Cooperation in Large Societies[*]

Alexander Wolitzky

MIT

July 28, 2022

Abstract

I survey models of cooperation in large populations. Topics include repeated games with public monitoring, random matching games, games on networks, infinite-population games, and models of legal or coercive enforcement. I emphasize how the scope of cooperation depends on the structure of social interactions, the information available to agents, and the feasibility of different types of rewards and punishments.

1 Introduction

Economists and other social scientists have long recognized that trust, cooperation, and reciprocity are critical determinants of societal outcomes. But what determines society's prospects for supporting cooperative behavior? Following seminal contributions by scholars such as Olson (1965), Williamson (1985), Ostrom (1990), North (1990), and Greif (1993), this question has been studied from a variety of perspectives, including (and often combining) game theory, historical and anthropological research, empirical economic analysis, and lab and field experiments. The current paper surveys the game theory literature in this area, focusing on repeated games and related models of long-run interaction.

I begin in Section 2 by discussing three empirical cases where societal cooperation is supported by long-run incentives: the provision of schools and wells in rural western Kenya (Miguel and Gugerty, 2005), informal lending and borrowing in Peruvian shantytowns (Karlan et al., 2009), and large-scale decentralized exchange on eBay (Resnick and Zeckhauser, 2002; Tadelis, 2016). These cases illustrate some real-world situations that repeated game models can shed light on—although the models I will focus on are not specifically tailored to these settings. They also highlight the wide range of conceptual issues that our models seek to understand, such as the structure of social interactions, the role of information, and the potential sanctions that are available to the community.

The bulk of this paper surveys the repeated game literature on cooperation in large societies. Traditional repeated game models are typically motivated by situations with a small number of relatively sophisticated players and "centralized" interactions and information, such as collusion in Cournot oligopoly (Green and Porter, 1984) or relational contracting (Bull, 1987; MacLeod and Malcomson, 1989; Levin, 2003). As we will see, the usual assumptions and results in this literature must be substantially modified to account for settings with a large number of diverse participants and "decentralized" interactions and information, like those discussed in Section 2. Nonetheless, it is useful to start by briefly reviewing the standard repeated game model, as well as some recent results for the large-population case—where, for example, players may have long time-horizons or high discount factors, but the discount factor may not be high "in comparison to the population size." As I show in Section 3, results in this case are quite different from standard folk theorems.

Section 4 covers the simplest and most canonical model of "decentralized" repeated games: games with uniform random matching. I offer a somewhat revisionist account of these games. Classic results in this area suggest that communities can sometimes rely on the threat of collective sanctions to sustain cooperation despite very limited monitoring. In contrast, I will argue that in a slightly richer and more realistic model that allows for a small amount of incomplete information, community enforcement instead requires personalized sanctions that rely on extensive information and communication.

Interactions in large groups are typically not only decentralized, but also non-random. Models with interaction structures other than uniform random matching are surveyed in

Section 5, under the heading of "repeated games on networks." Some of these models literally assume a fixed network structure of interactions or information flows, while others allow more flexible forms of non-random matching (the latter case is, however, relatively understudied). Here it is interesting to consider how standard results from fixed-partner repeated games or random matching games change when network structure is introduced. In addition, one can also ask new, network-specific questions, such as how network position determines which players take more cooperative actions.

In settings like E-commerce or modern credit and insurance markets, the population size is effectively infinite. There are essentially two ways to support cooperation in infinite populations: individual pairs or finite subsets of players can interact repeatedly (as in, e.g., long-term employment relationships), or players can have access to some centralized source of information about their current partner (as in, e.g., online ratings systems). I consider such models in Section 6. A theme here is that the scope for cooperation often depends on the type of information that players have access to—for example, information about one's partner's past actions only, as opposed to richer information about the context of these actions, or about the partner's "social standing."

Finally, Section 7 discusses a few models that go beyond the standard repeated game framework by introducing some kind of "institution" that can help support cooperation, either by providing some new type of information or by introducing new sanctioning instruments. These models aim to interpolate between repeated game models of informal enforcement and more traditional economic models, which implicitly assume that agreements are enforced by a powerful legal system. While it is still in its infancy, this approach may someday yield a more unified view of enforcement across different institutional settings.

The questions of how large societies support cooperation—and how researchers can productively model this behavior—are inherently interdisciplinary. This survey focuses on work by economists, but there are also related strands of research in other fields. The most active of these is probably the literature on "indirect reciprocity" in evolutionary biology (Alexander, 1987; Nowak and Sigmund, 1998; see Sigmund, 2010, for a survey). The models considered in this literature are essentially repeated games; however, this literature tends to focus on comparing the performance of a few pre-specified strategies (*Always Cooperate*,

Always Defect, Tit-for-Tat, etc.) and the resulting evolutionary dynamics, rather than analyzing the set of equilibria when players can choose any strategy. Evolutionary dynamics in this type of model have also been studied by physicists (Szabó and Fáth, 2007). Models of large-group cooperation also play a role in sociology (Granovetter, 1985; Coleman, 1990), political science (Axelrod, 1984; Keohane, 1984; Fearon and Laitin, 1996), anthropology (Boyd and Richerson, 1992; Henrich, 2004), law (Ellickson, 1991), and computer science (Friedman and Resnick, 2001).

2 Case Studies

I begin by reviewing some real-world examples of informal (repeated game) enforcement of public goods provision and bilateral exchange in large groups. These cases serve to (i) illustrate some applications we can bear in mind when considering the subsequent models, (ii) show how repeated game-type modeling can help guide empirical work in this area, and (iii) raise several theoretical issues that have not yet received much formal attention.

2.1 Public Goods in Rural Kenya

Why do some groups provide high levels of public goods, while others do not? A important empirical finding is that ethnically diverse groups generally provide fewer public goods (Easterly and Levine, 1997; Alesina, Baqir, and Easterly, 1999). This association is intriguing—Banerjee, Iyer, and Somanathan (2005) call it "one of the most powerful hypotheses in political economy"—but the underlying mechanism is controversial. Early explanations focused on commonalities of preferences or altruism among co-ethnics (Alesina and LaFerrara, 2000; Lutmer, 2001; Vigdor, 2004). However, some subsequent evidence favors an alternative interpretation based on repeated-game effects: it may be easier to apply sanctions for failing to contribute to public goods within ethnic groups rather than across groups, which leads to higher equilibrium contributions within groups.

Miguel and Gugerty (2005) argue for the importance of this social sanctioning mechanism in the context of public goods provision in rural Western Kenya. They find that ethnic diversity predicts not only lower primary school funding, worse school facilities, and worse

maintenance of water wells, but also fewer social sanctions and less verbal pressure directed at non-contributors, as recorded in the minutes of school committee meetings. For example, school committees in more ethnically diverse areas were less likely to exhort parents who were behind on their school fees to pay promptly, or to recommend that the village chief visit these parents to encourage payment. The authors' explanation for these findings is that "social sanctions and coordination are possible within groups due to the dense networks of information and mutual reciprocity that exist in groups but are not possible across groups."[1] However, their formal model simply assumes that sanctions are possible within groups but not across groups, rather than deriving this property as an implication of differences in network structure, information, or the existence of cooperative equilibria in a larger dynamic game. As we will see, explicitly modeling these features not only provides a foundation for the authors' assumption, but also leads to new, more detailed predictions, for example about how an individual's contributions may vary with her position in the social network.

Further evidence that causality runs from ethnic diversity to (lack of) social sanctions, and thence to (lack of) public goods comes from lab experiments run in Kampala, Uganda by Habyarimana et al. (2007). These authors find that subjects in a dictator game are equally generous toward co-ethnics and other participants when the recipient cannot observe the offerer's identity, but are substantially more generous toward co-ethnics when the offerer's identity is observable. The interpretation of the difference between the two treatments is that the non-anonymous treatment proxies for real-world situations where future punishment is possible (following Hoffman, McCabe, and Smith, 1996), so the results indicate that greater generosity toward co-ethnics is due to fear of punishment. Consistent with this interpretation, the authors obtain similar results in a prisoner's dilemma experiment where the possibility of punishment by a third-party is explicitly introduced. In contrast, little support is found for mechanisms based on common preferences or altruism within groups.[2]

---

[1]Besley, Coate, and Loury (1993) previously emphasized the importance of social sanctions in supporting collective action in their study of rotating savings and credit associations (Roscas).

[2]Additional experimental evidence consistent with Habyarimana et al.'s results come from Glaeser et al. (2000), who find that Harvard undergraduates with more mutual friends exhibit more trusting and trustworthy behavior; and from Leider et al. (2009) and Ligon and Schechter (2012), who use anonymous and non-anonymous variants of dictator games to distinguish between altruistic and repeated game motives for giving in different contexts.

2.2 Lending and Borrowing in Peru

Repeated game effects also play a role in informal contract enforcement, as in the large literature on risk-sharing without commitment (e.g., Coate and Ravallion, 1993; Udry, 1994; Foster and Rosenzwieg, 2001; Ligon, Thomas, and Worrall, 2002; Fafchamps and Lund, 2003). An example of such an analysis, which combines theory and empirics, is Karlan et al.'s (2009) study of "socially collateralized" lending and borrowing.[3] I will review Karlan et al.'s model in Section 5, but the key implication is that the maximum amount that individual i can credibly promise to repay to individual j—and hence the maximum amount that j is willing to lend to i—equals the "maximum flow" from i to j in a weighted social network, where the weight on the link between two individuals is the (exogenously given) value of maintaining their bilateral relationship beyond the current interaction. The interpretation is that players along a path from i to j can "vouch" that i will repay j, and that these relationships are lost if i fails to repay. The model is thus a reduced-form repeated game, where it is assumed that any deviation in a relationship causes the relationship to be permanently destroyed. This model makes strong predictions about the scope of lending and borrowing between any two agents. For example, since a path from i to j affects the maximum flow between these agents only through the weakest link in the path (i.e., the least valuable bilateral relationship), the model implies that the predictive value of adding a new path between i and j for lending and borrowing between them likewise depends only on the weakest link.

The authors test the model using survey data they collected from two shantytown communities with several hundred residents in Lima, Peru. Respondents were asked to name up to ten individuals with whom they spend the most time in an average week, as well as the set of friends from whom they borrowed money in the previous twelve months. The authors use the amount of time a pair of individuals spent together to proxy for the value of their relationship, and study how the resulting weighted social network predicts lending and borrowing. They find evidence in favor of some subtle predictions of the theory, such as a prediction that direct and indirect links matter roughly equally for borrowing, and a

---

[3] Other papers emphasizing the role of repeated game effects in supporting informal credit arrangements in developing economies include Karlan (2007), Feigenberg, Field, and Pande (2013), Giné and Karlan (2014), and de Quidt, Fetzer, and Ghatak (2016).

prediction that the importance of a path in the social network is largely determined by its weakest link. This is a nice example of how detailed, repeated game-type modeling can inform empirical work on cooperation in large groups.

2.3   Reputation Systems in Online Platform Markets

Informal enforcement is not confined to local communities like Kenyan villages or Peruvian shantytowns. In large online marketplaces like eBay or AirBnB, buyers and sellers can be separated by thousands of miles and often interact only once. These markets also feature a wide scope for opportunistic behavior, as when an eBay seller delivers a sub-standard product, or an AirBnB guest trashes an apartment. Such opportunism cannot be effectively constrained by the legal system, which involves itself in these transactions only in the most egregious cases of misbehavior. Instead, online marketplaces use centralized ratings systems to facilitate informal contractual enforcement: if individual i cheats individual j, j can give i a bad rating, which is observed by i's future potential partners and may discourage them from trading with i. A large literature (surveyed by Cabral, 2013, and Tadelis, 2016) documents the practical effectiveness of these ratings systems. Note that, as in the other case studies discussed in this section, the enforcement of cooperative behavior here is "multilateral," in that punishments are ultimately delivered by third parties. Just as a villager who fails to repay a loan may lose valuable relationships with individuals other than the lender, so an eBay seller who receives a bad rating may lose business from many potential buyers.

To successfully facilitate trade, online ratings systems must contend with several challenges, which interact with standard repeated-game considerations in interesting ways. For instance, an individual is usually free to quit a platform and then rejoin under a different name. This availability of "cheap pseudonyms" implies that an established agent's continuation payoff can never fall below that of a new entrant. Friedman and Resnick (2001) show that, if there is a positive chance of random "mistakes" in the players' interactions, this constraint imposes an upper bound on payoffs in the prisoner's dilemma with random matching, which applies for any discount factor. Payoffs close to this bound can be attained with simple "dues paying" strategies, where new entrants trade with established agents on unfavorable terms (e.g., new entrants always cooperate in the prisoner's dilemma, while es-

tablished agents who match with new entrants sometimes defect), which discourages entrants from cheating and then re-entering under a pseudonym.

Another important challenge for ratings systems is eliciting honest feedback. When a given pair of agents interact only once, why do they rate each other accurately, or even bother leaving a rating at all? A particular difficulty is that, if agents on both sides of the market rate each other, an agent who has been cheated may hesitate to give her partner a bad rating, for fear that the partner will retaliate by rating her badly in return. Bolton, Greiner, and Ockenfels (2013) find strong evidence of retaliatory behavior by sellers in eBay transactions in 2006 and 2007: ratings were highly correlated between the two parties to a transaction, and sellers disproportionately rated buyers shortly after buyers rated them, especially in instances where both parties left negative feedback. In 2008, eBay switched to a one-sided reputation system where only sellers are rated, likely in response to concerns about retaliatory rating (Tadelis, 2016).

3 Repeated Games with Many Players

Repeated game models are economists' main tool for understanding long-run cooperative relationships of the kind discussed in the previous section. Repeated games have been well-studied for many decades, and in principle can allow any number of players, N. However, some important assumptions and results that seem natural when N is small become unrealistic when N is large. For example:

1. Much research on repeated games concerns the limit case where players are very patient: that is, where the discount factor δ converges to 1. If we want to model cooperation among a large number of patient players, it is more natural to vary N and δ simultaneously, so that $\delta \to 1$ but $(1 - \delta) N$ may or may not converge to 0.

2. Classical repeated games often assume that players condition their behavior only on publicly available information. But information in large societies is typically decentralized. For example, each player may accurately observe her current trading partner's behavior, while knowing little about other agents' behavior.

3. Classical repeated games usually assume that all players are rational and play equilibrium strategies. But in large societies the possibility that at least a few players may be irrational cannot be ignored.

While I will eventually consider models that relax all of these assumptions, it is useful to start by quickly reviewing the classical repeated game model and some standard results, with a focus on the large-N case. This is done in Sections 3.1–3.3.[4] Section 3.4 then considers the implications of varying N and δ simultaneously in the classical model, which leads to quite different results. Subsequent sections introduce decentralized information (i.e., private monitoring) and irrational player-types.

3.1 The Standard Repeated Game Model

This section introduces the standard repeated game model and the accompanying notation. Readers familiar with repeated games may wish to skim or skip it.

A (finite) *stage game* $G = (I, A, u)$ consists of finite set of players $I = \{1, \ldots, N\}$, together with a finite set of actions A_i and a payoff function $u_i : A \rightarrow \mathbb{R}$ for each player $i \in I$. Here $A = \times_{i \in I} A_i$, and $u : A \rightarrow \mathbb{R}^N$ is given by $u(a) = (u_i(a))_{i=1}^N$ for each $a \in A$.[5] Payoff functions extend linearly to mixed actions: given a mixed action profile $\alpha \in \times_{i \in I} \Delta A_i$, we have $u(\alpha) = \sum_{a \in A} \Pr(a|\alpha) u(a)$. The *feasible payoff set* is

$$V = \operatorname{conv}\left\{v \in \mathbb{R}^N : v = u(a) \text{ for some } a \in A\right\},$$

where conv denotes convex hull. Player i's (mixed-action) *minmax payoff* is

$$\underline{v}_i = \min_{\alpha_{-i} \in \times_{j \in I \setminus \{i\}} \Delta A_j} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

A payoff vector is v is *individually rational* if $v_i \geq \underline{v}_i$ for all i. Thus, the set of feasible and

---

[4] See Mailath and Samuelson (2006) for a thorough textbook treatment of this material.

[5] In general, given a collection of sets $(X_i)_{i=1}^N$ with generic elements $x_i \in X_i$, their product is denoted $X = \times_{i \in I} X_i$ with generic element $x \in X$. I also write $X_{-i} = \times_{j \in I \setminus \{i\}} X_j$ with generic element $x_{-i} \in X_{-i}$.

individually rational payoff vectors is

$$V^* = \{v \in V : v_i \geq \underline{v}_i \text{ for all } i \in I\}.$$

For example, in the *prisoner's dilemma* (PD) stage game, given by

	C	D
C	$1, 1$	$-\ell, 1 + g$
D	$1 + g, -\ell$	$0, 0$

with $g, \ell > 0$ and $g < 1 + \ell$, each player's minmax payoff is 0, so $V^* = \text{conv}\left\{(0,0), \left(0, \frac{1+g+\ell}{1+\ell}\right), \left(\frac{1+g+\ell}{1+\ell}, 0\right), (1,1)\right\}$. I will return to this stage game repeatedly.

