Seeking Relationship Support: Strategic network formation and robust cooperation

David A. Miller and Xu Tan*

May 21, 2021

Draft in Progress

Abstract

We study cooperation on evolving social networks with private monitoring and communication. For arbitrary networks, we construct a class of *multilateral restitution* equilibria that attain high cooperation on all *supported* links, i.e., all links that are in triangles. These equilibria are both *robust*—preserving high cooperation between innocent players on and off the equilibrium path—and *local*—invariant to players' beliefs about the network outside their local neighborhoods. Guilty players are not ostracized; instead they remain involved to sustain cooperation on the network and while paying restitution by exerting high effort for their innocent partners. When new players arrive, they strategically form links that in aggregate lead to realistic "small worlds" network properties, including high support but relatively low clustering.

*Miller: University of Michigan (e-mail: econdm@umich.edu). Tan: University of Washington (e-mail: tanxu@uw.edu). We thank Lisa Bernstein and Gabriel Carroll for helpful suggestions, and participants at the Third Annual Conference on Network Science and Economics, National University of Singapore, Chinese University of Hong Kong, Western Ontario, University of Zurich, the 2018 Barcelona Summer Forum, Norwegian School of Economics, European Meeting of the Econometric Society at Cologne, Workshop on Relational Contracts at University of Chicago, Michigan Electrical Engineering & Computer Science, and the ThReD conference for helpful feedback.

1 Introduction

Consider a growing social network in which each link is an ongoing productive relationship. Within a relationship, the partners benefit from each others' efforts, but each has an individual temptation to shirk. Since these relationships are organized within a social network, there is scope for multilateral enforcement: if a player shirks on one partner, and that partner informs other players in the network, then eventually the deviator may be punished by multiple partners. Because the multilateral punishment from the community is harsher than a bilateral punishment, it can help sustain a higher degree of cooperation.

Since the nature of multilateral enforcement is shaped by the network, players may strategically choose which relationships to form, anticipating future cooperation. When new players arrive and find opportunities to connect, they may favor relationships that are well suited to multilateral enforcement. Over time, players' choices, aggregated, shape the long-run evolution of the network and its global properties.

In this paper we introduce a class of *multilateral restitution* equilibria that implement multilateral enforcement with private monitoring on an arbitrary network, with two important properties. First, it is *local*: each player's strategy depends only on the structure of her local neighborhood and the events that occur within it. Second, it is *robust*: when a punishment arises, its ramifications do not cascade through the network; instead the effect is contained within the deviator's neighborhood.

We apply multilateral restitution equilibria as a tool to study the evolution of the network when newly arriving players get random opportunities to meet "strangers" and become "friends", and then get opportunities to meet friends of friends (building on Jackson and Rogers 2007). Their strategic choices, anticipating cooperation equilibrium, result in a network with "small worlds" properties that are conducive to multilateral enforcement. In particular, the global network features high "support", but low "clustering". High *support* means that each pair of linked partners is likely to have at least one mutual neighbor (Jackson, Rodriguez-Barraquer, and Tan 2012); low *clustering* means that two players who share a mutual neighbor are unlikely to be linked to each other. Intuitively, under multilateral restitution a supported link enjoys higher cooperation, so players seek support for their relationships. In contrast, players do not have an incentive to seek clustering beyond what is needed to obtain support. In all, our results contribute to a unified understanding of how multilateral enforcement can both influence and be influenced by the evolution of social networks.

Our focus on strategies that are *local* contrasts with much of the theoretical literature on multilateral enforcement, which has assumed that the entire network is commonly known. In reality, the players may observe only a *local subset* of the entire network. For example, multilateral enforcement helps support joint production and risk sharing in rural villages of developing countries. Breza, Chandrasekhar, and Tahbaz-Salehi (2016) surveyed villages in Karnataka in southwest India, and found that the respondents' knowledge about the social network is highly localized.[1]

We focus on strategies that are *robust* since infractions arise and are punished in reality, and yet multilateral enforcement persists rather than breaking down. In the theoretical literature, early studies of multilateral enforcement either took truthful communication for granted (simplifying the task of making enforcement robust) or focused on "contagion" strategies in which a single infraction leads to widespread breakdown of social cooperation. As social networks and cooperative arrangements are durable in the real world, they must not rely on widespread breakdown of cooperation as punishment. More recent literature (particularly Ali and Miller 2016) has identified truthful communication as the key difficulty in constructing equilibria in which innocent players continue to cooperate. Our multilateral restitution equilibria guarantee that innocent players will not be harmed by truthfully conveying information about who has deviated.

Recent work on the theory of cooperation in social networks has largely taken a normative view on how to compare networks. Broadly speaking, the literature finds that denser networks yield more cooperation.[2] In contrast, the descriptive empirical side of the literature has found that social networks are not particularly dense. Compared to the dense networks that are optimal under many theories, real social networks have similarly high support, but dramatically lower clustering. As a result, for a given population and a given average degree, real social networks are much more "expansive" (Ambrus, Möbius, and Szeidl 2014) than is optimal under these theories. In this paper we take a positive theoretical approach, to study which networks players may form through their strategic

[1]For instance, 46% of respondents are not able to guess whether there is a link between a given pair of individuals, and conditional on making a guess, the accuracy is only 33%. Their findings are consistent with other surveys on people' knowledge of their networks. For example, Krackhardt (1990) finds that the accuracy of knowing other people's connections is 15%–48% within a startup firm with 36 employees; Casciaro (1998) finds the accuracy is around 45% in a research center of 25 people.

[2]See, for example, Ali and Miller (2013); Ambrus, Möbius, and Szeidl (2014); Jackson, Rodriguez-Barraquer, and Tan (2012); Wolitzky (2013); Karlan, Möbius, Rosenblat, and Szeidl (2009).

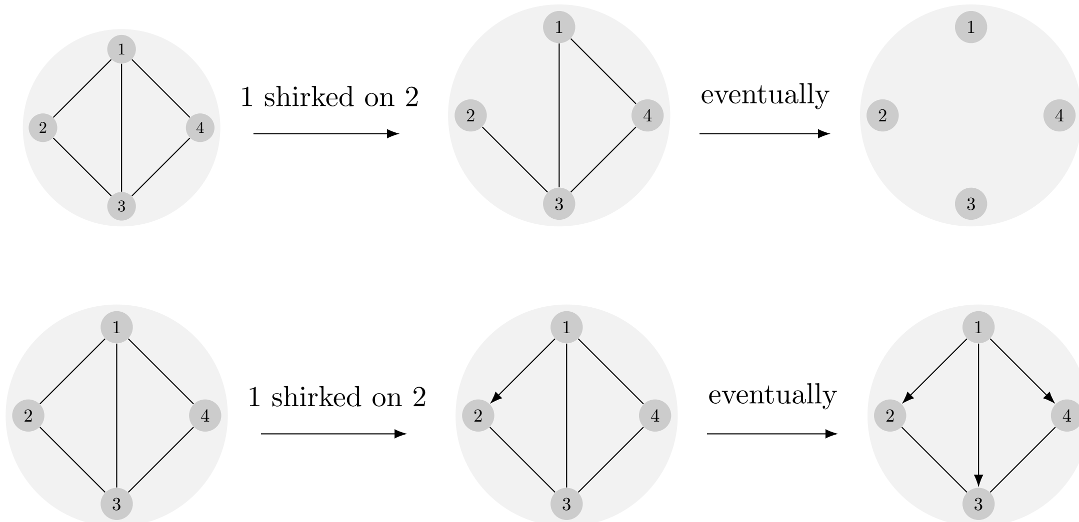

Figure 1: (Top) Network enforcement by ostracism; (Bottom) Network enforcement by multilateral restitution.

interactions. We find that anticipating cooperation, players seek out support for their relationships, generating networks with high support but low clustering, along with other "small world" characteristics like connectedness, small average distances, and realistic degree distributions.

Multilateral restitution The challenge in constructing robust equilibria with multilateral enforcement on an arbitrary network arises from interactions among overlapping neighborhoods. After player 1 deviates by shirking on player 2, for multilateral enforcement player 2 should inform player 3 of the deviation, and then both players 2 and 3 should punish player 1. However, these punishments may interfere with cooperation between players 2 and 3, since if their payoffs along links $1:2$ and $1:3$ are depressed, they have less to lose by shirking on link $2:3$. If both of them are also connected to player 4, then shirking may start to cascade through the network as a contagion and eventually all partnerships break down, as illustrated in the top panel of Figure 1.

Multilateral restitution strategies solve this problem by preserving the payoffs of innocent players following a deviation. The deviator is not simply ostracized from the community; instead he remains active in the network to help sustain cooperation among his neighbors. Specifically, after player 1 deviates on player 2, as a *restitution* punishment player 1 must exert high effort when meeting player 2 in the future, while player 2 exerts just enough effort to motivate player 1. At the same time, players 2 and 3 communicate truthfully and continue cooperating at a high level. Player 3 eventually starts punishing player 1 as well—either because player 1 shirks on player 3, or because player 3 learns about player 1's guilt from player 2. The maximal level of cooperation that a triangle can sustain on the equilibrium path is that which makes each player indifferent between working and shirking when facing the threat of maximal punishment by both neighbors.[3] Observe that since innocent players 2 and 3 always expect to get their equilibrium payoffs in their relationships on the $1:2:3$ triangle (it is a surprise for them if player 1 shirks), restitution punishments do not initiate a contagion. Similarly, if players 1, 2, and 3 are all connected to player 4, then 4 eventually also punishes 1 while preserving his cooperation with 2 and 3, as illustrated in the bottom panel of Figure 1. Accordingly, the equilibrium is robust.

