Cooperation in Partly Observable Networked Markets*

Itay P. Fainmesser† David A. Goldberg‡

July 12, 2012

Abstract

We present a model of repeated games in large two-sided networks between clients and agents in the presence of reputation networks via which clients share information about past transactions. The model allows us to characterize *cooperation networks* - networks in which each agent cooperates with every client that is connected to her. To this end, we show that: [1] the incentives of an agent a to cooperate depend only on her beliefs with respect to her local neighborhood - a subnetwork that includes agent a and is of a size that is independent of the size of the entire network; and [2] when an agent a observes the network structure only partially, the incentives of a to cooperate can be calculated *as if* the network was a random tree with agent a at its root. Our characterization sheds light on the welfare costs of relying only on repeated interactions for sustaining cooperation, and on how to mitigate such costs.

Keywords: Networks, trust, graph theory, repeated games.

1 Introduction

In many markets, successful execution of mutually beneficial economic transactions relies on informal contracts that are enforced by *social pressure* and *reputation*.[1] Informal enforcement

*We are grateful to Susan Athey, Drew Fudenberg, and Al Roth, for many discussions and invaluable comments. We also benefited from the help of Ilan Lobel and Theo Weber, and from the useful comments of Nageeb Ali, James Burns, Ben Golub, Roberto Serrano, Adrien Vigier, and seminar participants at Harvard University. This paper was significantly revised during a summer that Fainmesser spent in Microsoft Research New England. The paper was previously circulated under the title: "Bilateral and Community Enforcement in a Networked Market with Simple Strategies." A previous version is available on Fainmesser's website.

†Department of Economics, Brown University, Providence, RI 02906. E-mail: *Itay_Fainmesser@Brown.edu*

‡H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, E-mail: *dave.goldberg@isye.gatech.edu*

[1] Macaulay (1963) points out that social pressure and reputation are perhaps more widely used than formal contracts and enforcement.

mechanisms include *personal* and *community enforcement* mechanisms.[2] It is by now widely recognized that when transactions between two parties are sufficiently frequent, personal enforcement is highly effective.[3] Community enforcement can overcome the limitations of personal enforcement when transactions between two parties are infrequent, yet transactions in the population are frequent. In large markets, community enforcement is effective if third-party observability is available.[4] However, despite the abundance of research on repeated games and community enforcement, the frequency of interactions, as well as the level of third-party observability are mostly treated as 'black boxes' or modeled for highly specialized cases. For example, in much of the literature, either any two parties interact in every period or random matching is assumed. In contrast, it is well known that in many two-sided markets with buyers and sellers, or investors and entrepreneurs, each "client" has access to a different subset of the "agents" in the market, and chooses to interact with even a smaller subset. Moreover, each client often learns about the outcomes of a different subset of the interactions in the market.[5]

To address these issues, we develop a model of repeated interactions in networked markets with clients and agents. Consider groups of agents and clients, each with limited capacities. Initially, each client can interact with (e.g. purchase a good from, or make a loan to) only a subset of agents to whom he is *connected*. The initial connections between clients and agents define a two-sided *interaction network* G^0. Clients can also decide to eliminate their connections with agents who they do not trust. As a result, the interaction network may evolve over time. In every period, agents meet sequentially with clients who are connected to them and decide whether to cooperate or defect. The interaction outcome between an agent a and a client c is observable to a subset of the clients (including c); such clients are said to be *connected* to client c. The connections between clients define a *reputation network* R that captures the level of third-party observability in the market. The combined

[2] In personal enforcement mechanisms cheating triggers retaliation by the victim, whereas in community enforcement mechanisms dishonest behavior against one partner causes sanctions by several members in the society.

[3] For a good survey on long term relationships see Mailath and Samuelson (2006).

[4] See Kandori (1992), Greif (1993), and Ellison (1994).

[5] The economic literature offers extensive evidence for the presence of networks of cooperation and trust within markets. For example, see Fafchamps (1996), McMillan and Woodruff (1999), Hardle and Kirman (1995), Kirman and Vriend (2000), Weisbuch et al. (1996), and Karlan et al. (2009).

network, $N = (G, R)$ captures the market structure. We then ask the following questions for any fixed level of agents' patience: what structures of the network N can be sustained indefinitely in equilibria in which all interactions end in cooperation? For what structures of the interaction network G there exists a reputation network R such that $N = (G, R)$ can be sustained indefinitely and allows for an equilibrium with full cooperation? What is the *optimal* network structure that allows for the maximal number of mutually beneficial interactions, and can we do better with formal contracts? The answers define a set of network in which a connection between client c and agent a has the meaning that c is able and willing to interact with a, and that a always cooperates with c.

Analyzing a model of repeated games in networks poses several difficulties. In particular, the incentives of an agent to cooperate depend on the entire network structure, as well as on the strategies of all clients and agents in the market. This problem is exacerbated when the interaction network evolves over time and when the reputation network may be different from the interaction network. Another source of difficulty is that each agent can serve a limited number of clients in every period, and each client has demand for a limited number of services in every period. This implies that even on the path of a cooperative equilibrium, the future payoff of an agent depends on the entire network structure (as opposed to only the off path payoff). We alleviate these difficulties in two steps.

At the core of our methodological contribution is a new method for reducing questions about the global properties of a network (e.g. characterizing payoffs that depend on the entire network) to questions about the local properties of the network. This allows us to provide conditions under which the incentives of an agent a to cooperate with client c depend only on her beliefs with respect to her *local neighborhood* - a subnetwork that includes agent a and is of a size that is independent of the size of the entire network (Theorem 1). Thus, we are able to analyze large networks *as if* they were small. To derive these 'local conditions' we make use of recent results in the graph theory literature by Gamarnik and Goldberg (2010) - hereby GG - who study a simple matching algorithm and ask the following question: "When does the performance of the algorithm depend only on the local properties of the network?" By relating our market dynamics to the dynamics of the same algorithm, we are able to use the tools developed in GG to analyze market equilibria.[6]

A second step along the same lines proposes a natural model of beliefs over the network structure that captures the idea that: [1] there is a strong random component in the formation of networks, and [2] each agent (client) knows more about her immediate neighborhood than about the rest of the network. For this model of beliefs, we prove that if the network N is large and all other agents always cooperate, then the incentives of agent a to cooperate in N can be approximated by the incentives of a to cooperate in a *simpler network* - a random tree with known distributions over the numbers of connections of clients and agents in the network. The approximation improves as the network grows and the error goes asymptotically to zero (Theorem 2). This result is based on a key graph theoretic lemma that we prove: consider a large *bipartite graph* G that is chosen *uniformly at random* (u.a.r.) conditional on the (finite support) distributions of the number of links attached to nodes in the graph, then G is asymptotically *locally* like a random tree (Lemma 3). Although results of a similar flavor are known in the random-graph community (see Wormald 1999 and references therein), they have not received attention in the economics literature.[7]

We focus on equilibria in ostracizing strategies: on the equilibrium path, agents always cooperate and clients cut their links with (and only with) agents who they observe to defect. Using our random tree characterization, we provide conditions under which an asymptotically large network N can be sustained indefinitely and facilitate cooperation in all of the interactions in the network. As expected, we find that adding a large number of links to a reputation network R increases the set of interaction networks G that can be sustained indefinitely and facilitate full cooperation (Proposition 1). When R is sufficiently dense we find that networks in which there are fewer agents, each having more connections (in G), and more clients, each having fewer connections (in G), can be sustained indefinitely and

[6] The question of when the global properties of a graph are determined by the graph's local structure (a.k.a. a decay of correlation phenomenon) has a long history in the graph theory literature. E.g. Lauer and Wormald (2007) and Goring et al. (2009) study localization phenomenon for greedy algorithms, and Bayati et al. (2007) study the correlation decay phenomenon for the matchings of a graph. We also note that several recent works study related questions pertaining to when an (approximate) Nash equilibrium can be computed in a distributed / local manner using the tools of correlation decay, see e.g. Weber (2010) and Kanoria et al. (2010).

[7] An exception is Campbell (2011) who applies percolation theory (physics) to the study of monopoly pricing in the presence of WOM. This is related because percolation theory relies on insights that are directly related to the claim we prove in Lemma 3.

facilitate cooperation for a larger set of discount factors (Proposition 2). The implications of the latter are quite surprising. Consider an agent a and an agent a' that are parts of networks N and N' respectively. Suppose further that the immediate neighborhoods of a and a' are identical – i.e. they are connected to the same numbers of clients, each connected to the same number of agents. Now assume that the distribution of the numbers of connections in N and N' are such that agents (clients) are more (less) connected in N than in N'. Then, the expected on path payoff of a is higher than of a', and if the reputation network is sufficiently dense then the incentives of a to cooperate are stronger than the incentives of a' to cooperate.

Finally, we show that a sufficiently dense reputation network R guarantees that the fraction of interactions lost due to the incentive constraints goes to zero as the size of the market grows to infinity (Corollary 2). This is despite the fact that the optimal network that facilitates full cooperation achieves only a fraction (< 1) of the number of interactions that formal contracts could achieve in any finite market with significantly more agents than clients - a fact that is driven by the observation that in every network that facilitates full cooperation some agents are permanently excluded from the market.

Networks have been used to model market structure in many recent works in economics (see Jackson 2008 and Goyal 2007 for extensive surveys). When considering games in networks, much of the previous work analyzes static network games (e.g. Galeotti et al. 2010, Ballester et al. 2006, and Bramoullé, D'Amours, and Kranton 2010). In static network games a player's payoff depends only on the actions taken by her immediate neighbors. As a result, beliefs on the network structure are used by a player only to establish a prior over the actions that her neighbors will take, and Galeotti et al. (2010) find that assuming that a player has incomplete knowledge of the network structure simplifies the analysis. In contrast, in our framework, due to the dynamic nature of the interactions and the limited capacities, an agent's incentives to cooperate generally depend on the entire network structure and on the actions taken by all of the clients and agents in the network. In large markets this makes the problem prohibitively complex, and prior to establishing Lemma 3 there was no reason to expect that assuming incomplete knowledge of the network structure simplify the problem. In fact, one would expect that an agent who has incomplete knowledge of the network is required to compute her incentives in every network that has positive probability according to her prior.

This paper is also related to recent developments in the study of repeated games in networks (see Vega-Redondo 2006, Kinateder 2008, Lippert and Spagnolo 2006, Mihm et al. 2009, Jackson et al. 2011, Ali and Miller 2012, and Nava and Piccione 2012) and to the literature on trust and social collateral (see Karlan et al. 2009). An important difference from this literature is that we separate the analysis of the trade network from that of the communication network and allow both to vary in economically meaningful ways.

We note that the source and degree of third-party observability among traders in real-world markets varies widely across contexts.[8] We propose our notion of a reputation network as a reduce form that captures Word-Of-Mouth, reputation systems, or any other form of observability. A similar approach is taken by Balmaceda and Escobar (2011) who study a related enforcement problem in a market with one agent and many clients. They ask what reputation networks can sustain cooperation by the agent with all of the clients.

Finally, an application of the methodology developed in this paper can be found in Fainmesser (2012a). In particular, Fainmesser (2012a) characterizes the patterns of interactions in markets with buyers and sellers of experience goods in which third-party observability is nonexistent (in the language of our model: R is the empty reputation network), such as markets in which strategic considerations prevent the diffusion of information. In such markets, incentives for the provision of high quality goods require that networks be *sparse* and that there is a similar number of buyers and sellers with similar (and low) degrees. We present a more general framework and find that the limits on cooperation under *any* patterns of third-party observability are different. In particular, we are able to show that in the presence of third-party observability, *dense* networks facilitate cooperation better. The comparison between the results in the two papers provides insights on the limits of cooperation and on the types of inefficiencies that can (and those that cannot) be circumvented by improving observability in a market.

The following section offers two motivating examples. Section 3 follows with a model of a networked market, and the notion of a Totally Cooperative strict Bayes-Nash Equilibrium

[8] See also Esfahani and Salehi-Isfahani (1989) and Banerjee and Newman (1993).

with Ostracizing strategies (TCEO) is defined in section 4. In section 5, we derive our first main result and provide conditions under which the incentives of an agent to cooperate depend only on her local network structure. In section 6 we propose a specific model of beliefs with respect to the network structure, and in section 7 we characterize the structure of cooperation networks in this model. The welfare implications of our results are derived in section 8. Section 9 offers a discussion of the main methodological contributions of the paper, and suggests additional economic implications. Section 10 offers concluding remarks.

2 Motivating examples

To motivate our analysis, we briefly describe two examples of relevant applications.

Investors and entrepreneurs Consider a group of investors and a group of entrepreneurs who come up with risky investment opportunities over time. To realize her investment opportunity, an entrepreneur needs to take a loan which she might not be able to repay if the investment does not succeed. The realization of the investment is observable to the investor, but at the same time, the realization of the investment may not be verifiable, and the entrepreneur can choose to strategically default on her loan (see also Fainmesser 2012c). It is also reasonable to assume that entrepreneurs have a limited number of investment opportunities to offer, and investors have liquidity constraints. As a result, the patterns of interactions between investors and entrepreneurs play an important role: if the frequency of interactions between an investor and an entrepreneur is high enough, and if the outside option for the entrepreneur is low enough, a threat of punishment from the investor alone may provide the entrepreneur with the incentives to never strategically default on a loan. Otherwise, if investors have a way to share information credibly, or if strategic default is observable by other investors, a threat of ostracizing a defaulting entrepreneur may help to provide the appropriate incentives. In this paper we study the relationship between the observability network and the lending network between investors and entrepreneurs in enforcing repayment.

Experience goods Consider a group of buyers and a group of sellers of experience goods (e.g. services). A seller can decide whether to supply high or low quality goods, and may even provide some buyers with high quality goods and other with low (see also Kirman and Vriend 2000 and Fainmesser 2012a). If providing high quality goods costs more than providing low quality, then in the absence of sufficient future payoffs that are contingent on providing high quality, a seller may provide low quality. This paper studies an environment with many sellers and many buyers, and consider the possibility that sellers and buyers might not have complete knowledge of the patterns of interaction in the market.

3 The model

Consider a market with a set of clients $C \equiv \{1, 2, ..., n_c\}$ and a set of agents $A \equiv \{1, 2, ..., n_a\}$. Time is discrete ($t = 1, 2, ..., \infty$). Clients and agents live forever and have a common discount factor δ. In a given period, each client (agent) has the capacity to engage with one agent (client) in the following simple trust-based interaction with an outcome that depends only on the action of the agent. If the agent defects, the agent has a positive payoff of π and the client has a negative payoff of $-\varphi$. If the agent cooperates, the agent has a positive (but lower) payoff of $\pi - \gamma$ and the client has a positive payoff of β for some $\beta > \gamma$. If a client (agent) does not engage in any interaction in a given period, her payoff is zero.

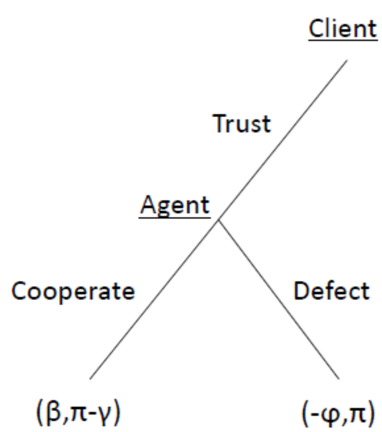

Figure 1: a trust-based interaction.

3.1 Interaction networks

The patterns of interaction in the market (i.e. which client interacts with which agent in every period) are determined by exogenous factors (i.e. which agents each client is able to interacts with) as well as clients' decision (which agents each client trusts). More specifically, in any period t, clients and agents are connected via a two-sided network of connections between clients and agents. Intuitively, a connection between a client and an agent implies that the client is *able and willing* to interact with the agent. We first introduce the notion of a network and then make the economic notion of a connection more precise by describing the patterns of interactions in the market given a network, as well as rules governing the evolution of the network structure over time. We conclude this section with a discussion of the assumptions imposed by the structure of the game.

