Francesco Nava
Repeated games and networks

Book section

http://eprints.lse.ac.uk

CHAPTER 6

REPEATED GAMES AND NETWORKS

FRANCESCO NAVA

6.1 INTRODUCTION

In many strategic environments, interaction is local and segmented. Competing neighborhood stores serve different yet overlapping sets of customers; informal lending and insurance arrangements often have to be fulfilled by relatives and friends; the behavior of the residents of an apartment block affects their contiguous neighbors to a larger extent than neighbors in a different block; a nation's foreign or domestic policy typically generates larger externalities for neighboring nations than for remote ones. One classic case is the private provision of local public goods. In addition to local interaction, one notable feature of these environments is local monitoring: whereas participants are aware of their own neighbors' identities and actions, they are not necessarily aware of the identity and actions of their neighbors' neighbors. Within these strategic environments, it is of particular interest to study long-run interaction, when incentives can only be provided locally in a decentralized manner. The main objective of this literature is analyzing such interactions within a repeated game framework that differs from the standard one in that actions can only be observed locally.

Three main lines of research have been developed in such environments. The first, and most classical, develops Folk Theorems for games with local monitoring, and establishes that network structure is usually irrelevant for enforcing cooperation when the frequency of interaction is sufficiently high. The other two explicitly study the link between network structure and the equilibrium payoffs by focusing on environments in which discount rates are fixed. One strand analyzes how the monitoring structure affects the maximal level of equilibrium cooperation, and broadly finds that larger and/or better connected groups are more cooperative. The other evaluates how different communication protocols affect the set of equilibrium payoffs and the incentives to cooperate in environments with local monitoring.

The analysis of community enforcement was initially developed in the context of repeated games with random matching. Pioneering studies by Kandori (1992) and Ellison (1994) focused on environments with pairwise matching, and established how collective punishments could sustain efficient equilibrium outcomes when bilateral punishments would fail. Subsequent and related contributions on random matching games include Harrington (1995), Takahashi (2010), and Deb (2014).[1] Although the matching literature and the literature focusing on stable local interactions (and therefore on networks) share several methodological insights, there are significant differences both in the assumptions on feasible interactions and in the broad aims. Whereas most random matching games assume all players potentially interact, and thus exchange information about deviant behavior, all network games constrain interactions, monitoring, and information exchange to take place on a stable network, which represents the topology of relationships in a society. Whereas most random matching games (with a few exceptions, including Harrington 1995) focus on Folk Theorems and seldom on optimal punishments, the study of network games aims to establish a relationship between the underlying network structure and the equilibrium correspondence (or alternatively the most efficient equilibrium payoff).

The chapter begins by presenting relevant definitions in the context of a baseline environment with local monitoring and local interaction. It proceeds with a survey of Folk Theorems for network games in Section 6.3. Sections 6.4 and 6.5 discuss community enforcement at a given frequency of interaction. In particular, Section 6.4 surveys results on optimal punishments and network structure, while Section 6.5 presents results on communication. Section 6.6 hints at related applications (on reciprocity, informal insurance, and lending), at relevant omissions, and at possible extensions. Static games with local interactions and the related literature are discussed in a separate chapter of this book (Chapter 8).

6.2 A Baseline Setup

Environments considered in the literature invoke different assumptions about information, matching, and the availability of individual punishments. This section introduces a baseline environment which nests a large number of possible setups to discuss contributions more transparently in the following sections.

The Stage Game: Consider a game, the stage game, played by a set N of n players in which any player i can interact with a subset of players $N_i \subseteq N \backslash \{i\}$, which is called the *neighborhood* of player i. As customary, assume that $j \in N_i$ if and only if $i \in N_j$. This

---

[1] Harrington (1995) shows that relationships with low frequencies of interaction can be supported using relationships that interact more frequently. Takahashi (2010) shows that cooperation can be sustained in repeated Prisoners' Dilemmas if all that is observed are partners' past play. Deb (2014) offers a general Folk Theorem for anonymous random matching environments.

structure of interaction defines an undirected graph (N, G) in which $ij \in G$ if and only if $j \in N_i$. Refer to G as the *interaction network*. In the stage game players interact with a possibly random subset of their neighbors. In particular, for any undirected subgraph $\bar{G} \subseteq G$, let $f(\bar{G}|G)$ denote the probability that the realized network of interactions is $\bar{G}$. Let $\bar{N}_i \subseteq N_i$ denote the realized neighborhood of i in this subgraph. Two extreme cases are generally considered in the literature. In the first case $f(G|G) = 1$, which I refer to as *local interaction*, while in the second $f(\bar{G}|G) > 0$ implies that $|\bar{N}_i| \leq 1$, which I refer to as *pairwise interaction*. The former scenario captures environments in which players interact with all of their neighbors in every period, while the latter environments in which interactions take place only between pairs of players. Refer to f as the *matching technology*.

Assumptions on information vary across setups, but consistently require that players know their neighborhood, N_i, their realized neighborhood, $\bar{N}_i$, and the matching technology, f. When players are privately informed about their neighborhood, their beliefs regarding the interaction network, conditional upon observing their neighborhood, are derived from a common prior distribution over the set of interaction networks. Beliefs regarding the realized interaction network are then constructed by simply applying Bayes' rule.

The action set of player i is denoted by A_i. Given a subset M of players, let A_M denote $\times_{j \in M} A_j$ and a_M an element of A_M. Also, let $-i$ denote the set $N \setminus \{i\}$. The stage game payoffs are common knowledge. The payoff of any player i depends only on actions chosen in his realized neighborhood, and is denoted by $v_i(a_i, a_{\bar{N}_i} | \bar{N}_i)$. As a convention, the payoff equals zero when $\bar{N}_i$ is empty. Payoffs are *separable*, if for any player $i \in N$ the stage game payoff satisfies

$$v_i(a_i, a_{\bar{N}_i} | \bar{N}_i) = \sum_{j \in \bar{N}_i} u_{ij}(a_i, a_j),$$

where $u_{ij}(a_i, a_j)$ is the payoff of player i from the relationship $ij \in G$.

The stage game is *separable* if: (a) payoffs are separable; (b) action sets have the product structure, $A_i = \times_{j \in N_i} A_{ij}$ for any $i \in N$; and (c) for any action profile $a_N \in A_N$, the stage game payoff on any link $ij \in G$ satisfies

$$u_{ij}(a_i, a_j) = x_{ij}(a_{ij}, a_{ji}),$$

for some map $x_{ij} : A_{ij} \times A_{ji} \to \mathbb{R}$. In a separable game, players choose actions that are specific to each interaction, and payoffs in an interaction depend only on actions chosen in that specific interaction. Any pairwise interaction game can be represented as a separable game, if the identity of players is known to their realized partners. If so, action sets have the product structure, as players can tailor behavior to every opponent. Thus, non-anonymous random matching games are separable. Anonymous random matching games, instead, are not separable, since action sets do not have the product structure as players cannot choose a different action for every realized interaction.