In such an equilibrium, player 2's effort level on each link depends only on whether the link belongs to a triangle, so she needs to know only her own neighborhood (her neighbors and the links among them). Moreover, she might be called upon to punish player 1 for shirking on a player who is not in her neighborhood, such as player 4 in Figure 1. But she does not need to learn about player 4's behavior or even existence; she needs only to hear from player 3 that player 1 is guilty. To determine how to behave, each player uses only information about her own neighbors' behavior.[4]

Network formation We introduce a dynamic network formation game, in which newly arriving players strategically form links, anticipating the benefits of multilateral enforcement from the network. We follow the setup of a growing network in Jackson and Rogers (2007). As each new player enters the community, in the first stage of her arrival she randomly meets some "strangers", and will link to each (become "friends") if their idiosyncratic

[3] In contract law, restitution damages are calculated to erase the "unjust enrichment" obtained by the party who breached the contract (see Thompson 1984). Within a triangle, this is precisely the punishment that a deviator faces, since the prospect of punishment makes him indifferent between cooperating and deviating on the equilibrium path. On a denser network, however, the deviator may suffer restitution punishments with each of many partners, akin to a legal remedy of paying a multiple of the restitution damages. Notice that victims are never compensated for being shirked on—restitution punishments operate in our context as deterrence, not for justice.

[4] There are further subtleties that we address in our analysis, such as determining which of a deviator's neighbors needs to punish him, how to work around the fact that a guilty player cannot be relied upon to shirk on a partner who doesn't yet know he is guilty, and what happens if two partners accuse each other.

linking cost is sufficiently low compared to the benefit they expect from cooperating in the absence of relationship support. In her second stage, the new player meets her “friends of friends”; she links to each if their idiosyncratic linking cost is sufficiently low compared to the benefit of cooperating in the presence of a supporting relationship, since every link formed in this second stage is supported.

In the third stage, which is new and motivated by multilateral restitution, the new player seeks support for her unsupported relationships. Specifically, she can revisit her decisions over links she elected not to form in the second stage. This time, she takes into account the externality that is provided—forming a new link brings support to one or two relationships that were unsupported after the second stage. In the first three stages, the arriving player pays all linking costs, and assumes that all her new friends are innocent. In the fourth stage, those new friends should pay her their shares of the linking costs; if they refuse, they become guilty. As a result, the network formation process is also robust, since the link formation behavior of arriving players is the same on and off the equilibrium path.

As the network grows, it tends toward small worlds properties. First, it inherits several properties from the Jackson and Rogers (2007) model, including small average distances, connectedness, and a fat-tailed degree distribution. But because arriving players seek support for their relationships, the network tends toward higher support, but not higher clustering. Specifically, we show that as the number of strangers and friends of friends that arriving players meet increases, the average clustering of the network tends to zero, while the average support tends to a positive limit. Moreover, the support that arriving players specifically seek out for their unsupported relationships (i.e., the support created in the third arrival stage) also tends to a positive limit. Quantitatively, our initial simulations suggest that this “support seeking” can be responsible for a sizable fraction of the the limiting support.

1.1 Related literature

Our model of cooperation builds on the private monitoring, variable effort models with asynchronous interactions, introduced by Ali and Miller (2013, 2016). We allow partners to exchange cheap-talk messages about their past histories, following Lippert and Spagnolo (2011) and Ali and Miller (2016). Similar to Ali and Miller (2021), we assume that partners act sequentially when they meet, enabling the first mover to be punished immediately. Our modeling innovation is to allow partners to endogenously select which of them is the first mover. By choosing who moves first as a function of the history, the players can attain higher equilibrium payoffs than if timing (either sequential or simultaneous) were imposed exogenously. In particular, they can punish the guilty partner by making him move first.

The equilibrium properties on which we focus build on two strands of the prior literature. We seek equilibria that are robust to contagion, so that a deviation does not cause a breakdown of cooperation outside the deviator's neighborhood. This kind of robustness was first formalized by Jackson, Rodriguez-Barraquer, and Tan (2012). We seek equilibria in which players need only local knowledge, as pioneered by Galeotti, Goyal, Jackson, Vega-Redondo, and Yariv (2010), and employed by Nava and Piccione (2014) and Campbell (2014). For both robustness and local knowledge, we introduce stronger notions than used in the prior literature. Our multilateral restitution equilibria, in which guilty players are not excluded, but rather must work hard to satisfy their punishers, build on the asymmetric, bilateral renegotiation proof punishments studied by Ali, Miller, and Yang (2016), which in turn built on the ideas of van Damme (1989).

Substantively, Jackson, Rodriguez-Barraquer, and Tan (2012) and Ali and Miller (2016) have studied questions that closely relate to ours. Jackson, Rodriguez-Barraquer, and Tan model each relationship as a fixed-effort repeated prisoners' dilemma, and assume that monitoring is public. They show that "social quilts"—networks in which cliques are arranged as nodes on trees—are optimal when community enforcement must be robust to contagion and invulnerable to renegotiation. The Ali and Miller (2016) model, like ours, has private monitoring, variable efforts, and communication; they study "ostracism" equilibria in which guilty players are excluded while innocent partners continue to cooperate. While one of their main results is negative (no "permanent ostracism" equilibrium can sustain more cooperation than bilateral enforcement), their other main result shows how to construct a "temporary ostracism" equilibrium on a complete network that supports strictly more cooperation than bilateral enforcement. We complement these works by demonstrating how multilateral enforcement can operate on an arbitrary network, and with agents holding only local knowledge of the network.

Much of the theoretical literature on social network formation has focused on Nash equilibrium or pairwise stable networks in symmetric, deterministic environments. Our network formation game is inspired instead by the stochastic but non-strategic model of Jackson and Rogers (2007), in which new players, arriving over time, are added to the network by first linking to random "strangers" (who become "friends"), and then linking to "friends of friends." At the stochastic limit, the resulting networks display "small worlds" properties. Closer to our network formation model are the stochastic *and* strategic models of Golub and Livne (2010) and Campbell (2014), where players first meet strangers and then meet friends of friends. Both models generate small worlds properties similar (Golub and Livne) or identical (Campbell) to the Jackson and Rogers model. However, Golub and Livne assume that each player makes an ex ante strategic decision regarding how intensively to socialize, and then the entire link formation process follows mechanically from these decisions. Campbell assumes that players cannot choose how many links to form or how to allocate them, other than to either form all of them with strangers or form a fixed fraction of them with friends of friends. Thus in both cases clustering and support are generated by the same mechanism. In contrast, we assume that players make case-by-case decisions over which links to form as the individual opportunities arrive; as a consequence our model contains not only a mechanism that generates both clustering and support, but also a mechanism that primarily generates support without clustering. As for payoffs, players in the Golub and Livne model value only direct connections; clustering and support arise mechanically because friends are valuable. Players in the Campbell model value support in order to signal high patience; once revealed to be patient they no longer need support to sustain cooperation in their relationships.[5]

The networks that form in our model can be interpreted as having both "strong ties" and "weak ties" (Granovetter 1973), where strong ties are links that are supported and exhibit high cooperation, while weak ties are links that are unsupported and exhibit low cooperation. However, in contrast to Granovetter's analysis it is not particularly likely that a player with a strong tie to a particular neighbor will seek to connect with others strongly tied to that neighbor. Instead, it is a player with an initially weak tie to a particular neighbor who is likely to seek connections to that neighbor's neighbors in order to strengthen their tie.

[5]Campbell (2014) features two other aspects that relate to our work. First, Campbell's players can choose the stakes of their relationships. However, these stakes are binary (high or low), and are used only as screening devices, not to take advantage of multilateral enforcement. Second, Campbell's players have only local knowledge of the network, which makes their inference problem of screening types non-degenerate.

2 The repeated interaction game

At first, we focus on how players can cooperate on a fixed network, in a way that is robust and local. We study a population of players $N = \{1, \ldots, n\}$ who are arranged on a network, and interact repeatedly along their network links. The network, G, is a collection of undirected bilateral links. If players i and j are linked in the network, we write $i : j \in G$, or ij for short; as synonym, we call them "partners" or "friends", and call their link a "partnership" or a "friendship". Later, in Section 4, we examine how the prospect of multilateral enforcement influences the formation of the network as new players arrive over time.

Following the social networks literature, we say a path of length k between player i and j is a sequence of distinct players $\{i_0, i_1, \ldots, i_k\}$ such that $i_0 = i$, $i_k = j$, and $i_l : i_{l+1} \in G$ for any $l \in \{0, \ldots, k-1\}$. Let $D(i, j)$ be the distance between player i and j, defined as the length of the shortest path between them. Let $D(i, i) = 0$ and $D(i, j) = \infty$ if i and j are not path-connected. Define player i's *D-neighborhood* as (g_i^D, G_i^D): $g_i^D = \{j : D(i, j) \leq D\}$ and $G_i^D = \{j : k \in G : j, k \in g_i^D\}$. For instance, a player's 1-neighborhood includes herself, all her friends, and all the links among herself and her friends.

