Journal of Theoretical Biology 398 (2016) 43–51

Contents lists available at ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/yjtbi

The evolutionary advantage of limited network knowledge

Jennifer M. Larson

New York University, 19 W. 4th St., New York, NY 10012, United States

HIGHLIGHTS

- I model the enforcement of cooperation when networks spread gossip relevant to punishment.
- Network structure matters for enforcing cooperation.
- Limited knowledge of network structure among the actors makes enforcing cooperation easier.
- Social sciences show real people have cognitive limits to network knowledge.
- Groups with imprecise network knowledge have an evolutionary fitness advantage.

ARTICLE INFO

Article history:
Received 30 September 2015
Received in revised form
2 February 2016
Accepted 11 March 2016
Available online 21 March 2016

Keywords:
Cooperation
Evolution
Networks
Cognitive limitations
Sanctioning

ABSTRACT

Groups of individuals have social networks that structure interactions within the groups; evolutionary theory increasingly uses this fact to explain the emergence of cooperation (Eshel and Cavalli-Sforza, 1982; Boyd and Richerson, 1988, 1989; Ohtsuki et al., 2006; Nowak et al., 2010; Van Veelen et al., 2012). This approach has resulted in a number of important insights for the evolution of cooperation in the biological and social sciences, but omits a key function of social networks that has persisted throughout recent evolutionary history (Apicella et al., 2012): their role in transmitting gossip about behavior within a group. Accounting for this well-established role of social networks among rational agents in a setting of indirect reciprocity not only shows a new mechanism by which the structure of networks is fitness-relevant, but also reveals that *knowledge* of social networks can be fitness-relevant as well. When groups enforce cooperation by sanctioning peers whom gossip reveals to have deviated, individuals in certain peripheral network positions are tempting targets of uncooperative behavior because gossip they share about misbehavior spreads slowly through the network. The ability to identify these individuals creates incentives to behave uncooperatively. Consequently, groups comprised of individuals who knew precise information about their social networks would be at a fitness disadvantage relative to groups of individuals with a coarser knowledge of their networks. Empirical work has consistently shown that modern humans know little about the structure of their own social networks and perform poorly when tasked with learning new ones. This robust empirical regularity may be the product of natural selection in an environment of strong selective pressure at the group level. Imprecise views of networks make enforcing cooperation easier.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Groups of individuals have social networks that structure interactions within the groups. Evolutionary theory increasingly turns to the fact that social networks can constrain who encounters whom to explain the emergence of cooperation (Eshel and Cavalli-Sforza, 1982; Boyd and Richerson, 1988, 1989; Ohtsuki et al., 2006; Nowak et al., 2010; Van Veelen et al., 2012). However, social networks perform another evolutionarily important function in groups of individuals: they transmit gossip about behavior, an activity that occupies approximately two-thirds of individuals' conversation time (Dunbar, 2004). When considering the role of networks in the spread of information, the relevant constraint becomes who *communicates* with whom.

I model a group of strategic agents who use gossip that spreads through their social network to identify and sanction uncooperative behavior via indirect reciprocity (Nowak and Sigmund, 1998, 2005; Wedekind and Milinski, 2000; Fishman, 2003; Mohtashemi and Mui, 2003; Brandt and Sigmund, 2004). Accounting for this function of social networks not only shows one mechanism by which networks are fitness-relevant, but also reveals that *knowledge* of social networks can be fitness-relevant as well. In fact,

E-mail address: jenn.larson@nyu.edu

http://dx.doi.org/10.1016/j.jtbi.2016.03.017
0022-5193/© 2016 Elsevier Ltd. All rights reserved.

groups comprised of individuals who possess precise information about their social networks would be at a fitness disadvantage relative to groups with individuals who hold a coarser picture of their networks.

Empirical work has consistently shown that modern humans know very little about the structure of their own social networks and perform poorly when tasked with learning new social networks (Bernard et al., 1980, 1982; Bondonio, 1998; Simpson et al., 2011; De Soto et al., 1968; Picek et al., 1975; Freeman, 1992; Kumbasar et al., 1994). This robust empirical regularity may be the product of natural selection in an environment of strong selective pressure at the group level. Imprecise views of networks make enforcing cooperation easier.

2. Networks and the evolution of cooperation

Social networks research is premised on the recognition that individuals have ties to some – often not all – others within a group of interest, and these ties connote affinity, or frequent interaction, or preferential treatment, or kinship, or a channel of information, or a source of peer pressure, or possibly all of these. Social networks structure relationships among individuals within groups.

Social networks appear to have been a part of group structure over long timescales (Apicella et al., 2012). Evolutionary theory increasingly turns to this structure as part of the explanation for the emergence of cooperation (Nowak et al., 2010): repeated interactions with social contacts allow for direct reciprocity (Boyd and Richerson, 1988), closed communities allow for indirect reciprocity (Boyd and Richerson, 1989), and the ability to form lasting associative connections fosters cooperation that persists over time (Eshel and Cavalli-Sforza, 1982).

Key to these explanations is a set of interactions that are structured – who encounters whom is limited and fixed (Ohtsuki et al., 2006). However, as groups become more sophisticated, complex and mobile over time, the structure may manifest itself not in constraints on who *encounters* whom but on who *communicates* with whom. The more collaborative activities a group undertakes, the more chances there are to encounter all others in the group while communication about behavior in these encounters may be confined to social contacts. In a modern market setting, for instance, members of a group may all have a chance of transacting with any other, but the social network determines whom they tell about their interactions.

The role of networks in information spread has been widely recognized. In modern societies, social networks have been found to structure communication relevant to finding jobs (Granovetter, 1973), learning about new agriculture technologies (Conley and Udry, 2010), making financial decisions (Duflo and Saez, 2002), making preventive health decisions (Rao et al., 2007), establishing social collateral (Karlan et al., 2009), becoming aware of opportunities for gain (Larson and Lewis, 2016), and, frequently, gossiping about the behavior of others (Dunbar, 2004; Gluckman, 1963).

Some argue that language may have evolved to facilitate group bonding (Dunbar, 1998), perhaps specifically to allow gossip as a means of social control as group size became large (Enquist and Leimar, 1993; Wilson et al., 2000). A host of non-evolutionary models show that social sanctioning by peers can enforce cooperative behavior (Kandori, 1992; Greif, 1993; Fearon and Laitin, 1996; Dixit, 2004; Wolitzky, 2013; Larson, 2016), and experimental subjects gossip in ways that ultimately support higher levels of cooperation (Sommerfeld et al., 2007).

Standard models of the evolution of cooperation leave little room for communication and gossip in the form observed among modern humans.[1] In these models, agents need know nothing about the environment they are in, the game they are playing, what individuals far away are doing, and they do not form a forward-looking strategy. Such approaches generate elegant explanations for the evolution of cooperation and its stability. However, given the empirical role that social networks play in spreading information, the utility of this information for sanctioning strategies that are able to promote cooperation, and the observation that evolution may also shape social networks (Apicella et al., 2012), a full understanding of the fitness implications of social networks requires accounting for realistic gossip.