The stage game is *binary-symmetric* if: (i) payoffs are separable; (ii) action sets are binary, $A_i = \{C,D\}$ for any $i \in N$; and (iii) payoffs are symmetric, for any link $ij \in G$,

$$u_{ij}(a_i, a_j) = \eta_{ij} u(a_i, a_j),$$

for some map $u : \{C,D\}^2 \to \mathbb{R}$ and a some scalar $\eta_{ij} \in \mathbb{R}_+$. In such games, players must choose the same action in every interaction and cannot discriminate across neighbors. For convenience, refer to action C as cooperation and action D as defection. Results on binary-symmetric games are generally developed for stage games in which: (iv) the payoff $u(a_i, a_j)$ of player i in an interaction with j satisfies

$i \backslash j$	C	D
C	1	$-l$
D	$1+g$	0

(v) mutual cooperation is efficient, $g - l < 1$; and (vi) defection is a best response when the opponent cooperates $g > 0$. The first assumption restricts the class of binary games by imposing a common payoff for mutual defection across relationships; the second uniquely pins down an efficient action profile; while the third rules out the trivial case in which mutual cooperation is an equilibrium of the stage game. Naturally, if $l > 0$, the stage game has a unique Bayes Nash equilibrium in which all players play D, and all pairwise interactions amount to a Prisoners' Dilemma. If instead $l < 0$, the stage game always possesses a mixed strategy Bayes Nash equilibrium, and all pairwise interactions amount to an anti-coordination game.[2]

Local interaction games with separable and symmetric payoffs capture environments in which behavior cannot be targeted to individual neighbors, while separable games capture environments in which players can make decisions contingent on the identity of their realized neighbor. For instance, decentralized competition between sellers, when prices set are independent of identity of buyers, fits in the class of local interaction games with separable payoffs; whereas non-anonymous negotiations between traders in an spatial economy fit in the class of separable games.

Repetition and Local Monitoring: The players play the infinite repetition of the stage game. The interaction network, G, is realized prior to the beginning of play and remains fixed throughout the game. Realized interactions, $\bar{G}(t)$, however, are drawn independently every period, t, from a distribution over the set of subnetworks of G.[3]

Monitoring is *local* implying that a player observes only the past play in his realized neighborhood. Local monitoring is a key assumption in the networks approach to community enforcement as it implies that realized interactions are not anonymous.

---

[2] When $l < 0$, pure strategy equilibria also exist in some networks, as miscoordinating with neighbors can be best reply. In particular, if beliefs are concentrated on bipartite graphs (which have only cycles of even length; Bramoullé 2007), pure equilibria exist, since all players can successfully miscoordinate their action with all their neighbors.

[3] A subnetwork G' of a networks G is a subset of G. That is, $G' \subseteq G$.

This differs from classical random matching models requiring anonymity, as players may now develop relationship-specific reputations to enforce good behavior. Formally, when the stage game is not separable, a history h_i^t of length t for player i consists of a sequence

$$h_i^t = \left\{ N_i, \bar{N}_i(0), \bar{a}_i(0), \bar{N}_i(1), ..., \bar{a}_i(t-1), \bar{N}_i(t) \right\}$$

that satisfies $N_i \subseteq N$, $\bar{N}_i(s) \subseteq N_i$, and $\bar{a}_i(s) \in \times_{j \in \bar{N}_i(s) \cup \{i\}} A_j$ for any value of s. When the stage game is separable, however, players monitor only the neighbor-specific actions played in their realized interactions and therefore $\bar{a}_i(s) \in \times_{j \in \bar{N}_i(s)} [A_{ij} \times A_{ji}]$ for any value of s. Denote by H_i^t the set of histories of length t for player i, and by H_i the corresponding set of possible histories, $H_i = \cup_{t=0}^{\infty} H_i^t$. A strategy for player i is a map that assigns to every history in H_i an action in A_i. A full history h^t of length t similarly consists of a sequence

$$h^t = \left\{ G, \bar{G}(0), \bar{a}(0), \bar{G}(1), ..., \bar{a}(t-1), \bar{G}(t) \right\}$$

satisfying $G \subseteq \{ij | i, j \in N\}$, $\bar{G}(s) \subseteq G$, and $\bar{a}(s) \in A_N$. Denote by H^t the set of full histories of length t and by H the set of possible full histories $H = \cup_{t=0}^{\infty} H^t$.

Players discount the future by a common factor $\delta \leq 1$. To construct the payoffs in the infinitely repeated game, fix a player $i \in N$ and a history $h_i \in H_i$, and let h_i^t denote the subhistory of length $t > 0$ of h_i. Define

$$w_i^t\left(h_i^t\right) = \sum_{s=0}^{t-1} \frac{v_i(a(s)|\bar{N}_i(s))}{t}$$

to be the average payoff up to period t and $w_i\left(h^i\right) = \{w_i^t\left(h_i^t\right)\}_{t=1}^{\infty}$ to be the sequence of average payoffs. Repeated game payoffs conditional on h_i are defined as

$$\mathcal{V}_i(h_i) = \begin{cases} (1-\delta) \sum_{t=0}^{\infty} \delta^t v_i(a(s)|\bar{N}_i(s)) & \text{if} \quad \delta < 1 \\ \Lambda\left(w_i\left(h_i\right)\right) & \text{if} \quad \delta = 1 \end{cases}$$

where $\Lambda(\cdot)$ denotes a suitable limit operator, such as the limit inferior or the Banach-Mazur limit of a sequence.[4]

A full history h uniquely pins down the history of play in the dynamic game. An observed history h_i is associated uniquely with an information set $\mathcal{I}(h_i)$ for player i and vice versa. A system of beliefs defines, at each information set $\mathcal{I}(h_i)$ of player i, the conditional probability of each full history $h \in \mathcal{I}(h_i)$.

Departures: Although the baseline setup allows for much flexibility, it does not capture the full range of environments considered in the literature. Some studies model the

---

[4] If ℓ_∞ denotes the set of bounded sequences of real numbers, a Banach-Mazur limit is a linear functional $\Lambda : \ell_\infty \to \mathbb{R}$ such that: (i) $\Lambda(e) = 1$ if $e = \{1, 1, ...\}$; (ii) $\Lambda(x^1, x^2, ...) = \Lambda(x^2, x^3, ...)$ for any sequence $\{x^t\}_{t=0}^{\infty} \in \ell_\infty$ (Aliprantis and Border 2005). It can be shown that, for any sequence $\{x^t\}_{t=0}^{\infty} \in \ell_\infty$,

$$\liminf_{t \to \infty} x^t \leq \Lambda\left(\{x^t\}_{t=1}^{\infty}\right) \leq \limsup_{t \to \infty} x^t.$$

interaction network as a directed graph. Others allow for *global interaction*, while assuming that monitoring is local. If so, players may be affected by the action chosen by every other player in the game, but only observe the behavior of a subset of players. Other frameworks have considered *imperfect local monitoring* by having players only observe a noisy signal of their neighbors' actions. Finally, many setups have focused on communication, by adding to the stage game a communication stage, modelled in one of many possible ways.

6.3 Limiting Results and Network Irrelevance

A significant body of literature provides conditions on the interaction network for a Folk Theorem to apply. These studies generally establish in many environments that a Folk Theorem obtains under very weak conditions on the network structure, and thus yield limited insights about the optimal monitoring structure. A key concern in these papers is ensuring that players do not cooperate off the equilibrium path, as grim trigger strategies may provide such strong incentives to cooperate on the equilibrium path that players prefer to cooperate even after observing a deviation. Ellison (1994) resolves this problem by introducing either a public randomization device or a milder version of grim trigger strategies tailored to make players indifferent between cooperating and defecting on-path, and then noting that cooperation is more appealing on-path than off-path (since off-path at least one opponent is already defecting). The literature on local monitoring addresses similar concerns with related approaches, either by allowing some form of communication, or by constructing suitable strategies with mild punishments. Further complications arise, however, with local monitoring as, upon observing defections, players try to infer the spread of defection and the beliefs of other players about future play.

All of the limiting results presented in this section apply to stage games with local monitoring that are not separable, since a Folk Theorem would trivially obtain otherwise. Most results are developed for stage games with local interaction and in which the network structure is common knowledge. Thus for expositional ease, restrict attention to such scenarios unless specified otherwise.

