Circles of Trust*

Samuel Lee Petra Persson

January 2010

Abstract

We study strategic transmission of rival information among individuals who share altruistic relations. Information sharing naturally takes the form of word-of-mouth communication. Yet, the social forces that promote communication can also undermine it. Indeed, larger social networks may share less information, and better-connected individuals may get less information. Therefore, when forming networks, individuals strategically cluster into groups that share information and trust by limiting other connections as a commitment to secrecy.

*S. Lee, NYU, slee@stern.nyu.edu; P. Persson, Columbia University, pmp2116@columbia.edu.

> "Of course I can keep secrets. It's the people I tell them to that can't keep them." – Anthony Haden Guest

1 Introduction

A focal point in the social network literature is the question of how knowledge, or information, diffuses. The conventional wisdom is that (more) social ties promote information sharing; information spreads farther when more individuals have ties, and better-connected individuals garner more information. This is simple to see in the case of information that is non-rival in use and mainly communicated through personal interaction, or cannot be concealed from "neighbors." It is somewhat less obvious in the case of information that is partly rival in use, privately concealable, but publicly communicable; though, it seems that social ties would not be harmful. This is the case analyzed in the present paper.

When those who are informed have incentives to withhold information, such as about scarce opportunities or personal secrets, social ties may foster information sharing due to non-pecuniary benefits that derive, for example, from sympathy. Indeed, we show that, even when such information is publicly communicable, it is naturally shared only through social ties, that is, "just between friends." This endogenous communication pattern accords with our sense that, in reality, certain types of information are only shared in private and by word of mouth. This seems to suggest that, also in this case, social ties are unequivocally conducive to information sharing.

On second thought, this is not true for the reason that social ties are not exclusive. One's friend may have other friends, friends of friends, to whom the friend, caught in a conflict of loyalties, may feel compelled to pass on the information. Since informing the friend may expose the information to social forces outside of one's control, one may withhold the information after all – albeit not for lack of sympathy but out of secrecy. Thus, while sympathy is a prerequisite for information sharing, it is not sufficient without trust that the shared information does not leak too far. Interestingly, sympathy derives from the (direct) friendship between two individuals, whereas trust depends on the overall structure of social relations in which their friendship is embedded.

For example, consider an unemployed person, 0, who knows about a company that searches for a couple of new employees. Even though sharing this information lowers each applicant's chances of being hired, 0 is inclined to tell an unemployed friend, 1, about this opportunity. However, whether 0 lets 1 in on this opportunity depends also on whether 1 would in turn leak the information to other job hunters (close to 1 but not to 0). Thus, it matters what social structure surrounds 0 and 1's friendship. (This motivating example is presented more formally in Section 2.)

There is irony in the logic that social ties can undermine (actual) information sharing precisely because they facilitate (counterfactual) information sharing. Indeed, the implications of this logic turn the conventional wisdom on its head; more social ties can lead to less information sharing, and better-connected individuals may receive less information. Furthermore, the fundamental need for trust imposes endogenous constraints on network formation, which are distinct from exogenous costs of forming social ties. We show that individuals cluster around information providers in *circles of trust*, shying away from forming ties outside of their circle. Such self-imposed limits to "networking" are a means to build trust by avoiding divided loyalties. In sum, agents form some ties to *access* information, but refrain from forming others to *commit* to secrecy; structural "holes" in the network are therefore, in part, a commitment device.

The above conceptual insights are the principal contribution of this paper. We derive these insights in a theoretical framework that has the following features: First, social ties are modeled as altruistic relations (as opposed to pure access). Second, information is partly rival in use, which creates private disincentives to share information. Third, communication is a strategic decision made sequentially by individuals upon receiving information. All three features represent significant departures from classic models of information diffusion in networks (see Chapters 7 and 8 in Jackson, 2009; Chapters 5 and 6 in Goyal, 2009).

Our paper is akin to a spate of new work on strategic network communication. Galeotti, Ghiglino, and Squintani (2009) analyze *cheap talk* between neighbors in a network; individuals send soft messages to influence their neighbors' actions. They show that whether individuals truthfully communicate to a neighbor depends on the neighbor's other communication relationships. Also, truthful communication occurs primarily among individuals with similar characteristics, thus leading to homogeneous communication circles. Lippert and Spagnolo (2010) examine how communication of soft information can support cooperative network relationships by facilitating sanctions aimed at deterring opportunistic behavior. Stein (2008) studies soft information exchange among competitors. He shows that communication can be truthful because a dynamic exchange of complementary ideas can be mutually beneficial. Moreover, vague ideas can travel far, whereas concrete ideas remain localized. His model, like ours, considers *sequential* communication.

The endogenous network structures in our paper are similar to those derived by Galeotti and Goyal (2010), who study strategic network formation by individuals that can either produce costly information or receive information by forming costly links to informed individuals. Their analysis relies on non-rival information (which makes communication non-strategic), and network formations are driven by cost-saving and free-riding considerations. In our analysis, network structures manifest a trade-off between access and secrecy, which arises from strategic communication issues. In fact, we derive these structures in settings without any exogenous costs (of information acquisition or link formation).

Finally, it is instructive to compare the nature of information in our model to the notion of public goods in networks (Bramoulle and Kranton, 2007). Public goods in networks cannot be excluded along social links; that is, they spill over to neighbors. By contrast, we consider information that is excludable. Yet, while an individual can exclude others, it cannot perfectly select whom to exclude; once information is released to some, it may also reach those whom the individual would like to have excluded. In this sense, exclusion is imperfect.

The paper proceeds as follows. Section 2 highlights the core insight through a simple example, which is then generalized in Section 3. Section 4 explores the implications of our results for network formation and information flow. Section 5 discusses a possible link between intellectual property rights and social network formation. Section 6 presents our concluding remarks, and mathematical proofs are in the Appendix.

2 Motivating example

Consider an individual who has a piece of information (henceforth agent 0). Anyone that owns the information enjoys private value $\pi(n) = 1/(5+n)$, where n is the number of agents who own the information. Note that $\pi'(n) < 0$, that is, sharing the information dilutes its value. Suppose there is one other individual, agent 1. The two agents are friends in the sense that they each internalize a share $\phi = 1/5$ of each other's payoff. Agent 0 shares its information with agent 1 if and only if $(1+\phi)\,\pi(2) \geq \pi(1)$, or

$$\phi\pi(2) \geq \pi(1) - \pi(2). \tag{1}$$

The inequality reflects agent 0's costs and benefits from sharing the information. Agent 0 internalizes part of agent 1's gains from the information (left-hand side), but conversely relinquishes some of its own gains (right-hand side). In our example, the inequality becomes $1/35 \geq 1/42$. So, the information is shared.

