

Cooperation in Partly Observable Networked Markets*

Itay P. Fainmesser[†] David A. Goldberg[‡]

June 24, 2015

Abstract

We present a model of repeated games in large two-sided networks between clients and agents in the presence of third-party observability 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 enforce-

*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.

[†]The Johns Hopkins Carey Business School and The Economics Department, The Johns Hopkins University, Baltimore, MD 21202. E-mail: *Itay_Fainmesser@jhu.edu*

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

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

ment mechanisms include *personal* and *community enforcement* mechanisms.² It is by now widely recognized that personal enforcement is highly effective when two parties interact frequently, and need each other in order to interact at all. That is, when *bilateral interaction opportunities are exclusive and frequent*. Third-party observability can facilitate community enforcement to overcome the limitations of personal enforcement.³ However, despite the abundance of research on repeated games and community enforcement, the frequency and especially the exclusivity of interaction opportunities are mostly treated as ‘black boxes.’ Similarly, the level of third-party observability is modeled for highly specialized cases.

To be specific, in much of the repeated games and community enforcement literature either any two parties interact in every period, or random matching is assumed.⁴ This is because the existing literature limits attention to setups in which interactions are non-competing. That is, whether two individuals have an opportunity to interact with each other is independent of whether and when each of them interacts with additional individuals. In contrast, we note that in many market environments interactions are competing—a buyer may want to purchase a limited quantity of a good, and a borrower may need a loan of a up to some given size, independent of how many sellers or lenders are in the market. As a result, the frequency and exclusivity of interactions in a two-sided markets with buyers and sellers, or investors and entrepreneurs are inseparable from the structure of the market—the subset of "agents" who each "client" has access to, and who he chooses to interact with.⁵ Second, much of the emerging literature on repeated games in networks consider either public monitoring, or a monitoring structure that is identical to the interaction structure (i.e., an individual observes interactions in which she participates). We note that in many markets each client learns about the outcomes of a different subset of the interactions in the market

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

³See also Kandori (1992) and Greif (1993).

⁴This is true also for papers which study repeated games in networks. For example, Nava and Piccione (2012) take the approach that any two connected individuals interact with each other at every period, whereas in Ali and Miller (2012), and Jackson et al. (2012), the meetings between any two connected individuals are governed by a random process that is independent of the network structure. See section 8.1 for a more detailed discussion of the related literature on repeated games in networks.

⁵The 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).

and that the patterns of third-party observability may not be subject to the same restrictions as the interaction structure.

This paper develops a new model and new techniques for the study of repeated games in a market with clients and agents, in which interactions are competing, and the observability structure is not tied to the interaction patterns (so that both can vary in economically meaningful ways). Initially, each client can interact (e.g. purchase a good from, or make a loan to) with 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. Each agent (client) *is able to interact with at most one* client (agent) in a given period, and 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 an *observability network* R that captures the level of third-party observability in the market (R can be thought of as a reduced form object that captures the diffusion of information that results from Word-Of-Mouth, reputation systems, or any other mechanism that facilitates third-party observability). The combined 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 an observability network R such that $N = (G, R)$ can be sustained indefinitely and allow 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 networks in which a connection between client c and agent a has the interpretation that c is able and willing to interact with a , and that a always cooperates with c .

To answer these questions we focus on a specific family of cooperative equilibria in simple strategies and ask how the interaction structure and third-party observability structure (G and R) affect the existence of such cooperative equilibria for any fixed level of patience. This

is different from much of the existing literature in which the equivalents of G and R are held fixed (and often identical to each other), and in which the goal is to construct strategies that sustain cooperation when individuals are sufficiently patient.⁶

Analyzing the model 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 because 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. Thus, even on the path of a cooperative equilibrium, the future payoffs of an agent depends on the entire network structure (as opposed to only the off path payoffs). 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.

We demonstrate the implications of the local neighborhood result for the study of cooperation by applying it to a specific model of beliefs over the network structure. The suggested model of beliefs 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 (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

⁶E.g. Kandori (1992), Ellison (1994), and more recently Deb and Gonzalez-Diaz (2012).

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.⁷

We note that our random tree characterization is quite general and can be extended to many settings. To apply it to the study of the effect of the interaction and observability structures on cooperation, we focus on cooperative equilibria in ostracizing strategies—equilibria in which agents always cooperate and clients cut their links with (and only with) agents who they observe to have defected. For this family of equilibria, we provide conditions under which an asymptotically large network N can be sustained indefinitely (and thus facilitate full cooperation). We find that adding a large number of links to an observability network R increases the set of interaction networks G that can be sustained indefinitely. 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 for a larger set of discount factors. Finally, we show that a sufficiently dense observability network R guarantees that the fraction of interactions lost due to the incentive constraints goes to zero as the size of the market grows. This is despite the fact that any 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.

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 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 characterize the structure of cooperation networks in this model. The welfare implications of our results are derived in section 7. Section 8 reviews related literature and offers a discussion of the main methodological contributions

⁷Although 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. An exception is Campbell (2012) who applies percolation theory (physics) to the study of monopoly pricing in the presence of Word-Of-Mouth. This is related because percolation theory relies on insights that are directly related to the claim we prove in Lemma 3.

of the paper, as well as additional economic implications.

2 Motivating examples

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

Risky investments Consider a market with investors and 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 may not be verifiable, and the entrepreneur can choose to strategically default on her loan (see also Fainmesser 2012b). 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, and if the entrepreneur is not able to borrow from alternative sources, a threat of not receiving future loans from the investor provides the entrepreneur with the incentives to never strategically default on a loan. Otherwise, if investors share information credibly, or if strategic default is observable by other investors, a threat of ostracizing a defaulting entrepreneur may provide the appropriate incentives. This paper sheds light on the relationship between the structures of the observability network and the lending network in enforcing repayment.

Experience goods and services Consider a market in which after receiving payment, a seller can decide whether to supply high or low quality goods, and may even provide some buyers with high quality goods or services and others 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 considers the possibility that sellers and buyers may not have complete knowledge of the patterns of interaction in the market.

3 Model

Consider a market with a set of clients $C \equiv \{1, 2, \dots, n_c\}$ and a set of agents $A \equiv \{1, 2, \dots, n_a\}$. Time is discrete, and clients and agents live forever and have a common discount factor δ . In any given period, each client (agent) has the capacity to engage with one agent (client) in the following 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 $\beta > \gamma - \varphi$ (hence, cooperation is efficient). If a client (agent) does not engage in any interaction in a given period, her payoff is zero.

