Massachusetts Institute of Technology

Communication and Community Enforcement

Takuo Sugaya

Stanford Graduate School of Business

Alexander Wolitzky

Massachusetts Institute of Technology

We study the repeated prisoner's dilemma with random matching, a ca-
nonical model of community enforcement with decentralized infor-
mation. We assume that (1) with small probability, each player is a "bad
type" who never cooperates, (2) players observe and remember their
partners' identities, and (3) each player interacts with others frequently
but meets any particular partner infrequently. We show that these as-
sumptions preclude cooperation in the absence of explicit communica-
tion but that introducing within-match cheap talk communication re-
stores cooperation. Thus, communication is essential for community
enforcement.

> He that filches from me my good name
> Robs me of that which not enriches him
> And makes me poor indeed. —*Othello*, 3.3

I.   Introduction

Everyday experience and a wealth of evidence from across the social sci-
ences indicate that communication about the reputation of third parties

is a key mechanism of social cooperation.[1] No one doubts that if they misbehave in a relationship with one (trading, business, romantic) partner, word might spread and they may end up being excluded from valuable future relationships. The threat of communication tomorrow and ostracism the day after keeps us on good behavior today.

While the role of communication in the community enforcement of cooperation seems familiar, it is not well captured by existing game theory models of cooperation in large societies. In the classic community enforcement models of Kandori (1992) and Ellison (1994), where players observe only their partners' actions, cooperation is supported without explicit communication by relying on contagion strategies, a form of collective punishment: whenever a player sees anyone defect, she starts defecting against everyone. Contagion strategies cause the fastest possible breakdown of cooperation following a defection and therefore the harshest punishment for defection. Why then do real-world societies often rely on communication and individualized punishment rather than contagion-like strategies? Are communication and individualized punishment truly necessary for supporting cooperation in large communities, or are they merely quirks of the particular cooperative equilibrium in which we happen to find ourselves?

This paper establishes the necessity of communication and individualized punishment in a version of the standard community enforcement model with more realistic assumptions. First, with small probability, each player is a "bad type" who always defects. In a companion paper (Sugaya and Wolitzky 2020), we show that in games with anonymous players—such as the anonymous prisoner's dilemma studied by Kandori and Ellison—this assumption completely precludes cooperation in large societies, intuitively because collective punishment is too likely to be triggered in the presence of bad types. We therefore consider here the more realistic case where players observe (and remember) their partners' identities, so that individualized punishments (i.e., strategies that condition on the partner's identity) are technologically feasible. With observable identities, the presence of a few bad types obviously poses no obstacle to cooperation when each pair of players interacts frequently, because players can treat the overall repeated game as a collection of two-player

---

[1] Many references can be given. For instance, see Grief (1993), Dixit (2003), and Tadelis (2016) in economics; Raub and Weesie (1990) in sociology; Ostrom (1990) and Ellickson (1994) in political science and law; Gluckman (1963) in anthropology; Noon and Delbridge (1993) in organizational behavior; Baumeister, Zhang, and Vohs (2004), Dunbar (2004), and Feinberg, Willer, and Schultz (2014) in psychology; and Sommerfeld et al. (2007) in evolutionary biology. Many of these papers refer to communication about third parties' behavior as "gossip." This terminology accords with the computer science literature on gossip protocols (Hedetniemi, Hedetniemi, and Liestman 1988; Shah 2009), which we draw on.

games, cooperating with each partner if and only if he has behaved well in their bilateral relationship. We instead assume that, while each player interacts with others frequently (i.e., the discount factor δ is close to one), she meets any particular partner infrequently (i.e., the population size N is much larger than $1/(1-\delta)$). In sum, we consider community enforcement with (1) a small chance of bad types, (2) observable identities, and (3) patient players but infrequent bilateral interactions. We show that cooperation in this environment is impossible in the absence of explicit communication (theorem 1) but becomes possible if within-match cheap talk—ordinary conversation between matched partners—is allowed (theorems 2 and 3).

