1}(k-1)} \subset N^{\pi_k,-}_{k+1}$ 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 $N^{\pi_k,-}_{\pi_k^{-1}(k-2)} \supseteq N^{\pi_k,-}_{k+1}$, 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, clearly, in the end it produces a permutation π_{k+1} satisfying (i) and (ii). Iterating the procedure yields π_n, and finally letting $\pi = \pi_n$ we obtain the required permutation. □

Lemma 5. *If $k \leq m$, then $(k+1)f(m+1) - kf(m)$ is:*
(i) weakly decreasing in k, given m, and

(ii) increasing in m, given $k \geq n/2$.

Proof. Part (i) follows from f being a weakly decreasing function.

To prove (ii) recall that by strong convexity, for each $m \geq 1$, we have $\Delta(m) \geq \Delta(m+1)(1 + \frac{1}{n/2})$. Hence, for each $m \geq 1$ and $k \geq n/2$ we have $\Delta(m) \geq \Delta(m+1)(1 + \frac{1}{k})$. Rewriting this we get:

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

□

Lemma 6. *Fix i and j such that $d_i = d_j = m$, $d_i^+ = d_j^+ = k$. If $2k < m$, then $g_{ij} = 1$.*

Proof. For the sake of contradiction suppose that $g_{ij} = 0$. Consider the change in the total reward 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 to weakly decreases a total reward if:

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

By the Benefit assumption we have

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

and hence the condition holds for $2k < m$, a contradiction to the optimality of G. □

Now we are ready to prove the main result.

Proof of Proposition 2. Take a nonincreasing permutation π satisfying (11), and let $\pi^{-1}(k) = s_k$ for $k = 1, \ldots, n$. Call agent s_i a *sink* if $N_{s_i}^{\pi,-} = \emptyset$. First, we show that each non-sink is connected to each sink. Note that each sink must have the same degree. For the sake of contradiction suppose that sinks x and y are such that $d_x > d_y$. Then by construction of π 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 increase the principal's expenses. 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. Now suppose that x and y are two sinks and non-sink s_j connects to x but not to y. Then by Lemma 2 each s_k, $k \leq j$ connects to x and each s_l , $l \geq j$ does not connect to y. Hence, there are strictly fewer agents connected to y than to x, a contradiction to the fact that x and y must have the same degree.

Second, we show that all non-sinks are connected. Let s_k be the last non-sink in sequence $(s_1, s_2, \ldots, s_n)$. First, we show that s_{k-1} connects to s_k. For the sake of contradiction suppose not. Then $N^{\pi,-}_{s_{k-1}} = N^{\pi,-}_{s_k}$ and $x \in N^{\pi,-}_{s_{k-1}}$ 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 $s_{k-1} \notin N^{\pi,-}_{s_j}$ and $s_k \in N^{\pi,-}_{s_j}$. Take all such agents. By symmetry adding links between these agents and s_{k-1} reduces the total reward, 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 4 s_{k-1} must be connected to s_k if $d^{\pi,+}_{s_k} < d^{\pi,-}_{s_k}$, where $d^{\pi,-}_{s_k}$ is also the number of sinks. Therefore instead suppose that $d^{\pi,+}_{s_{k-1}} = d^{\pi,+}_{s_k} \geq d^{\pi,-}_{s_{k-1}} = d^{\pi,-}_{s_k}$. Then

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

where the second inequality follows from $d^{\pi,+}_{s_k} \geq d^{\pi,-}_{s_k}$. 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 non-sink. 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 $(d_{s_k}^+ + 1)f(d_{s_k} + 1) - d_{s_k}^+ f(d_{s_h})$. We have:

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

where the first inequality follows because x connects to each non-sink, 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 3 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 non-sink s_j and s_{j+1} are not connected, and all non-sinks after j connect to the subsequent non-sinks. The argument above applies and hence the two non-sinks must be connected. □

