Enforcing Cooperation in Networked Societies

S. Nageeb Ali and David A. Miller[*]
UCSD

January 25, 2013

Abstract

Which social norms and networks maximize cooperation in bilateral relationships? We study a network of players in which each link is a repeated bilateral partnership with two-sided moral hazard. The obstacle to community enforcement is that each player observes the behavior of her partners in their partnerships with her, but not how they behave in other partnerships. We introduce a new metric for the rate at which information diffuses in a network, which we call *viscosity*, and show that it provides an operational measure for how conducive a network is to cooperation. We prove that clique networks have the lowest viscosity and that the optimal equilibrium of the clique generates more cooperation and higher average utility than any other equilibrium of any other network. This result offers a strategic foundation for the perspective that tightly knit groups foster the most cooperation. We apply this framework to analyze favor exchange arrangements, decentralized trade in networked markets, and social collateral.

---

[*] Contact us at snali@ucsd.edu and d9miller@ucsd.edu. We thank Joyee Deb, Matt Elliott, Itay Fainmesser, Ben Golub, Joe Harrington, Matt Jackson, Navin Kartik, Fahad Khalil, Markus Möbius, Kalle Moene, Paul Niehaus, Marcin Peski, Andres Santos, Joel Sobel, Max Stinchcombe, Adam Szeidl, Gaute Torsvik, Joel Watson, Alex Wolitzky, Muhamet Yildiz, Bill Zame, and numerous seminar and conference participants for helpful comments. We thank Aislinn Bohren for excellent research assistance. This research is financially supported by NSF grant SES–1127643. In addition, Ali gratefully acknowledges financial support from NSF grant NetSe–0905645, as well as financial support and hospitality from Harvard and Microsoft Research; Miller gratefully acknowledges financial support and hospitality from Microsoft Research, Yale, and the Cowles Foundation.

Contents

1 Introduction

Enforcement of cooperation through social and communal mechanisms is pervasive. A large literature on social norms emphasizes the role of networks in mechanisms of informal enforcement, with the understanding that the network structure influences how information diffuses and the incentives to cooperate. In particular, it is often argued that groups that are most socially connected or tightly knit induce the strongest incentives to cooperate.[1] In this paper, we develop a strategic foundation for analyzing the role of networks in supporting cooperation in bilateral relationships, and identify networks that promote the greatest cooperation and efficiency.

We model a networked society in which each link is an ongoing bilateral partnership with two-sided moral hazard. Each partnership meets at exponentially distributed arrival times to play a prisoners' dilemma at stakes that it chooses. Higher stakes generate greater payoffs for the partnership, but increase the temptation to shirk—and thus require stronger incentives. Two partners can ensure cooperation at modest stakes using simple bilateral grim trigger strategies between them, but involving other members of the community in their enforcement arrangement enables them to cooperate at higher stakes, using rewards and punishments in other relationships to enforce cooperation in their own. This is the essence of community enforcement.

A challenge central to community enforcement is *private monitoring*: each player knows what transpires within her own partnerships, but does not observe when or how her partners interact with others. As a consequence, she lacks direct access to information that would enable her to reward or punish her partners for what they have done in other relationships. The social network can partly mitigate this difficulty, since each link is not only a productive relationship, but also a conduit for information. We characterize networks that optimally address this issue.

Our main result is that the optimal network divides society into entirely disconnected *cliques*. The clique of degree d is a network of $d+1$ players that is completely connected; i.e., each of these players is linked to all the others. We prove in Theorem 3 that there exists an equilibrium on this clique that attains the highest utilitarian average payoffs among all equilibria on all networks in which each player has at most d partnerships. An implication of our result is that if no player can have degree exceeding d, then it is optimal to organize society into cliques of degree d (ignoring

---

[1] Within sociology, see Festinger, Schachter, and Back (1948), Granovetter (1985), Coleman (1988), Raub and Weesie (1990), and Granovetter (2005) for a survey. Within economics, the impact of social connectedness and networks on economic behavior is highlighted by Glaeser, Laibson, and Sacerdote (2002), and features prominently in discussions of risk-sharing (Besley, Coate, and Loury 1993; Udry 1994; Fafchamps and Lund 2003; Bloch, Genicot, and Ray 2008; Ambrus, Möbius, and Szeidl 2010), trading without enforcement (McMillan and Woodruff 1999; Dixit 2003), and political economy (Grossman and Helpman 2002).

remainder issues). When accounting for linking costs explicitly, we find that the complete network is optimal if linking costs are linear or concave, and that an optimally sized clique is Pareto efficient if linking costs are strictly convex. These network comparisons apply for every level of patience and frequency of interaction.

To appreciate why cliques are optimal, consider the incentives that Ann faces in deciding whether to work in her partnership with Bob. The key question that she must ask herself is, were she to shirk, would she be able to take advantage of her other partners before they punish her? At the core of her incentives is the speed at which information propagates through the network from her partnership with Bob to her other partnerships. Her calculation must account for the number, length, and arrangement of paths within the network from Bob to her other partners. To account for these incentives, we derive a new measure for networks, *network viscosity*, which is a sufficient statistic for Ann's incentive to work. The viscosity factor measures the discounted probability that Ann will be able to take advantage of another one of her partners (say, Carol) before Carol can punish Ann.

In Theorem 2, we prove that among all networks that have a maximal degree of d, the clique of degree d has the lowest viscosity. We prove this result by establishing that on the clique, the indirect paths between Bob and each of Ann's other partners are shorter than they could be on any other network with maximal degree d. Since computing viscosity for every possible network in which the maximal degree is d is infeasible, we develop a fundamental coupling argument that maps each path that propagates punishment on any of these networks with one on the clique that propagates it faster.

Merely because a clique maximizes the speed by which punishments may spread does not imply that its equilibria necessarily generate greater cooperation than those of other networks. The challenge is that incentives both on and off the equilibrium path vary greatly across networks and equilibria, so it is not feasible to compute the frontier of equilibrium payoffs and verify the incentive compatibility of punishments for every possible configuration of parameters. Accordingly, ranking networks by direct computation of equilibria is analytically intractable.

We introduce two conceptual insights that sidestep this technical hurdle. The first is that for a symmetric network—such as a clique—we can characterize the best stationary equilibrium payoffs for every discount factor. These payoffs, as we prove in Theorem 1, are those of a *binding contagion equilibrium*, in which each player is just willing to work on the equilibrium path but shirks once she has observed anyone else shirk. Of all the equilibria on a symmetric network, contagion makes the most effective use of viscosity in that network.

The second insight is that even if it is unclear what might constitute an optimal equilibrium on an arbitrary network G, all equilibria are bound by a common constraint: shirking forever cannot be a strictly profitable deviation. Perhaps surprisingly, this is all that is needed, in combination with our prior results, to find the optimal network. From it we derive a series of implications that apply our viscosity result to identify a clique whose binding contagion equilibrium outperforms every equilibrium on G.

Viscosity not only is an important technical element in our proof, but also provides a new operational way to measure how conducive a network is to cooperation. A rich literature on networks has argued that tightly knit networks offer the greatest cooperation (Granovetter 1985; Coleman 1988) but often has often focused on "local" measures such as clustering and support coefficients.[2] But networks may be identical by these measures and yet support different levels of cooperation. For instance, we show in an example in Section 3 that although all circles of more than three players have identical clustering and support coefficients of zero, nevertheless viscosity is strictly increasing in the circumference of the circle, and consequently the attainable level of cooperation is strictly decreasing. In addition to producing a closed form expressions for the viscosity of cliques, we offer a finite algorithm to compute the viscosity of any arbitrary network.

Viscosity is a pivotal consideration for networked relations that feature moral hazard in partnerships, and we demonstrate its relevance by considering three other such settings. First, we examine a "favor exchange" environment in which each player randomly encounters opportunities to do favors for his neighbors. Although no two players are ever called upon to do favors for each other simultaneously, the essence of our results remains: if each player can form up to d links then the clique of degree d maximizes average utility. Second, we consider networked markets that are segregated into buyers on one side and sellers on the other; no two players on the same side of the market can be linked. In such a *two-sided* network, viscosity can be computed separately for each side of the market. We show that if the maximal buyer degree is d_B and the maximal seller degree is d_S, then the "bipartite clique" (in which there are d_S buyers and d_B sellers, each linked to all the players on the opposite side of the market) supports the greatest gains from trade. Third, we highlight how low viscosity can endogenously generate *social collateral* that not only induces greater cooperation within the network, but improves relationships with outsiders to the network. We illustrate this effect by analyzing whether a firm would choose to hire from a spot market or through the referral of an existing worker when the challenge is that of moral

---

[2] Clustering coefficients measure whether two neighbors of a player are linked to each other; see Jackson (2008, p. 35). The support coefficient, defined by Jackson, Rodriguez-Barraquer, and Tan (2012), measures whether two linked players share a common neighbor.

hazard. We show that the firm would prefer to hire from the network because this permits it to reduce its offered wages, and that its incentives to do so are greater when the network has low viscosity.

Throughout our analysis, we study contagion equilibria on cliques because they attain the maximal level of cooperation, providing a bound on equilibrium payoffs. Other types of social norms could also potentially attain the same bound, and so network comparisons do not require any restriction to contagion equilibria. For example, suppose that players are able to communicate when they meet, and assume that "innocent" players and "third parties" communicate truthfully. A compelling alternative social norm is one in which "innocent" players continue cooperating with each other while *ostracizing* only those who are "guilty." We show that this alternative social norm generates the same network comparisons as in our main results.[3]

Related Literature Our work builds on ideas in both repeated games and networks. Here, we summarize research that is most closely related to our general framework.[4]

The issue of private monitoring that we tackle was first addressed in the literature on contagion, which seeks to understand how players interacting anonymously in a random-matching environment can be induced to cooperate in prisoners' dilemmas. Kandori (1992) and Ellison (1994) were the first to develop contagion in this setting, and Harrington (1995) develops a similar approach for a four player game in which behavior is non-anonymous. The key idea that we adopt from contagion is that relationships can be strategically interlinked to improve cooperation even in the absence of communication. In contrast to the prior literature, which makes restrictions or enriches the setting to guarantee existence, our variable-stakes setting cleanly affords existence and optimality of contagion equilibria on symmetric networks.

Our main interest is to compare networks from the standpoint of cooperation, and to understand the benefits of indirect paths. Within the recent theoretical literature on network-based cooperation, Karlan, Möbius, Rosenblat, and Szeidl (2009) articulate a static model of how indirect paths within a network can be leveraged towards social collateral between players who are not directly linked. Their work offers notions of trust within networks that, like our measure of viscosity, cannot be reduced to merely local information about a player's neighborhood.

Jackson, Rodriguez-Barraquer, and Tan (2012) also study contagion on networks, with a particular focus on how contagion can be locally contained rather than spreading globally. They

---

[3] We tackle the question of whether an ostracism social norm can be incentive compatible in our companion paper, (Ali and Miller 2012).

[4] We defer discussion of work related to our applications to Section 6.

study a complete information environment in which a player's neighbors all publicly and instantaneously observe his actions along each link. Since information propagation has no role in their framework, every network (including a tree) in which each player has some minimal degree generates the same level of cooperation. They find that, of all such networks, only "social quilts" (basically trees of cliques) are renegotiation-proof. Our framework focuses instead on the friction that each player observes only her interactions, and not those of others, and the role of the network is to foster endogenous information propagation. Thus, the force that makes cliques optimal in our framework is different from that which their work elucidates.

A different approach to community enforcement focuses on pooling incentive constraints with synchronous interaction, an approach reminiscent of multi-market collusion (Bernheim and Whinston 1990); this logic is applied to networks by Maggi (1999) and Lippert and Spagnolo (2011). Because pooling incentive constraints requires slack incentives to subsidize those that are not slack, this literature finds that community enforcement can improve upon bilateral enforcement in some, but not all, relationships.

Another strand of literature on repeated games on networks has focused on local interaction environments in each player takes a single action that affects all of his neighbors, rather than interacting with each of them bilaterally. Our approach (and those described above) apply to the community enforcement of *bilateral relations*, as in the trade between buyers and sellers, or the exchange of favors in a community; by contrast, the analysis of local interaction environments is applied most readily to incentives in local public good provision. Haag and Lagunoff (2006) study optimal network design and find that cliques are optimal in order to separate impatient players from those who are more patient. Although their setting also features private monitoring, a player's action is observed by all of her neighbors and so the force that pushes towards cliques in their setting is orthogonal to ours. Recently, Wolitzky (2012) constructs contagion equilibria to support public good provision with private monitoring, and Nava and Piccione (2012) construct "temporary" contagion equilibria for local interaction games in which players are uncertain about the network structure. These papers all focus on synchronous local interaction rather than asynchronous bilateral interactions.

2 Model

Network A society is a finite set of players, $N \equiv \{1, \ldots, n\}$, connected by an undirected *network* G, which is a set of cardinality-2 subsets of N. The network is commonly known by the players, and is fixed throughout the game. We use $\{ij\}$ to indicate a link, and define $|G|$ to be the number of links in G. Much of our analysis concerns the incentives of each player on the link: accordingly, we use ij to signify "player i on link $\{ij\}$" as distinct from ji ("player j on link $\{ij\}$").

Taking network G as given, player i's *neighborhood* N_i is the set of players to whom player i is linked: $N_i \equiv \{j \in N : \{ij\} \in G\}$. The cardinality of N_i is player i's *degree*, denoted by d_i. A *path* from player i to j is a sequence of nodes $i_1, \ldots, i_Z$ such that $\{i_z i_{z+1}\} \in G$ for each $z \in \{1, ..., Z-1\}$, $i_1 = i$, $i_Z = j$, and each node in the sequence is distinct. A *cycle* is a sequence of nodes $i_1, \ldots, i_{Z-1}, i_1$ such that $i_1, \ldots, i_{Z-1}$ is a path, and $\{i_1 i_{Z-1}\} \in G$. A *component* G' of a network G is a maximal connected subnetwork; i.e., if $\{ij\} \in G'$ then $\{k\ell\} \in G'$ if and only if there exists a path in G that contains both $\{ij\}$ and $\{k\ell\}$.

Social interactions Time is continuous, and players discount payoffs realized at time t in $\mathbb{R}_+$ by the common discount rate $r > 0$. Each link in the network is governed by an independent Poisson recognition process with the common rate $\lambda > 0$. Whenever link $\{ij\}$ is recognized, players i and j engage in a two-stage interaction that occurs in that instant:

1. *Stake selection stage*: Players i and j simultaneously propose the *stakes* at which they should interact. Player i's proposal is $\hat{\phi}_{ij} \in \mathbb{R}_+$; the stakes of the $\{ij\}$ relationship are set to the minimum of the two proposals. We use $\phi \in \mathbb{R}_+$ as a generic stakes parameter when the identities of the players along the link are unimportant.
2. *Action stage*: Each player simultaneously chooses an action from $A \equiv \{\text{work}, \text{shirk}\}$. Their stakes determine the payoffs; higher stakes increase the payoffs from mutual effort but strengthen the temptation to shirk. Specifically, given stakes $\phi = \min\{\hat{\phi}_{ij}, \hat{\phi}_{ji}\}$ they face the prisoners' dilemma in Figure 1.

		Player j	
		Work	Shirk
Player	Work	ϕ, ϕ	$-V(\phi), T(\phi)$
i	Shirk	$T(\phi), -V(\phi)$	$0, 0$

Figure 1. The prisoners' dilemma of stakes ϕ

The "temptation reward" T and the "victim's penalty" V are smooth functions satisfying $T(0) = V(0) = 0$, as well as $V(\phi) > 0$ and $T(\phi) > \phi$ for all $\phi > 0$. Thus if the stakes are positive, shirking is the strictly dominant action in the stage game for each player. Throughout the paper we assume that the temptation reward is increasing in the following manner:

Assumption 1 (Increasing Temptation)**.** *T is strictly increasing and strictly convex, with $T'(0) = 1$ and $\lim_{\phi\to\infty} T'(\phi) = \infty$.*

The important implication from Assumption 1 is that $T(\phi)/\phi$—the ratio of the payoffs from shirking vs. working while one's partner works—is close to 1 at low stakes but increases without bound as the stakes increase. As a consequence, the players require proportionally stronger incentives to work at higher stakes. We use the specification $T(\phi) = \phi + \phi^2$ in examples. Allowing players to set the stakes of their relationships is a key feature of our framework, on which we comment in Section 7.

For some of our results, we restrict attention to prisoners' dilemmas in which the incremental gain from working in the stage game is higher when one's partner works: this is a condition of supermodularity on the stage game, also referred to as strategic complementarity.

Definition 1. *The stage game satisfies **strategic complementarity** if $V(\phi) > T(\phi) - \phi$ for all $\phi > 0$.*

Monitoring and equilibrium Monitoring is pairwise: as play unfolds, each player observes only what transpires along his own links, and observes neither the meeting times nor the behavior along any other links. Each meeting between two players is an *interaction*, characterized by the link $\{ij\}$ that was recognized, the time t at which it was recognized, the stakes $(\hat{\phi}_{ij}, \hat{\phi}_{ji})$ that players i and j proposed in the stake selection stage, and the actions (a_{ij}, a_{ji}) that they chose in the action stage. For a player i, when one of his links is recognized, his *private history* h_i^t is an ordered list of all his interactions up to (but not including) time t, along with the identity of the partner he is interacting with at time t.[5] We denote the set of player i's private histories ending with a recognition of link $\{ij\}$ at time t as H_{ij}^t, and write $H_{ij} \equiv \bigcup_{t\in[0,\infty)} H_{ij}^t$, and $H_i \equiv \bigcup_{j\in N_i} H_{ij}$.

A (behavior) strategy for player i is a function $\sigma_i = (\sigma_i^{\mathrm{S}}, \sigma_i^{\mathrm{A}})$ such that $\sigma_i^{\mathrm{S}} : H_i \to \Delta[0,\infty)$ is his stake-selection strategy and $\sigma_i^{\mathrm{A}} : H_i \times \mathbb{R}_+^2 \to \Delta A$ is his action strategy. We study *weak perfect Bayesian equilibria*: each player's strategy is sequentially rational given a belief system constructed by Bayesian updating. We restrict attention to equilibria that are *stationary* on the equilibrium path: along each link, the partners' choices lead to the same distribution of stakes and actions at each of their interactions along the equilibrium path.

---

[5] Since player i acts only when meeting a partner, there is no need to define his private history at any other times. Also, since at most one link is ever recognized at a time (almost surely), we do not define histories for simultaneous meetings.

Definition 2. *A strategy profile* σ *is* ***stationary*** *if for every* $\{ij\} \in G$

1. *there exists* $\overline{\sigma}^{\mathsf{S}}_{ij} \in \Delta[0,\infty)$ *such that* $\sigma^{\mathsf{S}}_i(h) = \overline{\sigma}^{\mathsf{S}}_{ij}$ *for every equilibrium path history* h *in* H_{ij}*;*
2. *there exists* $\overline{\sigma}^{\mathsf{A}}_{ij} : \mathbb{R}^2_+ \to \Delta A$ *such that* $\sigma^{\mathsf{A}}_i(h, \hat{\phi}_{ij}, \hat{\phi}_{ji}) = \overline{\sigma}^{\mathsf{A}}_{ij}(\hat{\phi}_{ij}, \hat{\phi}_{ji})$ *for every equilibrium path history* h *in* H_{ij} *and proposals* $(\hat{\phi}_{ij}, \hat{\phi}_{ji})$ *in* $\mathrm{Supp}(\overline{\sigma}^{\mathsf{S}}_{ij}) \times \mathrm{Supp}(\overline{\sigma}^{\mathsf{S}}_{ji})$*.*

Throughout our analysis, we restrict attention to stationary equilibria; henceforth, we refer to them as equilibria. This is a restriction on equilibria (the importance of which we describe in Section 7 after presenting our results) but not on the set of feasible deviations for a player. For some of our efficiency results, it is useful to distinguish a particular class of equilibria that is often focal in applications: those in which players work on the equilibrium path.