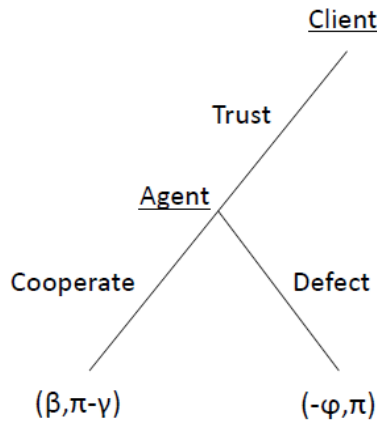


Figure 1: a trust-based interaction.

For ease of notation, whenever possible we keep implicit the dependencies on the parameters of the interaction $\gamma, \pi, \beta, \varphi$, and on the discount factor δ .

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 interact with) as well as clients' decision (which agents each client trusts). More specifically, in any period t , there exists a two-sided network of connections between clients and agents, which captures the expected interaction patterns. We first introduce the notion of a network and then make its economic meaning more precise.

At any time t there is a network captured by a bipartite graph $G^t = \langle C, A, E^t \rangle$ where $E^t \subseteq C \times A$ is the set of edges (or links) between clients and agents. 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 $\{(v_0, v_1), (v_1, v_2), \dots, (v_{l-1}, v_l)\}$ such that $v_0 = v$, $v_l = v'$, and for every $i \in \{1, 2, \dots, l\}$, $(v_{i-1}, v_i) \in G$. 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 , and let $\text{deg}(v) \equiv |N_1(v)|$ 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 $\{(v_0, v_1), (v_1, v_2), \dots, (v_{l-1}, v_l)\}$ such that $v_0 = v_l$. A graph that has no cycles 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 in the networked market. During any period t , all connected clients and agents meet at a random sequencing, i.e., 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.⁸ When a link (c, a) is chosen, c and a meet and engage in the aforementioned trust-based interaction unless either c or a has already interacted (with anyone else) during the same period. 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.

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 links he keeps and which he deletes permanently—the resulting network is G^t . That is, clients can decide who they interact with (or who they trust) by affecting the structure of

⁸Our analysis is independent of whether clients and agents learn the order after it is chosen.

the network.

Several restrictions are imposed on the patterns of interaction by the structure of the game. First, an agent cannot eliminate links (i.e., decline an interaction). It will become clear that for the family of equilibria that we consider, this assumption is without loss of generality. The same is true for the assumption that once a client disconnects from an agent the relationship cannot be revived.

More substantive is that 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.⁹ 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.

3.2 Observability networks

Each client has access to information about the outcomes of all of his past interactions, as well as limited information about other clients' past transactions as captured by the *observability network*. Formally, there is a graph R on the set of clients, where edge (c, c') is present in R if and only if client c observes the outcomes of all of the past interactions of client c' . For simplicity (and with no impact on our results), we assume that the graph R is undirected (i.e., $(c, c') \in R$ implies that $(c', c) \in R$). Let $N_1^R(c)$ denote the set of clients who are connected to client c in R , and $E_R^a(c)$ denote the set of edges (in G) between agent a and clients who observe past transactions of client c $\left(E_R^a(c) \triangleq \bigcup_{c' \in N_1^R(c) \cap N_1(a)} (a, c')\right)$.

We remain agnostic with respect to the underlying mechanism enabling observability, which could be direct observations, Word-Of-Mouth (WOM), reputation systems, or any other mechanism for third-party observability. The network R fully describes which transactions each client observes.

⁹See Jackson et al. (2012) and Fainmesser (2012b) for similar assumptions.

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

3.3 The network structure - knowledge and beliefs

We now develop a general framework of individuals' (clients or 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 for the study of the relative importance of local and global beliefs in calculating agents' expected payoffs.

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.¹² We refer to \mathcal{D}_v as the (first order) *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 $N_d(v)$. Figure 2 offers an example.




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.

For simplicity, we assume throughout that individuals' beliefs are stationary, and that individuals do not update their beliefs on the network structure. This allows us to make a first step towards the analysis of repeated games with incomplete knowledge of the network.¹³ To what extent this assumption is restrictive depends on the specific belief profile that we adopt (e.g., if we follow the most common assumption in the literature that individuals have

¹¹By *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 whose order is much smaller than that of the entire network, which may be arbitrarily large.

¹²A member of the set is a 2-vector $N = (G, R)$ consisting of both an interaction graph G and an observability graph R .

¹³For an alternative approach, see Fainmesser (2012b) who derives an upper bound on players' knowledge of the network when players learn about the network structure from their interactions.

accurate beliefs with respect to the entire network structure, then this assumption is not restrictive).¹⁴ Therefore, it is of interest to note that our main result on the importance of local beliefs (Theorem 1) holds for any beliefs specification.

Higher order beliefs are clearly important, but are not the focus of this paper. We therefore keep the discussion of high order beliefs mostly implicit using the following definition.

Definition 1 *Let V be the set of all individuals (clients and agents), and consider a vector $\mathbf{m} \triangleq (\gamma, \pi, \delta, \{\mathcal{D}_v\}_{v \in V})$. We say that an economy is consistent with \mathbf{m} if the economy can be described by \mathbf{m} and some high order beliefs. We let $\Omega_{\mathbf{m}}$ be the set of all economies consistent with \mathbf{m} , and denote by $\omega_{\mathbf{m}}$ an economy that belongs to the set $\Omega_{\mathbf{m}}$.*

4 Equilibrium

The common practice of *ostracism* is documented in economics, social psychology, anthropology, and biology.¹⁵ An ostracizing strategy requires an individual that observes a defection to defect in future interactions with the defector. In our setup, ostracism implies that if a client observes a defection by an agent, the client ceases to trust that agent. For any given set of parameters (including any fixed discount factor), we are interested in characterizing networks for which there exist cooperative equilibria that rely on the threat of ostracism.

Definition 2 *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 observability network R .*

Definition 3 *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 employ the strategy "always cooperate."¹⁶*

¹⁴For the more specific belief profiles considered in section 6 it is also never the case that an individual holds a belief given a history that is patently inconsistent with it.

¹⁵See 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).

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

If a TCEO exists then cooperation is possible. Therefore, Lemma 1 below provides a simple sufficient condition for cooperation to be possible. Focusing on Nash equilibrium is not without loss of generality. However, it will become clear that many of our results extend directly to the equivalent notion of perfect Bayesian equilibrium with ostracism.¹⁷

To see how focusing on TCEOs simplifies the analysis, 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 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 $\tau < t$. 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).