Ben-Porath and Kahneman's (1996) seminal contribution considers games with public randomization and in which players can make public announcements about the past behavior of other players whom they observed. The analysis characterizes the minimal level of observability required to obtain efficient outcomes for arbitrary stage games. The main result establishes that, when the discount factor tends to one, the limit set of sequential equilibrium payoffs contains the set of individually rational payoffs, whenever every player is observed by at least two other players. For arbitrary stage game payoffs, two monitors are required to guarantee that inconsistent public

announcements about past play can be sanctioned by the community. Results also establish that if payoffs are assessed by the limit inferior of the average payoff (that is if $\delta = 1$), every individually rational payoff is a sequential equilibrium payoff even when players are monitored by only one other player.

Renault and Tomala (1998) develop similar insights in a model with global interactions, local monitoring, no discounting, and no explicit communication. Their main finding establishes that a Nash Folk Theorem applies if and only if the monitoring network is 2-connected (that is, if there are two independent paths connecting any two players, or equivalently, if the subgraphs obtained by suppressing any one player are still connected). The result abstracts from sequential rationality, which considerably simplifies the problem as punishments need not be incentive compatible. Although explicit communication is ruled out, the no discounting assumption and the restriction to Nash equilibrium imply that players can use any finite number of future periods to privately communicate with neighbors at no cost. Tomala (2010) extends the analysis to partially known networks, in which players only know their neighbors and the number of players in the network, and derives a Nash Folk Theorem.

More recently, Laclau (2012) considers a local interaction setup analogous to Renault and Tomala (1998), while allowing for imperfect local monitoring and explicit communication between neighbors (private local cheap talk). Monitoring is imperfect, as players observe their payoff, but not the actions chosen in their neighborhood. Her main contribution identifies necessary and sufficient conditions on the network of interactions for a Nash Folk Theorem to hold when the payoff of every player is responsive to unilateral deviations (in that players monitor unilateral deviations in their neighborhood, despite local monitoring being imperfect). In a recent companion paper, Laclau (2014) extends conclusions to a model in which communication is global (players can communicate with all opponents), and can be either private or public. Contrary to Laclau (2012), where a Nash Folk Theorem is established, the analysis here applies to sequential equilibria of the infinitely repeated game with imperfect local monitoring. As before, payoffs are assumed to be sensitive to unilateral deviations. If so, a sequentially rational Folk Theorem holds provided that a joint pairwise identifiability condition regarding payoff functions is satisfied. The condition requires players to detect the identity of the deviating player, whenever they detect a unilateral deviation in their neighborhood. The analysis then shows that, when payoffs are sensitive to unilateral deviations, a necessary and sufficient condition on the network topology for the Folk Theorem to hold for all payoff functions is that no two players have the same set of neighbors (not counting each other). The main contribution of both papers consists in the analysis of imperfect local monitoring, which had been neglected by the earlier literature.

Three related studies, Xue (2004), Cho (2010), and Cho (2011), analyze cooperation in binary-symmetric Prisoners' Dilemma games. Even though it is not difficult to construct sequential equilibria supporting cooperation in these environments, the classical modification of a trigger strategy devised in Ellison (1994) to enforce a cooperative equilibrium has an undesirable feature. Namely, it is not stable to mistakes

in that defections spread over the network, and cooperation is never recovered whenever an agent defects by mistake in the repeated game. The main aim of these three studies, thus, consists of constructing equilibria that sustain cooperation and revert to cooperation after any history of play. The classical solution to this complication involves bounding the length of the punishment phase. That is, if an agent observes his neighbor playing defection, then he punishes his neighbor by defecting for finitely many periods. Local monitoring, however, may cause discrepancies in beliefs between agents about a neighbor's future actions (that is, the expected date at which a player ends a defection phase may not be common knowledge in his neighborhood). If there is such a discrepancy at some history, then an agent whose neighbors have different beliefs may not be able to satisfy the expectations of all his neighbors in any period which in turn may cause an infinite repetition of defection phases and, thus, a failure of stability. Furthermore, bounded punishment strategies may not even constitute a sequential equilibrium in a general networked setting. In order to prove the existence of a cooperative and stable sequential equilibrium, such discrepancies of beliefs may be resolved through some form of coordination in punishments. To this end, Xue (2004) restricts attention to line-shaped networks, and shows that in such graphs cooperation is a stable equilibrium when players comply with specific bounded punishment strategies. Cho (2010) establishes a similar results for acyclical networks by allowing agents to communicate locally with their neighbors. In contrast to Laclau (2012), the focus is on sequential equilibria; while in contrast to Laclau (2014), communication is only local, and not public, and therefore players cannot easily coordinate their punishments. Both Xue (2004) and Cho (2010) exploit the acyclicity of the network structure to simplify the inference problem associated with contagion, as players expect punishments to dissipate at the periphery of the network. Borrowing an idea from Ellison (1994), Cho (2011) instead shows that a cooperative and stable sequential equilibrium exists for any possible monitoring structure if players have access to a public randomization device. If so, the inference problem is solved through coordinating behavior rather than by restricting the class of network structures.

Nava and Piccione (2014) study a broader class of binary-symmetric games which satisfy the additional requirements (iv)–(vi) described in Section 6.2. In contrast to the earlier results, but similarly to Tomala (2011), the study allows for uncertainty about the interaction network. In particular, to capture behavior in large markets the analysis postulates that players are privately informed of their neighborhood. Their main result establishes that, for sufficiently high discount rates and any prior beliefs with full support about the network structure, sequential equilibria exist in which efficient stage-game outcomes are played in every period. Standard results do not apply in this framework because bilateral enforcement may not be incentive compatible when punishments in one relationship affect outcomes in all the others. For instance, punishing a neighbor indefinitely with a grim trigger strategy is not viable if cooperation in other relationships is disrupted (see Figure 6.1), and mild trigger strategies such as in Ellison (1994) work only for particular specifications of payoffs (e.g., Prisoners'

FIGURE 6.1 With trigger strategies, the central player prefers not to punish a single defection, as it would destroy cooperation in all his remaining relationships.

Dilemma). Equilibrium strategies supporting efficient outcomes are built so that players believe that cooperation will eventually resume, after any history of play.

The result is constructive, and exploits simple bounded-punishment strategies which are robust with respect to the players' priors about the monitoring structure. In particular, in the equilibria characterized only local information matters to determine players' behavior. Efficiency is supported by strategies that respond to defections with further defections. When the players' discount rate is smaller than one, the main difficulty in the construction of sequentially rational strategies that support efficiency is the preservation of short-run incentive compatibility after particular histories of play which involve several defections. When defections spread through a network, two complications arise. The first occurs when a player expects future defection coming from a particular direction. Suppose that somewhere in a cycle, for example, a defection has occurred and reaches a player from one direction. If this player does not respond, he may expect future defections from the opposite direction caused by players who are themselves responding to the original defection (see Figure 6.2). This player's short-term incentives then depend on the timing and on the number of future defections that he expects. In such cases, the verification of sequential rationality and the calculation of consistent beliefs can be extremely demanding. The analysis circumvents this difficulty by constructing consistent beliefs, which imply that a player never expects future defections to reach him (as unexpected behavior is always blamed on a neighbors' defection). Such beliefs are generated trivially when priors assign positive probability only to acyclic monitoring structures. More importantly, such beliefs can always be generated when priors have full support. The second complication arises when a player has failed to respond to a large number of defections. On the one hand, matching the number of defections of the opponent in the future may not be incentive compatible, say when this player is currently achieving efficient payoffs with a large number of different neighbors (as was the case with grim trigger strategies). The restriction that a player's action is common to all neighbors is of course the main source of complications here. On the other hand, not matching them may give rise to the circumstances outlined in the first type of complications, that is, this player may then expect future defections from a different direction. The former hurdle is circumvented by bounding the length of punishments, while the latter, as before, by constructing appropriate consistent beliefs.