Definition 3. *A stationary strategy profile* σ *is a* ***mutual effort profile*** *if* $\overline{\sigma}^{\mathsf{A}}_{ij}(\hat{\phi}_{ij}, \hat{\phi}_{ji})$ *assigns probability* 1 *to* work *for all* $(\hat{\phi}_{ij}, \hat{\phi}_{ji})$ *in* $\mathrm{Supp}(\overline{\sigma}^{\mathsf{S}}_{ij}) \times \mathrm{Supp}(\overline{\sigma}^{\mathsf{S}}_{ji})$*.*

In comparing strategy profiles, we say that a strategy profile σ *Pareto dominates* another strategy profile $\tilde{\sigma}$ if no player is worse off with σ and at least one player is strictly better off. The *value* of a strategy profile is the utilitarian average of players' expected payoffs that it delivers on the path of play. Two strategy profiles are *outcome equivalent* if their on-path behaviors coincide; they are *distinct* otherwise. They are *payoff equivalent* if they yield identical expected payoffs to each player.

All proofs for the main results ("theorems") are collected in Appendix A; proofs for other results ("propositions") are in Appendix B.

3 An Example

We begin with a simple example that highlights the essence of our approach. Consider a society in which each of Ann, Bob, and Carol is connected to the other two players. For this illustration, we suppose that $T(\phi) = \phi + \phi^2$, and restrict attention to equilibria in which all pairs coordinate on the same stakes at every equilibrium path history.

Bilateral enforcement: Consider the benchmark of *bilateral strategies*, in which behavior is strategically independent across links. Effectively, each pair plays an infinitely repeated prisoners' dilemma in isolation. Nash reversion is an optimal punishment in this class of strategies: consider strategies in which both players in a partnership work if and only if neither of them has ever

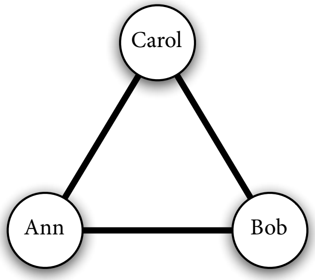

FIGURE 2. A triangle network

deviated. Under these strategies, when Ann meets Bob, her incentive constraint to work along the equilibrium path is

$$\phi + \phi^2 \leq \phi + \int_0^\infty e^{-rt} \lambda \phi \, dt.$$

The highest stakes at which working is incentive compatible is $\frac{\lambda}{r}$.

Instantaneous public monitoring: In contrast to bilateral enforcement, suppose that everyone in society observes all the meetings, stakes announcements, and actions along every link in real time. Then if Ann shirks on Bob, it immediately becomes common knowledge among Ann, Bob, and Carol that continuation play is off the equilibrium path. In this alternative environment, consider an equilibrium in which once anyone shirks, everyone subsequently shirks perpetually. Ann's incentive constraint when she meets Bob along the equilibrium path is:

$$\phi + \phi^2 \leq \phi + 2 \int_0^\infty e^{-rt} \lambda \phi \, dt.$$

The highest stakes at which working is incentive compatible are $2\frac{\lambda}{r}$, doubling what is attainable under bilateral enforcement. Ann is willing to cooperate with Bob at higher stakes because of the immediate punishment that she receives from Carol if she shirks on Bob. This benchmark is infeasible in our environment because each player observes only the activity along his or her own links, so Carol cannot instantaneously learn that Ann should be punished.

Contagion strategies: In *contagion strategies*, a player works so long as all of his partners have worked in the past; otherwise he shirks. Once Ann shirks on Bob, Bob will shirk on Carol at their next interaction, and from then on Carol will shirk on both Ann and Bob. So Ann's only chance for further gain is to meet Carol before Carol becomes "infected." According to her strategy, Ann should then shirk in her next interaction with Carol, so her cooperation phase incentive constraint is

$$\phi + \phi^2 + \int_0^\infty e^{-rt} e^{-\lambda t} \lambda e^{-\lambda t} (\phi + \phi^2)\, dt \leq \phi + 2 \int_0^\infty e^{-rt} \lambda \phi\, dt.$$

Here, $e^{-\lambda t}\lambda$ is the density of Ann's first meeting with Carol, and $e^{-\lambda t}$ is the probability that at that first meeting Carol will not yet have met Bob. The highest stakes at which working is incentive compatible are $\left(\frac{r+4\lambda}{r+3\lambda}\right)\frac{\lambda}{r}$, strictly greater than what is attainable under bilateral enforcement.

We must also verify that at these stakes the contagion phase incentive constraints are satisfied. In particular, consider the scenario in which Bob has very recently observed Ann shirk, and now Bob meets Carol. Bob is unsure whether Carol is "contagious" or not.[6] If Carol is still in the cooperation phase, then shirking yields Bob an immediate payoff of $\phi + \phi^2$ rather than ϕ, but at the cost of future payoffs. Proposition 1 shows that if the stakes are set to make the equilibrium path incentive constraints bind, then Bob strictly prefers to shirk in all his future interactions once one of his neighbors shirks, regardless of his beliefs. The same, of course, applies to Ann when she meets Carol after having shirked on Bob.

We pause to describe how contagion exploits the strategic interdependence of networked relationships. When Ann cheats Bob, she forfeits the opportunity to also cheat Carol if Bob should meet Carol first. This uncertainty dampens her motive to shirk and thereby enables her to cooperate at higher stakes. This effect relies on both asynchronous interaction and cycles in the network. Were behavior instead synchronous, Ann's best deviation from the equilibrium path would be to shirk on both Bob and Carol simultaneously, negating the usefulness of contagion. Similarly, if the network were a tree in which Ann was linked to both Bob and Carol (so Bob and Carol were not linked to each other), then no equilibrium could support mutual effort at stakes greater than under bilateral enforcement—shirking on Bob would not cause Carol to suffer any consequences in her relationship with Ann. A network must have cycles to transmit punishments.

Triangle vs. circles: Shorter cycles transmit punishment more quickly. To illustrate, consider a society comprising $n \geq 3$ individuals connected in a circle. Under contagion strategies, after Ann shirks on Bob, Bob will shirk on Carol, Carol will shirk on Dante, and the contagion will spread around the circle. Ann's only chance to capture another temptation reward is to meet her other

---

[6] If Ann was the first to shirk, she may have shirked on Carol as well as Bob. Another possibility is that Carol was the first to shirk, and Ann was spreading the contagion to Bob.

neighbor before he becomes contagious. We show that Ann's incentive constraint is

$$\phi + \phi^2 + (\phi + \phi^2) \sum_{z=1}^{n-2} \left(\frac{\lambda}{r + 2\lambda} \right)^z \leq \phi + 2 \int_0^\infty e^{-rt} \lambda \phi dt.$$

Evidently, Ann's other neighbor is less likely to be contagious for greater n, and so the maximal stakes on each link are attained for $n = 3$. At the same time, each player has the same degree in the circle as in the triangle. Therefore, if n is divisible by 3 then rearranging the network into $\frac{n}{3}$ triangles induces shorter indirect paths, and hence faster punishments, larger stakes, and higher payoffs. This force pushes optimal network design towards cliques.

4 Contagion

This section defines contagion equilibria, analyzes their incentives both on and off the path of play, and proves their existence and efficiency in symmetric networks.

While the focus of our paper is not on contagion, we describe it in detail for three reasons. First, its structure elucidates the role of viscosity. Second, it offers a tractable approach to finding an optimal equilibrium on symmetric networks, regardless of the discount factor. Third, the steps that we use to establish the optimality of contagion on cliques play an important role later in our main result (Theorem 3), which compares cliques to other networks. In Section 5.4, we establish that identical network comparisons emerge from social norms in which individuals communicate and ostracize guilty players.

4.1 Definition and Incentives

A contagion strategy profile can be represented by an automaton with two phases: cooperation and contagion. A player starts off in the *cooperation phase*, in which she proposes on-path stakes and works with all her partners. On the path of play, a cooperative player believes that all her partners are also cooperative. Should one of her partners shirk or propose non-equilibrium stakes, she switches permanently to the *contagion phase*. A contagious player propose the same stakes as on the path of play, but shirks on all her partners. A "contagion profile" is a strategy profile in which all players play contagion strategies; it is a stationary mutual effort profile in pure strategies. In a contagion profile, any deviation from the path of play initiates a contagion that eventually causes all cooperation within a network component to break down.

Formally, we denote by $\Phi_i = (\phi_{ij})_{j \in N_i}$ a profile of stakes in all of player i's relationships; we call this an *individual stakes profile.* We denote a *collective stakes profile* by $\Phi = (\Phi_i)_{i \in N}$. Each Φ is an element of the subset of $\mathbb{R}_+^{2|G|}$ in which $\phi_{ij} = \phi_{ji}$ for all $\{ij\} \in G$; we denote this space by $\mathcal{S}$.

Definition 4. *A strategy profile σ is a **contagion profile** if there exists a collective stakes profile Φ, such that for each link $\{ij\}$ in G and each history h in H_{ij}, player i plays according to which of the two phases she is in:*

1. *Cooperation phase: $\sigma_i^{\mathrm{S}}(h) = \phi_{ij}$, and $\sigma_i^{\mathrm{A}}(h, \hat{\phi}_{ij}, \hat{\phi}_{ji}) = \text{work}$ if and only if $\hat{\phi}_{ij} = \hat{\phi}_{ji} = \phi_{ij}$.*
2. *Contagion phase: $\sigma_i^{\mathrm{S}}(h) = \phi_{ij}$, and $\sigma_i^{\mathrm{A}}(h, \hat{\phi}_{ij}, \hat{\phi}_{ji}) = \text{shirk}$ for all $\hat{\phi}_{ij}, \hat{\phi}_{ji}$.*

Each player begins in the cooperation phase. If player i is in the cooperation phase at history h, then she stays in the cooperation phase if and only if both players i and j announced stakes ϕ_{ij} and worked; otherwise, she transitions to the contagion phase. The contagion phase is absorbing.

4.1.1 Incentives in the cooperation phase

We begin our analysis of contagion profiles by describing incentives in the cooperation phase. As illustrated by the difference between triangles and circles (Section 3), the rate at which contagion spreads determines the level of cooperation. When player i considers whether to shirk on link ij along the path of play, she needs to predict whether she will meet her other neighbors before they become infected; if contagion spreads quickly, she is not likely to do so.

Consider a contagion profile with collective stakes profile Φ. When link $\{ij\}$ is recognized at time t, if player i is in the cooperation phase then she must have an incentive to work. If she shirks, let $x_{ijk}(\tau)$ be the probability that, from the perspective of time t, she assigns to player k being cooperative at time $t + \tau$. Now we can write her incentive constraint as:

$$T(\phi_{ij}) + \sum_{k \in N_i \setminus \{j\}} T(\phi_{ik}) \int_0^\infty e^{-r\tau} \lambda x_{ijk}(\tau)\, d\tau \leq \phi_{ij} + \sum_{k \in N_i} \int_0^\infty e^{-r\tau} \lambda \phi_{ik}\, d\tau.$$

The right hand side of this inequality sums the payoff from mutual effort at time $t + \tau$ with the discounted expected payoff from future cooperation. The first term on the left hand side is what player i would earn immediately by shirking on player j, and the second term is what she would expect to earn in the future by shirking on her other neighbors. To simplify the incentive constraint, we combine the effects of x_{ijk}, λ, and r on the left hand side into a single term, the *ijk viscosity factor* $X_{ijk} \equiv \int_0^\infty e^{-r\tau} \lambda x_{ijk}(\tau)\, d\tau$. In Section 5.1, we show how to compute viscosity factors for arbitrary networks, but such computations are not yet needed. Viscosity factors simplify the cooperation phase incentive constraint to

$$T(\phi_{ij}) + \sum_{k \in N_i \setminus \{j\}} T(\phi_{ik}) X_{ijk} \leq \phi_{ij} + \frac{\lambda}{r} \sum_{k \in N_i} \phi_{ik} \,. \qquad (\mathrm{IC}_{ij}^{\mathrm{Coop}})$$

Let $\Psi_{ij}(G) \equiv \left\{ \Phi_i \in \mathbb{R}_+^{|N_i|} : \mathrm{IC}_{ij}^{\mathrm{Coop}} \text{ is satisfied} \right\}$ be the set of individual stakes profiles that satisfy the above inequality, i.e., player *i*'s cooperation phase incentive constraint on link $\{ij\}$. Collective stakes profiles associated with contagion equilibria must satisfy each player's cooperation incentive constraint on each link, and therefore any such collective stakes profile is an element of

$$\Psi(G) \equiv \left\{ \Phi \in \mathcal{S} : \Phi_i \in \Psi_{ij}(G) \text{ for every } \{ij\} \in G \right\}.$$

Denoting the interior of a Euclidean set $\mathcal{X}$ by $\operatorname{int} \mathcal{X}$, the following is true regarding these sets:

Lemma 1 (Convexity)**.** *For any network G, $\Psi(G)$ has a non-empty interior and is convex. Moreover, $\Psi_{ij}(G)$ is **strictly convex** for every link $\{ij\}$ in G; i.e., for every pair of individual stakes profiles Φ_i and Φ_i' in $\Psi_{ij}(G)$, if there exists a neighbor k in N_i such that $\phi_{ik} \neq \phi_{ik}'$, then $\alpha \Phi_i + (1-\alpha)\Phi_i' \in$* $\operatorname{int} \Psi_{ij}(G)$ *for any $\alpha \in (0,1)$.*

The notable property is strict convexity, which is a consequence of Assumption 1: because T is strictly convex, for any $\phi \neq \phi'$ and $\alpha \in (0,1)$, it follows from Jensen's inequality that $\alpha T(\phi) + (1-\alpha)T(\phi') > T(\alpha\phi + (1-\alpha)\phi')$. That is, player *i*'s incentives in the cooperation phase are slack along *all* her links if her individual stakes profile is a strictly convex combination of any two equilibrium individual stakes profiles that give her different stakes along *any one* link.

4.1.2 Incentives in the contagion phase

We turn to incentives in the contagion phase. Let $\pi_i(M)$ be player *i*'s continuation value when she believes that $M \subseteq N$ is the set of contagious players. In a contagion profile, if M is non-empty then it must include player i. A sufficient condition for player i to prefer shirking is

$$\phi_{ij} + \pi_i(M) \leq T(\phi_{ij}) + \pi_i(M \cup \{j\}). \qquad (\mathrm{IC}_{ij}^{\mathrm{Cont}})$$

If $M \subset N_i \cup \{i\}$, then this inequality embodies the incentives that player i faces when she knows (from her past history) that players in M are contagious, and believes that the remaining players are cooperative. If $\mathrm{IC}_{ij}^{\mathrm{Cont}}$ is satisfied, then player i prefers to shirk even if player j is not contagious.[7]

Working with contagion equilibria poses a well-known challenge (Kandori 1992; Ellison 1994): if the reward for working far exceeds the punishment for shirking, then a contagious player may prefer to delay infecting others and instead choose to work. We identify a class of contagion equilibria that solves this problem:

Definition 5. *A contagion profile is* ***binding*** *if* IC_{ij}^{Coop} *holds with equality for each neighbor* j *in* N_i*, for each player* i *in* N*.*

The virtue of making incentives bind on the path of play is that contagious players have an incentive to shirk. The logic is analogous to that of Lemma 1 of Ellison (1994): the marginal gain from working is decreasing in the number of contagious players and so indifference on the equilibrium path implies that a player strictly prefers to shirk off the equilibrium path.

Proposition 1. *Every binding contagion profile is an equilibrium.*

Since a binding contagion profile guarantees that each player has a strict incentive to shirk off the equilibrium path regardless of her beliefs about what phases others are in, every binding contagion profile also satisfies additional constraints that would be imposed by refinements such as perfect-extended Bayesian equilibrium (Fudenberg and Tirole 1991; Battigalli 1996).

4.2 Existence and Efficiency on Symmetric Networks

In a *symmetric network*, all links are isomorphic to each other, so each player views each link identically with respect to the network. While symmetric networks are special, their simplicity offers a useful starting point to investigate the existence and efficiency of contagion.

A permutation of the players $g : \mathcal{N} \to \mathcal{N}$ is a *graph automorphism* if link $g(\{ij\})$ is in G for each link $\{ij\}$ in G.[8] Graph automorphisms define symmetry:

Definition 6. *A network* G *is* ***symmetric*** *if for any two links* $\{ij\}$ *and* $\{k\ell\}$ *in* G *there exists a graph automorphism* $g : N \to N$ *such that* $g(i) = k$ *and* $g(j) = \ell$*.*

---

[7] Generally, contagious players need not hold such optimistic beliefs about their partners, but beliefs that attribute greater probability to others being contagious create a stronger incentive to shirk. Establishing incentives to shirk under the most optimistic beliefs about others ensures that contagious behavior is incentive compatible for all beliefs.

[8] When we apply a permutation g to an object containing player names, we mean to apply g to each player name within it. For instance, $g(ij) = g(i)g(j)$.

A symmetric network is necessarily *regular*; i.e., all non-isolated players have the same degree.[9] The convenience of symmetry is that it equates viscosity factors throughout the network.

Lemma 2. *For any symmetric network G, there exist* $x : \mathbb{R}_+ \to [0,1]$ *and* $X > 0$ *such that* $x_{ijk}(t) = x(t)$ *and* $X = X_{ijk}$ *for each pair of distinct neighbors j and k in* N_i*, for each player i in N.*

We show that a binding contagion equilibrium exists and is efficient.

Theorem 1. *Any symmetric network G with degree d has a symmetric binding contagion equilibrium,* σ^{C}*, in which the collective stakes profile is* $\Phi^{\mathrm{C}} = (\phi^{\mathrm{C}}, \ldots, \phi^{\mathrm{C}})$*, where* ϕ^{C} *solves*

$$\frac{T(\phi)}{\phi} = \frac{1 + \frac{d\lambda}{r}}{1 + (d-1)X}. \tag{1}$$

This equilibrium Pareto dominates every distinct mutual effort equilibrium. Moreover, if the stage game satisfies strategic complementarity, no other equilibrium has a higher value.

The overarching intuition is simple: contagion strategies offer the harshest punishments, and a binding contagion equilibrium should attain the highest payoffs available using contagion strategies. Yet several subtleties arise, indicating that the correct intuition is not so straightforward. Instead, an involved argument is needed. We sketch the logic below, and offered a detailed proof in Section 4.3:

1. Lemma 3 shows that σ^{C} is the unique non-trivial binding contagion equilibrium.
2. Lemma 4 shows that Φ^{C} Pareto dominates all other collective stakes profiles in $\Psi(G)$, including those that are asymmetric.[10] It is straightforward to see that Φ^{C} Pareto dominates any other symmetric collective stakes profile, since all symmetric collective stakes profiles are multiples of $(1, \ldots, 1)$. We use Lemma 1 to show that for any asymmetric collective stakes profile in $\Psi(G)$, there exists a symmetric collective stakes profile that makes no player worse off and satisfies all cooperation phase incentives with slack.
3. Lemma 5 extends the conclusion to all mutual effort equilibria. The key step is to recognize that a feasible deviation for any pure strategy mutual effort equilibrium is to shirk on all of one's partners without deviating in the stakes selection stage. Such a deviation is unprofitable if and only if the collective stakes profile is in $\Psi(G)$. We extend the argument

---

[9] A non-empty symmetric network may contain isolated players. A stronger definition imposing "vertex transitivity" could rule out isolated players, but would not otherwise affect our results.