The repeated interaction game proceeds over time $t \in [0, \infty)$. Each pair of linked players, $i : j \in G$, meets at random times generated by a Poisson process of rate $\lambda > 0$. Meetings are i.i.d. across links and over time. Whenever partners i and j meet, they play a stage game with three phases: *communication*, *sequencing*, and *effort*, as described below. Players outside the partnership cannot observe when the partners meet or how they behave when they meet. With this private monitoring, communication is crucial to implementing multilateral enforcement.

1. Communication: Each player i and j simultaneously sends a cheap-talk message. Each player's message space is the power set of the intersection of their 1-neighborhoods. (The message that player i sends will interpreted as the set of players in that intersection that player i considers guilty.)

2. Sequencing: Each player i and j simultaneously sends a message indicating which partner should go first. If they agree, the agreed-upon order takes place. If they disagree, then one of them is randomly selected to move first, with equal probabilities.[6]

[6] For the random selection, we assume that each pair has access to public randomization devices whose

3. Effort: Let i be the first mover. Then i and j sequentially choose effort levels x_i, y_j in $[0, \infty)$, where x_i indicates that i is the first mover and y_j indicates that j is the second mover.

Player i's stage game payoff function when partnership ij meets is $b(y_j) - c(x_i)$, where $b(y_j)$ is the benefit from her partner j's effort and $c(x_i)$ is the cost she incurs from her own. Similarly, player j's stage game payoff is $b(x_i) - c(y_j)$. All players share a common discount rate $r > 0$.

We normalize the net value of effort x to $b(x) - c(x) = x$. Our first assumption (following Ali, Miller, and Yang (2016)) articulates that higher effort levels increase the temptation to shirk.

Assumption 1. *The cost of effort c is smooth, strictly increasing, and strictly convex, with $c(0) = c'(0) = 0$ and $\lim_{x\to\infty} c'(x) = \infty$. The "relative cost" $c(x)/x$ is strictly increasing.*

Strict convexity with the limit condition guarantees that in equilibrium effort is bounded (as long as continuation payoffs are bounded, which we assume below). Increasing relative cost means a player requires proportionally stronger incentives to exert higher effort.

Solution concept Our solution concept is *plain perfect Bayesian equilibrium*, or PPBE (Watson 2016). This refinement of "weak perfect Bayesian" equilibrium Mas-Colell, Whinston, and Green (1995) imposes Bayesian updating on off-path beliefs, but is less restrictive and simpler to verify than sequential equilibrium or perfect extended-Bayesian equilibrium (Fudenberg and Tirole 1991; Battigalli 1996). Since we will construct a particular class of equilibria without making any claims about optimality, we could adopt a weaker solution concept. Our use of PPBE assures that our construction does not rely on the kinds of implausible off-path beliefs that are possible under weak perfect Bayesian equilibrium.

3 Robust community enforcement on a fixed network

In this section we show that community enforcement can sustain high cooperation in the repeated interaction game on whatever network arises from the network formation game, in a way that preserves cooperation among innocent players after a deviation, and without

realizations are observed only by that pair. We focus on sequential moves in the main model, because they yield strictly higher expected payoffs than simultaneous moves (see Remark 3 in the appendix).

requiring a player to know anything about the network structure or behaviors outside of his local neighborhood.

To begin with, we seek to sustain high levels of cooperation in society even off the equilibrium path.

Definition 1. *A strategy profile is* ***robust*** *if partners who have not deviated always cooperate at the same level, on and off the path of play.*

This property is stronger than the robustness criterion used by Jackson, Rodriguez-Barraquer, and Tan (2012), which allowed for cooperation to break down among a bounded set of innocent players following a deviation by one of their neighbors.

Next, the strategy profile only requires local knowledge, including both the network structure and the interactions.

Definition 2. *A strategy profile is* ***D-local*** *if each player i's strategy is invariant to her beliefs about interactions and network topology outside her D-neighborhood.*

We focus on strategy profiles that are 1-local, given a fixed network. It may be that an event in a player's 2-neighborhood, but not in his 1-neighborhood, may trigger a sequence of events that affects his behavior. However, he will not actually need to learn about that event. For example, in the network in Figure 1, player 4 ultimately punishes player 1 for having deviated on the $1{:}2$ link, but he does so because player 3 tells him that player 1 is guilty. That even off-path behavior is invariant to beliefs about the wider network makes this property stronger than the invariance criterion used by Nava and Piccione (2014), which allows players' behavior to depend on their beliefs about the global network in the course of play.[7]

3.1 Benchmark cooperation

Our benchmark for high cooperation is the maximum level of cooperation attainable by a stationary equilibrium on a triangle network. But first we introduce *bilateral cooperation*, the maximal cooperation attainable between two partners without the aid of community enforcement.

[7] The equilibrium Nava and Piccione construct (for a game of local interaction rather than bilateral interaction) exploits that possibility by inducing incorrect beliefs about the network following a deviation. The equilibrium we construct satisfies both our 1-locality criterion and their invariance criterion.

Bilateral cooperation Consider a strategy profile in which on the path of play first movers choose effort level x and second movers choose effort level y; off the equilibrium path each exerts zero effort. The equilibrium path incentive constraints are:

$$0 \leq -c(x) + b(y) + \int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x+y\big)dt \tag{1}$$

$$0 \leq -c(y) + \int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x+y\big)dt \tag{2}$$

The bilateral cooperation levels x^{B} and y^{B} are the effort levels that bind these incentive constraints. Since the grim trigger punishment is a minmax punishment and each partner's effort relaxes the other partner's incentive constraint, these are the maximum efforts that can be supported by any stationary equilibrium that does not involve community enforcement. Note that this implies $-c(y^{\mathrm{B}}) = -c(x^{\mathrm{B}}) + b(y^{\mathrm{B}})$; i.e., the gain from shirking is the same regardless of whether a player moves first or second, and on the equilibrium path the first mover receives a negative payoff in the stage game.[8]

Triangular cooperation Consider a triangle network, and a strategy profile in which on the path of play first movers choose effort level x and second movers choose effort level y. Off the equilibrium path, if the first mover deviates, the second mover chooses zero effort in the current interaction, and both then choose zero effort in all future interactions. If the second mover deviates, both choose zero effort in all future interactions. In such a strategy profile, if i deviates on j, both i and j will then shirk in their next meetings with k, so a "contagion" spreads until cooperation ceases over the whole network. Ali and Miller (2013) showed in a closely related model that such strategy profiles constitute equilibria if equilibrium-path incentive constraints bind, since then an induction argument guarantees the off-path incentive constraints. Moreover, since they implement minimax punishments, these equilibria maximize cooperation among all stationary equilibria. Here we focus only on the equilibrium-path constraints, to compute an upper bound on the cooperation that

[8] We could add another condition that no player receives negative payoff in the stage game, which would lower the level of efforts and the players' utilities, and the analysis of the model is analogous.

can be attained by any stationary equilibrium on a triangle network:

$$\int_0^\infty e^{-rt}e^{-2\lambda t}\lambda\frac{1}{2}b(x)\,dt \leq -c(x) + b(y) + 2\int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x+y\big)dt \tag{3}$$

$$\int_0^\infty e^{-rt}e^{-2\lambda t}\lambda\frac{1}{2}b(x)\,dt \leq -c(y) + 2\int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x+y\big)dt \tag{4}$$

These incentive constraints bind at effort levels x^{T} and y^{T}. As with bilateral cooperation, the gain from shirking is the same for the first mover and the second mover.

Lemma 1. *There exist triangular effort levels $x^{\mathrm{T}} > x^{\mathrm{B}}$ and $y^{\mathrm{T}} > y^{\mathrm{B}}$ that bind constraints (3) and (4).*

All omitted proofs are in Appendix A. The triangle can sustain higher levels of effort because each player faces more punishment if he deviates.

3.2 Robust cooperation with local knowledge

Our main result shows that high levels of cooperation can be sustained in a robust manner, with players needing only local information about the network and other players' behavior.

A link ij is *supported* if there exists k such that $ik \in G$ and $jk \in G$; i.e., if i and j have at least one common friend.

Theorem 1. *There exists a robust and 1-local multilateral restitution strategy such that on any network there exists a plain perfect Bayesian equilibrium in which all players employ this strategy, and triangular effort levels are attained on every supported link along the equilibrium path.*

Intuitively, robustness requires innocent partners to continue to cooperate with each other, in which case they do not have to worry about deviations that may occur outside their local knowledge. In order to achieve this robustness, innocent partners need the help of their guilty mutual friends to sustain cooperation at a high level off the equilibrium path. Accordingly, they should not use ostracism as illustrated in the top panel in Figure 1—instead they should punish the deviator in a less socially wasteful way. We describe the class of strategy profiles here, and provide the full proof in Section 3.4.