FIGURE 6.2 Cycling defections complicate incentives constraints, as beliefs about the timing of future defections become the key force driving behavior.

Some of these difficulties do not arise when players are patient (that is, if $\delta = 1$) as short-term incentives are irrelevant and punishments need not be bounded. Indeed, stronger results hold for the case of limit discounting in which payoffs are evaluated according to Banach-Mazur limit of the average payoff. If so, efficiency is resilient to histories of defections. In particular, there exists a sequential equilibrium such that, after any finite sequence of defections, paths eventually converge to the constant play of efficient actions in all neighborhoods in every future period. An essential part of the construction is that in any relationship in which defections have occurred the number of periods in which inefficient actions are played is "balanced": as the game unfolds from any history, both players will have played the inefficient action an equal number of times before resuming the efficient play. Remarkably, such balanced retaliations eventually extinguish themselves and always allow the resumption of cooperation throughout the network. Although the analysis is restricted to homogeneous discount rates and symmetric stage games with deterministic payoffs, the equilibria characterized are robust with respect to heterogeneity in payoffs and discount rates, and with respect to uncertainty in payoffs and population size, as long as the ordinal properties of the stage games are maintained across the players. These equilibria also persist as babbling equilibria in setups with communication. In addition, results extend to accommodate monitoring structures in which players interact with fewer players than they observe.

6.4 FIXED DISCOUNTING AND NETWORK AMPLIFICATION

Much of the literature on community enforcement (discussed in the introduction and in Section 6.3) focuses on the case of sufficiently high discount factors and does not characterize efficient equilibria at fixed discount factors. A major concern in these papers was ensuring that players did not cooperate off the equilibrium path. The literature on repeated networked games with fixed discount factors abstracts from such a concern by analyzing the most cooperative equilibrium in games with

continuous action sets. Such equilibria make players indifferent between cooperating and defecting on-path (as otherwise a player could be asked to cooperate more). By essentially the same argument as in Ellison (1994), this implies that players weakly prefer to defect off-path. Hence, a contribution of this literature is to show that grim trigger strategies provide the strongest possible incentives for cooperation on-path, not that they provide incentives for punishing off-path. The characterization of the most cooperative equilibrium has implications for the efficiency and stability of various network configurations which are the main objectives of this literature.

This approach was pioneered by several papers in public economics analyzing the effect of the size and structure of a group on the maximum equilibrium level of public good provision. Classical references, however, characterize maximum cooperation only for complete networks and public monitoring, and find few unambiguous relationships between group structure and maximum cooperation. Pecorino (1999) shows that public good provision is easier in large groups because a deviation causing everyone else to defect is more costly in large groups. Haag and Lagunoff (2007) consider a broader class of public goods games in a similar setup and characterizes the maximal average level of cooperation (MAC) over all stationary subgame perfect equilibrium paths. The MAC is shown to be increasing in monotone shifts of, and decreasing in mean preserving spreads of the distribution of discount factors. The latter suggests that more heterogeneous groups are less cooperative on average. Furthermore, in a class of Prisoners' Dilemma games, the MAC exhibits increasing returns to scale for a range of heterogeneous discount factors. That is, larger groups are more cooperative, on average, than smaller ones. By contrast, when the group has a common discount factor, the MAC is invariant to group size.

Haag and Lagunoff (2006) relax the public monitoring assumption and examine optimal network structure in a binary-symmetric Prisoners' Dilemma with local interactions and local monitoring, in which each individual's discount factor is randomly determined. A planner chooses a local interaction network before the discount factors are realized in order to maximize utilitarian welfare. A local trigger strategy equilibrium (LTSE) describes a sequential equilibrium in which each individual conditions his cooperation on the cooperation of at least a subset of neighbors. The main results restrict attention to the LTSE associated to the highest utilitarian welfare, and demonstrate a trade-off in the design problem between suboptimal punishments and social conflict. Potentially suboptimal punishments arise in designs with local interactions since monitoring is local. Owing to the heterogeneity of discount factors, however, greater social conflict may arise in more connected networks. When individuals' discount factors are known to the planner, the optimal network exhibits a cooperative core and an uncooperative fringe. Uncooperative players are impatient, and are connected to cooperative ones who are patient and tolerate their free riding so that social conflict is kept to a minimum. By contrast, when the planner knows only the ex-ante distribution over individual discount factors, in some cases the optimal design partitions individuals into isolated cliques, whereas in other cases incomplete graphs with small overlap are possible.

Two recent and related studies have addressed similar questions in the context of continuous action games with local monitoring, namely Wolitzky (2013) and Ali and Miller (2013). Both models feature smooth actions and payoffs so that, with grim trigger strategies, binding on-path incentive constraints imply slack off-path incentive constraints. Wolitzky (2013) studies cooperation in repeated networked games with a fixed and common discount factor. The setup displays local monitoring, while allowing for global interaction, and generalizes environments analyzed in Kandori (1993) and Ellison (1994). In particular, the analysis considers public goods games with continuous actions in which players choose of a level of cooperation (in that higher actions are privately costly but benefit everyone). Payoffs are separable, but depend on the action chosen by every other player in the game. In every period, a monitoring network is realized and players receive signals about the global structure of the realized network. Players perfectly observe the actions of their realized neighborhood, but observe nothing about any other player's action. A distinguishing feature of the environment analyzed is that in every period the monitoring network must be observed by players after actions are chosen. The assumption is substantial in the equilibrium construction, and results do not generally apply to alternative specifications in which uncertainty about the realized monitoring persists over time.[5]

The study characterizes the maximum level of cooperation that can be robustly sustained in Perfect Bayesian equilibrium (in that it can be sustained for any information that players may have about the realized monitoring network). The robustness criterion captures the perspective of an outside observer, who knows what information players have about each other's actions, but not what information players have about each other's information about actions, and who must make predictions that are robust to higher-order information. Determining the maximum level of cooperation for any specification of players' higher-order information appears intractable, as the strategies sustaining the maximum level of cooperation could in principle depend on players' private information in complicated ways. However, the main theoretical contribution establishes that the robust maximum level of cooperation is always sustained by simple grim trigger strategies, where each player cooperates at a fixed level unless he ever observes another player failing to cooperate at his prescribed level, in which case he stops cooperating forever. Grim trigger strategies also maximize cooperation when players have perfect knowledge of who observed whom in the past (as is the case when the monitoring network is fixed over time). Interestingly, it is when players have less information about the monitoring structure that more complicated strategies can do better than grim trigger. This is the case because the actions of different players are strategic complements when players know who observed whom in the past, as defecting makes every on-path history less likely when monitoring is local and strategies are

---

[5] The key role of the assumption is to ensure that stage game actions are one-dimensional (so that players simply choose a level of cooperation, rather than a map from the realized monitoring network to a level of cooperation).

grim triggers.[6] The strategic complementarity breaks down, however, when players can disagree about who has observed whom.

The analysis then compares different economies in terms of the maximal level of cooperation that can be achieved. Results are developed for two special cases: equal monitoring (when in expectation all players are monitored equally well); and fixed monitoring network (when the monitoring network is fixed over time). With equal monitoring, the effectiveness of a monitoring technology in supporting cooperation is completely determined by one simple statistic, its effective contagiousness, which captures the cumulative expected present discounted number of players who learn about a deviation. Naturally, higher levels of cooperation can be sustained if news about a deviation spreads throughout the network more quickly. Cooperation in the provision of pure public goods (when the marginal benefit of cooperating is independent of group size) is increasing in group size if the expected number of players who learn about a deviation is increasing in group size, while cooperation in the provision of divisible public goods (when the marginal benefit of cooperating is inversely proportional to group size) is increasing in group size if the expected fraction of players who learn about a deviation is increasing in group size. Hence, cooperation in the provision of pure public goods tends to be greater in larger groups, while cooperation in the provision of divisible public goods tends to be greater in smaller groups. In addition, there is a sense in which making monitoring more uncertain reduces cooperation. With a fixed monitoring network, instead, a novel notion of network centrality determines both which players cooperate more in a given network and which networks support more cooperation overall, thus linking the graph-theoretic property of centrality with the game-theoretic property of robust maximum cooperation. For example, adding links to the monitoring network necessarily increases all players' robust maximum cooperation, which formalizes the idea that individuals in better-connected groups cooperate more.