[10] A collective stakes profile Φ is symmetric if $\phi_{ij} = \phi_{kl}$ for every pair of links $\{ij\}$ and $\{kl\}$ in G; otherwise it is *asymmetric*.

to mixed strategy mutual effort equilibria using convexity and Jensen's inequality. These arguments together establish our first efficiency result, relying only on Assumption 1.

4. Lemmas 6–8 compare σ^{C} to equilibria in which players may shirk on the equilibrium path. The argument involves appropriately aggregating the players' utilities and incentive constraints. The key step shows that under strategic complementarity, if shirking occurs with positive probability on the equilibrium path, then there exists a mutual effort profile that attains the same aggregate payoffs without increasing any player's temptation to shirk. We then show that the non-trivial binding contagion equilibrium attains at least as much value.

4.3 Proof of Theorem 1

In this section, we describe the argument establishing Theorem 1 in detail. While we discuss and offer intuition for each lemma used below, formal proofs are relegated to Appendix A.

Throughout this subsection, fix a symmetric network G. In a mutual effort profile with collective stakes profile Φ, player i's expected payoff is $\frac{\lambda}{r}\sum_{j\in N_i}\phi_{ij}$. We say that a collective stakes profile $\Phi \succ_{PD} \Phi'$ if a mutual effort profile with collective stakes profile Φ Pareto dominates one with collective stakes profile Φ'.

We first characterize the binding contagion equilibria.

Lemma 3. *The only collective stakes profiles in $\mathcal{S}$ at which all cooperation phase incentive constraints bind are Φ^{C} and $(0,\ldots,0)$.*

Henceforth we ignore the degenerate collective stakes profile $(0,\ldots,0)$, and refer to σ^{C} as "the" binding contagion equilibrium. The following lemma compares its payoffs to all other contagion profiles that satisfy incentives on the equilibrium path, i.e., those that are *Nash equilibria*. It is straightforward to show that σ^{C} Pareto dominates all other symmetric contagion profiles, since each player's payoffs are increasing in the stakes and Φ^{C} is the highest symmetric collective stakes profile in $\Psi(G)$. We then use the convexity identified in Lemma 1 to prove that for any asymmetric contagion profile satisfying cooperation phase incentive constraints, there exists a symmetric contagion profile that makes no player worse off and makes at least one player better off—without violating the cooperation phase incentive constraints.

Lemma 4. *The binding contagion equilibrium σ^{C} Pareto dominates every other contagion profile that satisfies cooperation phase incentive constraints: $\Phi^{C} \succ_{PD} \Phi$ for every $\Phi \in \Psi(G)\backslash\{\Phi^{C}\}$.*

We expand the domain of comparison from contagion equilibria to all mutual effort equilibria.

Lemma 5. *For every mutual effort equilibrium, there exists a payoff-equivalent contagion profile with collective stakes profile* Φ *in* $\Psi(G)$*. It follows that every mutual effort equilibrium that is distinct from the binding contagion equilibrium* σ^{C} *is Pareto dominated by it.*

The intuition is most straightforward for pure-strategy mutual effort equilibria. One feasible deviation for player i is to shirk on every partner while announcing the stakes specified by the associated collective stakes profile Φ. Since such a deviation must not be profitable, her individual stakes profile in this pure strategy equilibrium, Φ_i, must satisfy $\mathrm{IC}_{ij}^{\mathrm{Coop}}$ for each of her links. Thus, Φ is in $\Psi(G)$ and can be supported by a contagion profile. Extending the argument to mixed strategy mutual effort equilibria (i.e., those in which the stakes in each relationship may be random) requires averaging the equilibrium path stakes along each link; necessarily, a pure strategy profile with these averaged stakes delivers the same payoffs for each player along the path of play, and Jensen's inequality implies that the deviation has an even lower payoff.

We turn to the next part of the theorem that compares the *values* of equilibria to argue that the non-trivial binding contagion equilibrium has the highest value. Throughout the discussion below, we fix an equilibrium σ and use the following steps to argue that the binding contagion equilibrium σ^{C} yields at least as much value. Since σ is stationary (Definition 2), we can denote by μ_{ij}^{S} the distribution of stakes that arises on link $\{ij\}$ along the equilibrium path. For a link $\{ij\}$ let $p_{ij}^{\mathrm{ww}}(\phi)$ be the on-path probability of mutual effort when stakes ϕ are realized, let $p_{ij}^{\mathrm{ws}}(\phi)$ be the on path probability that player i works while player j shirks, and let $p_{ij}^{\mathrm{ss}}(\phi)$ be the on-path probability of mutual shirking. Let u_{ij} be player i's expected equilibrium stage game payoff when link $\{ij\}$ is recognized, and let w_{ij} be her payoff from following the stake proposal strategy but deviating at the action stage to shirking regardless of the realized stakes; formally,

$$u_{ij} \equiv \int_0^\infty \left(p_{ij}^{\mathrm{ww}}(\phi)\phi + p_{ji}^{\mathrm{ws}}(\phi)T(\phi) - p_{ij}^{\mathrm{ws}}(\phi)V(\phi)\right)\, d\mu_{ij}^{\mathrm{S}},$$
$$w_{ij} \equiv \int_0^\infty T(\phi)\left(p_{ij}^{\mathrm{ww}}(\phi) + p_{ji}^{\mathrm{ws}}(\phi)\right)\, d\mu_{ij}^{\mathrm{S}}.$$

In these terms, the "total utility" of equilibrium σ is

$$U(\sigma) \equiv \frac{\lambda}{r}\sum_{i\in N}\sum_{j\in N_i} u_{ij} = \frac{\lambda}{r}\sum_{\{ij\}\in G}(u_{ij}+u_{ji}).$$

Our first claim is to argue that a counterpart of $\mathrm{IC}_{ij}^{\mathrm{Coop}}$ holds for every equilibrium σ.

Lemma 6. *For every equilibrium σ, every player i in N, and every neighbor j in N_i,*

$$w_{ij} + \sum_{k \in N_i \setminus \{j\}} X w_{ik} \le u_{ij} + \frac{\lambda}{r} \sum_{k \in N_i} u_{ik}. \tag{2}$$

The proof of Lemma 6 generalizes that of Lemma 5: one possible deviation for player i is to maintain the equilibrium path proposal strategies but to shirk on every partner. Since a contagion equilibrium has the lowest viscosity of any equilibrium, the left-hand side is a lower bound for the payoffs from this deviation.

The first step is to average the incentive constraints across links. We describe the "total deviation utility" by

$$W(\sigma) = \sum_{i \in N} \int_0^\infty e^{-rt} e^{-d\lambda t} \lambda \sum_{j \in N_i} \Big(w_{ij} + \sum_{k \in N_i \setminus \{j\}} X w_{ik} \Big)\, dt.$$

Observe that

$$U(\sigma) = \sum_{i \in N} \int_0^\infty e^{-rt} e^{-d\lambda t} \lambda \sum_{j \in N_i} \Big(u_{ij} + \frac{\lambda}{r} \sum_{k \in N_i} u_{ik} \Big)\, dt \ge W(\sigma),$$

where the equality follows from algebra and the inequality follows from (2).

Our second step is to eliminate shirking on the equilibrium path. We construct a mixed-strategy mutual effort profile that delivers as much value as σ and satisfies an analogous constraint.

Lemma 7. *If the stage game satisfies strategic complementarity, then for every equilibrium σ, there exists a mixed strategy mutual effort profile $\tilde{\sigma}$ such that $U(\tilde{\sigma}) \ge U(\sigma) \ge W(\sigma) \ge W(\tilde{\sigma})$.*

To convert σ into a mutual effort profile $\tilde{\sigma}$, replace every occurrence of work-shirk on the equilibrium path at stakes ϕ with mutual effort at stakes

$$f(\phi) = \max\left\{0, \frac{T(\phi) - V(\phi)}{2}\right\}.$$

By construction, action profile (work, work) at stakes $f(\phi)$ yields the same sum of payoffs as action profile (work, shirk) at stakes ϕ. Moreover, strict convexity of T (Assumption 1) and strategic complementarity (Definition 1) together guarantee that there is less total deviation utility in $\tilde{\sigma}$ than in σ. This is the only step that uses strategic complementarity.

Because $\tilde{\sigma}$ may not be an equilibrium, we cannot yet apply Lemma 5. Nevertheless, analogous arguments permit a comparison between the total utility of $\tilde{\sigma}$ and that of the non-trivial binding contagion equilibrium σ^{C}.

Lemma 8. *For any mixed-strategy mutual effort profile $\tilde{\sigma}$, if $U(\tilde{\sigma}) \geq W(\tilde{\sigma})$ then $U(\tilde{\sigma}) \leq U(\sigma^{\mathrm{C}})$.*

The argument constructs a symmetric contagion profile that attains the same value as $\tilde{\sigma}$ while satisfying the condition that total utility must exceed total deviation utility, and then shows that its stakes cannot be greater than ϕ^{C}. Together, these results imply that the non-trivial binding contagion equilibrium has at least as high a value.

5 Comparing Networks

Conventional wisdom holds that network architecture is important for cooperation that can be sustained in bilateral relationships, and existing studies have shed light on numerous important factors that contribute to cooperation. The critical friction in our framework is that players themselves are the only ones to observe their own relationships and our study of how a network may overcome this challenge focuses attention on a new determinant of cooperation, which we label as *viscosity*.

Section 5.1 describes viscosity, constructs an algorithm to compute it, and compares the viscosity of cliques to that of networks with a weakly lower maximal degree. Section 5.2 combines this result with Theorem 1 to compare binding contagion equilibria on cliques to equilibria on other networks. Section 5.3 characterizes networks that optimally balance the benefits from cooperation and the costs of linking.

5.1 The Viscosity of Networks

Viscosity measures how quickly behavior or information spreads in a network, as a function of players' patience. When player *i* shirks on player *j*, her consideration is the extent to which she can take advantage of player *k* prior to player *k* being infected. Unlike measures of propagation that would be used to study the spread of infectious disease or ideas, for example, a player's concern here is not how quickly the contagion spreads globally across the network,[11] but rather the speed with which it spreads—potentially via long and winding paths—back to her local neighborhood.

---

[11] In a study of public goods provision, Wolitzky (2012) develops a notion of "effective contagiousness" that offers an instructive comparison. Effective contagiousness measures how quickly the number of players who are contagious is expected to grow, following a deviation. That is, effective contagiousness is a measure of the global spread of contagion. Such a measure is appropriate to public goods environments because, in principle, all the players in the network can punish the deviator. Similarly, Golub and Jackson (2012) offer a measure of the speed of diffusion in their analysis of directed contagion processes on large random networks.

Viscosity is a measure of propagation that differs from those studied in other mechanisms of diffusion and depends in a subtle way on the structure of the network.

We develop some notation to describe paths by which contagion spreads. For a path ζ, let Z_ζ be its length, and for every $z \in \{1, \ldots, Z_\zeta\}$, let $\zeta(z)$ be the zth node in the path. For a set of players M such that $k \notin M$, we consider paths that begin in M, end at k, and do not travel through M:

$$s_k(M) \equiv \left\{\zeta : \zeta(1) \in M,\ \zeta(Z_\zeta) = k,\ \text{and } \zeta(z) \notin M \text{ for every } z > 1\right\}.$$

The set $s_k(M)$ contains all paths through which the contagion spreads from M to player k. We count the number of paths that infect a player outside M on the path to player k:

$$\tilde{s}_k(M) = \left\{i'j' \in G : i' \in M, \zeta(2) = j' \text{ for some } \zeta \in s_k(M)\right\}.$$

When the set of contagion players is M, one more player in the direction of player k is infected at the rate $|\tilde{s}_k(M)|\lambda$. Viscosity is computed recursively as $X_{ijk} = \chi_{ik}(\{i, j\})$, and for every M,

$$\begin{aligned}\chi_{ik}(M) &= \int_0^\infty e^{-rt} e^{-|\tilde{s}_k(M)|\lambda t}\left(\lambda \mathbb{1}(k \notin M) + \sum_{i'j' \in \tilde{s}_k(M)} \lambda \chi_{ik}(M \cup \{j'\})\right) dt \\ &= \frac{\lambda}{r + |\tilde{s}_k(M)|\lambda}\left(\mathbb{1}(k \notin M) + \sum_{i'j' \in \tilde{s}_k(M)} \chi_{ik}(M \cup \{j'\})\right).\end{aligned} \tag{3}$$

The term outside brackets represents the expected delay until a subsequent player becomes infected. The first term in brackets reflects that contagion may spread by means of player i meeting player k; the second term reflects that when contagion spreads to player j' the set of contagious players becomes $M \cup \{j'\}$. The recursive formulation highlights that viscosity depends only on λ/r and G. The example below calculates viscosity for simple four-player networks.

Example 1. *Consider the square network in Figure 3a, and suppose that $\lambda/r = 1$. Suppose that player 1 has deviated on 2 and is computing the discounted probability of meeting player 4 first. By (3), above,*

$$X_{124} = \frac{\lambda}{r + 2\lambda}\left(1 + \frac{\lambda}{r + 2\lambda}\right) = \frac{\lambda}{r + 2\lambda} + \left(\frac{\lambda}{r + 2\lambda}\right)^2 = \frac{4}{9}.$$

After player 1 shirks on player 2, $\frac{\lambda}{r+2\lambda}$ is the discounted probability that link $\{14\}$ is recognized before link $\{23\}$; conditional on link $\{23\}$ meeting first, $\left(\frac{\lambda}{r+2\lambda}\right)^2$ is the the probability link $\{14\}$ is

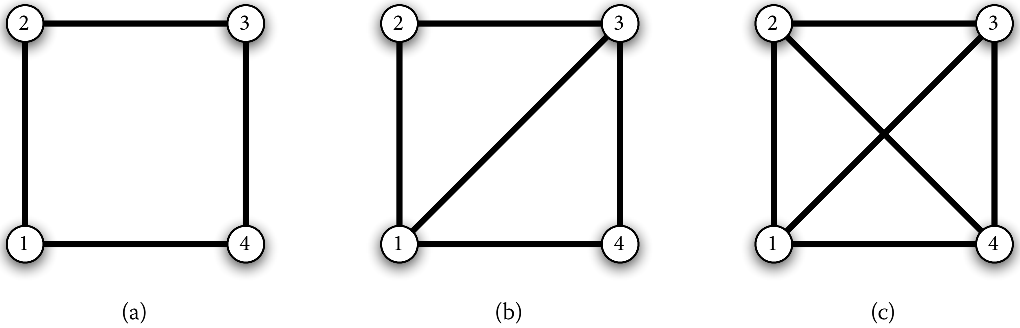

(a) (b) (c)

FIGURE 3

recognized before link $\{34\}$.

Adding the diagonal link $\{13\}$ *generates the network in Figure 3b; this link reduces viscosity throughout the network because there is an additional path by which information propagates. On this network we calculate the viscosity factors as follows:*

$$X_{124} = X_{123} = X_{213} = \frac{\lambda}{r+3\lambda}\left(1+\frac{\lambda}{r+3\lambda}\right)^2 = \frac{5}{16},$$

$$X_{132} = \frac{\lambda}{r+4\lambda}\left(1+2\frac{\lambda}{r+2\lambda}\right)^2 = \frac{1}{3}.$$

Completing the network by adding link $\{24\}$ *to construct the clique of degree 3 (Figure 3c) further reduces viscosity, but by a smaller margin:*

$$X_{124} = \frac{\lambda}{r+4\lambda}\left(1+\frac{\lambda}{r+3\lambda}\right)^2 = \frac{3}{10}.$$

While the above algorithm does not generate a general closed-form solution, it simplifies for the cases of trees and cliques, which yield the lower and upper bounds, respectively, on viscosity in all networks. On a *tree*—a network that lacks any cycles—the viscosity factor X_{ijk} is simply $\frac{\lambda}{r+\lambda}$, because the only way for player k to be infected by $\{i,j\}$ is when link $\{ik\}$ is recognized. Indeed, this is the greatest viscosity factor possible among a connected triple in any network.

Proposition 2. *For every network G and every pair of links $\{ij\}, \{ik\}$ in G, $X_{ijk} \leq \frac{\lambda}{\lambda+r}$. Moreover, $X_{ijk} < \frac{\lambda}{\lambda+r}$ if and only if G contains a cycle that includes both $\{ij\}$ and $\{ik\}$.*

In a clique, by contrast, every person in player i's component is a neighbor. We call a completely connected network in which there are $d+1$ players the *clique of degree d*; we denote such a network by $\overline{G}(d)$. For a clique, we can then dispense with proceeding recursively through the graph and instead consider only the number of infected neighbors in the contagion phase. This simplifies the viscosity factor to yield a closed form.

Proposition 3. *For the clique of degree d, and every pair of links* $\{ij\}$ *and* $\{ik\}$*,*

$$X_{ijk} = \overline{X}(d) \equiv \frac{1}{d-1} \sum_{m'=2}^{d} \left(\frac{1}{m'} \prod_{m=2}^{m'} \frac{\lambda m(d-m+1)}{r + \lambda m(d-m+1)} \right).$$

Though it may not be obvious from the expression above, larger cliques are less viscous than smaller cliques, the logic being that a larger clique presents more opportunities for the contagion to diffuse. Moreover, a clique of degree d is the least viscous among all networks in which no player has degree higher than d.

Theorem 2. *If the maximal degree in network G is no more than d, then for every pair of links* $\{ij\}, \{ik\}$ *in G,*

$$X_{ijk} \geq \overline{X}(d).$$

The inequality is strict if and only if the component of G that contains players i, j, and k is not a clique of degree d.

This result is intuitive insofar as a clique provides the shortest possible paths by which contagion can spread from player j to player k before player i meets player k. Since closed-form computation of viscosity factors for arbitrary networks is infeasible, the proof uses a more elementary coupling argument,[12] which, for every path in G that infects player k before link $\{ik\}$ is recognized, constructs a shorter path in $\overline{G}(d)$ that serves the same purpose. We should highlight that if G is a regular network of degree d, it has as many links per player as $\overline{G}(d)$ and so the lower viscosity emerges from the network architecture rather than by increasing link density.

5.2 Which Networks Support the Greatest Cooperation?

In this section, we combine the insights from analyzing the viscosity of a network with our characterization of contagion on symmetric networks to identify which networks support the most cooperation. A major challenge in ranking networks is that while contagion exists and is optimal

---

[12] We thank Ben Golub for encouraging this method of proof.

on symmetric networks, it is difficult to know which equilibria exist and are efficient on asymmetric networks for an arbitrary discount rate. From this perspective, it might seem infeasible to compare contagion on symmetric networks with an uncharacterized efficient equilibrium on asymmetric networks.

Yet, even without computing the set of equilibrium payoffs on asymmetric networks, we know one constraint that limits every equilibrium on every network: no player can do strictly better than her equilibrium payoff by shirking perpetually. Perhaps surprisingly, this simple fact is all that is needed to compare contagion on the clique of degree d with every equilibrium on any network in which no player has more than d partners.

We denote the stakes on each link in the (non-trivial) binding contagion equilibrium on the clique $\overline{G}(d)$ by $\overline{\phi}(d)$; the expected payoff for each player in this equilibrium is $\overline{u}(d) \equiv \left(\frac{d\lambda}{r}\right)\overline{\phi}(d)$.

Theorem 3. *Consider a network G in which the maximal degree is no more than d.*

(a) In every mutual effort equilibrium, no player attains a payoff that exceeds $\overline{u}(d)$; moreover, any player in a component of G that is not $\overline{G}(d)$ has a strictly lower payoff.

(b) If the stage game satisfies strategic complementarity, then the value of every equilibrium is less than $\overline{u}(d)$, and strictly less if at least one component of G is not $\overline{G}(d)$.