We model the network using a bipartite graph $G^t = \langle C, A, E^t \rangle$ where $E^t \subseteq C \times A$ is a set of client-agent pairs such that $(c, a) \in E^t$ if and only if there is an edge (or link) connecting client c and agent a in period t. We omit the superscript t when clear from the context. A *node* is an individual (client or agent) in the graph. A *path of length l in G* between node v and node v' is a sequence of edges $\{(v_0, v_1), (v_1, v_2), ..., (v_{l-1}, v_l)\}$ such that $v_0 = v$, $v_l = v'$, and for every $i \in \{1, 2, .., l\}$, $(v_{i-1}, v_i) \in G$. We say that the distance between v and v' in G is l if the length of the shortest path between node v and node v' in G equals l. For a given node v, let $N_1(v)$ be the set of nodes connected to v, let $d_v \equiv |N_1(v)|$ denote the degree (number of neighbors) of v in G. Similarly, let $N_2(v)$ denote the set that includes the set of nodes in $N_1(v)$ as well as the set of nodes connected to the nodes in $N_1(v)$. More generally, a node v' is in $N_d(v)$ if and only if the distance in G between v and v' does not exceed d. A *cycle* is a path $\{(v_0, v_1), (v_1, v_2), ..., (v_{l-1}, v_l)\}$ such that $v_0 = v_l$. A graph that has no cycle is also called a *tree*. A *rooted tree* is a tree in which one node is marked as the root. A node in a tree is called a *leaf* if its degree equals 1. The *depth* of a rooted tree is the largest distance between the root and any of the leafs in the tree. A node v is called a child of a node v' in a rooted tree if v and v' are connected AND v is at a larger distance from the root than v'.

We now describe the patterns of interaction given a network. We defer the discussion of what clients and agents observe with respect to the outcomes of bilateral interactions in the market and with respect to the network structure to sections 3.2 and 3.3 respectively. During period t, all connected clients and agents meet at a random sequencing – all of the links in E^t are ordered uniformly at random (u.a.r.) and then the links are chosen one by one according to that order.[9] When a link (c, a) is chosen, c and a meet and engage in

[9] Our analysis is independent of whether clients and agents learn the order after it is chosen.

the aforementioned trust-based interaction unless either c or a has already interacted (with anyone else) during the same period.

The network evolves over time in the following way. Before period 1, there is an initial network G^0. Subsequently, at the beginning of any period t, before any interaction takes place, clients make simultaneously the following decision: each client decides which of his connections (edges) he keeps and which connections he deletes permanently – the resulting network is G^t. Thus, clients can decide who they interact with by affecting the structure of the network, i.e. by deciding who they trust.

Several restrictions are imposed on the patterns of interaction by the structure of the game. First, an agent cannot decline an interaction, but rather only cooperate or defect. Allowing agents to refuse an interaction increases the complexity of the analysis without adding much insight. More specifically, there can be two reasons for an agent to refuse to interact: [1] in order to meet a client who she is planning to cheat; or [2] as a part of a collusive strategy with other agents. The earlier changes the timing of cheating, rather than introduce a new set of networks that can sustain permanent cooperation. The latter is unreasonable in our setup due to the large numbers of clients and agents and the incomplete knowledge of the network that we discuss below.

Second, we do not consider the formation of new links, but only the dissolution of links. This embodies the idea that the formation of new relationships is a longer-term process, and that the decision to cooperate and/or punish an agent (by disconnecting a trust relationship) can be taken more quickly (see Jackson et al. 2011 for a similar assumption in the context of favor exchange, as well as Fainmesser 2012c for a similar assumption in the context of financial lending networks). It is important to note that we do cover the case where the market starts with the initial network G^0 being the complete network as well as any other network, so we do not a priori restrict the links that might be formed, and so our results do make predictions about which networks can be sustained in a market.[10] The important restriction is that an agent who has lost a relationship cannot (quickly) replace it with a newly formed one.

In the same spirit, once a client disconnects from an agent the relationship cannot be

[10] In the complete network all of the clients are connected to all of the agents in the economy.

revived. This captures the idea that trust is more easily broken than restored. When we introduce the notion of reputation network below, it will become clear that this simplifies our analysis in that it eliminates complicated forms of punishment that take into account the possibility that an agent defects in order to make sure that she meets certain clients in the following period. Thus introducing via the back door the idea that an agent might choose not to interact with a given client in a given period.

Finally, we did not discuss the formation of the initial network G^0. Instead we take a different approach; we consider any initial network and ask whether it can be sustained, and sustain full cooperation indefinitely.

3.2 Reputation networks

Each client has access to information about the outcomes of all of his past transactions, as well as limited information about other clients' past transactions which he learns through his *reputation network* (e.g. via Word-Of-Mouth, reputation systems, or other mechanisms for third-party observability). Formally, there is a (reputation) graph R on the set of clients, where edge (c, c') is present in R if and only if client c is informed when any agent a defects in an interaction with client c'. Without loss of generality, we assume that the graph R is undirected (and thus $(c, c') \in R$ if and only if $(c', c) \in R$). Let $R(c, c')$ denote the indicator for whether or not $(c, c') \in R$. We also let $N_1^R(c) \triangleq \bigcup_{c':R(c,c')=1} c'$ denote the set of clients whose past transactions client c observes (or is informed of credibly), and $E_R^a(c) \triangleq \bigcup_{c' \in N_1^R(c) \cap N_1(a)} (a, c')$ denote the set of edges (in G) between agent a and clients who observe past transactions of client c.

3.3 The network structure - knowledge and beliefs

We now develop a general framework of clients' and agents' beliefs about the network structure. Note that an individual's *global beliefs* may be complicated, while her *local beliefs* may be more tractable.[11] Therefore, we require that our framework be flexible enough to allow

[11] By *global beliefs*, we generally refer to an individual's beliefs about the potential interactions between clients and agents separated from her in the network by a distance *on the same order* as the entire network. By *local beliefs*, we generally refer to her beliefs about the potential interactions between clients and agents separated from her in the network by a distance that is some small constant (e.g. 20) whose order is much smaller than that of the entire network, which may be arbitrarily large.

for the study of the relative importance of local and global beliefs in calculating clients' and agents' expected payoffs. Later we also suggest one particularly natural model in which individuals have incomplete knowledge of the network structure and stochastic beliefs with respect to their missing information.

Consider an individual v that can be a client or an agent. Individual v assigns some probability distribution $\mathcal{D}_v$ on the set of all possible networks $\mathcal{N}$, where a member of the set $\mathcal{N}$ is specified by a 2-vector $N = (G, R)$ consisting of both an interaction graph G and a reputation graph R. We call $\mathcal{D}_v$ a *belief* of individual v.

We now make the notion of local beliefs more formal. Recall that a node v' is in $N_d(v)$ if and only if the length of the shortest path in G between v and v' does not exceed d. For a given belief $\mathcal{D}_v$, we let $\mathcal{D}_v^d$ denote the distribution induced by $\mathcal{D}_v$ on $N_d(v)$. For example, if $\mathcal{D}_a$ (the belief of agent a) places probability 1 on the leftmost network in figure 2, then $\mathcal{D}_a^3$ would place probability 1 on the rightmost network in figure 2.

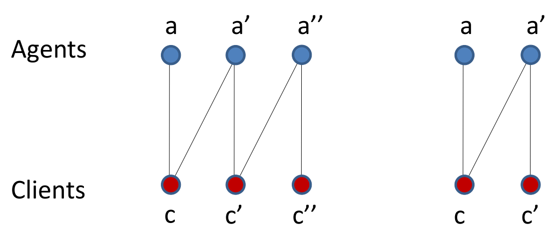

Figure 2: if $\mathcal{D}_a$ places probability 1 on the leftmost network, then $\mathcal{D}_a^3$ would place probability 1 on the rightmost network.

Finally, for simplicity, we assume throughout that agents' beliefs are stationary, and that agents do not update their beliefs on the network structure $\left(\{\mathcal{D}_a\}_{a \in A}\right)$.[12] While clearly restrictive, we believe that the analysis of repeated games with fixed beliefs on the network structure is an important first step and that our results are qualitatively robust. The analysis of learning the network structure through repeated interactions is beyond the scope of this paper and is left for future research.[13] In addition, Fainmesser (2012a) proposes an example of a network generating process in which small changes to the network structure over time

[12] Our analysis goes through without change regardless of whether clients update their beliefs or not.

[13] For a step in this direction, see Fainmesser (2012c) who studies repeated lending in financial networks with intermediaries and derives an upper bound on players' knowledge of the network when players learn about the network structure *only* from their financial interactions.

prevent agents from learning the exact network structure beyond certain comparative statics. A special case of the general framework that is consistent with the example in Fainmesser (2012a) is the GF model presented in section 6.

4 Equilibrium

Much of the literature on community enforcement achieves folk theorems using strategies that involve *contagion*.[14] A contagion strategy is a strategy in which an individual that was defaulted against defaults in all of her future interactions, regardless of the party with whom she interacts. While raising significant interest in the game theoretic literature, much of empirical evidence from economics, social psychology, anthropology, and even biology points towards the more common practice of *ostracism*.[15] An ostracizing strategy requires any individual that observes a defection to defect in future interactions with the initial defector. Ostracism is especially appealing in the large two-sided markets that motivate this paper. In such market, it is not clear what contagion entails, e.g. can a client spread contagion by avoiding the market? On the other hand, ostracism is intuitive; if a client observes a defection by an agent, the client stops to trust the agent and avoids interactions with that agent. To this end, we focus on Totally Cooperative strict Bayes-Nash Equilibria with Ostracism defined as follows.

Definition 1 *We say that client c uses* ***ostracizing strategies*** *if at any period t and for any connection $(c, a) \in E^t$, client c eliminates (c, a) at the beginning of period $t+1$ if and only if agent a defected in period t in an interaction with c or with any client c' who is connected to c in the reputation network $c' \in N_1^R(c)$.*

Definition 2 *A* ***Totally Cooperative strict Bayes-Nash Equilibrium with Ostracism (TCEO)*** *is a strict Bayes-Nash equilibrium in which all clients employ ostracizing strategies and all agents always cooperate.*[16]

[14] See also Kandori (1992), Ellison (1994), and Ali and Miller (2012).

[15] See also Durkheim (1933), Gordon (1975), Francis (1985), Goodall (1986), Gruter and Masters (1986), Lancaster (1986), Mahdi (1986), Greif (1993), Boehm (1999), Kurzban and Leary (2001), Williams (2001), Gaspart and Seki (2003), Wiessner (2005), and Karlan et al. (2009).

[16] In a strict Nash equilibrium, all players play a strict best response.

Focusing on TCEOs simplifies the analysis. To see how, note that a strategy of an agent a must specify the action taken by a in any period t in an interaction with any client $c \in N_1(a)$ as a function of agent a's beliefs with respect to the network structure ($\mathcal{D}_a$) and the entire history of play observed by a. The history of play can in turn depend on the entire network structure at any period $\tau < t$ ($\{G^\tau\}_{\tau=0}^{t-1}$ and R). Similarly, a strategy of a client c specifies the edges maintained by c at any period t as a function of client c's beliefs with respect to the network structure ($\mathcal{D}_c$) and the entire history of play observed by c. On the other hand, Lemma 1 implies that in a TCEO agent a's best response depends only on a's belief with respect to the network structure ($\mathcal{D}_a$). Clients' best responses follow immediately from agents' cooperation. Consequently, we can ask the following question: "Does a TCEO exist in an economy with γ, π, δ, and $\{\mathcal{D}_a\}_{a\in A}$)?"

Consider an vector $\mathbf{m} \triangleq (\gamma, \pi, \delta, \{\mathcal{D}_a\}_{a\in A})$ and let $\mathbf{m}_a \triangleq (\gamma, \pi, \delta, \mathcal{D}_a)$ be such that $\mathcal{D}_a$ puts probability 1 on the network $N = (G, R)$. Suppose that all agents $a' \neq a$ always cooperate and all clients use ostracizing strategies, and let u_N^{coop} be the expected present discounted value of all future payoff of agent a from the strategy "always cooperate." Similarly, let $\overline{u}_N$ be the expected present discounted value of all future payoff of agent a from her optimal strategy. Now consider any belief $\mathcal{D}_a$ and let $E_{\mathcal{D}_a}[u_N^{coop}]$ and $E_{\mathcal{D}_a}[\overline{u}_N]$ be the corresponding expectation given $\mathcal{D}_a$. Finally, recall that for a network $N = (G, R)$ and a set of edges E' of G, $N \setminus E'$ denotes $(G \setminus E', R)$, and let

$$IC(\mathbf{m}_a) \triangleq \min_{c \text{ s.t. } Pr(c\in N_1(a)|\mathcal{D}_a)>0} \delta(E_{\mathcal{D}_a}[u_N^{coop}] - E_{\mathcal{D}_a}[\overline{u}_{N\setminus E_R^a(c)}]) - \gamma,$$

and

$$IC(\mathbf{m}) \triangleq \min_{a\in A} IC(\mathbf{m}_a).$$

The proof of Lemma 1 is deferred to Appendix B.

Lemma 1 *(Incentives to Cooperate) In an economy that is consistent with the vector* $\mathbf{m} = (\gamma, \pi, \delta, \{\mathcal{D}_a\}_{a\in A})$, *there exists a TCEO in the economy if and only if* $IC(\mathbf{m}) > 0$.

Lemma 1 suggests that $IC(\mathbf{m}_a)$ is the (sufficient statistic for the) **Incentives of an agent** a **to Cooperate** with all of the clients connected to her, and that in a TCEO agent a's best response depends only on a's belief with respect to the network structure ($\mathcal{D}_a$). Nevertheless, computing the incentives of a seller to cooperate poses significant challenges.

The obvious difficulty is that $E_{\mathcal{D}_a}[\overline{u}_{N\setminus E_R^a(c)}]$ depends on the optimal strategy of agent a after deviating in an interaction with client c (when the underlying network is $N \setminus E_R^a(c)$). Moreover, even a direct computation of $E_{\mathcal{D}_a}[u_N^{coop}]$ is very complex for any belief $\mathcal{D}_a$ that puts positive probability on large networks. The source of the difficulty is in evaluating the probability that a given agent and a given client interact in a given period – a probability that depends on the entire network structure even given simple cooperative strategies. In the following sections we alleviate the difficulty in two stages. First, we derive conditions under which $E_{\mathcal{D}_a}[u_N^{coop}]$ and $IC(\mathbf{m}_a)$ can be approximated by considering the structure of only small parts of each of the networks that have positive probabilities under the belief $\mathcal{D}_a$. Second, we propose a plausible belief structure (a mapping from a network structure to a profile of agents' and clients' beliefs) such that $IC(\mathbf{m}_a)$ can be calculated as if the network was a simple-to-analyze network, i.e. a random tree.

Remark 1 *Definition 2 and Lemma 1 raise two related questions are worth discussing: [1] what is the role of the real underlying network when best responses depend only on agents' beliefs? And [2] how restrictive is the focus on TCEOs? Are there networks in which cooperation can be sustained using other strategies? The answer to the first question is in the connection between the belief profile and the actual network structure. For now, we imposed no such connection and our first main result (Theorem 1) holds for any connection (or disconnection) between the underlying network and the belief profile. In section 6 impose one specific model of beliefs that ties down the connection between a network structure and the corresponding belief profile. With respect to the second question, we note that in the most general case, the focus on TCEOs is restrictive. However, an implication of Proposition 1 is that if agents's have sufficient (local) knowledge of the real reputation network (R) then the largest set of initial networks (and corresponding beliefs) that can ever facilitate a TCEO is identical to the largest set of networks that can facilitate full cooperation in any strict Nash Equilibrium. Thus, the upper bound on the set of interaction networks that can be sustained in equilibrium is not affected by our focus on a ostracizing strategies. Since Nash equilibria are more permissive than perfect Bayesian equilibria, our results provides also an upper bound on the set of interaction networks that can be sustained in perfect Bayesian equilibrium.*

5 Cooperation based on local beliefs

Our first main result, which provides the foundation for later results, shows that whether or not there exists a TCEO in an economy that is characterized by a given vector $\mathbf{m}$ is *asymptotically independent* of the agents' *global beliefs*, and depends only on their *local beliefs*. This is quite surprising, since the fact that we focus on networks with bounded degree implies that the overwhelming majority of information about other clients and agents is not included in any agent's *local beliefs*.