In a (discounted) repeated game, the stage game is played repeatedly in periods $t = 1, 2, \ldots$, players get some information before taking actions in each period, and players maximize discounted expected payoffs. Formally, a *repeated game* $\Gamma = (G, \delta, Y, p)$ consists of a stage game G, a discount factor $\delta \in (0, 1)$ (assumed for simplicity to be common among the players), and a *monitoring structure* (Y, p), which summarizes the players' information. Here, $Y = \times_{i=1}^{N} Y_i$, where Y_i is the set of possible *signal realizations* for player i, and at the end of each period t a signal $y_t \in Y$ is drawn according to the probability distribution $p(\cdot|a_t) \in \Delta Y$, where $a_t \in A$ is the period-t action profile. Note that this notation implicitly assumes that the monitoring structure remains the same in every period, and that the distribution of signals depends only on current-period actions. Special monitoring structures include *perfect monitoring*, where $Y_i = A$ for all $i \in I$ and $p(y|a) = \mathbf{1}\{y_i = a \ \forall i \in I\}$ (where $\mathbf{1}\{\cdot\}$ is the indicator function); and *public monitoring*, where $Y_i = Y_j$ for all $i, j \in I$, and $p(y|a) > 0$ implies $y_i = y_j$ for all $i, j \in I$. The general, non-public case is called *private monitoring*.[6]

A *history* $h_i^t \in H_i^t$ for player i at the beginning of period t takes the form $h_i^t = (a_i^t, y_i^t)$, where $a_i^t = (a_{i,0}, \ldots, a_{i,t-1})$ and $y_i^t = (y_{i,0}, \ldots, y_{i,t-1})$, with $a_i^0 = y_i^0 = \emptyset$. Letting $H_i =$

---

[6]It is sometimes assumed that in every period the players can also observe the outcome of a *public randomizing device*: e.g., a uniform $[0, 1]$ random variable, independent of actions and signals. None of the results surveyed here require this assumption, although allowing public randomization simplifies some proofs.

$\bigcup_{t=1}^{\infty} H_i^t$, a (behavior) *strategy* $\sigma_i \in \Sigma_i$ for player i is a mapping $\sigma_i : H_i \to \Delta A_i$. That is, players can condition their play on their own past actions and observations. Player i's discounted expected payoff under strategy profile i (measured in per-period terms) is $U_i(\sigma) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}^\sigma [u_i(a_t)]$, where $\mathbb{E}^\sigma [\cdot]$ denotes expectation under the probability distribution over histories induced by strategy profile σ.[7] As usual, a *Nash equilibrium* is a strategy profile σ where $U_i(\sigma) \geq U_i(\sigma_i', \sigma_{-i})$ for all $i \in I, \sigma_i' \in \Sigma_i$. A *subgame perfect equilibrium* of a repeated game with perfect monitoring is a strategy profile where the pre-scribed continuation play at any history forms a Nash equilibrium. Further refinements of Nash equilibrium, such as perfect public equilibrium and sequential equilibrium, are defined below as needed.

3.2 The Folk Theorem with Perfect Monitoring

Perhaps the best-known result on repeated games is the "folk theorem," which says that if players are patient, monitoring is perfect, and V^* has non-empty interior (as a subset of $\mathbb{R}^N$), then every payoff vector in $\text{int} V^*$ is attainable in subgame perfect equilibrium. This is a classical formalization of the idea that the "shadow of the future" can incentivize cooperation when it weighs heavily and when behavior is well-observed.[8] The folk theorem is an important benchmark result that we will return to repeatedly under different sets of assumptions. Of course, it may be possible to support some cooperative behavior in equilibrium even if the full conclusion of the folk theorem fails, and we will also see some settings where this occurs.

Theorem 1 (Perfect Monitoring Folk Theorem) *Assume that* $\text{int} V^* \neq \emptyset$. *For any* $v \in \text{int} V^*$, *there exists* $\bar{\delta} < 1$ *such that, for every* $\delta > \bar{\delta}$, *there is a subgame perfect equilibrium of the repeated game with discount factor* δ *that gives payoff vector* v.

Proving the perfect monitoring folk theorem at this level of generality is not trivial, but a simple case that suffices for our purposes arises when $v = u(a)$ for some pure action profile

---

[7] Formally, this can be defined recursively starting from $t = 1$.

[8] Early versions of the perfect-monitoring folk theorem were developed by several authors, including Friedman (1971), Aumann and Shapley (1976), and Rubinstein (1979). The version reproduced here is due to Fudenberg and Maskin (1986); see also Abreu, Dutta, and Smith (1994) and Proposition 3.8.1 of Mailath and Samuelson (2006).

$a \in A$, and v strictly Pareto-dominates a static Nash equilibrium payoff vector, v^{NE}. In this case (which was addressed by Friedman, 1971), the theorem is proved by considering "Nash-reversion" strategies that prescribe action profile a so long as a has always been taken so far, and that otherwise prescribe a static equilibrium that gives payoff vector v^{NE}. If we denote player i's maximum static deviation gain at action profile a by $d_i = \max_{a_i' \in A_i} u_i\left(a_i', a_{-i}\right) - v$, it is easy to see that these strategies form a subgame perfect equilibrium whenever

$$d_i \leq \frac{\delta}{1-\delta}\left(v_i - v_i^{NE}\right) \quad \text{for all } i \in I.$$

Observe that, so long as the deviation gain is bounded from above and the desired payoff vector v is bounded away from v^{NE}, the discount factor δ does not need to be close to 1 for Nash reversion strategies to form an equilibrium. Moreover, the required discount factor is independent of N. There is thus no obvious reason why it is more difficult to support cooperation in larger groups under perfect monitoring.[9] However, the perfect monitoring assumption—which is already a strong assumption when N is small—seems even more extreme when N is large. For example, it is hard to imagine that all actions are observed by everyone in the community in any of the settings discussed in Section 2. For this reason, most models of cooperation in large groups assume some form of imperfect monitoring.

3.3 The Folk Theorem with Imperfect Public Monitoring

The simplest form of imperfect monitoring is *public monitoring*, where all players observe the same signal. Fudenberg, Levine, and Maskin (1994) show that the folk theorem extends to imperfect public monitoring, if the monitoring structure satisfies appropriate statistical conditions.[10] To give one of their sufficient conditions, let Y_0 denote the set of possible public signal realizations, and for any mixed action profile $\alpha \in \times_{i \in I} \Delta A_i$ and any player

---

[9]This statement of course assumes that the players play an equilibrium, and furthermore coordinate on the desired equilibrium. One might wonder if the problem of coordinating on a desired equilibrium—what Bowles and Gintis (2013) call the "missing choreographer" problem—is more difficult in larger groups. The literature on large-population coordination games following Kandori, Mailath, and Rob (1993) and Young (1993) bears on this type of question in static games. I am not aware of such work in repeated games.

[10]Related ideas appear in Matsushima (1989), Fudenberg and Levine (1994), and Kandori and Matsushima (1998). Mertens and Parthasarathy (1987) and Abreu, Pearce, and Stacchetti (1990) develop important results on the structure of the equilibrium payoff set in repeated games with imperfect public monitoring.

i, define the $|A_i| \times |Y_0|$ matrix $R_i(\alpha_{-i})$ with elements $[R_i(\alpha_{-i})]_{a_i,y} = p(y|a_i, \alpha_{-i})$, and define the $(|A_i| + |A_j|) \times |Y_0|$ matrix $R_{ij}(\alpha)$ given by stacking the matrices $R_i(\alpha_{-i})$ and $R_j(\alpha_{-j})$. Say that action profile α has *pairwise full rank for players i and j* if $\text{rank} R_{ij}(\alpha) = |A_i| + |A_j| - 1$.[11] Intuitively, this condition says that any deviation from α by player i or player j is statistically identifiable, and that moreover deviations by player i and player j are statistically distinguishable from each other. Fudenberg, Levine, and Maskin show that, to establish the folk theorem for payoffs that Pareto-dominate a static Nash equilibrium, it suffices to assume that all Pareto-efficient pure action profiles have pairwise full rank for all pairs of players.[12] Their result concerns *perfect public equilibria*, which are profiles of strategies that depend only on the history of public signals y^t, and that form a Nash equilibrium conditional on any such history.

Theorem 2 (Imperfect Public Monitoring Folk Theorem) *Assume that* $\text{int} V \neq \emptyset$ *and that all Pareto-efficient pure action profiles have pairwise full rank for all pairs of players. For any $v \in \text{int} V$ that Pareto-dominates a static Nash equilibrium, there exists $\bar{\delta} < 1$ such that, for every $\delta > \bar{\delta}$, there is a perfect public equilibrium of the repeated game with discount factor δ that gives payoff vector v.*

This important result influenced research on problems such as collusion among oligopolists (Athey and Bagwell, 2001) and relational contracts in organizations (Levin, 2003). Two issues should be borne in mind regarding its implications for cooperation in larger groups. The first is that, as I have already noted, monitoring in large groups is often far from public, as is the case when interactions are decentralized and players observe only their current partners' behavior. I discuss models of this type in Sections 4–6.

The second issue is that, even under public monitoring, Theorem 2 does not address how the value of the discount factor δ and the "precision" of the monitoring structure (Y, p) that are required to provide incentives vary with the population size N. In particular, note that Theorem 2 makes no reference to the signals' statistical informativeness. Intuitively,

---

[11]Note that $|A_i| + |A_j| - 1$ is the maximum possible rank of the matrix $R_{ij}(\alpha)$, because $\alpha_i R_i(\alpha_{-i}) = [p(y|\alpha)]_y = \alpha_j R_j(\alpha_{-j})$.

[12]They also show that a minmax-threat folk theorem holds if in addition every pure action profile a has *individual full rank*, which means that the matrix $R_i(a_{-i})$ has full row rank for every player i.

even if $p(y|a) \approx 1/|Y_0|$ for all y—so that signals are almost completely uninformative—strong incentives can still be provided on the basis of these signals when δ is sufficiently large, because large changes in continuation payoffs can offset noisy signals. Of course, the less informative are the signals, the higher δ must be.[13] For example, if the number of possible signal realizations $|Y_0|$ is held fixed as N grows, the informativeness of the signals for distinguishing between deviations by different players necessarily decreases: while pairwise full rank may hold for any N, as N grows large the full row rank submatrices of the matrix $R_{ij}(\alpha)$ must become close to singular for some pairs of players (i, j). A corresponding increase in δ is then required to preserve incentives.

This discussion indicates that the scope for cooperation in large-population repeated games with imperfect public monitoring depends on a three-way interaction among N, δ, and "monitoring precision." This interaction is the subject of the next subsection (where some results will also apply under private monitoring).

3.4   Anti-Folk Theorems with Many Players

A series of papers by Green (1980), Sabourian (1990), Levine and Pesendorfer (1995), Fudenberg, Levine, and Pesendorfer (1998), al-Najjar and Smorodinsky (2000, 2001), and Pai, Roth, and Ullman (2016) provide conditions under which, as the population size $N \to \infty$ for a fixed discount factor δ, players behave approximately myopically in every repeated game equilibrium. These negative results concerning the prospects for cooperation ("anti-folk theorems") establish continuity at the infinite population limit, where players act myopically.

The key feature of these models is that, in a large population, most players' individual actions have only a small effect on the signal distribution, and hence on other players' behavior. This condition is directly assumed in Green and Sabourian's models. Subsequent models (e.g., Fudenberg, Levine, and Pesendorfer, 1998; al-Najjar and Smorodinsky, 2000, 2001) instead derive this feature from an assumption that signals are generated through *individual noise*: each player's action a_i stochastically determines an *individual outcome* x_i (independently across players), and the distribution of the signal profile y depends on the action profile a only through the profile of individual-level outcomes $x = (x_i)_{i=1}^N$. For exam-

---

[13]Kandori (1992a) formalizes this point.

ple, when there are two individual outcomes for each player, and each individual outcome occurs with probability at least $\varepsilon > 0$ due to noise, al-Najjar and Smorodinsky (2000) show that the average probability that a player is pivotal for the realization of a binary signal $y \in \{0, 1\}$ is maximized when each player "votes" yea and nay with probability ε each (and abstains with probability $1 - 2\varepsilon$), and the signal takes value 1 iff the yea votes exceed the nays. This scheme yields a pivot probability of order $1/\sqrt{\varepsilon(1-\varepsilon)N}$, the reciprocal of the standard deviation of the number of yea votes. Since the impact on a player's continuation payoff in the event that her action is pivotal is of order at most $1/(1-\delta)$, it follows that non-trivial incentives cannot be provided if $N \to \infty$ while δ and the number of possible signal realizations $|Y|$ are held fixed, or more generally if $(1-\delta)\sqrt{N} \to \infty$ while $|Y|$ is fixed.[14]

Sugaya and Wolitzky (2022b) study the three-way interaction among N, δ, and $|Y|$ in more detail, under the same ε-individual noise assumption as in Fudenberg, Levine, and Pesendorfer, and al-Najjar and Smorodinsky. Let $K = \log_2 |Y|$, so the signal can be viewed as a K-dimensional binary random variable. Their main result implies that if $(1-\delta)N/K \to \infty$ along any sequence of games where N, δ, and K vary simultaneously (subject to a uniform upper bound on stage-game payoffs $|u_i|$ and a uniform lower bound on individual noise ε), then only vanishingly weak incentives can be provided in any Nash equilibrium. To state this result formally, for any player i and any distribution over action profiles $\alpha \in \Delta A$, denote player i's maximum expected deviation gain at action distribution α by

$$d_i(\alpha) = \max_{\tilde{a}_i : A_i \to A_i} \sum_r \alpha(r)\left(u_i(\tilde{a}_i(r_i), r_{-i}) - u_i(r)\right).$$

In addition, for any number $\eta > 0$, say that a payoff vector $v \in \mathbb{R}^N$ is *consistent with η-myopic play* if there exists an action distribution $\alpha \in \Delta A$ such that $u(\alpha) = v$ and $\sum_i d_i(\alpha)/N \leq \eta$: that is, if the per-player average deviation gain at α is less than η.

Theorem 3 (Limits of Individual Incentives) *Fix any upper bound on stage-game payoffs and any lower bound on individual noise. For any $\eta > 0$, there exists $k > 0$ such that, for any repeated game satisfying these bounds as well as the condition that $(1-\delta)N/K > k$, any Nash equilibrium payoff vector in this repeated game is consistent with η-myopic play.*

---

[14]Fudenberg, Levine, and Pesendorfer (1998) derive a similar result.

Under the stronger hypothesis that $(1 - \delta)\sqrt{N/K} \to \infty$, the theorem's conclusion follows from probability results which imply that the average "influence" of N independent binary random variables on a K-dimensional binary random variable is of order at most $\sqrt{K/N}$, and that therefore the maximum per-player incentive that can provided on the basis of such a random variable is of order at most $\sqrt{K/N}/(1 - \delta)$. The hypothesis can be weakened to $(1 - \delta)N/K \to \infty$ by combining this type of probability result with recursive methods developed in Sugaya and Wolitzky (2022a).

Theorem 3 can be further generalized to apply for any monitoring structure where the *mutual information* between signals and actions is bounded by K, which is a weaker condition on the monitoring structure than the requirement than $\log_2 |Y| = K$. Conversely, a folk theorem holds for some sequence of monitoring structures satisfying $(1 - \delta)N\log(N)/K \to 0$: for example, it suffices to perfectly monitor the actions of K randomly selected players in every period.[15] These results imply that the scope for cooperation in many-player repeated games is determined by the ratio of the discount rate $(1 - \delta)$ and the number of "bits of information per player," K/N (up to the $\log(N)$ gap between the necessary and sufficient conditions for cooperation). Thus, a high degree of patience and/or a large amount of information are required to incentivize cooperation in large groups.

How restrictive is the condition that discounting is low relative to per-player information? The answer depends on the type of game under consideration. The condition is restrictive in situations where a large population of players are monitored by a fixed "aggregate signal" that does not scale with the population size. For example, this case arises when the signal is the market price facing Cournot competitors, the level of pollution in a common water source, or the output of team production. On the other hand, the condition is much less restrictive in situations where players are monitored "separately," in a manner that scales with the population size. For example, this situation arises in the random matching games studied in Section 4, where each player observes her partner's action, as well as in games where a centralized rating system records a signal of each player's action, as in the models studied in Section 6. In sum, in public-monitoring models such as Fudenberg, Levine, and Pesendorfer

---

[15]The basic logic is that if a player is monitored with probability K/N, she can then be provided with an incentive of "size" $K/(N(1 - \delta))$. The extra $\log(N)$ term provides some slack that is used in the equilibrium construction.

(1998), al-Najjar and Smorodinsky (2000, 2001), and Sugaya and Wolitzky (2022a), the obstacle to supporting cooperation is that societal information is *insufficiently precise*; while in the community enforcement models covered in Sections 4 and 6, the obstacle is instead that societal information is *disaggregated*.

While the condition in Theorem 3 may already be restrictive, matters are far worse if we restrict attention to *strongly symmetric equilibria*, where all players have the same continuation payoff at every history, so that rewards and punishments are provided only "collectively."[16] Collective incentives have the well-known shortcoming that they encourage free-riding in large populations. In the repeated game context, the following result (also due to Sugaya and Wolitzky, 2022b) formalizes this intuition by showing that incentives can provided in strongly symmetric equilibria only if $\delta \to 1$ almost as fast as $\exp N \to \infty$, which is an extremely restrictive condition.