3. Network knowledge

The model below isolates a mechanism by which social networks bear on group fitness: groups of strategic actors use gossip transmitted through their network to identify and punish noncooperators. Accounting for gossip in a strategic setting reveals a surprising relationship between cooperation and network knowledge: knowing less about one's network can make enforcing full cooperation easier.

Limited network knowledge facilitates cooperation because individuals occupying certain positions within a network can be particularly tempting targets of misbehavior. Some are unable to spread gossip about others' misbehavior widely and quickly, and so misbehavior targeted at them is more difficult to sanction. When everyone perfectly knows the full network structure, these individuals can be perfectly identified and targeted. When instead individuals possess a coarser, limited view of their network, identifying the most vulnerable can be impossible, removing the temptation to act uncooperatively.

The advantage of limited network knowledge helps make sense of the robust empirical finding that modern humans know very little about their own social networks despite making constant use of them. Individuals perform consistently poorly when recalling the structure of their own social networks and when retaining information about new social networks (Bernard et al., 1980, 1982; Bondonio, 1998; Simpson et al., 2011). People rely on a series of imprecise compression heuristics to store the complicated object in memory (Brashears, 2013; Krackhardt and Kilduff, 1999; De Soto, 1960), recall networks with a high degree of error (De Soto et al., 1968; Picek et al., 1975; Freeman, 1992; Kumbasar et al., 1994), and assign low salience to the precise recall of links (Killworth and Bernard, 1976; Brewer and Webster, 2000). Early human groups would not have been much easier to keep track of since even the smaller social groups of the Pleistocene had around 50 members (Dunbar, 1992), yielding 2450 possible relationships.

The model below shows that precise knowledge of the network structure can confer a fitness *disadvantage* on the group – limited knowledge of the network makes enforcing cooperation easier, often substantially so. This suggests that under multilevel selection (Pacheco et al., 2006; Chalub et al., 2006), in the presence of high selective pressure at the group level (Bowles, 2006, 2009), groups of individuals with a limited capacity to know and recall the full network face an evolutionary advantage.

[1] While some argue that gossip may be a source of rapid agreement on a player's reputation (Ohtsuki et al., 2009; Nakamura and Masuda, 2011; Ghang and Nowak, 2015), with an attendant danger of manipulation through false gossip (Nakamaru and Kawata, 2004; Nowak and Sigmund, 2005; Sommerfeld et al., 2007), gossip has yet to be incorporated as the flow of information through links in a social network or among strategic, forward-looking players.

4. Strategic social sanctioning

The model below considers a group of individuals who can randomly interact but who share gossip with a constrained set of others determined by their social network. Individuals use the gossip to sanction misbehavior and enforce cooperation. When someone misbehaves, others hear about it from their social contacts, who spread the news to their social contacts, and so on. Knowledge of misbehavior allows others to punish it; the threat of this punishment can incentive cooperation. Specifically, an individual's expectation that many others will learn about his misdeed and punish him can incentivize him to behave cooperatively.

The approach follows the community enforcement literature (Kandori, 1992) and others who have accounted for networks in strategic play (Larson, 2016; Wolitzky, 2013). The game is not evolutionary, but has implications for group fitness. The total value experienced by the group is increasing in the number of cooperators and maximized under full cooperation. Given the distinctly lethal character of inter-group competition among humans (Bowles, 2006), this value can be thought of as group fitness.

Agents in the model use strategies that mimic strategies used by real groups of individuals who turn to in-group policing and threats of social sanctions to enforce cooperative behavior (Fearon and Laitin, 1996; Miguel and Gugerty, 2005; Gugerty, 2007). Moreover, a propensity to sanction uncooperative behavior is strong in laboratory settings; individuals will sanction even at cost to themselves (Fehr and Gächter, 1999, 2002). In the model, sanctioning is incentive compatible and over-sanctioning must be policed by peers. That cooperative behavior, including refraining from antisocial punishment, can be sustained in equilibrium is consistent with experimental evidence that subjects apply punishment in ways they believe to be consistent with norms shared among their group (Ertan et al., 2009).

I compare two scenarios of knowledge that the individuals are assumed to possess about the network in which they reside. The first assumes that all individuals know their full network with precision. The second (and more realistic) assumes that individuals know only a very coarse bit of information about the network. In the former, any individual could calculate "if I behave uncooperatively against i, exactly n people will know in the next period, n' people will know in the period that follows, n'' people will know in the period that follows…" and so on. In the latter, the coarsening is equivalent to knowing "if I behave uncooperatively against anyone, roughly n people will eventually know."

I show that coarse information about the network actually makes enforcing cooperation easier, often strictly easier.

5. Model setup

Let N be a group with players $\{1, \ldots, n\}$. Define an infinitely repeated game G such that all players play one round of prisoner's dilemma with an independently drawn random opponent each period, so that the probability of interacting with a given other player is $\frac{1}{n-1}$ in each period. Nature reveals to each player only his own pairing. Each round faces payoffs:

$$\begin{array}{cc} & \begin{array}{cc} c & d \end{array} \\ \begin{array}{c} c \\ d \end{array} & \begin{pmatrix} 1,1 & -\beta,\alpha \\ \alpha,-\beta & 0,0 \end{pmatrix} \end{array}$$

where $\alpha > 1$, $\beta > 0$, and $\frac{\alpha-\beta}{2} < 1$. Players discount future payoffs with a common discount factor $\delta < 1$.

Gossip within the group occurs according to a "communication network" defined by the pair (g^N, N) with $n \times n$ adjacency matrix g^N where $g^N_{i,j} = g^N_{j,i} = 1$ indicates a link between $i \neq j \in N$. For simplicity, I will refer to the network as "g^N," or simply as g when the group identity is unimportant. Links in the networks are undirected and unweighted.

In a standard setup with rational agents (as in Larson, 2016), all individuals would be assumed to have common knowledge of their network structure, which means all individuals would know the set of nodes N and the full set of links among them, g^N, and would know that others know this, would know that others know that others know this, and so on. This standard approach assumes one of many possible network knowledge environments in which individuals could be situated. I label it *PRECISE* since it entails all agents knowing their network structure precisely, and compare it to an alternate network knowledge environment defined below, *COARSE*, in which agents know only a coarse description of their network.

The communication network transmits relevant information about rounds (clarified below). Actions are unobservable to individuals not party to the interaction. Information spreads truthfully and deterministically through the network at a rate $r = \frac{\textit{Degrees Spread}}{\textit{Rounds Played}}$. The general results hold when information spread is stochastic, though require much more cumbersome notation. The appendix explores the potential problem of manipulation through untruthful gossip (Nakamaru and Kawata, 2004; Sommerfeld et al., 2007). When $r = 1$, news about rounds spreads one degree, to the immediate neighbors, before players are rematched and play again. When $r = 2$, news spreads two degrees, from person to person to person along links in the social network before the next round, and so on (Fig. 1).