More precisely, our results hold fixed the payoff parameters of the prisoner's dilemma stage game as well as an ε probability that each player is a "commitment type" who always defects (independent across players) and consider sequences of repeated games where the discount factor δ and the population size N change together. By viewing the overall repeated game as a collection of two-player games, it is trivial to support cooperation among pairs of rational players along any sequence where $(1-\delta)N \rightarrow 0$—that is, whenever bilateral interactions become frequent.[2]

In stark contrast, our first main result (theorem 1) shows that average payoffs converge to the mutual defection payoff along any sequence where $(1-\delta)N \rightarrow \infty$—that is, whenever bilateral interactions become infrequent. The logic of this result combines ideas from repeated game theory and information theory. Roughly speaking, the presence of bad types renders collective punishment ineffective, so incentives can be provided only by individualized punishment. When $(1-\delta)N \rightarrow \infty$, the population is too large for individualized punishment to be executed bilaterally—instead, a player's misbehavior against a partner must affect third parties' behavior toward her. Therefore, to support cooperation, players' actions must convey information about specific individuals' past behavior. This step is where information theory enters the picture: when actions are binary and explicit communication is not allowed, each player receives only one bit of information per period (i.e., the opponent's action). We show that $O(N)$ bits are required to provide significant information about N players' individual actions (lemma 3). Hence, $O(N)$ periods of communication via actions are required to monitor N players' actions. But this speed of communication is too slow to provide meaningful incentives

---

[2] To see why $(1-\delta)N \rightarrow 0$ corresponds to frequent bilateral interactions, suppose that players match once every Δ units of real time with fixed discount rate $r > 0$, so $\delta = e^{-r\Delta}$, and hence $(1-\delta) \approx r\Delta$. Since each pair of players interacts $1/(\Delta \times (N-1)) \approx r/((1-\delta)N)$ times per unit of real time on average, $(1-\delta)N \rightarrow 0$ means that each pair of players interacts frequently, while $(1-\delta)N \rightarrow \infty$ means that each of pair of players rarely interacts.

when $(1 - \delta)N \to \infty$. Thus, neither collective nor individualized punishment is effective, and cooperation is impossible.[3]

We then show that allowing within-match cheap talk restores the possibility of cooperation. Within-match cheap talk makes it technologically feasible for information about everyone's behavior to spread through the population exponentially quickly, reaching all players within $O(\log N)$ periods with high probability. Indeed, we establish that cooperation is possible along any sequence where $(1 - \delta) \log N \to 0$.[4] We first derive a relatively simple version of this result (theorem 2), which shows that cooperation can be achieved as an approximate Nash equilibrium using realistic strategies where (1) each player keeps track of a "blacklist" of opponents whom she believes have ever defected against a rational player, (2) players share their blacklists with each other before taking actions, and (3) each player defects against the opponents on her blacklist. However, such strategies form only an approximate equilibrium, because they break down in the low-probability event that a large fraction of the population consists of bad types; moreover, the possibility that this event can occur may unravel the equilibrium even in situations where no one assigns a high probability to this event.

Our final result (theorem 3) then shows that cooperation can be achieved as an exact sequential equilibrium by combining the simple blacklisting idea of theorem 2 with more complicated, "block belief–free" strategies that prevent unraveling. In our construction, players cooperate only after learning through communication that a large enough fraction of the population is rational. Furthermore, a player who does not learn that there are enough rational types can defect without fear of being punished (in the event that there are many rational types), because in this event subsequent communication will reveal that her defection was justified by her failure to learn.

Finally, the version of our model with communication can be extended by decoupling the rate at which players meet to engage in payoff-relevant interactions and the rate at which they meet to engage in cheap talk communication. In this extended model, we give a fairly complete characterization of how the prospects for cooperation depend on the population size and both meeting rates.

Related literature.—This paper contributes to the literatures on community enforcement, the folk theorem in repeated games, and the role of communication in supporting cooperation. Its most novel features are analyzing how the rates at which $\delta \to 1$ and $N \to \infty$ affect the scope for

---

[3] The same argument applies when each player receives K bits of information per period for any number K fixed independently of N. Thus, theorem 1 holds for many information structures besides the canonical one where players observe only their partners' actions.