Let $deg(G)$ denote the maximum degree of any client or agent in network G. For a given belief $\mathcal{D}_a$, let $deg(\mathcal{D}_a)$ denote the supremum, over all networks G to which $\mathcal{D}_a$ assigns strictly positive probability, of $deg(G)$.

Theorem 1 *(Local Beliefs Theorem)For any $\Delta > 0$ and $\epsilon > 0$, there exists a finite constant $d = d(\gamma, \pi, \delta, \Delta, \epsilon)$* ***independent*** *of the size of the entire network such that for any belief $\mathcal{D}_a$ satisfying $deg(\mathcal{D}_a) \leq \Delta$,*

$$\left| IC(\gamma, \pi, \delta, \mathcal{D}_a) - IC(\gamma, \pi, \delta, \mathcal{D}_a^d) \right| < \epsilon.$$

Theorem 1 implies that whenever we can make comparative statements about cooperation under beliefs $\{\mathcal{D}_a^d\}_{a \in A}$, we can also make (asymptotic) comparative statements about cooperation under beliefs $\{\mathcal{D}_a\}_{a \in A}$.

Consider a special case in which agents have complete knowledge of the true underlying network – i.e. consider a true network N, and for every $a \in A$, let $\mathcal{D}_a$ place probability 1 on N. An implication of Theorem 1 is that whenever we can make comparative statements about cooperation given the local neighborhoods of all agents (i.e. $\{N_d\,(a)\}_{a \in A}$), we can also make (asymptotic) comparative statements for the entire network N. This later interpretation of Theorem 1 highlights that when an agent a determines whether or not to cooperate she "discounts" links that are at a large distance from her and can asymptotically do as good by considering only her local neighborhood.

The proof of Theorem 1 builds on recent developments in graph theory, and in particular on GG who study randomized 'greedy' algorithms for matchings in a graph, and the relationship between the local and global properties of the set of matchings of a graph. We defer the complete proof to Appendix B. Instead, we state the following key lemma that sheds light on the intuition behind Theorem 1 as well as on the generality of the observation that local beliefs are sufficient to predict outcomes in a network.

Fix any network N, and let $N(a,d)$ denote the depth$-d$ neighborhood of agent a in N. For each client c and agent a, let $I_N^t(c,a)$ denote the indicator of the event that c interacted with a in period t, and let $Pr\big(I_N^t(c,a)\big)$ denote the probability that $I^t(c,a)=1$ in a network N. Note that one may interpret the quantity $Pr\big(I_N^1(c,a)\big)$ $\Big(Pr\big(I_{N(a,d)}^1(c,a)\big)\Big)$ as the probability that edge (c,a) is chosen to belong to the random graph matching constructed by examining the edges of N $\big(N(a,d)\big)$ in a random order (selected u.a.r.) and including an edge if no incident edges have already been examined. Noting that this randomized matching construction is exactly the matching algorithm studied in GG, it follows from Lemma 6 of GG that

Lemma 2 *(Locality Lemma) For any $\Delta>0$ and $\epsilon>0$, there exists a finite constant $d=d(\Delta,\epsilon)$* ***independent*** *of the size of the entire network such that for any network N satisfying $deg(G)\leq\Delta$,*

$$|Pr\big(I_N^1(c,a)=1\big)-Pr\big(I_{N(a,d)}^1(c,a)=1\big)|<\epsilon.$$

Lemma 2 highlights the observation that when interactions are mutually exclusive, and when there is a sufficiently strong stochastic element in the order of interactions, whether or not two individuals interact with each other depends heavily on the local patterns of interactions. On the other hand, in such environments, the global patterns of interactions may be less important. The intuition behind Theorem 1 and Lemma 2 is demonstrated in Example 1.

Example 1 *Consider the leftmost network in figure 2 and suppose that agent a has a belief D_a that puts probability 1 on the correct network. In order to decide whether to cooperate with client c, agent a evaluates the probability that she interacts with client c in a given period. Based on D_a^1 the corresponding probability is 1. To see why, note that there* *are no other agents in D_a^1. Now consider the belief D_a^2, because there are 2 agents in D_a^2 and only one client, and because the order of interactions is drawn u.a.r., the probability that a interact with c based on D_a^2 is $\frac{1}{2}$. Following the same logic, the corresponding probability based on D_a^3 is $\frac{2}{3}$. This is because the only orders of meetings in which c and a do not interact are those in which agent a' and client c meet before any other client and agent interact. Continuing the process in the same way we get the the probability that c and a interact in a given period based on beliefs D_a^4 and D_a^5 are $\frac{5}{8}$ and $\frac{19}{30}$ respectively. Notably, the sequence $\left(1, \frac{1}{2}, \frac{2}{3}, \frac{5}{8}, \frac{19}{30}\right)$ is monotonically converging (i.e. the deviation from the value that is based on the correct belief is monotonically decreasing).*

More generally, ordering the edges in a graph u.a.r. is equivalent to randomly and independently assigning each edge a real number distributed uniformly between 0 and 1, and then choosing the edges one by one from the low to the high value. Now, consider a fixed agent a and client c connected by an edge (a, c). Suppose that the edge (a, c) is contained within a "small" subgraph H such that every edge e in H has the following property: the value assigned to e is strictly less than the value assigned to all edges adjacent to e which are not contained in H. In this case, all inclusion/exclusion decisions (with respect to the random matching) about the edges on the boundary of H are made before any neighboring edges in $G \backslash H$ are even considered. The result is that no agent-client interactions external to H can have any influence on the agent-client interactions internal to H. Lemma 2 is based on the observation that as long as the maximum degree of any node in the overall network is uniformly bounded, with high probability any given agent-client edge (a, c) will be contained within such an influence-resistant subgraph, whose size is a small constant, independent of the size of the overall network. This leads to an asymptotic independence on the global structure of the network, and allows for a purely "local" analysis.

In the following section we apply Theorem 1 to an environment in which agents and clients have incomplete information with respect to the network structure. We show that in this environment the local beliefs of agents and clients with respect to the network structure are much simpler than their corresponding global beliefs. Thus, Theorem 1 offers a considerable simplification.

6 The *Global Fractions (GF)* model

We are interested in the following question: "For what structures of the interaction network G there exists a reputation network R such that there exists a TCEO with the network $N = (G, R)$?". So far we remained agnostic with respect to the relationship between the actual underlying network and the beliefs that agents (and clients) hold with respect to the network structure. This approach has the advantage of being the most general, but to answer our question of interest we must take a stand with respect to the role of the underlying network in generating agents' (and clients') beliefs.

Clearly, in any reasonable application, there is a close connection between the network structure and the underlying beliefs. For example, in any setting of repeated interactions, we expect any individual to know who she is connected to, and maybe have additional knowledge that relates to the outside options of individuals who are connected to her. An individual may also have some aggregate information on the global network structure. Whatever this knowledge may be, an individual's belief must be consistent with her knowledge.

We now propose a specific model of individuals' knowledge and beliefs with respect to the network structure, which we call the *Global Fractions (GF)* model. The GF model is based on the idea that the underlying process of the formation of the network has a significant random component, but that (in large networks) the fraction of clients and agents with a given degree is more or less constant and therefore known.[17] Due to the random component, individuals have only partial information on the network structure, and each individual may hold private information about her own local area of the network.

Formally, assume that before period 1 the network N is drawn u.a.r. from all networks with a set of agents A, a set of clients C, and a given degree distribution that specifies for all d: [1] the *fraction* of clients that have degree d, and [2] the *fraction* of agents that have degree d. Let the underlying network selection process be common knowledge. In addition, each agent a has access to *private local information* including: the set of clients connected to her ($N_1(a)$), the degree of each client connected to her (d_c for all $c \in N_1(a)$), and which of her clients observes the outcomes of her interactions with any of her other clients ($R(c, c')$

[17] For evidence on consistency in degree distributions in large networks, see also Barabási and Albert (1999).

for all $c, c' \in N_1(a)$). The (Bayesian) posterior of agent a is denoted by $\mathcal{D}_{GF}(a|N)$. We note that $\mathcal{D}_{GF}(a|N)$ assigns *equal probability* to any network that satisfies [1] and [2], has sets of clients and agents C and A respectively, and is consistent with the agent's private local information.[18]

Note that specifying the aforementioned fractions is equivalent (under a simple transformation) to instead specifying the probability that the client (agent) in an edge (c, a) selected u.a.r. from all edges of the network has degree d (for all d).[19] Furthermore, specifying the number of clients n_c along with these probabilities fully determines the number of agents n_a. Thus, for any $d_a \in \mathbb{Z}^+, \bar{b}(a) \in (\mathbb{Z}^+)^{d_a}, n_c \in \mathbb{Z}^+$ and probability distributions $\mathcal{G}^a, \mathcal{G}^c$ each with *bounded support* on $\mathbb{Z}^+$ and assigning each integer a *rational* probability, we let $\mathcal{N}(d_a, \bar{b}(a), R, n_c, \mathcal{G}^a, \mathcal{G}^c)$ denote the set of all networks in which agent a has degree d_a, the sorted vector of degrees of the clients that are connected to a is equal to $\bar{b}(a)$, the reputation network $R(c, c')$ for all $c, c' \in N_1(a)$ is given by R, the total number of clients in the network is n_c, and the probability that the agent (client) $a'(c')$ in an edge (c', a') selected u.a.r. from all edges of the network has degree d is equal to $Pr(\mathcal{G}^a = d)$ $\big(Pr(\mathcal{G}^c = d)\big)$.[20] Consequently, if the true underlying network is $N \in \mathcal{N}(d_a, \bar{b}(a), R, n_c, \mathcal{G}^a, \mathcal{G}^c)$, then the posterior of agent a is that the network is selected u.a.r. from all networks in $\mathcal{N}(\cdot)$. We denote this belief by $\mathcal{D}_{GF}(a|N) = \mathcal{D}_{GF}(d_a, \bar{b}(a), R, n_c, \mathcal{G}^a, \mathcal{G}^c)$. Using the GF model, we relate any underlying network structure to a corresponding belief profile.

Definition 3 *(Cooperation Network) Let* $\mathbf{m}^{GF}(N) = (\gamma, \pi, \delta, \{\mathcal{D}_{GF}(a|N)\}_{a \in A})$. *We say that a network N is a* GF cooperation network *if and only if there exists a TCEO in the economy* $\mathbf{m}^{GF}(N)$.

In the following section we develop a set of results that, combined with Theorem 1, relate the question of whether a large network N is a GF cooperation network to a simpler question.

[18] Our analysis goes through with individuals that are less informed, and can also be modified to allow for more informed individuals (e.g. knowing the degrees of neighbors of neighbors).

[19] Let $P_a(d)$ the proportion of agents with degree d, and let $\bar{d}_a = \sum_d P_a(d) \cdot d$ be the average agent's degree. Then the probability that an agent a in an edge that is chosen u.a.r. has degree d is, $\frac{P_a(d) \cdot d}{\bar{d}_a}$.

[20] Suppose further that $Pr(\mathcal{G}^a = d_a) > 0$, and $Pr(\mathcal{G}^c = \bar{b}_i(a)) > 0$ for all i. Then, it is well-known that for any fixed $d_a, \bar{b}(a), R, \mathcal{G}^a, \mathcal{G}^c$ there exists an infinite strictly increasing sequence of integers $\{n_c\}$ s.t. $\mathcal{N}(\cdot|n_c, a) \neq \emptyset$. This follows from the Gale-Reyser Theorem – see e.g. Krause (1996), and (in our particular setting) Theorem 1.3 of Greenhill et al. (2006). All statements should be read as holding only for n_c s.t. the aforementioned set is non-empty.

Namely, we show that it is sufficient to examine the incentive of each agent to cooperate *as if* the agent is a part of a simpler network that we define – *a simple random tree.*

Remark 2 *The economic literature offers several models of network formation (see also Goyal 2007 and Jackson 2008 and references therein). By construction, any such process can be captured by some* $\{\mathcal{D}_a\}_{a\in A}$. *This implies that our analysis is 'formation process free'. The GF model adds structure to captures a scenario in which individuals 'have no clue' how networks are formed, but have some information on their attributes.*[21]

7 Cooperation and network structure in the GF model

In this section, we focus on agents whose knowledge and beliefs are consistent with the GF model and derive conditions on a network N such that N is a *GF cooperation network.* A key result is that for all finite support distributions $\mathcal{G}^a, \mathcal{G}^c$ and for every agent a in any asymptotically large network N with degree distributions $\mathcal{G}^a, \mathcal{G}^c$, the belief of agent a is asymptotically identical to the belief that the network looks locally like a corresponding *simple random tree.* We then combine this result with Theorem 1 to derive the main result of this section: a large network N is a GF cooperation network if and only if the incentives of each agent to cooperate in her corresponding random tree are sufficiently large. In the following section we rely on this result and characterize the set of interaction networks G for which there exists a reputation network R such that the network $N = (G, R)$ is a GF cooperation network.

For any given $d_0 \in \mathbb{Z}^+$, $\bar{b} \in (\mathbb{Z}^+)^{d_0}$, distributions $\mathcal{G}^a, \mathcal{G}^c$ with finite support, and $d \geq 1$, let $T(d_0, \bar{b}, R, \mathcal{G}^a, \mathcal{G}^c, d)$ denote the random depth$-d$ tree such that the root r has degree d_0, the sorted vector of degrees of the children of r is $\bar{b}$, all subsequent non-leaf nodes at an even depth have a number of children drawn i.i.d. from $\mathcal{G}^a - 1$, all subsequent non-leaf nodes at an odd depth have a number of children drawn i.i.d. from $\mathcal{G}^c - 1$, the underlying reputation network for the clients connected to the node $(N_1(r))$ is R, and for all other clients pairs c, and c', the probability that c and c' are connected is $\frac{1}{2}$ $\left(\Pr\left(R(c, c') = 1\right) = \frac{1}{2}\right)$.

[21] An algorithm for generating valid large random graphs with arbitrary degree distributions exists and is suggested in the proof of Lemma 3.

To measure how *different* two beliefs are we use the notion of *total variational distance.* For two random variables (r.v.) X, Y with support on some countably infinite set $\mathcal{X}$, the *total variational distance* between X and Y, $TVD(X,Y)$, is defined as $\sum_{x\in\mathcal{X}} |Pr(X = x) - Pr(Y = x)|$. Then for two belief distributions $\mathcal{D}_a, \mathcal{D}'_a$, $TVD(\mathcal{D}^d_a, \mathcal{D}'^d_a) < \epsilon$ implies that the belief agents a has about her depth-d neighborhood under $\mathcal{D}_a$ is 'within ϵ' of the belief agent a has about her depth-d neighborhood under $\mathcal{D}'_a$.

Lemma 3 *(Locally Tree-Like Lemma) For all $d_0, \bar{b}, R$, and finite support $\mathcal{G}^a, \mathcal{G}^c$, and for all $d > 0$,*

$$\lim_{n_c\to\infty} TVD\big(\mathcal{D}^d_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c), T(d_0, \bar{b}, R, \mathcal{G}^a, \mathcal{G}^c, d)\big) = 0$$

The proof of Lemma 3 is presented in Appendix A and employs the so-called *configuration method* (see Wormald 1999). Using this technique, a random graph is related to a different random object - the configuration model. In the configuration model, each client (agent) is viewed as a bucket, and each bucket is endowed with a number of points equal to the desired degree of the corresponding client (agent). The points in the buckets are then matched randomly, and an agent a and a client c are connected if a point from a's bucket is matched to a point of c's bucket. By starting this construction with a given agent node and continuing sequentially by connecting at every step all of the points in the buckets who where connected to in the previous step, we show that in asymptotically large networks: [1] the number of steps that it takes until a cycle is closed is arbitrarily large; and [2] after any finite number of steps, the degree distribution of the buckets that are still unmatched is asymptotically (on the size of the network) identical to the degree distribution in the entire network.[22]

Lemma 3 implies that in an asymptotically large network N, for any fixed d, the belief $\mathcal{D}^d_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c)$ converges to a belief on a random tree. Thus, to make use of the convenient structure of a random tree, one must establish that for the measure of interest, an agent's depth-d belief is a sufficient statistic to determine the measure of interest. Theorem 1 completes this gap with respect to agents' incentives to cooperate and we are able to derive the following key result.