Let us add another individual, agent 2, who is a friend of agent 1, but only indirectly a friend of agent 0 (a friend of a friend). Assume that indirect friends internalize only a share $\phi^2 = 1/25$ of each other's payoff. Given the information, agent 1 would share it with agent 2 if $(1+2\phi)\,\pi(3) \geq (1+\phi)\,\pi(2)$, or

$$\phi\pi(3) \geq (1+\phi)\left[\pi(2) - \pi(3)\right]. \tag{2}$$

Similar to above, on one hand, agent 1 internalizes some of agent 2's gain when passing on the information; on the other hand, it suffers from the loss of private information value both by itself and by its already informed friend, agent 0. In our example, the inequality becomes $1/40 \geq 3/140$. Thus, agent 1 would pass on the information. By backward induction, agent 0 anticipates that the information, if communicated to agent 1, would travel to agent 2. Hence, agent 0 shares the information only if $\left(1+\phi+\phi^2\right)\pi(3) \geq \pi(1)$, or

$$\left(\phi+\phi^2\right)\pi(3) \geq \pi(1)-\pi(3). \tag{3}$$

In our example, this becomes $3/100 \geq 1/24$, which is not true. Thus, the information is not shared at all. Interestingly, we know that agent 0 would want agent 1 to have the information. The reason it does not share the information with agent 1 is that agent 1 would feel compelled to give it to agent 2.

Now let us add one more individual, agent 3, who is a friend of agent 2, but only indirectly a friend of agent 1 (a friend of a friend) and agent 0 (a friend of a friend of a friend). Again, we use backward induction: Agent 2 would transmit the information to agent 3 if $\phi\pi(4) \geq [\pi(3)-\pi(4)]\left(1+\phi+\phi^2\right)$. This condition holds in our example ($0.0\bar{2} \geq 0.017\bar{2}$). Anticipating this, agent 1 would transmit the information to agent 2 if $\left(\phi+\phi^2\right)\pi(4) \geq [\pi(2)-\pi(4)](1+\phi)$. In our example, this condition is violated ($2/75 \geq 4/105$)., So agent 1 would withhold the information. Anticipating this, agent 0 gives the information to agent 1. That is, agent 0 *confides in* agent 0. The next section generalizes this example.

3 Word of mouth, secrecy, and trust

Consider a *chain* of N friends, indexed $i \in \{0, 1, \ldots\}$. Agent 0 is endowed with a piece of hard information that it can share. Let I denote the set of agents that have the information. An individual agent's payoff from owning the information, $\pi(n)$, depends on the total number of agents that own it, $n = |I|$.

Assumption 1 (*A1*) $\pi(\cdot)$ *satisfies the following properties:*

- $\pi'(\cdot) \leq 0$
- $\lim_{n\to\infty} \pi(n) = 0$
- $\frac{\partial}{\partial n}[n\pi(n)] \geq 0$ *for all* n*.*

In words, an agent's payoff from owning the information decreases, and vanishes in the limit, as the number of informed agents grows. The third property ensures that the collective payoff of all informed agents increases with the number of informed agents. It is not necessary for our results, but implies that sharing the information (with as many agents as possible) is socially optimal.

We model the strength of friendship between two arbitrary agents, i and i', in reduced form as a function of the distance between them, $\phi(|i' - i|)$.

Assumption 2 (*A2*) $\phi(\cdot)$ *satisfies the following properties:*

- $\phi'(\cdot) < 0$
- $\lim_{n\to\infty} \phi(n) = 0$
- $\sum_{k=1}^{\infty} \phi(k) < \infty$*.*

In words, the friendship between two agents becomes weaker, and vanishes in the limit, as the distance between them grows. The third property says that an agent never cares *infinitely* much about all of its friends.[1]

[1] It is straightforward to verify that both Assumptions 1 and 2 are satisfied in the motivating example, where $\phi(k) = \phi^k$ and $\pi(n) = 1/(5 + n)$.

The utility of agent i can be written as

$$\begin{aligned} u_i &= \mathbb{I}_{i\in I}\pi(n) + \sum_{k=1}^{i} \mathbb{I}_{i-k\in I}\phi(k)\pi(n) + \sum_{l=1}^{N-i-1} \mathbb{I}_{i+l\in I}\phi(l)\pi(n) \\ &= \pi(n)\left[\mathbb{I}_{i\in I} + \sum_{k=1}^{i} \mathbb{I}_{i-k\in I}\phi(k) + \sum_{l=1}^{N-i-1} \mathbb{I}_{i+l\in I}\phi(l)\right]. \end{aligned} \tag{4}$$

For use below, we define

$$\Phi\left(i,N\right) \equiv \sum_{k=1}^{i} \mathbb{I}_{i-k\in I}\phi(k) + \sum_{l=1}^{N-i} \mathbb{I}_{i+l\in I}\phi(l). \tag{5}$$

so that

$$u_i = \pi(n) + \pi(n)\Phi(i,N) = \pi(n)\left[1 + \Phi(i,N)\right]. \tag{6}$$

Recall that $\pi(n)$ denotes the information payoff to an individual informed agent. $\Phi(i,N)$ reflects to what extent agent i derives additional utility from the fact that (some of) its *friends* are also informed and hence receive this payoff. We refer to $\Phi(i,N)$ as i's "community" factor.

Clearly, no information would be transmitted in the absence of social ties $[\phi(\cdot) = 0]$ because the private payoff from owning the information decreases with the number of informed agents $[\pi'(\cdot) < 0]$. The role of social ties in our model *is* to facilitate information transmission. The following assumption implies that agents want, all else equal, their *closest* friend to have the information – regardless of who else is already informed.

Assumption 3 (*A*3) $\frac{\phi(1)}{1+\Phi(i,n)} \geq \frac{\pi(n)-\pi(n+1)}{\pi(n+1)}$ *for all* i *and* n*.*

This assumption accentuates the “paradox” that we want to illustrate: Social ties provides strong incentives to share information but can, by the same token, become an impediment to information sharing.