Multilateral restitution strategies operate as follows. Partners who have never deviated communicate "truthfully"; i.e., when partners i and j meet, partner i reveals who she deems guilty among the mutual friends she has with partner j. In the sequencing phase, they never nominate themselves to move first, with the result that the first mover is always randomly selected.[9] If their link is supported, they cooperate at the triangular levels x^{T} and y^{T}. If their link is unsupported, they cooperate at the bilateral levels x^{B} and y^{B}.[10]

Now consider a supported link ij and suppose that player i has deviated on player j. Then we say that i and j both deem i "guilty." If i's deviation occurs in the effort phase while moving first, j immediately chooses zero effort when moving second in the same interaction. In either case, thereafter they punish i on the ij link by requiring i to always be the first mover (that is, both nominate i in the sequencing phase), and choosing effort levels x^{P} and y^{P} that are calibrated to deliver the equilibrium path payoff to j while delivering a zero payoff to i:

$$-c(x^{\mathrm{P}}) + b(y^{\mathrm{P}}) = 0, \text{ and } b(x^{\mathrm{P}}) - c(y^{\mathrm{P}}) = \frac{1}{2}\big(x^{\mathrm{T}} + y^{\mathrm{T}}\big). \tag{5}$$

Lemma 2 shows that there exists a solution satisfying $0 < x^{\mathrm{P}} < x^{\mathrm{T}}$ and $0 < y^{\mathrm{P}} < y^{\mathrm{T}}$.

In addition to player i's punishment on the ij link, to support triangular levels of cooperation i must also be punished by at least one other neighbor, such as player k whose jk link supports ij link. However, once player i is being punished by both players j and k, i may no longer have an incentive to cooperate at triangular levels with other neighbors l, m, and so on. To accommodate this potential collapse of incentives, the equilibrium specifies that guilty i should also eventually be punished by every neighbor reachable via a path from j that is contained in i's 1-neighborhood but does not pass through i. The set of such players—including j—is denoted Σ_{ij}. Formally,

$$\Sigma_{ij} = \{k : \exists \text{ path } \{j_0, j_1, \ldots, j_l\} \subset g_i^1 \text{ s.t. } j_0 = j, j_l = k\}. \tag{6}$$

The punishments within individual relationships may of course be delayed, since information about i's deviation must be passed through the network. When i meets such a neighbor $k \in \Sigma_{ij}$, in the communication phase k will reveal if she deems i guilty, in which case i's punishment on the ik link starts immediately. If instead k still thinks i is innocent,

[9] As Section 3.1 noted, the first mover receives a lower payoff than the second mover, so no punishment is needed to deter deviations on the equilibrium path in the sequencing phase.

[10] Cooperation at the bilateral level on an unsupported link is motivated by the threat of letting the guilty player receive zero utility. Since play along unsupported links is measurable with respect to the interactions along their link, we may ignore them for the remainder of this discussion.

i may either shirk on k or pretend to be innocent by working. Regardless, eventually all players in Σ_{ij} will learn that i is guilty. Since i is punished along at least one other link after deviating on j, the punishment is sufficiently severe to support triangular cooperation.

Multilateral restitution equilibria are by construction robust and 1-local, and attain triangular effort on supported links. It remains for us to establish that there exists a multilateral restitution equilibrium. For this we need to tackle several difficulties stemming from the fact that the communication is via the network and needs to be incentive compatible, and the fact that the set of players Σ_{ij} that punishes player i for deviations on the ij link will generally intersect, but not coincide with, the similarly defined sets for other supported partnerships. Before proving the theorem in Section 3.4, we first illustrate the difficulty during the communication phase, and how a multilateral restitution equilibrium handles it.

3.3 Diffusion of information

With private monitoring, it is important to establish that information of a player's deviation diffuses through the network fast enough to provide sufficient punishment. Truthful communication cannot be taken for granted: Ali and Miller (2016) show that players would not communicate truthfully under permanent ostracism if they were cooperating above the bilateral effort. In this section, we investigate players' incentives for truthful communication.

First, a guilty player does not have incentives to communicate truthfully. This lack of incentive would not be a problem if the guilty player were guaranteed to shirk on every partner in the punishing set (Σ_{ij}) as soon as possible, so each partner would learn of his guilt in a timely manner. However, as shown in the following example, the guilty player has incentives to slow down the diffusion of information about his guilt, so he may work with some partners rather than shirking.

Example 1. *Consider a "generalized diamond" with "diagonal" players i and k and common neighbors $j, l_1, \ldots, l_m$, as shown in Figure 2. In a multilateral restitution equilibrium, if m is sufficiently large then i, after shirking on j along the equilibrium path, will not immediately start shirking on k.*

Proof. Consider the scenario that player i meets player k immediately after he shirks on

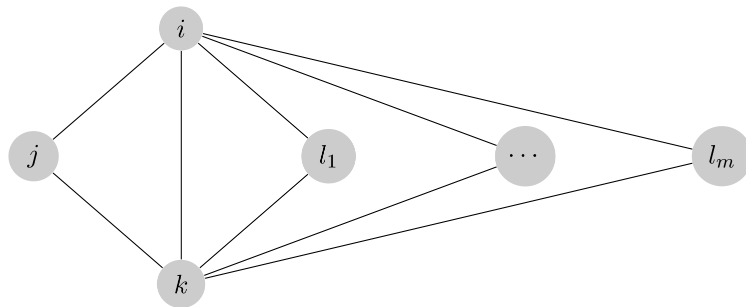

Figure 2: A generalized diamond

player j, so i has not met any of $l_1, \dots, l_m$ since deviating. In the communication phase, k indicates she does deem i guilty; then suppose nature chooses i to move first. If i shirks, his expected payoff is

$$0 + m \int_0^\infty e^{rt} e^{-\lambda t} \lambda \frac{1}{2} b(x) dt. \tag{7}$$

Now, we consider another strategy for i, which we will show guarantees a strictly higher payoff than (7) for m sufficiently large. In this strategy, i cooperates with k in the current meeting, and shirks on all his neighbors in the future. First, by cooperating with k he earns $b(y^T) - c(x^T) < 0$ in the current meeting, and at least zero thereafter on link ik. But on each link il_z, for $z = 1, \dots, m$, i expects a strictly higher payoff than he would if k knew about his shirking. It suffices to show that for m sufficiently large, the improvement is strictly greater than $c(x^T) - b(y^T)$.

We claim that when m is sufficiently large, with a probability above e^{-1}, player i meets at least $\sqrt{m} + 1$ players from the set $\{l_1, \dots, l_m\}$ before k knows about i's shirking. The probability that i meets at least $\sqrt{m} + 1$ of them before k meets i (again), j, or, any other player i has shirked on, is

$$\frac{m}{m+2} \cdot \frac{m-1}{m+2} \cdots \frac{m-\sqrt{m}}{m+2} \geq \left(\frac{m-\sqrt{m}}{m+2} \right)^{\sqrt{m}},$$

which converges to e^{-1} as $m \to \infty$. Conditional on i meeting these $\sqrt{m} + 1$ players before k learns i is guilty, he expects his first meeting with each of them to occur earlier on average than the unconditional expected first meeting time. Therefore i's expected payoff, from the perspective of his initial meeting with k after shirking on j, is at least $\sqrt{m}\, e^{-1} \int_0^\infty e^{-rt} e^{-\lambda t} \frac{1}{2} b(x^{\mathrm{T}})\, dt$. Since this lower bound is strictly increasing and linear in $\sqrt{m}$, the claim is proven. $\diamond$

Based on this example, we cannot assume that information always diffuses through every feasible channel. However, in a multilateral restitution equilibrium, innocent players are always willing to share information truthfully. The restitution punishments are specifically calibrated to deliver them their equilibrium path payoffs, even off the equilibrium path. Therefore they cannot benefit from slowing down the diffusion of information. Similarly, a guilty player is willing to communicate truthfully about the guilt of others, since it does not affect how other players treat him.

The truthful communication of innocent players puts a lower bound on the speed of information diffusion. In the generalized diamond network shown in Figure 2, while i cannot be relied upon to shirk on k after shirking on j, knowledge of i's deviation will be passed from j to k then to each l_z. This lower bound is sufficient to show that if players cooperate at triangular levels on all supported links along the equilibrium path, then i's continuation payoff on each link il_z is greater on the equilibrium path than it is off the equilibrium path after i shirks on j (see Lemma 3). This implies that even if player i hides his guilt from k, diffusion of knowledge about his guilt still suffices to deter his deviations on generalized diamond networks. In Lemma 4, we then use an induction argument to extend this conclusion to arbitrary networks.

3.4 Proof of Theorem 1

We prove that a multilateral restitution equilibrium exists and satisfies the desired properties, first on a triangle network, then on a "diamond" network, and then on a "generalized diamond". Ultimately, we prove by induction that it works on any arbitrary network.

1. Triangle network
2. Diamond network
3. Generalized diamond
4. Arbitrary network

Triangle network. Consider a triangle $\{i, j, k\}$. The incentives on the equilibrium path have been verified when deriving the effort levels x^{T} and y^{T}. We need to examine incentives off the equilibrium path. Start with verifying that player i, after initially deviating on player j, wants to shirk on player k when k does not know he is guilty, even if he moves first:

$$-c(x^{\mathrm{T}}) + b(y^{\mathrm{T}}) + \int_0^\infty e^{-rt} e^{-2\lambda t} \lambda \frac{1}{2} b(x^{\mathrm{T}})\, dt \leq 0. \tag{8}$$

This is implied by Equations (3) and (4). It follows that i also wants to shirk on k when moving second.