Ali and Miller (2013) analyze community enforcement in a pairwise interaction game in which the network is common knowledge. Their analysis compares interaction networks in terms of maximal level of cooperation in variable-stakes Prisoners' Dilemmas. Results establish that cliques are optimal network structures when players' equilibrium path behavior is stationary. Results are developed in the context of a continuous time model in which all players discount the future at a common fixed rate. Every link of the network is governed by an independent Poisson recognition process with a common recognition rate. Whenever a link is recognized, an instantaneous two-period interaction is played within the selected relationship. In the first subperiod, both players propose stakes at which they intend to interact, and the smallest of the two proposals determines the actual stakes in the relationship. In the second subperiod, players engage in a Prisoners' Dilemma. If both cooperate, each receives a payoff which coincides with the agreed stakes; if both defect, each receives a payoff equal to zero; whereas when one defects while the other cooperates, the cooperating player incurs a

---

[6] Actions are strategic complements if a player is willing to cooperate more at any on-path history whenever another player cooperates more at any on-path history.

cooperation-loss which may depend on the agreed stakes, while the defecting player receives a deviation-gain which may also depend on the agreed stakes.[7] The stage game is said to satisfy strategic complementarity whenever stakes exceed the difference between deviation gains and cooperation losses.[8] Monitoring is local. Thus, players ignore both the actions chosen in interactions to which they did not belong and the time at which these interactions took place. The analysis restricts attention to stationary strategies in which behavior is independent of the history of play at any equilibrium path. A stationary equilibrium in which players always cooperate for any possible equilibrium-path history is said to be a mutual effort equilibrium. Any stationary grim trigger strategy profile that prescribes stakes such that incentive constraints bind at every equilibrium path history is therefore a Perfect Bayesian equilibrium by the arguments above.

The main result establishes that any symmetric network[9] with degree d possesses a symmetric contagion equilibrium that Pareto dominates every distinct mutual effort equilibrium (and thus identifies the optimal stakes). The result also implies that no other stationary equilibrium has a higher value if the stage game satisfies strategic complementarity. The argument relies on a measure of network viscosity (which is minimal in the clique) that captures the incentives to comply with equilibrium strategies. This measure differs from the effective contagiousness in Wolitzky (2013) because in public goods games every player may punish a deviator upon receiving news of a defection, whereas in separable games only neighbors can effectively punish a deviation.

Results exploit the characterization of the optimal stakes to analyze how network structure affects aggregate welfare. Adding links has two roles in the model: it helps information diffusion through contagion; and it increases the number of interactions (as links are recognized at the same rate) and consequently the expected surplus when cooperating. The main welfare implication of the model is the optimality of cliques. In particular, for any network in which the maximal degree is no more than d, no player attains a mutual effort equilibrium payoff that exceeds his optimal equilibrium payoff in the symmetric network of degree d. Moreover, if the stage game satisfies strategic complementarity, then the value in every equilibrium is less than the optimal value in the symmetric network of degree d.

---

[7] The deviation-gain is assumed: to exceed stakes; to be zero when stakes are equal to zero; to be strictly increasing and strictly convex in stakes; to have a first derivative that is greater than 1 at zero and diverging to infinity as stakes diverge. The cooperation-loss is also assumed to be zero when stakes are equal to zero.

[8] When strategic complementarity holds, mutual cooperation is efficient in the stage game. But the assumption is stronger. Even with two players, efficiency of mutual cooperation does not ensure that optimal equilibria are mutual effort, which is why the stronger assumption is invoked.

[9] A permutation of the players $\pi : N \to N$ is a graph automorphism if $ij \in G$ implies $\pi(i)\pi(j) \in G$. A network G is symmetric if for any two links $ij, kl \in G$ there exists a graph automorphism π such that $\pi(i) = k$ and $\pi(j) = l$. In a symmetric network, all links are isomorphic to each other.

Results also extend such logic to a model in which players incur an additively separable cost of forming links which depends only on the number of neighbors they have. Costs are said to be concave if the average link cost weakly decreases with the number of links; while costs are convex otherwise. When costs are concave, there exists a symmetric Perfect Bayesian equilibrium on the complete network that yields to each player a payoff that is higher than his payoff in any mutual effort equilibrium of any other incomplete network. Moreover, when the game satisfies strategic complementarity, the claim holds for every equilibrium (not just for mutual effort equilibria). When costs are convex, the welfare maximizing network may no longer be the complete network, and the analysis applies only to regular networks. In such cases it is possible to find the clique size that maximizes the payoff of a player in the welfare maximizing equilibrium, and the associated optimal value. No mutual effort equilibrium on any regular network attains payoffs that exceed such value. Moreover, if the stage game satisfies strategic complementarity, no equilibrium on any regular network attains a higher value.

All the papers discussed provide novel and interesting insights linking interaction and monitoring networks to measures of aggregate welfare. These observations can in principle explain why community enforcement may lead to substantially different levels of cooperation across societies. The main limitation of these studies, however, is the restrictive class of games to which results apply, as results are generally developed for Prisoners' Dilemma type stage games possessing a mutual minmax Nash equilibrium. Generalizing techniques to arbitrary stage game does not seem straightforward, as the characterization of the equilibrium with the highest utilitarian welfare may become intractable.

6.5 Fixed Discounting and Communication

A separate strand of the literature analyzes how equilibrium outcomes are affected by the availability of different communication technologies. These studies include Lippert and Spagolo (2011), Wolitzky (2014), and Ali and Miller (2014).

Lippert and Spagnolo (2011) consider environments with local interaction and separable stage games. In particular, they focus on stage games in which every pair of players plays an asymmetric Prisoners' Dilemma in which the interaction network may be direct, but is necessarily common knowledge. In this setup, they first consider two benchmark cases: public monitoring (when each agent observes the full history of play); and local monitoring. The main focus, however, is a variant of the local monitoring model in which players have access to a fixed number of rounds of private cheap talk in every period of the game. The communication network coincides with the interaction

network (as was implicitly the case in Renault and Tomala 1998; Cho 2010; and Laclau 2012).

With cheap talk, the possibility of transmitting soft information about privately observed defections to other agents may foster cooperation in games with fixed discount factors. Grim trigger strategies, which are optimal (in the sense of Abreu 1988) under public monitoring and which correspond to the contagion strategies studied in random matching games, are no longer optimal when information transmission is endogenous and players account for their incentives to communicate truthfully. When cooperation in the network is disciplined by such strategies, cheap talk is never used in equilibrium, as an agent reverts to non-cooperative play forever after observing a defection. This triggers a contagious process that eliminates all prospects of future cooperation in the network, thereby removing any motive for truthful communication. When forgiving strategies are used, instead, agents do have incentives to transmit information truthfully to avoid the collapse of cooperation as upon observing a defection, non-defecting agents continue cooperating and spread information on the deviation until only the initial deviator can be punished by a neighbor who benefits from such punishment. As information transmission within the network speeds up punishment phases, forgiving equilibria strictly dominate contagious equilibria.