[4] Conversely, it is straightforward to show that cooperation is impossible if there exists $\rho > 0$ such that $(1 - \delta)^{1+\rho} \log N \to \infty$.

cooperation in a repeated random matching game, as well as showing that introducing explicit communication dramatically affects the race between δ and N.

The literature on community enforcement in repeated games originates with Kandori (1992) and Ellison (1994).[5] These authors assume complete information (no bad types) and show that the threat of collective punishment via contagion supports cooperation whenever $(1 - \delta) \log N \to 0$. Our companion paper (Sugaya and Wolitzky 2020) shows that collective punishment breaks down in the presence of bad types, which motivates the current paper's focus on individualized punishment.[6]

Bad types are also considered in the community enforcement models of Ghosh and Ray (1996) and Heller and Mohlin (2018), but in these papers bad types help support cooperation, by making players less tempted to cheat their current partners and return to the matching pool (in Ghosh and Ray's voluntary separation model) or by stabilizing grim trigger-like strategies by making the observation that a partner defected in the past informative of his being a bad type (in Heller and Mohlin's model). These papers are therefore less closely related to ours.[7]

In many papers on the folk theorem in repeated games, implicit communication through actions is found to be just as effective as explicit communication through cheap talk when δ is close to one. For example, this is the case in Hörner and Olszewski's (2006) folk theorem with almost perfect monitoring and in Deb, Sugaya, and Wolitzky's (2020) folk theorem for anonymous random matching games. In contrast, implicit and explicit communication are not equivalent in our model, because we take $\delta \to 1$ and $N \to \infty$ simultaneously (so communication speed matters) and explicit communication allows more information to be transmitted in each meeting.[8]

Several papers on community enforcement and repeated games on networks fix $\delta < 1$ and show that cooperation is easier to support when the news that a defection occurred spreads more quickly (Raub and Weesie 1990; Klein 1992; Ahn and Suominen 2001; Dixit 2003; Lippert and Spagnolo 2011; Ali and Miller 2013; Wolitzky 2013; Balmaceda and

---

[5] See also Harrington (1995) and Okuno-Fujiwara and Postlewaite (1995).

[6] Kandori and Ellison were well aware of the importance of bad types but did not include them in their models. For example, Ellison wrote, "If one player were 'crazy' and always played D [defect] . . . contagious strategies would not support cooperation. In large populations, the assumption that all players are rational and know their opponents' strategies may be both very important to the conclusions and fairly implausible" (578).

[7] One result closer in spirit to ours is Heller and Mohlin's theorem 1, which shows that cooperation is impossible in the "offensive" (submodular) prisoner's dilemma with bad types, while Takahashi (2010) showed that cooperative "belief-free" equilibria exist in this setting without bad types. Dilmé (2016) considers a similar model where cooperation is robust to introducing a small measure of bad types.

[8] Deb (2020) establishes a folk theorem for anonymous random matching games with explicit communication. In these games, communication has the distinctive role of serving to relax anonymity, as players can identify each other via endogenous "names." This role does not arise in our model with observable player identities.

Escobar 2017). This force differs from the role of communication in our model, where introducing within-match communication does not increase the speed at which the community learns that *someone* defected but rather enriches the information that can be transmitted in each match, so that the community learns faster *which* players defected. The value of communication is thus tied to the need to use individualized rather than collective punishment, which in turn is necessitated by the presence of bad types (which are absent in the above papers).[9] Finally, a number of papers consider settings where the need to provide incentives for honest communication constrains community enforcement (Bowen, Kreps, and Skrzypacz 2013; Wolitzky 2015; Ali and Miller 2016, 2020; Barron and Guo 2021). While we of course also insist that communication is incentive compatible, theorem 3 shows that this constraint is ultimately not binding in our model.