Formally, suppose that a 's belief puts probability 1 on a network N . 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 \bar{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 $E_{\mathcal{D}_a}[\bar{u}_N]$ be the corresponding expected values given \mathcal{D}_a . Recall also that γ is the immediate additional payoff for an agent who defects, and that $E_R^a(c)$ denotes the set of edges (in G) between agent a and clients who observe past transactions of client c . In other words, $E_R^a(c)$ is the set of edges that a will immediately lose in a TCEO if she defects in an interaction with c . Finally, let

$$IC(\mathcal{D}_a) \triangleq \min_{c \in N_1(a)} \delta(E_{\mathcal{D}_a}[u_N^{coop}] - E_{\mathcal{D}_a}[\bar{u}_N \setminus E_R^a(c)]) - \gamma,$$

and for $\mathbf{m} = (\gamma, \pi, \delta, \{\mathcal{D}_v\}_{v \in V})$

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

The proofs of Lemma 1 and all subsequent results are deferred to the Appendix.

Lemma 1 (*Incentives to Cooperate*) *There exists an economy $\omega_{\mathbf{m}}$ for which there exists a TCEO if and only if $IC(\mathbf{m}) > 0$.*

¹⁷In particular, our results with respect to the network structures that support cooperation when the observability network is the complete one (Corollary 1, Proposition 2, Theorem 3, and Corollary 3), can be restated with the stronger equilibrium notion.

The proof of Lemma 1 consists of showing that in a TCEO, $IC(\mathcal{D}_a)$ is a sufficient statistic for the **Incentives of an agent a to Cooperate** with all of the clients connected to her. Nevertheless, Lemma 1 does not tell us how to compute $IC(\mathcal{D}_a)$, which poses a significant challenge. The obvious difficulty is that $E_{\mathcal{D}_a}[\bar{u}_{N \setminus E_R^a(c)}]$ 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. The main results of this paper, which are presented in sections 5 and 6 alleviate this difficulty.

5 Cooperation based on local beliefs

Our first main result, which provides the foundation for later results, shows that the existence of a TCEO 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*.

Theorem 1 (*Local Beliefs Theorem*) *Consider any $\Delta > 0$ and $\epsilon > 0$, and let \mathcal{D}_a be any belief that puts probability zero on any individual having a degree greater than Δ . Then, there exists a finite constant $d = d(\gamma, \pi, \delta, \Delta, \epsilon)$ **independent** of the size of the entire network of such that*

$$|IC(\mathcal{D}_a) - IC(\mathcal{D}_a^d)| < \epsilon.$$

To better understand the implications of Theorem 1, 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. for any agent a the neighborhood that includes $N_d(a)$), we can also make (asymptotic) comparative statements for the entire network N . Put differently,

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 algorithms for matchings in a graph, and the relationship between the local and global properties of the set of matchings of a graph. The following key lemma 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 $I^t(c, a)$ denote the indicator of the event that c interact with a in period t , and let $Pr(I_N^t(c, a))$ denote the probability that $I^t(c, a) = 1$ in a network N . Note that one may interpret the quantity $Pr(I_N^1(c, a)) \left(Pr(I_{N(a,d)}^1(c, a)) \right)$ as the probability that edge (c, a) is chosen to belong to the random graph matching constructed by examining the edges of $N(N(a, d))$ in a random order (selected u.a.r.) and including an edge if no incident edges have already been included. Noting that this randomized matching construction is 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 $d = d(\Delta, \epsilon)$ such that for any network N in which no individual has a degree greater than Δ ,*

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

Lemma 2 implies that when interactions are competing, 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 are less important.

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 $\frac{1}{2}$. 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. Similarly, the probabilities 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, in the sequence $(1, \frac{1}{2}, \frac{2}{3}, \frac{5}{8}, \frac{19}{30})$ the deviation from the value that is based on the correct belief is monotonically decreasing.

To see how example 1 generalizes to Lemma 2, consider the following. First, note that 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 an agent a and a client c who are connected. 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 $G \setminus H$ 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.

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 an observability 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. To this end, we now consider a specific model of individuals' knowledge and beliefs with respect to the network structure, and use it to demonstrate how Theorem 1 can be used to answer our question.

Our example, a model of beliefs that we call the *Global Fractions (GF)* model, formalizes 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.¹⁸ Each individual has private information about her own local areas of the network. However, due to the random component, an individual has only partial information on the global network structure.

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 r : [1] the *fraction* of clients that have degree r , and [2] the *fraction* of agents that have degree r . Each agent a knows the network formation process and 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 ($\deg(c)$ for all $c \in N_1(a)$), and which of her clients observes the outcomes of her interactions with any of her other clients ($E_R^a(c)$ for all $c \in N_1(a)$). 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.¹⁹

Note that specifying the aforementioned fractions is equivalent (under a simple transfor-

¹⁸For evidence on consistency in degree distributions in large networks, see Barabási and Albert (1999).

¹⁹Our 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).

mation) 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 r (for all r).²⁰ Thus, agents' posteriors can be described as follows. Let $\mathcal{G}^a, \mathcal{G}^c$ be random variables each with *bounded support* on \mathbb{Z}^+ and taking each integer value with a *rational* probability. Then, for any network N in which 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 r is equal to $Pr(\mathcal{G}^a = r)$ ($Pr(\mathcal{G}^c = r)$), the posterior of an agent a (i.e., $\mathcal{D}_{GF}(a|N)$) is that the network is selected u.a.r. from all networks in which there are the same numbers of clients and agents as in N , the degree of a and the degrees (and identities) of clients connected to a are the same degrees (and identities) as in N , the degree distribution in the interaction network (G) is captured by the same $\mathcal{G}^a, \mathcal{G}^c$ as in the network N , and the connections in the observability network (R) between clients who are connected to a are the same as in N .

Furthermore, note that in any network, specifying the number of clients n_c along with the distributions of $\mathcal{G}^a, \mathcal{G}^c$ fully determines the number of agents n_a . Consequently, in order to consider large networks with bounded degrees, we can denote by $\mathcal{D}_{GF}(a|N, n_c)$ the belief of an agent a in a network in which there are n_c clients, the degree of a and the degrees (and identities) of clients connected to a are the same degrees (and identities) as in N , the degree distributions in the interaction network (G) are captured by the same $\mathcal{G}^a, \mathcal{G}^c$ as in the network N , and the connections in the observability network (R) between clients who are connected to a are the same as in N .²¹

Remark 1 *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 capture a scenario in which individuals ‘have no clue’ how networks are formed, but have some information on their attributes.²²*

²⁰Let $P_a(r)$ the proportion of agents with degree r , and let $\bar{r}_a = \sum_r P_a(r) \cdot r$ be the average agent's degree. Then the probability that an agent a in an edge that is chosen u.a.r. has degree r is, $\frac{P_a(r) \cdot r}{\bar{r}_a}$.

²¹Suppose further that $Pr(\mathcal{G}^a = \deg(a)) > 0$, and $Pr(\mathcal{G}^c = \deg(c)) > 0$ for all $c \in N_1(a)$. Then, it is well-known that for any fixed $\mathcal{G}^a, \mathcal{G}^c$ there exists an infinite strictly increasing sequence of integers $\{n_c\}$ s.t. there exists at least one network as required. 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.