Proof of Corollary

Proof. First, note that assumption SC is satisfied. Then by Proposition 2 an optimal network is a galaxy, and hence we need to maximize over $n-1$ possible galaxies. Given a galaxy with s stars, $1 \leq s \leq n$, the total social benefit extracted by the principal is

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

Substituting the expression for f, we get a quadratic

$$As^2 + Bs + n,$$

where $A = \frac{1}{2(n-1)} - \frac{\alpha}{2}$, $B = \alpha(n - \frac{1}{2}) - \frac{2n-1}{2(n-1)}$. If $\alpha < \frac{1}{n-1}$, then the function is concave and the maximum is achieved either when $s = 1$ or $s = n$. 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 convex and is maximized at $s^* = -\frac{B}{2A}$. Some algebra reveals that optimal value s^* is constant in α and equal to $n - \frac{1}{2}$, and hence the maximum in integer values is achieved when $s = n$ or $s = n-1$, both cases corresponding to a complete network. □

Proof of Proposition 3

The proof of the result is a straightforward consequence of Lemma 4, which we used in the proof of Proposition 2 above.

Proof of Proposition 3. Fix an optimal network G. Take a nonincreasing permutation π satisfying (11), and let $\pi^{-1}(k) = s_k$ for $k = 1, \ldots, n$. Let l_1 be the smallest k such that $\{s_k, \ldots, s_n\}$ is an independent set. Then using the arguments similar to those in the proof of Proposition 2, we find that each node outside of $\{s_l, \ldots, s_n\}$ must be connected to each node in $\{s_l, \ldots, s_n\}$. Let $C_1 = \{s_l, \ldots, s_n\}$. Let l_2 be the smallest k such that $\{s_k, \ldots, s_{l_1-1}\}$ is an independent set. Similarly, we find that each node outside of $\{s_{l_2}, \ldots, s_{l_1-1}\}$ must be connected to each node in $\{s_{l_2}, \ldots, s_{l_1-1}\}$. Let $C_2 = \{s_{l_2}, \ldots, s_{l_1-1}\}$. Continue the process to get a sequence $C_1, \ldots, C_{k+1}$. Note that C_{k+1} must be a clique because node s_1 must connect to all nodes. Hence, we have shown that G must be a k-cluster galaxy. □

Subgame dominant equilibrium

In this section, we introduce a sequential move version of the model. Consider a perfect information extensive form game with n stages. At every stage each agent i observes the actions of all the agents who moved at the previous stages, and decides in turn whether to take the action or refrain from doing so. The payoffs of the agents are given by (1), as before. Notice that for any sequence in which the agents decide to take the action, when $t_i = c_i - d_i f(d_i)$ for each i, there is a SPE where each agent acts. However, this equilibrium is fragile, as it requires agents to believe that everyone else will also act.[16] We can strengthen the equilibrium condition requiring that each agent choose a (weakly) dominant strategy. Specifically, a profile of strategies is a *subgame dominant equilibrium* (SDE) if a strategy of each player induces a weakly dominant strategy of this player in every subgame (Halac et al., 2020a). An *extensive form mechanism* consists of an order of moves of agents, and the rewards t. An extensive form mechanism is *optimal* if all agents

[16]Note also that there is an equivalent equilibrium in the simultaneous move game. However, this equilibrium is not unique.

choose to take an action in SDE of the induced extensive form game, and the sum of rewards is minimized. The following result effectively establishes an equivalence between the sequential and the simultaneous game approaches.

Proposition 6. *An extensive form mechanism is optimal if and only if the order of moves of agents is given by a nonincreasing permutation, and the rewards are determined by equation (2).*

Below we sketch the proof of the result as it is very similar to the one for Proposition 1.

Proof. Fix any order of moves π. It is clear that the rewards that induce SDE where each agent acts must be given by (2). Indeed, consider the last agent to move and assume that everyone else acts. Then this agent must be paid at least (2), and so on. But then for any order π the rewards are given by (9), as before. Hence, the same argument as that in the proof of Proposition 1 yields the result. □