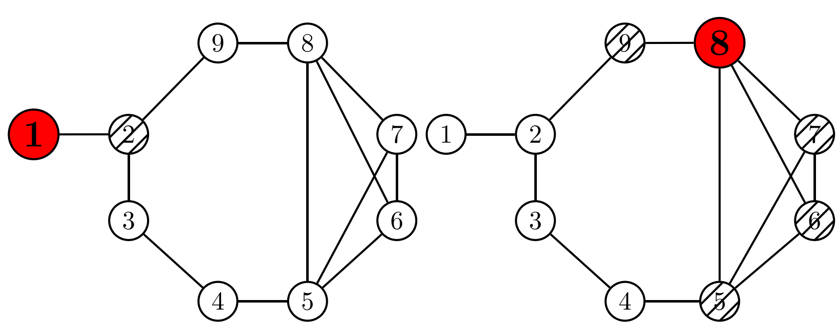

Fig. 1. The reach of messages by $t+1$ when sent by player 1 (left) in t and player 8 (right) in t when $r=1$ in an example network. Gossip sent by player 8 reaches more people quickly than gossip sent by player 1.

6. Strategies

Here I focus on in-group policing strategies that make use of gossip and threaten to punish defectors for T rounds, where punishment takes the form of capitulation (the actor being punished plays c while the punisher is allowed to play d, akin to extracting a fine).[2] These strategies are the rational agent analogue to strategies with second-order norms like stern-judging (Pacheco et al., 2006; Scheuring, 2010) and "Kandori" (Ohtsuki and Iwasa, 2007): they regard those who cooperate with "good" individuals and defect against "bad" individuals as "good", and those who defect against "good" individuals and cooperate with "bad" individuals as "bad."

The strategies rely on gossip spread through the network: they are a mapping from messages received into actions. While others have considered the role of gossip as a general source of coordination (Ohtsuki et al., 2009), here gossip is specifically a set of messages sent by victims of misbehavior to neighbors in the network who pass it along to their neighbors who pass it along and so on.

Suppose that victims of misbehavior send messages to their network neighbors about the offender, including the offender's identity and the round in which the offense occurs.[3] Strategies are then a mapping from the content of these messages received into the action set $\{c, d\}$.

Specifically, call $M_{i,t}$ the set of individuals about whom i has received messages by the start of time t about rounds that have occurred since $t-T$. Then we can define an in-group policing strategy profile σ^{NWIGP} on a network that maps $M_{i,t} \rightarrow \{c, d\}$:

All players play c in the first round. Subsequently, all players i in round t play c if matched with $j \notin M_{i,t}$ and d if matched with $j \in M_{i,t}$. When player j plays player i and $j \in M_{i,t}$ but $i \notin M_{j,t}$, (i knows j deserves punishment but not vice versa), j plays c while i plays d. When j plays i and $j \in M_{i,t}$ and $i \in M_{j,t}$ (both know each other deserves punishment), both play d.

Define the messages as follows: when i deviates from σ^{NWIGP} in a round with in-group member j, j sends a message $m_{j,i,t} = \{i, t\}$ to himself and his neighbors which spreads through the network g at rate r.

The appendix elaborates on the behavior prescribed by the strategies both in and out of equilibrium. Incentives to deviate depend on the number of others who receive messages from any prospective victim and how quickly, which is a function of the network. Let $\ell(i,j)$ be the length of the shortest path from i to j. Define i's **k-neighborhood** in network g^A, $N_i^k(g^A)$ to be the set of all j such that the shortest path from i to j is less than or equal to k. That is,

$$N_i^k(g^A) = \{j \in A : \ell(i,j) \le k, i \ne j\}.$$

The reach of messages is a function of the network structure and the size of k-neighborhoods within the network.

7. Optimality of limited network knowledge

Given the bounds on α and β, the globally efficient outcome is full cooperation. Any deviation destroys value, so although an evolutionary process is not modeled here, group fitness would be increasing in the number of cooperators in equilibrium. The conditions under which σ^{NWIGP} supports full cooperation depends on the network knowledge environment.

Define network knowledge environment *PRECISE* as that in which all players perfectly know g– the set of nodes and all links among them– and this knowledge is common knowledge. Then we have the following conditions for full cooperation in equilibrium, with the proof and a discussion of beliefs that extend the behavior to sequential equilibrium in the appendix:

Proposition 1 (*Perfect Network Knowledge*). *Under PRECISE,* σ^{NWIGP} *is sequentially rational given r, T, and g iff*

$$\delta^T \ge \frac{(\alpha-1)(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)} \tag{1}$$

and

$$\delta^T \ge \frac{\beta(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)}. \tag{2}$$

Now consider an information environment in which individuals have a much more limited view of their network. Define network knowledge environment *COARSE* as that in which all individuals know N, the set of nodes in the network, and the average Tr-neighborhood in the network, $\frac{1}{n}\sum_{j \in N} \#N_j^{Tr}(g)$. This exhausts the information that each player possesses about the network g, and this knowledge is common knowledge. Then we have the following conditions under which σ^{NWIGP} supports full cooperation in equilibrium, again with the proof and discussion of consistent beliefs in the appendix:

Proposition 2 (*Limited Network Knowledge*). *Under COARSE,* σ^{NWIGP} *is sequentially rational given r, T, and g iff*

$$\delta^T \ge \frac{(\alpha-1)(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j \in N}\#N_j^{Tr}(g)\right)} \tag{3}$$

and

$$\delta^T \ge \frac{\beta(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j \in N}\#N_j^{Tr}(g)\right)}. \tag{4}$$

Consider the comparison between the ease of supporting full cooperation under precise and limited network knowledge. The key result is as follows:

Proposition 3 (*Cooperation Easier with Limited Network Knowledge*). *Full cooperation can be enforced more easily under COARSE than under PRECISE. Specifically, full cooperation can be supported by*

[2] The strategies of interest are symmetric in-group policing strategies with finite, symmetric capitulation punishment that can support full efficiency in sequential equilibrium. Symmetry in the strategies represents a group that has coordinated on a shared norm of enforcement. Finite punishment strategies are considered because real groups appear to use strategies like these when enforcing cooperation (Fearon and Laitin, 1996) and because finite punishments destroy less value out of equilibrium and so have desirable efficiency properties in an environment with errors or mistakes. Capitulation punishments, those that require an offender to capitulate and play c while the punisher plays d, resemble real world punishment schemes in which an offender can atone or pay a fine in exchange for forgiveness, and are renegotiation proof (punishers like to punish). Symmetric punishment means that the punishment warranted for an offense is invariant to the identity of the offender. This property is essentially egalitarianism, that like offenses are punished alike.

[3] We could add to realism by further supposing that the offenders neither participate in initially spreading nor passing on information about offenses they committed. Adding the latter stipulation dramatically increases the amount of notation that must be carried around to account for potential bottle-necks without changing the character of the results.

a lower minimum discount factor under limited network knowledge than under perfect network knowledge.

The proof is in the appendix. The intuition for the difference is as follows: some individuals on a network would be tempting targets of defection, particularly those with the smallest *Tr*-neighborhoods. These individuals would spread messages about defections, but these messages have the narrowest reach. Consequently, the expected punishment for defecting against these individuals would be the weakest. When the precise network structure is common knowledge, these individuals can be identified, and so the strong incentives to defect against them must be countered.

When instead only a coarser view of the network structure is common knowledge, the network may offer tempting opportunities for defection, *but they cannot be identified*. If there is an individual who has a very small *Tr*-neighborhood, that person would be a tempting target if the precise network structure were common knowledge. However, since all individuals know only the mean and lack a more precise view of the structure, this opportunity remains concealed. So long as everyone knows that everyone knows that the mean *Tr*-neighborhood size is all anyone knows, full cooperation can be enforced.[4] The same general result holds under a variety of coarsenings, discussed in the appendix.