Another central finding of the analysis is that with asymmetric stage-games, interaction networks display a rather general end-network effect that occurs under any informational assumption. Network structures such as trees may not sustain cooperative behavior, as agents with only outgoing links cannot be sanctioned if they defect. This end-network effect is a special case of gatekeeping and characterizes those gatekeepers as key players to cooperation in the network. Circular networks overcome this problem, ensuring that all defections can be met with punishment and that networks of relations are sustainable in equilibrium. The results provide an intuitive explanation for the importance of "closure" and "density." When monitoring is local and agents play according to grim trigger strategies, the enforceability of cooperation in bilateral relationships may hinder global cooperation in the larger networks, as a pair may not be willing to sacrifice their bilateral relationship to be part of the multilateral punishment mechanism which could sustain cooperation in the larger network. This argument extends from bilateral relations to larger subnetworks and establishes why coalitional agreements may undermine global cooperation ones by softening third-party punishments. This problem, however, can be overcome by forgiving strategies.

Wolitzky (2014) analyzes games with fixed discounting under different communication protocols. His main contribution establishes a direct relationship between different communication technologies and the set of sequential equilibrium payoffs. Results apply to separable stage games with local interaction, in which monitoring is local and imperfect, and in which the interaction network is common knowledge. In particular, the actions chosen by the two players in a relationship determine a signal realization pinning down payoffs in that relationship. Signals are random variables that depend only on the actions chosen by the two players on a link (and are thus independent

across relationships). Signals are locally public, but local monitoring is imperfect, as players observe only their action and the signal realizations in the interactions to which they belonged (but not necessarily the actions of their opponents). As the stage game is separable, the study aims at characterizing the community enforcement for a given discount factor. Results apply to stage games which possess a mutual minmax Nash equilibrium in every realized interaction.

The analysis first establishes that different communication protocols replicate any sequential equilibrium of a corresponding game with public information. The public information benchmark analyzed here is one in which all players observe the signal realizations on every link (but not necessarily the actions chosen by players other than themselves). In every period of the game, communication is modelled as an infinite number of rounds in which messages can be sent. Three communication technologies are considered: public cheap talk, private cheap talk, and tokens. The first result extends contributions in Ben-Porath and Kahneman (1996), and establishes that any equilibrium payoff of the game with public monitoring is also an equilibrium payoff of a corresponding game with local monitoring and public cheap talk. The second result builds on the contribution by Renault and Tomala (1998), and considers environments in which cheap talk is private and constrained to take place only on the interaction network (that is, when the interaction and communication networks coincide). The result establishes that any public monitoring equilibrium payoff is also an equilibrium payoff of a game with local monitoring and private cheap talk if and only if the network is 2-connected. The main departures from Renault and Tomala (1998) are: (a) that 2-connectedness is not only sufficient, but also necessary (in that for any network that is not 2-connected there exists a game in which private cheap talk cannot replicate public monitoring); and (b) that 2-connectedness is sufficient for replication even when the frequency of interaction is low. A final replication result considers environments with private cheap talk in which tokens can be exchanged in every relationship at each communication round. The main difference between cheap talk and tokens is that players must own tokens before transferring them. Results establish that public monitoring outcomes can always be replicated as sequential equilibria with tokens. Although the equilibrium construction relies both on tokens and private cheap talk, the same conclusions would hold if cheap talk were ruled out, since infinitesimal amounts of valueless tokens could be used to communicate. Message spaces and monetary endowments need not to be tailored to the specific game provided that a spanning tree exists in which all non-leaf players have a positive token endowment.

The final contribution presents sufficient conditions for tokens to expand the set of equilibrium payoffs compared both to games without communication, and to games with private cheap talk. Sufficient conditions require: (a) the network to possess a subtree; (b) every game played by two linked players to have a product structure; (c) the set of public information equilibria to include the convex hull of the locally public equilibria of the game with private information. These conditions simplify in many common environments, and only require the existence of a small subtree in which a strategy with tokens expands the equilibrium set. The essentiality of tokens then follows

since tokens expand the equilibrium payoff hull in the entire game when they do so in a subtree (as the remaining players can always comply with a strategy with private cheap talk in which tokens play no role).

The analysis of tokens builds on and is closely related to the literature on microfoundations of money. One of the most important themes in that literature asks when letting individuals exchange inherently valueless tokens can expand the equilibrium payoff sets in dynamic decentralized economies; for instance, Kocherlakota (1998, 2002). Results here carry out a similar exercise in the context of a more general setting.

Ali and Miller (2014) analyze the same environments discussed in their 2013 paper (in Section 6.4) while allowing for pre-play communication. In particular, before selecting stakes, partners may communicate to their neighbors information about the behavior of other players. The analysis studies both evidentiary communication (when players can conceal information, but cannot falsify it), and cheap talk. The analysis focuses on ostracism strategies in which players target punishments toward defecting players while cooperating with those they believe to be cooperative. To understand the impact of strategic communication, the analysis first characterizes two classical benchmarks. The first is bilateral enforcement, which identifies equilibria that abstract from community enforcement or communication (which in this setup amounts to bilateral grim trigger strategies played independently in each relationship). The second benchmark is mechanical communication, which characterizes settings in which players are constrained to reveal all their information truthfully. Permanent ostracism is an equilibrium with mechanical communication, since defectors must reveal themselves as such in all their future interactions. As permanent ostracism employs the harshest feasible punishment against defectors, it supports at least as much cooperation as any other equilibrium, and it coincides with the most cooperative equilibrium of a model with public monitoring.

When communication is strategic, one may conjecture that, while defecting players have a strong incentive to conceal their own misdeeds, cooperating players should have aligned interests in revealing and punishing the guilty. The main result establishes that this intuition is wrong. If defecting players are permanently ostracized, then their victims have a strong incentive to conceal such defections and to defect on other cooperating players. This strategic motive implies that permanent ostracism cannot be optimal with strategic communication and that the players are no better off than under bilateral enforcement. In other words, truthful communication is incentive compatible with permanent ostracism only if community enforcement is redundant. This stark negative conclusion applies to every network, even when communication is evidentiary. In fact, consider a permanent ostracism equilibrium and a relationship between two neighbors. Suppose by contradiction that they cooperate at stakes that would not be attainable under bilateral enforcement. Each player's incentives to cooperate must then be driven by the threat of punishments from others. Now consider a history at which one of them knows that everyone, except the two of them, has defected and should be ostracized. Because all the other players are defecting, this player's only incentive to cooperate arises from his continuation play with the one cooperative neighbor he

has left, just as under bilateral enforcement. Thus, he must strictly prefer to conceal his information and to defect at the equilibrium stakes, rather than telling the truth and reducing their stakes to keep on cooperating.

This result is most pronounced in Prisoners' Dilemmas, but analogues apply to general separable stage games. In any symmetric permanent ostracism equilibrium, each player's equilibrium payoff in a relationship is bounded above by the highest payoff attainable in a bilateral enforcement equilibrium in that relationship. Asymmetric permanent ostracism equilibria allow for more flexibility, but a bound on payoffs, arising from bilateral enforcement equilibria, still applies regardless of the network structure. Thus, the incentives to conceal information generally constrain the surplus that can be attained through permanent ostracism.

The negative theoretical conclusion on permanent ostracism contrasts with the prevalence of ostracism in communities and markets. Observed community enforcement norms, however, often involve forgiveness, in that players are only ostracized temporarily. The analysis provides a rationale for such norms by showing how forgiveness may encourage truthful communication between cooperative victims. In particular, when ostracism is temporary and players are forgiven at random times, innocent players communicate truthfully and cooperate with each other at levels beyond those attainable under bilateral enforcement (if players are sufficiently patient or society is sufficiently large). Temporary punishments may thus facilitate community enforcement by maintaining social collateral that fosters communication and cooperation among non-defecting players in the wake of defections.