[22] More generally, there is a vast literature that both analyzes the configuration model, and relates it back to many random-graph models of interest (see e.g. Greenhill et al. 2006, Wormald 1999). Such relations often involve subtle counting and conditioning arguments - for more details the reader is referred to the proof in Appendix A.

Theorem 2 *(Asymptotic Characterization of Cooperation Networks I) Fix γ, π, δ and consider a sequence of networks $(N^1, N^2 ...)$ with identical finite support degree distributions $\mathcal{G}^a, \mathcal{G}^c$ and an increasing size (i.e. the numbers of clients and agents in network N^{i+1} are larger than the corresponding numbers in network N^i). Then, there exists a number $\bar{i}$ such that for all $i > \bar{i}$ the network N^i is a GF cooperation network if and only if for any agent a,*

$$\lim_{d \to \infty} IC(\gamma, \pi, \delta, T(d_a, \bar{b}(a), R, \mathcal{G}^a, \mathcal{G}^c, d)) > 0,$$

where $\bar{b}(a)$ is the sorted vector of the degrees of the clients that are connected to agent a.

Proof. We prove the theorem by proving the following result: For all $\gamma, \pi, \delta, d_0, \bar{b}$, and finite support $\mathcal{G}^a, \mathcal{G}^c$, $\lim_{d \to \infty} IC(\gamma, \pi, \delta, T(d_0, \bar{b}, R, \mathcal{G}^a, \mathcal{G}^c, d))$ and $\lim_{n_c \to \infty} IC(\gamma, \pi, \delta, \mathcal{D}_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c))$ both exist, and equal one-another. This follows from Theorem 1 and Lemma 3. ■

Lemma 3 implies that in a large network N, the sequence of beliefs $\mathcal{D}^d_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c)$, $d = 1, 2, \ldots$, converges in a sense to a belief on an 'infinite random tree'. As a result, in a large network N, and as long as agents' beliefs are consistent with the GF model, many relevant quantities can be described in terms of the associated limits. Most relevant for the characterization of large cooperation networks, Theorem 2 implies an explicit asymptotic characterizations of cooperation networks in terms of a dynamic program that is based on the belief that the network is a random tree.

In a particularly interesting limit case we can give an especially simple characterization. Let R^1 be the complete reputation network in which any two clients are connected and each client is informed of transaction outcomes of all other clients (for all c' and c, $R^1(c', c) = 1$). Recall that $\Pr\left(I^1_N(c, a)\right)$ is the probability that c and a interact in period 1 given the network N.

Corollary 1 *(Asymptotic Characterization of Cooperation Networks II) Let $\gamma, \pi, \delta, d_0, \bar{b}, \mathcal{G}^a, \mathcal{G}^c$ be fixed and $\mathcal{G}^a, \mathcal{G}^c$ have finite support. For a agent a, let $N_{d,a}$ denote the belief $T(d_0, \bar{b}, R, \mathcal{G}^a, \mathcal{G}^c, d)$. Then, $\lim_{d \to \infty} \sum_{c \in N_1(a)} \Pr\left(I^1_{N_{d,a}}(c, a)\right)$ exists, and*

$$sign\left(\lim_{n_c \to \infty} IC(\gamma, \pi, \delta, \mathcal{D}_{GF}(d_0, \bar{b}, R^1, n_c, \mathcal{G}^a, \mathcal{G}^c))\right) = sign\left(\frac{\delta(\pi - \gamma)}{1 - \delta} \lim_{d \to \infty} \sum_{c \in N_1(a)} \Pr\left(I^1_{N_{d,a}}(c, a)\right) - \gamma\right). \tag{1}$$

Proof. That the necessary limit exists follows from Lemma 2. To derive (1) note that if the reputation network is captured by R^1 then $\overline{u}_{N \setminus E_R^a(c)} = 0$ for all c. Therefore, it follows from (9) that for any fixed belief N and strategy Q,

$$\left(\overline{u}_N\left(Q\right) - u_N^{coop}\right) = \frac{\left(\gamma - \delta \cdot u_N^{coop}\right) \cdot \sum_{c \in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)}{1 - \delta + \delta \cdot \sum_{c \in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)} \tag{2}$$

The Corollary then follows from (2), (10), and Theorem 2. ■

In the following section we use Corollary 1 and the corresponding tree structure described in Lemma 3 and Theorem 2 to characterize the interaction networks G for which there exists a reputation network R such that the network $N = (G, R)$ is a GF cooperation network.

7.1 Third-party observability and cooperation

In this section, we focus on two questions: [1] how do the patterns of third-party observability, as captured by the reputation component of the network (R), affect the ability of different patterns of repeated interactions, as captured by the interaction network (G), to facilitate cooperation? And [2] what is the largest set of interaction networks G for which there exists a reputation network R such that the network $N = (G, R)$ is a GF cooperation network? We first establish that third-party observability helps cooperation and allows a larger set of interaction networks to sustain cooperation. Then, we focus on the analysis of the incentives of agents to cooperate as a function of the structure of the interaction network when the reputation network is complete.

The intuition that the complete reputation network (R^1) allows for the largest set of interactions networks to be to sustain cooperation is consistent with much of the literature on community enforcement; when more clients are aware of an agent's defection, the agent faces a larger punishment for defecting.

Proposition 1 *(Weak Monotonicity of Cooperation in Third-Party Observability)Fix γ, π, δ and consider a sequence of interaction networks $(G^1, G^2 ...)$ with identical finite support degree distributions $\mathcal{G}^a, \mathcal{G}^c$ and an increasing size (i.e. the numbers of clients and agents in network G^{i+1} is large than the corresponding numbers in network G^i). Then, there exists a number $\bar{i}$ such that for all $i > \bar{i}$ and for any reputation network R, if the network (G^i, R) is a GF cooperation network then the network (G^i, R^1) is also a GF cooperation network.*

Proof. We prove Proposition 1 by proving the following claim: Let $\gamma, \pi, \delta, d_0, \bar{b}, \mathcal{G}^a, \mathcal{G}^c$ be fixed and $\mathcal{G}^a, \mathcal{G}^c$ have finite support, and let $IC_{\mathcal{D}_{GF}}(R, n_c) = IC(\gamma, \pi, \delta, \mathcal{D}_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c))$. Then, for any reputation network R, $IC_{\mathcal{D}_{GF}}(R, n_c) \leq IC_{\mathcal{D}_{GF}}(R^1, n_c)$.

To prove this claim, note that by definition,
$IC_{\mathcal{D}_{GF}}(R, n_c) = \min_{c \in N_1(a)} \delta \left(E_{\mathcal{D}_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c)}[u_N^{coop}] - E_{\mathcal{D}_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c)}[\overline{u}_{N \setminus E_R^a(c)}] \right) - \gamma$.
The lemma then follows from: [1] for all R, $E_{\mathcal{D}_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c)}[\overline{u}_{N \setminus E_R^a(c)}] \geq 0$;
[2] $E_{\mathcal{D}_{GF}(d_0, \bar{b}, R^1, n_c, \mathcal{G}^a, \mathcal{G}^c)}[\overline{u}_{N \setminus E_{R^1}^a(c)}] = 0$; and [3] $E_{\mathcal{D}_{GF}(d_0, \bar{b}, R, n_c, \mathcal{G}^a, \mathcal{G}^c)}[u_N^{coop}] = E_{\mathcal{D}_{GF}(d_0, \bar{b}, R^1, n_c, \mathcal{G}^a, \mathcal{G}^c)}[u_N^{coop}]$.
■

Fainmesser (2012a) relies on the results established in this paper (especially Theorems 1 and 2) and studies the structure of cooperation networks in markets in which third-party observability is extremely limited (when R is the empty network, i.e. no two clients are connected in R). Fainmesser (2012a) shows that the incentives to cooperate are strongest in *sparse* interactions networks in which there are similar numbers of clients and agents (with similar and low degrees).

Proposition 1 teaches us that in order to study the limits of cooperation we are required to focus on networks that take the form $N = (G, R^1)$ and make comparative statements with respect to the incentives of agents to cooperate as a function of G. We find that when $R = R^1$, the incentives to cooperate are strongest in *dense* interaction networks in which there are more clients then agents (so agents' degrees are high and clients' degrees are low). In the next section, we extend our model to allow for stochastic elements in the production function of agents and show that our results imply that third-party observability enhances efficiency in markets that rely on trust and cooperation, but that nevertheless, when there are more agents than clients in a market, an efficient outcome is not guaranteed.

A fundamental observation that is captured by Corollary 1 is that when $R = R^1$, agents who expect to interact with higher overall probability (larger $\sum_{c \in N_1(a)} \Pr\left(I_N^1(c, a)\right)$) have stronger incentives to cooperate. As a result, cooperation is better sustained when agents' (clients') degrees are large (small). To see why, note that [1] the probability that an agent a interacts in a given period is increasing in her degree; and [2] the probability that a interacts in a given period is an increasing function of the degree of any agent that is connected to a client that a is connected to. The intuition for the latter is subtle: because agents with high degrees are less likely to interact with each one of the clients to whom they are connected, they are less likely to interact with each client before other agents who are connected to the same client get a chance to interact with him.[23]

More generally, when agents' degrees are large and clients' degrees are small, there are more clients and less agents who interact. Consequently, agents interact with high probability and expect large payoffs. Figure 3 provides a simple deterministic example of the general rule that is captured by Proposition 2.

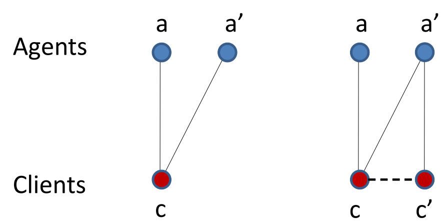

Figure 3: Consider the two networks above and let the dashed line between clients c and c' in the rightmost network represent a link between them in the reputation network ($R(c, c') = 1$). In the leftmost network, agent a expects to interact with probability $\frac{1}{2}$ in any given period. As a result, conditional on cooperation between a' and c, agent a cooperates with c if and only if $\frac{\delta}{1-\delta}(\pi - \gamma) \cdot \frac{1}{2} > \gamma$. In the rightmost network, agent a' has an additional link. Clearly, this increases the probability that agent a' interacts in a given period from $\frac{1}{2}$ to 1. However, this also increases the probability that agent a interacts in a given period from $\frac{1}{2}$ to $\frac{2}{3}$. As a result, conditional on cooperation between a' and c, as well as between a' and c', agent a cooperates with c if and only if $\frac{\delta}{1-\delta}(\pi - \gamma) \cdot \frac{2}{3} > \gamma$.

Theorem 2 shows that the incentive effect of any change in the degree of any client or agent in the network can be captured by considering the effect of the corresponding change in the corresponding random tree. In turn, any change to the degree of a node in a random tree can be captured by appending or removing subtrees (A tree T' is subtree of a tree T if $T' \subseteq T$). Thus, Proposition 2 shows that the incentive effect of any change in the degree of any client or agent in the network can be determined by identifying whether the corresponding change in the corresponding random tree involves adding or removing (as children) subtrees from a client or an agent node. E.g. appending (as children) subtrees to agents nodes in the corresponding random tree (rooted with agent a) can capture: [1] adding links between a and some clients; and/or [2] increasing the degree distribution of agents in the network as a whole. Thus, Proposition 2 shows that the effects of [1] and [2] on the

[23] The intuition applies most directly to trees. However, our random tree characterization highlights the connection between trees and large networks.

incentives of agent a are qualitatively the same. The proof of Proposition 2 is deferred to Appendix B.

Proposition 2 *(Monotonicity of Cooperation in Degree) Let γ, π, δ be fixed. Suppose that for all $d \geq 1$, the random tree $T^2 = T(d_0^2, \bar{b}^2, R^1, \mathcal{G}^{2^a}, \mathcal{G}^{2^c}, d)$ can be constructed (on the same probability space) from the random tree $T^1 = T(d_0^1, \bar{b}^1, R^1, \mathcal{G}^{1^a}, \mathcal{G}^{1^c}, d)$ by performing only the two operations: [1] appending (as children) subtrees to agent nodes in an arbitrary way, and [2] removing (as children) subtrees from client nodes in an arbitrary way. Then,*
$\lim_{n_c \to \infty} IC\left(\gamma, \pi, \delta, \mathcal{D}_{GF}\left(d_0^1, \bar{b}^1, R^1, n_c, \mathcal{G}^{1^a}, \mathcal{G}^{1^c}\right)\right) > 0$ *implies*
$\lim_{n_c \to \infty} IC\left(\gamma, \pi, \delta, \mathcal{D}_{GF}\left(d_0^2, \bar{b}^2, R^1, n_c, \mathcal{G}^{2^a}, \mathcal{G}^{2^c}\right)\right) > 0.$

Proposition 2 implies that networks in which agents are well connected and clients have only low degrees maximize the incentives to cooperate. However, it is just one implication of a more general rule that is due to the perfect alignment between an agent's probability of interacting and her incentives to cooperate: consider two networks $N' = (G', R^1)$ and $N'' = (G'', R^1)$, and assume that conditional on full cooperation, the minimal probability of interaction of any agent is weakly higher in N'' than in N' $\left(\min_a \sum_{c \in N_1''(a)} \Pr\left(I^1_{N''}(c,a)\right) \geq \min_a \sum_{c \in N_1'(a)} \Pr\left(I^1_{N'}(c,a)\right)\right)$. Then N' being a cooperation network implies that N'' is a cooperation network. This is true regardless of the size of the network and the model of beliefs that we consider. In networks in which all agents are symmetric, this rule implies that the expected aggregate number of interactions and the incentives of agents to cooperate are perfectly aligned. To fully demonstrate the implications of this observation, we generalize our model to allow for stochastic elements in the production technology. Considering an environment in which agents might have periods in which they are not able to interact (e.g. entrepreneurs might not have an investment idea, and sellers may have stochastic shocks to their production) highlights the limits on efficiency that are imposed by the need to rely on trust and cooperation – even in the presence of perfect third-party observability.

8 Optimal networks

In this section we study the constraints imposed on the number of (mutually beneficial) bilateral interactions by the need to enforce cooperative behavior in an incentive compatible manner. To this end, we extend our model in two ways. First, we allow for stochastic elements in the production technology of agents. Formally, we assume that in every period, each agent is active with some probability $\mu \in [0,1]$ and inactive otherwise. The realization of whether an agent is active in a given period is i.i.d. across agents and periods. An active agent can interact as described above, whereas an inactive agent cannot interact. The probability that an agent is active (μ) captures any stochastic element in the production technology. For example, if the agent is an entrepreneur then she may not be able to come up with a profitable investment opportunity in every period. Similarly, if the agent is a producer, she may suffer from exogenous shocks to her production process. Notably, all of our results above extend without change to this more general environment.[24] Moreover, introducing instead a stochastic element on the demand side yield qualitatively identical results. We show below that stochastic elements in the market provide a key reason to prefer fully connected interaction networks in order to maximize the number of trades. We further show that in some markets, for a fully connected interaction network G, there is no reputation network R such that the network $N = (G, R)$ is a GF cooperation network.

A second way we extend our model is by considering a network design problem that is useful for comparing the first best (i.e. how many cooperative interactions are possible without the need to satisfy the incentive constraints for cooperation?) with the second best (i.e. how many cooperative interactions are possible when the incentive constraints for cooperation must be satisfied?). The difference between the first and second best captures the limits of cooperation, as well as the limits of the effectiveness of third-party observability in facilitating cooperation.

Consider a market with $\overline{n}_a$ agents and $\overline{n}_c$ clients and parameters γ, π, μ, δ. Let ΔN be a probability distribution over network structures $N = (G, R)$. In the *unconstrained design problem*, a planner chooses ΔN and *compels* all agents to follow strategy $Q_{\mathcal{D}}^{coop}$ (always cooperate). In the *cooperation constrained design problem*, the planner chooses ΔN and *recommends* that all agents follow strategy $Q_{\mathcal{D}}^{coop}$; agents are then informed of ΔN and follow the planner's recommendation only if $(\gamma, \pi, \mu, \delta, \Delta N)$ admits a TCEO. Let $n_c(N)$

[24] The interested reader is referred to Fainmesser and Goldberg (2011) which is an older draft of this paper and is available on Fainmesser's websites.