The extent to which enforcing cooperation is easier under limited network knowledge depends on the difference between the smallest *Tr*-neighborhood and the mean *Tr*-neighborhood– the extent to which the most peripheral network position is peripheral.

8. Conclusion

This paper accounts for the role that social networks play in spreading gossip among strategic agents who use social sanctioning to enforce cooperation in equilibrium. Doing so reveals that the social network structure and individuals' *knowledge* of the social network structure are strategically relevant and bear on a group's ability to enforce cooperation. It also reveals that precise knowledge of networks makes enforcing full cooperation harder, often strictly harder, than when group members merely have coarse representations of the networks in mind.

In an environment of multilevel selection (Pacheco et al., 2006; Chalub et al., 2006), groups comprised of individuals who know and recall their network with low precision would face a fitness advantage. This would be especially valuable in cases of strong selective pressure at the group level (Bowles, 2006, 2009), where an ability to ensure the cooperation of all group members in tasks like defense initiatives or emergency measures could mean the difference between fending off an invasion or surviving an environmental disaster and the demise of the group. This suggests that the regular finding that humans are quite poor at reproducing their social networks with accuracy and use heuristics and highly compressed representations to store their networks in memory may be evolutionarily relevant.

The evolutionary relevance of network knowledge is hinted at by a number of sources. Intuitively, knowledge of a network structure helps an individual determine how to navigate the social environment (Janicik and Larrick, 2005). Empirical evidence suggests that knowing a network's structure can offer advantages to individuals (Krackhardt, 1992; Krackhardt and Hanson, 1993; Morris, 1997), and theoretical models suggest that knowing the network structure allows social sanctioning that promotes group cooperation (Larson, 2016). Moreover, there is a relationship between physical properties of the brain and social networks in general (Bickart et al., 2011; Meyer et al., 2012; Sallet et al., 2011; Zahn et al., 2007), between memory capacity and social network size (Stiller and Dunbar, 2007), and these relationships hold for online social networks as well (Kanai et al., 2011). Given that social networks may be shaped by natural selection (Apicella et al., 2012) and that knowledge of networks determines a social networks' impact on group fitness, it is plausible that natural selection also shapes the cognitive ability to store and recall networks.

That cognitive limits to knowing networks may be the product of natural selection suggests an additional implication: when individuals err in reproductions of their networks or of newly presented networks, they should err in ways that make cooperation more likely. Indeed, this appears to be the case. Individuals are more likely to add new links than to forget existing links, and the added links are those that alter the network structure in ways that would reduce the strongest incentives to behave uncooperatively (De Soto et al., 1968; Freeman, 1992).

Appendix A

A.1. Elaboration of strategies

Call $M_{i,t}$ the set of individuals about whom i has received messages by the start of time t about rounds that have occurred since $t-T$. Then we can define an in-group policing strategy profile σ^{NWIGP} on a network that maps $M_{i,t} \to \{c, d\}$:

All players play c in the first round. Subsequently, all players i in round t play c if matched with $j \notin M_{i,t}$ and d if matched with $j \in M_{i,t}$. When player j plays player i and $j \in M_{i,t}$ but $i \notin M_{j,t}$, (i knows j deserves punishment but not vice versa), j plays c while i plays d. When j plays i and $j \in M_{i,t}$ and $i \in M_{j,t}$ (both know each other deserves punishment), both play d.

Define the messages as follows: when i deviates from σ^{NWIGP} in a round with in-group member j, j sends a message $m_{j,i,t} = \{i, t\}$ to himself and his neighbors which spreads through the network g at rate r.

This strategy profile is analogous to behavioral strategies that instruct players to cooperate with those with a "good" reputation and defect against those with a "bad" reputation combined with higher-order assessments that those who comply with this rule maintain good standing and those who do not enter bad standing (Ohtsuki and Iwasa, 2007).

Here players use gossip sent to them in the form of messages from victims of misbehavior to determine standing. $M_{i,t}$ is the set of players known to i to be in bad standing at the start of period t. $N \backslash M_{i,t}$ is then the set of players thought to i to be in good standing at the start of period t. Note that $N \backslash M_{i,t}$ may contain someone who in fact misbehaved, but i simply has not heard gossip about this by the beginning of t and so regards that player to be in good standing. In this sense, a strategy that has i only punish those in $M_{i,t}$ is "generous" (Scheuring, 2010). The conditions below ensure that rational agents playing this generous strategy are still incentivized to refrain from misbehaving.

[4] The network knowledge environments considered here assume that all individuals have access to the same information: under *PRECISE*, all individuals know the network precisely, and under *COARSE*, all individuals know the same coarse measure of the network. It could be that individuals instead have heterogeneous knowledge of the network structure, some knowing the network more precisely than others. The character of the results holds for heterogeneous states of network knowledge: access to precise information reveal opportunities to defect that may make defecting more profitable *for those who know about them* than it would be if this information were concealed. Additionally, heterogeneous network knowledge also opens the possibility of differential abilities to exploit others since those with access to precise network knowledge would be able to exploit the peripheral network positions in ways that those without access to this knowledge could not, an intriguing subject of future research. Thanks to an anonymous reviewer for this suggestion.

All players are instructed to play c in the first round. Any time a player i plays someone in $N\backslash M_{i,t}$ (about whom he has heard nothing bad), he is instructed to play c. Any time he plays someone in $M_{i,t}$, he is instructed to play d. When a player has misbehaved, he is to play c for the next T rounds, unless he plays someone in $M_{i,t}$, in which case he plays d. During these T rounds, anyone who has received gossip about his misbehavior is to play d against him.

Informally, a player i identifies misbehavior as any defection against him or resistance of punishment by him while i is in good standing. More precisely, a player i identifies misbehavior committed against him as the following: anyone playing d against him when i has not misbehaved – that is, when i has not played d against anyone in $N\backslash M_{i,t}$ – in the last T rounds. This is true for anyone playing d against him, both in the case that it is played by someone in $N\backslash M_{i,t}$, and in the case that it is played by someone in $M_{i,t}$ – someone resisting punishment. Any d played against i when i *has* misbehaved– has played d against at least one person in $N\backslash M_{i,t}$ – in the last T rounds is not misbehavior. When i experiences misbehavior by opponent j, he sends a message saying so – who misbehaved and when – to his network neighbors which spreads as gossip through the network.

Key to the success of these strategies is that a player i can always identify misbehavior committed against himself. If he has misbehaved in the last T rounds, no action against him counts as misbehavior. If he has not misbehaved in the last T rounds, any d played against him counts as misbehavior. While i cannot be sure that, given that he misbehaved, any particular player playing d against him actually knew this, the conditions below ensure that this ambiguity does not have consequences for cooperation (essentially because no one would have the right information to exploit this).

Consequently, here cooperation can be supported even when different individuals have different access to gossip and hold different assessments of the same individuals. The conditions in Propositions 1 and 2 ensure that nonetheless, full cooperation can be supported in sequential equilibrium.