²²An algorithm for generating valid large random graphs with arbitrary degree distributions exists and is

6.1 Cooperation and network structure

We now focus on our leading example in which agents' knowledge and beliefs are consistent with the GF model and derive conditions for a network N to be a *GF cooperation network*.

Definition 4 (*Cooperation Networks*) Let $\mathbf{m}^{GF}(N) \triangleq (\gamma, \pi, \delta, \{\mathcal{D}_{GF}(a|N)\}_{a \in A})$. We say that a network N is a **GF cooperation network** if and only if there exists an economy that is consistent with $\mathbf{m}^{GF}(N)$ ($\omega_{\mathbf{m}^{GF}(N)}$) for which there exists a TCEO.

The main result of this section states that 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.

Formally, for a given network N with degree distribution captured by $\mathcal{G}^a, \mathcal{G}^c$, agent a , and integer $d \geq 1$, let $T(a, N, d)$ denote the random depth- d tree such that the root r has the same degree as agent a in N , the degrees (and identities) of the children of r are the same as the degrees (and identities) of the clients connected to a in $G(N_1(a))$, all subsequent non-leaf nodes at an even depth have a number of children drawn i.i.d. from the distribution governing $\mathcal{G}^a - 1$, all subsequent non-leaf nodes at an odd depth have a number of children drawn i.i.d. from the distribution governing $\mathcal{G}^c - 1$, the underlying observability network for the clients connected to the root ($N_1(r)$) is R , and for all other clients pairs c and c' , the probability that c and c' are connected in R is $1/2$.

Theorem 2 (*Asymptotic Characterization of Cooperation Networks*) Consider a sequence of networks (N^1, N^2, \dots) with identical finite support degree distributions captured by $\mathcal{G}^a, \mathcal{G}^c$ 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 ,

$$\lim_{d \rightarrow \infty} IC(T(a, N^i, d)) > 0.$$

At the heart of the proof of Theorem 2 is the observation that the local belief of any agent a in a large network is asymptotically identical to the belief that the network looks suggested in the proof of Lemma 3.

locally like a corresponding *simple random tree*.

To be precise, for two random variables X, Y with support on some countable 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 $\mathcal{D}_a, \mathcal{D}'_a$, $TVD(\mathcal{D}_a^d, \mathcal{D}'_a^d) < \epsilon$ 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 networks N and agents a in N , and for all $d > 0$,*

$$\lim_{n_c \rightarrow \infty} TVD(\mathcal{D}_{GF}^d(a|N, n_c), T(a, N, d)) = 0$$

Lemma 3 implies that in an asymptotically large network, for any fixed d , the belief $\mathcal{D}_{GF}^d(a|N, n_c)$ converges to a belief on a random tree. Thus, to make use of the convenient structure of a random tree in our setup, one must establish that an agent’s depth- d belief is a sufficient statistic to determine the incentives to cooperate. Theorem 1 completes this gap with respect to agents’ incentives to cooperate and we are able to derive Theorem 2.

The proof of Lemma 3 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 were 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.²³

²³More generally, the configuration model can be related to many random-graph models of interest (see e.g. Greenhill et al. 2006, Wormald 1999). For more details the reader is referred to the proof in Appendix A.

6.2 Third-party observability and cooperation

In this section we make use of the simplification offered by Theorem 2 and exploit the tree structure to characterize the set of interaction networks G for which there exists an observability network R such that the network $N = (G, R)$ is a GF cooperation network. To this end, 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 observability network is complete—that is, when the observability network is such that any two clients are connected, and each client is informed of transaction outcomes of all other clients. We denote the complete observability network by R^1 .

Proposition 1 (*Weak Monotonicity of Cooperation in Third-Party Observability*) *Consider any interaction network G . Then, for any observability network R , if the network (G, R) is a GF cooperation network then the network (G, R^1) is also a GF cooperation network.*

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 .

To state our next result we recall that Theorem 2 shows that the incentive effect of any change in the degree of any client or agent in a large network can be approximated by considering the effect of the corresponding change in the corresponding random tree. It is also the case that 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 $T' \subseteq T$ and T' is connected). 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.

Proposition 2 (*Monotonicity of Cooperation in Degree*) Suppose that for all $d \geq 1$, the random tree $T' = T(a, (G', R^1), d)$ can be constructed (on the same probability space) from the random tree $T = T(a, (G, R^1), d)$ 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,

$$\lim_{n_c \rightarrow \infty} IC(\mathcal{D}_{GF}(a | (G, R^1), n_c)) > 0 \text{ implies}$$

$$\lim_{n_c \rightarrow \infty} IC(\mathcal{D}_{GF}(a | (G', R^1), n_c)) > 0.$$

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 observation: when $R = R^1$, agents who expect to interact with higher overall probability (larger $\sum_{c \in N_1(a)} \Pr(I_N^1(c, a))$) have stronger incentives to cooperate. To see how this observation leads to Proposition 2, 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 also increasing in the degree of any agent that is connected to any client who is connected to a . The intuition for the latter is subtle: because agents with high degrees are less likely to interact with any 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.²⁴ Put differently, when agents' degrees are large and clients' degrees are small, there are more clients and less agents in the network. 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 observability network. In the leftmost network, agent a expects to interact with probability $\frac{1}{2}$ in any given period. As a

²⁴The intuition applies most directly to trees. However, our random tree characterization highlights the connection between trees and large networks.

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 $\frac{\delta}{1-\delta}(\pi - \gamma) \cdot \frac{2}{3} > \gamma$.

7 Welfare

In this section we show that third-party observability enhances total welfare in markets that rely on trust and cooperation. It does so by allowing for the existence of dense cooperation network. As a result, asymptotic (on the size of the network) efficiency is achievable. However, in finite markets with significantly more agents than clients, efficiency is always bounded away from the optimal outcome.

Since cooperative interactions in our setup are mutually beneficial, total expected welfare can be measured by the expected number of trades in any given period. Formally, for a given network N , let $V(N) = E \left[\sum_{a \in A} \sum_{c \in N_1(a)} Pr(I_N^1(c, a)) \right]$ 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 $V(\Delta N)$ the corresponding expected value given a probability distribution ΔN over networks.

To make the effect of the network structure on the expected number of trades more salient, we generalize our model to allow for stochastic elements in the production technology. 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. For example, if an agent is an entrepreneur then she may not be able to come up with a profitable investment idea in every period. Similarly, if the agent is a producer, she may suffer from exogenous shocks to her production process. The role of the added stochastic elements is to capture a key reason that highly connected interaction networks maximize the number of trades. Notably, all of our results above extend without change to this more general environment.²⁵ Introducing instead a stochastic element on the

²⁵The interested reader is referred to Fainmesser and Goldberg (2011) which is an older draft of this paper and is available on Fainmesser's webpage.

demand side yield qualitatively identical results.