The results on communication and ostracism should be contrasted with community enforcement schemes without information transmission, such as contagion equilibria introduced for anonymous random matching environments by Kandori (1992) and Ellison (1994), and applied to social networks by Wolitzky (2013), Ali and Miller (2013), and others. Contagion offers a useful benchmark for attainable payoffs in the absence of institutions or communication, but it also represents a fragile form of collective reputation, in that a single defection destroys a player's trust in the entire community. Ostracism, by contrast, reflects the principle that players ought to trust those partners who have never defected to their knowledge, while punishing those who have done so. Thus, with ostracism, reputations are entirely at the individual level. Hybrid community enforcement norms can be envisioned in which cooperative players communicate truthfully to other cooperators while ostracizing those who have defected in the past so long as they know of no more than d defecting players, and defect on all their partners otherwise. Such equilibria improve upon permanent ostracism, but average stakes are bounded by contagion with $n - d$ players.

Results in Ali and Miller (2014) rely on several modeling assumptions and innovations. Players interact at random privately observed times, which contrasts with classical repeated games in which all players are known to have interacted in every period. This generates non-trivial incentives at the communication stage, as players may now conceal an interaction from their partners. Incentives would differ if the timing of interactions were public. Unraveling would compel a player to reveal all details of

his past interactions, since a partner could rationally consider his failure to disclose as evidence of a deviation. If so, strategic communication would be as effective as permanent ostracism with mechanical communication. However, equilibria would be fragile, and even the slightest chance of interactions happening at privately observed times would again undermine any incentive for truthful communication.

The variable stakes model allows for a tight comparison of equilibria at a fixed discount rate and offers more scope for cooperation. Prisoners' Dilemma games with fixed stakes partly obscure incentives to ostracize by limiting the extent to which players can tailor their actions to the environment. In fact, if the stakes in each relationship were fixed, permanent ostracism would do no better than bilateral enforcement (as players, who are unwilling to cooperate under bilateral enforcement, would be unwilling to cooperate when only two cooperating players remain). In contrast, variable stakes enable partners to adjust the terms of their relationship based on their mutual history (for instance, by reducing their stakes once some players have been ostracized), shifting focus from technological constraints to the incentives for truthful communication.

6.6 COMMENTS: APPLICATIONS AND OMISSIONS

Applications: The use of implicit social sanctions to deter misconduct has been widely documented in economics (Milgrom, North, and Weingast 1990; Greif 1993), political science (Ostrom 1990; Fearon and Laitin 1996), sociology (Coleman 1990; Raub and Weesie 1990), and law (Bernstein 1992). Some of these studies have stressed the importance of community cohesion for attaining socially desirable outcomes in trust-based transactions, for example, Coleman (1990), Greif (1993), McMillan (1995), Fearon and Laitin (1996), Uzzi (1996), Dixit (2006). Coleman seminal's contribution identifies a notion of social capital, and relates such notion to the underlying social architecture. In Coleman's findings, the enforcement of cooperation is more effective in networks with high closure and cohesion, as cohesion facilitates the implementation of social sanctions, thereby increasing welfare. Other studies highlight the importance of information dissemination within the community for the effectiveness of such community-based sanctions. Greif (2006) finds that contract enforcement between medieval Maghribi traders is effective only when a close-knit community disseminates information so to align its members' incentives to comply with the community-based sanctions against deviant behavior.

Coleman's notion of social capital has motivated many of the more applied theoretical contributions in this field. For instance, Vega-Redondo (2006) considers a novel approach to network formation in the context of a repeated binary-symmetric Prisoners' Dilemma with random payoffs. The social network specifies not only the local interaction structure, but also the diffusion of information about past play, and

the availability of new cooperation opportunities. Search plays an important role in this environment, as agents always look for new partners when relationship-specific payoffs are volatile. In this context, the analysis develops a notion of social capital and shows how the social network adapts to changes in the environment. Network effects are important in enhancing cooperation; and the social network endogenously adapts by displaying more cohesiveness whenever the environment deteriorates. Conclusions are obtained by numerical simulations and supported by approximate mean-field analysis.

More recently, Balmaceda and Escobar (2014) build on results from Haag and Lagunoff (2006, 2007), discussed in Section 6.4, to show that cohesive communities (in which players are partitioned into isolated cliques) emerge as welfare-maximizing network structures. Cohesive communities generate local common knowledge, which allows players to coordinate their punishments, and, as a result, yield high equilibrium payoffs. Results provide an additional theoretical rationale for Coleman's link between cohesion and social capital, but apply only to environments in which monitoring is local, while interactions are centralized (in that all community members interact with a single player who knows the full history of play). The analysis also establishes that optimal networks are minimally connected, when players monitor every other community member in their component of the social network. If so, as in Burt (1992, 2001), bridging structural holes in the monitoring network becomes the sole consideration identifying the optimal social network (as cohesion within a component is imposed by assumption).

Other recent studies have theoretically analyzed and empirically documented the impact of network structure on different kinds of cooperation, such as favor exchange (Möbius 2001; Hauser and Hopenhayn 2004; Karlan et al. 2009; Jackson, Rodriguez-Barraquer, and Tan 2011) and risk-sharing (Ambrus, Möbius, and Szeidl 2010; Bramoullé and Kranton 2007; Bloch, Genicot, and Ray 2008). These studies are survey and discussed in Chapter 28 of this handbook. Although much empirical work remains to be done, empirical findings hint at different measures of centrality as determinants of cooperation within social interactions. For example, Karlan et al. (2009) find that indirect network connections between individuals in Peruvian shantytowns support lending and borrowing, consistent with findings showing that more central players cooperate more. More subtly, Jackson, Rodriguez-Barraquer, and Tan (2011) find that favor-exchange networks in rural India exhibit high support (the property that linked players share at least one common neighbor).

Endogenizing Networks: General results on network formation are discussed in several chapters of this handbook (Chapters 5–7). Most studies on repeated interactions have focused on optimal network design, rather than network formation, as in a repeated setup many well-documented network formation games generate large multiplicity of equilibrium networks (often including efficient networks). To see this, consider a pairwise linking process in which players simultaneously propose the partnerships they wish to engage in, and in which a partnership forms if and only if both players propose it. Consider a Prisoners' Dilemma game, in which the formed network is

common knowledge. It is straightforward to see that any network G can arise in an equilibrium of this game if it yields an individually rational net-payoff to each player, via the following strategy profile: if network G arises, then players follow the prescribed equilibrium, but if any other network forms then each player perpetually defects. This simple punishment deters players from deviating in the network formation stage. A similar logic applies to more complex games in which the network may not be common knowledge since any link remains local common knowledge among the two neighbors.

Separating Monitoring from Interaction: Most studies analyze environments in which the monitoring network and the network of interactions coincide (as was the case in the baseline setup presented in Section 6.2). However, conclusions generally carry over to the case in which players monitor more individuals than they interact with (as payoffs in any interaction can always be set to zero). Models with local monitoring and global interaction have only been analyzed in a limited number of studies which include Renault and Tomala (1998), Laclau (2012), and Wolitzky (2013).

Omissions: Some notable contributions to the literature have been omitted from the main discussion to streamline exposition. Ahn (1997) and Ahn and Souminen (2001) are precursors to several subsequent, but more general, contributions. They analyze cooperation in the context of binary-symmetric seller-buyer games with local monitoring and cheap talk, and present somewhat strong conditions for efficient outcomes to obtain. Kinateder (2008) considers a particular Prisoners' Dilemma game with global interaction, local monitoring, and in which players can truthfully communicate information to neighbors over time. The Folk Theorem extends to this setup, although the set of sequential equilibria and the corresponding payoff set may be reduced for discount factors strictly below 1. If players are allowed to communicate strategically, truthful communication arises endogenously only under additional assumptions. An additional implication of his analysis is that, when the discount factor is below 1, the viability of cooperation depends on the network's diameter, but not on its clustering coefficient. Mihm, Toth, and Lang (2009) consider strategic interaction in separable stage games with local monitoring. Their main contributions establish why strategic interdependencies between relationships on a network may facilitate efficient outcomes, and derive necessary and sufficient conditions to characterize the efficient equilibria of the network game in terms of the architecture of the underlying network.