A.2. Proof of Proposition 1

Proof. Let C_t be the set of players in N who would incur no punishment according to σ^{NWIGP} in t if the network were complete, and $\overline{C}_t$ be the set of players in N who would incur punishment in t according to σ^{NWIGP} if the network were complete. Hence C_t are the "cooperators" in t, $\overline{C}_t$ are the "defectors" in t, $C_t \cup \overline{C}_t = N$ and $\#(C_t \cup \overline{C}_t) = n$.

If network g is incomplete, a player $i \in N$ may be unaware of at least one defection for at least some of the time. Let $C_{i,t}$ be the set of players that i does not know to be defectors at time t, and $\overline{C}_{i,t}$ be the set of players that i knows to be defectors at time t. Clearly $\overline{C}_{i,t} \subseteq C_t$ and once again $C_{i,t} \cup \overline{C}_{i,t} = N$.

Let $K_{i,j,t,l}$ be the set of players who have received a message about i's round with j that occurred in period t by period $t+l$. Let $\overline{K}_{i,j,t,l}$ be the set of players in i's group that do not know about his round with j that occurred in t by period $t+$l. For $i \in N$, $K_{i,j,t,l} \cup \overline{K}_{i,j,t,l} = N$ for any j.

If player k does not know about any of player i's defections, then we can say that to player k, i is an unknown defector (which is to say that k treats i as if he were a cooperator).

In the strategy profile σ^{NWIGP}, there are 6 possible ways to deviate from the strategy and "defect." A player i can defect while he is a cooperator by playing d against a cooperator/unknown defector (c1), or playing c against a known defector (c2). A player i can defect when he is a defector by playing d against a cooperator/unknown defector when i is unknown (d1), playing d against a cooperator/unknown defector when i is known (d2), playing c against a defector when i is unknown (d3), or playing c against a defector when i is known (d4).

To establish sequential rationality, I will show that for any history and at any information set, a player prefers to comply with σ^{NWIGP} given the conditions above.

First consider a player i who has defected most recently in $t-t^d$ and is contemplating defecting against opponent j whom i does not know to be a defector in period t. In other words, i is considering defecting via (d1). A defection in t would result in punishment until $t+T$. He, being a defector, already expected punishment until $t+T-t^d$. Cooperating in t is preferred iff

$$\alpha-1 \le \sum_{l=T-t^d+1}^{T-1} \delta^l \left[\frac{\#(C^{\star}_{i,t+l} \cap K_{i,j,t,l})}{n-1}(1+\beta)+\frac{\#(\overline{C}^{\star}_{i,t+l} \cap K_{i,j,t,l})}{n-1}\alpha\right] + \delta^T \left[\frac{\#(C_{i,t+T} \cap K_{i,j,t,T})}{n-1}(1+\beta)+\frac{\#(\overline{C}_{i,t+T} \cap K_{i,j,t,T})}{n-1}\alpha\right]$$

for any $j \in N$ where $C^{\star}_{i,t+l}$ is the set of cooperators that i *believes* he will know about in period $t+$l.[5] This condition is hardest to satisfy in the set of histories in which i's defection was in the previous period, i.e. where $t^d = 1$. The condition then becomes

$$\alpha-1 \le \delta^T \left[\frac{\#(C_{i,t+T} \cap K_{i,j,t,T})}{n-1}(1+\beta)+\frac{\#(\overline{C}_{i,t+T} \cap K_{i,j,t,T})}{n-1}\alpha\right]$$

for any $j \in N$. In $t+T$, $C_{i,t+T} = A\backslash i$ and $\overline{C}_{i,t+T} = \varnothing$ so the condition reduces to:

$$\frac{\alpha-1}{\delta^T} \le \min\left\{\frac{\#K_{i,j,t,T}}{n-1}(1+\beta), \frac{\#K_{i,j,t,T}}{n-1}\alpha\right\} \tag{5}$$

for any $j \in N$.

To safeguard against (c1) defections, note that expected punishment from defecting via (c1) is the same as the expected punishment from defecting via (d1) when $t^d = T^p$, i.e. when the most recent defection was far enough in the past. The above makes clear that condition (7) is sufficient to prevent defections via (c1).

Defections via (d2) incur the same expected punishment but gain β from playing d rather than c against an opponent's d. The condition is then

$$\frac{\beta}{\delta^T} \le \min\left\{\frac{\#K_{i,j,t,T}}{n-1}(1+\beta), \frac{\#K_{i,j,t,T}}{n-1}\alpha\right\} \tag{6}$$

for any $j \in N$.

Deviations according to (c2), (d3) and (d4) are trivially not preferred. Deviating according to (c2) and (d3) would entail foregoing the gains from issuing punishment and in doing so, earn punishment in return. Deviating according to (d4) would entail foregoing the chance to not incur a loss during punishment and earning punishment as a consequence.

Given the way information spreads along g, $\#K_{i,j,t,T} = \#N_j^{Tr}(g)$. Hence if conditions (7) and (8) are satisfied for all players, no player has an incentive to deviate in any history given any beliefs and so σ^{NWIGP} is sequentially rational. Since the binding conditions for sequential rationality are independent of beliefs, any consistent beliefs trivially extend the behavior to sequential equilibrium.□

A.3. Proof of Proposition 2

Proof. The proof is nearly identical under limited network information to the proof under perfect network information. The only difference is that individuals do not know $\#K_{i,j,t,l}$ with precision. Instead, they hold beliefs over the number of others in this set for

[5] Strictly speaking, the condition as written holds for $t^d \ge 2$, and should be written without the first sum for $t^d = 1$.

$l < T$, and the know the expected number in this set for $l = T$. As in the above proof, any consistent beliefs will do because the binding case will again be independent of beliefs.

Use the same notation as above, and define the analogous term $K^{\star}_{i,j,t,l}$ to be i's beliefs about who will know about an interaction with j in t by $t+l$.

Once again to establish sequential rationality, I will show that for any history and at any information set, a player prefers to comply with σ^{NWIGP} given the conditions above.

First consider a player i who has defected most recently in $t-t^d$ and is contemplating defecting against opponent j whom i does not know to be a defector in period t. In other words, i is considering defecting via (d1). A defection in t would result in punishment until $t+T$. He, being a defector, already expected punishment until $t+T-t^d$. Cooperating in t is preferred iff

$$\alpha-1 \le \sum_{l=T-t^d+1}^{T-1} \delta^l \left[\frac{\#(C^{\star}_{i,t+l} \cap K^{\star}_{i,j,t,l})}{n-1}(1+\beta) + \frac{\#(\overline{C}^{\star}_{i,t+l} \cap K^{\star}_{i,j,t,l})}{n-1}\alpha\right] + \delta^T \left[\frac{AVG\#(C_{i,t+T} \cap K_{i,j,t,T})}{n-1}(1+\beta) + \frac{AVG\#(\overline{C}_{i,t+T} \cap K_{i,j,t,T})}{n-1}\alpha\right]$$

for any $j \in N$ where $C^{\star}_{i,t+l}$ is the set of cooperators that i *believes* he will know about in period $t+l$.[6] This condition is hardest to satisfy in the set of histories in which i's defection was in the previous period, i.e. where $t^d = 1$. The condition then becomes

$$\alpha-1 \le \delta^T \left[\frac{AVG\#(C_{i,t+T} \cap K_{i,j,t,T})}{n-1}(1+\beta) + \frac{AVG\#(\overline{C}_{i,t+T} \cap K_{i,j,t,T})}{n-1}\alpha\right]$$

for any $j \in N$. In $t+T$, $C_{i,t+T} = A \backslash i$ and $\overline{C}_{i,t+T} = \varnothing$ so the condition reduces to:

$$\frac{\alpha-1}{\delta^T} \le \min\left\{\frac{AVG\#K_{i,j,t,T}}{n-1}(1+\beta),\ \frac{AVG\#K_{i,j,t,T}}{n-1}\alpha\right\} \tag{7}$$

for any $j \in N$.