Theorem 4 (Limits of Collective Incentives) *Fix any upper bound on stage-game pay-offs and any lower bound on individual noise. For any $\eta > 0$ and $\rho > 0$, there exists $k > 0$ such that, for any repeated game satisfying these bounds as well as the condition that $(1 - \delta) \exp\left(N^{1-\rho}\right) > k$, any strongly symmetric equilibrium payoff vector in this repeated game is consistent with η-myopic play.*

To understand this result, suppose we wish to enforce a symmetric pure action profile a at which the maximum static deviation gain equals η. Once we restrict ourselves to strongly symmetric equilibria, it is without further loss to restrict attention to "tail tests," where continuation values take one (low) value whenever the number n of players who take action a falls below a cutoff n^*, and another (high) value whenever $n \geq n^*$. Under ε individual noise, the number n follows a binomial distribution with mean $(1 - \varepsilon) N$ and standard deviation $\sqrt{\varepsilon (1 - \varepsilon) N}$. Thus, the probability that each player is pivotal for passing the tail test is maximized by setting $n^* \approx (1 - \varepsilon) N$, which yields a pivot probability of order $1/\sqrt{\varepsilon (1 - \varepsilon) N}$, as in Fudenberg, Levine, and Pesendorfer, or al-Najjar and Smorodinsky. Now, to enforce action profile a, the static deviation gain η must be less than the product

---

[16]Strongly symmetric equilibria are defined only in symmetric games. However, the same results hold for a class of "linear equilibria" in general games, where all players' continuation payoffs are required to lie on a line.

of the pivot probability (a term of order $1/\sqrt{N}$) and the maximum difference between any two continuation payoffs (a term of order $1/(1-\delta)$). Hence, if η is to remain bounded away from 0 as $N \to \infty$ and $\delta \to 1$, the product $(1-\delta)\sqrt{N}$ cannot explode.

The more restrictive condition of Theorem 4—that $(1-\delta)\exp(N^{1-\rho})$ cannot explode for any $\rho > 0$—follows because if $n^* \approx (1-\varepsilon)N$, then the tail test is failed with probability bounded away from 0 in every period, which leads to the destruction of all continuation value when $\delta \to 1$. To prevent this value destruction while still providing incentives, the *statistical score* of the tail test—that is, the ratio $\Pr(n = n^*)/\Pr(n \leq n^*)$—must increase linearly with the variance of n, and hence with N.[17] Since the distribution of n is approximately normal when N is large, and the density of the normal distribution decreases exponentially with the score,[18] it follows that $\Pr(n = n^*)$—the pivot probability—must decrease exponentially with N. Theorem 4 now follows because the product of the pivot probability and $1/(1-\delta)$ cannot explode if η is to remain bounded away from 0.[19]

To summarize, in large-population repeated games with imperfect monitoring, collective incentives cannot be provided unless δ is extremely (likely implausibly) high, and even individual incentives can be provided only if δ is high and/or signals are informative. While this is a rather negative result for settings where a large population of players is monitored by a fixed aggregate signal, sufficiently informative monitoring structures arise naturally in "decentralized" settings such as random-matching games or games on networks. These settings are the subject of the following two sections.

4   Repeated Games with Random Matching

First introduced by Rosenthal (1979), repeated games with random matching are a canonical model of cooperation in decentralized groups. There are several versions of the model, which can all be formally nested as special cases of the general repeated game model presented in Section 3.1. In this section, I consider games with *uniform random matching*, meaning that

---

[17]Sannikov and Skrzypacz's (2007) result that collusion in Cournot duopoly is impossible with frequent actions and Brownian noise depends on a similar observation; see also Fudenberg and Levine (2007, 2009).

[18]Because $\phi(z)/\Phi(z)$ is of order $|z|$, where ϕ and Φ are the standard normal pdf and cdf.

[19]Conversely, tail tests can support a Nash-threat folk theorem if there exists $\rho > 0$ such that $(1-\delta)\exp(N^{1+\rho}) \to 0$.

in each period the population of N players (here assumed to be even) breaks into pairs to play a 2-player stage game $\tilde{G} = \left(\tilde{A}, \tilde{u} \right)$, where each player's partner is selected uniformly at random from among the other $N - 1$ players.

Matching is *anonymous* if players choose actions without observing their current partner's identity, so that A (the action set in each period of the N-player repeated game) is the same as $\tilde{A}$ (the action set in the 2-player stage game). Matching is *non-anonymous* if players observe their current partner's identity before choosing actions, in which case $A = \tilde{A}^{(N-1)}$, the set of mappings from the partner's identity to the stage game action. Either way, I assume that each player perfectly observes her current partner's action at the end of each period, but learns nothing about the actions taken in other matches.[20] The model thus abstracts from monitoring imperfections within each match but completely rules out public signals, in stark contrast to the public monitoring structures considered in Section 3.

4.1   Anonymous Matching

Because players receive so little information under anonymous random matching, a natural first question is whether non-myopic play can be supported at all. Early results by Kandori (1992b), Ellison (1994), and Harrington (1995) show that cooperation can be supported in the PD by relying on a form of collective punishment called *contagion strategies*. Contagion strategies are simply grim trigger strategies adapted to anonymous random matching games: each player starts by playing C (cooperate), but switches to D (defect) whenever she observes a single play of D. Note that each player i is incentivized to play C on path (i.e., when she has observed only C so far) because playing D causes her current partner j to switch to D in the next period, which subsequently causes j's future partners to switch to D, and so on until the contagion of playing D spreads throughout the population, eventually reducing i's payoff. We obtain the following result of Kandori's.

Theorem 5 (Anonymous Cooperation in Nash Equilibrium) *In the repeated PD with anonymous random matching, there exists $\bar{\delta} < 1$ such that, for every $\delta > \bar{\delta}$, there is a Nash*

---

[20]Monitoring structures of this type, where each action is observed either perfectly or not at all, are sometimes called *semi-standard monitoring* (Lehrer, 1990), *partial monitoring* (Ben-Porath and Kahneman, 1996), or *network monitoring* (Wolitzky, 2013).

equilibrium where all players take C in every period along the equilibrium path.

This result depends on the order of limits between N and δ. For any N, contagion spreads throughout the population in finite time with high probability, so deviating from C to D is unprofitable whenever δ is sufficiently high (and hence contagion strategies form a Nash equilibrium, since C is prescribed at every on-path history). However, the amount of time that it takes for contagion to spread goes to infinity as $N \to \infty$, so for any $\delta < 1$ deviating from C to D is profitable when N is sufficiently high. Nevertheless, the joint condition on N and δ required for contagion strategies to form a Nash equilibrium is actually quite permissive. Since the diffusion of contagion is well-approximated by a logistic curve, the number of "infected" players increases exponentially until a substantial fraction of the population is infected, so it can be shown that contagion strategies form a Nash equilibrium whenever $(1 - \delta) \log N \to 0$.[21]

Another apparent limitation of Theorem 5 is that it applies only to Nash equilibrium. A more natural solution concept in this model is sequential equilibrium (Kreps and Wilson, 1982), which requires sequential rationality with respect to beliefs that arise as the limit of conditional probabilities derived from completely mixed strategies. The simplest version of contagion strategies—where switching from C to D is permanent—may be inconsistent with sequential equilibrium, because a patient player who observes D in period 1 (say) may benefit from playing C for several periods to slow the spread of contagion.[22] However, Ellison (1994) shows that contagion strategies can be modified to form a sequential equilibrium.[23] The key observation is that, under contagion strategies, switching to D is more attractive when a greater number of other players have already switched to D. Hence, if the severity of the punishment incurred by starting contagion is reduced to the point where a player is indifferent to starting contagion on path (i.e., when only C has been observed), then

---

[21]That contagion spreads approximately logistically in an N-player random matching model can be shown by arguments akin to those in Frieze and Grimmett (1985), who consider a similar process whereby a "rumor" spreads through a finite population in discrete time.

[22]If such a player believes that not only her period-1 partner but also all of the other $N - 2$ players in the population deviated to D in period 1, then playing D from period 2 on as prescribed by contagion strategies is sequentially rational. Since seeing one's partner take D in period 1 is off path, this belief is consistent with Bayes' rule, and hence with (weak) perfect Bayesian equilibrium. However, it is not consistent with the belief that the opponents "tremble" independently, and hence is not consistent with sequential equilibrium.

[23]Kandori (1992b) showed that standard contagion strategies already form a sequential equilibrium for some parameters.

spreading contagion is always optimal off path (i.e., once D has been observed). Ellison describes two ways of reducing the severity of contagion: occasionally restarting cooperation according to a public randomizing device, and conditioning play on calendar time in such a way that a given occurrence of contagion only periodically affects play (a technique now known as "threading"). While the realism of such strategies is debatable, they suffice to establish the following important result.

Theorem 6 (Anonymous Cooperation in Sequential Equilibrium) *In the repeated PD with anonymous random matching, there exists $\bar{\delta} < 1$ such that, for every $\delta > \bar{\delta}$, there is a sequential equilibrium where all players take C in every period along the equilibrium path.*

While suitably modified contagion strategies can be used to support cooperation as a sequential equilibrium in the repeated PD, it is not straightforward to extend these strategies to support non-myopic behavior in other games (or to support outcomes in the PD other than mutual cooperation, such as profiles where some players cooperate while others defect). For example, consider the *product choice game*, where player 1 produces a high or low quality product, player 2 buys or does not buy, and payoffs are given by

	B	$\neg B$
H	$1, 1$	$-\ell, 0$
L	$1 + g, -c$	$0, 0$

with $g, \ell, c > 0$. This is an asymmetric game, so let us suppose that there are separate, equal-sized populations of player 1's and player 2's, where every period each player 1 meets a random player 2. Suppose we try to support the outcome (H, B) with the threat of contagion to the Pareto-inferior Nash equilibrium $(L, \neg B)$: that is, each player 1 takes H so long as (H, B) is observed, and otherwise switches to L; and each player 2 takes B so long as (H, B) is observed, and otherwise switches to $\neg B$. These strategies form a Nash equilibrium whenever δ is sufficiently high, but modifying them to form a sequential equilibrium is challenging. To see the problem, suppose that some player 2 sees his partner deviate to L in period 1. By Kreps-Wilson consistency, this player 2 must believe no one other than his period 1 partner deviated in period 1, and that therefore (in a large population) his

period 2 partner is very likely to play H. But then player 2 should play B rather than $\neg B$ in period 2, since this yields a higher expected stage game payoff, as well as slowing contagion. In general, the problem is that since the "punishment" profile $(L, \neg B)$ is not in dominant strategies (in contrast to the punishment profile (D, D) in the PD), players who observe deviations do not want to switch to the punishment profile unless they believe that sufficiently many other players have already switched, and this belief may not be consistent under contagion strategies.

Deb and González-Díaz (2019) use a clever and intricate argument to extend contagion strategies to establish a partial folk theorem (in sequential equilibrium, for symmetric payoffs that Pareto-dominate a static Nash equilibrium) for a class of stage games that includes the product choice game. The key idea is to specify that deviations from the desired action profile (e.g., (H, B) in the product choice game) during a long "trust-building" phase at the beginning of the game do not spark contagion until after the end of the trust-building phase. Due to the existence of the trust-building phase, once players are supposed to switch to the punishment profile, they are free to believe that the earliest deviation actually occurred a long time ago, and that consequently many other players have already switched (or will switch at the same time). This belief then rationalizes the prescribed switch.[24]

The strategies considered by Kandori, Ellison, and Deb and González-Díaz all rely on collective punishment: if a single player deviates, the entire community is eventually punished. While such strategies seem natural in anonymous games, individualized incentives can also be provided even under anonymity. For example, suppose that some particular player—say player 1—is prescribed to take D in each of T consecutive periods in the repeated PD, while the other $N-1$ players are prescribed to take C during this T-period block. This behavior is potentially enforceable, because if some player other than player 1 deviates to taking D during the block, then there will be two players taking D rather than one, which is statistically detectable by the other players. Asymmetric behavior of this type features in a construction by Deb, Sugaya, and Wolitzky (2020), which establishes the full folk theorem for anonymous random matching games. The construction is rather complicated, and it resembles equilibria used to establish the folk theorem for general private monitoring structures (in particular, the

---

[24]Deb (2020) proves a more general folk theorem when matched players can communicate by cheap talk.

"block belief-free" strategies introduced by Matsushima, 2004, and Hörner and Olszewski, 2006) more than it does contagion strategies.[25] In the following statement of their theorem, V^* refers to set of feasible and individually rational payoffs in the N-player stage game induced by anonymous random matching. For example, if the underlying 2-player stage game is the PD, then the payoff vector that results when player 1 takes D while everyone else takes C is feasible (and also individually rational, if N is large), while the payoff vector that results when player 1 takes D while everyone else takes C iff they match with player 1 is infeasible, as it cannot be attained under anonymity.

Theorem 7 (Anonymous Folk Theorem) *Assume that* $\mathrm{int}V^* \neq \emptyset$. *For any* $v \in \mathrm{int}V^*$, *there exists* $\bar{\delta} < 1$ *such that, for every* $\delta > \bar{\delta}$, *there is a sequential equilibrium of the anonymous random matching repeated game with discount factor* δ *that gives payoff vector* v.

The positive results reported in Theorems 5–7 may strike the reader as being "too good to be true." Even if we are willing to assume that players can coordinate on the most efficient sequential equilibrium, however complicated, is it really possible that large, decentralized groups can support cooperation when players are completely anonymous and have no information beyond their own personal experiences? If not, what explains the discrepancy between the model's prediction that cooperation is possible despite minimal information, and our intuition that this prediction is likely too optimistic?

One possible explanation is that the models considered in this subsection do not allow "noise": actions are implemented without error, and within-match monitoring is perfect. At first glance, contagion strategies seem extremely non-robust to noise, since a single error would spark contagion.[26] However, Ellison shows that the modified contagion strategies used to prove Theorem 6 are robust to introducing sufficiently small noise. To see the intuition, note that under the modified contagion strategies that make players indifferent to starting contagion on path, the expected number of periods before cooperation "restarts" following an episode of contagion converges to a finite number as $\delta \to 1$. Hence, contagion is rarely

---

[25] Sugaya (2020) extends block belief-free strategies to prove a general folk theorem for repeated games with private monitoring under standard full-support and statistical identifiability conditions. These conditions are robustly violated in random matching games.

[26] Fearon and Laitin (1996) take this observation as a starting point for considering when relations between ethnic groups lead to contagion-like "spirals," rather than being contained through "within-group policing."

triggered when noise is small, and the payoff loss when it is triggered is finite. Consequently, in the iterated limit where first $\delta \to 1$ and then $\varepsilon \to 0$, Ellison's contagion strategies support mutual cooperation in almost every period of the repeated game.

While this result provides some reassurance that Theorems 5–7 are not entirely non-robust, it still faces some limitations (as Ellison himself observed). First, the result is sensitive to the order of limits between N and ε: it applies when $\varepsilon \to 0$ for fixed N, but not if we instead consider a sequence where $N\varepsilon$ (the expected number of "errors" in the population each period) remains fixed, much less when $N \to \infty$ for fixed ε. Second, it applies only when noise takes the form of implementation or monitoring errors that are independent across periods. A different type of "noise" that is likely realistic in large populations arises when a small fraction of players may be irrational, and thus deviate from equilibrium play in a manner that is correlated across periods. As we will see, introducing a small amount of this kind of incomplete information is quite devastating for contagion strategies, as well as for any other strategies that might attempt to provide non-myopic incentives in anonymous repeated games.

4.2   Incomplete Information and Non-Anonymous Matching

Following a suggestion in Ellison (1994), Sugaya and Wolitzky (2020) investigate the effect of introducing a slight probability of irrationality in anonymous repeated games.[27] They derive an anti-folk theorem for a large class of symmetric, anonymous games, which specializes as follows to the repeated PD with anonymous random matching and independent player-types. (Note that the theorem's negative conclusion holds for all Nash equilibria, and thus does not depend on equilibrium refinements.)

Theorem 8 (Anonymous Anti-Folk Theorem with Irrational Types) *Consider the repeated PD with anonymous random matching where each player is committed to the strategy "play D in every period" with probability ε, independently across players. For any $\eta > 0$, there exists $\bar{N}$ such that, for any population size $N > \bar{N}$ and any discount factor δ, in every*

---

[27]A stage game is *anonymous* if each player's payoff depends only on her own action and the number of other players taking each action (and not their identities). Anonymous random matching games are an important special case.

Nash equilibrium the population average payoff is less than $u(D, D) + \eta$.

That is, if each player is committed to the strategy *Always D* with some fixed independent probability $\varepsilon > 0$ (which may be arbitrarily small), equilibrium population payoffs are always close to the mutual defection payoff $u(D, D)$ whenever N is sufficiently large (even if δ is arbitrarily high, so that for instance $(1 - \delta) N$ can be arbitrarily small). The proof is rather simple. Note that, if a rational (non-committed) player deviates by playing *Always D* and the realized number of committed players is some number n, the population distribution of actions will be the same as it would have been if this player had not deviated and the realized number of committed players had instead been $n + 1$. When players' types are independent and N is large, the probability that there are n or $n + 1$ committed players are similar, so a deviation by a single rational player to *Always D* has only a small impact on the population action distribution. Since a rational player must prefer following her equilibrium strategy to deviating to *Always D*, it follows that *Always D* cannot perform much better than the rational-player equilibrium strategy against the equilibrium population action distribution. But, since D is strictly dominant in the stage game, this implies that the rational-player equilibrium strategy must itself almost always take D. Finally, if committed players always take D and rational players almost always take D, population payoffs must be close to $u(D, D)$.[28]

Theorem 8 brings discouraging news about the prospects for cooperation under anonymity. However, anonymous matching may be viewed as a theoretical benchmark rather than a realistic model of large-group cooperation: in most real-world settings, it seems likely that players can usually recognize and remember their partners. It is thus natural to ask whether the presence of irrational types of the sort considered in Theorem 8 still obstructs cooperation under non-anonymous matching.[29]

---

[28] This argument is related to Mailath and Postlewaite's (1990) theorem on the impossibility of large-population public good provision. An intuition for their theorem is that a player will not reveal that she values a public good unless her reported value has either a large impact on the probability of provision or a small impact on her expected payment. Since the former is impossible when N is large, payments must be insensitive to values, and hence the public good cannot be provided unless every possible type is willing to pay a $1/N$ share of the cost.