In what follows, we analyze how far the information travels through the chain of friends. Communication choices are sequential and strategic: An agent chooses whether to share information only upon receiving it, while considering the choices of those who, as a result, would become informed. A subtle issue is that agents (must) form conjectures about which agents are already informed. Hence, we look for Perfect Bayesian Equilibria (henceforth, equilibria).

Pattern of communication Our first step is to show that we can restrict attention to communication between *direct* friends without loss of generality. We start by considering whom among the uninformed an informed agent most wants to share the information with.

Lemma 1 *Any agent prefers to transmit information to a closer friend.*

The proof of this result is as straightforward as the intuition behind it. Since gains from transmitting information arise from sympathy towards the receivers, agents prefer to share information with their closest friends. A direct implication of Lemma 1 is that the members of I always forms an uninterrupted chain, with no uninformed agents in-between. That is, $I = \{0, 1, \ldots, n-1\}$. This allows us to rewrite the community factor in (5), with a slight abuse of notation, as

$$\Phi(i, n) = \sum_{k=1}^{i} \phi(k) + \sum_{l=1}^{n-i-1} \phi(l). \tag{7}$$

It also pins down how the community factor varies across informed agents.

Lemma 2 $\Phi(n-1,n) \leq \Phi(i,n)$ *for all* $i < n-1$.

That is, among the informed agents, it is the agent furthest away from agent 0 – agent $n-1$ – who has the smallest community factor. As such, agent $n-1$ is the informed agent who least internalizes the loss incurred by the already *informed* agents when the information is transmitted further. We can now determine from whom an uninformed agent is most likely to receive information.

Lemma 3 *The incentive to transmit information to some uninformed agent* i *is strongest for its closest informed friend.*

This implies that, in order to determine whether an uninformed agent becomes informed, we only need consider the incentives of agent $n-1$. Relative to the other informed agents, agent $n-1$ has the least to lose from diluting the information value (Lemma 2) but, being closest to the uninformed agents, the most to gain. Taken together, Lemmas 1 and 3 imply the following.

Proposition 1 (Word of mouth) *Under* ***A1****-****A2****, no equilibrium requires (direct) communication between indirect friends.*

Here, communication *endogenously* occurs along direct ties. On one hand, it is a dominant choice for all agents to inform, first of all, their closest friend. On the other hand, the closest friend is most eager to inform an agent. This implies that communication between indirect friends is either unwanted or redundant and, moreover, that any agent i – when choosing whether to pass on the information – believes all and only agents $i' < i$ to be (already) informed. As a result, we can restrict attention to equilibria in which information flows only through direct ties.

Extent of communication Suppose the information has reached agent i. Agent i might like to share the information with some uninformed friends, even though that dilutes its payoff from the information. At the same time, i will not want to share it with too many agents, since its payoff from the information vanishes in the limit as more agents become informed.

To derive i's preferences over who else should be informed, consider its utility when the number of informed agents reaches $n > i$: $u_i(n) = \pi(n)\,[1 + \Phi(i, n)]$. From $\lim_{n\to\infty} \pi(n) = 0$ (***A1***) and $\lim_{n\to\infty} [1 + \Phi(i, n)] < \infty$ (***A2***), it follows that $\lim_{n\to\infty} u_i(n) = 0$. This in turn implies the next result.

Lemma 4 *For every agent i, there exists a unique finite $N_i^* \geq i$ such that*

$$u_i(n) < u_i(i) \quad \text{for all } n > N_i^*. \tag{8}$$

In words, N_i^* denotes the maximum number of informed agents that i prefers to be included in I over withholding the information, that is, being the last informed agent. This is merely to say that agents' willingness to share information is limited. For the next proposition, a weaker condition would suffice; all we need is that some agent is unwilling to freely share information, lest the equilibrium be trivially that everyone becomes informed.

We are interested in how the number of informed agents, n, evolves as a function of the length of the chain, N. Let $n(N)$ denote this function. Our next result shows that $n(N)$ is non-monotonic.

Proposition 2 (Secrecy) *Under **A1**-**A3**, there exists a unique finite $N^* > 2$ such that $n(N) = N$ for all $N < N^*$ but $n(N^*) < N^* - 1$.*

Initially, information is always transmitted through the whole chain, and growth of the chain increases the number of informed agents one-to-one. However, when the chain reaches a certain length, the number of informed agents suddenly decreases; the information not only ceases to travel *further* but it travels *less far* than before.

The intuition is the same as in the example of Section 2. Although all agents $i \in \{0, \ldots, N^* - 3\}$ would be willing to share the information with agent $N^* - 2$, some of them do not want to share the information with agent $N^* - 1$. They draw a line there because the friendship they feel for $N^* - 1$ is not strong enough to compensate for the further erosion in information value. However, the information is not even transmitted to agent $N^* - 2$, because the other agents know that $N^* - 2$, who is closest to $N^* - 1$, feels differently, that is, would feel compelled to pass on the information. In fact, this implies that the information is not shared with those who would pass it on to $N^* - 2$, nor with those who would pass it on to those who would pass it on to $N^* - 2$, and so forth. In equilibrium, the information travels only to those who can contain its spread, or put differently, keep it *secret*.

Proposition 2 tells us that an increase in the chain of friends above some size $N^* - 1$ can erode the incentives to transmit the information. The next result shows that a further increase in the chain of friends can partly restore these incentives, though it does not make the incentives stronger than for $N = N^* - 1$.

Proposition 3 (Trust) *Under **A1**-**A3**, $n(N)$ oscillates between $n(N^*)$ and $N^* - 1$ for $N \geq N^*$.*

As N increases over and above N^*, the incentives to transmit information are partly restored. The reason is that, as the chain of friends becomes even larger, more of the agents—especially those who would have passed on the information previously—-become sufficiently concerned about secrecy to withhold the information. Hence, from the perspective of those who previously withheld the information, these agents become trustworthy again. As a result, they are (re-)admitted to the circle of trust.

But even as N becomes very large, the size of the circle of trust never exceeds $N^* - 1$; it is intuitive that the group of confidants must be bounded when secrecy is the key concern. It is noteworthy that the variation in $n(N)$ for $N \geq N^*$ is driven by changes in the *trustworthiness* of various agents rather than changes in their original *preferences*. Note also that, for the purposes of information transmission, increasing the chain length beyond $N = N^* - 1$ is futile.