To safeguard against (c1) defections, note that expected punishment from defecting via (c1) is the same as the expected punishment from defecting via (d1) when $t^d = T^p$, i.e. when the most recent defection was far enough in the past. The above makes clear that condition (7) is sufficient to prevent defections via (c1).

Defections via (d2) incur the same expected punishment but gain β from playing d rather than c against an opponent's d. The condition is then

$$\frac{\beta}{\delta^T} \le \min\left\{\frac{AVG\#K_{i,j,t,T}}{n-1}(1+\beta),\ \frac{AVG\#K_{i,j,t,T}}{n-1}\alpha\right\} \tag{8}$$

for any $j \in N$.

Deviations according to (c2), (d3) and (d4) are trivially not preferred. Deviating according to (c2) and (d3) would entail foregoing the gains from issuing punishment and in doing so, earn punishment in return. Deviating according to (d4) would entail foregoing the chance to not incur a loss during punishment and earning punishment as a consequence.

Given the way information spreads along g, $AVG\#K_{i,j,t,T} = \frac{1}{n}\sum_j \#N_j^{Tr}(g)$. Hence if conditions (7) and (8) are satisfied for all players, no player has an incentive to deviate in any history given any beliefs and so σ^{NWIGP} is sequentially rational. Since the binding conditions for sequential rationality are independent of beliefs, any consistent beliefs trivially extend the behavior to sequential equilibrium. □

[6] Strictly speaking, the condition as written holds for $t^d \ge 2$, and should be written without the first sum for $t^d = 1$.

A.4. Proof of Proposition 3

Enforcing full cooperation is supported by a lower minimum value of the discount factor δ under limited network knowledge than under compete network knowledge.

Proof. σ^{NWIGP} is sequentially rational under network knowledge environment *PRECISE* given r, T, and g iff

$$\delta^T \ge \frac{(\alpha-1)(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)} \tag{9}$$

and

$$\delta^T \ge \frac{\beta(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)}. \tag{10}$$

σ^{NWIGP} is sequentially rational under network knowledge environment *COARSE* given r, T, and g iff

$$\delta^T \ge \frac{(\alpha-1)(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)} \tag{11}$$

and

$$\delta^T \ge \frac{\beta(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}. \tag{12}$$

Restating the above conditions, the minimum discount factor that could support full cooperation under *PRECISE* is

$$\delta_{precisemin} = \left[\max\left\{\frac{(\alpha-1)(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)}, \frac{\beta(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)}\right\}\right]^{\frac{1}{T}}$$

and under *COARSE* is

$$\delta_{limitedmin} = \left[\max\left\{\frac{(\alpha-1)(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}, \frac{\beta(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}\right\}\right]^{\frac{1}{T}}.$$

For given α, β, g, r and T, if the left term is the maximum under precise network knowledge, the left term is the maximum under limited network knowledge, and vice versa. Since

$$\min_j\left\{\#N_j^{Tr}(g)\right\} \le \frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g),$$

$$\frac{(\alpha-1)(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)} \ge \frac{(\alpha-1)(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}$$

$$\Longleftrightarrow \left[\frac{(\alpha-1)(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)}\right]^{\frac{1}{T}} \ge \left[\frac{(\alpha-1)(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}\right]^{\frac{1}{T}}$$

which, if the left terms are the maximum, implies that

$\delta_{precisemin} \ge \delta_{limitedmin}$.

Likewise,

$$\min_j\left\{\#N_j^{Tr}(g)\right\} \le \frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)$$

$$\Longleftrightarrow \frac{\beta(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)} \ge \frac{\beta(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}$$

$$\Longleftrightarrow \left[\frac{\beta(n-1)}{(1+\beta)\left(\min_j\left\{\#N_j^{Tr}(g)\right\}\right)}\right]^{\frac{1}{T}} \geq \left[\frac{\beta(n-1)}{(1+\beta)\left(\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)\right)}\right]^{\frac{1}{T}}$$

which, if the right terms are the maximum, implies that

$\delta_{precisemin} \geq \delta_{limitedmin}$.□

The difference between $\delta_{precisemin}$ and $\delta_{limitedmin}$ is increasing in the difference between $\min_j\left\{\#N_j^{Tr}(g)\right\}$ and $\frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)$. $\delta_{precisemin} = \delta_{limitedmin}$ only when $\min_j\left\{\#N_j^{Tr}(g)\right\} = \frac{1}{n}\sum_{j\in N}\#N_j^{Tr}(g)$. This is only the case when either the network is a very rare $Tr-$*regular* graph, or when $Tr > diam(g)$. Otherwise, $\delta_{precisemin} > \delta_{limitedmin}$ and so full cooperation is strictly easier to enforce under limited network information. Any true coarsening that conceals the identity of the most tempting victim has the same consequence.

A.5. False gossip

A widely recognized drawback of punishing in response to gossip is the potential that the gossip contains false information. This concern arises in models in which myopic agents respond automatically to gossip they receive (Nakamaru and Kawata, 2004; Nowak and Sigmund, 2005; Ohtsuki et al., 2009; Nakamura and Masuda, 2011) and is supported by experimental evidence which shows that subjects respond to gossip even when they also possess verifiable information (Sommerfeld et al., 2007). The worry is that individuals could manipulate others through false gossip to the advantage of themselves, the disadvantage of others, and this gossip could lead to the unraveling of cooperation and the evolutionary success of uncooperative and or deceitful traits (though in some cases the presence of false gossip does not have pernicious consequences for cooperation in these models (Ghang and Nowak, 2015).)

A model of strategic agents allows identification of the precise incentives to manipulate others with false gossip, and also reveals options for incentivizing truth-telling in equilibrium.

The conditions above ensure the successful enforcement of full cooperation in sequential equilibrium, assuming that gossip, specifically the messages sent from victims of misbehavior to their neighbors in network g, is truthful. If instead this assumption was relaxed so that the players were allowed to strategically choose the content of their messages, a related set of strategies could enforce both full cooperation and truth-telling in equilibrium (related to a nested model to do so among non-strategic agents in Scheuring (2010)).

Relaxing the truth-telling assumption would admit two ways to issue false gossip: (1) falsely claim to have been the victim of a defection, or (2) withhold gossip when victimized by a defection. When strategic agents choose their actions, as they do in the above model, as well as the content of the gossip about their actions, which by assumption they do not do in the above model, manipulative lying is possible. However, only some instances of it would be profitable.

Lying via option (2) would not be individually profitable. Withholding information that an opponent j defected offers no gain to victim i. Spreading this gossip allows future punishment of j, withholding it prevents future punishment of j; preventing punishment of j offers no payoff gain to i.