Once player i has deviated on player j, j's immediate gain from a deviating on the ij link decreases, as shown in the lemma below.[11]

Lemma 2. *$x^{\mathrm{P}} < x^{\mathrm{T}}$ and $y^{\mathrm{P}} < y^{\mathrm{T}}$.*

This lemma implies that innocent players prefer not to shirk even off the equilibrium path, when some of their partners are guilty. Moreover, innocent players cannot gain by being untruthful, and guilty players by construction cannot gain by deviating.

Diamond network. Consider a diamond network of 4 players, $\{i, j, k, l\}$ (as $\{1, 2, 3, 4\}$ in the network in Figure 1); i.e., all pairs are connected except for link jl. In this network, we call link ik the "diagonal."

Lemma 3. *On a diamond network, there exists a 2-local multilateral restitution equilibrium that supports triangular efforts on every link along the path of play.*

We focus on the deviation of i on j, because the analysis of other deviations are identical to either that of the ij deviation or a deviation on a triangle network. From the analysis of the triangle network, the loss of utility from the ik link is sufficient to deter i from shirking on j. Now with the additional link il which could be an additional shirking opportunity for i, the key is to show that i is also weakly worse off on the il link. Specifically, we want to show that the payoff i gets on the il link is weakly higher on the equilibrium path than

[11] Here we can invoke the restrictions on off-path beliefs imposed by PPBE: all players' beliefs must accord with the behavior specified by the strategy profile after i's deviation. In this case, j must believe that if she deviated on i and became guilty, k would learn of her guilt according to the stochastic process generated by the strategy profile. Henceforth we invoke these restrictions without further note.

after i has deviated on j. The payoff from the il link on the equilibrium path is

$$\int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)dt. \tag{9}$$

Off the equilibrium path, as discussed in Section 3.3, we consider the lower bound on information transmission which relies only on players other than i. That is, the news about i's deviation goes only from j to k then to l, but not from i to k then to l. Then, a potentially loose upper bound on i's expected payoff on the link il after his deviation is

$$\int_0^\infty e^{-rt}e^{-2\lambda t}\Big(\lambda\frac{1}{2}b(x^{\mathrm{T}})+\lambda\int_0^\infty e^{-r\tau}e^{-2\lambda\tau}\lambda\frac{1}{2}b(x^{\mathrm{T}})\,d\tau\Big)dt. \tag{10}$$

That is, i can shirk on l and gain $b(x^T)$ if i moves second when meeting l (with 1/2 probability), provided that l does not know i is guilty. The news will not reach l in time in two scenarios: i meets l before j meets k, or otherwise i meets l before k meets l. Then, we prove in the appendix that (10) is weakly lower than (9).

Generalized diamond. A generalized diamond has a diagonal link ik and two or more common neighbors of players i and k, $\{l_1,\ldots,l_m\}$, as shown in Figure 2. It is straightforward to see that i does not benefit from deviating on j along the equilibrium path in any generalized diamond: first, he is punished by k, and then, by Lemma 3, he gets a weakly lower payoff on each link il_z after shirking on j. Similarly, if he deviates on k then he will be punished by all $l_1,\ldots,l_m$.

Arbitrary network. Now we complete the proof of Theorem 1 by the following lemma.

Lemma 4. *Consider any arbitrary neighborhood Σ_{ij}. Triangular efforts on supported link ij are sustained by a multilateral restitution equilibrium.*

Although the punishing set Σ_{ij} can be large and complex in an arbitrary network, we can divide Σ_{ij} into subsets and prove players' incentives to sustain high cooperation by induction. To begin with, because ij is supported, $|\Sigma_{ij}|\geq 2$. When $|\Sigma_{ij}|=2$, the local neighborhood is a triangle, and i's incentive to cooperate with j in a multilateral restitution strategy profile has been verified. Next, suppose i has incentives to cooperate with j under a multilateral restitution strategy profile when facing any punishing set with size $|\Sigma_{ij}|\leq m$. Then we consider when the punishing set has size $|\Sigma_{ij}|=m+1$. We can divide Σ_{ij} into two subsets: the first one containing j and at least one of her partners is a triangle, a diamond or a generalized diamond, and the second one is the rest of Σ_{ij}. Player i receives sufficiently punishment from the first subset of partners as shown in the previous three cases, and he does not have incentives to shirk on any partner in the second subset by the induction hypothesis. Thus, his cooperation with j can be sustained.

3.5 To be added

- Formal definition of multilateral restitution strategies and associated beliefs

4 Implications for evolving networks

Since partnerships supported by multilateral enforcement are particularly valuable, players may strategically seek them out. In this section we introduce a dynamic network formation game with random linking opportunities and random link formation costs. We identify a simple equilibrium in this game in which players form a network with realistic *small worlds* properties, because they anticipate playing in a multilateral restitution equilibrium after forming their links. We focus in particular on measures of "clustering" and "support": *clustering* is the fraction of paths of length two that are contained in triangles; *support* is the fraction of links that are contained in triangles.

4.1 Dynamic network game

We follow the setup of a growing network in Jackson and Rogers (2007). According to an independent clock process (Poisson or otherwise), at each arrival a new player $i \in \mathbb{N}$ is born, and, in that instant, forms new network links according to the following four *arrival stages*. Players are numbered according to their order of arrival. Each existing link continues to be recognized for partnership meetings according to its independent Poisson clock of rate λ as described in Section 2, and players' actions at those meetings follow multilateral restitution strategies.

- Arrival stage 0: An i.i.d. random cost γ_{ij} is drawn for each pair of players $\{i, j\}$ with $j < i$, from a distribution with CDF F. These costs are not revealed to the players.

- Arrival stage 1: m_r players* are uniformly randomly selected from a pool of players born before i. Each selected pair $\{i, j\}$ jointly learns their cost γ_{ij}. Player i decides which, if any, of these links to form. If i forms link ij, he pays $2\gamma_{ij}$.

- Arrival stage 2: m_n players* are uniformly randomly selected from the union of i's stage-1 friends' friends. Each selected pair $\{i, k\}$ jointly learns their cost γ_{ik}; in addition, player k learns about any links that formed in her 1-neighborhood in arrival stage 1. Player i decides which, if any, of these links to form. If i forms link ik, he pays $2\gamma_{ik}$.[12]

- Arrival stage 3: In uniform random sequence, each triple $\{i, j, k\}$ such that ij was formed in arrival stage 1, jk is connected, and $\{i, k\}$ was selected in arrival stage 2 but the link ik did not form, is recognized again. When triple $\{i, j, k\}$ is recognized, the players jointly learn γ_{jk} and the union of their 1-neighborhoods, including links formed in arrival stages 1–2. Player i decides whether to form the jk link. If i forms the link, he pays $2\gamma_{ik}$.

- Arrival stage 4: All players to whom i linked simultaneously pay non-negative transfers to player i. These transfers are observed only by player i.

* If the pool of eligible players is smaller than the number to be selected, then all eligible players are selected.

At the end of these arrival stages, each player in the network observes any new links that formed in their own 1-neighborhood.

4.2 Network formation equilibrium

While a multilateral restitution equilibrium exists in the repeated interaction game for any network, in this section, we show that players who arrive anticipating joining a multilateral restitution equilibrium will form a network with certain realistic characteristics. For tractability, we focus on an equilibrium that is somewhat naive from the players' collective

[12] It is worth noting that our network-based search differs slightly from Jackson and Rogers (2007) because in their model these m_n players are chosen from the union of all m_r stage-1 players' friends regardless of whether i forms a link to them or not in Stage 1. In our setup, we assume these m_n players are chosen from friends of those who i forms links to in Stage 1. This is because anticipating cooperation, it is more beneficial to search through linked friends.

perspective, due to some coordination failures. Let $u^{\mathrm{B}} = \frac{1}{2}(x^{\mathrm{B}} + y^{\mathrm{B}})$ be the utility of a link with bilateral cooperation, and let $u^{\mathrm{T}} = \frac{1}{2}(x^{\mathrm{T}} + y^{\mathrm{T}})$ be the utility of a link with triangular cooperation. We define a "greedy linking" strategy profile for player i's arrival as follows.

- Arrival stage 1: When a pair $\{i, j\}$ is recognized to meet, i forms link ij iff $\gamma_{ij} \leq u^{\mathrm{B}}$.
- Arrival stage 2: When a pair $\{i, k\}$ is recognized to meet, i forms link jk iff $\gamma_{ik} \leq u^{\mathrm{T}}$.
- Arrival stage 3: When a triple $\{i, j, k\}$ is recognized to meet, i forms link ik iff $\gamma_{ik} \leq u^{\mathrm{T}} + z_{ijk}(u^{\mathrm{T}} - u^{\mathrm{B}})$, where $z_{ijk} \in \{0, 1, 2\}$ is the number of unsupported links in the triple at the time they meet.
- Arrival stage 4: If a link ij was formed but should not have been according to i's strategy profile, then player j pays zero to player i. Otherwise, if a link ij was formed in arrival stage 1–2, then player j pays γ_{ij} to player i. In addition, if the link ik is formed for the triple $\{i, j, k\}$, then players j and k each pay player i $\frac{1}{3}z_{ijk}(u^{\mathrm{T}} - u^{\mathrm{B}})$.