Cliques maximize the level of cooperation because they are the least viscous. If we restrict attention to mutual effort equilibria on arbitrary networks, then the clique "Pareto dominates" every other network that has the same maximal degree: each player in the clique is at least as well off, and some are strictly better off.[13] Our proof actually establishes a stronger property than that stated above: each player obtains greater payoffs even on a per-link basis in the contagion equilibrium in the clique than in any distinct equilibrium on any network.[14]

If we allow for equilibria in which players shirk on the equilibrium path, then non-clique networks are not necessarily Pareo dominated. Nonetheless, if the game satisfies strategic complementarity then the binding contagion equilibrium on the clique generates the highest average

---

[13] This comparison ignores any remainder that might arise if the population of players is not divisible by $d + 1$. Formally, we could make Pareto comparisons between two networks by replicating them both to form two "replica networks," each containing the least common multiple of the original populations.

[14] One might conjecture an even stronger result: adding a link to any arbitrary network should increase the level of cooperation attainable on each link in equilibrium, since each new link makes the network less viscous. While this is true for Nash equilibria, in which incentive constraints off the path of play are ignored, it need not hold for perfect Bayesian equilibria, in which the addition of links influences incentives both on and off the path of play. Adding a link may relax on-path incentives but it may impose more severe off-path constraints, and thereby shrink the set of attainable per-link payoffs.

value. As in the case of mutual effort equilibria, we further prove that the contagion equilibrium on the clique has a higher value per link than any other equilibrium on any other network.

The proof builds on ideas from our previous results. For simplicity, consider a pure strategy mutual effort equilibrium σ on an arbitrary network G of maximal degree d, and suppose that player i's component is not a clique of degree d. That player i would not do strictly better by perpetually shirking generates the incentive constraint $\mathrm{IC}_{ij}^{\mathrm{Coop}}$ (on p. 13) in which the viscosity X_{ijk} corresponds to that of G. Replacing this viscosity with that of the clique, $\overline{X}(d)$, introduces greater slack into the incentive constraint, since by Theorem 2, $\overline{X}(d) < X_{ijk}$. Put differently, the clique induces faster punishment. We use this property to show that there exists an individual stakes profile for the clique that attains the same payoffs for player i as equilibrium σ and satisfies the cooperation phase incentives on the clique. Thus, as argued in Theorem 1, the binding contagion equilibrium of the clique must generate higher payoffs for player i. The comparison of contagion to equilibria on G in which players shirk on the equilibrium path makes parallel observations, and builds on ideas from Lemmas 6–8 to show that the binding contagion equilibrium on the clique must generate a higher value.

The conclusion to draw from Theorem 3 is that regardless of players' patience or the frequency with which they interact, the clique of degree d is the optimal network among all networks in which the maximal degree is d. Thus, if the population is divisible by $d+1$, organizing society into cliques of degree d is a better social architecture than any other in which the maximal degree is d.

5.3 Balancing the Benefits and Costs of Linking

Establishing relationships can be costly, so a network should optimally balance the benefits of linking with its costs. This section characterizes networks that achieve the optimal tradeoff.

We model linking costs in a reduced-form manner: if player i has d_i links, then she pays the linking cost $c(d_i)$ at time 0. A player's *net payoff* is the sum of her expected equilibrium payoffs from interactions (henceforth *interaction payoffs*) minus the linking costs that she incurs at time 0. An equilibrium's *net value* is the average net payoff in society.

The linking cost function c is non-decreasing and, as a normalization, satisfies $c(0) = 0$. We study linking costs that belong in one of two categories below: for $d \geq 1$, linking costs are

1. *Concave* if $c(d)/d$ is non-increasing.
2. *Strictly convex* if $c(d)/d$ is strictly increasing.

A special case of concave linking costs is that in which linking costs are linear.

There are two benefits of forming a link: it creates a new relationship, and it indirectly benefits other relationships if it completes a cycle (ignoring potential off-path complications). The first benefit is internalized by the partners who form the link, while the latter is a positive externality. Since the marginal linking cost is either flat or decreasing when $c(\cdot)$ is concave, an extreme solution dominates in which if it is worthwhile to link to anyone then it is better to link to everyone. Since it is symmetric, the complete network has a binding contagion equilibrium in which off-path incentives are guaranteed. In this equilibrium, each player earns a net payoff of $\overline{u}(n-1) - c(n-1)$; by symmetry this is also the equilibrium's net value.

Theorem 4. *Suppose that linking costs are concave.*

1. *Consider any mutual effort equilibrium on any incomplete network. If a player has a non-negative net payoff, then her net payoff is strictly less than $\overline{u}(n-1) - c(n-1)$.*

2. *Consider any equilibrium on any incomplete network, and suppose the stage game satisfies strategic complementarity. If the equilibrium's net value is non-negative, then its net value is strictly less than $\overline{u}(n-1) - c(n-1)$.*

Thus the complete network Pareto dominates other networks if players follow a mutual effort equilibrium, and is utilitarian optimal even when shirking is allowed on the equilibrium path.

Comparisons across networks are more subtle if $c(\cdot)$ is convex. Now a player may not find it in her own interest to link to all other players, but others always benefit from her doing so. Since a player's preference is no longer aligned with others' preferences, we focus on Pareto efficiency, and prove that the "optimal" clique is Pareto efficient: if linking costs are sufficiently convex, there exists some (generically unique) *optimal clique* $\overline{G}(d^*)$ such that d^* uniquely maximizes $\overline{u}(d) - c(d)$. If there is no interior optimum, then the optimal clique is the complete network.

Theorem 5. *Suppose that linking costs are strictly convex. For every network G and every mutual effort equilibrium, if there exists a player whose net payoff is strictly greater than $\overline{u}(d^*) - c(d^*)$, then there exists another player whose net payoff is strictly less than $\overline{u}(d^*) - c(d^*)$.*

The logic is straightforward: consider a network G in which there exists a component in which some player earns a strictly higher net payoff than $\overline{u}(d^*) - c(d^*)$. Necessarily, this component cannot be a clique since by definition no clique obtains a net value greater than $\overline{u}(d^*) - c(d^*)$. Within this component, consider the player with the highest degree, d'. By Theorem 3, her interaction payoff is strictly less than $\overline{u}(d')$, the payoff she would obtain in a clique of her own degree,

$\overline{G}\left(d'\right)$; meanwhile, her linking costs in G and $\overline{G}\left(d'\right)$ are identical. It follows that her net payoff is strictly less than $\overline{u}(d') - c(d')$.

Stronger comparisons emerge when contrasting contagion on cliques to smaller classes of networks. Consider the class of networks that are *regular*, i.e., those in which all players share the same degree. Since the linking costs on a regular network are identical to those on the clique with the same degree, Theorem 3 has a direct corollary.

Corollary 1. *No mutual effort equilibrium on any regular network attains payoffs that exceed $\overline{u}(d^*) - c(d^*)$. Moreover, if the stage game satisfies strategic complementarity, no equilibrium on any regular network attains a higher value than $\overline{u}(d^*) - c(d^*)$.*

While our analysis of this case is limited to either mutual effort equilibria or regular networks, it has direct implications for the behavior we might expect on potentially efficient irregular networks: if there is a network in which the player with the highest degree has payoffs that strictly exceed $\overline{u}(d^*) - c(d^*)$, then her payoffs are incompatible with her having to work in every relationship on the equilibrium path. If the stage game satisfies strategic complementarity, then she must be gaining a larger share of the surplus in one of her relationships by being permitted to shirk on the equilibrium path.

5.4 An Alternative Social Norm with Communication

So far, we have compared networks using repeated game equilibria that do not rely upon explicit communication. As with the prior literature, we use contagion to characterize an optimal social norm for which the off-path propagation of punishment is incentive compatible. The same network comparisons also apply for alternative enforcement mechanisms considered by the literature (e.g. Dixit 2003), in which word of mouth communication plays an important role.

Specifically, suppose that which whenever a pair meets, the two first exchange truthful messages that describe their prior histories before setting stakes and making their effort choices. Consider any *ostracism norm* in which innocent players punish guilty players, but not others who are innocent. When a player shirks, her partner truthfully communicates the news to others, and this information percolates through the network. As in contagion, a guilty player's only opportunity for further gain is to meet her other partners before they learn of her guilt. Indeed, a player's incentive to work under such a social norm is identical to that under contagion: when facing an individual stakes profile Φ_i, player i's incentives to work on $\{ij\}$ are captured exactly by $\mathrm{IC}_{ij}^{\mathrm{Coop}}$ on p. 13. Thus our network comparison results hold for for any ostracism social norm in which innocent players are assumed to communicate truthfully.[15]

6 Applications

While we have framed our results in the context of a prisoners' dilemma with variable stakes, our findings are relevant for other settings that feature relationships with moral hazard.

6.1 Favor Exchange

Prior to the advent of modern market institutions, one prominent role for community enforcement and social capital was to ensure that players could rely on reciprocity norms to ensure that they would perform favors for each other. The role of social networks in a *trading favors* context has been emphasized by earlier sociological analyses (e.g. Coleman 1988), and studied in recent works in economics (e.g. Möbius 2001; Jackson, Rodriguez-Barraquer, and Tan 2012).

Consider a network G of connected players in N. If players i and j are connected, then with Poisson density λ player i will be called upon to do a favor for player j. When player i performs a favor of level $\phi > 0$ for player j, she incurs a cost of $C(\phi)$ that satisfies $C'(0) = 0$ and $\lim_{\phi \to \infty} C'(\phi) = \infty$; player j's benefit from this favor is simply ϕ. As before, players i and j are the only ones to observe what transpires in their bilateral relationship, and all players share a common discount rate r. We restrict attention to equilibria that are stationary on the equilibrium path.[16]

Unlike the model of Section 2, efforts along each link are asynchronous, triggered only when a player is called upon to help her partner.[17] The optimal favor, ϕ^*, solves $C'(\phi) = 1$; were favors contractible, such favors could be enforced without community involvement. Our interest is in studying the extent to which favor exchanges can be self-enforcing in a networked community.

In a pure strategy equilibrium, we denote the favor performed by player i for player j by ϕ_{ij}, and let $\Phi = (\phi_{ij}, \phi_{ji})_{\{ij\} \in G}$ be the profile of equilibrium path favors. In a contagion equilibrium, a player stops doing favors for all her partners when she observes that any of her partners has failed to do a favor for her. Thus, when player i fails to perform a favor for partner j, she can count on partner k to perform favors for her only so long as he has not observed any deviation from

---

[15] Our companion paper, Ali and Miller (2012), studies ostracism equilibria in which truthful communication must be incentive compatible.

[16] In other words, we look at equilibria in which for every on-path history, whenever player i is called upon to do a favor for player j, her randomization over favor sizes is independent of the history.

[17] With asynchronous efforts, players no longer need to agree explicitly to the stakes of their interaction obviating the stake selection stage in Section 2.

on-path behavior. Let $\vec{x}_{ijk}$ be the *ijk* viscosity factor, analogous to x_{ijk}; then the equilibrium path incentive constraint is

$$\sum_{k\in N_i\setminus\{j\}} \phi_{ki} \int_0^\infty e^{-r\tau}\vec{x}_{ijk}(\tau)\lambda\, d\tau \leq -C(\phi_{ij}) + \frac{\lambda}{r}\sum_{k\in N_i}\left(\phi_{ki} - C(\phi_{ik})\right). \tag{$\vec{IC}_{ij}$}$$

This incentive constraint is similar to the cooperation phase incentive of the prisoners' dilemma ($\text{IC}_{ij}^{\text{Coop}}$ on p. 13). Let $\vec{X}_{ijk} \equiv \int_0^\infty e^{-r\tau}\vec{x}_{ijk}(\tau)$ denote the viscosity factor of favor exchanges.

Proposition 4. *For every network G, the viscosity factor in favor exchanges is identical to that of the basic model.*

This similarity permits us to port results to this application, although the different setting generates a subtle variation on Theorem 1.

Proposition 5. *In the favor exchange model, every symmetric network G with degree d has a symmetric binding contagion equilibrium in which the favors performed on each link,* ϕ^{F}, *solve*

$$\frac{C(\phi)}{\phi} = \frac{d\lambda - r(d-1)X}{1 + d\lambda}. \tag{4}$$

If $\phi^{\text{F}} < \phi^*$, *this equilibrium has higher value than every distinct equilibrium on G.*

Importantly, the solution to (4) is decreasing in the viscosity factor, so networks that have lower viscosity increase the volume of favors and the payoffs of each player.

Proposition 6. *In the favor exchange model, so long as* $\phi^{\text{F}} < \phi^*$, *the binding contagion equilibrium on the clique of degree d has a higher value than every distinct equilibrium on every network with maximal degree no more than d.*

We omit the proof of this result because it is virtually identical to that of Theorem 3. As in the basic model, the friction of pairwise monitoring is best overcome by a network with low viscosity, so cliques foster the highest level of favor provision. As in our discussion in Section 5.4, these network comparisons are not limited to a contagion equilibrium: ostracism equilibria would generate the same network comparison.

Within the context of informal risk-sharing, the literature has documented patterns that describe how more risk-sharing is associated with lower viscosity. Udry (1990, 1994) and Fafchamps and Lund (2003) describe how a large fraction of risk-sharing is done by within the same village or within kinship groups despite the potential diversification gains of sharing risk more broadly.

Jackson, Rodriguez-Barraquer, and Tan (2012) find that favor exchange networks in southern India have a high fraction of links in which the linked players have a common neighbor, thereby fostering short indirect paths.

6.2 Enforcing Trade in Networked Markets

In many market contexts, traders rely on informal enforcement mechanisms to support cooperative behavior. While many have emphasized the importance of networks in the informal enforcement of commerce (e.g. Greif 1993; McMillan and Woodruff 1999; Banerjee and Duflo 2000; Dixit 2003), there are few models of self-enforcing trading networks.[18] This section applies our framework to two-sided markets in which trade emerges from each buyer's and seller's consideration of future trading opportunities.

In a networked market, there is a group $\mathcal{N}^B = \{1, \dots, b\}$ of buyers and a group $\mathcal{N}^S = \{1, \dots, s\}$ of sellers. Every link in the network connects a buyer to a seller. For notational convenience, define the networked market $\vec{G}$ to be a subset of $\mathcal{N}^B \times \mathcal{N}^S$; i.e., treat every link as being directed from a buyer to a seller, and denote a link (b, s) as simply bs. We let N_b be the set of sellers to which a buyer b is connected and N_s the set of buyers that are connected to seller s.

When a link is selected, a buyer and seller can trade: the seller chooses what quality (or quantity) q of good to produce, and the buyer chooses a payment p in exchange for the good. We call the product quality of trade that arises on the equilibrium path the "level of trade," and interpret $q = 0$ as the absence of trade. Their resulting payoffs are $q - p$ for the buyer and $p - C(q)$ for the seller, where $C(\cdot)$ is a strictly increasing, strictly convex cost function, with $C'(0) = 0$ and $\lim_{q \to \infty} C'(q) = \infty$.

In each interaction, the buyer and seller make choices simultaneously.[19] The buyer has an incentive to renege on making the payment while the seller has an incentive to produce a low quality good. Were external enforcement available, the buyer and seller could write a contract that forces the seller to produce at the quality q^* that solves $C'(q) = 1$, and forces the buyer to make a payment in $[C(q^*), q^*]$. We study trading relations that rely purely on self-enforcement.

---

[18] Ahn and Suominen (2001) study a monopolist firm facing a sequence of consumers who observe their neighbors' past interactions with the firm. Deb and González-Díaz (2011) study the product choice game with anonymous random matching, and construct subtle dynamic equilibria in which the community builds trust over time. Fainmesser (2012) studies a variant of the product choice game in which the tension is a capacity constraint, and characterizes networks that foster the greatest cooperation under the assumption that all enforcement is bilateral; Fainmesser and Goldberg (2012) expand the analysis to community enforcement, under the assumption that information diffuses exogenously.

[19] E.g., even if the seller must deliver the product before the buyer pays, the buyer cannot verify its quality before paying.

As before, we focus on equilibria in which behavior on the equilibrium path is stationary.

In a pure strategy equilibrium, let $(q_{bs}, p_{bs})_{bs \in \vec{G}}$ be the profile of trades on the equilibrium path. In a contagion equilibrium, trade continues so long as no player has deviated in the past, but once a player deviates, every infected player deviates and so market cooperation unravels. For a buyer b and two sellers s and s', we let $x_{bss'}$ denote the analogue of x_{ijk}: if players b and s are infected at time 0, $x_{bss'}(\tau)$ is the probability that s' is not infected; we denote the analogue for a seller s and two buyers b and b' by $x_{sbb'}$. The equilibrium path incentive constraints for a buyer b and seller s are

$$q_{bs} + \sum_{s' \in N_b \setminus \{s\}} q_{bs'} \int_0^\infty e^{-r\tau} x_{bss'}(\tau) \lambda \, d\tau \leq q_{bs} - p_{bs} + \frac{\lambda}{r} \sum_{s' \in N_b} (q_{bs'} - p_{bs'}) \,, \tag{$\vec{IC}_{bs}$}$$

$$p_{bs} + \sum_{b' \in N_s \setminus \{b\}} p_{bs'} \int_0^\infty e^{-r\tau} x_{sbb'}(\tau) \lambda \, d\tau \leq p_{bs} - C(q_{bs}) + \frac{\lambda}{r} \sum_{b' \in N_s} (p_{b's} - C(q_{b's})) \,. \tag{$\vec{IC}_{sb}$}$$

These constraints bound the level of trade, as well as how much value can be extracted from each buyer through prices. Since the computation of viscosity is unaffected by whether the network is directed or undirected, the viscosity factors are identical to those of the undirected representation of the network. However, the bipartite structure demands a different notion of symmetry:

Definition 7. *A networked market $\vec{G} \subset \mathcal{N}^B \times \mathcal{N}^S$ is* ***symmetric on each side*** *if for any two links $bs, b's' \in \vec{G}$ there exists a directed graph automorphism*[20] *f such that $f(s) = s'$ and $f(b) = b'$.*

If the networked market is symmetric on each side ("symmetric" for short), then analogous to Lemma 2, there exists a common viscosity factor on each side—X_B for the buyers and X_S for the sellers. We let d_B and d_S denote the degrees of the buyers and sellers, respectively, in a symmetric networked market. The analogue of Theorem 1 is:

Proposition 7. *Every symmetric networked market $\vec{G}$ has a symmetric binding contagion equilibrium, in which the quality of trade, $\vec{q}$, on each link solves*

$$\frac{C(q)}{q} = \left(\frac{\lambda d_S - r(d_S - 1) X_S}{r + \lambda d_S} \right) \left(\frac{\lambda d_B - r(d_B - 1) X_B}{r + \lambda d_B} \right), \tag{5}$$

and the price paid by each buyer on the equilibrium path is

$$\vec{p} = \left(\frac{\lambda d_B - r(d_B - 1) X_B}{r + \lambda d_B} \right) \vec{q}.$$

---

[20] A directed graph automorphism is a bijection $f : \mathcal{N}^B \cup \mathcal{N}^S \to \mathcal{N}^B \cup \mathcal{N}^S$ such that if $bs \in \vec{G}$, then $f(b)f(s) \in \vec{G}$.

This equilibrium generates a higher average level of trade than every distinct equilibrium. Moreover, if $\vec{q} < q^*$, then this equilibrium generates higher value than any other equilibrium.

It follows that decreasing the viscosity on either the buyers' side or the sellers' side increases not only the level of trade, but also equilibrium prices. Interestingly, decreasing the buyer-side viscosity X_B decreases $(\vec{q} - \vec{p})/\vec{q}$, which measures the fraction of the buyer's payoff that she keeps in each trade. With lower viscosity, she can be counted on to pay more to each seller without violating her cooperation phase incentive constraint. Analogously, each seller charges a lower "markup," $\frac{\vec{p}-C(\vec{q})}{C(\vec{q})}$, when the seller-side viscosity X_S is lower, because lower viscosity reduces his temptation to skimp on quality.