Consider lies via option (1). These could only be profitable if the lies allow the liar to defect on the subject of the lies in the future with reduced punishment. Consider an example: i falsely claims that j defected. This has two consequences. One, those who act on the false message and randomly encounter j punish j; this is payoff-irrelevant to i. Two, if i then encounters j and defects on j, the lie would affect i's payoff if someone who would have punished i for this defection in response to j's message instead believes the action to have been punishment for j's supposed earlier defection and refrains from punishing i. If the set of individuals who refrain from punishing i is large enough, i could find lying profitable, and could prefer to behave uncooperatively.

The size of this set, and so the incentives i faces to lie, depends on i's network position (which affects how widely i's lie spreads), j's network position (which affects how widely j's message that i defected spreads), and what players who receive conflicting information do in response. Any player who receives both $i$aposs lie and j's later message that i defected can detect a conflict – if i was telling the truth, j's message cannot be correct; if j was telling the truth, i's message cannot have been correct.

In order to incentivize i to tell the truth, it must be that his expected gain from lying does not exceed the expected cost. For some parameter values, the chance of being able to defect against the subject of a lie in the near future and suffer sufficiently reduced consequences for it is too remote for the lie to be profitable, so i is automatically incentivized to tell the truth. If lying is tempting enough, though, one way to ensure truth-telling is to have players play strategies that respond negatively to the presence of conflicting information (which is consistent with observed human reactions to false gossipers (Kroupa, 2014). For instance, upon receiving conflicting information regarding players i and j, players could switch to a grim strategy of always play d.[7] So long as this will result in a prospective liar suffering enough in expectation from the unraveling of cooperation he would trigger with his lie relative to his expected gain from lying in the first place, cooperation and truth-telling could be supported in equilibrium and the results about network structure and knowledge would continue to hold.

This mechanism for incentivizing truth-telling nests the cooperation game within a truth-telling game. Of course real human agents may use a variety of methods in addition to or instead of socially sanctioning liars to deal with the risks of false gossip. Experimental studies show that individuals feel naturally compelled to share information about misbehavior with others (Feinberg et al., 2012), which may counteract an inclination to construct lies. Experiments also show that a greater volume of gossip may attenuate the negative affects of false gossip (Sommerfeld et al., 2008). Moreover, individuals may have developed psychological adaptations early on to suss out the veracity of gossip (Hess and Hagen, 2006). Whatever exact method they use to deal with liars, human groups can sustain cooperation with gossip that spreads through social networks.

References

Apicella, C.L., Marlowe, F.W., Fowler, J.H., Christakis, N.A., 2012. Social networks and cooperation in hunter-gatherers. Nature 481 (7382), 497–501.

Bernard, H.R., Killworth, P.D., Sailer, L., 1980. Informant accuracy in social network data IV: a comparison of clique-level structure in behavioral and cognitive network data. Social Netw. 2 (3), 191–218.

Bernard, H.R., Killworth, P.D., Sailer, L., 1982. Informant accuracy in social-network data V. An experimental attempt to predict actual communication from recall data. Social Sci. Res. 11 (1), 30–66.

Bickart, K.C., Wright, C.I., Dautoff, R.J., Dickerson, B.C., Barrett, L.F., 2011. Amygdala volume and social network size in humans. Nat. Neurosci. 14 (2), 163–164.

Bondonio, D., 1998. Predictors of accuracy in perceiving informal social networks. Social Netw. 20 (4), 301–330.

[7] This requires a condition on the presence of cycles in a network to ensure that for every possible false message, at least one person would receive both the false message and a conflicting message.