$(n_a(N))$ be the number of clients (agents) whose degree in G is at least 1. If $n_c(N) < \overline{n}_c$ $(n_a(N) < \overline{n}_a)$ we say that $\overline{n}_c - n_c(N)$ clients $(\overline{n}_a - n_a(N)$ agents) are *excluded* from the market in N.

For a given network N, let $E[V(N)] = E\left[\sum_{a\in A}\sum_{c\in N_1(a)} Pr\left(I_N^1(c,a)\right)\right]$ denote the expected number of cooperative interactions that are achieved in a given period if all agents follow strategy $Q_{\mathcal{D}}^{coop}$. Denote by $E[V(\Delta N)]$ the corresponding value given a probability distribution ΔN over networks. Let $N \in \mathcal{N}(\overline{n}_c, \overline{n}_a)$ if the network N can be constructed with $\overline{n}_c$ clients and $\overline{n}_a$ agents (i.e. $n_c(N) \leq \overline{n}_c$ and $n_a(N) \leq \overline{n}_a$), and let $N^{uc}(\cdot)$ $(N^c(\cdot))$ be the solution to the unconstrained (constrained) design problem. Then,

$$N^{uc}(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta) = \underset{\Delta N | supp(\Delta N) \subseteq \mathcal{N}(\overline{n}_c, \overline{n}_a)}{argmax} E[V(\Delta N)]$$

and

$$N^c(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta) = \underset{\Delta N | supp(\Delta N) \subseteq \mathcal{N}(\overline{n}_c, \overline{n}_a)}{argmax} \underset{, \min_{a\in A} IC(\gamma,\pi,\mu,\delta,\Delta N) > 0}{E[V(\Delta N)]},$$

where $\min_{a\in A} IC(\gamma, \pi, \mu, \delta, \Delta N) > 0$ guarantees that $(\gamma, \pi, \mu, \delta, \Delta N)$ admits a TCEO. Thus, the proportion of welfare loss due to the constraints on the structure of cooperation networks is

$$WL(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta) = 1 - \frac{E[V(N^c(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta))]}{E[V(N^{uc}(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta))]}$$

If $WL(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta) = 0$, then cooperation networks can achieve the first best in a market with $(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta)$.

We now revisit our conclusion following Proposition 2. Namely, that the expected aggregate number of interactions and the incentives of agents to cooperate are perfectly aligned. We demonstrate this insight using a special family of interaction networks which we call *semi-regular networks*. In a semi-regular network all agents have the same degree d_A and all clients have the same degree d_C. We also show that adding a sufficiently large number of links to an interaction network G guarantees that: [1] if $R = R^1$, agents have (asymptotically) the maximal possible incentives to cooperate given n_a, n_c; and [2] conditional on N being a GF cooperation network, the expected number of interactions in every period is (asymptotically) maximal given n_a, n_c. The proof of Theorem 3 is presented in Appendix B.

Theorem 3 *(Dense Networks Maximize Welfare and Incentives to Cooperate) Let* $\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)$

be an agent's belief according to the GF model when the underlying reputation network is R^1, there are n_c clients in the market, and the degrees of all agents and clients in the interaction network G are d_A and d_C respectively. Then,

$$sign\Big(IC\left(\gamma,\pi,\mu,\delta,\mathcal{D}_{GF}\left(R^1,n_c,d_A,d_C\right)\right)\Big)=sign\left(\frac{\delta(\pi-\gamma)}{1-\delta}\cdot\frac{1}{n_a}\cdot E\left[V\left(\mathcal{D}_{GF}(R^1,n_c,d_A,d_C)\right)\right]-\gamma\right)$$

and

$$\lim_{n_c\to\infty}\frac{E\left[V\left(\mathcal{D}_{GF}(R^1,n_c,d_A,d_C)\right]\right.}{\min(\mu\cdot n_c\cdot\frac{d_C}{d_A},n_c)}\geq 1-\big(\max(d_A,d_C)-1\big)^{-1}.$$

Moreover, $\frac{n_c}{n_a}=\frac{d_A}{d_C}$. Thus, $\min(\mu\cdot n_c\cdot\frac{d_C}{d_A},n_c)=\min(\mu\cdot n_a,n_c)$, which equals the maximal volume of trade possible conditional on n_c and n_a.[25]

Theorem 3 implies an asymptotic lower bound on the incentives to cooperate.

$$\lim_{n_c\to\infty}IC\left(\gamma,\pi,\mu,\delta,\mathcal{D}_{GF}(R^1,n_c,d_A,d_C)\right)\geq\frac{\delta(\pi-\gamma)}{1-\delta}\cdot\min(\mu,\frac{d_A}{d_C})\cdot\Big(1-\big(\max(d_A,d_C)-1\big)^{-1}\Big)-\gamma. \quad (3)$$

Expression (3) highlights the importance of the ratio $\frac{d_A}{d_C}$ (or $\frac{n_c}{n_a}$) in determining whether repeated interactions can sustain cooperation in a market that gives positive probability for any of n_a agents and n_c clients to interact. If the ratio $\frac{n_c}{n_a}$ is large enough, cooperation can asymptotically be sustained and maximal number of interactions facilitated even with limited (yet large) degrees. Corollary 2 shows that Theorem 3 has implications to welfare that go beyond the semi-regular setup.

Recall that $WL\left(\overline{n}_c,\overline{n}_a,\gamma,\pi,\mu,\delta\right)$ is the proportion of welfare loss due to the constraints on the structure of networks that can sustain cooperation. Recall further that if $\frac{\delta(\pi-\gamma)}{1-\delta}\cdot\mu-\gamma<0$ then no network (apart from the empty network) sustains cooperation. This is true because even an agent who is guaranteed to interact whenever she is active, and expects to lose her entire future payoff if she defects, will still defect. Consequently, $\frac{\delta(\pi-\gamma)}{1-\delta}\cdot\mu-\gamma<0$ implies that $WL\left(n_c,n_a,\gamma,\pi,\mu,\delta\right)>0$ for any n_c,n_a. Corollary 2 covers the more interesting case where $\frac{\delta(\pi-\gamma)}{1-\delta}\cdot\mu-\gamma>0$. The proof is deferred to Appendix B.

[25] From an algorithmic perspective it is interesting to note that in addition to being easy to implement and leading to a tractable analysis, the simple matching mechanism governing our market is also *asymptotically welfare maximizing* w.r.t. the number of interactions when all agents cooperate, and all agents (clients) have the same (large) degree d_A (d_C). This is surprising, since the mechanism is quite simplistic, and corresponds better to a random decentralized market than to known algorithms for constructing optimal matchings.

Corollary 2 *(Barrier to Entry and Asymptotic Efficiency) Consider a market with* $\overline{n}_c$ *clients and* $\kappa\overline{n}_c$ *agents, and fixed* γ, π, δ *s.t.* $\delta < 1$ *and* $\frac{\delta(\pi-\gamma)}{1-\delta} - \gamma > 0$*. Then,*

1. *For any* κ, $WL\left(\overline{n}_c, \kappa\overline{n}_c, \gamma, \pi, \mu = 1, \delta\right) = 0$.
2. *Let* $\mu < 1$*. There exists* $\overline{\kappa}(\mu)$ *such that* $WL\left(\overline{n}_c, \kappa\overline{n}_c, \gamma, \pi, \mu, \delta\right) > 0$ *for any* $\kappa > \overline{\kappa}$.
3. *For any* $\mu \in \left[\gamma \cdot \frac{1-\delta}{\delta(\pi-\gamma)}, 1\right]$ *and* $\kappa \in \mathbb{Q}^+$, $lim_{\overline{n}_c \to \infty} WL\left(\overline{n}_c, \kappa\overline{n}_c, \gamma, \pi, \mu, \delta\right) = 0$.

Part 1 of Corollary 2 shows that when there is no stochastic element in the production technology ($\mu = 1$), the incentive constraints do not restrict welfare. In particular, when $\mu = 1$, a network that consists of pairs of clients and agents and some excluded clients or agents (but not both) provides the maximal number of interactions as well as the maximal incentives to cooperate. On the other hand, part 2 of the Corollary addresses the case of stochastic production technology. If $\mu < 1$ the maximal number of interactions cannot be achieved if any agent is excluded from the market. At the same time, if there are many more agents than clients, all cooperation networks exclude some agents from the market. This leads to a welfare loss. Figure 4 provides an example.

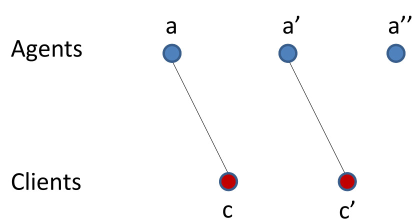

Figure 4: Assume that $\mu = 1$. In the above network, agent a cooperates with client c (and agent a' cooperates with client c'), if and only if $\frac{\delta}{1-\delta}(\pi - \gamma) > \gamma$. Moreover, conditional on cooperation between every client and agent that are connected, two interactions will take place in every period. This is the maximal number of interactions that can take place in one period in a network with three agents and two clients. Now assume that $\mu < 1$. There exists positive probability that in a given period only agents a and a'' are active. Thus, any network in which agent a'' is not connected to any client limits the number interactions to less than two even though two agents are active.

Part 3 of Corollary 2 is encouraging; in large markets (asymptotic) efficiency is restored. Theorem 3 provides the necessary intuition: let $d_A = \frac{d_C}{\kappa'}$ and fix κ'. Then as long as $1 \geq \kappa' \cdot \mu$,

$$\lim_{d_C \to \infty} \lim_{n_c \to \infty} IC\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right) = \frac{\delta(\pi - \gamma)}{1-\delta} \cdot \mu - \gamma$$

and as long as $1 \leq \kappa' \cdot \mu$,

$$\lim_{d_C \to \infty} \lim_{n_c \to \infty} \frac{E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right]}{n_c} = 1$$

Going back to part 3 of Corollary 2, no matter how large is $\frac{\overline{n}_a}{\overline{n}_c}$, a planner can choose ΔN in the following way: [1] set $R = R^1$; and [2] pick large positive integers d_A, d_C s.t. $\frac{d_A}{d_C} = \mu$ and choose G u.a.r. from the set of interaction networks s.t. $Pr(d_c = d_C) = 1$, $Pr(d_a = d_A) = \frac{\overline{n}_c}{\overline{n}_a} \cdot \frac{1}{\mu}$, and $Pr(d_a = 0) = 1 - \frac{\overline{n}_c}{\overline{n}_a} \cdot \frac{1}{\mu}$. Then, the planner achieves (asymptotically in d_A, d_C) both high incentives to cooperate and maximal number of interactions. This is interesting because the planner does not need to create a complete network. In fact, d_C (d_A) does not need to be in the order of n_a (n_c) and can be much smaller. The implications of our results in the context of barriers to entry and efficiency are discussed further in section 9.3.

9 Discussion

In this section, we first highlight the implications of our methodological contribution to social networks analysis. Then, we provide an interpretation of our characterization of GF cooperation networks in the more traditional context of market structure and discuss the implications for barriers to entry and efficiency.

9.1 The (un)importance of global beliefs

The recent literature on static network games (e.g. Galeotti et al. 2010) suggests that when players have incomplete knowledge of the networks structure, the analysis of the induced (Bayesian) game is much simpler than the corresponding analysis when players know the entire network structure. However, this does not mean that global knowledge of the network is not important. In fact, Galeotti et al. provide several examples in which changing the information structure changes the set of equilibria significantly. Kets (2011) shows that when a game is local (a player's payoff depends only on her action and the actions of her immediate neighbors) and when players have a common prior, there are weak conditions under which *small changes* to the priors do not change the equilibrium payoffs.

In our model, for the family of TCEOs, any change to a belief of an agent that keeps the agent's belief over her local neighborhood intact does not affect the agent's best response correspondence. This is especially surprising given that our game is not local – an agent's payoff generally depend on the entire network structure. The methodology we use can be applied to other setups as long as the strategic influence of one individual on another decays with the distance between the individuals. In static network games this occurs due to an assumed decay of influence, whereas in our setup this is due to the stochastic component in the order of interactions within a period.

9.2 Random network formation and random trees

The following three ideas raise separate interest in economics, sociology, and psychology: [1] the formation process of social networks has a stochastic component; [2] individuals do not know the exact structure of the (social) network in which they are embedded; and [3] individuals often consider separate interactions as independent (even when they are not).[26] Lemma 3 (and to some extent Theorem 2) offers a connection between these three observations: if the stochastic element in the underlying process of the network formation is sufficiently salient, and if individuals cannot observe perfectly or learn the entire network structure, then in a large network the *correct* prior of an individual is that her local environment is a random tree. In a random tree separate observations of an individual are independent. Naturally, one would like to explore more realistic mental models that generate similar results. To that extent, this is only a first stab at an important question: to what extent can simplified heuristics that people use to deal with incomplete knowledge of the network be explained as 'averaging' over a stochastic prior?

A by-product of Lemma 3 is the provision of sufficient conditions under which a network is expected to exhibits no degree correlation.[27] We note that this provides a microfoundation to previous reduced form assumptions used in the networks literature. For example, Jackson and Yariv (2007) assume that each player in a network has expectations on the number

[26] E.g. DeMarzo, Vayanos, and Zwiebel (2003) propose a model in which individuals learn from their neighbors about the state of the world. In their model, individuals experience *persuasion bias* - each individual i continuously updates her prior based on her neighbors' opinions ignoring the fact that her neighbors' opinions depend on the network structure and on information that was previous accessible to i. Golub and Jackson (2010) develop a similar model that allows for more flexibility in the updating rule, but maintains the assumption that an individual updates her prior ignoring the network structure.

[27] For an application, see Fainmesser (2012b).

of connections of each of the other players connected to her that are captured by a fixed degree distribution. Our result provide sufficient conditions under which this assumption is consistent with a common prior.

9.3 Third-party observability, barriers to entry, and efficiency

Proposition 1 suggests that third-party observability enhance efficiency by allowing for a larger set of interaction networks to be GF cooperation networks. Theorem 3 provides a more direct positive result: in the presence of perfect third-party observability, for *large* n_c and n_a, networks that maximize the number of (mutually beneficial) interactions and networks that maximize the incentives to cooperate are approximately identical. In such 'optimal' networks, the degrees of clients and agents are large. However, as illustrated in Example 2, Theorem 3 and Corollary 2 also suggest that there are some non-degenerate scenarios in which even with perfect third-party observability there is no network $N = (G, R)$ such that: [1] all agents in A have an opportunity to interact, and [2] N facilitates full cooperation. In Example 2, some agents are excluded permanently from the market in any network that facilitates full cooperation. Depending on the (unmodeled) network formation mechanism in a given market, this observation lends itself to several interpretations: either that the need to sustain cooperation may create a barrier to entry, or alternatively that the existence of barriers to entry may be necessary to facilitate cooperation in some markets.

Example 2 *Consider $\gamma, \pi, \mu, \delta, n_c$, and n_a such that $\frac{n_c}{n_a} < \mu$,*

$$\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \frac{n_c}{n_a} - \gamma < 0, \tag{4}$$

and

$$\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \mu - \gamma > 0. \tag{5}$$

Condition 4 guarantees that no network N in which $Pr(d_a = 0|N) = 0$ admits a TCEO. At the same time, condition 5 assures us that there exists a non-empty network that admits a TCEO. For example, a network in which $Pr(d_c = 1|N) = 1$, $Pr(d_a = 1|N) = \frac{n_c}{n_a}$, and $Pr(d_a = 0|N) = 1 - \frac{n_c}{n_a}$ is a cooperation network.