We now consider which networked market maximizes the social gains from trade. The *bipartite clique* of degrees (d_B, d_S) is a bipartite network if there are d_B sellers and d_S buyers, and each buyer is connected to each seller; we denote it by $\overline{G}\,(d_B, d_S)$, its buyer side viscosity by $\overline{X}_B(d_B, d_S)$, its seller side viscosity by $\overline{X}_S(d_B, d_S)$, and the level of trade in its binding contagion equilibrium by $\overline{q}(d_B, d_S)$. We prove the analogue of Theorem 2 and Theorem 3. (Although a different coupling argument is needed, the overall proof is similar.)

Proposition 8. *Consider a bipartite networked market $\vec{G}$ in which no seller is connected to more than d_B buyers, and no buyer is connected to more than d_S sellers. The lowest viscosity factor for each side is at least that of the bipartite clique of degrees (d_B, d_S); i.e., for every triple of links $bs, b's, bs'$ in $\vec{G}$,*

$$X_{bss'} \geq \overline{X}_B(d_B, d_S), \quad \text{and} \quad X_{sbb'} \geq \overline{X}_S(d_B, d_S)$$

Each inequality is strict if the component of G that contains these players is not the bipartite clique. Every equilibrium of $\vec{G}$ generates a lower average level of trade $\overline{q}(d_B, d_S)$, and generates a lower value if $\overline{q}(d_B, d_S) < q^*$.

This result highlights how our result on the importance of cliques generalizes in a straightforward manner to networked markets. While we have focused on a two-sided incentive issue in which both buyer and seller would myopically wish to deviate, the results are similar if the incentive issue is only one-sided. For example, suppose that players on one side of the market (e.g. buyers) can commit to the terms of trade but those on the other (e.g. sellers) cannot: by maximizing the diffusion of information, cliques generate the strongest self-enforcement for each side of the market.[21]

---

[21] A one-sided incentive issue also arises if in each interaction, the buyer and seller were making choices sequentially

This analysis highlights how viscosity is better suited to this kind of analysis than measures that focus on connected triples (e.g. clustering or support coefficients), which cannot distinguish between different bipartite graphs since no bipartite graph has a connected triple. Viscosity, by contrast, uses the global architecture to measure the transmission of punishments that influence each player's incentives to fulfill trades as planned.

6.3 Social Collateral

Viscosity offers a natural measure for the *social collateral* within a network of partnerships: in each relationship, players can trust each other to a greater degree when they recognize that a player's deviation incurs the loss of social collateral embodied in other relationships. Decomposing cooperation phase incentives makes this clear: notice that $\mathrm{IC}_{ij}^{\mathrm{Coop}}$ on p. 13 can be re-written as

$$T(\phi_{ij}) - \phi_{ij}\left(1 + \frac{\lambda}{r}\right) \leq \sum_{k \in N_i \setminus \{j\}} \left(\frac{\lambda}{r}\phi_{ik} - T(\phi_{ik})X_{ijk}\right). \tag{6}$$

The left-hand side measures the extent to which cooperation on link $\{ij\}$ exceeds that supportable by simply bilateral enforcement—i.e., the bilateral stakes ϕ^{B}. Were ϕ_{ij} equal to ϕ^{B}, this term would be zero; at stakes greater than ϕ^{B}, it represents how much of the cooperation in link $\{ij\}$ is being subsidized by the other relationships. The right-hand side measures the level of social collateral available to subsidize cooperation on link $\{ij\}$. For each of player i's other relationships, the measure aggregates how much player i gains from subsequent cooperation minus how much she would lose were she to be punished before she could shirk. Lower levels of viscosity increase the level of social collateral, and as we have seen, the clique uses social collateral in the strongest possible way.

Social collateral strengthens both relationships within the network and those with outsiders who would trade with, lend to, or hire members of the network. We illustrate this effect by considering employee referrals and hiring choices. Hiring workers through referrals can mitigate moral hazard by taking advantage of social collateral.[22]

Suppose that a firm F must fill a vacancy for which effort, not skill, is important. If the hired worker exerts effort e, she bears the convex cost $\gamma(e)$ whereas the firm accrues the output e. The challenge is that the firm cannot contract on output even if it is observable. When hiring from the

---

so that one of the sides could not gain from its deviation.

[22] We thank Kalle Moene and Gaute Torsvik for encouraging this application.

spot market, the firm's only available punishment for shirking is to fire the worker. If the worker is not fired, the firm and worker meet at Poisson distributed times with rate parameter λ_f. In the absence of social collateral, the firm's problem is to select "efficiency wages" (Shapiro and Stiglitz 1984):

$$\max_{e,w\geq 0} e - w \quad \text{s.t. } w \leq (w - \gamma(e))\left(1 + \frac{\lambda_f}{r}\right).$$

Network referrals offer the firm an additional instrument to induce effort in the form of social collateral. Suppose that the recruit, player *i*, is referred to the firm by an existing worker, player *j*. Suppose for simplicity that both are part of a symmetric favor-trading network in which the common viscosity factor is X and each player does favors of level ϕ. It is understood that if player *i* shirks, the firm will immediately inform player *j*, after which information about player *i*'s deviation propagates through the network. The recruit's incentives to work arise from both her efficiency wages and her ongoing benefits from favor exchange. Thus, the firm's problem is

$$\max_{e,w\geq 0} e - w \quad \text{s.t. } w \leq (w - \gamma(e))\left(1 + \frac{\lambda_f}{r}\right) + \underbrace{\frac{d\lambda}{r}(\phi - C(\phi)) - (d-1)X\phi}_{\text{Social Collateral}}.$$

Social collateral enables the firm to reduce wages without discouraging effort. Reductions in viscosity increase social collateral through two channels. First, a reduction in viscosity increases the chance that those in her network will not do favors for her in their subsequent interactions. Second, communities with lower viscosity enforce a higher level of favors, and so being excluded from favor exchange is a harsher punishment.[23]

This analysis has direct implications for hiring choices: a firm would prefer to recruit on the basis of referrals, especially from tightly-knit communities, and will pay those hired through networks less than those hired on the spot market. These predictions are broadly consistent with patterns that have been found in low-skill labor markets in developing countries. Moreover, our explanation resonates with that suggested by prior research: informal peer punishment can mitigate moral hazard and make hiring through a network preferable.[24] A tightly-knit community creates the social collateral that makes it an attractive source of employees.

---

[23] If the level of favors is selected according to Proposition 5, the social collateral term is simply $C(\phi^{\text{F}})$, which is higher in less viscous networks.

[24] See Kugler (2003), Iversen, Sen, Verschoor, and Dubey (2009), Heath (2011), and Dhillon, Iversen, and Torsvik (2012).

Our focus on moral hazard complements the existing understanding of network referrals. Most closely related, Karlan, Möbius, Rosenblat, and Szeidl (2009) elucidate how network-based trust can mitigate issues of *adverse selection* by inducing workers to refer only recruits with high ability. Their analysis focuses on the market for skilled labor and they find that referred recruits are paid more than counterparts hired from the spot market. Together, their framework and ours suggest that network referrals are important for both high and low skill labor, for different reasons, and the impact of networks on wages depends on whether adverse selection or moral hazard is the more paramount consideration.

7 Discussion

This paper characterizes networks that optimally sustain cooperation. Our main result (Theorem 3) compares the binding contagion equilibrium on a clique of degree d to equilibria on all networks that have maximal degree d. The binding contagion equilibrium on the clique Pareto dominates all mutual effort equilibria on these other networks, and has higher average payoffs than any equilibrium on them. As we emphasize in Section 5.4, though contagion offers one mechanism to achieve the highest possible cooperation, other social norms may also implement the same equilibrium path payoffs. In the remainder of this section, we comment on various features of our framework.

Variable Stakes We represent the level of cooperation in a mutual effort equilibrium by the endogenously selected stakes at which cooperation is incentive compatible. In our view, permitting individuals to select the stakes of their relationship is a realistic formulation of partnerships in which individuals choose the terms at which to cooperate. Such cooperative arrangements are ubiquitous: in risk-sharing arrangements, individuals choose how much self-insurance they can attain; in trading and employer-employee relationships, the seller of a good or service chooses how much effort to exert, and the the buyer chooses how much to pay.

Ghosh and Ray (1996) and Kranton (1996) were the first to note the relevance of variable stakes in community enforcement, but to elucidate a different force: building cooperation over time helps screen out myopic players and deters patient players from shirking and re-matching with a new partner.[25] Our stylized framework for stakes and stake selection departs from theirs,

---

[25] See Watson (1999, 2002) and Athey, Calvano, and Jha (2010) for related insights. Haag and Lagunoff (2007) and Wolitzky (2012) use a continuous action environment in a local interaction setting in which a player takes a single action with respect to her entire community to measure cooperation.

but identical results would hold for many different formulations of variable stakes. In terms of the stake selection protocol, all that is needed is that it be sufficiently permissive: for every $\phi > 0$, there must be some strategy profile of the partnership that implements ϕ with probability 1.[26] It also would suffice for players to select stakes only at time 0, rather than at each interaction. Moreover, similar results would also hold if, as in Ghosh and Ray (1996), players in each partnership simultaneously chose actions from a continuum in which higher actions benefit the partner but come at a greater cost.[27]

Apart from realism, the inherent flexibility of variable stakes simplifies analysis and exposition considerably. In contrast to the standard repeated games approach of fixing the stage game payoffs and then identifying sets of discount factors for which cooperation arises, we can identify the maximal level of cooperation given a fixed level of patience and then directly compare payoffs across equilibria and networks at the same discount rate. Were the stakes ϕ fixed, we would be compelled to distinguish networks and equilibria by the sets of parameters for which incentive conditions are satisfied, which is both indirect and less transparent. Moreover, a technical challenge that emerges in fixed stakes environments is in verifying both equilibrium path incentives and the credibility of punishments off the equilibrium path.

Stationary Equilibrium An important limitation of our work is our restriction to equilibria in which behavior is stationary on the equilibrium path. Although these equilibria correspond most closely to those studied in applications and simplify the comparison of network architectures, the restriction is *with* loss of generality. We consider it important to characterize rich history-dependent equilibria for each level of patience and note that there are few, if any, techniques at this stage to find optimal history-dependent equilibria for each level of patience in this rich environment with private monitoring.[28] Insofar as our focus is on the significance of *viscosity* across networks, we expect our qualitative insights regarding cliques to extend to nonstationary equilibria.

---

[26] For example, similar results would hold if players were required to both propose ϕ for ϕ to be selected, or the average of two proposals were selected.

[27] This would be analogous to the favor exchange environment of Section 6.1 in which both players simultaneously performed favors for each other.

[28] Our restriction to stationarity rules out belief-free equilibria in which players' randomizations are influenced by their experiences on the equilibrium path. Prior work (Takahashi 2010; Deb 2011) has illustrated how such history-dependence can foster the incentive to cooperate in similar community enforcement settings, but focusing on the behavior of very patient players in fixed-stakes environments.

Contagion on Asymmetric Networks From the standpoint of our analysis, it suffices to show existence and optimality of contagion equilibria on cliques. Because we can compare contagion on cliques to other equilibria on other networks without knowing the precise form that those equilibria might take, we do not need to prove the existence of contagion more broadly.

Nevertheless, a generalized version of contagion equilibria do exist in asymmetric networks: an asymmetric network can be partitioned into symmetric subnetworks and contagion equilibria from these different subnetworks can be patched.

Definition 8. *The* ***symmetric decomposition*** *of a network G is a partition $\mathcal{P}$ of G such that each partition element is a symmetric graph.*

A *generalized contagion profile* applies a contagion profile to each subnetwork in $\mathcal{P}$ and treats the interactions in different subnetworks as if they were in different games. Since this is a straightforward notion, we relegate its technical description to Appendix B. Using generalized contagion, we can guarantee existence on asymmetric networks.

Proposition 9. *Every network G has a generalized contagion equilibrium with strictly positive stakes on each link. For every cycle in G, there exists a generalized contagion equilibrium in which the equilibrium path stakes on that cycle strictly exceed ϕ^{B}, the stakes from bilateral enforcement.*

Generalized contagion profile may be appealing from the perspective of off-path behavior. In particular, keeping each subnetwork strategically independent of the rest of the network bounds the spread of contagion after a deviation. In this way, cooperation throughout most of the network can be preserved off the equilibrium path. Accordingly, in our framework—in contrast to that of Jackson, Rodriguez-Barraquer, and Tan (2012)—containing contagion does not require any restructuring of the network *per se*, but can instead be achieved through strategic behavior.

Equilibrium Networks We focus on optimal networks rather than equilibrium networks throughout this paper. Many well-studied network formation games generate multiple equilibrium networks, often including efficient networks. To see this most transparently, consider a two-sided linking process in which each player simultaneously proposes the partnerships she wishes to engage in, and the $\{ij\}$ partnership forms if and only if both players i and j propose it. Once formed, the network is common knowledge. It is straightforward to see that any network G can arise in an equilibrium of this game if it yields a non-negative net payoff for each player, via the following strategy profile: if network G arises then players follow an equilibrium that guarantees each of them non-negative net payoffs, but if any other network forms then each player perpetually shirks. This simple punishment deters players from deviating in the network formation stage.

References

Illtae Ahn and Matti Suominen. Word-of-mouth communication and community enforcement. *International Economic Review*, 42(2):399–415, 2001.

S. Nageeb Ali and David A. Miller. Ostracism. Working paper, 2012.

Attila Ambrus, Markus Möbius, and Adam Szeidl. Consumption risk-sharing in social networks. Working paper, February 2010.

Susan Athey, Emilio Calvano, and Saumitra Jha. A theory of community formation and social hierarchy. Working paper, 2010.

Abhijit V. Banerjee and Esther Duflo. Reputation effects and the limits of contracting: A study of the Indian software industry. *Quarterly Journal of Economics*, 115(3):989–1017, 2000.

Pierpaolo Battigalli. Strategic independence and perfect Bayesian equilibria. *Journal of Economic Theory*, 70(1):201–234, July 1996.

B. Douglas Bernheim and Michael Whinston. Multimarket contact and collusive behavior. *Rand Journal of Economics*, 21(1):1–26, 1990.

Timothy Besley, Stephen Coate, and Glenn Loury. The economics of rotating savings and credit associations. *American Economic Review*, 83(4):792–810, 1993.

Francis Bloch, Garance Genicot, and Debraj Ray. Informal insurance in social networks. *Journal of Economic Theory*, 143(1):36–58, November 2008.

James S. Coleman. Social capital in the creation of human capital. *American Journal of Sociology*, 94(S1):S95–S120, 1988.

Joyee Deb. Cooperation and community responsibility: A folk theorem for repeated matching games with names. Working paper, May 2011.

Joyee Deb and Julio González-Díaz. Community enforcement beyond the prisoner's dilemma. Working paper, July 2011.

Amrita Dhillon, Vegard Iversen, and Gaute Torsvik. Employee referral, social proximity and worker discipline. Working paper, September 2012.

Avinash K. Dixit. Trade expansion and contract enforcement. *Journal of Political Economy*, 111 (6):1293–1317, 2003.

Glenn Ellison. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61(3):567–588, July 1994.

Marcel Fafchamps and Susan Lund. Risk-sharing networks in rural Philippines. *Journal of Development Economics*, 71:261–287, 2003.

Itay P. Fainmesser. Community structure and market outcomes: A repeated games-in-networks approach. *American Economic Journal: Microeconomics*, 4(1):1–38, 2012.

Itay P. Fainmesser and David A. Goldberg. Cooperation in partly observable networked markets. Working paper, July 2012.

Leon Festinger, Stanley Schachter, and Kurt Back. *Social Pressures in Informal Groups*. MIT Press, Cambridge, M.A., 1948.

Drew Fudenberg and Jean Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, April 1991.

Parikshit Ghosh and Debraj Ray. Cooperation in community interaction without information flows. *Review of Economic Studies*, 63(3):491–519, 1996.

Edward L. Glaeser, David Laibson, and Bruce Sacerdote. An economic approach to social capital. *Economic Journal*, 112(483):F437–F458, 2002.

Benjamin Golub and Matthew O. Jackson. How homophily affects the speed of learning and best-response dynamics. *Quarterly Journal of Economics*, 127(3):1287–1338, August 2012.

Mark S. Granovetter. Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, pages 481–510, 1985.

Mark S. Granovetter. The impact of social structure on economic outcomes. *Journal of Economic Perspectives*, 19(1):33–50, 2005.

Avner Greif. Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition. *American Economic Review*, 83(3):525–548, 1993.

Gene M. Grossman and Elhanan Helpman. *Special Interest Politics*. MIT Press, Cambridge, MA, 2002.

Matthew Haag and Roger Lagunoff. Social norms, local interaction, and neighborhood planning. *International Economic Review*, 47(1):265–296, February 2006.

Matthew Haag and Roger Lagunoff. On the size and structure of group cooperation. *Journal of Economic Theory*, 135(1):68–89, 2007.

Joseph E. Harrington, Jr. Cooperation in a one-shot prisoners' dilemma. *Games and Economic Behavior*, 8:364–377, 1995.

Rachel Heath. Why do firms hire using referrals? Evidence from Bangladeshi garment factories.

Working paper, October 2011.

Vegard Iversen, Kunal Sen, Arjan Verschoor, and Amaresh Dubey. Job recruitment networks and migration to cities in India. *Journal of Development Studies*, 45(4):522–543, 2009.

Matthew O. Jackson. *Social and economic networks*. Princeton University Press, Princeton, NJ, 2008.

Matthew O. Jackson, Tomas Rodriguez-Barraquer, and Xu Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, 102(5):1857–1897, 2012.

Michihiro Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59 (1):63–80, 1992.

Dean Karlan, Markus Möbius, Tanya Rosenblat, and Adam Szeidl. Trust and social collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, August 2009.

Rachel E. Kranton. The formation of cooperative relationships. *Journal of Law, Economics, and organization*, 12(1):214–233, 1996.

Adriana D. Kugler. Employee referrals and efficiency wages. *Labour Economics*, 10(5):531–556, 2003.

Steffen Lippert and Giancarlo Spagnolo. Networks of relations and word-of-mouth communication. *Games and Economic Behavior*, 72(1):202–217, May 2011.

Giovanni Maggi. The role of multilateral institutions in international trade cooperation. *American Economic Review*, 89(1):190–214, 1999.

John McMillan and Christopher Woodruff. Interfirm relationships and informal credit in Vietnam. *Quarterly Journal of Economics*, 114(4):1285–1320, 1999.

Markus Möbius. Trading favors. Working paper, May 2001.

Francesco Nava and Michele Piccione. Efficiency in repeated games with local interaction and uncertain monitoring. Working paper, July 2012.

Werner Raub and Jeroen Weesie. Reputation and efficiency in social interactions: An example of network effects. *American Journal of Sociology*, 96(3):626–654, 1990.

Carl Shapiro and Joseph E. Stiglitz. Equilibrium unemployment as a worker discipline device. *American Economic Review*, 74(3):433–444, 1984.

Satoru Takahashi. Community enforcement when players observe partners' past play. *Journal of Economic Theory*, 145(1):42–62, January 2010.

Christopher Udry. Credit markets in northern Nigeria: Credit as insurance in a rural economy. *World Bank Economic Review*, 4(3):251–269, 1990.

Christopher Udry. Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria. *Review of Economic Studies*, 61(3):495–526, 1994.

Joel Watson. Starting small and renegotiation. *Journal of Economic Theory*, 85(1):52–90, 1999.

Joel Watson. Starting small and commitment. *Games and Economic Behavior*, 38(1):176–199, 2002.

Alexander Wolitzky. Cooperation with network monitoring. *Review of Economic Studies*, 2012.