7.1 Welfare in the GF model

As a first step, we establish the relationship between network density and efficiency for a simple family of networks and using the GF model of beliefs. We derive the more general connection between network structure and efficiency in the following section.

For networks in which all agents are symmetric, our discussion in the previous section implies that the expected aggregate number of interactions and the incentives of agents to cooperate are perfectly aligned. Theorem 3 shows that a result of this alignment is that given extensive third-party observability, dense interaction networks are efficient as long as there are sufficiently many clients and sufficiently few agents.

Theorem 3 (*Dense Networks Maximize Welfare and Incentives to Cooperate*) *Consider a sequence of networks $\{N^i = (G^i, R^1)\}_{i=1,2,\dots}$ with increasing sizes (i.e. the numbers of clients and agents in network N^{i+1} are larger than the corresponding numbers in network N^i) in which the degrees of all agents and clients in any interaction network G^i are d_A and d_C respectively. Let n_a^i and n_c^i be the number of agents and clients in network N^i . Then, for any i , N^i is a GF cooperation network if and only if*

$$\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \frac{1}{n_a^i} \cdot V(\mathcal{D}_{GF}(N^i)) - \gamma > 0$$

and

$$\lim_{i \rightarrow \infty} \frac{V(\mathcal{D}_{GF}(N^i))}{\min(\mu \cdot n_c^i \cdot \frac{d_C}{d_A}, n_c^i)} \geq 1 - (\max(d_A, d_C) - 1)^{-1}.$$

Moreover, for all i , $\frac{n_c^i}{n_a^i} = \frac{d_A}{d_C}$. Thus, $\min(\mu \cdot n_c^i \cdot \frac{d_C}{d_A}, n_c^i) = \min(\mu \cdot n_a^i, n_c^i)$, which equals the maximal number of trades possible between n_c^i clients and n_a^i agents when $i \rightarrow \infty$.

Theorem 3 shows 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 . Moreover, The lower bound on the expected number of trades guarantees that 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.

In the next section we show that the observation that the ratio $\frac{n_c}{n_a}$ determines whether repeated interactions can sustain cooperation is more general and goes beyond the GF model.

7.2 Optimal networks

We now derive more general welfare implications that go beyond the example suggested by the GF model. In particular, we compare the first-best (i.e. how many cooperative interactions are possible between \bar{n}_c clients and \bar{n}_a agents?) with the second-best (i.e. how many cooperative interactions are possible in a network which satisfies the incentive constraints for cooperation?). This difference captures the limits of the effectiveness of third-party observability in facilitating cooperation.

Consider a market with \bar{n}_a agents and \bar{n}_c . In the *unconstrained design problem*, a planner chooses a probability distribution over network structures 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 there exists a TCEO with ΔN . Let $n_c(N)$ ($n_a(N)$) be the number of clients (agents) whose degree in G is at least 1. If $n_c(N) < \bar{n}_c$ ($n_a(N) < \bar{n}_a$) we say that $\bar{n}_c - n_c(N)$ clients ($\bar{n}_a - n_a(N)$ agents) are *excluded* from the market in N . Let $N^{uc}(\cdot)$ ($N^c(\cdot)$) be the solution to the unconstrained (constrained) design problem. Then,

$$N^{uc}(\bar{n}_c, \bar{n}_a) = \underset{\Delta N | \bar{n}_c, \bar{n}_a}{argmax} V(\Delta N)$$

and

$$N^c(\bar{n}_c, \bar{n}_a) = \underset{\Delta N | \bar{n}_c, \bar{n}_a}{argmax} \underset{IC(\Delta N) > 0}{V(\Delta N)},$$

where $IC(\Delta N) > 0$ guarantees that there exists a TCEO with ΔN . Thus, the proportion of welfare loss due to the constraints on the structure of cooperation networks is²⁶

$$WL(\bar{n}_c, \bar{n}_a) = 1 - \frac{V(N^c(\bar{n}_c, \bar{n}_a))}{V(N^{uc}(\bar{n}_c, \bar{n}_a))}$$

Now recall that if $\frac{\delta(\pi-\gamma)}{1-\delta} \cdot \mu - \gamma < 0$ then no network sustains cooperation. This is true because even an agent who is guaranteed to interact whenever she is active, and expects to

²⁶If $WL(\bar{n}_c, \bar{n}_a) = 0$, then cooperation networks achieve the first-best in a market with (\bar{n}_c, \bar{n}_a) .

lose her entire future payoff if she defects, will still defect. Consequently, $\frac{\delta(\pi-\gamma)}{1-\delta} \cdot \mu - \gamma < 0$ trivially implies that $WL(n_c, n_a) > 0$ for any n_c, n_a . Corollary 3 covers the more interesting case where $\frac{\delta(\pi-\gamma)}{1-\delta} \cdot \mu - \gamma > 0$.

Proposition 3 (*Asymptotic Efficiency*) *Let $\delta < 1$ and $\frac{\delta(\pi-\gamma)}{1-\delta} \cdot \mu - \gamma > 0$ and consider a market with \bar{n}_c clients and $\kappa\bar{n}_c$ agents. Then,*

1. For any κ , $WL(\bar{n}_c, \kappa\bar{n}_c | \mu = 1) = 0$.
2. Let $\mu < 1$. There exists $\bar{\kappa}(\mu)$ such that $WL(\bar{n}_c, \kappa\bar{n}_c) > 0$ for any $\kappa > \bar{\kappa}$.
3. For any $\kappa \in \mathbb{Q}^+$, $\lim_{\bar{n}_c \rightarrow \infty} WL(\bar{n}_c, \kappa\bar{n}_c) = 0$.

Part 1 of Proposition 3 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 Proposition addresses the case of stochastic production technology. If $\mu < 1$ the maximal number of interactions cannot be achieved if some agents are 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 Proposition 3 is encouraging; in large markets (asymptotic) efficiency is restored. Theorem 3 provides the necessary intuition: no matter how large is $\frac{\bar{n}_a}{\bar{n}_c}$, 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. $Pr(\deg(c) = d_C) = 1$, $Pr(\deg(a) = d_A) = \frac{\bar{n}_c}{\bar{n}_a} \cdot \frac{1}{\mu}$, and $Pr(\deg(a) = 0) = 1 - \frac{\bar{n}_c}{\bar{n}_a} \cdot \frac{1}{\mu}$. Then, the planner achieves (asymptotically in d_A, d_C) both high incentives to cooperate and maximal number of interactions. Notably, 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 8.1.