The arriving player i deems guilty any player j who deviates from the transfer he or she should have paid in arrival stage 4 if innocent; player i deems all other new neighbors innocent. The network formed upon player i's arrival (assuming player i does not deviate) is denoted G_i.

It is clear that each link that should form is strictly beneficial to the (innocent) players who form it. However, the players suffer from some coordination failure in Stages 1 and 3 that causes them to forego profitable linking opportunities. In Stage 1, they fail to jointly anticipate links that may form in Stages 2 and 3; optimally anticipating later links would lead them to form some links in Stage 1 for which $\gamma_{ij} > u^{\mathrm{B}}$. In Stage 3, they fail to jointly anticipate links that may form later in Stage 3: since supporting links are formed by a greedy algorithm rather than an optimal algorithm, a less beneficial link may be chosen early over a more beneficial one that is recognized later. Focusing on an equilibrium with these coordination failures aids in finding closed form expressions for network statistics like support and clustering, and in simulations.

The next result shows that equilibrium cooperation behaviors in Theorem 1 are not disrupted by the link formation process.

Theorem 2 (loosely stated)**.** *In the dynamic network game, there exists a multilateral restitution and greedy linking equilibrium that attains triangular effort levels on every supported link along the path of play.*

Proof idea. We observe that no player has incentives to shirk and also all innocent players share information truthfully.

First, consider players $i > j > k$. Suppose j has shirked on k before player i is born. We claim that player j gets less value from i's entry compared to that when j is innocent. First, the same set of links are formed by i regardless of j's guilt, because j's guilt is not anticipated by i. If j pays his share of the cost for link ij, then the link formation is identical to that when j is innocent. However, j faces weakly more punishment because $\Sigma_{jk}(G_{i-1}) \subset \Sigma_{jk}(G_i)$. If j does not pay his share of the cost for link ij, j immediately becomes guilty to i and thus gets zero utility from the link ij. Moreover, the set of players to punish j becomes strictly larger: $\Sigma_{jk}(G_{i-1}) \subsetneq \Sigma_{jk}(G_i) \cup \Sigma_{ji}(G_i)$.

Second, each innocent player $l \in \Sigma_{jk}(G_{i-1})$ shares the information about j's shirking truthfully. This is straightforward because player i forms the same set of links, and each innocent player pays the same amount of her share of the cost for the link to i. As l's utility does not depend on whether other players know about j's guilt, she will not benefit from withholding such information. ◇

In arrival stage 1, each possible link forms with probability $p_r \equiv F(u^{\mathrm{B}})$. In arrival stage 2, players can form clusters by closing triples; each eligible link forms with probability $p_n \equiv F(u^{\mathrm{T}})$. In arrival stage 3, they can revisit links not formed in arrival stage 2 to support their relationships.

Remark 1. *As the network grows, it inherits several properties from the results of Jackson and Rogers (2007), including:*

- *Connectedness: The limiting probability that any two players are connected by a path is 1.*
- *Fat-tailed degree distribution: The tail of the limiting degree distribution is heavier than the exponential distribution.*

These observations arise from the fact that without arrival stage 3 the network formed is essentially a Jackson and Rogers (2007) network, and then arrival stage 3 adds additional links.

Most importantly, we address clustering and support. Let $m = p_r m_r + p_n m_n$ and $\rho = p_r m_r / p_n m_n$.

Remark 2. *As the network grows,*

- *Low clustering: The limiting probability that two neighbors of the same player are linked is strictly positive if $\rho > 1$, and it converges to 0 as $m_r \to \infty$.*
- *High support: The limiting probability that two linked players share a common neighbor is strictly positive, and it converges to a value above $1 - e^{-1/\rho^*} > 0$ as $m_r \to \infty$, where $\rho^* = \lim_{m_r \to \infty} \rho$.*

We begin with examining these two measures when the network is formed absent of stage 3, denoted as C_2 and S_2 for clustering and support, respectively. Then, we build stage 3 on top of it to bound clustering and support in the network, denoted as C_3 and S_3.

Support and clustering from arrival Stages 1–2 When the network is formed without arrival stage 3, our model is very similar to Jackson and Rogers (2007). By analysis analogous to theirs, the clustering measure is[13]

$$C_2 = \begin{cases} 0 & \text{if } \rho \leq 1; \\ \frac{6}{(1+\rho)[(3m-2)(\rho-1)+2m\rho]} & \text{if } \rho > 1. \end{cases} \tag{11}$$

We now calculate the support measure. All $p_n m_n$ links formed by network-based search must be supported, so we focus on the $p_r m_r$ links formed randomly. Say ij is a random link. (While we consider undirected links, we use the order ij to indicate the link is formed by i arriving when player j is already present). There are three possibilities to support ij. First, it could be supported by another random link ik formed in player i's arrival stage 1, such that jk is linked. As the population increases, the probability of such a link jk goes to zero. Second, it could be supported by a later player h, who forms a random link hi in her arrival stage 1 and then forms hj in her arrival stage 1 or 2. Again, as the population increases, the probability of link hi goes to zero. So, we are left with the third case in which ik is formed by network-based search through the link jk, in player i's arrival stage 2.

For each of m_n players, with the probability $\frac{1}{p_r m_r}$ it is search through player j, say through link jk, and i forms a link to k with probability p_n. Thus the probability link ij

[13]Their clustering expression features an extra term in the numerator, to account for the fact that they allow player i to meet neighbors of a player j he met in Stage 1, even if he did not form link ij in arrival stage 2. In our model, player i meets player j's neighbors in arrival stage 3 only if the ij link was formed in arrival stage 2.

is supported

$$\beta_2 = 1 - \left(1 - \frac{p_n}{p_r m_r}\right)^{m_n}. \tag{12}$$

In other words, the link ij is not supported if none of her links to all m_n friends of friends she meets supports ij. To sum up,

$$S_2 = \frac{p_r m_r \beta_2 + p_n m_n}{p_r m_r + p_n m_n}. \tag{13}$$

Support and clustering with arrival stage 3 We first calculate the probability, denoted as β_3, that a random-search link ij formed in player i's arrival stage 1 does not become supported in arrival stage 2, but does become supported in arrival stage 3. Recall that each link formed in stage 3 must support one or two other links. Let $\alpha_{3L} \equiv F(2u^T - u^B)$ be the probability of a link cost below the value of both cooperating at the triangular level (u^T) and supporting one other link ($u^T - u^B$). And let $\alpha_{3H} \equiv F(3u^T - 2u^B)$ be the the probability of a link cost below the value of both cooperating at the triangular level and supporting two other links. Using (12), the lower bound on β_3 is

$$\beta_3 \geq \beta_{3L} = \left(1 - \frac{p_n}{p_r m_r}\right)^{m_n} - \left(1 - \frac{\alpha_{3L}}{p_r m_r}\right)^{m_n}. \tag{14}$$

Let K_3 be the expected number of links i forms in stage 3. Thus, $K_3 \geq \frac{1}{2}\beta_{3L} p_r m_r$. From right to left, this is because i has in expectation $p_r m_r$ links formed in arrival stage 1, each of which is unsupported in arrival stage 2 but supported in arrival stage 3 with probability at least β_{3L}, but the link that brings the support may also support a second link and thus for a given i represents at least half of an additional link.

Putting them together, the lower bound on support in the network is

$$S_3 \geq \frac{(\beta_2 + \beta_{3L}) p_r m_r + p_n m_n + \frac{1}{2}\beta_{3L} p_r m_r}{p_r m_r + p_n m_n + \frac{1}{2}\beta_{3L} p_r m_r}. \tag{15}$$

Similarly, the upper bound on β_3 is

$$\beta_3 \leq \beta_{3H} = \left(1 - \frac{p_n}{p_r m_r}\right)^{m_n} - \left(1 - \frac{\alpha_{3H}}{p_r m_r}\right)^{m_n}. \tag{16}$$

The number of links i forms in arrival stage 3 is at most $K_3 \leq \beta_{3H} p_r m_r$. Thus, the upper bound on support in the final network is

$$S_3 \leq \frac{(\beta_2 + \beta_{3H}) p_r m_r + p_n m_n + \beta_{3H} p_r m_r}{p_r m_r + p_n m_n + \beta_{3H} p_r m_r}. \tag{17}$$

Notice that $S_3 > S_2 > \beta_2$, and $\beta_2 \to 1 - e^{-1/\rho^*}$ as $m_r \to \infty$. Thus, we have shown the second point of Remark 2.

When $\rho \leq 1$, $C_3 = 0$. When $\rho > 1$,

$$\begin{aligned} C_3 &\leq \frac{3m^2 \frac{1}{m(1+\rho)} + 3\beta_{3H} \frac{\rho m}{\rho+1}}{\frac{m(m-1)}{2} + m^2 + \frac{m(2m\rho+1-\rho)}{2(\rho-1)} + 2\beta_{3H} \frac{\rho m}{\rho+1}}, \\ &= \frac{6 + 6\beta_{3H}\rho}{(1+\rho)[3m - 2 + \frac{2m\rho}{\rho-1}] + 4\beta_{3H}\rho}. \end{aligned} \tag{18}$$

When $m_r \to \infty$, $m \to \infty$, and thus C_3 converges to zero.