Appendix A Main Proofs

A.1 Proofs for Section 4

Proof of Lemma 1 on p. 13. First consider Φ_i and Φ_i' such that for some ik, $\phi_{ik} \neq \phi_{ik}'$. Consider the incentives for player i along link $\{ij\}$ that are induced by the individual stakes profile $\alpha\Phi_i + (1-\alpha)\Phi_i'$ in which $\alpha \in (0,1)$:

$$\begin{aligned}
&T(\alpha\phi_{ij} + (1-\alpha)\phi_{ij}') + \sum_{k\in N_i\setminus\{j\}} X_{ijk} T(\alpha\phi_{ik} + (1-\alpha)\phi_{ik}') \\
&\quad < \alpha T(\phi_{ij}) + (1-\alpha)T(\phi_{ij}') + \sum_{k\in N_i\setminus\{j\}} X_{ijk}\left(\alpha T(\phi_{ik}) + (1-\alpha)T(\phi_{ik}')\right) \\
&\quad = \alpha\Big(T(\phi_{ij}) + \sum_{k\in N_i\setminus\{j\}} X_{ijk}T(\phi_{ik})\Big) + (1-\alpha)\Big(T(\phi_{ij}') + \sum_{k\in N_i\setminus\{j\}} X_{ijk}T(\phi_{ik}')\Big) \\
&\quad \leq \alpha\Big(\phi_{ij} + \frac{\lambda}{r}\sum_{k\in N_i}\phi_{ik}\Big) + (1-\alpha)\Big(\phi_{ij}' + \frac{\lambda}{r}\sum_{k\in N_i}\phi_{ik}'\Big) \\
&\quad = \alpha\phi_{ij} + (1-\alpha)\phi_{ij}' + \frac{\lambda}{r}\sum_{k\in N_i}\Big(\alpha\phi_{ik} + (1-\alpha)\phi_{ik}'\Big),
\end{aligned} \tag{7}$$

where the strict inequality follows from the strict convexity of T and Jensen's inequality, the first equality is obtained by re-arranging, the weak inequality is from $\Phi_i, \Phi_i' \in \Psi_{ij}(G)$, and the final equality is from rearranging. Therefore, $\alpha\Phi_i + (1-\alpha)\Phi_i' \in \operatorname{int}\Psi_{ij}(G)$ if $\Phi_i, \Phi_i' \in \Psi_{ij}(G)$.

As the intersection of convex sets, $\Psi(G)$ is convex. Because $\Psi(G)$ contains $(\phi^{\mathrm{B}}, \ldots, \phi^{\mathrm{B}})$ and $(0, \ldots, 0)$, it follows from above that for every $\alpha \in (0,1)$, $(\alpha\phi^{\mathrm{B}}, \ldots, \alpha\phi^{\mathrm{B}}) \in \operatorname{int}\Psi_{ij}(G)$ for every ij. Therefore, $(\alpha\phi^{\mathrm{B}}, \ldots, \alpha\phi^{\mathrm{B}})$ is in the interior of $\Psi(G)$. □

Proof of Lemma 2 on p. 15. Consider two triples $\{i,j,k\}$ and $\{i',j',k'\}$ such that $\{j,k\}$ are distinct neighbors of i and $\{j',k'\}$ are distinct neighbors of i'. It suffices to show that $x_{ijk}(t) = x_{i'j'k'}(t)$ for every t. Because the network is symmetric, there exists a graph automorphism g such that $g(ij) = i'j'$ and $g(ik) = i'k'$.

Let $(\tau_z)_{z=1}^{\infty}$ be an ordered list of link recognition times and $(l_z)_{z=1}^{\infty}$ be the list of links in their order of recognition. Suppose there exist $\bar{z}$ such that $\{ik\} = l_{\bar{z}}$; i.e., link $\{ij\}$ is realized at time $\tau_{z'}$. Suppose also that at time 0 the set of contagious players is $\{i,j\}$, and that all players play according to an arbitrary contagion profile σ. The realization of link recognitions $\{\tau_z, l_z\}_{z=1}^{\bar{z}}$ can be coupled with the permuted recognition realization $\{\tau_z, g(l_z)\}_{z=1}^{\bar{z}}$. It follows that given the permuted recognition realization, if at time 0 the set of contagious players is $\{i', j'\}$ and all players play according to any contagion profile σ', then player k' is contagious at time $\tau_{\bar{z}}$ if and only if player k was contagious at time $\tau_{\bar{z}}$ given the original recognition realization. Since recognition realization $\{\tau_z, l_z\}_{z=1}^{\bar{z}}$ has the same probability density as recognition realization $\{\tau_z, g(l_z)\}_{z=1}^{\bar{z}}$ regardless of the equilibrium, it follows that $x_{ijk}(t) = x_{i'j'k'}(t)$. □

Proof of Lemma 3 on p. 16. Every symmetric collective stakes profile $(\phi, \ldots, \phi)$ must satisfy the following simplified version of $\mathrm{IC}_{ij}^{\mathrm{Coop}}$:

$$T(\phi)\,(1 + (d-1)X) \le \phi\left(1 + d\frac{\lambda}{r}\right). \tag{8}$$

First we show that every binding contagion equilibrium is symmetric. For any $i \in N$ and distinct $j, k \in N_i$, if the incentive constraints bind on both link $\{ij\}$ and link $\{ik\}$, it follows that

$$T(\phi_{ij}) - \phi_{ij} + T(\phi_{ik})X = T(\phi_{ik}) - \phi_{ik} + T(\phi_{ij})X. \tag{9}$$

Since T is strictly convex, $\phi_{ij} = \phi_{ik}$. Therefore, in a binding contagion equilibrium, every player must have the same stakes across all her links, and these must set (8) to equality. By Assumption 1, the only solutions to that equation are $\phi = 0$ and $\phi = \phi^{\mathrm{C}}$. □

Proof of Lemma 4 on p. 16. The equilibrium payoff of a symmetric contagion profile in which the stakes are ϕ is $d\frac{\lambda}{r}\phi$, and every such profile in $\Psi(G)$ satisfies (8). Because ϕ^{C} are the greatest stakes that satisfy (8), it follows that the binding contagion equilibrium σ^{C} Pareto dominates every other symmetric contagion profile in $\Psi(G)$.

Consider an asymmetric collective stakes profile Φ in $\Psi(G)$ and consider the set of players

$$Y = \Big\{ i \in N : \ \Phi_i \text{ is asymmetric and } \sum_{k \in N_i} \phi_{ik} \ge \sum_{k \in N_j} \phi_{jk} \text{ for every } j \in N \Big\}.$$

First suppose that Y is non-empty and consider $i \in M$. From Lemma 2, it follows that each of the $d!$ permutations of Φ_i is in $\Psi_{ij}(G)$ for every $j \in N_i$. The equally weighted convex combination of the set of permutations of Φ_i is $\overline{\Phi}_i = (\overline{\phi})_{j \in N_i}$ such that $\overline{\phi} = \frac{1}{d}\sum_{j \in N_i} \phi_{ij}$. Since Φ_i is asymmetric, Lemma 1 implies that $\overline{\Phi}_i \in \operatorname{int} \Psi_{ij}(G)$ for every $j \in N_i$. Since G is symmetric, each player's cooperation phase incentives are strictly satisfied if her individual stakes profile correspond to $\overline{\Phi}_i$ and therefore, the collective stakes profile $\overline{\Phi} = (\overline{\phi}, \ldots, \overline{\phi})$ is in $\operatorname{int} \Psi(G)$. Thus, we have constructed a symmetric collective stakes profile $\overline{\Phi}$ in the interior of $\Psi(G)$ that makes no player worse off relative to the asymmetric collective stakes profile Φ. Since $\Phi^{\mathrm{C}} >_{PD} \overline{\Phi}$, it follows that $\Phi^{\mathrm{C}} >_{PD} \Phi$.

Now suppose that Y is empty: then there exists player i such that Φ_i is symmetric and player i's payoffs are at least as high as those of any other player. Consider the collective stakes profile $\Phi' = (\Phi_i, \ldots, \Phi_i)$.

Therefore $\Phi' >_{PD} \Phi$, and, since either $\Phi' = \Phi^{\mathrm{C}}$ or $\Phi^{\mathrm{C}} >_{PD} \Phi'$, it follows that $\Phi^{\mathrm{C}} >_{PD} \Phi$. □

Proof of Lemma 5 on p. 17. As in the text, it is helpful to begin with a pure strategy mutual effort equilibrium. Consider a collective stakes profile Φ associated with a pure strategy mutual effort equilibrium and consider player i who interacts with partner j on the equilibrium path. Following equilibrium strategies generates the same expected payoff for player i as the LHS of $\mathrm{IC}_{ij}^{\mathrm{Coop}}$. One possible deviation for player i is to maintain her equilibrium proposal strategy but to shirk in each interaction. The payoff from this deviation is at least that on the RHS of $\mathrm{IC}_{ij}^{\mathrm{Coop}}$. This deviation is unprofitable if and only if $\mathrm{IC}_{ij}^{\mathrm{Coop}}$ is satisfied. Since this holds for every player i and link $\{ij\}$, it follows that Φ is in $\Psi(G)$.

Now consider a mixed strategy mutual effort equilibrium σ that is distinct from every contagion profile, i.e., there exists at least one partnership such that the equilibrium path stakes on that link are stochastic. For each link $\{ij\} \in G$, let μ_{ij}^{S} be the distribution of stakes on link $\{ij\}$ associated with σ. As above, one possible deviation for player i is to follow the mixed stakes announcement specified by σ but to then shirk in every interaction. For this to be unprofitable in an on-path history in H_{ij}, the following must be satisfied for every ϕ in the support of μ_{ij}^{S}:

$$T(\phi) + \sum_{k \in N_i \setminus \{j\}} \left(\int_0^\infty T(\phi_{ik})\, d\mu_{ik}^{\mathrm{S}} \right) X \leq \phi + \frac{\lambda}{r} \sum_{k \in N_i} \int_0^\infty \phi_{ik}\, d\mu_{ik}^{\mathrm{S}}. \tag{10}$$

Let $\overline{\phi}_{ij} = \int_0^\infty \phi_{ij}\, d\mu_{ij}^{\mathrm{S}}$, and consider the individual stakes profile $\overline{\Phi}_i = (\overline{\phi}_{ij})_{j \in N_i}$. Re-writing (10) yields

$$\begin{aligned}
\overline{\phi}_{ij} + \frac{\lambda}{r} \sum_{k \in N_i} \overline{\phi}_{ik} &= \int_0^\infty \phi_{ij}\, d\mu_{ij}^{\mathrm{S}} + \frac{\lambda}{r} \sum_{k \in N_i} \int_0^\infty \phi_{ik}\, d\mu_{ik}^{\mathrm{S}} \\
&\geq \int_0^\infty T(\phi)\, d\mu_{ij}^{\mathrm{S}} + \sum_{k \in N_i \setminus \{j\}} \left(\int_0^\infty T(\phi_{ik})\, d\mu_{ik}^{\mathrm{S}} \right) X \\
&\geq T(\overline{\phi}_{ij}) + \sum_{k \in N_i \setminus \{j\}} T(\overline{\phi}_{ik}) X,
\end{aligned}$$

where the equality follows by construction, the first inequality arises from integrating (10) over all ϕ in the support of μ_{ij}^{S} and simplifying, and the second inequality follows from the convexity of T and Jensen's inequality. Note that the second inequality is strict if μ_{ik}^{S} is non-degenerate for any $k \in N_i$. Therefore, the contagion profile with collective stakes profile $\overline{\Phi} = (\overline{\Phi}_i)_{i \in N}$ is payoff equivalent to σ, and $\overline{\Phi}$ is in $\Psi(G)$. Since by assumption σ has a non-degenerate mixture, at least one incentive constraint in the payoff equivalent contagion profile $\overline{\Phi}$ is slack.

Up to this point in the proof we have not used the fact that the network is symmetric. Now, using symmetry, we apply Lemma 4 to find that σ^{C} Pareto dominates σ. □

Proof of Lemma 6 on p. 17. Consider a deviation by player i in which she maintains her equilibrium stakes proposal strategy but shirks in every relationship. When she shirks along some link $\{ik\}$, by definition she earns a stage payoff of w_{ik} if player k still doesn't know that she has deviated from the equilibrium path; otherwise she earns a stage payoff of no less than zero.

The second step of the argument is to compare the viscosity of equilibrium σ to that of contagion. Suppose that player i shirks on parter j at time 0. We argue that the probability with which player k knows that behavior is off the equilibrium path is no greater than $x_{ijk}(t)$. Fix a sequence of link recognitions that takes place in $(0, \infty)$ such that by time t, no player has shirked on player k by time t in a contagion equilibrium. That implies that player k has not met player i before time t or someone who has met … who has met player i at any time in $[0, t)$. Thus, even in equilibrium σ, player k cannot know that behavior is off the equilibrium path for that sequence of link recognitions. Thus, the viscosity of equilibrium σ is at least that of contagion.

The above observations establish the result: the payoff from this deviation is at least that on the LHS of (2), and σ being an equilibrium implies that this deviation cannot be profitable. □

Proof of Lemma 7 on p. 18. Since $f(\phi) \equiv \max\left\{0, \frac{1}{2}\left(T(\phi) - V(\phi)\right)\right\}$ need not be a bijection, for a set $A \subset \mathbb{R}_+$, define $f^{-1}(A) \equiv \left\{\phi \in \mathbb{R}_+ : f(\phi) \in A\right\}$. If the stage game satisfies strategic complementarity, then $f(\phi) < \phi/2$. Therefore

$$T(f(\phi)) < T\left(\frac{\phi}{2}\right) < \frac{T(\phi)}{2}, \tag{11}$$

where the first inequality follows from T being strictly increasing, and the second inequality from T being strictly convex.

Using f, we construct a new distribution of stakes ρ^{S}_{ij} from μ^{S}_{ij}. For every measurable subset A of $\mathbb{R}_+$, let

$$\rho^{\mathrm{S}}_{ij}(A) = \int_{\phi\in A} p^{\mathrm{ww}}_{ij}(\phi)\, d\mu^{\mathrm{S}}_{ij} + \int_{\phi\in A} \int_{\hat{\phi}\in f^{-1}(\{\phi\})} \left(p^{\mathrm{ws}}_{ij}(\hat{\phi}) + p^{\mathrm{ws}}_{ji}(\hat{\phi})\right) d\mu^{\mathrm{S}}_{ij}\, d\phi + \mathbb{1}(0 \in A) \int_{\phi} p^{\mathrm{ss}}_{ij}(\phi)\, d\mu^{\mathrm{S}}_{ij}.$$

Consider a mutual effort profile in which if $i < j$, for every history $h \in H_{ij}$, $\tilde{\sigma}^{\mathrm{S}}_i(h) = \rho^{\mathrm{S}}_{ij}$, and if $i > j$, $\tilde{\sigma}^{\mathrm{S}}_i(h) = \sup_{\phi\in\mathrm{Supp}(\rho^{\mathrm{S}}_{ij})} \phi$. Such a stake proposal strategy profile implements the distribution ρ^{S}_{ij} in each link $\{ij\}$. Since $\tilde{\sigma}$ is a mutual effort profile, players work on the equilibrium path.

To argue that $U(\tilde{\sigma}) \geq U(\sigma)$, let $\tilde{u}_{ij}$ and $\tilde{w}_{ij}$ be the analogues of u_{ij} and w_{ij}. By construction of f, u, and $\tilde{u}$,

$$\begin{aligned}
\tilde{u}_{ij} + \tilde{u}_{ji} &= \int_0^\infty 2\phi\, d\rho^{\mathrm{S}}_{ij} \\
&= \int_0^\infty 2\phi p^{\mathrm{ww}}_{ij}(\phi)\, d\mu^{\mathrm{S}}_{ij} + \int_0^\infty 2\phi \int_{\hat{\phi}\in f^{-1}(\{\phi\})} \left(p^{\mathrm{ws}}_{ij}(\hat{\phi}) + p^{\mathrm{ws}}_{ji}(\hat{\phi})\right) d\mu^{\mathrm{S}}_{ij}\, d\phi \\
&= \int_0^\infty 2\phi p^{\mathrm{ww}}_{ij}(\phi)\, d\mu^{\mathrm{S}}_{ij} + \int_0^\infty \int_{\hat{\phi}\in f^{-1}(\{\phi\})} 2f(\hat{\phi})\left(p^{\mathrm{ws}}_{ij}(\hat{\phi}) + p^{\mathrm{ws}}_{ji}(\hat{\phi})\right) d\mu^{\mathrm{S}}_{ij}\, d\phi \\
&\geq \int_0^\infty 2\phi p^{\mathrm{ww}}_{ij}(\phi)\, d\mu^{\mathrm{S}}_{ij} + \int_0^\infty \int_{\hat{\phi}\in f^{-1}(\{\phi\})} \left(T(\hat{\phi}) - V(\hat{\phi})\right)\left(p^{\mathrm{ws}}_{ij}(\hat{\phi}) + p^{\mathrm{ws}}_{ji}(\hat{\phi})\right) d\mu^{\mathrm{S}}_{ij}\, d\phi \\
&= u_{ij} + u_{ji}.
\end{aligned}$$

Since this holds for every $\{ij\}$ in G, it follows that $U(\tilde{\sigma}) \geq U(\sigma)$.

We take the analogous steps for W. Notice that

$$W(\sigma) = \frac{\lambda}{r + d\lambda}(1 + (d-1)X) \sum_{\{ij\} \in G} (w_{ij} + w_{ji}),$$

and so to prove that $W(\tilde{\sigma}) \leq W(\sigma)$, it suffices to establish that for every $\{ij\}$ in G, $\tilde{w}_{ij} + \tilde{w}_{ji} \leq w_{ij} + w_{ji}$:

$$\begin{aligned}
\tilde{w}_{ij} + \tilde{w}_{ji} &= \int_0^\infty 2T(\phi)\, d\rho_{ij}^{\mathrm{S}} \\
&= \int_0^\infty 2T(\phi) p_{ij}^{\mathrm{ww}}(\phi)\, d\mu_{ij}^{\mathrm{S}} + \int_0^\infty 2T(\phi) \int_{\hat{\phi} \in f^{-1}(\{\phi\})} \big(p_{ij}^{\mathrm{ws}}(\hat{\phi}) + p_{ji}^{\mathrm{ws}}(\hat{\phi})\big)\, d\mu_{ij}^{\mathrm{S}}\, d\phi \\
&= \int_0^\infty 2T(\phi) p_{ij}^{\mathrm{ww}}(\phi)\, d\mu_{ij}^{\mathrm{S}} + \int_0^\infty \int_{\hat{\phi} \in f^{-1}(\{\phi\})} 2T(f(\hat{\phi}))\big(p_{ij}^{\mathrm{ws}}(\hat{\phi}) + p_{ji}^{\mathrm{ws}}(\hat{\phi})\big)\, d\mu_{ij}^{\mathrm{S}}\, d\phi \\
&\leq \int_0^\infty 2T(\phi) p_{ij}^{\mathrm{ww}}(\phi)\, d\mu_{ij}^{\mathrm{S}} + \int_0^\infty \int_{\hat{\phi} \in f^{-1}(\{\phi\})} T(\hat{\phi})\big(p_{ij}^{\mathrm{ws}}(\hat{\phi}) + p_{ji}^{\mathrm{ws}}(\hat{\phi})\big)\, d\mu_{ij}^{\mathrm{S}}\, d\phi \\
&= w_{ij} + w_{ji},
\end{aligned}$$

where the inequality follows from (11), and the rest from construction. □

Proof of Lemma 8 on p. 19. We first use an argument analogous to that of Lemma 5. Let ρ_{ij}^{S} be the distribution of stakes along link $\{ij\}$ in the mutual effort profile $\tilde{\sigma}$. Let $\overline{\phi}_{ij} = \int_0^\infty \phi_{ij}\, d\rho_{ij}^{\mathrm{S}}$, and consider the contagion profile $\overline{\sigma}$ with collective stakes profile $(\overline{\phi}_{ij})_{i \in N, j \in N_i}$. Let $\overline{u}_{ij}$ and $\overline{w}_{ij}$ be the analogous stage game payoffs from working and shirking on the path of play. By construction, $\overline{u}_{ij} = \tilde{u}_{ij}$. It follows from the convexity of T and Jensen's inequality that $\overline{w}_{ij} \leq \tilde{w}_{ij}$ for each link $\{ij\}$.