In an environment in which $\mu < 1$, for any finite n_a and n_c, the exclusion of agents from the market lowers the expected number of interactions. To see why, note that even if $n_a > n_c$, in periods in which more than n_c agents are active, the number of interactions is at most n_c whereas if less than n_c of the connected agents are active, the number of interactions is bounded above by the number of the connected agents who are active. In (asymptotically) large markets, some agents might still be excluded from any network that facilitates full cooperation. However, the welfare problem is resolved. As long as condition (5) holds, there exists an (asymptotically) welfare maximizing network that facilitates full cooperation.

Example 2 (cont.) *Suppose there exist positive integers d_C and d_A, and $\phi_a \in (0,1)$ s.t. $Pr(d_a = d_A) = \phi_a$, $Pr(d_a = 0) = 1 - \phi_a$, and $Pr(d_c = d_C) = 1$. Let $\widetilde{n_a}$ be the number of agents who have degree d_A. By construction, $\widetilde{n_a} = \phi_a \cdot n_a = \frac{n_c \cdot d_C}{d_A}$* and *as long as $\frac{n_c}{\widetilde{n_a}} \geq \mu$, condition 5 implies that $\frac{\delta(\pi-\gamma)}{1-\delta} \cdot \frac{n_c}{\widetilde{n_a}} - \gamma > 0$. Now consider ϕ_a^* such that $\widetilde{n_a} = \frac{n_c}{\mu}$ and note that fixing ϕ_a implies a fixed ratio $\frac{d_C}{d_A}$. Then by Theorem 3,* $\lim_{d_C \to \infty} \lim_{n_c \to \infty} \left(\frac{E[V(\mathcal{D}_{GF}(n_c))]}{n_c} | \phi_a^* \right) = 1$, *and*

$$\lim_{d_C \to \infty} \lim_{n_c \to \infty} \left(IC\left(\mathcal{D}_{GF}(n_c)\right) - \left[\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \mu - \gamma \right] | \phi_a^* \right) \geq 0$$

which guarantee that given large enough number of clients and agents, there exists a network that facilitates full cooperation and (asymptotically) the maximal number of interactions possible.

10 Conclusion

This paper presents a model of repeated games in two-sided networks with reputation networks that allow clients to share information about past transactions. The model allows us to vary separately the interaction network between clients and agents and the reputation network between clients, and examine how the quality of the reputation network affects the ability to sustain cooperation in any given interaction network.

More broadly, we make both a new *methodological* contribution in the form of a novel method for moving beyond the assumption that a player's payoff function depends only on the actions taken by her immediate neighbors, and an *applied* contribution in the form of a study of repeated interactions and community enforcement in networked markets with clients and agents. At the core of our methodological contribution is a new method for reducing questions about the global properties of a networked market to questions about the local properties of the network. This allows us to analyze large networks *as if* they were small. Our methodology can be applied to various economically important network interactions as long as (on the equilibrium path) the effect of one economic agent on another is a decreasing function of the network distance between the two agents. In this environment, the introduction of incomplete knowledge of the network structure allows us to approximate any small neighborhood of the network with a simple representative network - a random tree. Notably, both Theorem 1 and Lemma 3 can be modified to include networks that are not two-sided; allowing for a richer set of interactions.

By applying our methodology to the study of repeated games in networks, we show that while good reputation networks allow for cooperation in markets that could not sustain cooperation otherwise, they do not guarantee cooperation in every market. Surprisingly, a market with high quality reputation networks between clients can exhibit 'barriers to entry' because the number of agents that can be a part of any cooperation network is a bounded function of the number of clients in the market. The exclusion of agents from the market can hinder efficiency in any small market. However, as a market grows, a high quality reputation network that allows for optimal welfare emerges. Such a network facilitates the maximal number of (mutually beneficial) interactions as well as sustains cooperation between every client and agent who are connected.

11 Appendix A: proof of Lemma 3

Lemma 3 has implications that go beyond its role in the analysis of repeated games in networks. For example, Fainmesser (2012b) employs a variant of the Lemma for simplifying the analysis of networked labor markets in static settings. Results of a similar flavor have also been found useful in other disciplines.[28] To this end, we present the proof of Lemma 3 as a stand-alone section and follow the conventions of the graph theoretic literature with respect to notation and definitions. We hope that this will make it easier for our more technical readers to appreciate the generality of the result and to be able to adopt the result or parts of it to be used in further applications.

11.1 Notations and definitions

A graph $\Gamma = (V, E)$ is a set of nodes V and a set of edges E, where each edge $e = (v_1, v_2)$ specifies that there is a connection between nodes v_1 and v_2. To prove Lemma 3 we introduce a particular randomization scheme (which we will soon describe in depth). We first formalize

[28] See Richardson and Urbanke (2008) for an example from coding theory.

the class of graphs over which we randomize, and the different notions of degree distribution (d.d.) that we will use. A graph Γ is bipartite if and only if Γ can be partitioned into two sets (e.g. $A(gents)$ and $C(lients)$) such that all edges contain exactly one node from A and one node from C. A bipartite graph is said to be bicolored if the nodes of the one partite are distinguished from the nodes of the other partite. For example, the bicolored property guarantees that the graph on three nodes in which one agent node is connected to two client nodes is distinguished from the graph on three nodes in which one client node is connected to two agent nodes. We say that a bicolored bipartite graph is labeled if each node in partite A have a distinct label from the set $\{1, ..., n_a\}$, and each node in partite C have a distinct label from the set $\{1, ..., n_c\}$. A graph Γ is rooted if one of the nodes on Γ is labelled in a special way to distinguish it from the graph's other nodes. This special node is called the root of the graph. For two rooted graphs Γ_1, Γ_2, we say that $\Gamma_1 = \Gamma_2$ if the two graphs are isomorphic with respect to the root. For a node v in a graph Γ, recall that d_v denotes the degree (number of neighbors) of v in Γ. Sometimes to make the underlying graph explicit, we use the notation d_v^Γ.

For a graph Γ and a subset of nodes V' of Γ, the subgraph induced by V' will refer to the subgraph of Γ consisting of the nodes V' and all edges in Γ that connect nodes in V'. Recall that for a given node v and depth d, $N_d(v)$ was earlier defined as the set of nodes whose graphical distance from v is at most d. For the remainder of Appendix A, $N_d(v)$ should be read as referring not just to the given set of nodes, but the subgraph induced by that set of nodes. Sometimes, to make the reference graph explicit, we use the notation $N_d^\Gamma(v)$. Also, for a given node v in a graph G, we let $F_G(v)$ denote the set of degrees of the nodes adjacent to v in G. Recall that the set of degrees of a given bipartite graph Γ may be defined in two distinct ways. Let $\mathcal{H}_\Gamma^a(\mathcal{H}_\Gamma^c)$ denote the random variable (r.v.) representing the degree of an agent (client) node selected u.a.r. from all agent (client) nodes. Alternatively, let $\mathcal{G}_\Gamma^a(\mathcal{G}_\Gamma^c)$ denote the r.v. representing the degree of the agent (client) belonging to an edge selected u.a.r. from all edges of Γ.

For concreteness, let us fix some given degree distributions $\mathcal{H}^a, \mathcal{H}^c$ with finite, non-negative support and rational probabilities. We let $m_\mathcal{H}$ denote some integer bound on the support of both $\mathcal{H}^a$ and $\mathcal{H}^c$. Let $\mathcal{G}^a, \mathcal{G}^c$ denote the corresponding degree distributions under the random edge interpretation. Let $\mathcal{G}(n_c)$ denote the set of labeled bicolored bipartite graphs that satisfy d.d. $\mathcal{G}^a, \mathcal{G}^c$, and in which the client partite has n_c nodes. We let n_a denote the corresponding number of nodes in the agent partite (determined uniquely by n_c and $\mathcal{G}^a, \mathcal{G}^c$). Let $\mathcal{R}(n_c)$ denote a graph selected u.a.r. from $\mathcal{G}(n_c)$. Let $\mathcal{R}^A(n_c)\big(\mathcal{R}^C(n_c)\big)$ denote the set of nodes in the agent (client) partite of $\mathcal{R}(n_c)$. Let $\mathcal{F}$ denote the set of vectors $\mathbf{f}$ s.t. $Pr(F_{\mathcal{R}(n_c)}(v) = \mathbf{f}) > 0$ for some $v \in \mathcal{R}(n_c)$ (note that $\mathcal{F}$ is dictated by $\mathcal{H}^a, \mathcal{H}^c, n_c$).

Note that the random graph $\mathcal{R}(n_c)$ has some non-trivial dependencies. Indeed, if one conditions on there being an edge between nodes a and c, the precise effect of this conditioning on the degrees of the other nodes is difficult to characterize exactly; large-scale dependencies are introduced by the condition that the graph has the global structure dictated by $\mathcal{H}^a, \mathcal{H}^c$. In spite of this, we prove that the local structure of $\mathcal{R}(n_c)$ is quite simple, namely that of a tree in which the degrees are chosen i.i.d. Let $\mathcal{T}(d, r)$ denote a rooted depth-d tree generated as follows. The degree of the root equals r. Each node at an even depth $k \leq d-1$ is given an i.i.d. number of children distributed as $\mathcal{G}^c - 1$, and each node at odd depth $k \leq d-1$ is given an i.i.d. number of children distributed as $\mathcal{G}^a - 1$.

Note that to prove Lemma 3, it suffices to show the following.

Lemma 4 *For all* $\mathbf{f} \in \mathcal{F}$ *and trees* T*,*

$$\lim_{n_c \to \infty} \sup_{v \in \mathcal{R}^A(n_c)} \left| Pr(N_d^{\mathcal{R}(n_c)}(v) = T | F_{\mathcal{R}(n_c)}(v) = \mathbf{f}) - Pr(\mathcal{T}(d, d_v) = T | F_{\mathcal{T}(d,d_v)}(v) = \mathbf{f}) \right| = 0.$$

11.2 Configuration method

To analyze $\mathcal{R}(n_c)$ and prove Lemma 3, it will be convenient to analyze the well-known pairing (a.k.a. configuration) method for generating $\mathcal{R}(n_c)$ (see e.g. Greenhill et al. 2006, Section 2). First, construct n_a agent buckets $A_1, A_2, ..., A_{n_a}$ and n_c client buckets $C_1, C_2, ..., C_{n_c}$. Second, for each $d \geq 1$, populate a $Pr(\mathcal{H}^a = d)\big(Pr(\mathcal{H}^c = d)$ fraction of agent (client) buckets with exactly d indistinguishable points. Here we let $|A_i|(|C_j|)$ denote the number of points assigned to bucket $A_i(C_j)$, and $n_{c,p}$ $(n_{a,p})$ denote the total number of client (agent) points as dictated by n_c, $\mathcal{H}^a$, and $\mathcal{H}^c$. Third, select a matching $\mathcal{M}(n_c)$ between the agent points and the clients points u.a.r. Fourth, construct a labeled bicolored bipartite graph $R'(n_c)$ such that there are n_c client nodes, n_a agent nodes, and an edge connecting agent node a_i and client node c_j iff at least one point belonging to agent bucket A_i was matched to a point belonging to client bucket C_j. Note that it is possible that in $\mathcal{M}(n_c)$, there exist buckets A_i, C_j such that two points in A_i are connected to two points in C_j, in which case the d.d. of $\mathcal{R}'(n_c)$ need not be the same as that of $\mathcal{R}(n_c)$.

Our approach to proving Lemma 4 will be to first prove an analogue (but without the conditioning involving $\mathcal{F}$) for $\mathcal{R}'(n_c)$.

Lemma 5 *For all trees* T*,* $\lim_{n_c \to \infty} \sup_{v \in \mathcal{R}'^A(n_c)} |Pr(N_d^{\mathcal{R}'(n_c)}(v) = T) - Pr(\mathcal{T}(d, d_v) = T)| = 0$*.*

Proof. Note that we may construct the random matching $\mathcal{M}(n_c)$ in the following manner. First, we pick an arbitrary agent or client point p_1 of our choice. Then, if p_1 was an agent point, we select a point p_2 u.a.r. from all client points. Alternatively, if p_1 was a client point, we select a point p_2 u.a.r. from all agent points. We then add edge (p_1, p_2) to $\mathcal{M}(n_c)$; eliminate p_1 and p_2 from the set of remaining points; and repeat until all points are matched. It follows that we may construct $\mathcal{M}(n_c)$ by selecting the points in an order such that for any bucket A_i of our choosing, $N_d(a_i)$ is 'generated first'. Roughly speaking, we first pair off those points whose buckets will eventually correspond to neighbors of an agent a_i in $\mathcal{R}'(n_c)$; we then pair off those points whose buckets will eventually become neighbors of neighbors of a_i in $\mathcal{R}'(n_c)$, etc. More precisely, we may construct the matching $\mathcal{M}(n_c)$ using the following algorithm. We proceed through a series of stages, indexed by k. We will decide which point we pair off next (more precisely the bucket containing that point) by assigning the buckets labels as the algorithm proceeds.

$RANDGEN$:

Initialize: $k = 1$. Assign bucket A_i the label 1.
While there exists at least one unmatched point:

While there exists at least one bucket with label k:

Select a bucket U u.a.r. from all buckets with label k:
Select an unmatched point p u.a.r. from U:
Select an unmatched point p' u.a.r. from all unmatched agent (client) points;
Add edge (p, p') to $\mathcal{M}(n_c)$;
Remove points p, p' from the set of remaining points;
Assign the bucket containing point p' the label $k+1$;
If there does not exist a bucket with label $k+1$ containing at least one unmatched point:
Select a bucket U u.a.r. from all agent (client) buckets with ≥ 1 unmatched point;
Assign bucket U label $k+1$;
$k = k + 1$;

A simple proof by contradiction shows that $RANDGEN$ always terminates, and a simple induction shows that no bucket is ever assigned two different labels. Note that since each time we pick a point we match it u.a.r. to a remaining point of the 'other' partite, $RANDGEN$ indeed returns a matching distributed u.a.r.

Let $E_{i,\Delta}$ be the event that no bucket with label $k \leq \Delta + 1$ was assigned its label more than once.

Observation 1 *Conditional on the event $E_{i,\Delta}$, $N_{\Delta}^{\mathcal{R}'(n_c)}(a_i)$ is acyclic.*

By a simple induction, at most $2(m_{\mathcal{H}}+m_{\mathcal{H}}(m_{\mathcal{H}}-1)+m_{\mathcal{H}}(m_{\mathcal{H}}-1)^2+...+m_{\mathcal{H}}(m_{\mathcal{H}}-1)^{\Delta-1}) \leq 2\Delta m_{\mathcal{H}}^{\Delta}$ points are matched while $k \leq \Delta$. Let p_1, p_2 be any two points belonging to agents' buckets matched during stage $k \leq \Delta$ for k even. Then the probability that p_1, p_2 were matched to points q_1, q_2 belonging to the same client's bucket is at most $\frac{m_{\mathcal{H}}-1}{n_{c,p}-2\Delta m_{\mathcal{H}}^{\Delta}}$. Indeed, w.l.o.g. assuming p_1 was matched first (with q_1), there are at most $m_{\mathcal{H}}-1$ points out of at least $n_{c,p} - 2\Delta m_{\mathcal{H}}^{\Delta}$ remaining points which q_2 could be matched to so that q_1, q_2 belong to the same bucket. Since there are at most $\binom{2\Delta m_{\mathcal{H}}^{\Delta}}{2}$ pairs of points such that both are matched during stage $k \leq \Delta$, it follows from a union bound that[29]

$$Pr(E_{i,\Delta}) \geq 1 - \binom{2\Delta m_{\mathcal{H}}^{\Delta}}{2} \frac{m_{\mathcal{H}}-1}{n_{c,p} - 2\Delta m_{\mathcal{H}}^{\Delta}} = 1 - O(\frac{1}{n_c}). \tag{6}$$

Let U be any agent bucket assigned label $k \leq \Delta$, and p any point in U that is matched during stage k. It follows from (6) and the previous discussion that for any i, regardless of the value of d_{a_i} and the actions taken by $RANDGEN$ before p was matched, the probability that p is matched to a point q contained in a bucket C_i satisfying $d_{c_i}^{\mathcal{R}'(n_c)} = l$ is at least $Pr(\mathcal{H}^c = l) - O(\frac{1}{n_c})$. Similarly, the probability that p is matched to a point q contained in a bucket C_i satisfying $d_{c_i}^{\mathcal{R}'(n_c)} = l$ is at most $Pr(\mathcal{H}^c = l) + O(\frac{1}{n_c})$. We note that corresponding bounds hold with the role of clients and agents interchanged. It follows that the number of points in the bucket chosen next by $RANDGEN$ is asymptotically independent and identically distributed, where the associated distributions (which depend only on whether the current bucket is a client or agent bucket) correspond to $\mathcal{H}^a, \mathcal{H}^c$. Lemma 5 then follows

[29]The union bound is also known as Boole's inequality: for any finite or countable set of events, the probability that at least one of the events happens is no greater than the sum of the probabilities of the individual events.

from a standard coupling argument, in which we construct $\mathcal{T}(d, d_v)$ and $N_d^{\mathcal{R}'(n_c)}(v)$ on the same probability space. ■

11.3 Relating the configuration model back to the original model

We now relate $\mathcal{R}(n_c)$ to $\mathcal{R}'(n_c)$ probabilistically. Namely, it is well-known (see e.g. Greenhill et al. 2006) that

Lemma 6 $\mathcal{R}(n_c)$ *is distributed exactly as* $\mathcal{R}'(n_c)$ *conditioned to belong to the set* $\mathcal{G}(n_c)$.