Lastly, we can obtain the level of "support-seeking" by comparing the support in the network with and without arrival stage 3. In particular, "seeking for support" is bounded between $s_{3L} - s_2$ and $s_{3H} - s_2$. Some algebra shows that as $m_r \to \infty$, if $\rho^* > 0$ then $s_{3L} - s_2$ converges to a strictly positive limit. Indeed, support-seeking can be quantitatively quite significant. For example, consider $\rho = 2$ and $\alpha_{3L}/p_n = \alpha_{3H}/\alpha_{3L} = 2$; that is, the probability of forming a supported link that supports another link is twice as that of forming a supported link that supports no other link, and is half of that of forming a supported link that supports two other links. Then, $s_2 = 0.60$, $s_{3L} = 0.77$, and $s_{3H} = 0.92$, so that support-seeking increases the support measure by 30–55%.

4.3 To be added

- Formalize Theorem 2 and proof
- Formalize Remark 2 and proof

- Simulation results on clustering and support: What is “low” and “high” when $m_r \ll \infty$? =

5 Concluding remarks

We introduce a class of multilateral restitution equilibria, which implements multilateral enforcement on an arbitrary network, while avoiding cascading punishments and allowing the players to know only their local neighborhoods. The key component of multilateral restitution equilibria is that guilty players are not ostracized from the community; instead they work hard with their partners to preserve the stability of the network. Then, anticipating playing the multilateral restitution equilibria, arriving players who form new links seek support for their relationships. We show that as the network grows, it exhibits high support and low clustering.

References

S. Nageeb Ali and David A. Miller. Enforcing cooperation in networked societies. Working paper, 2013.

S. Nageeb Ali and David A. Miller. Ostracism and forgiveness. *American Economic Review*, 106(8):2329–2348, August 2016.

S. Nageeb Ali and David A. Miller. Designing markets to foster communication and cooperation. *American Economic Journal: Microeconomics*, 2021.

S. Nageeb Ali, David A. Miller, and David Yilin Yang. Renegotiation-proof multilateral enforcement. Working paper, 2016.

Attila Ambrus, Markus Möbius, and Adam Szeidl. Consumption risk-sharing in social networks. *American Economic Review*, 104(1):149–182, January 2014.

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

Emily Breza, Arun G. Chandrasekhar, and Alireza Tahbaz-Salehi. Seeing the forest for the trees? an investigation of network knowledge. *working paper*, 2016.

Arthur Campbell. Signaling in social network and social capital formation. *Economic Theory*, 57(2):303–337, October 2014.

Tiziana Casciaro. Seeing things clearly: social structure, personality, and accuracy in social network perception. *Social Networks*, 20(4):331–351, 1998.

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

Andrea Galeotti, Sanjeev Goyal, Matthew O. Jackson, Fernando Vega-Redondo, and Leeat Yariv. Network games. *Review of Economic Studies*, 77(1):218–244, January 2010.

Benjamin Golub and Yair Livne. Strategic random networks and tipping points in network formation. Working paper, November 2010.

Mark S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78: 1360–1380, 1973.

Matthew O. Jackson and Brian W. Rogers. Meeting strangers and friends of friends: How random are social networks? *American Economic Review*, 97(3):890–915, June 2007.

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.

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

David Krackhardt. Assessing the political landscape: Structure, cognition, and power in organizations. *Administrative Science Quarterly*, 35(2):342–369, 1990.

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

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

Francesco Nava and Michele Piccione. Efficiency in repeated games with local interaction and uncertain monitoring. *Theoretical Economics*, 9(1):279–312, January 2014.

Robert B. Thompson. Measure of recovery under rule 10b-5: A restitution alternative to tort damages. *Vanderbilt Law Review*, 37(2):349–398, March 1984.

Eric van Damme. Renegotiation-proof equilibria in repeated prisoners' dilemma. *Journal of Economic Theory*, 47(1):206–217, 1989.

Joel Watson. Perfect Bayesian equilibrium: General definitions and illustrations. Working paper, January 2016. URL http://econweb.ucsd.edu/~jwatson/PAPERS/WatsonPBE2015.pdf.

Alexander Wolitzky. Cooperation with network monitoring. *Review of Economic Studies*, 80(1):395–427, January 2013.

A Proofs

Proof of Lemma 1. Evaluating the integrals in equation (4) and rearranging, we have

$$c(y) \leq \frac{\lambda}{r}(x+y) - \frac{\lambda}{2(r+2\lambda)}(y+c(y)).$$

Rearrange it,

$$\frac{r(2r+5\lambda)}{2\lambda(r+2\lambda)}c(y) - \frac{r+4\lambda}{2(r+2\lambda)}y \leq x.$$

Take the cost of both sides of the inequality, we have

$$c\left(\frac{r(2r+5\lambda)}{2\lambda(r+2\lambda)}c(y) - \frac{r+4\lambda}{2(r+2\lambda)}y\right) \leq c(x) = y + 2c(y). \tag{19}$$

The last equality uses the observation that in the binding case $-c(y^{\mathrm{T}}) = -c(x^{\mathrm{T}}) + b(y^{\mathrm{T}})$.

By Assumption 1, $c'(0) = 0$, $c'(x)$ increases in x, and $\lim_{x\to\infty} c'(x) = \infty$. So, there exists a unique $\underline{y} > 0$ such that

$$\frac{r(2r+5\lambda)}{2\lambda(r+2\lambda)}c(\underline{y}) = \frac{r+4\lambda}{2(r+2\lambda)}\underline{y}.$$

The LHS of (19) is zero when $y = \underline{y}$, while the RHS must be positive, so (19) holds. Then, as y increases to infinity, the cost increases much faster than the value. As a result, when y is sufficiently large, the LHS of (19) is always higher than the RHS. Specifically, we can first identify the threshold y' such that when $y > y'$, $c(y) > y$. By convexity, $c(3y) \geq 3c(y) > y + 2c(y)$. Next we can find the threshold y'' such that when $y > y''$, $\frac{r(2r+5\lambda)}{2\lambda(r+2\lambda)}c(y) - \frac{r+4\lambda}{2(r+2\lambda)}y > 3y$. Then when $y > \max(y', y'')$, LHS of (19) is always higher than the RHS. Thus, there must exists a value y^{T} such that when $y = y^{\mathrm{T}}$, LHS is equal to RHS, and when $y > y^{\mathrm{T}}$, LHS is always higher than RHS. x^{T} can be calculated by $c(x^{\mathrm{T}}) = y^{\mathrm{T}} + 2c(y^{\mathrm{T}})$. ◇

Proof of Lemma 2. We prove $y^{\mathrm{P}} < y^{\mathrm{T}}$ by contradiction. Suppose $y^{\mathrm{P}} \geq y^{\mathrm{T}}$. Summing

$b(y^{\mathrm{P}}) - c(x^{\mathrm{P}}) = 0$ and $b(x^{\mathrm{P}}) - c(y^{\mathrm{P}}) = \frac{1}{2}\left(x^{\mathrm{T}} + y^{\mathrm{T}}\right)$, it leads to

$$2(x^{\mathrm{P}} + y^{\mathrm{P}}) = x^{\mathrm{T}} + y^{\mathrm{T}} \leq x^{\mathrm{T}} + y^{\mathrm{P}}.$$

Then $x^{\mathrm{T}} > x^{\mathrm{P}} + y^{\mathrm{P}}$. On the other hand, by $c(x^{\mathrm{T}}) = c(y^{\mathrm{T}}) + b(y^{\mathrm{T}})$, we have $c(y^{\mathrm{P}}) + b(y^{\mathrm{P}}) > c(x^{\mathrm{T}})$. Together with $b(y^{\mathrm{P}}) - c(x^{\mathrm{P}}) = 0$, it implies $c(x^{\mathrm{P}}) + c(y^{\mathrm{P}}) > c(x^{\mathrm{T}})$. Since $c(x)$ is strictly convex, it must be that $x^{\mathrm{T}} < x^{\mathrm{P}} + y^{\mathrm{P}}$. It is a contradiction. So $y^{\mathrm{P}} < y^{\mathrm{T}}$ must hold.

Lastly, observing that $c(x^{\mathrm{P}}) = b(y^{\mathrm{P}}) < c(y^{\mathrm{T}}) + b(y^{\mathrm{T}}) = c(x^{\mathrm{T}})$, so $x^{\mathrm{P}} < x^{\mathrm{T}}$. ⋄

Proof of Lemma 3. Note that the analysis after a deviation by j or l is identical to the analysis on a triangle network. Moreover, innocent players are always willing to implement the prescribed punishments (since they still get their equilibrium path payoffs, but would be punished themselves for any deviation), and to communicate truthfully. So it suffices to consider only equilibrium path deviations by i.

First, consider whether player i could gain by shirking on player j, on the equilibrium path. Once player k knows of i's deviation, incentives on the ik link are straightforwardly similar to the triangle case; this is also true for the il link. So we focus on what happens when i meets a partner (either k or l) who does not know of i's deviation. (Note that a partner who does know of i's deviation should demonstrate that knowledge in the pre-play communication phase.) Our class of strategy profiles does not specify whether i should work or shirk in such meetings. What is important is that whatever happens should not lead to payoffs for i that are high enough to justify shirking on j in the first place.