Information acquisition So far, we have assumed that agent 0 is endowed with information. For completeness, assume instead that agent 0 must acquire information at some fixed cost C. We are concerned with the situation in which agent 0 is not willing to incur this cost *unless* its can share the benefits with *part* of the community, namely its closest friends. (Otherwise, the information is always or never acquired.)

Proposition 4 (Innovation) *Suppose that* $N_0^* \leq N_i^*$ *for all* $i > 0$ *and* $u_0(1) < C \leq u_0(n)$ *for some* n. *Under* **A1**-**A3**, *there exist a non-empty interval* $\left[\underline{N}, \overline{N}\right] \subset (1, N^*)$ *such that the information is acquired only if* $n(N) \in \left[\underline{N}, \overline{N}\right]$.

To sustain information acquisition, the degree of information sharing must strike a delicate balance. When information reaches but a few, the "community benefits" may be too small to motivate agent 0 to acquire information. At the same time, when information travels far, information acquisition may be undermined because agent 0 cannot control precisely to whom the information is transmitted. The anticipation that the information either leaks out too far or cannot be shared out of secrecy discourages agent 0 from acquiring the information. The conclusion remains that some, but not too many, social ties are desirable.

4 Network formation and information flow

The above results have a salient implication for strategic network formation: In order to garner information, individuals (must) build a personal network so as to draw *sympathy* and *trust* from those who are in the position to share information. Abstracting from particulars, it is straightforward to see that this implication is generally robust whenever informed individuals are willing to share information only with friends but, even then, only a few of them.

A fully-fledged analysis of network formation that allows for general $\pi(\cdot)$ and $\phi(\cdot)$ is prohibitively cumbersome because there are too many parametric cases to consider. For instance, one fundamental layer of complexity would stem from the following question (and all of its permutations): How many friends with distance k would an individual want to confide information to when $n_{k'}$ friends with distance $k' \in \mathbb{N} \setminus \{k\}$ are already informed? Yet, such complications are orthogonal to the basic insight that an individual receives information only through social ties and only if it can be trusted not to spread the information too much farther, which in turn can limit the number of social ties that the individual wishes to establish.

To explore this idea more simply, we digress from the (details of the) model in Section 3, but instead incorporate its key implications into the following parsimonious network formation game. Consider N, initially unconnected, individuals, all of whom want information. Only a few of them, $N_o \ll N$, are *originators*, one of whom nature randomly chooses to endow with information; the other $N_\iota = N - N_o$ individuals are *imitators*, who obtain information only through communication with others.

Nature confers either of two types of information on the selected originator. With some probability, it is information that any individual is willing to share with any friend, irrespective of how far it travels; otherwise, it is partly rival information that an individual is not willing to share freely. We distill the spirit of the preceding analysis into a simple but effective assumption: The selected originator is inclined to share rival information with *at most* τ (direct or indirect) friends, where $\tau \ll N_\iota/N_o$, whereas everyone else, upon receiving the information, wants to *pass it on* to at least one uninformed friend. Furthermore, we assume that individuals randomize between equally attractive information recipients, and that obtaining the partly rival information is everyone's overriding aim.

Networks form as follows. All individuals invite others, as many as they want, into a relationship, and invited individuals can accept or decline invitations. Forming a relationship imposes no (exogenous) costs. Thus, relationships form simultaneously and merely require bilateral consensus. We focus on Nash equilibrium (henceforth, equilibrium). In describing equilibrium outcomes, we restrict attention to social ties that serve the purpose of information sharing, thereby abstracting from such that are solely formed to internalize others' payoffs without any material benefit.[2] We call such outcomes, accordingly, *information networks*.

[2]This restriction merely serves to economize on the exposition of the analysis, and to focus on social ties that are avoided because they would undermine information transmission.

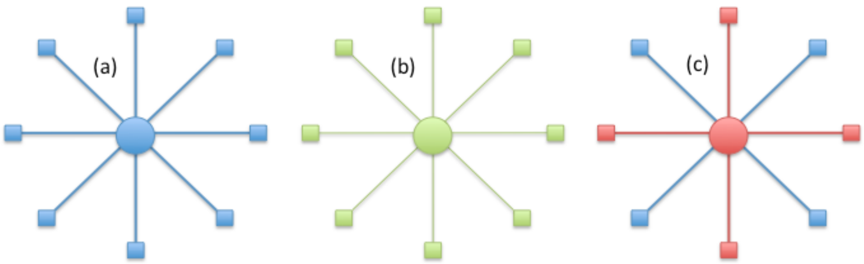

Figure 1: Single-originator. The figure depicts the unique equilibrium information network for $N_o = 1$, $N_\iota = 8$, and $\tau = 4$. Circles represent originators, while squares represent imitators. Uninformed individuals are blue, those with non-rival information green, and those with partly rival information red. Part (a) depicts the network, part (b) how non-rival information is shared, and part (c) how partly rival information is shared.

Single originator Let us start with the simplest case, namely that of a *single* originator: $N_o = 1$.

Proposition 5 (Star network) *The* ***unique*** *information network with a single originator is a star: Each imitator is connected with the originator. The network is pairwise stable and constrained efficient.*

It is noteworthy that, despite the absence of explicit formation costs, the star is the *only* equilibrium outcome. This manifests a simple but fundamental point: The precedent objective of all imitators is to get *access* to partly rival information. Since the originator prefers sharing information with closer friends (cf. Lemma 1), it is critical to be as *close* as possible to originators. Consequently, every imitator establishes a direct relationship with the originator.