We now bound the probability that $\mathcal{R}'(n_c)$ belongs to $\mathcal{G}(n_c)$. In particular, it follows from Theorem 1.3 and Lemma 2.1 of Greenhill et. al. (2006) that for the fixed degree distributions $\mathcal{H}^a, \mathcal{H}^c$,

Lemma 7 $\lim_{n_c \to \infty} Pr\left(\mathcal{R}'(n_c) \in \mathcal{G}(n_c)\right) > 0$.

11.4 Completing the proof of Lemma 4

The only remaining hurdle to proving Lemma 4 is to 'reincorporate' the conditioning involving $\mathcal{F}$. This can be proven directly by computing the relevant conditional probabilities. However, we offer an alternative proof that is more general. We show that for almost all graphs in $\mathcal{G}(n_c)$, the fraction of nodes whose neighborhood is isomorphic to any given tree T is approximately the same as the probability that a corresponding i.i.d. random tree is isomorphic to T. Therefore, the fact that an agent knows her degree and the degrees of clients connected to her does not affect the agent's posterior over the global network structure, or even over her local network structure that is not included in her explicit knowledge. We do that by proving a concentration result, namely that for any tree T, the variance of the number of agents whose neighborhood looks like T in $\mathcal{R}'(n_c)$ goes to zero as n_c goes to infinity.

Lemma 8 *For any rooted tree* T, $Var[n_a^{-1} \sum_{a_i} I(N_\Delta^{\mathcal{R}'(n_c)}(a_i) = T)] = O(\frac{1}{n_a})$.

Proof. After expanding the variance using its definition as the difference between the expected value of the square and the square of the expectation, the only non-trivial step in proving Lemma 8 is bounding the covariance of the indicators $I(N_\Delta^{\mathcal{R}'(n_c)}(a_i) = T)$ and $I(N_\Delta^{\mathcal{R}'(n_c)}(a_j) = T)$ for (arbitrary) nodes a_i, a_j. To analyze this covariance, we consider implementing $RANDGEN$ in a slightly modified manner- namely, we generate 'both' $N_\Delta^{\mathcal{R}'(n_c)}(a_i)$ and $N_\Delta^{\mathcal{R}'(n_c)}(a_j)$ 'first'. More precisely, let $RANDGEN'$ be the algorithm that is equivalent to $RANDGEN$, except at initialization both buckets A_i and A_j are assigned the label 1. The covariance of $I(N_\Delta^{\mathcal{R}'(n_c)}(a_i) = T)$ and $I(N_\Delta^{\mathcal{R}'(n_c)}(a_j) = T)$ is then bounded by analyzing $RANDGEN'$ to show that $N_\Delta^{\mathcal{R}'(n_c)}(a_i)$ and $N_\Delta^{\mathcal{R}'(n_c)}(a_j)$ are asymptotically independent (in an appropriate sense). The analysis proceeds very similarly to our proof of (6), and we omit the details. ■

12 Appendix B: additional proofs

Lemma 1 - Proof. We prove the Lemma for the deterministic case where an agent a has a particular belief $\mathcal{D}_a$ that places probability 1 on the network $N = (G, R)$. The extension for stochastic beliefs follows simply by adding the expectation operation when applicable.

Denote by σ^t the order in which edges are chosen in period t. We can represent agents' strategies in the following way. At the start of each period t each active agent a constructs a *quality function* $Q_a^t : N_1(a) \to \{0, 1\}$, where $Q_a^t(c) = 1(0)$ implies that conditional on a not having already interacted with another client *and* client c not having already interacted by the time the edge (c, a) is chosen in σ^t, a will cooperate with c.

Assume that all other agents $a' \neq a$ always cooperate. Without loss of generality, any strategy of a can be described as a mapping $f^a(t, \{\sigma^\tau\}_{\tau=1}^{t-1}) = Q_a^t$, $t \geq 1$. This is true independent of whether agents observe $\{\sigma^\tau\}_{t=1}^{\infty}$ or not. This follows by a simple induction since: [1] the only freedom agent a has is to set her function Q_a^t; [2] Q_a^t must be a function of the information available to agent a through stage $t-1$; [3] conditional on all other agents $a' \neq a$ always cooperating this information is fully captured by $\{Q_a^\tau\}_{\tau=1}^{t-1}$ and $\{I^\tau(a, c)\}_{c \in N_1(a), \tau=1\ldots t-1}$; [4] $\{I^\tau(c, a)\}_{c \in N_1(a), \tau=1\ldots t-1}$ is deterministic given $\{Q_a^\tau\}_{\tau=1}^{t-1}$ and $\{\sigma^\tau\}_{\tau=1}^{t-1}$; and [5] Q_a^1 must be a function of $\mathcal{D}_a$ alone.

In fact, we can say more. Note that the periods of the repeated game are probabilistically identical until agent a defects in some interaction. Hence, there always exists an optimal strategy in which $Q_a^t = Q_a^1$ up until the smallest t such that a defects for the first time in period t (which we denote by t_a^1). Similarly, denoting by t_a^k the period in which agent a defects for the k-th time, it follows that there always exists an optimal strategy in which Q_a^t is constant for $t \in [1, t_a^1], (t_a^1, t_a^2], \ldots$.

Let $\mathcal{O}_N$ denote some strategy for a that maximizes the expectation of her total payoff conditional on her having a particular belief $\mathcal{D}_a$ that places probability 1 on the network $N = (G, R)$, and assuming that all other agents $a' \neq a$ always cooperate and all clients use ostracizing strategies. Let $\mathcal{O}_N(Q)$ denote some strategy such that $Q_a^\tau = Q$ for every $\tau \leq t_a^1$, and such that for every $\tau > t_a^1$ the strategy maximizes the expectation of her total payoff conditional on her having a particular belief $\mathcal{D}_a$ that places probability 1 on the network $N = (G, R)$, and assuming that all other agents $a' \neq a$ always cooperate. Let Q^{coop} denote the strategy in which a always cooperates with all clients in all periods. Let $u_N(Q)$ denote the expected total payoff for a due to playing strategy Q conditional on her having belief $\mathcal{D}_a = N$, and assuming that all other agents $a' \neq a$ always cooperate and all clients use ostracizing strategies. For ease of notation, let $\overline{u}_N(Q) \triangleq u_N(\mathcal{O}_N(Q))$, $\overline{u}_N \triangleq u_N(\mathcal{O}_N)$, and $u_N^{coop} \triangleq u_N(Q^{coop})$. For each client c and agent a, let $I^t(c, a)$ denote the indicator of the event that c interacted with a in period t. Let $Pr\big(I_N^t(c, a)\big)$ denote the probability that $I^t(c, a) = 1$ in a network N. Then by the stationarity of the game (until period t_a^1), for any belief N and strategy Q,

$$\begin{aligned}\overline{u}_N(Q) &= \sum_{c\in N_1(a):Q(c)=1} Pr\left(I_N^1(c,a)\right)\left(\pi-\gamma+\delta\cdot\overline{u}_N(Q)\right) \\ &\quad+\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)\left(\pi+\delta\cdot\overline{u}_{N\setminus E_R^a(c)}\right) \\ &\quad+\left(1-\sum_{c\in N_1(a)} Pr\left(I_N^1(c,a)\right)\right)\cdot\delta\cdot\overline{u}_N(Q)\end{aligned} \tag{7}$$

In particular,

$$u_N^{coop}=(\pi-\gamma)\cdot\sum_{c\in N_1(a)} Pr\left(I_N^1(c,a)\right)+\delta\cdot u_N^{coop} \tag{8}$$

It follows that,

$$\left(\overline{u}_N(Q)-u_N^{coop}\right)=\frac{\left(\gamma+\delta\cdot\left(\overline{u}_{N\setminus E_R^a(c)}-u_N^{coop}\right)\right)\cdot\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)}{1-\delta+\delta\cdot\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)} \tag{9}$$

Since $\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)>0$, it follows from (9) that when for every agent a, $\mathcal{D}_a$ consists of a single and fixed network N, the vector $\mathbf{m}=(c,\pi,\delta,\{\mathcal{D}_a\}_{a\in A})$ admits a TCEO if and only if for every agent a and each client $c\in N_1(a)$,

$$\gamma<\delta\cdot\left(u_N^{coop}-\overline{u}_{N\setminus E_R^a(c)}\right).$$

■

Theorem 1 - Proof. We prove Theorem 1 for the following restricted domain of agents' beliefs: for any agent a, the belief $\mathcal{D}_a$ is such that $Pr\left(c\in N_1(a)|\mathcal{D}_a\right)\in\{0,1\}$. Namely, for any client $c\in C$, either agent a believes that she is connected to c with probability 1, or she believes (with probability 1) that she is not connected to c. The proof of the case that $Pr\left(c\in N_1(a)|\mathcal{D}_a\right)\in[0,1]$ follows the same argument but requires additional notation and is omitted.

Equations (7) and (8) imply that

$$u_N^{coop}=(\pi-\gamma)\cdot\frac{\sum_{c\in N_1(a)} Pr\left(I_N^1(c,a)\right)}{1-\delta} \tag{10}$$

and

$$\overline{u}_N = \max_Q \left[\overline{u}_N(Q)\right] = \tag{11}$$
$$= \max_Q \left(\frac{\pi \sum_{c\in N_1(a)} Pr\left(I_N^1(c,a)\right) - \gamma \sum_{c\in N_1(a):Q(c)=1} Pr\left(I_N^1(c,a)\right)}{1-\delta+\delta\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)} + \frac{\delta\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)}{1-\delta+\delta\sum_{c\in N_1(a):Q(c)=0} Pr\left(I_N^1(c,a)\right)} \cdot \overline{u}_{N\setminus E_R^a(c)} \right)$$

Finally, noting that in a network of maximum degree Δ, $Pr\big(I_N^1(c,a)\big) \in [\frac{1}{2\Delta-1}, 1]$ for all edges (c,a), Theorem 1 follows by interpreting Equation (11) as a dynamic program (for computing $\overline{u}_N$), combined with Lemma 2 and a simple induction, and applying the same logic to (10). ■

Proposition 2 - Proof. Consider a randomized matching algorithm that progresses by examining the edges of a network in a random order (selected u.a.r.) and including an edge if no incident edges have already been examined. GG study the properties of exactly this algorithm, which they name $GREEDY$. We first state an important monotonicity property of $GREEDY$, which follows from Proposition 1 of GG and a straightforward induction/coupling argument.

Lemma 9 *Suppose that $\widehat{G}, G$ are (rooted) tree networks, and $\widehat{G}$ can be constructed from G by performing only the two operations: [1] appending (as children) subtrees to nodes at even depth in G in an arbitrary way, and [2] removing (as children) subtrees from nodes at odd depth in G in an arbitrary way (where the depth of the root is 0 by default). Then the probability that $GREEDY$ matches the root of $\widehat{G}$ when run on $\widehat{G}$ is at least the probability that $GREEDY$ matches the root of G when run on G.*

The proof of Proposition 2 then follows from Lemma 9 and interpreting $Pr\big(I_N^1(c,a)\big)$ as the probability that edge (c,a) is selected by $GREEDY$. ■

Theorem 3 - Proof. Since all agents are symmetric, we have that $E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right] = n_a \sum_{c\in N_1(a)} \Pr(I^t(c,a)=1)$, where a is any agent. Furthermore, because only those agents who are active can trade, we have the further refinement

$$E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right] = \mu n_a \sum_{c\in N_1(a)} \Pr(I^t(c,a)=1 |\ a \text{ is active at } t). \tag{12}$$

Let $T^1(d, d_A, d_C)$ denote the rooted depth-d tree (with root r^1) s.t. the root has d_A children, each non-leaf node at odd depth has $d_C - 1$ children, and each non-leaf node at even depth has $d_A - 1$ children. For $0 < \mu < 1$, let $T^1(d, d_A, d_C, \mu)$ denote the random rooted depth-d tree (with root r^1) constructed by taking $T^1(d, d_A, d_C)$ and deleting each agent (other than r^1) w.p. μ (i.i.d. across agents). For a graph G, let $\mathcal{M}(G)$ denote the random greedy graph matching (on G) constructed by examining the edges of G in a u.a.r. permutation, always including an edge iff no incident edge has already been included. For a node $v \in G$, let $I\big(v \in \mathcal{M}(G)\big)$ denote the indicator for the event that v is matched in G (equivalently v is incident to a selected edge). Then it follows from Lemma 3, Lemma 6 of GG, and (12) that for any $\epsilon > 0$, there exist $N_{\epsilon,d_A,d_C,\mu}$, and $d_{\epsilon,d_A,d_C,\mu}$ (depending only on ϵ, d_A, d_C, μ) s.t. for all $n_a, n_c \geq N_{\epsilon,d_A,d_C,\mu}$ and $d \geq d_{\epsilon,d_A,d_C,\mu}$,

$$|\frac{E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right]}{n_a} - \mu \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C, \mu)\big)\right)| < \epsilon, \tag{13}$$

and thus

$$\frac{E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right]}{n_a} \geq \mu \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C, \mu)\big)\right) - \epsilon. \tag{14}$$

It follows from Lemma 9 that $\Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C, \mu)\big)\right) \geq \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C)\big)\right)$, since deleting agents is equivalent to removing (as children) subtrees from client nodes. Combining with (14), we find that for all $n_a, n_c \geq N_{\epsilon,d_A,d_C,\mu}$ and $d \geq d_{\epsilon,d_A,d_C,\mu}$

$$\frac{E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right]}{n_a} \geq \mu \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C)\big)\right) - \epsilon. \tag{15}$$

We now treat two cases. First, suppose $d_A \geq d_C$. Then it follows from Lemma 9 that

$$\Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C)\big)\right) \geq \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_A)\big)\right). \tag{16}$$

But it follows from Corollary 6 of GG (in light of Lemma 6 of GG) that for all $d_A \geq 3$,

$$\lim_{d \to \infty} \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_A)\big)\right) = 1 - (d_A - 1)^{-\frac{d_A}{d_A - 2}}. \tag{17}$$

Thus since $\frac{d_A}{d_A-2} \geq 1$, we have that

$$\lim_{d \to \infty} \Pr\left(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_A)\big)\right) \geq 1 - \frac{1}{d_A - 1}. \tag{18}$$

Combining (15),(16), and (18) demonstrates the Theorem for the case $d_A \geq d_C$.

Now, suppose $d_C \geq d_A$. It follows from Lemma 3 that for any fixed d, ϵ, d_C, d_A, there exists a bipartite graph $G(d, \epsilon, d_C, d_A)$ (with partites C, A) s.t.: 1. all nodes in partite C have degree d_C and all nodes in partite A have degree d_A, and 2. a $1-\epsilon$ fraction of nodes in partite A (partite C) have depth-d neighborhoods isomorphic to $T^1(d, d_A, d_C)\big(T^1(d, d_C, d_A)\big)$. By Lemma 6 of GG, for any fixed ϵ, d_C, d_A we may select a sufficiently large $d \triangleq d(\epsilon, d_C, d_A)$ s.t. for any node a belonging to the (at least) $(1-\epsilon)|A|$ nodes of partite A with depth-d neighborhoods isomorphic to $T^1(d, d_A, d_C)$, $|\Pr\left(a \in \mathcal{M}\big(G(d, \epsilon, d_C, d_A)\big)\right) - \Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C)\big)\right| < \epsilon$.