8 Discussion

8.1 Competing interactions, third-party observability, and barriers to entry

The study of repeated games in networks has seen many recent developments (E.g., Vega-Redondo 2006, Lippert and Spagnolo 2011, Mihm et al. 2009, Jackson et al. 2012, Ali and Miller 2012, and Nava and Piccione 2012). This paper is different from much of this literature along two important dimensions. First, we study interactions that are competing—that is, there is a capacity constraint on the number of interactions of each individual agent (or client) in any given period. As a result, having more links does not lead to an increase in the maximal number of interactions an individual can engage in within a period, and can only increase the probability of an interaction. In this environment, the value that an agent has for a connection is determined by the network structure, even conditional on full cooperation. This allows us obtain predictions on the network structure without imposing an exogenous cost of creating a link, which is commonly done in the network formation literature.²⁷

An second key departure from the existing literature is that we separate the analysis of the trade network from that of the observability network and allow both to vary in economically meaningful ways. Varying separately the interaction network and the observability network, we show that third-party observability enhance efficiency by allowing for a larger set of interaction networks to be cooperation networks. A deeper understanding of the effect of third-party observability on cooperation can be achieved by comparing our results

²⁷See Goyal (2007) and Jackson (2008) for good surveys of the network formation literature. See also Ali and Miller (2012) for an example in the context of repeated games in networks.

to the results in Fainmesser (2012a) who provides an application of our methodology to an environment in which third-party observability is nonexistent. That is, Fainmesser (2012a) makes use of our random tree characterization and studies the existence of TECOs when the observability network R is the empty network. Fainmesser (2012a) finds that in the absence of third-party observability, sparse interaction networks in which the degrees of clients and agents are similar, are GF cooperation networks for the largest range of parameters. This restricts the number (mutually beneficial) interactions in the absence of third-party observability. In contrast, we show that when the observability network is complete (R^1), dense networks in which the degrees of clients are small and the degrees of agents are large are GF cooperation networks for an even larger range of parameters.

Consequently, in the presence of perfect third-party observability, for *large* n_c and n_a , dense networks that maximize the number of (mutually beneficial) interactions and networks that maximize the incentives to cooperate are approximately identical. However, as illustrated in Example 2, Theorem 3 and Corollary 3 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 Let $\frac{n_c}{n_a} < \mu$,

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

and

$$\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \mu - \gamma > 0. \tag{2}$$

Condition (1) guarantees that no network N in which $Pr(d_a = 0|N) = 0$ admits a TCEO. At the same time, condition (2) assures us that there exists a non-empty network that admits a TCEO.²⁸

As demonstrated in figure 4, 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. However, in (asymptotically) large markets, and as long as condition (2) holds, there exists an (asymptotically) welfare maximizing network that facilitates full cooperation.

Example 2 (cont.) *Suppose that there exist positive integers d_C and d_A , and $\phi_a \in (0, 1)$ s.t. $Pr(\deg(a) = d_A) = \phi_a$, $Pr(\deg(a) = 0) = 1 - \phi_a$, and $Pr(\deg(c) = d_C) = 1$. 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 $\frac{n_c}{n_a} \geq \mu$, condition (2) implies that $\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \frac{n_c}{n_a} - \gamma > 0$. 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, $\lim_{d_C \rightarrow \infty} \lim_{n_c \rightarrow \infty} \left(\frac{V(\mathcal{D}_{GF}(n_c))}{n_c} \mid \phi_a^* \right) = 1$, and*

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

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.

8.2 The (un)importance of global beliefs

Much of the previous work on games in networks analyzes static network games (e.g. Galeotti et al. 2010, Ballester et al. 2006, and Bramoullé, D'Amours, and Kranton 2013). 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. 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.

In contrast, in our model and 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 –

²⁸For 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.

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.

8.3 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).²⁹ 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. In this sense, this paper raises 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 exhibit no degree correlation.³⁰ 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 of connections of each of the other players connected to her that are captured by a fixed degree distribution. Lemma 3 provides sufficient conditions under which this assumption is consistent with a common prior.

²⁹E.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 previously 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.

³⁰For an application, see Fainmesser (2013).

9 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 (2013) 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.³¹ To this end, we now 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 first restate the Lemma in graph theoretic term in Lemma 4, and then prove Lemma 4). 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.

9.1 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 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. $A(\text{gents})$ and $C(\text{lients})$) 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 $\{1, \dots, n_a\}$, and each node in partite C have a distinct label from the set $\{1, \dots, n_c\}$. 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 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 random variable representing the degree of the agent (client) belonging to an edge selected u.a.r. from all edges of Γ .

³¹See Richardson and Urbanke (2008) for an example from coding theory.

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 $\mathcal{R}^A(n_c)(\mathcal{R}^C(n_c))$ 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 ,*

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

9.2 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, \dots, A_{n_a} and n_c client buckets C_1, C_2, \dots, C_{n_c} . Second, for each $d \geq 1$, populate a $Pr(\mathcal{H}^a = d)(Pr(\mathcal{H}^c = d)$ 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 $\mathcal{R}'(n_c)$ 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 , $\lim_{n_c \rightarrow \infty} \sup_{v \in \mathcal{R}'^A(n_c)} |Pr(N_d^{\mathcal{R}'(n_c)}(v) = T) - Pr(\mathcal{T}(d, d_v) = T)| = 0$.*

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 $2(m_{\mathcal{H}} + m_{\mathcal{H}}(m_{\mathcal{H}} - 1) + m_{\mathcal{H}}(m_{\mathcal{H}} - 1)^2 + \dots + m_{\mathcal{H}}(m_{\mathcal{H}} - 1)^{\Delta-1}) \leq 2\Delta m_{\mathcal{H}}^{\Delta}$ 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 that³²

$$Pr(E_{i,\Delta}) \geq 1 - \binom{2\Delta m_{\mathcal{H}}^{\Delta}}{2} \frac{m_{\mathcal{H}} - 1}{n_{c,p} - 2\Delta m_{\mathcal{H}}^{\Delta}} = 1 - O\left(\frac{1}{n_c}\right). \quad (3)$$

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 (3) 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 $Pr(\mathcal{H}^c = l) - O\left(\frac{1}{n_c}\right)$. 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 $Pr(\mathcal{H}^c = l) + O\left(\frac{1}{n_c}\right)$. 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 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. ■

9.3 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 $\lim_{n_c \rightarrow \infty} Pr(\mathcal{R}'(n_c) \in \mathcal{G}(n_c)) > 0$.

9.4 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

³²The 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.

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 (3), and we omit the details. ■

10 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 a network $N = (G, R)$. The extension for stochastic beliefs follows simply by adding the expectation operation when applicable.