In equilibrium, non-rival information reaches everyone, whereas partly rival information is shared only with τ randomly chosen imitators (see Figure 1). This is efficient in the sense that no other network structure leads to more information

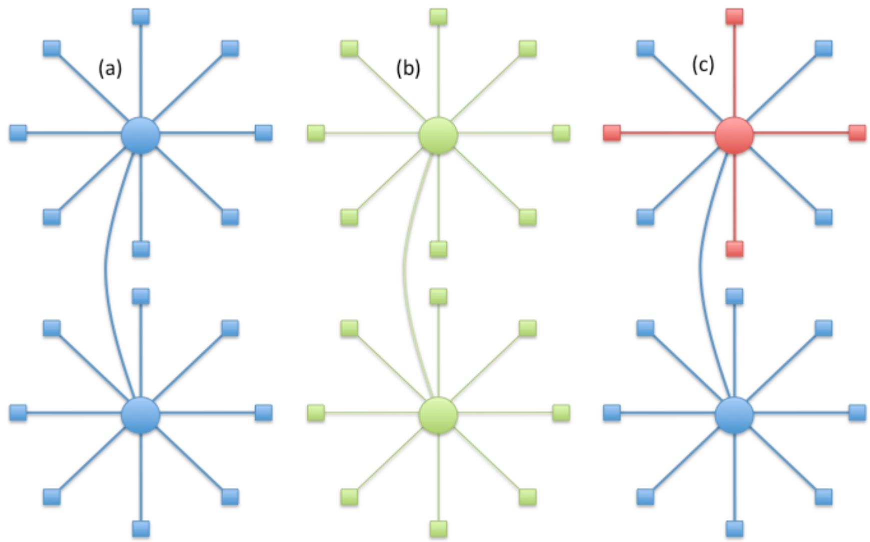

Figure 2: Star meeting. The figure depicts a network of linked stars for $N_o = 2$, $N_\iota = 16$, and $\tau = 4$. Circles represent originators, while squares represent imitators. The top circle represents the originator endowed with information. Uninformed individuals are blue, those with non-rival information green, and those with partly rival information red. Part (a) depicts the network, part (b) how non-rival information is shared, and part (c) how partly rival information is shared.

sharing. Furthermore, no individual obtains an informational benefit from adding or severing a relationship, which is to say that the star is pairwise stable.

Star meeting Before analyzing the network formation game with multiple originators, we consider what happens when separately formed star networks "meet" each other. The central question is whether they become connected and, if so, in what manner. More precisely, we consider unilateral relationship changes to study whether a network of disjoint stars is pairwise stable and, if not, which additional relationships emerge to connect them.

Proposition 6 (Circles of trust) *Separately formed star networks connect only through originators, forming a pairwise stable and efficient network of linked stars:*

Each imitator is connected to one originator, while originators form a line.

This result emphasizes the tension between access and trust. All else equal, imitators want to access the potential information flow in other stars. However, they are constrained by the fact that, by establishing other relationships, they forfeit the trust of their current originators. In fact, imitators connected to several originators fall between stools. Thus, they refrain from entering into new alliances, not because of exogenous costs but to safeguard their trustworthiness; the lack of networking is a commitment device.

Originators have no such concerns; if anything, they gain better access to non-rival information by connecting to other stars. In the resulting network of linked stars, non-rival information reaches everyone, through originators acting as hubs, whereas partly rival information is shared only within stars (Figure 2). This is the sense in which each star embodies a *circle of trust*.

Many originators A peculiar feature of the information flow in Figure 2 is that the originators never receive partly rival information from each other. This is because they are too well-connected; being at the center of a star actually harms them. This suggests that they prefer a different network structure, which requires more radical changes since the above network is pairwise stable. Indeed, a different information network emerges when individuals form social ties in the presence of several originators.

Proposition 7 (Inner circles) *In information networks with $N_o < \tau$, all originators are completely connected, and some imitators are excluded. In information networks with $N_o > \tau$, there are disjoint small worlds of at most τ completely con-*

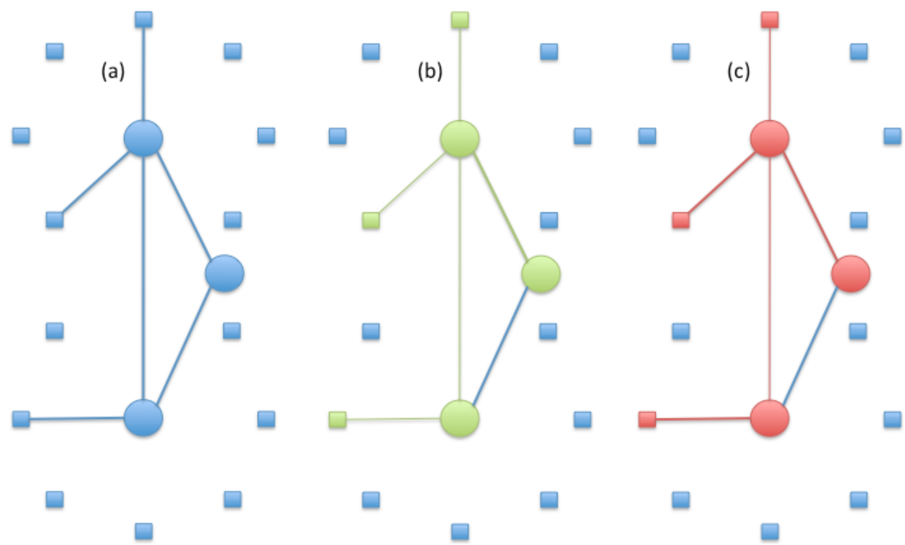

Figure 3: Inner circle. The figure depicts completely connected originators for $N_o = 3$, $N_\iota = 14$, and $\tau = 6$. Circles represent originators, while squares represent imitators. The top circle represents the originator endowed with information. Uninformed individuals are blue, those with non-rival information green, and those with partly rival information red. Part (a) depicts the network, part (b) how non-rival information is shared, and part (c) how partly rival information is shared.

nected originators, and some imitators are excluded. In either case, the networks are pairwise stable and contrained efficient.

Originators first and foremost establish relationships with each other in order to ensure mutual exchange of partly rival information. Yet, in order to sustain trust, these complete subnetworks cannot comprise more than τ originators. For the same reason, originators may decline invitations from imitators, some of whom eventually remain excluded from the information flow. Intuitively, the originators form circles of trust among themselves. As a result, information circulates exclusively in close-knit "small worlds" of well-informed agents (Figure 3). Again, the lack of network structure is strategic and serves to uphold secrecy.