The next step is to argue that there exists a symmetric contagion profile that satisfies the same constraint. Let $\phi^* = \frac{1}{|G|} \sum_{\{ij\} \in G} \overline{\phi}_{ij}$, and consider a contagion profile σ^* in which those are the stakes in each link. Clearly, $U(\sigma^*) = U(\overline{\sigma})$. That $W(\sigma^*) \leq W(\overline{\sigma})$ follows from

$$\sum_{\{ij\} \in G} (w_{ij}^* + w_{ji}^*) = 2|G|T(\phi^*) = 2|G|T\Big(\sum_{\{ij\} \in G} \overline{\phi}_{ij}/|G|\Big) \leq 2 \sum_{\{ij\} \in G} T(\overline{\phi}_{ij}) = \sum_{\{ij\} \in G} (\overline{w}_{ij} + \overline{w}_{ji}),$$

where the inequality is again by convexity of T and Jensen's inequality.

The final step is to show that $\phi^* \leq \phi^{\mathrm{C}}$. Observe that $U(\sigma^*) \geq W(\sigma^*)$ implies that

$$\frac{2|G|\lambda}{r}\phi^* \geq \frac{2|G|\lambda}{r + d\lambda}\big(1 + (d-1)X\big)T(\phi^*) \quad \implies \quad \frac{T(\phi^*)}{\phi^*} \leq \frac{r + d\lambda}{r + rX(d-1)} = \frac{T(\phi^{\mathrm{C}})}{\phi^{\mathrm{C}}}.$$

Since T is strictly convex, the conclusion follows. □

A.2 Proofs for Section 5

Proof of Theorem 2 on p. 22. We say that a sequence of link recognitions $(\tau_z, l_z)_{z=1}^{\infty}$ contains a path ζ if there exists a sub-string of $(l_z)_{z=1}^{\infty}$ that corresponds to ζ.

Consider a network G in which the maximal degree is d, and fix a triple $\{i, j, k\}$ such that $\{ij\}, \{ik\} \subset G$. Let G_{-i} be the network that results from deleting all of player i's links. If there is no path between players j and k in G_{-i}, then i, j, k are not part of a cycle and the result is implied by Proposition 2. So suppose that there is at least one path between players j and k in G_{-i}.

Consider all the paths in G_{-i} from player j to player k: let ζ be a generic such path, and let S be the set of all such paths. We consider a partition of S, such that two paths are in the same partition element if and only if $\zeta(2) = \zeta'(2)$ and $\zeta(Z_{\zeta} - 1) = \zeta'(Z_{\zeta'} - 1)$. In other words, the second and second to last players in the path coincide. Since player j and player k each has at most $(d-1)$ neighbors in G_{-i}, there are at most $(d-1)^2$ partition elements. We denote a partition element by S_{uv} if $\zeta(2) = u$ and $\zeta(Z_{\zeta} - 1) = v$ for every $\zeta \in S_{uv}$.

Now, consider an arbitrary triple $\{\bar{i}, \bar{j}, \bar{k}\}$ in $\overline{G}(d)$. Let $\overline{N}_m$ be player m's neighborhood in $\overline{G}(d)$. Consider injective functions $g : N_j \to \overline{N}_{\bar{j}}$ and $h : N_k \to \overline{N}_{\bar{k}}$ such that (i) if $\{jk\} \in G$, then $g(k) = \bar{k}$ and $h(j) = \bar{j}$; (ii) $g(v) = h(v)$ for every $v \in N_j \cap N_k$; and (iii) $\bar{i} = g(i) = h(i)$.

Suppose that at time 0, players j and $\bar{j}$ are contagious in networks G and $\overline{G}(d)$ respectively, and contagion spreads whenever the link between an infected player and an uninfected player is recognized. We couple the processes on G and $\overline{G}(d)$ as follows. Given a sequence of link recognitions $(l_z)_{z=1}^{\infty}$ on G,

1. Player $\bar{j}$ meets player $g(u) \in \overline{N}_{\bar{j}}$ whenever player j meets player $u \in N_j$.
2. Player $\bar{k}$ meets player $h(v) \in \overline{N}_{\bar{k}}$ whenever player k meets player $v \in N_k$.
3. For any path $\zeta \in S_{uv}$ contained in $(l_z)_{z=1}^{\infty}$ for which $Z_{\zeta} \geq 4$, player $g(u) \in \overline{N}_{\bar{j}}$ meets player $h(v) \in \overline{N}_{\bar{k}}$ when player u meets player $\zeta(3)$ in ζ.

Consider any sequence of link recognitions $(l_z)_{z=1}^{\infty}$ on G such that at time 0 the set of contagious players is $\{i, j\}$, and such that player k becomes contagious before link $\{ik\}$ is recognized. Such a sequence contains some path ζ that is completed (all its links are recognized) before link $\{ik\}$ is recognized. By considering all possible such paths, we argue that in the coupled link recognition process on $\overline{G}(d)$, player $\bar{k}$ must become contagious before meeting $\bar{i}$. The most straightforward case is that of $Z_{\zeta} = 2$: it follows that link $\{jk\} \in G$, and the unique path $(\{jk\}) \in S_{kj}$ is completed before link $\{ik\}$ meets; by the coupling, player $\bar{j}$ must meet player $\bar{k}$ (and so player $\bar{k}$ is contagious) before link $\{\bar{i}\bar{k}\}$ is recognized. The next case is that of $Z_{\zeta} = 3$: it follows that there exists a player $u \in N_j \cap N_k \backslash \{i\}$, and the path is $\zeta \in S_{uu}$; by the coupling, first player $\bar{j}$ must meet player $g(u) \neq \bar{i}$, and then player $g(u)$ must meet player $\bar{k}$ (and so player $\bar{k}$ is contagious), both before link $\{ik\}$ is recognized. Finally, consider any path $\zeta \in S_{uv}$ with $Z_{\zeta} \geq 4$: by the coupling, first player $\bar{j}$ must meet player $g(u) \neq \bar{i}$, next player $g(u)$ must meet player $h(v) \neq \bar{i}$, then player $h(v)$ must meet player $\bar{k}$ (and so player $\bar{k}$ is contagious), all before link $\{ik\}$ is recognized.

Thus, for any sequence of link recognitions $(l_z)_{z=1}^{\infty}$ on G in which link $\{ik\}$ meets only after player k has already infected by a path from player j, there is a coupled sequence of link realizations on $\overline{G}(d)$ in which the analogous event occurs. Note that the density of link realizations on $\overline{G}\,(d)$ is at least as high as that in G.

The other form of link realization by which player k may be infected before $\{ik\}$ is recognized is if player i meets a neighbor $m \in N_i \backslash \{j, k\}$, and then a path from player m infects player k before link $\{ik\}$ is recognized. The proof applies *mutatis mutandi* with player m taking the stead of player j. Since player i has at most d neighbors in network G, the density of these link realizations is no greater in network G than in $\overline{G}\,(d)$.

Finally, we argue that the inequality is strict if the component of G that contains players i, j, and k is anything other than $\overline{G}\,(d)$. If $|N_j \cap N_k| < d$, or if $\{jk\} \notin G$, then $\overline{G}\,(d)$ has additional paths by which player $\bar{k}$ may become infected before link $\{\overline{ik}\}$ is recognized, and similarly if $|N_i \cap N_k| < d$. ☐

Proof of Theorem 3 on p. 23. Consider a network G with maximal degree d. First consider a mixed strategy mutual effort equilibrium. By the argument of Lemma 5, it follows that there exists a payoff equivalent pure strategy mutual effort profile whose on path collective stakes profile Φ is in $\Psi(G)$.

Observe that for every player i and link $\{ij\}$ in network G,

$$T(\phi_{ij}) + \sum_{k \in N_i \backslash \{j\}} T(\phi_{ik}) \overline{X}(d) \leq T(\phi_{ij}) + \sum_{k \in N_i \backslash \{j\}} T(\phi_{ik}) X_{ijk} \leq \phi_{ij} + \frac{\lambda}{r} \sum_{k \in N_i} \phi_{ik}, \tag{12}$$

in which the first inequality follows from $X_{ijk} \geq \overline{X}(d)$ (established by Theorem 2) and the second inequality from $\Phi_i \in \Psi_{ij}(G)$. There are d_i inequalities of this form for player i, one for each of her links. Averaging over them yields

$$\begin{aligned} \frac{1}{d_i} \sum_{j \in N_i} \phi_{ij} + \frac{\lambda}{r} \sum_{k \in N_i} \phi_{ik} &\geq \frac{1}{d_i} \sum_{j \in N_i} \Big(T(\phi_{ij}) + \sum_{k \in N_i \backslash j} T(\phi_{ik}) \overline{X}(d) \Big) \\ &= \Big(\frac{1}{d_i} \sum_{j \in N_i} T(\phi_{ij}) \Big) \big(1 + (d_i - 1) \overline{X}(d) \big). \end{aligned}$$

Consider an individual stakes profile in which player i's stakes on each link are $\phi'_i = \frac{1}{d_i} \sum_{j \in N_i} \phi_{ij}$. Jensen's inequality implies that

$$\phi'_i + \frac{d_i \lambda}{r} \phi'_i \geq T(\phi'_i) \big(1 + (d_i - 1) \overline{X}(d) \big).$$

and therefore,

$$\frac{T(\phi'_i)}{\phi'_i} \leq \frac{1 + d_i \frac{\lambda}{r}}{1 + (d_i - 1) \overline{X}(d)}.$$

Our aim is to show that $\phi'_i \leq \overline{\phi}(d)$. If $d_i = d$, then this follows from Assumption 1 since the RHS corresponds to $T\big(\overline{\phi}(d)\big)/\overline{\phi}(d)$. Otherwise, if $d_i < d$ it suffices to show that

$$\frac{1 + d_i \frac{\lambda}{r}}{1 + (d_i - 1)\overline{X}(d)} < \frac{1 + d\frac{\lambda}{r}}{1 + (d - 1)\overline{X}(d)},$$

which is equivalent to $\overline{X}(d) < \frac{\lambda}{r+\lambda}$. Since this was already established in Proposition 2, it follows that $\phi_i' \leq \overline{\phi}(d)$. Since player i's payoff in equilibrium σ is $\frac{d_i \lambda}{r}\phi_i'$, it follows from $d_i \leq d$ and $\phi' \leq \overline{\phi}(d)$ that her payoff in equilibrium σ is less than $\overline{u}(d)$—that of the binding contagion equilibrium on $\overline{G}(d)$. Observe that if the component that contains player i is not $\overline{G}(d)$, then the first inequality in (12) is strict, which translates into every subsequent inequality in this line of reasoning being strict.

Now consider an equilibrium profile σ such that shirking occurs on the equilibrium path with strictly positive probability, and assume that the stage game satisfies strategic complementarity. We use a series of arguments analogous to Lemmas 6-8 to argue that the binding contagion equilibrium on $\overline{G}(d)$ has the greatest average payoff. Using the notation of Lemma 6, let $U(\sigma)$ be the total utility, which can be written in different forms:

$$U(\sigma) \equiv \frac{\lambda}{r}\sum_{i\in N}\sum_{j\in N_i} u_{ij} = \frac{\lambda}{r}\sum_{\{ij\}\in G}(u_{ij} + u_{ji}) = \sum_{i\in N}\int_0^\infty e^{-rt}e^{-d_i\lambda t}\lambda\sum_{j\in N_i}\Big(u_{ij} + \frac{\lambda}{r}\sum_{k\in N_i} u_{ik}\Big)\,dt.$$

Applying Lemma 6 and Theorem 2 implies that for every player i and link $\{ij\}$,

$$w_{ij} + \sum_{k\in N_i\setminus\{j\}} w_{ik}\overline{X}(d) \leq w_{ij} + \sum_{k\in N_i\setminus\{j\}} w_{ik}X_{ijk} \leq u_{ij} + \frac{\lambda}{r}\sum_{k\in N_i} u_{ik}. \tag{13}$$

It follows that for every i,

$$\begin{aligned}
\int_0^\infty e^{-rt}e^{-d_i\lambda t}\lambda\sum_{j\in N_i}\Big(u_{ij} + \frac{\lambda}{r}\sum_{k\in N_i} u_{ik}\Big)\,dt &\geq \int_0^\infty e^{-rt}e^{-d_i\lambda t}\lambda\sum_{j\in N_i}\Big(w_{ij} + \sum_{k\in N_i\setminus\{j\}}\overline{X}(d)w_{ik}\Big)\,dt \\
&= \frac{\lambda}{r + d_i\lambda}\big(1 + (d_i - 1)\overline{X}(d)\big)\sum_{j\in N_i} w_{ij}, \\
&\geq \frac{\lambda}{r + d\lambda}\big(1 + (d - 1)\overline{X}(d)\big)\sum_{j\in N_i} w_{ij},
\end{aligned}$$

where the first inequality follows from (13) and the second inequality follows from $\overline{X}(d) < \frac{\lambda}{r+\lambda}$. Define

$$\overline{W}(\sigma) = \frac{\lambda}{r + d\lambda}\big(1 + (d - 1)\overline{X}(d)\big)\sum_{\{ij\}\in G}(w_{ij} + w_{ji});$$

it follows that $U(\sigma) \geq \overline{W}(\sigma)$. Analogous to Lemma 7, there is a mixed strategy mutual effort profile $\tilde{\sigma}$ such that $U(\tilde{\sigma}) \geq U(\sigma) \geq \overline{W}(\sigma) \geq \overline{W}(\tilde{\sigma})$: construct $\tilde{\sigma}$ in exactly the same way as in the proof of Lemma 7. The inequality follows from the analogous observation that for every link $\{ij\}$,

$$\tilde{u}_{ij} + \tilde{u}_{ji} \geq u_{ij} + u_{ji}, \text{ and } \tilde{w}_{ij} + \tilde{w}_{ji} \geq w_{ij} + w_{ji}.$$

Now we proceed to apply the argument of Lemma 8: transforming the mixed-strategy mutual effort profile $\tilde{\sigma}$ to the pure-strategy mutual effort profile $\overline{\sigma}$ preserves the relationship: naturally, $\tilde{u}_{ij} = \overline{u}_{ij}$ and $\overline{w}_{ij} \leq \tilde{w}_{ij}$, and so $U(\overline{\sigma}) \geq \overline{W}(\overline{\sigma})$. Consider the average stakes $\phi^* = \sum_{\{ij\}\in G} \overline{\phi}_{ij}/|G|$, and the contagion profile σ^* in which those are the stakes in each link. By construction

$$U(\sigma) \leq U(\sigma^*) = |G| \left(\frac{\lambda}{r} \phi^* \right).$$

Applying an argument virtually identical to that of Lemma 8 yields $\overline{W}(\sigma^*) \leq \overline{W}(\overline{\sigma})$, which in turn implies that $U(\sigma^*) \geq \overline{W}(\sigma^*)$. Therefore, it follows that $\phi^* \leq \overline{\phi}(d)$ for exactly the same reason as in Lemma 8. Summarizing, this implies that the average utility

$$\frac{U(\sigma)}{n} \leq \frac{|G|}{n} \left(\frac{\lambda}{r} \overline{\phi}(d) \right) \leq \frac{d\lambda}{r} \overline{\phi}(d) = \overline{u}(d), \tag{14}$$

in which the second inequality follows from $|G|/n \leq d$. Notice that if G includes a component that is not $\overline{G}(d)$, the first inequality in (13) is strict for some player i and link $\{ij\}$, and thus the first inequality in (14) is also strict. □

Proof of Theorem 4 on p. 25. Consider a non-empty incomplete network G in which some player i obtains interaction payoff u_i in a mutual effort equilibrium, and $u_i \geq c(d_i)$. From the argument in Theorem 3, it follows that

$$\frac{u_i}{d_i} < \frac{\overline{u}(n-1)}{n-1},$$

and, because linking costs are concave,

$$\frac{c(d_i)}{d_i} \geq \frac{c(n-1)}{n-1}.$$

Combining these two inequalities and multiplying by $n-1$ yields

$$\overline{u}(n-1) - c(n-1) > \left(\frac{n-1}{d_i} \right) (u_i - c(d_i)) \geq u_i - c(d_i).$$

Now consider an equilibrium in which shirking occurs on the equilibrium path. From the argument in Theorem 3, it follows that the average interaction payoff is strictly less than $\overline{u}(n-1)$. Because linking costs are concave, it also follows that the average linking cost is at least $c(n-1)/(n-1)$, and therefore, it follows as above that the net value is strictly less than $\overline{u}(n-1) - c(n-1)$.

□

Proof of Theorem 5 on p. 25. First restrict attention to mutual effort equilibria. Suppose towards a contradiction that there exists network G in which every player obtains at least $\overline{u}(d^*) - c(d^*)$ and at least one obtains a strictly higher payoff. G must have at least one component in which there exists a player who obtains a payoff strictly exceeding $\overline{u}(d^*) - c(d^*)$. By our reasoning above, this component cannot be a clique. Consider a player, say player i, who has the highest degree, d' in this component: by Theorem 3, her interaction payoff is strictly less than $\overline{u}(d')$, which is what she would obtain in $\overline{G}\left(d'\right)$. Therefore, her total payoff is strictly less than $\overline{u}(d') - c(d')$, and therefore strictly less than $\overline{u}(d^*) - c(d^*)$. □

Appendix B Supplementary Appendix

B.1 Additional Proofs for Section 4

We first prove the analogue of Lemma 1 of Ellison (1994):

Lemma 9. *For every non-empty set of players $M \subseteq N$ with $i \in M$,*

$$\pi_i(M\backslash\{j\}) - \pi_i(M \cup \{j\}) \leq \pi_i(\{i\}) - \pi_i(\{i,j\}). \tag{15}$$

Proof. We establish this claim for every generic sequence of link recognitions (in which no two links meet simultaneously) and then take expectations over them. Let ξ be a generic sequence of link recognitions that take place in $[0,\infty)$; let $(\tau_z)_{z=1}^{\infty}$ be an ordered list of link recognition times and $(l_z)_{z=1}^{\infty}$ be the list of links in their order of recognition.