[29] Sugaya and Wolitzky (2022c) extends the analysis of the 2020 paper to stage games which are symmetric but not anonymous. The negative result of the 2020 paper extends for symmetric, non-anonymous games under a certain "population dominance" condition. However, this condition is violated in the repeated PD with non-anonymous random matching.

This question is addressed by Sugaya and Wolitzky (2021). A first observation is that the stark negative conclusion of Theorem 8 does not fully extend to non-anonymous matching: if players are sufficiently patient, they can simply treat the N-player random matching game as a collection of 2-player games and play grim trigger in each of them, and these strategies support cooperation among rational players when ε is sufficiently small. More specifically, these "bilateral trigger strategies" form an equilibrium if $(1 - \delta) N$ is small, which corresponds to situations where each *pair* of players meet frequently.[30] To understand whether multilateral, community-based enforcement of cooperation is effective, it therefore makes more sense to consider the case where δ is close to 1 but $(1 - \delta) N$ is large, so that each player interacts with other players frequently but meets each particular partner only infrequently. Somewhat unexpectedly, the result here parallels the negative conclusion of Theorem 8.

Theorem 9 (Non-Anonymous Anti-Folk Theorem with Irrational Types) *Consider the repeated PD with non-anonymous random matching where each player is committed to the strategy "play D in every period" with probability ε, independently across players. For any $\eta > 0$, there exists $k > 0$ such that, if the population size N and the discount factor δ satisfy $(1 - \delta) N > k$, then in every Nash equilibrium the population average payoff is less than $u (D, D) + \eta$.*

A rough intuition for this result can be given as follows. In light of Theorem 8, collective incentives—strategies where players to do not condition their play on their opponent's identity—are ineffective in the presence of irrational types. So the only hope for supporting cooperation lies with individualized incentives—strategies that somehow monitor all N players' actions. In the repeated PD with non-anonymous matching, in every period, each player observes one bit of information concerning the profile of her opponents' actions: the current opponent's action, C or D. However, it can be shown that monitoring N players' actions with non-vanishing precision as $N \to \infty$ requires $O (N)$ bits of information,[31] and

---

[30] To see why $(1 - \delta) N$ is an inverse measure of the meeting rate for each pair of players, fix the real-time discount rate r and suppose matching occurs every Δ unit of real time, so $\delta = e^{-r\Delta}$, which is close to $1 - r\Delta$ when Δ is small. Then on average each pair of players interact $1/(\Delta (N - 1)) \approx r/((1 - \delta) N)$ times per unit of real time.

[31] This observation extends the pivot probability bounds of Fudenberg, Levine, and Pesendorfer (1998) and al-Najjar and Smorodinsky (2000), which were discussed in Section 3.4.

so it takes $O(N)$ periods of play to accumulate enough information to monitor everyone's actions. But if $(1 - \delta) N$ is large, rewards or punishments delivered $O(N)$ periods in the future are insufficient to motivate cooperate, so defection prevails in every equilibrium.

As this intuition suggests, Theorem 9 continues to hold if players observe any additional K-dimensional signal each period—such as the most recent action of each of K randomly chosen players, or K signals about the average level of cooperation is the population—for any number K fixed independently of N.[32] While this is a stark negative result, there do exist realistic mechanisms by which players can obtain enough information to support cooperation. For instance, suppose that matched partners can communicate freely via cheap talk before taking actions, so that (for example) each player can tell her partner her opinion of the "standing" of each of the other $N-2$ players in the population. This form of "gossip" provides a number of bits of information that grows linearly with N, and indeed Sugaya and Wolitzky (2021) show that allowing such gossip supports a folk theorem under permissive conditions.[33] One reason why this is interesting is that gossip seems to help support cooperation in reality through a roughly similar mechanism as in the theory.[34]

Overall, Theorems 8 and 9 suggest that the view of large-group cooperation under minimal information expressed in Theorems 5–7 is probably too optimistic. In reality, cooperation in large societies likely requires substantially richer information that these minimalist, first-generation community enforcement models allowed. Fortunately, sufficient information can sometimes be obtained in practice, depending on the communication opportunities and other "informational institutions" available to the group. I will return to this theme in Section 6, which focuses on the question of what kind of information players need to support cooperation in infinite populations.

---

[32] A subtlety here is that signals that arrive at the end of period t are assumed to be independent of the period $t + 1$ match realizations. For example, the signal where each player observes only her *next* partner's most recent action is excluded. That kind of signal can sometimes support cooperation even in infinite populations, and is thus considered in Section 6.

[33] The required sufficient condition is $(1 - \delta) \log N \to 0$. This condition is also almost necessary for non-myopic incentives to be provided in any model where information is exchanged through pairwise meetings, as if $(1 - \delta)^{1+\rho} \log N \to \infty$ for any $\rho > 0$ then with high probability a player never meets anyone who met anyone who... met her within a payoff-relevant time frame.

[34] Pertinent discussions of the role of gossip include Gluckman (1963), Ellickson (1994), Noon and Delbridge (1993), Baumeister, Zhang, and Vohs (2004), Dunbar (2004), Sommerfeld et al. (2007), and Feinberg, Willer, and Schultz (2014).

5   Repeated Games on Networks

I now turn to models with interaction structures that are richer than uniform random matching. This section focuses on "repeated games on networks," meaning that a network determines the payoff or information structure.[35] Some models in this literature assume that the network is fixed over time, while others allow it to be redrawn in every period of the repeated game—note that the latter assumption allows random pairwise matching as a special case. Important new issues here include how the network structure determines the scope for cooperation, and how an individual's network position determines her behavior.

5.1   Multilateral Punishment and Contagion in Networks

A key question in any model of community enforcement of cooperation is how social interactions facilitate multilateral punishment: that is, punishment by third parties. The canonical example of multilateral punishment in random matching games is the contagion strategy, discussed in Section 4.1. For games played on a fixed network (which, depending on the specifics of the model, determines the payoff and/or monitoring structure of the game), multilateral punishment arises when a player's misbehavior toward one player in the network eventually leads to an adverse reaction by other players. In the typical case where a player's payoff is determined only by the actions of her network neighbors, this means that misbehavior toward one neighbor eventually causes a negative response by other neighbors.

While not explicitly concerned with networks per se, Bernheim and Whinston's (1990) early paper on "multimarket contact" is an important starting point for understanding multilateral punishment on networks. The paper asks whether it is easier for oligopolists to collude when they interact in many markets rather than just one, and models this by considering a population of players who play m PD stage games simultaneously (of course, the m games can be formally viewed as one big stage game). They first establish an "irrelevance theorem": if (i) the m stage games are identical, (ii) payoffs are additive across the stage games, and (iii) actions in all stage games are perfectly observed (so the repeated game has perfect monitoring), then the equilibrium condition for cooperation to be supportable in the

---

[35]Repeated games on networks are also surveyed by Nava (2016).

repeated game is independent of the number of stage games, m. The intuition is simply that a player's most profitable deviation involves defecting in all m games simultaneously, while the harshest punishment for any deviation is permanent mutual defection in all m games. So, both the deviation gain and the harshest punishment are proportional to m, and hence the equilibrium condition is independent of m.

By the same logic as in Bernheim and Whinston's irrelevance theorem, if N players on a network simultaneously play identical 2-player PD's with each of their neighbors under perfect (network-wide) monitoring with additive payoffs, the condition for cooperation to be supportable in the overall repeated game is the same as the condition for cooperation to be supportable in a single 2-player PD. Consequently, if network structure is to affect the possibility of cooperation, at least one of the conditions of the irrelevance theorem must be violated. As Bernheim and Whinston note, the irrelevance theorem breaks down if the stage games are not identical, monitoring is imperfect, or payoffs are not additive across the stage games. Many papers on repeated games on networks can be organized according to which of these violations they involve.

If the stage games are not identical, any slack in the equilibrium conditions in one game can be used to relax those in the other games. For example, suppose that two PD's are played every period, but the games differ in the value of the parameter g in the PD payoff matrix from Section 3.1. Suppose that the value of this parameter is g_1 in the first game and g_2 in the second, where

$$g_1 < \frac{\delta}{1-\delta}, \quad g_2 > \frac{\delta}{1-\delta}, \quad \text{and} \quad g_1 + g_2 < \frac{2\delta}{1-\delta}.$$

If these games are played separately then cooperation is supportable in the first game but not the second, while if the games are played together then cooperation can be supported in both of them by specifying that a deviation in either game leads to punishment in both. This and related points are developed further in the network context by Spagnolo (1999), Lippert and Spagnolo (2011), Mihm, Toth, and Lang (2009), and Mihm and Toth (2020).

Imperfect monitoring arises naturally in games on networks, for example in the much-studied case where players observe only their neighbors' actions (so the network determines

the monitoring structure, perhaps in addition to the payoff structure). An important special case of imperfect monitoring arises when m stage games are played asynchronously, as in this case we can consider a player's "action" to be a mapping from the stage game to be played in the current period to her action in that game, so that the player's realized action is an imperfect signal of this mapping. Indeed, multilateral interaction clearly helps support cooperation with asynchronous stage games, because when the games are played one-by-one the maximum deviation gain is independent of m, while the harshest punishment is still increasing in m, as a deviation in one game can eventually lead to punishments in other games. Note that the extent to which multilateral interaction facilitates cooperation depends on how quickly information about a deviation in one stage game spreads to the other stage games. This information spreads instantaneously under perfect monitoring, but more generally it spreads through the network according to a contagion-type process.

Understanding how the percolation of information through a network affects the scope for cooperation is the subject of papers by Wolitzky (2013) and Ali and Miller (2013). In the former paper, in every period each player i chooses a "cooperation level" $a_i \in \mathbb{R}_+$, and payoffs are given by

$$u_i(a) = \left(\sum_{j \neq i} f_{ij}(a_j) \right) - a_i,$$

where f_{ij} is an increasing, concave, and bounded function that measures the benefit that i receives from j's cooperation. The stage game is thus a continuous-action version of the PD. At the end of every period, each player observes some other players' actions perfectly, while observing nothing about the remaining players' actions, where the network of "who observes whom" is drawn iid across periods and is itself publicly observed.[36] For example, if players randomly match in pairs and observe only their partner's action, we recover the standard PD with anonymous random matching studied by Kandori and Ellison (albeit with continuous actions, and assuming that the realized matching process is observable). On the

---

[36]In particular, player i does not observe her own payoff $u_i(a)$. This assumption is unproblematic in several cases. First, since $u_i(a)$ is additively separable in the opponents' actions, it can be interpreted as player i's expected payoff at action profile a, where player i's realized payoff depends only on actions that she observes (e.g., on the actions of her current-period partners). Second, the infinitely repeated game can be interpreted as a game of uncertain finite duration, where δ is the probability that the game continues for another period, and all payoffs are received at the end of the game.

other hand, if each player observes the actions of the same set of opponents in every period, we have a repeated game with a fixed (directed) monitoring network.

Define player i's "average cooperation level" under a strategy profile σ by $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}^{\sigma} [a_{i,t}]$. Note that, since the functions f_{ij} are concave, if there happens to exist a stationary strategy profile that maximizes each player's average cooperation level, Jensen's inequality implies that this profile must also maximize utilitarian social welfare, so long as each player i's average cooperation level is below the "first-best" level defined by $\sum_{j \neq i} f'_{ji} \left(a_i^{FB} \right) = 1$. The paper's main result is that such a strategy profile does exist, and it is given by contagion strategies. To state this result, fix a period T, and for any two players $i, j \in I$ and any number t, let π_{jit} denote the probability that, by the beginning of period $T + t$, player j has observed a player who has observed a player who... observed player i's period-T action.[37] For example, if the monitoring network is fixed across periods then $\pi_{jit} = 1$ if the distance from j to i is at most t, and otherwise $\pi_{jit} = 0$.

Theorem 10 (Maximal Cooperation with Network Monitoring) *There exists a component-wise maximal action profile $a^* \in \mathbb{R}^N$ that satisfies*

$$a_i^* = (1 - \delta) \sum_{t=1}^{\infty} \delta^t \pi_{jit} f_{ij} \left(a_j^* \right) \qquad for\ all\ i \in I. \tag{1}$$

In every Nash equilibrium, each player i's average cooperation level is at most a_i^*. Moreover, the following contagion strategy profile gives a perfect Bayesian equilibrium where each player i's average cooperation level equals a_i^*: in every period, each player i takes action a_i^* so long as she has never observed another player $j \neq i$ taking any action $a_j \neq a_j^*$; otherwise, she takes action 0.

Since the functions f_{ij} are increasing and actions can be bounded from above by individual rationality, the existence of a maximal vector a^* satisfying equation (1) follows from Tarski's fixed point theorem. That the contagion strategies described in the theorem form a Nash equilibrium is true almost by construction: the left-hand side of equation (1) is player

---

[37]This probability does not depend on T, by the assumption that the monitoring network is iid across periods.

i's per-period cost of taking a_i^* in every period rather than deviating to the myopically dominant action 0, while the right-hand side is her expected, discounted lost benefits from others' foregone cooperation if she deviates. Moreover, equation (1) implies that each player i is indifferent between action a_i^* and action 0 at every on-path history, and hence she strictly prefers action 0 at every off-path history, by a similar argument as in Ellison (1994): consequently, the strategy profile is also sequentially rational.[38] Finally, Tarski's theorem also implies that cooperation levels above a^* cannot be supported in any Nash equilibrium where players use stationary strategies (i.e., where each player's actions are constant on path). The most subtle part of Theorem 10 is therefore the claim that such cooperation levels also cannot be supported by non-stationary strategies.[39]

Equation (1) is useful for characterizing which network structures support higher levels of cooperation, as well as which players in a fixed network have greater maximum cooperation levels than others. For example, suppose that the functions f_{ij} are all identical (so each agent benefits equally from every other agent's effort), and say that player i is *recursively more central* than player j if i has more distance-t neighbors than j does (for each t), i's distance-t neighbors have more distance-t' neighbors than j's distance-t neighbors do (for each t and t'), and so on. In this case, it is not hard to see that (1) implies that $a_i^* \geq a_j^*$, for any δ and f. That is, recursively more central players in a network take higher actions in all equilibria that support maximal cooperation.[40]

Ali and Miller (2013) consider a related model, where instead of taking a single cooperative action each period that benefits every opponent, players interact with their neighbors asynchronously in continuous time and take actions that benefit only one's current partner. They show that symmetric contagion strategies on a clique graph of degree d maximize utilitarian welfare among all stationary Nash equilibria on all graphs when each player has at

---

[38]The feature that "maximal" Nash equilibria are always sequentially rational in a continuous-action PD also arises by Ali and Miller (2013). In contrast, Haag and Lagunoff (2006) consider a repeated local-interaction model with binary actions, where off-path incentives are non-trivial.

[39]This last claim is also the only part of the theorem that uses the assumption that the monitoring network—"who observes whom"—is itself observable. The proof involves averaging incentives across a player's information sets and applying monotone methods.

[40]Recursive centrality is similar to the "ordinal centralities" studied by Sadler (2020). Sadler shows that if i is ordinally more central than j, then i takes a higher action than j in the maximal equilibrium of any game with strategic complements.

most d neighbors. The intuition is that cliques maximize a quantity that Ali and Miller term *viscosity*, which is similar the right-hand side of equation (1), and likewise captures a player's future payoff loss when she deviates from equilibrium play. Other versions of the point that cooperative incentives are strengthened when information percolates through the network faster—and that some types of networks support faster percolation than others, and hence support stronger incentives—have been made by Raub and Weesie (1990), Klein (1992), Ahn and Suominen (2001), Dixit (2003), Vega Redondo (2006), Bloch, Genicot, and Ray (2008), Lippert and Spagnolo (2011), Fainmesser (2012), Balmaceda and Escobar (2017), Larson (2017), and Fainmesser and Goldberg (2018). For example, Lippert and Spagnolo consider a model where cheap talk communication can potentially spread information more quickly than can contagion via actions, and show that this feature can rationalize "forgiving" strategies where contagion is temporary, because players who observe a deviation under permanent contagion strategies have no incentive to tell others about the deviation, whereas such communication can be incentivized under temporary contagion strategies. Similarly, Ali and Miller (2016, 2021) show that strategies where deviators are "ostracized" only temporarily can induce truthful communication about each player's standing, while "permanent ostracism" strategies cannot.[41]

The papers considered in this section so far study how the threat of multilateral punishment can enforce cooperation when the discount factor δ is too low for bilateral punishment to do so. A related literature derives folk theorems in network games where bilateral incentives are insufficient even when δ is close to 1. For example, this situation arises when payoffs are not additively separable across bilateral interactions (while monitoring is nonetheless determined by a network). A key obstacle to establishing the folk theorem in this setting is that decentralized information about each player's action must be aggregated across the network, but individual players may sometimes benefit from misrepresenting or concealing information. Thus, a typical sufficient condition for the folk theorem is that the network is *2-connected*—meaning that the network remains connected when any single player is removed—as this condition ensures that a single player cannot prevent the rest of

---

[41]Bowen, Kreps, and Skrzypacz (2013) and Barron and Guo (2021) discuss other mechanisms by which the need to induce honest communication constrains community enforcement.

the community from aggregating information.[42] Papers on this topic include Ben-Porath and Kahneman (1996), Renault and Tomala (1998), Tomala (2011), Laclau (2012, 2014), Kinateder (2013), and Cho (2014). Nava and Piccione (2014) establish a folk theorem for binary-action local interaction games, using strategies that are independent of players' beliefs about the global network structure. Wolitzky (2015) asks the related question of when the equilibrium payoff set in a network game is the same when players observe only their neighbors' actions as when they observe everyone's actions. The answer again hinges on 2-connectedness when players communicate with their neighbors through cheap talk, while a positive answer is given for any network if players can exchange a form of hard evidence akin to fiat money.