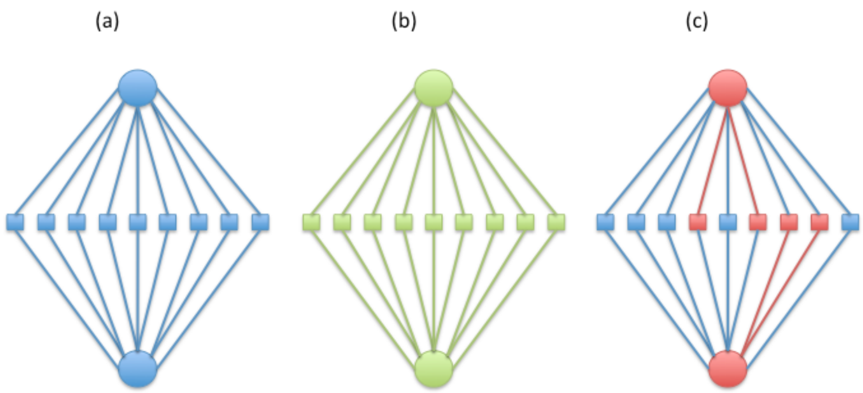

Figure 4: Endogenous origination. The figure depicts two interlinked stars $N = 11$ and $\tau = 5$. Circles represent originators, while squares represent imitators. Uninformed individuals are blue, those with non-rival information green, and those with partly rival information red. Part (a) depicts the network, part (b) how non-rival information is shared, and part (c) how partly rival information is shared.

Endogenous origination Last but not least, we consider a different situation in which, as in Galeotti and Goyal (2010), all individuals are ex ante identical and information is costly to produce: Network formation proceeds as before, but any individual can now produce the partly rival information by incurring a cost C. Crucially, in the spirit of Proposition 4, we assume that it is worthwhile to produce information only if it can be shared with at least one but at most τ friends.

Proposition 8 (Interlinked stars) *Every pairwise stable information network with endogenous origination includes at least one but at most $\tau/2 \ll N$ originators, and comprises interlinked stars: Every imitator is connected to all and only originators, while originators are not connected to each other.*

In this setting, each imitator connects to all originators. As before, the intention is to improve access to information; tapping more originators increases the chances of becoming informed, since originators randomize among equally attractive imitators. But unlike before, connections to multiple originators do no longer erode trust because all originators are informed. By the same token, originators need not connect to each other. Despite access to multiple originators, an imitator does not always become informed (Figure 4). In fact, an imitator's probability of receiving partly rival information is bounded by τ/N and vanishes for $N \to \infty$.

Due to the small chance of receiving information, individuals sometimes resort to producing information such that there can be more than one originator in equilibrium. Still, the number of originators is bounded by $\tau/2$ in equilibrium because, given that information is ultimately (expected to be) owned by τ individuals, it cannot be worthwhile for everyone to produce the information. Thus, the reason that not all choose to be originators is not that they expect to eventually get the information via communication. Rather, the dilution in information value renders *autarkic* information acquisition by everyone unprofitable.

Apart from resorting to origination, an imitator can, at least in principle, increase its chances of becoming informed by forming connections to other imitators. That way, if the originator transmits information to the latter, the imitator can tap the information flow indirectly. However, in equilibrium, such strategies must prove futile. Such information circulation among imitators undermines trust. If information were freely shared among imitators, originators would not share their information; in fact, they would not find it worthwhile to produce information.

To summarize, Propositions 5 to 8 derive information networks that, though distinct, all exhibit *core-periphery* structures. Intuitively, these structures emerge from a trade-off between gaining access to information, on one hand, and ensuring trust and secrecy, on the other hand. This basic insight appears robust.

5 Property rights and social networks

Preferences over how far information should travel depend on how much the private payoffs from owning the information are diluted as the information is shared with more agents. It is hence natural to ask whether the dilution can be mitigated for already informed agents who must be motivated to share their information.

By ***A1***, the total payoff from information, $n\pi(n)$, is strictly increasing. One possibility is therefore to redistribute payoffs among informed agents, for example, through monetary transfers. One could compensate existing informed agents whenever the information is passed on to the next uninformed agent. In particular, it is possible to keep every informed agent's net payoff constant, that is, just as high as when the agent first became informed. Consider the following transfers:

- $n = 1$: Agent 0's payoff is $\pi(1)$.
- $n = 2$: The total payoff is $2\pi(2)$, of which agent 0 gets $\pi(1)$ and agent 1 gets $2\pi(2) - \pi(1) > 0$.
- $n = 3$: The total payoff is $3\pi(3)$, of which agent 0 gets $\pi(1)$, agent 1 gets $2\pi(2) - \pi(1)$, and agent 2 gets $3\pi(3) - [2\pi(2) - \pi(1) + \pi(1)] = 3\pi(3) - 2\pi(2) > 0$.
- $n = 4$: The total payoff is $4\pi(4)$, of which agent 0 gets $\pi(1)$, agent 1 gets $2\pi(2) - \pi(1)$, agent 2 gets $3\pi(3) - 2\pi(2)$, and agent 3 gets $4\pi(4) - 3\pi(3) > 0$.
- Etc.

Given such transfers, no informed agent ever objects to transmitting the information to more uninformed agents. In fact, each agent's utility strictly increases as more agents become informed. This becomes clear when we look at $\partial u_i(n)/\partial n$ in the case with transfers. Since the private *cum-transfer* payoff of all already informed agent remains constant at $\pi_i(i)$, we get $\partial u_i(n)/\partial n = \Phi_n\,(i,n)\,\pi_n(n) > 0$. We summarize this instance of the Coase Theorem as follows.

Proposition 9 *Suppose each newly informed agent i pays fees to compensate existing informed agents for the decrease in their payoff, $\pi(i+1)-\pi(i)$. Then the information is always transmitted to all N agents, and larger N strictly improve incentives to acquire information.*

Intellectual property rights are one interpretation of the above transfers. It is not surprising that such rights increase information diffusion. Yet, Proposition 9 implies more than that: Better protection of intellectual property changes the impact of social ties on information transmission. Rather than being a threat to secrecy, large(r) social networks become a channel for word-of-mouth "marketing," thereby promoting information sharing and acquisition. Empirically, this suggests that better *individual* property rights over intellectual assets lead to the formation of larger networks, which in turn improve the *collective* use of information.

6 Concluding remarks

We have studied communication of partly rival information among individuals who are connected through altruistic ties. We show that larger social networks can reduce information sharing, which in turn imposes endogenous constraints on network formation. When forming social ties, individuals cluster into *circles of trust*, on one hand, to access information and, on the other hand, to ensure secrecy. Limits to personal networking, and thus the "missing links" in a network, are commitment devices to create trust.

A central departure from related work on communication in networks is our view of social ties not as merely providing access or utility but as relationships that involve *incentive spillovers*. The impact of a social network on information diffusion is *endogenous* and operates through its impact on incentives. We believe this approach to social networks harbors promising avenues for future research.