Large Bipartite Networks: More recently, two studies have considered a novel and interesting approach to analyzing repeated networked games with a large number of players, namely Fainmesser and Goldberg (2011), and Fainmesser (2012). Fainmesser and Goldberg (2011) analyze repeated games in large bipartite networks with local monitoring and incomplete information about the network structure (players are informed of their neighbors and of several additional characteristics about the underlying graph). The model characterizes networks in which each agent cooperates in some equilibrium with every client to whom he is connected. To this end, the analysis establishes that in the proposed game: (a) the incentives of an agent to cooperate depend only on her beliefs with respect to her local neighborhood (a subnetwork

whose size is independent of the size of the entire network); and (b) when an agent observes the network structure only partially, his incentives to cooperate can be calculated as if the network was a random tree with him at its root. The characterization sheds light on the welfare costs of relying only on repeated interactions for sustaining cooperation, and on how to mitigate such costs. Fainmesser (2012) builds on this analysis by considering buyer-seller games in large bipartite networks, in which sellers have the option to cheat their buyers, and buyers decide whether to repurchase from different sellers. While endowing sellers with incomplete knowledge of the network, the analysis derives conditions that determine whether a network is consistent with cooperation between every buyer and seller that are connected. Three network features reduce the minimal discount factor sufficient for cooperation: moderate and balanced competition, sparseness, and segregation.

REFERENCES

Abreu, Dilip J. (1988). "On the theory of infinitely repeated games with discounting." *Econometrica* 56(2), 383–396.

Ahn Iltae (1997). "Three essays on repeated games without perfect information." PhD Thesis, University of Pennsylvania.

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

Ali, Nageeb S. and Miller David A. (2013). "Enforcing cooperation in networked societies." Mimeo.

Ali, Nageeb S. and Miller David A. (2014). "Ostracism." Mimeo.

Aliprantis, Charalambos D. and Border Kim C. (2005). *Infinite Dimensional Analysis*. Springer.

Balmaceda Felipe and Escobar Juan F. (2014). "Trust in cohesive communities." Mimeo.

Ben-Porath Elchanan and Kahneman Michael (1996). "Communication in repeated games with private monitoring." *Journal of Economic Theory* 70(2), 281–297.

Bernstein Lisa (1992). "Opting out of the legal system: Extralegal contractual relations in the diamond industry." *Journal of Legal Studies* 21(1), 115–157.

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

Bramoulle Yann (2007). "Anti-coordination and social interactions." *Games and Economic Behavior* 58(1), 30–49.

Bramoulle Yann and Kranton Rachel (2007). "Risk-sharing networks." *Journal of Economic Behavior and Organization* 64, 275–294.

Burt, Roland S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, Roland S. (2001). "Structural holes versus network closure as social capital." *Social Capital: Theory and Research* Chapter 2, first issue, 31–56, https://books.google.co.uk/books?hl=en&lr=&id=u_KTkBHY_kgC&oi=fnd&pg=PR7&dq=Burt+Structural+Holes+versus+Network+Closure+as+Social+Capital,+Social+Capital:+Theory+and+Research,+31-56,+2001.&ots=PLXeYmfxMO&sig=W2shYlshCP_1ftF0YhBepv9sdBc#v=onepage&q&f=false.

Cho Myeonghwan (2011). “Public randomization in the repeated prisoner's dilemma game with local interaction.” *Economic Letters* 112(3), 280–282.

Cho Myeonghwan (2010). “Cooperation in the Prisoner's Dilemma game with local interaction and local communication.” Mimeo.

Coleman, James S. (1990). *Foundations of Social Theory*. Cambridge, MA: Belknap Press.

Deb Joyee (2014). “Cooperation and community responsibility: A Folk Theorem for random matching games with names.” Mimeo.

Dixit Avinash (2006). *Lawlessness and Economics. Alternative Modes of Governance*. Oxford: Oxford University Press.

Ellison Glenn (1994). “Cooperation in the Prisoner's Dilemma with anonymous random matching.” *Review of Economic Studies* 61(3), 567–588.

Fainmesser, Itay P. (2012). “Community structure and market outcomes: A repeated games in networks approach.” *American Economic Journal: Microeconomics* 4(1), 32–69.

Fainmesser, Itay P. and Goldberg David A. (2012). “Cooperation in partly observable networked markets.” Mimeo.

Fearon, James D. and Laitin David D. (1996). “Explaining interethnic cooperation.” *American Political Science Review* 90(4), 715–735.

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

Greif Avner (2006). *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge, UK: Cambridge University Press.

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

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

Harrington, Joseph E. (1995). “Cooperation in a one-shot Prisoners' Dilemma.” *Games and Economic Behavior* 8(2), 364–377.

Hopenhayn, Hugo A. and Hauser Christine (2004). “Trading favors: Optimal exchange and forgiveness.” Meeting Papers 125, *Society for Economic Dynamics*.

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

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

Karlan Dean, Mobius Markus, Rosenblat Tanya and Szeidl Adam (2009). “Trust and social collateral.” *Quarterly Journal of Economics* 124, 1307–1361.

Kinateder Markus (2009). “Repeated games played on a network.” Mimeo.

Kocherlakota, Narayana R. (1998). “Money is memory.” *Journal of Economic Theory*, Elsevier, 81(2), 232–251.

Kocherlakota, Narayana R. (2002). “The two-money theorem.” *International Economic Review* 43, 333–346.

Laclau Marie (2012). “A Folk Theorem for repeated games played on a network.” *Games and Economic Behavior* 76(2), 711–737.

Laclau Marie (2014). “Communication in repeated network games with imperfect monitoring.” *Games and Economic Behavior* 87(C), 136–160.

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

McMillan John (1995) "Reorganizing vertical supply relationships." *Trends in Business Organization*, Tübingen: Mohr, ISBN 3161463919, 203–222.

Mihm Maximilian, Toth Russell and Lang Corey (2009). "What goes around comes around: A theory of indirect reciprocity in networks." Mimeo.

Milgrom, Paul R., North Douglas C. and Weingast Barry R. (1990). "The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs." *Economics and Politics*, 2(1), 1–23.

Mobius Markus (2001). "Trading favors." Mimeo.

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

Ostrom Elinor (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge, UK: Cambridge University Press.

Pecorino Paul (1999). "The effect of group size on public good provision in a repeated game setting." *Journal of Public Economics* 72, 121–134.

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

Renault Jerome and Tomala Tristan (1998). "Repeated proximity games." *International Journal of Game Theory* 27(4), 539–559.

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

Tomala Tristan (2010). "Fault reporting in partially known networks and Folk Theorems." *Operations Research* 59(3), 754–763.

Uzzi Brian (1996). "The sources and consequences of embeddedness for the economic performance of organizations: The network effect." *American Sociological Review* 61(4), 674–698.

Vega-Redondo Fernando (2006). "Building social capital in a changing world." *Journal of Economic Dynamics and Control* 30(11), 2305–2338.

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

Wolitzky Alexander (2015). "Communication with tokens in repeated games on networks." *Theoretical Economics* 10(1), 67–101.

Xue Jun (2004). "Essays on cooperation, coordination, and conformity." PhD Thesis, Pennsylvania State University.