5.2   Robustness and Renegotiation in Networks

As we discussed in the context of random matching games in Section 4.1, an important concern regarding contagion strategies is their possible non-robustness to various types of noise. The basic point that the simplest version of contagion strategies are not robust—so that some type of error tolerance or forgiveness is necessary in the presence of noise—applies whether interactions are organized by random matching or a network. Other considerations are specific to the network context, such as a possible desire for the repercussions of a single deviation not to spread too far in the network, or a concern that a subnetwork of players might jointly deviate by splitting off from the rest of the group.[43]

Jackson, Rodriguez-Barraquer, and Tan (2012) consider an asynchronous PD-type game on a network, where in each a period a random player i gets an opportunity to do a favor for a random neighbor of hers, j. The arrival of favor-opportunities, as well as whether the favor was done, are publicly observed. The authors assume that if i ever declines to do a favor for

---

[42]The earliest reference to this type of condition in the repeated games literature that I am aware of is Fudenberg and Levine (1991).

[43]Another issue is that—whether the interaction structure is given by random matching or a network—robustness concerns differ in binary-action games (like those of Kandori and Ellison) and continuous-action games (like those of Wolitzky and Ali and Miller). In binary-action games, on-path incentive constraints are often slack when δ is high, so softening punishments by introducing some error tolerance can be costless. In continuous-action games, on-path incentive constraints bind in the maximial equilibrium, so softening punishments reduces the maximal incentive-compatible level of cooperation. There is thus a tradeoff between robustness and incentives in continuous-action games.

j, then i and j never do favors for each other in the future: equivalently, the link between them is permanently severed. The main case of interest is when this bilateral punishment is insufficient to motivate favors, and that instead a player provides a favor iff she loses at least $m \geq 2$ links if she declines to provide it. Clearly, the strongest incentives to do favors are provided by multilateral punishment strategies, where a player loses all her links if she ever fails to do a favor. Under these strategies, a favor-exchange network is sustainable iff each player has at least m links.

However, the authors argue that these strategies are not renegotation-proof: while it is assumed that i and j can never repair their link if i declines to do a favor for j, this presumption does not extend to links between i and her other neighbors. This reasoning leads the authors to introduce a recursive definition of renegotation-proof networks, where a network is renegotiation-proof if, whenever a player i fails to do a favor for a neighbor j (and hence the ij link is deleted), i ends up with at least m fewer links in the largest renegotiation-proof network that is a subnetwork of the original network without the ij link. Many networks are renegotiation-proof in this sense, but the authors then additionally impose a form of "robustness against contagion," which says that if a link ij is deleted and the largest renegotiation-proof subnetwork forms, players other than mutual neighbors of i and j cannot lose any links. It turns out that robustness against contagion and renegotiation-proofness jointly imply that the network must take a particular form that the authors call a "social quilt," which is a union of m-player cliques.[44] Finally, one implication of this structure is that each ij link is "supported"—meaning that i and j have at least one mutual neighbor, k—and the authors verify that most links are supported in survey data on social networks in a sample of Indian villages.

Karlan et al. (2009) and Ambrus, Möbius, and Szeidl (2014) consider a different type of renegotiation in models of informal lending and risk-sharing, respectively. In Karlan et al.'s model (which was briefly discussed in Section 2.2), an agent i needs a loan from another agent j. If i gets the loan, she can then obtain some gain g by strategically defaulting; however,

---

[44]Jackson, Rodriguez-Barraquer, and Tan's approach depends on assuming that some punishments—in particular, deleting link ij when i fails to do a favor for j—cannot be renegotiated. Ali, Miller, and Yang (2017) consider a related model without this assumption and show that the resulting notion of renegotiation-proofness is much less restrictive. Notions of renegotiation for general repeated games (where all punishments can be renegotiated) were previously developed by Farrell and Maskin (1989) and others.

doing so may cause her to lose her relationship with j, as well as possibly her relationships with some third parties. These losses are costly because each relationship is assumed to have some value, which is given exogenously but can be interpreted as a reduced form for some continuation play.

Without renegotiation, the harshest possible punishment would specify that i loses all of her relationships if she defaults on the loan. However, the authors assume that i can organize a coalitional deviation by offering to pay third parties to maintain some of their relationships. For example, if the set of players is $\{i, j, k\}$ and the relationship values satisfy $0 = v_{ij} + v_{jk} < g < v_{ij} + v_{ik}$, then the threat to destroy the valuable ik relationship (in addition to the ij relationship) if i fails to repay j is deemed non-credible, because if i fails to repay j while paying k an amount v_{jk} (which compensates k for the potential loss of his relationship with j), then i nets a profit of $g - v_{ij} - v_{jk} > 0$. Extending this logic, the authors show that the maximum amount that i will repay to j when only renegotiation-proof punishments are allowed is equal to the maximum flow from i to j in the weighted network where the weight on each link equals the corresponding relationship value. This result follows from the max-flow min-cut theorem (Ford and Fulkerson, 1956), because coalitional deviations correspond to network "cuts" that separate i and j, where the amount i must pay third parties to go along with the deviation equals the value of the cut (i.e., the total value of all severed relationships), and the max-flow min-cut theorem says that the value of the least expensive cut that separates i and j equals the maximum flow from i to j. As discussed in Section 2.2, the authors find some empirical support for the model's predictions.

Ambrus, Möbius, and Szeidl (2014) consider a similar model, where now many agents may be called upon simultaneously to make bilateral transfers (which are similar to the repayments in Karlan et al.) as part of a risk-sharing arrangement. Such an arrangement is coalition-proof if no set of agents $S \subset N$ is called upon to jointly transfer to the remaining agents $S^c = N \backslash S$ an amount of money that exceeds the total value of all relationships between agents in S and agents in S^c. In the special case where all relationships are equally valuable, this observation has interesting implications about the geometry of networks that support a high degree of risk-sharing. For example, if society tries to support perfect risk-sharing, the most tempting deviation for coalition S arises when everyone in S has the

maximum possible income (say 1) and everyone else has the minimum possible income (say 0). In this case, the total on-path consumption of agents in S is given by $|S|^2/N$ (which equals the size of coalition S, $|S|$, multiplied by per-capita consumption under perfect risk-sharing, $|S|/N$), while the maximum total consumption of these agents if they deviate equals $|S|$. Therefore, the coalition S will share risk as prescribed iff the difference $|S| - |S|^2/N = |S|(1 - |S|/N)$ is less than the per-relationship value v multiplied by the size of the *perimeter* of S, denoted $c(S)$, which is the number of links between S and S^c. This condition reduces to $a(S)v \geq 1 - |S|/N$, where $a(S) = c(S)/|S|$ is the *perimeter-area ratio* of the set S. This condition implies in particular that perfect risk-sharing is sustainable only if the perimeter-area ratio of every set of agents consisting of less than half the community is at least $1/2$, which is a restrictive condition. For example, if agents are arranged on a two-dimensional grid and S is rectangle, then $c(S)$ is of order $\sqrt{|S|}$, so $a(S)$ is of order $1/\sqrt{|S|}$, which is much less than $1/2$ when $|S|$ is large. However, the paper goes on characterize conditions for partial risk-sharing (i.e., relaxing the assumption that risk is shared perfectly even after extreme shocks like the one where everyone in S has income 1 and everyone else has income 0), and shows that these conditions are satisfied in "plane-like" networks of the kind often encountered in actual village economies.[45]

6   Repeated Games in Infinite Populations

In order to understand, say, informal contract enforcement on an online platform like eBay, or the cooperative norms that prevail in large modern societies more broadly, we must consider situations where the population size is effectively infinite. If matching in an infinite population is uniformly random, then strategies based only on players' personal experiences—such as contagion strategies—are ineffective, because it is very unlikely that an individual meets someone who met someone who... met her within a payoff-relevant time frame.[46]   There

---

[45] A related kind of coalitional deviation arises in some models of repeated games where players choose their partners. Such models are considered by Jackson and Watts (2010) in the context of general finitely-repeated games, and by Xing (2020) in the context of repeated risk-sharing games. These models are also related to the endogenous matching games discussed in Section 6.1. Relatedly, Bendor et al. (2022) analyze how frequently to "re-shuffle" players across parallel repeated games.

[46] By some reasoning laid out in Section 4.1, the same argument applies when N is finite but larger than $\exp(1/(1-\delta))$. So there is no discontinuity between large finite populations and infinite ones.

are thus two broad ways in which cooperation can be supported in an infinite population: matching can be non-random (as when pairs of players can form long-term relationships), or players can have access to some additional information about their current partner's history (such as a user's rating on an online platform, or a consumer's credit rating). I consider the former type of model in Section 6.1 and the latter type in Section 6.2

6.1 Endogenous Matching

Many economic relationships are naturally modeled as long-term partnerships that can be ended voluntarily by either party, where separation triggers a search process for both parties that eventually leads to rematching with new partners. The canonical example is the employment relationship between a worker and a firm, but business and romantic partnerships also fit the model well. Such voluntarily severable relationships are a hybrid of standard, fixed-partner repeated games, and random-matching games of the kind discussed in Section 4. As compared to fixed-partner games, the ability to separate and rematch poses a new obstacle to cooperation, as it can allow a deviating agent to escape punishment. As compared to random matching games, the ability to stick with one's current partner helps support cooperation, as it can allow an agent who behaves well to enjoy future rewards.

The seminal voluntary separation model is Shapiro and Stiglitz's (1984) work on unemployment as a worker disciplining device. In their model, employees who are caught shirking are fired and return to the pool of unemployed workers, where it takes time to match with a new firm. To incentivize workers to exert effort, firms pay "efficiency wages" that exceed the present value of unemployment. Shapiro and Stiglitz's key observation—that incentive provision depends on matched agents' continuation payoffs strictly exceeding unmatched agents'—is fundamental in voluntary separation models.

To see how these models work, suppose there is a continuum of identical agents, and consider the frictionless benchmark case where each agent is always matched with a partner. In each period $t = 1, 2, \ldots$, fraction $1 - \gamma$ of agents die and are replaced by newborns who immediately form matches, and each matched pair of players play the PD stage game from Section 3.1. Note that the survival probability γ replaces the discount factor δ: the per-period expected lifetime payoff of a player whose expected payoff in her t^{th} period of life is u_t

equals $(1 - \gamma) \sum_{t=1}^{\infty} \gamma^{t-1} u_t$. Assume for simplicity that matched partners always die at the same time. At the end of every period t, each player can sever her match, which sends both partners back to the pool, whereupon they immediately rematch with random unmatched agents at the beginning of period $t + 1$. Assume that players are anonymous, so when a player forms a match she does not know if her partner is a new agent or an old agent whose match has just broken up.[47] Players perfectly observe behavior in their own matches, but get no information about behavior in other matches.

In this model, there is no Nash equilibrium where players cooperate in every period. For, if there were, players would take C in the first period of every new match; but then a deviant who in every period takes D and then separates from her current partner and rematches would obtain a payoff of $1 + g$ in every period, which exceeds the mutual cooperation payoff of 1. This observation just restates the point that incentive provision depends on matched agents' continuation payoffs strictly exceeding unmatched agents'.

While cooperation cannot be supported in every period, it can be supported in most periods if γ is sufficiently high.[48]

Theorem 11 (Cooperation with Voluntary Separation) *For any $T \in \mathbb{N}$ such that*

$$\gamma^{T+1} \leq \gamma - (1 - \gamma) g, \tag{2}$$

there exists a sequential equilibrium where each player's per-period expected payoff equals γ^{T+1}. In particular, for any $\varepsilon > 0$, there exists $\bar{\gamma} < 1$ such that, for every $\gamma > \bar{\gamma}$, there is a sequential equilibrium of the voluntary separation game with survival probability γ where each player's per-period expected payoff exceeds $1 - \varepsilon$.

This result follows from considering *T-period trust-building strategies*, where newly matched partners play D for the first T periods of their relationship and then switch to playing C forever, with any deviation punished by immediate break-up. This strategy profile—which I

---

[47]The difference between the voluntary separation models covered in the current section and models where players have the option of exiting and re-entering under a pseudonym (as in Friedman and Resnick, 2001) is that in the latter case the partner of a player who exits can carry his reputation over to his next match—his record is not wiped clean just because his partner's is. Thus, while the two models are similar, there are more equilibria in "cheap pseudonym" models than in standard voluntary separation models.

[48]This type of result seems to have been first noted by Datta (1993) and Kranton (1996a).

will denote by σ^T—gives each player an expected payoff of γ^{T+1}, since on path $u_t = 0$ for for the first T periods of a player's life and $u_t = 1$ thereafter. To see when strategy σ^T forms an equilibrium, note that there is nothing to be gained by deviating from D to C during the first T periods of a relationship, or by breaking up a relationship prematurely (since continuation payoffs increase over the course of each relationship). So, the only tempting deviation is playing D rather than C in a relationship that has already lasted more than T periods. A player who deviates in this way gets a current-period payoff of $1 + g$ and then rematches, for a total continuation payoff of $(1 - \gamma)(1 + g) + \gamma^{T+1}(1)$. Therefore, σ^T is an equilibrium iff this deviation payoff is less than the equilibrium continuation payoff of a player whose relationship has already lasted more than T periods, which equals 1. This inequality reduces to (2). Finally, the right-hand side of (2) converges to 1 as $\gamma \to 1$, so T can be chosen as a function of γ so that $\sigma^{T(\gamma)}$ is an equilibrium yielding payoffs that converge to 1 as $\gamma \to 1$. This last step reflects the fact that, when $\gamma \approx 1$, approximately g periods of trust building are required to deter deviations, and waiting g periods before starting to cooperate is almost costless when $\gamma \approx 1$.

The essential feature of trust-building strategies is that newly matched players do not immediately start cooperating. This feature is reflected in many real-world social and economic relationships, presumably for much the same reason as in the model: if strangers are trusted as thoroughly as long-term partners, the ability to repeatedly exploit new partners destroys cooperation. This and related points are discussed by Carmichael and MacLeod (1997), who also note that other mechanisms besides delay can serve to build trust, such as ceremonial gift exchanges at the start of relationships that amount to a form of "money-burning." The exchange of wedding rings is a standard example.[49]

While trust-building strategies are approximately efficient when the survival probability γ is high, they can generally be improved upon for a given value of γ. Eeckhout (2006) shows that one way to do this is through "endogenous segregation," where a payoff-irrelevant trait like skin color determines the time at which cooperation commences in each relationship. In his model, skin color serves as a correlating device, the availability of which is shown to

---

[49]The relationship-building role of gift exchanges has also been intensively studied by anthropologists, dating back to the foundational works of Malinowski (1922) and Mauss (1925).

yield a strict Pareto improvement for generic values of the parameters (g, ℓ, γ). To see the basic idea, suppose that γ is low enough that 1-period trust-building strategies (i.e., σ^1) form an equilibrium, and that moreover so do strategies where new partners mix between C and D with some probabilities in the first period of their relationship (before taking C with probability 1 beginning in the second period). Such independent mixing between C and D typically entails some inefficiency as compared to the case where the players can jointly mix between (C, C) and (D, D), which becomes possible if a correlating device is available. Fujiwara-Greve and Okuno-Fujiwara (2009) show that it is also possible to improve on the T-period trust-building strategy σ^T without exogenously introducing correlating devices, by instead considering equilibria where players mix between σ^T and σ^{T-1}.

It is also interesting to consider the effects of introducing adverse player types in voluntary separation models, such as a fraction of players who are committed to always playing D (like the commitment types in Section 4.2). Ghosh and Ray (1996) show that the presence of such types can sometimes *help* support cooperation. This is because the presence of adverse types in the matching pool discourages players from exploiting and then separating from their current partners. For example, in this model σ^0 (i.e., immediate cooperation in each relationship) is an equilibrium iff the population share of adverse types is intermediate. Indeed, if the share of adverse types is close to 0, then rational player payoffs under σ^0 are close to 1, while a deviator who always takes D gets a payoff close to $1 + g$; while if the share of adverse types is close to 1, then rational player payoffs under σ^0 are close to $-\ell$, while a deviator who always takes D gets a payoff close to 0; however, it is not hard to show that σ^0 is sometimes an equilibrium for an intermediate share of adverse types.[50]

6.2 Uniform Matching and Record-Keeping

The applications just discussed notwithstanding, there are also many economic relationships which are necessarily short-term: for instance, a consumer is typically interested in trans-

---

[50]Conversely, features of the environment that improve payoffs for the pool of anonymous, unmatched players can undermine cooperation by making separation more attractive. Kranton (1996b) argues that the development of market exchange can play such a role (see also Baker, Gibbons, and Murphy, 1994). Gagnon and Goyal (2017) analyze a related model where long-run relationships are governed by a network, and Banerjee et al. (2021) develop an empirical application to credit markets in rural India.

acting at most once with any one eBay seller or AirBnB host. In an effectively infinite population with short-term relationships, incentives can be provided only if players have access to some information about their current partners' histories. We will see that the richness of this information is a key determinant of the scope for cooperation. In particular, if players receive information about their partners' past actions but not about the context of these actions—such as, for example, what actions their partners' past partners took toward them—then whether non-myopic play can be supported may depend on the details of the stage game. In contrast, the folk theorem holds quite generally when this type of contextualizing information is available.