Appendix

Proof of Lemma 1

Consider agent i. Suppose i has the information and considers whether to pass it on to either agent i' or agent i'', where agent i' is s' steps away from agent i and agent i'' is $s'' > s'$ steps away from agent i. Suppose that n agents, including i but excluding i' and i'', have the information. If i transmits the information only to i', her utility is

$$u_i = \pi(n+1) + \phi(s')\pi(n+1) + \hat{\Phi}(i, N)\pi(n+1) \tag{9}$$

where $\hat{\Phi}(i, N)$ is i's community factor vis-a-vis every other agent except for i' and i''. Conversely, if i transmits the information only to i'', her utility is

$$u_i = \pi(n+1) + \phi(s'')\pi(n+1) + \hat{\Phi}(i, N)\pi(n+1). \tag{10}$$

Since $\phi(s') > \phi(s'')$, the expression in (9) is larger than the expression in (10). Hence, if an informed agent can transmit the information to only one more agent, she informs the closest uninformed friend (or none). (Agent i may transmit information to both agents $i+1$ and $i+2$. The above implies that, if she chooses to transmit information to $i+2$, she surely transmits information also to $i+1$.)

Proof of Lemma 2

Using (7), we can write

$$\begin{aligned}\Phi(n-1,n) &= \sum_{k=1}^{n-1}\phi(k)+\sum_{l=1}^{0}\phi(l)\\ &= \sum_{k=1}^{n-1}\phi(k)\\ &= \phi(1)+\phi(2)+\cdots+\phi(n-1). \qquad (11)\end{aligned}$$

and

$$\begin{aligned}\Phi(i,n) &= \sum_{k=1}^{i}\phi(k)+\sum_{l=1}^{n-i-1}\phi(l)\\ &= [\phi(1)+\phi(2)+\cdots+\phi(i)]+[\phi(1)+\phi(2)+\cdots+\phi(n-1-i)] \qquad (12)\end{aligned}$$

Note that

$$\Phi(i,n) - \Phi(n-1,n) =$$
$$[\phi(1) + \phi(2) + \cdots + \phi(i)] + [\phi(1) + \phi(2) + \cdots + \phi(n-1-i)] -$$
$$\phi(1) + \phi(2) + \cdots + \phi(n-1) =$$
$$[\phi(1) - \phi(i+1)] + [\phi(2) - \phi(i+2)] + \cdots + [\phi(n-1-i) - \phi(n-1)] \, .$$

Since $\phi'(\cdot) < 0$, the differences in the last expression are all positive for $0 \leq i < n-1$, which implies that $\Phi(i,n) \geq \Phi(n-1,n)$.

Proof of Lemma 3

Consider an uninformed agent i and suppose there are n informed agents. By Lemma 1, we know that $I = \{0, 1, \ldots, n-1\}$. This implies that $i > n-1$ and, if i' denotes i's closest informed friend, that $i' = n-1$. Suppose that i' is s' steps away from i. Also, consider another informed agent i'', who is $s'' > s'$ steps away from i.

Suppose that only i' can inform i. If i' transmit the information, it utility is

$$u_{i'} = \pi(n+1) + \phi(s')\pi(n+1) + \Phi(i,n)\pi(n+1) \tag{13}$$

where $\Phi(i,n)$ is i''s community factor vis-à-vis the other informed agents, $I \setminus \{i, i'\}$. If i' does not transmit the information, its utility is

$$u_{i'} = \pi(n) + \Phi(i,n)\pi(n).$$

Thus, it transmits the information to i (assuming that no one else does) if and only if

$$\pi(n+1)\phi(s') \geq [1+\Phi(i',n)]\,[\pi(n)-\pi(n+1)]\,. \tag{14}$$

Now suppose instead that only i'' can inform i. Using the same steps as above, it can be shown that i'' transmits the information to i (assuming that no one else does) if and only if

$$\pi(n+1)\phi(s'') \geq [1+\Phi(i'',n)]\,[\pi(n)-\pi(n+1)] \tag{15}$$

where $\Phi(i'',n)$ is i'''s community factor vis-à-vis the other informed agents, $I\backslash\{i,i''\}$.

We make the following observations. First, in both preceding inequalities, the right-hand and left-hand sides are positive. Second, the left-hand side of (14) is greater than the left-hand side of (15), as $\phi(s') > \phi(s'')$. Third, the right-hand side of (14) is smaller than the right-hand side of (15), as $\Phi(i',n) < \Phi(i'',n)$ by Lemma 2. Thus, (15) implies (14) but not vice versa. In other words, i' has a stronger incentive than i'' to transmit information to i.

Proof of Proposition 2

For every N, consider the set $\{N_i^*\}_{i\in\{0,\ldots,N-1\}}$. We know that $N_{N-1}^* > N$; agent $N-2$ wants to inform agent $N-1$ (by ***A3***). But this need not be true for the other informed agents, $i < N-2$. As long as $\min\{N_i^*\}_{i\in\{0,\ldots,N-3\}} \geq N$, information is transmitted to everyone in the chain, that is, $n(N) = N$. However, it follows from Lemma 4 that eventually $\min\{N_i^*\}_{i\in\{0,\ldots,N-3\}} < N$, as N grows further. Let N^* denote the smallest chain length at which this happens. Also, define $M(N^*) \equiv \{i \in \{0, \ldots, N^* - 3\} : N_i^* < N^*\}$, the set of agents who would have information transmission rather stop with themselves than continue to N^*. Then, at N^*, information only travels up to $\max M(N^*)$. If the information were transmitted to *any*one else, the information would be transmitted to everyone in the chain. Thus, $n(N^*) = \max M(N^*) + 1 \leq N^* - 3 + 1 < N^* - 1$.

Proof of Proposition 3

We continue from the proof of Proposition 2. For $N = N^*$, the set of agents who receive the information is a sequence $\{0, 1, \ldots, N'(N^*)\}$ where $N'(N) \equiv \max M(N)$. (The fact that it is an uninterrupted sequence follows from Proposition 1, or more precisely, Lemmas 1 and 3.) Accordingly, we have $n(N^*) = 1 + N'(N^*)$. Now consider $N \geq N^*$. From Lemma 4, it follows that the set of agents who do not want the information to travel to the end of the chain (weakly) increases in N. This implies that $N'(N)$ (weakly) increases with N. Since information is transmitted to at most $N'(N)$, we can think of $N'(N)$ as a *virtual* chain length (as opposed to the actual chain length N). We can now repeat the arguments from the proof of Proposition 2, though using the virtual length $N'(N)$ instead of the actual length N. From this, it immediately follows that, $n(N) = N'(N)$ for all N such that $N'(N) < N^*$. When $N'(N) = N^*$, or $N = N'^{-1}(N^*)$, then $n(N) = 1 + N'(N^*)$. At this point, we define a new virtual length $N''(N'(N))$ and reiterate the arguments, etc.