II.   Model

A set $I = \{1, ..., N\}$ of N players interacts in discrete time, $t = 1, 2, ...,$ with N even. Each period, the players match in pairs, uniformly at random and independently across periods, to play the prisoner's dilemma:

	C	D
C	1, 1	$-L, 1 + G$
D	$1 + G, -L$	0, 0

where $G, L > 0$ and $G < 1 + L$, so D is strictly dominant but (C, C) maximizes the sum of stage-game payoffs.

Each player is rational with probability $1 - \varepsilon$ and *bad* with probability ε, for some $\varepsilon \in (0, 1)$, independently across players.[10] The number of bad players thus follows a binomial distribution. Rational players maximize expected discounted payoffs with discount factor $\delta \in (0, 1)$. Bad players always play D.

Matching is nonanonymous. That is, at the beginning of each period t, every player i observes the identity (but not the type) of her period t partner, which we denote by $\mu_{i,t} \in I \setminus \{i\}$. A player then chooses her own action $a_{i,t} \in \{C, D\}$ and finally observes her partner's action $a_{\mu_{i,t},t}$ at the end of the period. Thus, player i's history at the beginning of period t is $h_i^t = ((\mu_{i,\tau}, a_{i,\tau}, a_{\mu_{i,\tau},\tau})_{\tau=1}^{t-1}, \mu_{i,t})$, with $h_i^1 = \mu_{i,1}$. In section IV, we augment the game by allowing preplay cheap talk communication within each match. The description of a history for player i will then also include

---

[9] A very different role for explicit communication with high δ (and small N) is analyzed by Awaya and Krishna (2016, 2019). They show that communication can in effect improve monitoring by exploiting correlation between players' signals.

[10] We discuss generalizations to multiple "commitment types" and to correlated types in sec. V.A.

the history of messages sent by player i to her partners and received by player i from her partners.

A strategy σ_i for player i maps histories h_i^t to $\Delta(\{C, D\})$ for each t. The interpretation is that player i plays $\sigma_i(h_i^t)$ at history h_i^t when rational; when bad, she always plays D. Given a strategy profile $\sigma = (\sigma_i)_i$, denote player i's expected discounted per-period payoff, conditional on the event that she is rational, by U_i. Our measure of population average payoffs is $U = \Sigma_i U_i / N$, which is the average over players of their expected payoffs conditional on the event that they are rational.[11] Note that since the minmax payoff is zero, bad types always take D, and the maximum sum of stage-game payoff is 2, we have $U \in [0, 1]$ in any Nash equilibrium. Since the payoff from mutual defection is $(0, 0)$, we say that (some) cooperation arises if and only if $U \neq 0$.

For all of our results, we fix the stage-game payoff parameters G and L and the "commitment probability" ε and simultaneously vary the population size N and the discount factor δ. It is fairly trivial to see that cooperation can occur in a Nash (or sequential) equilibrium if $(1 - \delta)N \to 0$, even without cheap talk communication, and that cooperation cannot occur in any Nash equilibrium if there exists $\rho > 0$ such that $(1 - \delta)^{1+\rho} \log N \to \infty$, even with within-match cheap talk.[12] In contrast, our main results show that without cheap talk cooperation cannot occur if $(1 - \delta) N \to \infty$ and that with cheap talk a folk theorem holds if $(1 - \delta) \log N \to 0$.

III.   No Cooperation without Communication

We first show that cooperation without communication is impossible when bilateral interactions are infrequent.

THEOREM 1.   For any sequence (N, δ) where $(1 - \delta)N \to \infty$ and any corresponding sequence of Nash equilibrium population payoffs (U), we have $U \to 0$.

The intuition is that (1) under incomplete information ($\varepsilon > 0$), cooperation requires separately monitoring the actions of $O(N)$ players, (2) implicit communication via actions can convey only one bit of information per period, (3) $O(N)$ bits—and hence $O(N)$ periods of communication—are required to monitor the actions of $O(N)$ players, and (4) the promise of reward or punishment $O(N)$ periods in the future is insufficient to motivate cooperation when $(1 - \delta)N$ is large.