To state these results precisely, we must first specify a model of repeated games with random matching in an infinite population, given an underlying symmetric, 2-player stage game (A, u). There are two main ways of doing this. The first adapts the fixed finite-population models discussed in Section 4.1 to a fixed infinite population, and in particular assumes that all players start the game at the same time, share a common notion of calendar time, and are infinitely-lived with discount factor δ. This approach is taken by Okuno-Fujiwara and Postlewaite (1995) and Takahashi (2010). The second approach modifies the models with entry and exit covered in Section 6.1 by replacing voluntary separation with uniform random matching, and thus considers steady states in a model where each player has a geometric lifespan with survival probability γ. This approach is taken by Clark, Fudenberg, and Wolitzky (2020). As we will see, many results in these two classes of models are similar, although there are also some differences.

In either model, one assumes that each player carries a *record r* that summarizes some information about her history, and that when two players meet, each observes the other's record but no further information.[51] At the end of each period, a player's record is updated according to an exogenous rule that typically depends on at most the current partner's record and both players' actions. If the updating rule depends only on a player's own action, the record is called *first-order*. A leading example of first-order records arises when a player's record consists of the entire history of her own actions, so that players perfectly observe their

---

[51]The term *record* is from Clark, Fudenberg, and Wolitzky. Okuno-Fujiwara and Postlewaite call the same object a *status level*; Kandori (1992b) calls it a player's *state*; and the evolutionary biology literature surveyed by Sigmund (2010) often calls it a player's *standing*.

partners' past play but get no further information—this is the case considered by Takahashi (2010). If the updating rule additionally depends on the partner's action—as for example in the case where a player's record consists of the entire history of the *(own action, opponent's action)* outcomes of her past matches—the record is called *second-order*. In the general case where the updating rule can also depend directly on the partner's record—as in certain models of "social standing" (Sigmund, 2010)—records are called *interdependent*.

Second-order records turn out to be rich enough to support a general folk theorem. Let A_0 denote each player's stage-game action set, and let $\underline{v}^P = \min_{a_j \in A_0} \max_{a_i \in A_0} u\left(a_i, a_j\right)$ denote the pure-strategy minmax payoff.

Theorem 12 (Folk Theorem with Second-Order Records) *There exists a second-order record system such that the following holds: for any action a_0 satisfying $u\left(a_0, a_0\right) > \underline{v}^P$, there exists $\bar{\delta} < 1$ such that, for every $\delta > \bar{\delta}$, there is a sequential equilibrium of the repeated game with discount factor δ where all players take a_0 in every period along the equilibrium path.*

This result combines ideas from Okuno-Fujiwara and Postlewaite and Clark, Fudenberg, and Wolitzky.[52] It applies equally for fixed-population games and games with geometric player lifespans (with γ in place of δ). To prove it, fix a minmaxing action $a^- \in \operatorname{argmin}_{a_j \in A_0} \max_{a_i \in A_0} u\left(a_i, a_j\right)$ and a best response $a^+ \neq a_0$ to a^-, and let (α, α) be a symmetric stage-game Nash equilibrium.[53] Suppose that each player's record is *Good* if the outcome of each of her past matches is either (a_0, a_0) or (a^-, a^+), and is *Bad* otherwise: note that, by definition, such records are second-order. Finally, suppose that when two players meet, both take a^0 if both of their records are *Good*; both take α if both of their records are *Bad*; and otherwise the *Good*-record player takes a^- and the *Bad*-record player takes a^+. To see that these strategies form an equilibrium when δ is sufficiently high, note that the *Bad* record is absorbing and that *Bad*-record players always take myopic best responses. Since $a^+ \neq a_0$, a *Good*-record player who meets a *Good*-record opponent (and thus anticipates that the opponent will play a_0) must play a_0 to maintain her *Good* record; similarly, a

---

[52]Sugden (1986) and Kandori (1992b) established related results.

[53]Such (a^-, a^+) always exist unless a_0 is the unique best response to every minmaxing action. In this case, Theorem 12 remains true, but the proof requires a more complicated construction. See Theorem 2(i) of Clark, Fudenberg, and Wolitzky for the details.

Good-record player who meets a *Bad*-record opponent must play a^- to maintain her record. Note that in equilibrium all players' records are *Good*, so a *Good*-record player's continuation payoff is $u(a_0, a_0)$, while that of a *Bad*-record is $\underline{v}^P$. Therefore, when δ is sufficiently high, it is optimal for *Good*-record players to play a_0 against *Good*-record opponents and to play a^- against *Bad*-record opponents, as prescribed.

This argument invites the question of whether Theorem 12 is robust to introducing noise—e.g., constraining players to take each action with at least ε probability—which would guarantee that a positive share of players has every possible record. Clark, Fudenberg, and Wolitzky show that the result is indeed robust, in that there exists a sequence of equilibrium steady states that converges to the steady state where all players have *Good* records in the double limit where $(\gamma, \varepsilon) \to (1, 0)$. Moreover, the corresponding equilibria are strict, which ensures robustness to various further perturbations of the model.

Theorem 12 shows that a relatively small amount of contextualizing information is enough to support cooperation in large populations: it suffices to observe summary statistics of one's partner's past *outcomes*, meaning both their own actions and what actions were taken toward them. However, in some settings contextualizing information is entirely absent, such as when a platform user's rating depends only on the quality of the products she has sold, so that only first-order records are available. In this case, whether cooperation is supportable for high δ turns out to depend on the details of the stage game.

To see the idea, consider again the PD stage game introduced in Section 3.1, and suppose that records are first-order, so that a player's continuation payoff at the beginning of period $t+1$ depends only on her own period t record and action. If a player is supposed to take C rather than D against an opponent whom she expects to take C, her continuation payoff must be *at least* g utils higher when she takes C rather than D (because she would gain g in the current period by taking D). On the other hand, if she is supposed to take D against an opponent whom she expects to take D, her continuation payoff must be *at most* ℓ utils higher when she takes C rather than D (because she would lose ℓ in the current period by taking C). In particular, if $g > \ell$—so that the game is *strictly submodular*, meaning that C is relatively more attractive when the opponent takes D—then it is impossible to motivate a player to take C against opponents who take C while taking D against opponents who

take D. Moreover, in the absence of such "reciprocating" behavior, it can be shown that cooperation cannot be supported in any strict equilibrium.[54]

Theorem 13 (Anti-Folk Theorem with First-Order Records) *For any first-order record system and any discount factor δ, if the stage game is the submodular PD then the unique strict equilibrium in the repeated game is* Always Defect.

By contrast, in the strictly supermodular PD (where $g < \ell$) with a fixed population of infinitely-lived players, Takahashi shows that there exists a strict equilibrium where players cooperate in every period along the equilibrium path. This result follows by using Ellison's "threading" technique to tune the effective discount factor so that the difference in a player's continuation payoff when she takes C rather than D lies in between g and ℓ. However, if one introduces noise in the model, then eventually play in such an equilibrium converges to mutual defection in every period, so these equilibria do not seem very robust. Consistent with this intuition, Clark, Fudenberg, and Wolitzky establish necessary and sufficient conditions for mutual cooperation to be supported in a robust equilibrium steady state with geometric player lifespans and noise when $(\gamma, \varepsilon) \to (1, 0)$, and find that these conditions are strictly more restrictive than $g < \ell$.[55]

While Theorem 13 shows that cooperation with first-order records in the PD is delicate, cooperation becomes much more robust if the game is augmented with an "exit option," which yields some payoff between 0 and 1 for both players whenever either player takes it. Intuitively, the problem with first-order records in the PD is that they cannot distinguish between "justified" defections that are taken to punish an opponent with a bad record, and "unjustified" defections that are taken to exploit an opponent with a good record. If punishments are instead delivered by "exiting," this distinction does not need to be drawn, because a player who expects her opponent to cooperate is not tempted to exit. Clark,

---

[54]This result was established for fixed-population games by Takahashi, and for games with geometric player lifespans by Clark, Fudenberg, and Wolitzky. Heller and Mohlin (2018) derive the same result in a model with infinitely-lived players and doubly-infinite time. Takahashi shows that cooperation can be supported in a non-strict, belief-free equilibria when $g \geq l$; however, Heller and Mohlin exhibit a sense in which these equilibria are non-robust. Dilmé (2016) establishes related results, focusing on the case where $g = l$.

[55]Heller and Mohlin (2018) derive related conditions for the existence of cooperative equilibria in their doubly-infinite time model. Their equilibria are not strict but do satisfy a stability condition.

Fudenberg, and Wolitzky formalize this point by deriving a folk theorem with first-order records for games where such "unprofitable punishments" are available.[56]

The models discussed so far involve normal-form stage games. Bhaskar and Thomas (2018) consider a sequential-move stage game, where a lender decides whether to extend a loan, and then a borrower decides whether to repay. They show that cooperative equilibria unravel when records are of any bounded length T (under a "purifiability" refinement), because lenders cannot be prevented from lending to borrowers who last defaulted exactly T periods ago, which effectively reduces records to length $T - 1$, and so on. However, cooperation is restored if lenders learn only whether or not the borrower defaulted at some point in the last T periods. This result provides a rationalization for coarse credit reports.[57]

Finally, a number of important papers analyze how specific institutional mechanisms provide enough information for large groups to cooperate in particular contexts. In addition to the work on online ratings systems and credit ratings already discussed, these include Milgrom, North, and Weingast (1990) on private judges in medieval trade fairs, Greif, Milgrom, and Weingast (1994) and Clay (1997) on merchant coalitions in medieval Europe and 19th-century California, respectively, Kiyotaki and Wright (1993) on fiat money, Kandori and Obayashi (2014) on Japanese labor unions, and Budish (2018), Abadi and Brunnermeier (2019), and Biais et al. (2019) on blockchains.

7 Institutionalized Enforcement

In all of the models I have considered so far, cooperative behavior is enforced informally through repeated game effects. In contrast, traditional economic models implicitly assume the existence of an all-powerful legal system that perfectly enforces any agreement that agents may make. In reality, incentives derive from both informal, relational sources and formal, institutionalized ones. A number of recent papers have tried to understand how these

---

[56]This result resonates with a debate in evolutionary biology as to whether the inability to distinguish warranted punishments from unwarranted deviations is a significant impediment to supporting cooperation. Feher and Gächter (2002) and Boyd et al. (2003) argue for the importance of costly punishments, while Baumard (2010) and Guala (2012) suggest that human cooperation is supported by the threat of simply avoiding deviators (an example of an "unprofitable punishment"), rather than actively punishing them.

[57]Other studies of credit reporting as an information system in a repeated game include Klein (1992) and Padilla and Pagano (2000).

different sources of incentives interact. For example, what types of behaviors are, or should be, enforced only informally, and when instead does the legal system or another institution with coercive power involve itself in regulating behavior? Are formal and informal incentives "complements" or "substitutes"? Given that judges, police officers, and armed goons are presumably all rational agents in their own rights, should we even distinguish between formal and informal enforcement at all, rather than viewing all of society as being constantly engaged in a single "game of life" (Binmore, 1994)? No standard modeling approach for addressing these questions has yet emerged, but I will nonetheless try to give a snapshot of the state of the literature.

I will distinguish between two general perspectives on how institutionalized enforcement relates to standard repeated game enforcement of the kind discussed in the earlier sections. The first is that institutionalized enforcement is that subset of repeated game enforcement which is carried out by special agents whom we call "judges," "police," or "goons," where these agents differ from "regular agents" in that they have access to special types of punishing actions, such as imprisonment or corporal punishment. This perspective of "institutional enforcement as special sanctions" is taken by Ramey and Watson (2002), Dixit (2003a,b, 2004), Sánchez-Pagés and Straub (2010), Aldashev et al. (2012), Andreoni and Gee (2012), Masten and Prüfer (2014), Acemoglu and Jackson (2016), Levine and Modica (2016), Aldashev and Zanarone (2017), and Acemoglu and Wolitzky (2020, 2021). The second perspective is that institutionalized enforcement refers to a special type of payoff-irrelevant signal or message that serves to coordinate the players in the game, such as a legal ruling or royal decree. This perspective of "institutional enforcement as special communication" (which is related to the analysis of information systems in Section 6.2) is taken by Myerson (2004), Benabou and Tirole (2011), Hadfield and Weingast (2012), McAdams (2015), Mailath, Morris, and Postlewaite (2017), Basu (2018), and Jackson and Xing (2020). A thorough survey of this literature would require its own paper, so I will just discuss a couple articles from each of these two lines of research that are especially closely related to standard repeated game models of community enforcement.[58]

---

[58]There are also other perspectives on institutional enforcement besides these two. For example, a third perspective, emphasized by North, Wallis, and Weingast (2009), is that laws are distinguished by being "impersonal": that is, by treating agents differently based only on their societal positions, rather than their

Acemoglu and Wolitzky (2020) add agents specialized in coercive enforcement to a standard community enforcement model, so the population consists of N "regular agents" and M "enforcers," where for simplicity I here take $M = 1$. Every period, each regular agent i chooses a "cooperation level" $a_i \in \mathbb{R}_+$, which provides a benefit of $f(a_i)$ to every other agent (both the other regular agents and the enforcer) at a cost of a_i to agent i, where f is an increasing, concave, and bounded function.[59] At the end of every period, after observing each regular agent's chosen cooperation level, the enforcer chooses a "punishment level" $b_i \in \mathbb{R}_+$ for each regular agent i, which inflicts a disutility of $g(b_i)$ upon agent i at a cost of b_i to the enforcer, where g is another increasing and concave function (which may or may not be bounded). Thus, just as regular agents must be incentivized to operate their "production technology" f, so the enforcer must be incentivized to operate his "punishment technology" g. The game is repeated under perfect monitoring, and all agents (both regular agents and the enforcer) share a common discount factor δ.

The paper asks what subgame perfect equilibrium maximizes cooperation among regular agents.[60] Since monitoring is assumed to be perfect, an obvious candidate is grim trigger strategies, where regular agents take some prescribed action a^* so long as no regular agent has yet deviated, and otherwise take the myopically optimal action, 0. These strategies always form an optimal equilibrium without an enforcer, but they are *never* optimal when an enforcer is present (though they remain *an* equilibrium, since everyone can always ignore the enforcer). The reason is that grim trigger strategies give the enforcer no incentives to operate the punishment technology: under these strategies, cooperation completely breaks down after any deviation, so the enforcer has no incentive to exert "punishment effort" himself.

Instead, consider "repentance strategies," where if some regular agent i deviates in period t, she is supposed to be punished by the enforcer at the end of period t, and then is supposed to cooperate in period $t + 1$ while all other regular agents exert 0 effort; and then (assuming that the enforcer punishes agent i at the end of period t and agent i exerts effort in period

---

personal identities.

[59] The modeling of regular agents here is thus as in Wolitzky (2013), which I discussed in Section 5.1.

[60] This equilibrium also maximizes utilitarian social welfare if the maximum cooperation level is below the first-best level.

$t+1$) all regular agents resume cooperation at the original on-path level a^* in period $t+2$. It is not hard to see that in the absence of an enforcer, repentance strategies support the same maximum cooperation level as grim trigger strategies.[61] Hence, since repentance strategies provide some incentive to the enforcer, they support a strictly higher cooperation level in the game with an enforcer. Indeed, the optimal equilibrium with an enforcer can be shown to take a form similar to repentance strategies, with the difference that if the punishment technology is sufficiently effective, the optimal equilibrium may involve no reduction in effort by regular agents following a deviation by another regular agent. In this case, the threat of reduced effort by regular agents is used exclusively to motivate the enforcer; and, conversely, regular agents are motivated solely by the threat of being punished by the enforcer. A sufficient condition for such "enforcer punishment strategies" to be optimal is that $g'(b) \geq 1$ for all $b \in \mathbb{R}_+$.

Acemoglu and Wolitzky (2021) extend this model to develop a theory of the emergence of legal equality. Suppose that the coercive punishment technology is now in the hands of a "state," and that there are two types of productive agents: "regular agents" and "elites." The difference between these two types of agents is that elites are partially immune to coercive punishment: while the state can inflict a punishment of g on a regular agent, it can only inflict a punishment of ρg on an elite, where $\rho \in [0, 1]$ is a parameter that measures the extent to which elites are "above the law." The paper asks what level of ρ is collectively preferred by elites. Intuitively, a higher value of ρ (corresponding to greater legal equality) commits elites to exerting more effort, which is myopically costly but has the benefit of increasing continuation values for both regular agents and elites, which in turn increases the maximum incentive-compatible effort level of regular agents. A key comparative static is that the elite-preferred level of ρ is decreasing in g: that is, the greater is the "coercive capacity" of the state, the less inclined are elites to move toward legal equality. The logic of this result is that, since there are diminishing returns to effort (i.e., f is concave), trading off greater own effort for greater normal-agent effort by increasing ρ is less attractive for elites

---

[61]Intuitively, under repentance strategies the deviator "burns her own utility" instead of facing future punishment. The observation that optimal equilibria are stationary when unbounded money-burning is available is similar to results that arise in repeated games with transferrable utility (Levin, 2003; Goldlücke and Kranz, 2012).

when the threat of coercive punishment already induces a higher level of normal-agent effort.