Also, for any node c belonging to the (at least) $(1-\epsilon)|C|$ nodes of partite C with depth-d neighborhoods isomorphic to $T^1(d, d_C, d_A)$, $|\Pr\Big(c \in \mathcal{M}\big(G(d, \epsilon, d_C, d_A)\big)\Big) - \Pr\big(r_1 \in \mathcal{M}\big(T^1(d, d_C, d_A)\big)| < \epsilon$. Combining the above, we find that for the graph $G(d, \epsilon, d_C, d_A)$,

$$|E[\sum_{c\in C} I\Big(c \in \mathcal{M}\big(G(d, \epsilon, d_C, d_A)\big)\Big)] - |C|\Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_C, d_A)\big)\Big)| \leq 2\epsilon|C|, \tag{19}$$

and

$$|E[\sum_{a\in A} I\Big(a \in \mathcal{M}\big(G(d, \epsilon, d_C, d_A)\big)\Big)] - |A|\Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C)\big)\Big)| \leq 2\epsilon|A|. \tag{20}$$

Note that since the number of matched nodes in partite A always equals the number of matched nodes in partite C, one has $E[\sum_{c\in C} I\Big(c \in \mathcal{M}\big(G(d, \epsilon, d_C, d_A)\big)\Big)] = E[\sum_{a\in A} I\Big(a \in \mathcal{M}\big(G(d, \epsilon, d_C, d_A)\big)\Big)]$. It thus follows from (19) and (20) that

$$|\Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_C, d_A)\big)\Big) - \frac{d_C}{d_A}\Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C)\big)\Big)| \leq 2\epsilon(1 + \frac{d_C}{d_A}), \tag{21}$$

and

$$\lim_{d\to\infty} \Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_C, d_A)\big)\Big) = \frac{d_C}{d_A}\lim_{d\to\infty} \Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_A, d_C)\big)\Big). \tag{22}$$

Combining with (15), we find that for any fixed ϵ, d_A, d_C, μ there exist $N'_{\epsilon,d_A,d_C,\mu}$, and $d'_{\epsilon,d_A,d_C,\mu}$ s.t. for all $n_a, n_c \geq N'_{\epsilon,d_A,d_C,\mu}$, $d \geq d'_{\epsilon,d_A,d_C,\mu}$,

$$\frac{E\left[V\left(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C)\right)\right]}{n_a} \geq \mu\frac{d_A}{d_C}\Pr\Big(r_1 \in \mathcal{M}(T^1(d, d_C, d_A))\Big) - \epsilon. \tag{23}$$

It follows from Lemma 9 that $\Pr\Big(r_1 \in \mathcal{M}\big(T^1(d, d_C, d_A)\big)\Big) \geq \Pr\big(r_1 \in \mathcal{M}\big(T^1(d, d_C, d_C)\big)\Big)$. Combining with (18) (replacing d_A by d_C) and taking limits demonstrates the theorem for the case $d_C \geq d_A$. ■

Corollary 2 - Proof. Part 1: Let $\mu = 1$ and $\kappa \geq 1$. Let $N = (G, R)$ be any network that satisfy the following: [1] for every $c \in C$, $d_c = 1$; and [2] $\max\{d_a\}_{a\in A} = 1$. The network N consists of $\overline{n}_c$ client-agent pairs and $\kappa\overline{n}_c - \overline{n}_c$ agents that are not connected to any client (R can be chosen arbitrarily). Let ΔN put probability 1 on the network N. Then,

$$E[V(\Delta N)] = \overline{n}_c \text{ and } \min_{a\in A} IC(\gamma, \pi, \mu, \delta, \Delta N) = \frac{\delta(\pi-\gamma)}{1-\delta}\cdot\mu - \gamma > 0. \tag{24}$$

Plugging (24) into the definition of $WL\left(\overline{n}_c, \overline{n}_a, \gamma, \pi, \mu, \delta\right)$ completes the proof. The proof for the case where $\kappa < 1$ is symmetric.

Part 2: Assume by contradiction that for every $\overline{\kappa}$ there exists $\kappa > \overline{\kappa}$ such that $WL(\overline{n}_c, \kappa\overline{n}_c, \gamma, \pi, \mu, \delta) = 0$. Let $\overline{n}_a^t$ be the number of agents that are active in period t. The contradiction assumption implies that there exists ΔN such that in every period $\min\{\overline{n}_c, \overline{n}_a^t\}$ interactions take place and that $\min_{a\in A} IC(\gamma, \pi, \mu, \delta, \Delta N) > 0$. However, given that $\mu < 1$, to satisfy that in every period $\min\{\overline{n}_c, \overline{n}_a^t\}$ interactions take place, ΔN must provide each agent with a positive probability of interacting in every period that she is active. Thus, for $\kappa > \frac{1}{\mu}$, $\min_{a\in A} IC(\gamma, \pi, \mu, \delta, \Delta N) < \frac{\delta(\pi-\gamma)}{1-\delta} \cdot \frac{1}{\kappa} - \gamma$ which is guaranteed to be negative for any $\kappa > \frac{\delta(\pi-\gamma)}{(1-\delta)\gamma}$. This completes the proof by contradiction to $\min_{a\in A} IC(\gamma, \pi, \mu, \delta, \Delta N) > 0$.

Part 3: Let $\kappa \geq \frac{1}{\mu}$ and $n_a = \frac{1}{\mu} \cdot \overline{n}_c$. Let ΔN assign identical probability to any network that is possible conditional on the following: [1] $d_a = d_A$ for exactly n_a agents and $d_a = 0$ for $\overline{n}_a - n_a$ agents; and [2] $d_c = d_C = \frac{1}{\mu} \cdot d_A$ for every $c \in C$. Combining Theorem 3 and Equation (3) we get that

$$\lim_{d_A\to\infty} \lim_{\overline{n}_c\to\infty} \frac{E\left[V\left(\Delta N\right)\right]}{\overline{n}_c} \geq 1 \ , \ \text{and} \ \lim_{d_A\to\infty} \lim_{\overline{n}_c\to\infty} IC\left(\gamma, \pi, \mu, \delta, \Delta N\right) \geq \frac{\delta(\pi-\gamma)}{1-\delta} \cdot \mu - \gamma > 0,$$

which completes the proof. The proof for the case where $\kappa < \frac{1}{\mu}$ is much simpler and follows a similar logic and therefore omitted. ■

13 References

Ali S. Nageeb and David A. Miller (2012), "Enforcing Cooperation in Networked Societies," working paper.

Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou (2006), "Who's Who in Networks. Wanted: The Key Player," Econometrica, 74, 1403-1417.

Balmaceda, Felipe and Juan F. Escobar (2011), "Self Governance in Social Networks of Information Transmission," working paper.

Banerjee, Abhijit V. and Andrew F. Newman (1993), "Occupational Choice and the Process of Development," Journal of Political Economy, vol. 101(2), pages 274-98.

Barabási, Albert-László and Réka Albert (1999), "Emergence of Scaling in Random Networks," Science, Vol. 286. no. 5439, pp. 509 - 512.

Bayati, Mohsen, David Gamarnik, Dimitriy A. Katz, Chandra Nair, and Prasad Tetali (2007), "Simple Deterministic Approximation Algorithms for Counting Matchings," Proceedings of the thirty-ninth annual ACM symposium on Theory of computing (STOC '07), pp. 122-127.

Boehm, Christopher (1999), "Hierarchy in the Forest: The Evolution of Egalitarian Behavior," Harvard University Press, Cambridge.

Bramoullé, Yann, Martin D'Amours, and Rachel Kranton (2010), "Strategic Interaction and Networks," working paper.

Campbell, Arthur (2011), "Tell your Friends! Word of Mouth and Percolation in Social Networks," working paper.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003), "Persuasion Bias, Social Influence, and Unidimensional Opinions" The Quarterly Journal of Economics, 118, 909-968.

Durkheim, Emile (1933), "The Division of Labor in Society," Free Press, New York.

Ellison, Glenn (1994), "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," Review of Economic Studies, 61, 567–588.

Esfahani, Hadi S. and Djavad Salehi-Isfahani (1989), "Effort Observability and Worker Productivity: Towards an Explanation of Economic Dualism," Economic Journal, vol. 99(397), pages 818-36.

Fafchamps, Marcel (1996), "The Enforcement of Commercial Contracts in Ghana," World Development, 24, 427-448.

Fainmesser, Itay P. (2012a), "Community Structure and Market Outcomes: A Repeated Games in Networks Approach," American Economic Journal: Microeconomics, forthcoming.

Fainmesser, Itay P. (2012b), "Social Networks and Unraveling in Labor Markets," working paper.

Fainmesser, Itay P. (2012c), "Intermediation in (Un)observable Financial Networks," working paper.

Fainmesser, Itay P. and David A. Goldberg (2011), "Bilateral and Community Enforcement in a Networked Market with Simple Strategies," working paper.

Francis, Hywel (1985), "The law, oral tradition and the mining community," Journal of Law and Society. 12, 267–271.

Goodall, Jane (1986), "Social rejection, exclusion, and shunning among the Gombe chimpanzees," Ethnology and Sociobiology, 7, 227–236.

Gordon, Michael A. (1975), "The labor boycott in New York City, 1880–1886," Labor History, 16, 184–229.

Gruter, Margaret and Roger D. Masters (1986), "Ostracism as a social and biological phenomenon: an introduction," Ethnology and Sociobiology, 7, 149–158.

Galeotti, Andrea, Sanjeev Goyal, Matthew O. Jackson, Fernando Vega-Redondo, and Leeat Yariv (2010), "Network Games," Review of Economic Studies, Vol. 77 (1), pp 218-244.

Gamarnik, David and David Goldberg (2010), "Randomized Greedy Algorithms for Independent Sets and Matchings in Regular Graphs: Exact Results and Finite Girth Corrections," Combinatorics, Probability and Computing, vol. 19, Issue 1, pp.61-85.

Gaspart, Frédéric and Erika Seki (2003), "Cooperation, status seeking and competitive behaviour: theory and evidence," Journal of Economic Behavior and Organization, 51, 51–77.

Golub, Benjamin and Matthew O. Jackson (2010), "Naive Learning in Social Networks: Convergence, Influence and the Wisdom of Crowds," American Economic Journal: Microeconomics, 2(1), 112–49.

Goring, Frank, Jochen Harant, Dieter Rautenbach, and Ingo Schiermeyer (2009), "Locally Dense Independent Sets in Regular Graphs of Large Girth - An Example of a New Approach," Research Trends in Combinatorial Optimization, pp. 163-183.

Goyal, Sanjeev (2007), "Connections: an introduction to the economics of networks," Princeton University Press.

Greenhill, Mckay, and Wang (2006), "Asymptotic enumeration of sparse 0-1 matrices with irregular row and column sums," J. Combinatorial Theory, Ser. A, 113, 291-324.

Greif, Avner (1993), “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition,” American Economic Review, 83, 525-548.

Hardle, Wolfgang, and Alan P. Kirman (1995), "Nonclassical Demand: A Model-Free Examination of Price-Quantity Relations in the Marseille Fish Market," Journal of Econometrics, 67, 227-257.

Jackson, Matthew O. and Leeat Yariv (2007), "Diffusion of Behavior and Equilibrium Properties in Network Games," The American Economic Review (Papers and Proceedings), Volume 97(2), pages 92-98.

Jackson, Matthew O. (2008), "Social and Economic Networks," Princeton University Press.

Jackson, Matthew O., Tomas Rodriguez-Barraquer, and Xu Tan (2011), "Social Capital and Social Quilts: Network Patterns of Favor Exchange," American Economic Review, forthcoming.

Kandori, Michihiro (1992), “Social Norms and Community Enforcement,” Review of Economic Studies, 59, 63–80.

Kanoria, Yashodhan, Mohsen Bayati, Christian Borgs, Jennifer Chayes, and Andrea Montanari (2010), "Fast Convergence of Natural Bargaining Dynamics in Exchange Networks," forthcoming ACM-SIAM Symposium on Discrete Algorithms (SODA), 2011.

Karlan, Dean, Tanya Rosenblat, Markus M. Mobius, and Adam G. Szeidl (2009), "Trust and Social Collateral," Quarterly Journal of Economics, Vol. 124, No. 3, pp. 1307-1361.

Kets, Willemien (2011), "Robustness of Equilibria in Anonymous Local Games," Journal of Economic Theory, 146, pp. 300–325.

Kinateder, Markus (2008), "Repeated Games Played in a Network," FEEM Working Paper No. 22.2008.

Kirman, Alan P., and Nicolaas J. Vriend (2000), "Learning to Be Loyal. A Study of the Marseille Fish Market", in Domenico Delli Gatti, Mauro Gallegati, and Alan P. Kirman "Interaction and Market Structure: Essays on Heterogeneity in Economics," Springer.

Krause, Manfred (1996), "A simple proof of the Gale-Ryser Theorem," The American Mathematical Monthly, Vol. 103, No. 4 , 335-337.

Kurzban, Robert and Mark R. Leary (2001), "Evolutionary origins of stigmatization: the functions of social exclusion," Psychological Bulletin, 127, 187–208.

Lancaster, Jane B. (1986), "Primate social behavior and ostracism," Ethnology and Sociobiology, 7, 215–225.

Lauer, Joseph and Nicholas C. Wormald (2007), "Large Independent Sets in Regular Graphs of Large Girth," Journal of Combinatorial Theory, Series B, 97, pp. 999-1009.

Lippert, Steffen and Giancarlo Spagnolo (2006), "Networks of Relations and Social Capital," SSE/EFI Economics and Finance Working Paper No. 570.

Macaulay, Stewart (1963), "Non-Contractual Relations in Business: A Preliminary Study," American Sociological Review, Vol. 28, No. 1, pp. 55-67.

Mahdi, Niloufer Q. (1986), "Pukhtunwali: ostracism and honor among the Pathan Hill tribes," Ethnology and Sociobiology, 7, 295–304.

Mailath, George J. and Larry Samuelson (2006), "Repeated games and reputations : long-run relationships," Oxford University Press, New York, NY.

McMillan, John and Christopher Woodruff (1999), "Interfirm Relationships and Informal Credit in Vietnam" Quarterly Journal of Economics, 114, 1285-1320.

Mihm, Maximilian, Russell Toth, and Corey Lang (2009), "What Goes Around Comes Around: A theory of indirect reciprocity in networks," Cornell Center for Analytical Economics working paper #09-07.

Nava, Francesco and Michele Piccione (2012), "Efficiency in Repeated Two-Action Games with Local Monitoring " working paper.

Richardson, Tom, and Ruediger Urbanke (2008), "Modern Coding Theory," Cambridge University Press.

Vega-Redondo, Fernando (2006), "Building Up Social Capital in a Changing World," Journal of Economic Dynamics and Control Vol. 30, pp. 2305–2338.

Weber, Theophane (2010), "Correlation Decay and Decentralized Optimization in Graphical Models," Ph. D. dissertation, Massachusetts Institute of Technology. URI: http://hdl.handle.net/1721.1/58079.

Weisbuch Gerard, Alan P. Kirman, and Dorothea Herreiner (1996), "Market Organisation and Trading Relationships", The Economic Journal, 110, 411 - 436.

Wiessner, Polly (2005), "Norm enforcement among Ju/'hoansi bushmen. A case of strong reciprocity?" Human Nature, 16, 115–145.

Williams, Kipling D. (2001), "Ostracism: The power of Silence," Guilford Publications, New York.

Wormald, Nicholas C. (1999), "Models of random regular graphs," in Surveys in Combinatorics, J.D. Lamb and D.A. Preece eds. London Mathematical Society Lecture Note Series, vol 276, pp. 239-298. Cambridge University Press, Cambridge.