---

[11] By focusing on players' payoffs conditional on being rational, we avoid the need to specify utility functions for bad types. An earlier version of this paper instead assumed that bad types have the same utility function as rational types (while being constrained to always play D) and considered ex ante expected payoffs. Our results hold verbatim for this alternative notion of average payoffs, except that the payoffs supported in theorem 3 must be adjusted by an $O(\varepsilon)$ term.

[12] We establish the former result in sec. III.C and the latter (which is similar to proposition 3 of Kandori 1992) in sec. IV.

We prove theorem 1 in section III.A, deferring some details to appendix A. In section III.B, we discuss how the theorem extends when players observe additional information beyond their current partner's action. We then present a partial converse—a folk theorem when $(1 - \delta)N \to 0$—in section III.C.

A.   Proof of Theorem 1

We establish the stronger conclusion that $U \to 0$ for any sequence of strategy profiles σ such that each player i obtains a higher expected payoff from playing σ_i than from always playing D. That is, we relax the requirement that σ is a Nash equilibrium to the weaker requirement that each player i prefers σ_i to the strategy *Always Defect*. Furthermore, we show that this conclusion holds even when we consider "extended" strategy profiles $\sigma = (\sigma_i)_{i \in I}$, where σ_i can condition player i's action not only on her own partners' identities but on the entire match realization. We thus allow more potential equilibrium strategies than are actually available to the players while requiring that fewer potential deviations are unprofitable.

Slightly abusing notation, let $\mu^t = (\mu_{i,\tau})_{i \in I, \tau \leq t}$ denote the first t periods of the match realization, let $h_i^t = (a_{i,\tau}, a_{\mu_{i,\tau}, \tau})_{\tau=1}^{t-1}$ denote the history of player i's own actions and past opponents' actions at the beginning of period t, and let σ_i denote a mapping from (h_i^t, μ^t) to a mixed action. That is, $\sigma_i(h_i^t, \mu^t)$ is the (possibly mixed) action taken by player i in period t at history (h_i^t, μ^t). (Note that μ^t includes the identity of i's period t partner.) Fix such an (extended) strategy profile $\sigma = (\sigma_i)_i$.

Let 0_i (respectively, 1_i) denote the event that player i is rational (respectively, bad). For any $x_i \in \{0_i, 1_i\}$ and $x_j \in \{0_j, 1_j\}$, let $\Pr(h_i^t, h_j^t | x_i, x_j, \mu^t)$ denote the probability that, under strategy profile σ, h_i^t and h_j^t are the period t histories of player i and player j, conditional on the event (x_i, x_j) and the event that the first t periods of the match realization are given by μ^t.

When player i's opponents play $(\sigma_j)_{j \neq i}$, the distribution over paths of play of the repeated game when $x_i = 0_i$ but i deviates to *Always Defect* is the same as that when $x_i = 1_i$ (in which case i is forced to play *Always Defect*). We call the condition that the rational type of player i prefers her equilibrium strategy σ_i to the strategy *Always Defect* "incentive compatibility," since it requires precisely that this type prefers to follow her own equilibrium strategy rather than the bad type's strategy. This condition appears in appendix A as equation (A1). The first step of the proof of theorem 1 puts this condition in a more convenient form and averages it over players $i \in I$. (Proofs of lemmas are deferred to app. A. In the following, for $x \in \mathbb{R}$, $(x)_+ := \max\{x, 0\}$.)

Lemma 1.   If each player i prefers strategy σ_i to *Always Defect* (i.e., if eq. [A1] holds for all $i \in I$), then

$$(1 - \delta)\sum_t \delta^{t-1} \sum_{\mu^t} \Pr(\mu^t) \sum_i \frac{1}{N} \sum_{h_i^t} \Pr(h_i^t | 0_i, \mu^t) \Pr(\sigma_i(h_i^t, \mu^t) = C) \min\{G, L\}$$

$$\leq (1 - \varepsilon)(1 - \delta)\sum_t \delta^{t-1} \sum_{\mu^t} \Pr(\mu^t)$$

$$\times \sum_{i,j \neq i} \frac{1}{N(N-1)} \sum_{h_j^t} \left(\Pr(h_j^t | 0_i, 0_j, \mu^t) - \Pr(h_j^t | 1_i, 0_j, \mu^t)\right) + (1 + G).$$