Turning from enforcement-based theories to informational ones, Hadfield and Weingast (2012) articulate a view of legal decisions as mechanisms for coordinating multilateral punishment. In their model, a seller interacts repeatedly with each of two buyers. The seller has occasional opportunities to short-change each buyer, and can be deterred from doing so by the threat of a coordinated boycott, where both buyers break off trade; however, the threat of a unilateral boycott is assumed to be insufficient to deter opportunism by the seller. Each buyer receives a private signal of whether she was short-changed in the current period, so it is possible for each buyer to threaten to boycott the seller when she is short-changed. But, since it is assumed that the buyers cannot communicate with each other, a buyer cannot threaten to boycott the seller when the *other* buyer is short-changed, and hence seller opportunism cannot be deterred.[62] To this model, Hadfield and Weingast add a "classification institution," where each buyer reports her signal to the institution at the end of each period, and the institution then issues a public "ruling" regarding the seller's behavior. The authors show that, if the institution's rulings are sufficiently accurate, they can form the basis for coordinated boycotts, and can thereby deter seller opportunism.

Jackson and Xing (2020) develop another model where legal decisions coordinate multilateral punishment, with the novel feature that these punishments can be carried out only within an individual's own community. In the model, each individual needs a task to be completed in every period. A binding societal convention dictates which tasks are to be completed within the community, and which are to be completed in the market. Due to gains from specialization, tasks can always be completed more efficiently in the market. However, agents are assumed to interact in the market anonymously, whereas their identities are observed within their community. This implies that an agent who fails to complete another community member's task can be completely ostracized within her community, while an agent who fails to complete a market task can only be forced to pay a fine (which is assumed to be capped at a certain level). There is thus a familiar tradeoff between more-efficient

---

[62]This situation is similar to the models of multilateral enforcement on networks that we encountered in Section 5.1. View the seller as the center node in a three-node line, and suppose that simultaneous PD's are played on each link, and that the peripheral nodes cannot directly communicate with each other. Then cooperation by the seller can be supported iff it can be supported in each bilateral relationship independently. This result is formalized by Proposition 1 in Wolitzky (2015).

market exchange and better-monitored and enforced community exchange. The model adds to this standard situation the assumption that if an agent shirks a market task, then not only is she fined, but with some probability the information that she shirked in the market spreads to her community, in which case she may also be ostracized by her community. Due to this effect, allocating more tasks to the community effectively makes market exchange more enforceable, as it increases the penalty for shirking in market exchange (i.e., the fine plus the product of the probability of ostracism and the loss from being excluded from community tasks). Together with the assumption that market exchange is more efficient, this effect implies that it is typically optimal to allocate some tasks to the market and others to the community. The model thus uncovers a novel complementarity between community and market exchange.

8   Conclusion

Our theoretical understanding of large-group cooperation has already come a long way. A generation ago, economists were just beginning to understand the importance of informal contractual enforcement and collective action, as well as the role of information and repeated interaction in supporting these behaviors. We now have a much more sophisticated under-standing of, for example, the interaction among group size, patience, and information; the distinct roles of collective and individual sanctions; the significance of whether behavior is monitored through centralized, public signals or decentralized, private ones; the role of social structures such as networks and voluntarily separable relationships; and the importance of various types of informational and coercive institutions.

Nonetheless, in reviewing this literature I am mostly struck by how much we still have to learn. In almost every area I have touched on, only a couple papers have been written (or at most a small handful), and there are many wide open questions with prima facie relevance to important applied questions. For example, while we are beginning to appreciate the significance of various model features, such as whether signals are public or private, whether noise is modeled as independent shocks or the presence of adverse types, or whether matching is random or network-like, we seem to know extremely little about how these

different features interact with each other. Since theory research in this area aspires to influence more applied thinking on social cooperation as well as empirical work (which are, of course, proceeding apace), constructing more unified models that take these interactions into account seems to be a relatively urgent task.

Finally, I have said very little about other areas of economic theory that bear on social cooperation, such as behavioral economics (e.g., Sobel, 2005), organizational economics (Kreps, 1990; Baker, Gibbons, and Murphy, 2002), and political economy (Myerson, 2008; North, Wallis, and Weingast, 2009; Acemoglu and Robinson, 2012). Synthesizing some of the core models in these areas—such as models of social preferences, corporate culture, and organizational or political hierarchies—with those considered in this survey is another promising direction for future research.

References

[1] Abadi, Joseph, and Markus Brunnermeier. Blockchain economics. No. w25407. National Bureau of Economic Research, 2019.

[2] Abreu, Dilip, Prajit K. Dutta, and Lones Smith. "The folk theorem for repeated games: a NEU condition." Econometrica: Journal of the Econometric Society (1994): 939-948.

[3] Abreu, Dilip, David Pearce, and Ennio Stacchetti. "Toward a theory of discounted repeated games with imperfect monitoring." Econometrica: Journal of the Econometric Society (1990): 1041-1063.

[4] Acemoglu, Daron, and Matthew O. Jackson. "Social norms and the enforcement of laws." Journal of the European Economic Association 15.2 (2017): 245-295.

[5] Acemoglu, Daron, and James A. Robinson. Why nations fail: The origins of power, prosperity, and poverty. Currency, 2012.

[6] Acemoglu, Daron, and Alexander Wolitzky. "Sustaining Cooperation: Community Enforcement versus Specialized Enforcement." Journal of the European Economic Association 18.2 (2020): 1078-1122.

[7] Acemoglu, Daron, and Alexander Wolitzky. "A theory of equality before the law." *Economic Journal*, 131:1429-1465. (2021).

[8] Ahn, Illtae, and Matti Suominen. "Word-Of-Mouth Communication and Community Enforcement." *International Economic Review* 42.2 (2001): 399-415.

[9] Aldashev, G., Chaara, I., Platteau, J. P., & Wahhaj, Z. (2012). Using the law to change the custom. Journal of Development Economics, 97(2), 182-200.

[10] Aldashev, Gani, and Giorgio Zanarone. "Endogenous enforcement institutions." Journal of Development Economics 128 (2017): 49-64.

[11] Alesina, Alberto, Reza Baqir, and William Easterly. "Public goods and ethnic divisions." The Quarterly journal of economics 114.4 (1999): 1243-1284.

[12] Alesina, Alberto, and Eliana La Ferrara. "Participation in heterogeneous communities." The quarterly journal of economics 115.3 (2000): 847-904.

[13] Alexander, Richard D. The biology of moral systems. Transaction Publishers, 1987.

[14] Ali, S. Nageeb, and David A. Miller. "Enforcing Cooperation in Networked Societies." *Working Paper* (2013).

[15] Ali, S. Nageeb, and David A. Miller. "Ostracism and Forgiveness." *American Economic Review* 106.8 (2016): 2329-48.

[16] Ali, S. Nageeb, and David A. Miller. "Communication and Cooperation in Markets." *AEJ:Microeconomics*, Forthcoming (2021).

[17] Ali, S. Nageeb, David A. Miller, and David Yang. "Renegotiation-proof multilateral enforcement." *Working paper* (2017).

[18] Al-Najjar, Nabil I., and Rann Smorodinsky. "Pivotal Players and the Characterization of Influence." *Journal of Economic Theory* 92.2 (2000): 318-342.

[19] Al-Najjar, Nabil I., and Rann Smorodinsky. "Large nonanonymous repeated games." Games and Economic Behavior 37.1 (2001): 26-39.

[20] Ambrus, Attila, Markus Möbius, and Adam Szeidl. "Consumption risk-sharing in social networks." American Economic Review 104.1 (2014): 149-82.

[21] Andreoni, James, and Laura K. Gee. "Gun for hire: Delegated enforcement and peer punishment in public goods provision." Journal of Public Economics 96.11-12 (2012): 1036-1046.

[22] Aumann, Robert J., and Lloyd S. Shapley. "Long-term competition—a game-theoretic analysis." 1976.

[23] Athey, Susan, and Kyle Bagwell. "Optimal collusion with private information." RAND Journal of Economics (2001): 428-465.

[24] Axelrod, Robert. *The Evolution of Cooperation*. Basic Books. 1984.

[25] Baker, George, Robert Gibbons, and Kevin J. Murphy. "Subjective performance measures in optimal incentive contracts." The quarterly journal of economics 109.4 (1994): 1125-1156.

[26] Baker, George, Robert Gibbons, and Kevin J. Murphy. "Relational Contracts and the Theory of the Firm." The Quarterly Journal of Economics 117.1 (2002): 39-84.

[27] Balmaceda, Felipe, and Juan F. Escobar. "Trust in Cohesive Communities." *Journal of Economic Theory* 170 (2017): 289-318.

[28] Banerjee, A. V., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2021). Changes in social network structure in response to exposure to formal credit markets. Available at SSRN 3245656.

[29] Banerjee, Abhijit, Lakshmi Iyer, and Rohini Somanathan. "History, social divisions, and public goods in rural India." Journal of the European Economic Association 3.2-3 (2005): 639-647.

[30] Barron, Daniel, and Yingni Guo. "The Use and Misuse of Coordinated Punishments." *Quarterly Journal of Economics*, 136.1 (2021): 471-504.

[31] Basu, Kaushik. The republic of beliefs: A new approach to law and economics. Princeton University Press, 2018.

[32] Baumard, Nicolas. "Has punishment played a role in the evolution of cooperation? A critical review." Mind & Society 9.2 (2010): 171-192.

[33] Baumeister, Roy F., Liqing Zhang, and Kathleen D. Vohs. "Gossip as Cultural Learning." *Review of General Psychology* 8.2 (2004): 111-121.

[34] Bendor, Jonathan, Lukas Bolte, Nicole Immorlica, and Matthew O. Jackson. "Interactions across multiple games: cooperation, corruption, and organizational design." *Working Paper* (2022).

[35] Ben-Porath, E. and Kahneman, M. (1996), "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 281-298.

[36] Benabou, Roland, and Jean Tirole. Laws and norms. No. w17579. National Bureau of Economic Research, 2011.

[37] Bernheim, B. Douglas, and Michael D. Whinston. "Multimarket contact and collusive behavior." The RAND Journal of Economics (1990): 1-26.

[38] Besley, Timothy, Stephen Coate, and Glenn Loury. "The economics of rotating savings and credit associations." The American Economic Review (1993): 792-810.

[39] Bhaskar, V., and Caroline Thomas. "Community Enforcement of Trust with Bounded Memory." *Review of Economic Studies* 86.3 (2019): 1010-1032.

[40] Biais, B., Bisiere, C., Bouvard, M., & Casamatta, C. (2019). The blockchain folk theorem. The Review of Financial Studies, 32(5), 1662-1715.

[41] Binmore, Kenneth George. *Game theory and the social contract.* MIT press, 1994.

[42] Bloch, F., Genicot, G. and Ray, D. (2008), "Informal Insurance in Social Networks," *Journal of Economic Theory*, 143, 36-58.

[43] Bolton, Gary, Ben Greiner, and Axel Ockenfels. "Engineering trust: reciprocity in the production of reputation information." Management science 59.2 (2013): 265-285.

[44] Bowen, T. Renee, David M. Kreps, and Andrzej Skrzypacz. "Rules with Discretion and Local Information." *Quarterly Journal of Economics* 128.3 (2013): 1273-1320.

[45] Bowles, Samuel, and Herbert Gintis. A cooperative species: Human reciprocity and its evolution. Princeton University Press, 2013.

[46] Boyd, R., Gintis, H., Bowles, S., & Richerson, P. J. (2003). The evolution of altruistic punishment. Proceedings of the National Academy of Sciences, 100(6), 3531-3535.

[47] Boyd, Robert, and Peter J. Richerson. "Punishment allows the evolution of cooperation (or anything else) in sizable groups." Ethology and sociobiology 13.3 (1992): 171-195.

[48] Budish, Eric. The economic limits of bitcoin and the blockchain. No. w24717. National Bureau of Economic Research, 2018.

[49] Bull, Clive. "The existence of self-enforcing implicit contracts." The Quarterly journal of economics 102.1 (1987): 147-159.

[50] Cabral, Luis. "Reputation on the Internet." The Oxford handbook of the digital economy (2012): 343-354.

[51] Carmichael, H. Lorne, and W. Bentley MacLeod. "Gift giving and the evolution of cooperation." International Economic Review (1997): 485-509.

[52] Cho, Myeonghwan. "Cooperation in the repeated prisoner's dilemma game with local interaction and local communication." International Journal of Economic Theory 10.3 (2014): 235-262.

[53] Clark, Daniel, Drew Fudenberg, and Alexander Wolitzky. "Record-Keeping and Cooperation in Large Societies." *Review of Economic Studies* 88:2179-2209 (2021).

[54] Clay, Karen. "Trade without law: private-order institutions in Mexican California." The Journal of Law, Economics, and Organization 13.1 (1997): 202-231.

[55] Coate, Stephen, and Martin Ravallion. "Reciprocity without commitment: Characterization and performance of informal insurance arrangements." Journal of development Economics 40.1 (1993): 1-24.

[56] Coleman, J. (1990), *Foundations of Social Theory*, Cambridge: Harvard University Press.

[57] Datta, Saikat. Building trust. No. 1996/305. Suntory and Toyota International Centres for Economics and Related Disciplines, LSE, 1996.

[58] Deb, Joyee (2020), "Cooperation and Community Responsibility" *Journal of Political Economy* 128(5), 2020.

[59] Deb, Joyee, and Julio González-Díaz. "Enforcing social norms: Trust-building and community enforcement." Theoretical Economics 14.4 (2019): 1387-1433.

[60] Deb, Joyee, Takuo Sugaya and Alexander Wolitzky. "The Folk Theorem in Repeated Games with Anonymous Random Matching." *Econometrica* 88.3 (2020): 917-964.

[61] De Quidt, Jonathan, Thiemo Fetzer, and Maitreesh Ghatak. "Group lending without joint liability." Journal of Development Economics 121 (2016): 217-236.

[62] Dilmé, Francesc. "Helping Behavior in Large Societies." *International Economic Review* 57.4 (2016): 1261-1278.

[63] Dixit, Avinash. "Trade Expansion and Contract Enforcement." *Journal of Political Economy* 111.6 (2003a): 1293-1317.

[64] Dixit, Avinsh. "On modes of economic governance." Econometrica 71.2 (2003b): 449-481.

[65] Dixit, Avinash K. Lawlessness and Economics: Alternative Modes of Governance. Princeton University Press, 2004.

[66] Dunbar, Robin. "Gossip in Evolutionary Perspective," *Review of General Psychology* 8.2 (2004): 100-110.

[67] Easterly, William, and Ross Levine. "Africa's growth tragedy: policies and ethnic divisions." The quarterly journal of economics 112.4 (1997): 1203-1250.

[68] Eeckhout, Jan. "Minorities and endogenous segregation." The Review of Economic Studies 73.1 (2006): 31-53.

[69] Ellickson, Robert C. *Order Without Law: How Neighbors Settle Disputes*. Harvard University Press, 1994.

[70] Ellison, Glenn. "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching." *Review of Economic Studies* 61.3 (1994): 567-588.

[71] Fafchamps, Marcel, and Susan Lund. "Risk-sharing networks in rural Philippines." Journal of development Economics 71.2 (2003): 261-287.

[72] Fainmesser, Itay P. "Community structure and market outcomes: A repeated games-in-networks approach." American Economic Journal: Microeconomics 4.1 (2012): 32-69.

[73] Fainmesser, Itay P., and David A. Goldberg. "Cooperation in partly observable networked markets." Games and Economic Behavior 107 (2018): 220-237.

[74] Farrell, Joseph, and Eric Maskin. "Renegotiation in repeated games." Games and economic behavior 1.4 (1989): 327-360.

[75] Fearon, James D., and David D. Laitin. "Explaining interethnic cooperation." American political science review 90.4 (1996): 715-735.

[76] Fehr, Ernst, and Simon Gächter. "Altruistic punishment in humans." Nature 415.6868 (2002): 137-140.

[77] Feigenberg, Benjamin, Erica Field, and Rohini Pande. "The economic returns to social interaction: Experimental evidence from microfinance." Review of Economic Studies 80.4 (2013): 1459-1483.

[78] Feinberg, Matthew, Robb Willer, and Michael Schultz. "Gossip and Ostracism Promote Cooperation in Groups." *Psychological Science* 25.3 (2014): 656-664.

[79] Ford, Lester Randolph, and Delbert R. Fulkerson. "Maximal flow through a network." Canadian journal of Mathematics 8 (1956): 399-404.

[80] Foster, Andrew D., and Mark R. Rosenzweig. "Imperfect commitment, altruism, and the family: Evidence from transfer behavior in low-income rural areas." Review of Economics and Statistics 83.3 (2001): 389-407.

[81] Friedman, James W. "A non-cooperative equilibrium for supergames." The Review of Economic Studies 38.1 (1971): 1-12.

[82] Friedman, Eric J., and Paul Resnick. "The social cost of cheap pseudonyms." Journal of Economics & Management Strategy 10.2 (2001): 173-199.

[83] Frieze, Alan M., and Geoffrey R. Grimmett. "The Shortest-Path Problem for Graphs with Random Arc-Lengths." *Discrete Applied Mathematics* 10.1 (1985): 57-77.

[84] Fudenberg, Drew, and David K. Levine. "An approximate folk theorem with imperfect private information." Journal of Economic Theory 54.1 (1991): 26-47.