First, let $\omega_{\mathbf{m}}$ be such that $\{\mathcal{D}_a\}_{a \in A}$ is common knowledge. Thus, to prove that a TCEO with $\omega_{\mathbf{m}}$ exists, it is sufficient to prove that all agents have an incentive to cooperate in a TCEO.

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* $Q_a^t : N_1(a) \rightarrow \{0, 1\}$, 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 (defect) with c . We now show that conditional on $\omega_{\mathbf{m}}$, Q_a^t depends only on t and $\{\sigma^\tau\}_{\tau=1}^{t-1}$ (this is true independent of whether agents observe $\{\sigma^\tau\}_{t=1}^{\infty}$ or not).

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 $Pr(I_N^t(c, a))$ denote the probability that $I_N^t(c, a) = 1$ in network N . Now assume that all other agents $a' \neq a$ always cooperate. Then, that Q_a^t depends only on t and $\{\sigma^\tau\}_{\tau=1}^{t-1}$ 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 $\{I^\tau(a, c)\}_{c \in N_1(a), \tau=1 \dots t-1}$; [4] $\{I^\tau(c, a)\}_{c \in N_1(a), \tau=1 \dots t-1}$ 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 $t \in [1, t_a^1], (t_a^1, t_a^2], \dots$

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 and all clients use ostracizing strategies. 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 $\bar{u}_N(Q) \triangleq u_N(\mathcal{O}_N(Q))$, $\bar{u}_N \triangleq u_N(\mathcal{O}_N)$, 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 $Pr(I_N^t(c, a))$ 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 ,

$$\begin{aligned} \bar{u}_N(Q) &= \sum_{c \in N_1(a): Q(c)=1} Pr(I_N^1(c, a)) (\pi - \gamma + \delta \cdot \bar{u}_N(Q)) \\ &+ \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a)) (\pi + \delta \cdot \bar{u}_{N \setminus E_R^a(c)}) \\ &+ \left(1 - \sum_{c \in N_1(a)} Pr(I_N^1(c, a)) \right) \cdot \delta \cdot \bar{u}_N(Q) \end{aligned} \quad (4)$$

In particular,

$$u_N^{coop} = (\pi - \gamma) \cdot \sum_{c \in N_1(a)} Pr(I_N^1(c, a)) + \delta \cdot u_N^{coop} \quad (5)$$

It follows that,

$$(\bar{u}_N(Q) - u_N^{coop}) = \frac{(\gamma + \delta \cdot (\bar{u}_{N \setminus E_R^a(c)} - u_N^{coop})) \cdot \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))}{1 - \delta + \delta \cdot \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))} \quad (6)$$

Since $\sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a)) > 0$, it follows from (6) that when for every agent a , \mathcal{D}_a consists of a single and fixed network N , then ω_m admits a TCEO if and only if for every agent a and each client $c \in N_1(a)$,

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

■

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 $Pr(c \in N_1(a) | \mathcal{D}_a) \in \{0, 1\}$. Namely, for any client $c \in \mathcal{C}$, 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 $Pr(c \in N_1(a) | \mathcal{D}_a) \in [0, 1]$ follows the same argument but requires additional notation and is omitted.

Equations (4) and (5) imply that

$$u_N^{coop} = (\pi - \gamma) \cdot \frac{\sum_{c \in N_1(a)} Pr(I_N^1(c, a))}{1 - \delta} \quad (7)$$

and

$$\begin{aligned} \bar{u}_N &= \max_Q [\bar{u}_N(Q)] = \\ &= \max_Q \left(\frac{\pi \sum_{c \in N_1(a)} Pr(I_N^1(c, a)) - \gamma \sum_{c \in N_1(a): Q(c)=1} Pr(I_N^1(c, a))}{1 - \delta + \delta \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))} \right. \\ &\quad \left. + \frac{\delta \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))}{1 - \delta + \delta \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))} \cdot \bar{u}_{N \setminus E_R^a(c)} \right) \end{aligned} \quad (8)$$

Finally, noting that in a network of maximum degree Δ , $Pr(I_N^1(c, a)) \in [\frac{1}{2\Delta-1}, 1]$ for all edges (c, a) , Theorem 1 follows by interpreting Equation (8) as a dynamic program (for computing \bar{u}_N), combined with Lemma 2 and a simple induction, and applying the same logic to (7). ■

Theorem 2 - Proof. We prove the theorem by proving the following result: For all networks N and agents a in N , $\lim_{d \rightarrow \infty} IC(T(a, N, d))$ and $\lim_{n_c \rightarrow \infty} IC(\mathcal{D}_{GF}(a | N, n_c))$ both exist, and equal one-another. This follows from Theorem 1 and Lemma 3. ■

Recall that R^1 is the complete observability network and that $Pr(I_N^1(c, a))$ is the probability that c and a interact in period 1 given the network N . Then,

Corollary 1 Consider a network $N = (G, R^1)$ and an agent a in the network. Then, $\lim_{d \rightarrow \infty} \sum_{c \in N_1(a)} Pr(I_{T(a, N, d)}^1(c, a))$ exists, and

$$sign \left(\lim_{n_c \rightarrow \infty} IC(\mathcal{D}_{GF}(a | N, n_c)) \right) = sign \left(\frac{\delta(\pi - \gamma)}{1 - \delta} \lim_{d \rightarrow \infty} \sum_{c \in N_1(a)} Pr(I_{T(a, N, d)}^1(c, a)) - \gamma \right). \quad (9)$$

Proof. That the necessary limit exists follows from Lemma 2. To derive (9) note that if the observability network is captured by R^1 then $\bar{u}_{N \setminus E_R^a(c)} = 0$ for all c . Therefore, it follows from (6) that for any fixed belief N and strategy Q ,

$$(\bar{u}_N(Q) - u_N^{coop}) = \frac{(\gamma - \delta \cdot u_N^{coop}) \cdot \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))}{1 - \delta + \delta \cdot \sum_{c \in N_1(a): Q(c)=0} Pr(I_N^1(c, a))} \quad (10)$$

The Corollary then follows from (10), (7), and Theorem 2. ■

Proposition 1 - Proof. We prove Proposition 1 by proving the following claim: Consider a network $N = (G, R)$, and let $IC_{\mathcal{D}_{GF}}(R') = IC(\mathcal{D}_{GF}(a|(G, R')))$. Then, for any observability network \widehat{R} , $IC_{\mathcal{D}_{GF}}(\widehat{R}) \leq IC_{\mathcal{D}_{GF}}(R^1)$.

To prove this claim, note that by definition,

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

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, Corollary 1, and interpreting $Pr(I_N^1(c, a))$ as the probability that edge (c, a) is selected by *GREEDY*. ■

Theorem 3 - We prove the theorem by proving the following equivalent claim.

Lemma 10 *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 observability 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,*

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

and

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

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 number of trades possible conditional on n_c and n_a .