$$(1)$$

Intuitively, the left-hand side of (1) is a lower bound on the average over players i of the "cooperation cost" that player i incurs by following strategy σ_i rather than *Always Defect*, and the right-hand side is an upper bound on the average over players i of the "benefit from averted punishment" that player i gains by following σ_i rather than *Always Defect*. The heart of the proof of theorem 1 consists of showing that the average benefit from averted punishment goes to zero if $(1 - \delta)N \to \infty$. Roughly speaking, this amounts to showing that the (expected, discounted, per-player average) influence of player i's type on the histories of players $j \neq i$ is small.

LEMMA 2. If $(1 - \delta)N \to \infty$, then the average benefit from averted punishment (the right-hand side of [1]) goes to zero.

To see that lemmas 1 and 2 imply the theorem, note that

$$U_i \leq (1 - \delta)\sum_t \delta^{t-1} \sum_{\mu^t} \Pr(\mu^t) \sum_{j \neq i} \frac{1}{N-1} \sum_{h_j^t} (1 - \varepsilon)$$

$$\times \Pr(h_j^t | 0_i, 0_j, \mu^t) \Pr(\sigma_j(h_j^t, \mu^t) = C)(1 + G),$$

because player i's payoff would equal the right-hand side of this inequality if her stage-game payoff were $1 + G$ whenever her partner takes C and zero whenever her partner takes D. By Bayes' rule, $\Pr(h_j^t | 0_i, 0_j, \mu^t) \leq \Pr(h_j^t | 0_j, \mu^t)/(1 - \varepsilon)$. Hence,

$$U_i \leq (1 + G)(1 - \delta)\sum_t \delta^{t-1} \sum_{\mu^t} \Pr(\mu^t) \sum_{j \neq i} \frac{1}{N-1}$$

$$\times \sum_{h_j^t} \Pr(h_j^t | 0_j, \mu^t) \Pr(\sigma_j(h_j^t, \mu^t) = C).$$

Averaging over $i \in I$, we have

$$U \leq (1 + G)(1 - \delta)\sum_t \delta^{t-1} \sum_{\mu^t} \Pr(\mu^t) \sum_j \frac{1}{N} \sum_{h_j^t} \Pr(h_j^t | 0_j, \mu^t) \Pr(\sigma_j(h_j^t, \mu^t) = C).$$

Since lemmas 1 and 2 imply that the right-hand side of this inequality goes to zero if $(1 - \delta)N \to \infty$, we have $U \to 0$ as well.

The proof of lemma 2 relies on the following information theory result, which implies that the average influence of N players' types on a k-dimensional binary random variable is nonnegligible only if $k/N \nrightarrow 0$.[13]

LEMMA 3.   Let X_1, X_2, $\ldots$, X_N be i.i.d. binary random variables, with $\Pr(X_i = 1) = \Pr(1_i) = \varepsilon$, and let S be a k-dimensional binary random variable defined on the same probability space. Let $\underline{\varepsilon} = \min\{\varepsilon, 1 - \varepsilon\}$. Then

$$\sum_{i=1}^{N} \sum_{s \in \{0,1\}^k} (\Pr(s|0_i) - \Pr(s|1_i))_+ \; \leq \; \sqrt{\frac{\log(2)kN}{\underline{\varepsilon}}}. \tag{2}$$

Let us provide some intuition and a proof sketch for lemma 3. The $k = 1$ case is the relatively familiar result that the average influence of N independent "votes" on a binary outcome is maximized by majority rule: if we let $s = 1$ if and only if $\{i : X_i = 1\} \geq \varepsilon N$, each voter i is pivotal with probability approximately $1/\sqrt{N}$, and summing over voters gives a "total influence" of $\sum_{i=1}^{N}\sum_{s \in \{0,1\}}(\Pr(s|0_i) - \Pr(s|1_i))_+ \approx \sum_{i=1}^{N} 1/\sqrt{N} = \sqrt{N}$.[14] The lemma asserts that in general the total influence of N independent votes on k binary outcomes is bounded by approximately $\sqrt{kN}$. This bound can be attained by splitting the population into k equal-sized groups and running majority rule within each group; each voter is then pivotal with probability approximately $1/\sqrt{N/k}$, and summing over voters gives a total influence of approximately $\sum_{i=1}^{N} 1/\sqrt{N/k} = \sqrt{kN}$. Intuitively, the bound is tight because more complex signals introduce correlation between the different signal dimensions, which reduces the average influence of a vote.