Proof of Proposition 4

If $N_0^* \leq N_i^*$ for all $i > 0$, then $0 \in M(N^*)$. By the definitions of $M(N^*)$ and N_0^*, this implies that $u_0(N^* + 1) < u_0(1)$, which in turn implies that $u_0(N^* + 1) < C$. Hence, if agent 0 expects the information to travel all the way to N^*, she will not acquire the information. By assumption, there exist some $n > 1$ such that $u_0(n) \geq C$.

Proof of Proposition 5

Access to partly rival information requires a direct or indirect connection to the originator. The single originator has no incentives to decline an invitation. Taken together, this implies that everyone is somehow connected in equilibrium. For the purpose of this proof, let us refer to those that are directly connected to the originator as "friends," and those who are only indirectly connected to the originator as "friends of friends." We now establish that it is a strictly dominant strategy for every imitator to become a friend rather than a friend of friend. Suppose that some imitator i is a friend of friend, only connected to the originator only through some friend $i' \neq i$. Recall that the originator prefers to circulate the information among its closest friends, and also that i' would share information with i (under the assumptions of the network formation game). So, if there are (more than) τ friends *that are not connected to friends of friends*, neither i' nor, by extension, i receive the information. Hence, i would rather (deviate to) become a friend. Alternatively, if there are fewer than τ friends *that are not connected to friends of friends*, the other friends may also receive the information with some probability smaller than 1. Again, i would rather (deviate to) become a friend because it would then receive the information with certainty.

Proof of Proposition 6

Consider N_o separately formed stars, each with an originator at its core, who is selected by nature (to be endowed with information) with probability $1/N_o$. In principle, each imitator has incentives to connect to another originator, and hence another star, in order to receive information with a higher probability. However, under our assumptions, such an imitator would relay information received from one star to (the originator in) the other star. Hence, an imitator that connects to another originator forfeits the trust of both originators, each of whom would rather confide in its other friends (of whom there are more than τ). It is straightforward to see that originators can only gain from connecting to each other, since they can then share non-rival information while withholding partly rival information from each other. Thus, they will connect.

Proof of Proposition 7

Recall that the overriding objective of all individuals, imitators and originators alike, is to gain access to partly rival information. Originators obtain no *informational* benefits from connecting to imitators. Therefore, they first and foremost want to be friends with each other, thereby entering into a mutual information "insurance." If there are no more than τ originators, all of them connect to each other. The rationale for connecting *directly* to another originator is similar to that in Proposition 5; it is a (here, weakly) dominant strategy. If there are less than τ originators, there is scope for at least some, but not all, imitators to connect to this "core" of originators. Some imitators are excluded to ensure that, ultimately, only τ individuals share the information. For the same reason, if there are more than τ originators, the originators split up into several groups, each consisting of no more than τ individuals. This is because, if one such group would connect, and hence spread the information among, more than τ individuals, a selected originator in that group would rather withhold than share the information.

Proof of Proposition 8

First, there can be no pairwise stable equilibrium without origination. Suppose there are some invitations, and everyone expects others not to produce information. Then, at least one individual would find it optimal to decline all but one (or up to τ) invitations and to produce information. Suppose there are no invitations, and everyone expects others not to produce information. Then, at least two individuals would find it optimal to create a relationship, with one of them producing information. Second, there can be no equilibrium with more than $\tau/2$ originators. Recall that information production is worthwhile only if it can be shared with at least one but at most τ friends. A fortiori, nor is it worthwhile if the information is owned by more than τ *non*-friends. If there are more than $\tau/2$ originators, there is either at least one originator who does *not* share its information or more than τ individuals end up having the information. In either case, at least one originator would not find it worthwhile to produce the information. Third, originators do not need to connect to each other, since they need not obtain information from each other. Fourth, every imitator is *directly* connected to *some* originator, for the same reason as in Proposition 5. Fifth, imitators are not connected to each other. Suppose two imitators i and $i' \neq i$ are connected. If both of them are connected to the same (sub)set of originators, their connection is either redundant (because both of them get information directly from some originator) or suboptimal (because, say, i does not get the information from some originators only because i would share the information with i'). If they are not connected to the same (sub)set of originators, at least one of the originators connected to, say, i is unwilling to share the information with i because i would share the information with i'. Sixth, every imitator is connected to all originators. Connecting to several originators is an individually rational strategy for an imitator to improve its access to information, since the originators randomize between equally attractive imitators. At the same time, being connected to other *originators* is not costly in terms of trust; originators are not concerned about imitators communicating information to other originators, because the latter are already informed. Seventh, there can be an equilibrium with more than one originator. Let $\tau > 3$. Suppose there are two originators; each shares its information with at least one imitator, and both together share information with τ imitators. Consider whether one of the originators would rather become an imitator. As an imitator, it would receive the information for free with probability $1/(N-1)$, since the remaining originator would randomize between all imitators. Clearly, originating the information and having it shared with τ imitators becomes more attractive (at some point) as $N \to \infty$.

References

Bramoulle, Yann and Rachel Kranton, 2007, "Public Goods in Networks," Journal of Economic Theory, 135, 478-494.

Galeotti, Andrea, Christia Ghiglino, and Francesco Squintani, 2009, "Strategic Information Transmission in Networks," Discussion Paper No. 668, Department of Economics, University of Essex.

Galeotti, Andrea and Sanjeev Goyal, 2010, "The Law of the Few," American Economic Review, 100, 1468-1492.

Goyal, Sanjeev, 2007, Connections: An Introduction to the Economics of Networks, Princeton University Press, Princeton, NJ.

Jackson, Matthew O., 2008, Social and Economics Networks, Princeton University Press, Princeton, NJ.

Lippert, Steffen and Giancarlo Spagnolo, 2010, "Networks of Relations and Word-of-Mouth Communication," SSE/EFI Working Paper in Economics and Finance No. 570.

Stein, Jeremy C., 2008, "Conversations among Competitors," American Economic Review, 98, 2150-2162.