Fix a player i and suppose that M_0 is the set of players who are contagious at a time normalized to zero. If the subsequent set of link recognitions follows ξ, then the set of contagious players at time τ_z is

$$C_z(M_0,\xi) \equiv \begin{cases} M_o & \text{if } z = 0, \\ C_{z-1}(M_0,\xi) & \text{if } z > 0 \text{ and either } l_z \subseteq C_{z-1}(M_0,\xi) \text{ or } l_z \subseteq N\backslash C_{z-1}(M_0,\xi), \\ C_{z-1}(M_0,\xi) \cup l_z & \text{otherwise.} \end{cases} \tag{16}$$

When two players who are both cooperative or both contagious meet, no player changes phase; it is only when a contagious player meets a cooperative player that the latter also becomes contagious. Define $\pi_i\left(M_0|\xi\right)$ to be the equilibrium continuation value of player i when players in M_0 (including player i)

are in the contagion phase at time zero, the realization of recognition times is $\{\tau_z\}_{z=1}^{\infty}$. By calculation,

$$\begin{aligned}
&\pi_i\left(M_0|\xi\right) - \pi_i\left(M_0 \cup \{j\}|\xi\right) \\
&\quad= \sum_{z=1}^{\infty} e^{-r\tau_z} \sum_{k\in N_i} T\left(\phi_{ik}\right) \mathbb{1}\Big(\{k\} = l_z \cap \Big(C_z(M_0 \cup \{j\}, \xi) \backslash C_z(M_0, \xi)\Big)\Big) \\
&\quad\leq \sum_{z=1}^{\infty} e^{-r\tau_z} \sum_{k\in N_i} T\left(\phi_{ik}\right) \mathbb{1}\Big(\{i,k\} = l_z \cap \Big(C_z(\{i,j\}, \xi) \backslash C_z(\{i\}, \xi)\Big)\Big) \\
&\quad= \pi_i\left(\{i\}|\xi\right) - \pi_i\left(\{i,j\}|\xi\right)
\end{aligned} \tag{17}$$

where $\mathbb{1}$ is the indicator function. The weak inequality follows from

$$C_z\left(M_0 \cup \{j\}, \xi\right) \backslash C_z\left(M_0, \xi\right) \subseteq C_z\left(\{i,j\}, \xi\right) \backslash C_z\left(\{i\}, \xi\right). \tag{18}$$

Since (17) holds for every generic ξ, taking the expectation over ξ yields (15). □

Proof of Proposition 1 on p. 14. Consider a binding contagion profile with collective stakes profile Φ. Since this profile satisfies cooperation phase incentive constraints by construction, it suffices to establish contagion phase incentives.

Clearly, if $\hat{\phi}_{ij}$ or $\hat{\phi}_{ji}$ differ from ϕ_{ij}, or if player i or j has shirked with each other in a previous interaction, player i has no incentive to work since player j will shirk. Instead, suppose as in $\text{IC}_{ij}^{\text{Cont}}$ on p. 13 that neither of these have occurred and that player i knows that players M are in the contagion phase. Let $\pi_i(\varnothing)$ represent player i's equilibrium continuation value. It follows from binding $\text{IC}_{ij}^{\text{Coop}}$ on p. 13 that

$$T(\phi_{ij}) - \phi_{ij} = \pi_i(\varnothing) - \pi_i(\{i,j\}) = \pi_i(\{i\}) - \pi_i(\{i,j\}) \geq \pi_i(M) - \pi_i(M \cup \{j\}), \tag{19}$$

in which the first equality is a binding $\text{IC}_{ij}^{\text{Coop}}$ being re-written, the second inequality follows from all of player i's cooperation phase incentive constraints binding and so $\pi_i(\varnothing) = \pi_i(\{i\})$, and the inequality follows from Lemma 9. Adding $\phi_{ij} + \pi_i\left(M \cup \{j\}\right)$ to each side yields $\text{IC}_{ij}^{\text{Cont}}$. □

B.2 Additional Proofs for Section 5

Proof of Proposition 2 on p. 21. First, suppose there is no cycle that includes both ij and ik. Then player k becomes contagious in a contagion profile only through meeting player i on link ik, which occurs at rate λ. Hence $x_{ijk}(t) = e^{\lambda t}\lambda$ so $X_{ijk} = \int_0^{\infty} e^{-r\tau} e^{\lambda t} \lambda \, d\tau = \frac{\lambda}{\lambda + r}$. Now suppose there is a cycle including both ij and ik. Then for all $\tau > 0$ there is a path from j to k that does not include i, and there is strictly positive probability that contagion has spread from j to k along this path, independent of whether link $\{ik\}$ has been recognized. Therefore $x_{ijk}(\tau) < e^{\lambda\tau}\lambda$, so $X_{ijk} < \frac{\lambda}{\lambda+r}$. □

Proof of Proposition 3 on p. 22. We use various properties of a clique to simplify the algorithm. Since a clique is symmetric, Lemma 2 holds, and so there exists $\overline{X}(d)$ such that if $\{ij\}$ and $\{ik\}$ are both in G, then $X_{ijk} = \overline{X}(d)$. The term $(d-1)\overline{X}(d)$ corresponds to player i being able to shirk on neighbors other than player j before they are infected by others; we can derive a closed-form expression for this term by recursing on the number of infected neighbors. Suppose that there are $m-1$ of player i's neighbors that are currently infected: then there are $m(d+1-m)$ links by which the contagion spreads to an uninfected neighbor of player i; of these $d+1-m$ correspond to links of player i. Thus, it follows that $(d-1)\overline{X}(d) = \chi(2)$ in which for $m \geq 2$,

$$\begin{aligned}\chi(m) &= \int_0^\infty e^{-rt} e^{-\lambda m(d-m+1)t} \lambda m(d-m+1) \left(\frac{1}{m} + \chi(m+1) \right) dt \\ &= \frac{\lambda m(d-m+1)}{r + \lambda m(d-m+1)} \left(\frac{1}{m} + \chi(m+1) \right).\end{aligned}$$

Setting $\chi(d+1) = 0$ generates the expression in the statement. □

B.3 Proofs for Section 6

Proof of Proposition 4 on p. 28. The argument proceeds by showing that the recursive formula for $\vec{X}_{ijk}$ is identical to that for X_{ijk}. Let $\vec{ij}$ represent the interaction in which player i is called upon to do a favor for player j: given a network G, one can then consider the directed network $\vec{G}$ that consists of all the possible directed links. A directed path $\vec{\zeta}$ from i to k is a series of directed links that begin at node i and end at node k. We define the analogues of $s_k(M)$ and $\tilde{s}_k(M)$:

$$\begin{aligned}\vec{s}_k(M) &= \left\{ \vec{\zeta} : \vec{\zeta}(1) \in M,\ \vec{\zeta}(Z_{\vec{\zeta}}) = k,\ \text{and } \vec{\zeta}(z) \notin M \text{ for every } z > 1 \right\}. \\ \vec{\tilde{s}}_k(M) &= \left\{ \{i'j'\} \in \vec{G} : i' \in M, \vec{\zeta}(2) = j' \text{ for some } \vec{\zeta} \in \vec{s}_k(M) \right\}.\end{aligned}$$

In the favor exchange environment, contagion spreads at a rate of $|\vec{\tilde{s}}_k(M)|\lambda$. Thus, $\vec{X}_{ijk}$ can be computed recursively as $\vec{X}_{ijk} = \vec{\chi}_{ik}(\{i,j\})$, and for every M that contains $\{i,j\}$ but not k,

$$\begin{aligned}\vec{\chi}_{ij}(M) &= \int_0^\infty e^{-rt} e^{-|\vec{\tilde{s}}_k(M)|\lambda t} \left(\lambda + \sum_{i'j' \in \vec{\tilde{s}}_k(M)} \lambda \vec{\chi}_{ij}(M \cup \{i', j'\}) \right) dt \\ &= \frac{\lambda}{r + |\vec{\tilde{s}}_k(M)|\lambda} \left(1 + \sum_{i'j' \in \vec{\tilde{s}}_k(M)} \overrightarrow{\chi}_{ij}(M \cup \{i', j'\}) \right).\end{aligned}$$

For every M that contains k, $\vec{\chi}_{ij}(M) = 0$. Therefore, $\vec{X}_{ijk} = X_{ijk}$. □

Proof of Proposition 5 on p. 28. Straightforward algebra verifies that the stakes ϕ^{F} are the stakes of a symmetric binding contagion equilibrium, and exists because C is strictly convex and satisfies Inada conditions. We sketch below how to apply previous arguments to verify that it has a higher average utility than every distinct equilibrium. Let $\Phi_i \equiv (\phi_{ij}, \phi_{ji})_{\{ij\} \in G}$ be the individual profile of favors that player i performs or a neighbor performs for her in a pure strategy profile, and let Φ be the collective profile of favors $(\Phi_1, \ldots, \Phi_n)$.

Let $\vec{\Psi}_{ij}(G)$ be the set of player i's favor profiles that satisfy $\vec{IC}_{ij}$ on 28, and $\vec{\Psi}(G)$ be the set of collective favor profiles that satisfy all cooperation phase incentives.

Because C is strictly convex, it suffices to consider pure strategy favor profiles in $\vec{\Psi}(G)$: for every mixed strategy favor equilibrium, the pure strategy profile in which each player i performs the averaged favor for player j generates identical payoffs with a lower temptation to not perform the favor. The total utility for a pure strategy favor profile is

$$U(\sigma) = \frac{\lambda}{r} \sum_{\{ij\}\in G} \left(\phi_{ij} - C(\phi_{ij}) + \phi_{ji} - C(\phi_{ji})\right).$$

Consider the total deviation utility for a pure strategy favor profile:

$$\begin{aligned} W(\sigma) &= \sum_{i\in N} \int_0^\infty e^{-rt} e^{-d\lambda t} \lambda \sum_{j\in N_i} \Big(C(\phi_{ij}) + \sum_{k\in N_i\setminus\{j\}} X\phi_{ik}\Big)\, dt. \\ &= \frac{\lambda}{r + d\lambda} \sum_{\{ij\}\in G} \left(C(\phi_{ij}) + (d-1)X\phi_{ij} + C(\phi_{ji}) + (d-1)X\phi_{ji}\right). \end{aligned}$$

Because Φ is in $\vec{\Psi}(G)$, $W(\sigma) \le U(\sigma)$. Consider a symmetric strategy profile $\tilde{\sigma}$ in which for every $\{ij\} \in G$,

$$\tilde{\phi}_{ij} = \tilde{\phi}_{ji} = \sum_{\{i'j'\}\in G} \frac{\phi_{i'j'} + \phi_{j'i'}}{2|G|}.$$

Because C is strictly convex, it follows that $W(\tilde{\sigma}) < W(\sigma) \le U(\sigma) < U(\tilde{\sigma})$. By algebra, the implication of $W(\tilde{\sigma}) < U(\tilde{\sigma})$ is that $\tilde{\phi}_{ij} < \phi^{\mathrm{F}}$. If $\phi^{\mathrm{F}} < \phi^*$, it follows that $\phi - C(\phi)$ is strictly increasing in the interval $[\tilde{\phi}_{ij}, \phi^{\mathrm{F}}]$, and therefore, the value is higher in the binding contagion equilibrium than in any other equilibrium. ☐

Proof of Proposition 7 on p. 30. To argue that $\vec{q}$ and $\vec{p}$ are part of a binding contagion equilibrium, we rewrite $\vec{IC}_{bs}$ and $\vec{IC}_{sb}$ for the symmetric case in which $q_{bs} = q$ and $p_{bs} = p$ for every bs in $\vec{G}$:

$$\begin{aligned} q + (d_B - 1)qX_B &\le q - p + \frac{d_B\lambda}{r}(q - p), \\ p + (d_S - 1)pX_S &\le p - C(q) + \frac{d_S\lambda}{r}\left(p - C(q)\right). \end{aligned}$$

Setting both inequalities to bind generates $\vec{q}$ and $\vec{p}$. Now we argue that this equilibrium has the greatest average trade. For a buyer b, let $\Phi_b \equiv (q_{bs}, p_{bs})_{s\in N_b}$ be her profile of trades, and for a seller s, let $\Phi_s \equiv (q_{bs}, p_{bs})_{b\in N_s}$ be his profile of trades. Let $\Psi_{bs}(\vec{G})$ be the set of trade profiles Φ_b that satisfies the buyer's incentive constraint, $\vec{IC}_{bs}$; notice that because of the linearity, this is a convex set. Similarly, let $\Psi_{sb}(\vec{G})$ be the set of trade profiles that satisfy the seller's incentive constraint, $\vec{IC}_{sb}$. The strict convexity of $\Psi_{sb}(\vec{G})$ follows from the strict convexity of C. This strict convexity is sufficient to argue the analogue of Lemma 5: for every mixed strategy equilibrium, there exists a payoff equivalent pure strategy profile. Notice that the incentive constraints can be re-written as

$$p_{bs} + \sum_{s'\in N_b\setminus\{s\}} q_{bs'}X_B \le \frac{\lambda}{r}\sum_{s'\in N_b}(q_{bs'} - p_{bs'}),$$

$$C(q_{bs}) + \sum_{b'\in N_s\setminus\{b\}} p_{bs'}X_S \le \frac{\lambda}{r}\sum_{b'\in N_s}(p_{b's} - C(q_{b's})).$$

The total utility for a pure strategy profile is

$$U(\sigma) = \frac{\lambda}{r}\sum_{bs\in\vec{G}}(q_{bs} - C(q_{bs})).$$

Define the total deviation utility to be

$$W(\sigma) = \frac{\lambda}{r + d_B\lambda}\sum_{b\in\mathcal{N}^B}\sum_{s\in N_b}\left(p_{bs} + \sum_{s'\in N_b\setminus s} q_{bs'}X_B\right) + \frac{\lambda}{r + d_S\lambda}\sum_{s\in\mathcal{N}^S}\sum_{b\in N_s}\left(C(q_{bs}) + \sum_{b'\in N_s\setminus b} p_{bs'}X_S\right).$$

It follows from the incentive constraints that $W(\sigma) \le U(\sigma)$. Consider a symmetric strategy profile $\tilde{\sigma}$ in which for every $bs \in \vec{G}$,

$$\tilde{q}_{bs} = \sum_{bs'\in\vec{G}} \frac{q_{bs'}}{|G|},$$

$$\tilde{p}_{bs} = \sum_{bs'\in\vec{G}} \frac{p_{bs'}}{|G|}.$$

Because C is strictly convex, it follows that $W(\tilde{\sigma}) < W(\sigma) \le U(\sigma) < U(\tilde{\sigma})$. By algebra, the implication of $W(\tilde{\sigma}) < U(\tilde{\sigma})$ is that $\tilde{q}_{bs} < \vec{q}$. Therefore, every distinct equilibrium generates lower average trade than the binding contagion equilibrium. If $\vec{q} < q^*$, it follows that the value is greater in the binding contagion equilibrium than in any other equilibrium. □

Proof of Proposition 8 on p. 31. For the purposes of viscosity, and so as to use the prior notation, it is helpful to transform the bipartite networked market $\vec{G}$ into the undirected network G such that $\{bs\} \in G$ if and only if $bs \in \vec{G}$.

We compare $X_{bss'}$ to $\overline{X}_B(d_B, d_S)$. Let G_{-i} be the network that results from deleting all of buyer b's links. If there is no path between sellers s and s' in G_{-i}, then b, s, s' are not part of a cycle and the result is implied by Proposition 2. So suppose that there is at least one path between sellers s and s' in G_{-i}. As in the proof of Theorem 2, let ζ be a generic path from seller s to seller s' in G_{-i}, and let S be the set of all such paths. We consider a partition of S, such that two paths are in the same partition element if and only if $\zeta(2) = \zeta'(2)$; in other words, the second player coincides. Since seller s has at most $(d_S - 1)$ neighbors in G_{-i}, there are at most $(d_S - 1)^2$ partition elements. We denote a partition element by S_u if $\zeta(2) = u$ for every $\zeta \in S_u$.

Now, consider an arbitrary buyer $\overline{b}$ and pair of sellers $\overline{s}$ and $\overline{s}'$ in the bipartite clique $\overline{G}(d_b, d_S)$. Let $\overline{N}_m$

be player m's neighborhood in the bipartite clique. Consider an injective function $g : N_s \to \overline{N}_{\overline{s}}$ such that $\overline{b} = g(b)$. Suppose that at time 0, players s and $\overline{s}$ are contagious in networks G and $\overline{G}(d_b, d_S)$ respectively, and contagion spreads whenever the link between an infected player and an uninfected player is recognized. We couple the processes on G and $\overline{G}(d_b, d_S)$ as follows. Given a sequence of link recognitions $(l_z)_{z=1}^{\infty}$ on G,

1. Seller $\overline{s}$ meets buyer $g(u) \in \overline{N}_{\overline{s}}$ whenever seller s meets buyer $u \in N_s$.
2. Buyer $\overline{b}$ meets seller $\overline{s}' \in \overline{N}_{\overline{b}}$ whenever buyer b meets seller $s' \in N_j$.
3. For any path $\zeta \in S_u$ contained in $(l_z)_{z=1}^{\infty}$ for which $Z_\zeta \geq 3$, seller $\overline{s}'$ meets buyer $g(u) \in \overline{N}_{\overline{s}'}$ when seller s' meets buyer $\zeta(Z_\zeta - 1)$ in ζ.

Consider any sequence of link recognitions $(l_z)_{z=1}^{\infty}$ on G such that at time 0 the set of contagious players is $\{b, s\}$, and such that seller s' becomes contagious before link $\{bs'\}$ is recognized. Such a sequence contains some path $\zeta \in S$ that is completed (all its links are recognized) before link $\{bs'\}$ is recognized. It immediately follows that in the coupled link recognition process on $\overline{G}(d_b, d_S)$, seller $\overline{s}'$ must become contagious before meeting buyer $\overline{b}$. Thus, for any sequence of link recognitions $(l_z)_{z=1}^{\infty}$ on G in which link $\{bs'\}$ meets only after seller s' has already infected by a path from seller s, there is a coupled sequence of link realizations on $\overline{G}(d_b, d_S)$ in which the analogous event occurs. Note that the density of link realizations on $\overline{G}(d_b, d_S)$ is at least as high as that in G.

The other form of link realization by which seller s' may be infected before $\{bs'\}$ is recognized is if buyer b meets another seller $m \in N_b \backslash \{s, s'\}$, and then a path from seller m infects seller s' before link $\{bs'\}$ is recognized. The proof applies *mutatis mutandi* with seller m taking the stead of seller s. Since buyer b has at most d_B neighbors in network G, the density of these link realizations is no greater in network G than in $\overline{G}(d_b, d_S)$.

Finally, we argue that the inequality is strict if the component of G that contains buyer b and sellers s and s' is anything other than $\overline{G}(d_b, d_S)$. If $|N_s \cap N_{s'}| < d_S$ then $\overline{G}(d)$ has additional paths by which seller $\overline{s}'$ may become infected before link $\{\overline{b}\overline{s}'\}$ is recognized.

An analogous argument establishes the same claim for $X_{ssb'}$. Finally, having established the claim for viscosity, the latter statements follow from application of the argument in Theorem 3.

□

B.4 Proofs for Section 7

We define a generalized contagion profile: given a network G, link $\{ij\}$ in G, and a partition $\mathcal{P}$, let $P(ij)$ be the member of the partition that contains $\{ij\}$. We define $N(P)$ to be the set of players who have at least one link in P, a generic member of the partition. For a history h, we let the history restricted to member P_k of the partition, denoted by $h_{|P_k}$, be the list of all interactions in h in that involve only links in P_k. We let $\sigma_{i|P}$ be a strategy for player i of the game in which the network is P_k and the set of players correspond to $N(P)$.

Definition 9. *A strategy profile σ is a* ***generalized contagion profile*** *if there exists some symmetric decomposition $\mathcal{P}$ such that*

1. *for every link* $\{ij\}$ *in* G*, for every* h *and* h' *in* H_{ij} *in which* $h_{|P(ij)} = h'_{|P(ij)}$, $\sigma_i(h) = \sigma_i(h') = \sigma_{i|P(ij)}(h'_{|P(ij)})$ *for some* $\sigma_{i|P(ij)}$.
2. *for every member* P *of* $\mathcal{P}$*, the profile* $(\sigma_{i|P})_{i \in N(P)}$ *is a contagion profile of the game with* $N(P)$ *players and the network* P.

Proof of Proposition 9 on p. 36. For every G, consider the partition in which each member is a singleton. For each member, bilateral enforcement at stakes ϕ^{B} coincides with contagion and is an equilibrium; therefore, there exists a generalized contagion profile in which the stakes on each link is ϕ^{B}. If G contains a cycle, from i to i, let P' be the set of links on that cycle. Since P' is a circle, consider the non-trivial binding contagion profile on P' (which exists by Theorem 1); by earlier, there exists a generalized contagion equilibrium with strictly positive stakes for the network $G \backslash P'$, and so combining the two generates a generalized contagion equilibrium for the network G. ☐