Bowles, S., 2006. Group competition, reproductive leveling, and the evolution of human altruism. Science 314 (5805), 1569–1572.
Bowles, S., 2009. Did warfare among ancestral hunter-gatherers affect the evolution of human social behaviors? Science 324 (5932), 1293–1298.
Boyd, R., Richerson, P.J., 1988. The evolution of reciprocity in sizable groups. J. Theor. Biol. 132 (3), 337–356.
Boyd, R., Richerson, P.J., 1989. The evolution of indirect reciprocity. Soc. Netw. 11 (3), 213–236.
Brandt, H., Sigmund, K., 2004. The logic of reprobation: assessment and action rules for indirect reciprocation. J. Theor. Biol. 231 (4), 475–486.
Brashears, M.E., 2013. Humans Use Compression Heuristics to Improve the Recall of Social Networks, Scientific Reports 3.
Brewer, D.D., Webster, C.M., 2000. Forgetting of friends and its effects on measuring friendship networks. Soc. Netw. 21 (4), 361–373.
Chalub, F.A., Santos, F.C., Pacheco, J.M., 2006. The evolution of norms. J. Theor. Biol. 241 (2), 233–240.
Conley, T., Udry, C., 2010. Learning about a new technology: pineapple in Ghana. Am. Econ. Rev. 100 (1), 35–69.
De Soto, C.B., 1960. Learning a social structure. J. Abnorm. Soc. Psychol. 60 (3), 417.
De Soto, C.B., Henley, N.M., London, M., 1968. Balance and the grouping schema. J. Personal. Soc. Psychol. 8 (1p1), 1.
Dixit, A.K., 2004. Lawlessness and Economics: Alternative Modes of Governance. Princeton University Press, Princeton.
Duflo, E., Saez, E., 2002. Participation and investment decisions in a retirement plan: the influence of colleagues' choices. J. Public Econ. 85 (1), 121–148.
Dunbar, R.I., 1992. Neocortex size as a constraint on group size in primates. J. Hum. Evol. 22 (6), 469–493.
Dunbar, R., 1998. Grooming, Gossip, and the Evolution of Language. Harvard University Press, Cambridge.
Dunbar, R.I., 2004. Gossip in evolutionary perspective. Rev. Gen. Psychol. 8 (2), 100.
Enquist, M., Leimar, O., 1993. The evolution of cooperation in mobile organisms. Anim. Behav. 45 (4), 747–757.
Ertan, A., Page, T., Putterman, L., 2009. Who to punish? individual decisions and majority rule in the solution of free rider problems. Eur. Econ. Rev. 3, 495–511.
Eshel, I., Cavalli-Sforza, L.L., 1982. Assortment of encounters and evolution of cooperativeness. Proc. Natl. Acad. Sci. 79 (4), 1331–1335.
Fearon, J., Laitin, D., 1996. Explaining interethnic cooperation. Am. Polit. Sci. Rev. 90 (4), 715–735.
Fehr, E., Gächter, S., 1999. Cooperation and Punishment in Public Goods Experiments, Institute for Empirical Research in Economics. Working Paper 10.
Fehr, E., Gächter, S., 2002. Altruistic punishment in humans. Nature 415 (6868), 137–140.
Feinberg, M., Willer, R., Stellar, J., Keltner, D., 2012. The virtues of gossip: reputational information sharing as prosocial behavior. J. Personal. Soc. Psychol. 102 (5), 1015.
Fishman, M.A., 2003. Indirect reciprocity among imperfect individuals. J. Theor. Biol. 225 (3), 285–292.
Freeman, L.C., 1992. Filling in the blanks: a theory of cognitive categories and the structure of social affiliation. Soc. Psychol. Q., 118–127.
Ghang, W., Nowak, M.A., 2015. Indirect reciprocity with optional interactions. J. Theor. Biol. 365, 1–11.
Gluckman, M., 1963. Gossip and scandal. Curr. Anthropol. 4 (3), 307–316.
Granovetter, M., 1973. The strength of weak ties. Am. J. Sociol., 1360–1380.
Greif, A., 1993. Contract enforceability and economic institutions in early trade: the Maghribi traders' coalition. Am. Econ. Rev. 83 (3), 525–548.
Gugerty, M.K., 2007. You can't save alone: commitment in rotating savings and credit associations in Kenya. Econ. Dev. Cult. Change 55 (2), 251–282.
Hess, N.H., Hagen, E.H., 2006. Psychological adaptations for assessing gossip veracity. Hum. Nat. 17 (3), 337–354.
Janicik, G.A., Larrick, R.P., 2005. Social network schemas and the learning of incomplete networks. J. Personal. Soc. Psychol. 88 (2), 348.
Kanai, R., Bahrami, B., Roylance, R., Rees, G., 2011. Online social network size is reflected in human brain structure. Proc. R. Soc. B: Biol. Sci., rspb20111959
Kandori, M., 1992. Social norms and community enforcement. Rev. Econ. Stud. 59 (1), 63.
Karlan, D., Mobius, M., Rosenblat, T., Szeidl, A., 2009. Trust and social collateral*. Q. J. Econ. 124 (3), 1307–1361.
Killworth, P.D., Bernard, H.R., 1976. Informant accuracy in social network data. Hum. Organ. 35 (3), 269–286.
Krackhardt, D., 1992. The Strength of Strong Ties: The Importance of Philos.
Krackhardt, D., Hanson, J., 1993. Informal networks: the company behind the charts. Harv. Bus. Rev. 71, 104–112.
Krackhardt, D., Kilduff, M., 1999. Whether close or far: social distance effects on perceived balance in friendship networks. J. Personal. Soc. Psychol. 76 (5), 770.
Kroupa, S., 2014. Love or fear: can punishment promote cooperation?. Evolut. Anthropol.: Issues News Rev. 23 (6), 229–240.
Kumbasar, E., Rommey, A.K., Batchelder, W.H., 1994. Systematic biases in social perception. Am. J. Sociol., 477–505.
Larson, J., 2016. Networks and Interethnic Cooperation. Working Paper.
Larson, J.M., Lewis, J.I., 2016. Ethnic Networks. Working Paper.
Meyer, M.L., Spunt, R.P., Berkman, E.T., Taylor, S.E., Lieberman, M.D., 2012. Evidence for social working memory from a parametric functional mri study. Proc. Natl. Acad. Sci. 109 (6), 1883–1888.
Miguel, E., Gugerty, M., 2005. Ethnic diversity, social sanctions, and public goods in kenya. J. Public Econ. 89 (1112), 2325–2368.
Mohtashemi, M., Mui, L., 2003. Evolution of indirect reciprocity by social information: the role of trust and reputation in evolution of altruism. J. Theor. Biol. 223 (4), 523–531.
Morris, A.J., 1997. Perceptions of Social Structure, Others, and Self: The Role of Social Intelligence in Managerial Effectiveness (Ph.D. thesis). University of California, Irvine.
Nakamaru, M., Kawata, M., 2004. Evolution of rumours that discriminate lying defectors. Evol. Ecol. Res. 6 (2), 261–283.
Nakamura, M., Masuda, N., 2011. Indirect reciprocity under incomplete observation. PLOS Comput Biol 7.7, e1002113. http://dx.doi.org/10.1371/journal.pcbi.1002113.
Nowak, M.A., Sigmund, K., 1998. Evolution of indirect reciprocity by image scoring. Nature 393 (6685), 573–577.
Nowak, M.A., Sigmund, K., 2005. Evolution of indirect reciprocity. Nature 437 (7063), 1291–1298.
Nowak, M.A., Tarnita, C.E., Wilson, E.O., 2010. The evolution of eusociality. Nature 466 (7310), 1057–1062.
Ohtsuki, H., Iwasa, Y., 2007. Global analyses of evolutionary dynamics and exhaustive search for social norms that maintain cooperation by reputation. J. Theor. Biol. 244 (3), 518–531.
Ohtsuki, H., Hauert, C., Lieberman, E., Nowak, M.A., 2006. A simple rule for the evolution of cooperation on graphs and social networks. Nature 441 (7092), 502–505.
Ohtsuki, H., Iwasa, Y., Nowak, M.A., 2009. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. Nature 457 (7225), 79–82.
Pacheco, J.M., Santos, F.C., Chalub, F., 2006. Stern-judging: a simple, successful norm which promotes cooperation under indirect reciprocity. PLoS Comput. Biol. 2 (12), e178.
Picek, J.S., Sherman, S.J., Shiffrin, R.M., 1975. Cognitive organization and coding of social structures. J. Personal. Soc. Psychol. 31 (4), 758.
Rao, N., Mobius, M., Rosenblat, T., 2007. Social Networks and Vaccination Decisions. Federal Reserve Bank of Boston. Working Papers.
Sallet, J., Mars, R., Noonan, M., Andersson, J., Oreilly, J., Jbabdi, S., Croxson, P., Jenkinson, M., Miller, K., Rushworth, M., 2011. Social network size affects neural circuits in macaques. Science 334 (6056), 697–700.
Scheuring, I., 2010. Coevolution of honest signaling and cooperative norms by cultural group selection. BioSystems 101 (2), 79–87.
Simpson, B., Markovsky, B., Steketee, M., 2011. Power and the perception of social networks. Soc. Netw. 33 (2), 166–171.
Sommerfeld, R., Krambeck, H., Semmann, D., Milinski, M., 2007. Gossip as an alternative for direct observation in games of indirect reciprocity. Proc. Natl. Acad. Sci. 104 (44), 17435.
Sommerfeld, R.D., Krambeck, H.-J., Milinski, M., 2008. Multiple gossip statements and their effect on reputation and trustworthiness. Proc. R. Soc. London B: Biol. Sci. 275 (1650), 2529–2536.
Stiller, J., Dunbar, R.I., 2007. Perspective-taking and memory capacity predict social network size. Social Netw. 29 (1), 93–104.
Van Veelen, M., García, J., Rand, D.G., Nowak, M.A., 2012. Direct reciprocity in structured populations. Proc. Natl. Acad. Sci. 109 (25), 9929–9934.
Wedekind, C., Milinski, M., 2000. Cooperation through image scoring in humans. Science 288 (5467), 850–852.
Wilson, D.S., Wilczynski, C., Wells, A., Weiser, L., 2000. Gossip and other aspects of language as group-level adaptations. Evol. Cognit., 347–365.
Wolitzky, A., 2013. Cooperation with network monitoring. Rev. Econ. Stud. 80 (1), 395–427.
Zahn, R., Moll, J., Krueger, F., Huey, E.D., Garrido, G., Grafman, J., 2007. Social concepts are represented in the superior anterior temporal cortex. Proc. Natl. Acad. Sci. 104 (15), 6430–6435.