Suppose, after shirking on player j, player i meets player l; l does not know of i's deviation, and i does not know whether k knows of i's deviation. Observe that on the equilibrium path, if j were not present then i would have been just indifferent between working (at effort x^{T} or y^{T}, depending on whether he moved first) and shirking on l. Now off the equilibrium path, i expects zero future payoffs on the ij link, but j's presence means k and ultimately l will learn of i's original deviation sooner in expectation. This loss of future social collateral strictly reduces i's incentive to work with l, compared to the equilibrium path on a triangle network. Hence i strictly prefers to shirk on l, regardless of i's belief about the probability that k knows of i's deviation. (Moreover, it follows from analysis of the triangle that after shirking on l, i subsequently strictly prefers to shirk on k.)

Next consider player i (after shirking on player j) meeting player k, when k does not know of i's deviation, and when i has not yet shirked on l. Now matters are a bit more complicated—by working rather than shirking on k, i can slow down the rate at which l learns that i has deviated. The most i can slow down the rate at which l learns of the deviation is to work with k until either k learns of i's deviation from j, or i shirks on l. We already know this makes i strictly worse off on the ik link than shirking immediately (since i would be indifferent on the equilibrium path of a triangle, but here j will spread the news to k).

The key is to show that i is also weakly worse off on the il link. That is, the payoff i gets on the il link on the equilibrium path is weakly higher than when i has deviated on j; this is satisfied if

$$\int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)dt \geq \int_0^\infty e^{-rt}e^{-2\lambda t}\Big(\lambda\frac{1}{2}b(x^{\mathrm{T}})+\lambda\int_0^\infty e^{-r\tau}e^{-2\lambda\tau}\lambda\frac{1}{2}b(x^{\mathrm{T}})\,d\tau\Big)dt, \tag{20}$$

which we verify below. Note that we consider the slowest information transmission, such that the news about i's deviation goes only from j to k then to l, but not from i to k then to l. So the value on the RHS of (20) is a loose upper bound on i's expected payoff on the link il after his deviation.

We now verify that (20) holds. From summing the binding incentive constraints in (3) and (4) with $b(x^{\mathrm{T}})$, we have

$$b(x^{\mathrm{T}})+\int_0^\infty e^{-rt}e^{-2\lambda t}\lambda b(x^{\mathrm{T}})\,dt = x^{\mathrm{T}}+y^{\mathrm{T}}+2\int_0^\infty e^{-rt}\lambda\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)dt. \tag{21}$$

Then, we simplify the RHS of equation (20):

$$\begin{aligned}
&\int_0^\infty e^{-rt}e^{-2\lambda t}\Big(\lambda\frac{1}{2}b(x^{\mathrm{T}})+\lambda\int_0^\infty e^{-r\tau}e^{-2\lambda\tau}\lambda\frac{1}{2}b(x^{\mathrm{T}})\,d\tau\Big)dt\\
&=\frac{\lambda}{r+2\lambda}\cdot\frac{1}{2}\left(b(x^{\mathrm{T}})+\int_0^\infty e^{-rt}e^{-2\lambda t}\lambda b(x^{\mathrm{T}})\,dt\right)\\
&=\frac{\lambda}{r+2\lambda}\cdot\frac{1}{2}\left(x^{\mathrm{T}}+y^{\mathrm{T}}+2\int_0^\infty e^{-rt}\lambda\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)dt\right)\\
&=\frac{\lambda}{r+2\lambda}\cdot\frac{1}{2}\cdot\frac{r+2\lambda}{r}\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)\\
&=\frac{\lambda}{r}\cdot\frac{1}{2}\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)\\
&=\int_0^\infty e^{-rt}\lambda\frac{1}{2}\big(x^{\mathrm{T}}+y^{\mathrm{T}}\big)dt,
\end{aligned}$$

where the second equality is from (21). Thus equation (20) holds with equality.

Then, consider whether player i could gain by shirking on player k along the equilibrium path. Since k then spreads the news to both players j and l, this is strictly worse for i than shirking on j on the equilibrium path.

Finally, consider players other than the original deviator i: as in our analysis of the triangle network above, they expect to receive equilibrium-path payoffs on all their links (being shirked on is always a surprise) and therefore have no incentive to deviate on or off the equilibrium path. ◇

Proof of Lemma 4. Because ij is supported, $|\Sigma_{ij}| \geq 2$. We prove the lemma by induction on the number of players in Σ_{ij}. When $|\Sigma_{ij}| = 2$, the local neighborhood is a triangle, and i's incentive to cooperate with j in a multilateral restitution strategy profile has been verified. Suppose i has incentives to cooperate with j under a multilateral restitution strategy profile when facing any punishing set such that $|\Sigma_{ij}| \leq m \geq 2$. Then we consider when the punishing set has size $|\Sigma_{ij}| = m + 1 \geq 3$.

We consider two separate cases. First, if i has is at least one partner k whose link to i is uniquely supported by j (i.e., if the ij link were removed, the ik link would be unsupported), then we partition Σ_{ij} as follows. Σ_{ij}^1 includes j and all players whose links to i are uniquely supported by j, and $\Sigma_{ij}^2 = \Sigma_{ij} \setminus \Sigma_{ij}^1$. Note that players in $\Sigma_{ij}^1 \setminus \{j\}$ do not have any links to Σ_{ij}^2. Therefore once i deviates on j, there is no incentive for i to delay shirking on neighbors in Σ_{ij}^1 in order to slow the rate at which neighbors in Σ_{ij}^2 learn he is guilty, nor vice versa. Hence it suffices to show that i does not gain from deviating on j separately on each subnetwork $\{i\}\cup\Sigma^1_{ij}$ and $\{i,j\}\cup\Sigma^2_{ij}$. As for $\{i\}\cup\Sigma^1_{ij}$, it is a generalized diamond with ij being its diagonal, so by our previous analysis Σ^1_{ij} itself suffices to deter i from deviating on j. As for $\{i,j\}\cup\Sigma^2_{ij}$, it is a punishing set of size $|\Sigma^2_{ij}|\leq m$, which by the induction hypothesis itself suffices to deter i from deviating on j.

The second case is when there is no player whose link to player i is uniquely supported by ij. Then we arbitrarily choose player k (a common neighbor of i and j, at least one of which exists since ij is supported), and partition Σ_{ij} as follows. Σ^1_{ij} includes j, k, and any neighbor of i whose link to i is unsupported in $\Sigma_{ij}\setminus\{i,j\}$, while $\Sigma^2_{ij}=\Sigma_{ij}\setminus\Sigma^1_{ij}$. Then $\{i\}\cup\Sigma^1_{ij}$ is a generalized diamond combined with zero or more links among Σ^1_{ij}. Since the added links (none of which directly involve i) aid in distributing information about i's deviation on j among the players in Σ^1_{ij}, Σ^1_{ij} itself suffices to deter i from deviating on j. As in the first case, Σ^2_{ij} is a punishing set of size $|\Sigma^2_{ij}|\leq m$, which by the induction hypothesis itself suffices to deter i from deviating on j. So the overall punishment from Σ_{ij} deters i from deviating on j.

Off the equilibrium path, by Lemma 2 the same analysis applies to any innocent player, even if she has guilty neighbors.

Evidently the strategies form an equilibrium, which by construction is robust and 2-local. ◇

B Extensions and discussion

B.1 Sequential vs. simultaneous moves

We show that players earn strictly higher expected payoffs from sequential moves. We prove this result for bilateral cooperation, and it is straightforward to extend it to triangular cooperation. Recall that in a bilateral cooperation, players use x^{B} and y^{B} when they move sequentially, and let them both use z^B when they move simultaneously.

Remark 3. *In a bilateral cooperation:* $x^{\mathrm{B}}>z^B$ *and* $y^{\mathrm{B}}>z^B$*.*

Proof. If players move simultaneously, their incentive constraint is

$$0\leq -c(z)+\int_0^\infty e^{-rt}\lambda z dt=-c(z)+\frac{\lambda}{r}z$$

So z^B satisfies $c(z^B) - \frac{\lambda}{r} z^B = 0$. By Assumption 1, there is a unique solution of $z^B > 0$.

If players move sequentially, x^{B} and y^{B} bind the constraints (1) and (2). In particular, summing up (1) and (2), we have $c(x^{\mathrm{B}}) - \frac{\lambda}{r} x^{\mathrm{B}} = (1 + \frac{\lambda}{r}) y^{\mathrm{B}}$. Since $y^{\mathrm{B}} > 0$, it is clear that $x^{\mathrm{B}} > z^B$. Next, (2) implies $c(y^{\mathrm{B}}) - \frac{\lambda}{r}\frac{1}{2}(x^{\mathrm{B}} + y^{\mathrm{B}}) = 0$. When $y^{\mathrm{B}} = z^B$, $c(z^B) - \frac{\lambda}{r}\frac{1}{2}(x^{\mathrm{B}} + z^B) < 0$ because $x^{\mathrm{B}} > z^B$. Thus, $y^{\mathrm{B}} > z^B$. $\diamond$