[85] Fudenberg, Drew, and David K. Levine. "Efficiency and observability with long-run and short-run players." Journal of Economic Theory 62.1 (1994): 103-135.

[86] Fudenberg, Drew, and David K. Levine. "Continuous time limits of repeated games with imperfect public monitoring." Review of Economic Dynamics 10.2 (2007): 173-192.

[87] Fudenberg, Drew, and David K. Levine. "Repeated games with frequent signals." The Quarterly Journal of Economics 124.1 (2009): 233-265.

[88] Fudenberg, Drew, David Levine, and Eric Maskin. "The Folk Theorem with Imperfect Public Information." Econometrica 62.5 (1994): 997-1039.

[89] Fudenberg, Drew, David Levine, and Wolfgang Pesendorfer. "When are Nonanonymous Players Negligible?" *Journal of Economic Theory* 79.1 (1998): 46-71.

[90] Fudenberg, Drew, and Eric Maskin. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." Econometrica (1986-1998) 54.3 (1986): 533.

[91] Fujiwara-Greve, Takako, and Masahiro Okuno-Fujiwara. "Voluntarily separable repeated prisoner's dilemma." The Review of Economic Studies 76.3 (2009): 993-1021.

[92] Gagnon, Julien, and Sanjeev Goyal. "Networks, markets, and inequality." American Economic Review 107.1 (2017): 1-30.

[93] Ghosh, Parikshit, and Debraj Ray. "Cooperation in Community Interaction without Information Flows." *Review of Economic Studies* 63.3 (1996): 491-519.

[94] Giné, Xavier, and Dean Karlan, "Group Versus Individual Liability: Short and Long Term Evidence from Philippine Microcredit Lending Groups," *Journal of Development Economics*, 107:65-83 (2014).

[95] Glaeser, Edward L., David I. Laibson, José A. Scheinkman, and Christine L. Soutteret "Measuring trust." The quarterly journal of economics 115.3 (2000): 811-846.

[96] Gluckman, Max. "Papers in Honor of Melville J. Herskovits: Gossip and Scandal." *Current Anthropology* 4.3 (1963): 307-316.

[97] Goldlücke, Susanne, and Sebastian Kranz. "Infinitely repeated games with public monitoring and monetary transfers." Journal of Economic Theory 147.3 (2012): 1191-1221.

[98] Granovetter, Mark. "Economic action and social structure: The problem of embeddedness." American journal of sociology 91.3 (1985): 481-510.

[99] Green, Edward J. "Noncooperative price taking in large dynamic markets." Journal of Economic Theory 22.2 (1980): 155-182.

[100] Green, Edward J., and Robert H. Porter. "Noncooperative collusion under imperfect price information." Econometrica: Journal of the Econometric Society (1984): 87-100.

[101] Greif, Avner. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition." *American Economic Review* (1993): 525-548.

[102] Greif, A., Milgrom, P., and Weingast, B. (1994), "Coordination, Commitment, and Enforcement: The Case of the Merchant Guild," *Journal of Political Economy*, 102, 745-776.

[103] Guala, Francesco. "Reciprocity: Weak or strong? What punishment experiments do (and do not) demonstrate." Behavioral and Brain Sciences 35 (2012): 1-59.

[104] Habyarimana, James, et al. "Why does ethnic diversity undermine public goods provision?." American Political Science Review (2007): 709-725.

[105] Hadfield, Gillian K., and Barry R. Weingast. "What is law? A coordination model of the characteristics of legal order." Journal of Legal Analysis 4.2 (2012): 471-514.

[106] Harrington Jr, Joseph E. "Cooperation in a One-Shot Prisoners' Dilemma." *Games and Economic Behavior* 8.2 (1995): 364-377.

[107] Heller, Yuval, and Erik Mohlin. "Observations on Cooperation." *Review of Economic Studies* 85.4 (2018): 2253-2282.

[108] Henrich, Joseph. "Cultural group selection, coevolutionary processes and large-scale cooperation." Journal of Economic Behavior & Organization 53.1 (2004): 3-35.

[109] Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith. "Social distance and other-regarding behavior in dictator games." The American economic review 86.3 (1996): 653-660.

[110] Hörner, Johannes, and Wojciech Olszewski. "The Folk Theorem for Games with Private Almost-Perfect Monitoring." *Econometrica* 74.6 (2006): 1499-1544.

[111] Jackson, Matthew O., Tomas Rodriguez-Barraquer, and Xu Tan. "Social capital and social quilts: Network patterns of favor exchange." American Economic Review 102.5 (2012): 1857-97.

[112] Jackson, Matthew O., and Alison Watts. "Social games: Matching and the play of finitely repeated games." Games and Economic Behavior 70.1 (2010): 170-191.

[113] Jackson, Matthew O., and Yiqing Xing. "The complementarity between community and government in enforcing norms and contracts, and their interaction with religion and corruption." Johns Hopkins Carey Business School Research Paper 18-07 (2020).

[114] Kandori, Michihiro. "The use of information in repeated games with imperfect monitoring." *Review of Economic Studies* 59.3 (1992a): 581-593.

[115] Kandori, Michihiro. "Social Norms and Community Enforcement." *Review of Economic Studies* 59.1 (1992b): 63-80.

[116] Kandori, Michihiro, George J. Mailath, and Rafael Rob. "Learning, mutation, and long run equilibria in games." Econometrica: Journal of the Econometric Society (1993): 29-56.

[117] Kandori, Michihiro, and Hitoshi Matsushima. "Private Observation, Communication and Collusion." *Econometrica* (1998): 627-652.

[118] Kandori, Michihiro, and Shinya Obayashi. "Labor union members play an OLG repeated game." Proceedings of the National Academy of Sciences 111.Supplement 3 (2014): 10802-10809.

[119] Karlan, Dean S. "Social connections and group banking." The Economic Journal 117.517 (2007): F52-F84.

[120] Karlan, D., Möbius, M., Rosenblat, T. and Szeidl, A. (2009), "Trust and Social Collateral," *Quarterly Journal of Economics*, 124, 1307-1361.

[121] Keohane, Robert. *After Hegemony.* 1984.

[122] Kinateder, Markus. "Delayed perfect monitoring in repeated games." International Journal of Game Theory 42.1 (2013): 283-294.

[123] Kiyotaki, Nobuhiro, and Randall Wright. "On money as a medium of exchange." Journal of political Economy 97.4 (1989): 927-954.

[124] Klein, Daniel B. "Promise keeping in the great society: A model of credit information sharing." Economics & Politics 4.2 (1992): 117-136.

[125] Kranton, Rachel E. "The formation of cooperative relationships." The Journal of Law, Economics, and Organization 12.1 (1996a): 214-233.

[126] Kranton, Rachel E. "Reciprocal exchange: a self-sustaining system." The American Economic Review (1996b): 830-851.

[127] Kreps, David M. "Corporate culture and economic theory." Perspectives on positive political economy 90.109-110 (1990): 8.

[128] Laclau, Marie. "A folk theorem for repeated games played on a network." Games and Economic Behavior 76.2 (2012): 711-737.

[129] Laclau, Marie. "Communication in repeated network games with imperfect monitoring." Games and Economic Behavior 87 (2014): 136-160.

[130] Larson, Jennifer M. "Networks and interethnic cooperation." The Journal of Politics 79.2 (2017): 546-559.

[131] Leider, S., Möbius, M. M., Rosenblat, T., & Do, Q. A. (2009). Directed altruism and enforced reciprocity in social networks. The Quarterly Journal of Economics, 124(4), 1815-1851.

[132] Lehrer, Ehud. "Nash equilibria of n-player repeated games with semi-standard information." International Journal of Game Theory 19.2 (1990): 191-217.

[133] Levin, Jonathan. "Relational incentive contracts." American Economic Review 93.3 (2003): 835-857.

[134] Levine, David K., and Salvatore Modica. "Peer discipline and incentives within groups." Journal of economic behavior & organization 123 (2016): 19-30.

[135] Levine, David K., and Wolfgang Pesendorfer. "When are agents negligible?." The American Economic Review (1995): 1160-1170.

[136] Ligon, Ethan, and Laura Schechter. "Motives for sharing in social networks." Journal of Development Economics 99.1 (2012): 13-26.

[137] Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. "Informal insurance arrangements with limited commitment: Theory and evidence from village economies." The Review of Economic Studies 69.1 (2002): 209-244.

[138] Lippert, Steffen, and Giancarlo Spagnolo. "Networks of Relations and Word-of-mouth Communication." Games and Economic Behavior 72.1 (2011): 202-217.

[139] Luttmer, Erzo FP. "Group loyalty and the taste for redistribution." Journal of political Economy 109.3 (2001): 500-528.

[140] MacLeod, W. Bentley, and James M. Malcomson. "Implicit contracts, incentive compatibility, and involuntary unemployment." Econometrica: Journal of the Econometric Society (1989): 447-480.

[141] Mailath, George J., and Andrew Postlewaite. "Asymmetric information bargaining problems with many agents." The Review of Economic Studies 57.3 (1990): 351-367.

[142] Mailath, George J., Stephen Morris, and Andrew Postlewaite. "Laws and authority." Research in Economics 71.1 (2017): 32-42.

[143] Mailath, George J., and Larry Samuelson. Repeated games and reputations: long-run relationships. Oxford university press, 2006.

[144] Malinowski, B. (1922). Argonauts of the Western Pacific: An account of native enterprise and adventure in the Archipelagoes of Melanesian New Guinea. London: Routledge and Kegan Paul

[145] Masten, Scott E., and Jens Prüfer. "On the evolution of collective enforcement institutions: communities and courts." The Journal of Legal Studies 43.2 (2014): 359-400.

[146] Matsushima, Hitoshi. "Efficiency in repeated games with imperfect monitoring." Journal of Economic Theory 48.2 (1989): 428-442.

[147] Matsushima, Hitoshi. "Repeated games with private monitoring: Two players." Econometrica 72.3 (2004): 823-852.

[148] Mauss, Marcel. The gift: The form and reason for exchange in archaic societies. Routledge, 1925.

[149] McAdams, Richard H. The expressive powers of law. Harvard University Press, 2015.

[150] Mertens, Jean-François, and T. Parthasarathy. Equilibria for discounted stochastic games. No. 1987050. Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 1987.

[151] Mihm, Maximilian, and Russell Toth. "Cooperative networks with robust private monitoring." Journal of Economic Theory 185 (2020): 104974.

[152] Mihm, Maximilian, Russell Toth, and Corey Lang. What goes around comes around: A theory of indirect reciprocity in networks. No. 09-07. CAE Working Paper, 2009.

[153] Miguel, E., and Gugerty, M.K. (2005), "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya," *Journal of Public Economics*, 89, 2325-2368.

[154] Milgrom, P.R., North, D.C., and Weingast, B.R. (1990), "The Role of Institutions in the Revival of Trade: the Law Merchant, Private Judges, and the Champagne Fairs," *Economics and Politics*, 2, 1-23.

[155] Myerson, Roger B. "Justice, institutions, and multiple equilibria." Chi. J. Int'l L. 5 (2004): 91.

[156] Nava, Francesco. "Repeated Games and Networks." The Oxford Handbook of the Economics of Networks.

[157] Nava, Francesco, and Michele Piccione. "Efficiency in repeated games with local interaction and uncertain local monitoring." Theoretical Economics 9.1 (2014): 279-312.

[158] Noon, Mike, and Rick Delbridge. "News from Behind my Hand: Gossip in Organizations," *Organization Studies* 14.1 (1993): 23-36.

[159] North, Douglass Cecil. "Institutions, institutional change, and economic performance." (1990).

[160] North, Douglass C., John Joseph Wallis, and Barry R. Weingast. Violence and social orders: A conceptual framework for interpreting recorded human history. Cambridge University Press, 2009.

[161] Nowak, Martin A., and Karl Sigmund. "Evolution of indirect reciprocity by image scoring." Nature 393.6685 (1998): 573-577.

[162] Okuno-Fujiwara, Masahiro, and Andrew Postlewaite. "Social Norms and Random Matching Games." *Games and Economic Behavior* 9.1 (1995): 79-109.

[163] Olson, Mancur. *The logic of collective action: public goods and the theory of group.* Harvard University Press, 1965.

[164] Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press, 1990.

[165] Padilla, A. Jorge, and Marco Pagano. "Sharing default information as a borrower discipline device." European Economic Review 44.10 (2000): 1951-1980.

[166] Pai, Mallesh M., Aaron Roth, and Jonathan Ullman. "An antifolk theorem for large repeated games." ACM Transactions on Economics and Computation (TEAC) 5.2 (2016): 1-20.

[167] Ramey, Garey, and Joel Watson. "Contractual intermediaries." Journal of Law, Economics, and organization 18.2 (2002): 362-384.

[168] Raub, Werner, and Jeroen Weesie. "Reputation and Efficiency in Social Interactions: An Example of Network Effects." *American Journal of Sociology* 96.3 (1990): 626-654.

[169] Renault, J., and Tomala, T. (1998), "Repeated Proximity Games," *International Journal of Game Theory*, 27, 539-559.

[170] Resnick, Paul, and Richard Zeckhauser. "Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system." The Economics of the Internet and E-commerce 11.2 (2002): 23-25.

[171] Rosenthal, Robert W. "Sequences of games with varying opponents." Econometrica: Journal of the Econometric Society (1979): 1353-1366.

[172] Rubinstein, Ariel. "Equilibrium in supergames with the overtaking criterion." Journal of economic theory 21.1 (1979): 1-9.

[173] Sabourian, Hamid. "Anonymous repeated games with a large number of players and random outcomes." Journal of Economic Theory 51.1 (1990): 92-110.

[174] Sadler, Evan. "Ordinal Centrality." *Journal of Political Economy* 130.4 (2022): 926-955.

[175] Sanchez-Pages, Santiago, and Stéphane Straub. "The emergence of institutions." The BE Journal of Economic Analysis & Policy 10.1 (2010).

[176] Sannikov, Yuliy, and Andrzej Skrzypacz. "Impossibility of collusion under imperfect monitoring with flexible production." American Economic Review 97.5 (2007): 1794-1823.

[177] Shapiro, Carl, and Joseph E. Stiglitz. "Equilibrium unemployment as a worker discipline device." The American Economic Review 74.3 (1984): 433-444.

[178] Sigmund, Karl. The calculus of selfishness. Princeton University Press, 2010.

[179] Sobel, Joel. "Interdependent preferences and reciprocity." Journal of economic literature 43.2 (2005): 392-436.

[180] Sommerfeld, Ralf D., Hans-Jürgen Krambeck, Dirk Semmann, and Manfred Milinski. "Gossip as an Alternative for Direct Observation in Games of Indirect Reciprocity." *Proceedings of the National Academy of Sciences* 104.44 (2007): 17435-17440.

[181] Spagnolo, Giancarlo. "Social relations and cooperation in organizations." Journal of Economic Behavior & Organization 38.1 (1999): 1-25.

[182] Sugden, Robert. The economics of rights, co-operation and welfare. Basingstoke: Palgrave Macmillan, 1986.

[183] Sugaya, Takuo, "Folk theorem in repeated games with private monitoring." *Review of Economic Studies*, Forthcoming (2022).

[184] Sugaya, Takuo, and Alexander Wolitzky, "A Few Bad Apples Spoil the Barrel: An Anti-Folk Theorem for Anonymous Repeated Games with Incomplete Information." *American Economic Review* 110: 3817-3835 (2020).

[185] Sugaya, Takuo, and Alexander Wolitzky, "Communication and Community Enforcement." *Journal of Political Economy* 129:2595-2628 (2021).

[186] Sugaya, Takuo, and Alexander Wolitzky, "Informational Requirements for Cooperation." *working paper* (2022a).

[187] Sugaya, Takuo, and Alexander Wolitzky, "Repeated Games with Many Players." *working paper* (2022b).

[188] Sugaya, Takuo, and Alexander Wolitzky, "Bad Apples in Symmetric Repeated Games." *working paper* (2022c).

[189] Szabó, György, and Gabor Fath. "Evolutionary games on graphs." Physics reports 446.4-6 (2007): 97-216.

[190] Tadelis, Steven. "Reputation and Feedback Systems in Online Platform Markets." *Annual Review of Economics* 8 (2016): 321-340.

[191] Takahashi, Satoru. "Community Enforcement when Players Observe Partners' Past Play," *Journal of Economic Theory* 145.1 (2010): 42-62.

[192] Tomala, Tristan. "Fault reporting in partially known networks and folk theorems." Operations Research 59.3 (2011): 754-763.

[193] Udry, Christopher. "Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria." The Review of Economic Studies 61.3 (1994): 495-526.

[194] Vega-Redondo, Fernando. "Building up social capital in a changing world." Journal of Economic Dynamics and Control 30.11 (2006): 2305-2338.

[195] Vigdor, Jacob L. "Community composition and collective action: Analyzing initial mail response to the 2000 census." Review of Economics and Statistics 86.1 (2004): 303-312.

[196] Williamson, Oliver. *The economic institutions of capitalism. Firms, markets, relational contracting.* 1985.

[197] Wolitzky, Alexander. "Cooperation with Network Monitoring." *Review of Economic Studies* 80.1 (2012): 395-427.

[198] Wolitzky, Alexander. "Communication with Tokens in Repeated Games on Networks." *Theoretical Economics* 10.1 (2015): 67-101.

[199] Xing, Yiqing, "Who Shares Risk with Whom and How? Endogenous Matching and Informal Risk Sharing." *working paper* (2020).

[200] Young, H. Peyton. "The evolution of conventions." Econometrica: Journal of the Econometric Society (1993): 57-84.