The proof of lemma 3 proceeds as follows. First, we use Pinsker's inequality and some manipulations to show that the influence of i's vote, $\sum_{s \in \{0,1\}^k}(\Pr(s|0_i) - \Pr(s|1_i))_+$, is at most $\sqrt{\log(2)I(S; X_i)/\underline{\varepsilon}}$, where $I(S; X_i)$ denotes the mutual information between S and X_i, measured in bits. Elementary properties of mutual information, together with independence of the X_i's, imply that $\sum_i I(S; X_i)$ is at most the entropy of S, which in turn is at most k since S is a k-dimensional binary random variable. Therefore, the sum of the squared influences is at most $\log(2)\sum_i I(S; X_i)/\underline{\varepsilon} \leq \log(2)k/\underline{\varepsilon}$, and hence the sum of the influences is at most $\sqrt{\log(2)kN/\underline{\varepsilon}}$ by the $\ell_1 - \ell_2$ norm inequality.[15]

---

[13] We thank Omer Tamuz for suggesting a proof of this lemma.

[14] See, e.g., lemma A of Fudenberg, Levine, and Pesendorfer (1998) and theorem 2 of Al-Najjar and Smorodinsky (2000).

[15] One could also try to prove lemma 3 by induction on k. This approach easily gives a bound of order $k\sqrt{N}$. This approach is used by Awaya and Krishna (2016, lemma 4; 2019, lemma A.1) and is also reminiscent of proposition 1 of Fudenberg, Levine, and Pesendorfer (1998). However, lemma 3 requires a bound of order $\sqrt{kN}$, which seems difficult to establish by induction.

The proof of the theorem is completed by showing that lemma 3 implies lemma 2. This final step uses one more mathematical fact, which is that $\sum_{t=1}^{\infty}\delta^t\sqrt{t} \leq (1-\delta)^{-3/2}$ for all $\delta \in (0,1)$.[16] Now, since $h_j^t = (a_{j,\tau}, a_{\mu_{j,\tau},\tau})_{\tau=1}^{t-1}$ is a $2(t-1)$-dimensional binary random variable whose distribution, conditional on $x_j = 0_j$ and μ^t, depends on the $N-1$ binary random variables $(X_i)_{i\neq j}$ (which are themselves independent conditional on $x_j = 0_j$ and μ^t), we have

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{i,j\neq i}\frac{1}{N(N-1)}\sum_{h_j^t}\left(\Pr\left(h_j^t|0_i,0_j,\mu^t\right) - \Pr\left(h_j^t|1_i,0_j,\mu^t\right)\right)_+$$

$$= (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_j\frac{1}{N(N-1)}$$
$$\sum_{i\neq j}\sum_{h_j^t}\times\left(\Pr\left(h_j^t|0_i,0_j,\mu^t\right) - \Pr\left(h_j^t|1_i,0_j,\mu^t\right)\right)_+$$

$$\leq (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\frac{N}{N(N-1)}\sqrt{\frac{2\log(2)(t-1)(N-1)}{\underline{\varepsilon}}} \quad \text{(by [2])} \tag{3}$$

$$= \sqrt{\frac{2\log(2)}{\underline{\varepsilon}}}\frac{1}{\sqrt{N-1}}(1-\delta)\sum_{t=1}^{\infty}\delta^t\sqrt{t}$$

$$\leq \sqrt{\frac{2\log(2)}{\underline{\varepsilon}}}\frac{1}{\sqrt{N-1}}\frac{1}{\sqrt{1-\delta}}.$$

Hence, if $(1-\delta)N \to \infty$, then the right-hand side of (1) goes to zero.