Proof. Since all agents are symmetric, we have that

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

$$E[V(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C))] = \mu n_a \sum_{c \in N_1(a)} \Pr(I^t(c, a) = 1 \mid a \text{ is active at } t). \quad (11)$$

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 $I(v \in \mathcal{M}(G))$ 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 (11) 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 $d \geq d_{\epsilon, d_A, d_C, \mu}$,

$$\left| \frac{E[V(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C))]}{n_a} - \mu \Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C, \mu))\right) \right| < \epsilon, \quad (12)$$

and thus

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

It follows from Lemma 9 that $\Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C, \mu))\right) \geq \Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C))\right)$, since deleting agents is equivalent to removing (as children) subtrees from client nodes. Combining with (13), we find that for all $n_a, n_c \geq N_{\epsilon, d_A, d_C, \mu}$ and $d \geq d_{\epsilon, d_A, d_C, \mu}$

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

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

$$\Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C))\right) \geq \Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_A))\right). \quad (15)$$

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

$$\lim_{d \rightarrow \infty} \Pr\left(r_1 \in \mathcal{M}(T^1(d, d_A, d_A))\right) = 1 - (d_A - 1)^{-\frac{d_A}{d_A - 2}}. \quad (16)$$

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

$$\lim_{d \rightarrow \infty} \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_A, d_A)) \right) \geq 1 - \frac{1}{d_A - 1}. \quad (17)$$

Combining (14),(15), and (17) 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 $T^1(d, d_A, d_C)$ ($T^1(d, d_C, d_A)$). 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 $T^1(d, d_A, d_C)$, $|\Pr \left(a \in \mathcal{M}(G(d, \epsilon, d_C, d_A)) \right) - \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C)) \right)| < \epsilon$. Also, for any node c belonging to the (at least) $(1 - \epsilon)|C|$ nodes of partite C with depth- d neighborhoods isomorphic to $T^1(d, d_C, d_A)$, $|\Pr \left(c \in \mathcal{M}(G(d, \epsilon, d_C, d_A)) \right) - \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_C, d_A)) \right)| < \epsilon$. Combining the above, we find that for the graph $G(d, \epsilon, d_C, d_A)$,

$$|E[\sum_{c \in C} I \left(c \in \mathcal{M}(G(d, \epsilon, d_C, d_A)) \right)] - |C| \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_C, d_A)) \right)| \leq 2\epsilon|C|, \quad (18)$$

and

$$|E[\sum_{a \in A} I \left(a \in \mathcal{M}(G(d, \epsilon, d_C, d_A)) \right)] - |A| \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C)) \right)| \leq 2\epsilon|A|. \quad (19)$$

Note that since the number of matched nodes in partite A always equals the number of matched nodes in partite C , one has $E[\sum_{c \in C} I \left(c \in \mathcal{M}(G(d, \epsilon, d_C, d_A)) \right)] = E[\sum_{a \in A} I \left(a \in \mathcal{M}(G(d, \epsilon, d_C, d_A)) \right)]$. It thus follows from (18) and (19) that

$$|\Pr \left(r_1 \in \mathcal{M}(T^1(d, d_C, d_A)) \right) - \frac{d_C}{d_A} \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C)) \right)| \leq 2\epsilon \left(1 + \frac{d_C}{d_A} \right), \quad (20)$$

and

$$\lim_{d \rightarrow \infty} \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_C, d_A)) \right) = \frac{d_C}{d_A} \lim_{d \rightarrow \infty} \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_A, d_C)) \right). \quad (21)$$

Combining with (14), 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 $n_a, n_c \geq N'_{\epsilon, d_A, d_C, \mu}$, $d \geq d'_{\epsilon, d_A, d_C, \mu}$,

$$\frac{E[V(\mathcal{D}_{GF}(R^1, n_c, d_A, d_C))]}{n_a} \geq \mu \frac{d_A}{d_C} \Pr \left(r_1 \in \mathcal{M}(T^1(d, d_C, d_A)) \right) - \epsilon. \quad (22)$$

It follows from Lemma 9 that $\Pr\left(r_1 \in \mathcal{M}(T^1(d, d_C, d_A))\right) \geq \Pr\left(r_1 \in \mathcal{M}(T^1(d, d_C, d_C))\right)$. Combining with (17) (replacing d_A by d_C) and taking limits demonstrates the theorem for the case $d_C \geq d_A$. ■

Proposition 3 - Proof. Part 1: Let $\mu = 1$ and $\kappa \geq 1$. Let $N = (G, R)$ be any network that satisfy the following: [1] for every $c \in C$, $\deg(c) = 1$; and [2] $\max\{\deg(a)\}_{a \in A} = 1$. The network N consists of \bar{n}_c client-agent pairs and $\kappa\bar{n}_c - \bar{n}_c$ agents that are not connected to any client (R can be chosen arbitrarily). Let ΔN put probability 1 on the network N . Then,

$$V(\Delta N) = \bar{n}_c \text{ and } \min_{a \in A} IC(\Delta N) = \frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \mu - \gamma > 0. \quad (23)$$

Plugging (23) into the definition of $WL(\bar{n}_c, \bar{n}_a)$ completes the proof. The proof for the case where $\kappa < 1$ is symmetric.

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

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

$$\lim_{d_A \rightarrow \infty} \lim_{\bar{n}_c \rightarrow \infty} \frac{V(\Delta N)}{\bar{n}_c} \geq 1,$$

and that $\frac{\delta(\pi - \gamma)}{1 - \delta} \cdot \mu - \gamma > 0$ implies that

$$\lim_{d_A \rightarrow \infty} \lim_{\bar{n}_c \rightarrow \infty} IC(\Delta N) > 0,$$

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. ■

11 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.

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

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

Bramoullé, Yann, Martin D'Amours, and Rachel Kranton (2013), "Strategic Interaction and Networks," *American Economic Review*, forthcoming.

Campbell, Arthur (2012), "Word of Mouth and Percolation in Social Networks," *American Economic Review*, forthcoming.

Deb, Joyee and Julio González-Díaz (2012), "Community Enforcement Beyond the Prisoner's Dilemma," 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.

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*, 4(1), 32-69.

Fainmesser, Itay P. (2012b), "Intermediation and Exclusive Representation in Financial Networks," working paper.

Fainmesser, Itay P. (2013), "Social Networks and Unraveling in Labor Markets," *Journal of Economic Theory*, 148 (1), 64-103.

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.

Gaspard, 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.

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 (2012), "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *American Economic Review*, 102(5), 1857-1897.

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

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.

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.

Lippert, Steffen and Giancarlo Spagnolo (2011), "Networks of Relations and Word-of-Mouth Communication," *Games and Economic Behavior*, 72 (1), 202–217.

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.

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 Games with Uncertain 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.