B. Additional Signals

While theorem 1 concerns the canonical repeated prisoner's dilemma with nonanonymous random matching, where players observe only their current partner's identity and action, it also extends to richer settings where players observe additional signals. Suppose that at the beginning of every period t, each player i privately observes an additional K-dimensional binary signal $s_{i,t} \in \{0,1\}^K$, where the distribution of $s_{i,t}$ depends on the entire history of matches, actions, and signals up to the end of period $t-1$, $(\mu_{i,\tau}, a_{i,\tau}, s_{i,\tau})_{i\in I, \tau\leq t-1}$, and these signals can be arbitrarily correlated across players. For example, each player might observe the most recent action of each of K randomly chosen players in the population, or she might observe in which one out of 2^K possible "bins" lies the total number of

---

[16] This follows because, since $\sqrt{\cdot}$ is a concave function, Jensen's inequality gives $((1-\delta)/\delta)\sum_{t=1}^{\infty}\delta^t\sqrt{t} \leq \sqrt{((1-\delta)/\delta)\sum_{t=1}^{\infty}\delta^t t} = \sqrt{((1-\delta)/\delta)(\delta/((1-\delta)^2))} = 1/\sqrt{1-\delta}$, and hence $\sum_{t=1}^{\infty}\delta^t\sqrt{t} \leq \delta(1-\delta)^{-3/2} \leq (1-\delta)^{-3/2}$.

players who cooperated last period, $|\{j \in I : a_{j,t-1} = C\}|$.[17] Theorem 1 extends to this more general model as follows.

THEOREM $1'$. Suppose that each player observes an additional K-dimensional binary signal in every period, for some $K \in \mathbb{N}$ that may vary together with N and δ. For any sequence (N, δ, K) where $(1 - \delta)N/K \to \infty$ and any corresponding sequence of Nash equilibrium population payoffs (U), we have $U \to 0$.

Proof. If we redefine h_i^t as $(a_{i,\tau}, a_{\mu_{i,\tau},\tau}, s_{i,\tau})_{\tau=1}^{t-1}$ rather than $(a_{i,\tau}, a_{\mu_{i,\tau},\tau})_{\tau=1}^{t-1}$, the proof of theorem 1 goes through as written, except that h_i^t now has dimension $(2 + K)(t - 1)$ rather than $2(t - 1)$. This means that (3) must be replaced by

$$\sqrt{\frac{(2 + K)\log(2)}{\varepsilon}} \frac{1}{\sqrt{N - 1}} \frac{1}{\sqrt{1 - \delta}}.$$

If $(1 - \delta)N/K \to \infty$, then this expression goes to zero, which implies the conclusion of theorem 1. QED

Since bilateral interactions between any two players become infrequent if $(1 - \delta)N \to \infty$, a corollary of theorem $1'$ is that, if players observe K-dimensional signals for any $K \in \mathbb{N}$ fixed independently of N and δ, cooperation is impossible when bilateral interactions are infrequent.

Theorem $1'$ also implies that the cheap talk communication we introduce in section IV can make a difference only if communication can convey an unbounded amount of information in a single period. This can be achieved by, for example, having each player convey a binary summary of the "reputation" of every player in the population to her current partner. The same conclusion also applies in situations (which are outside our model) where players have rich action sets and implicitly communicate through the fine details of their actions, as when bidders in an auction communicate via the trailing digits of their bids (Cramton and Schwartz 2000).

C. A Converse: Cooperation with Frequent Bilateral Interactions

We provide a partial converse to theorem 1, which shows that the presence of bad types does not hinder cooperation when bilateral interactions

---

[17] Note that the signal is not allowed to depend on the identity of player i's period t partner, $\mu_{i,t}$. This seemingly small point is actually very important. For example, if each player is exogenously informed of the social standing of her current partner, then cooperation can be supported even if $(1 - \delta)N$ is large (and indeed, even if the population is infinite). Models of this type, which are quite different from ours, are studied by Kandori (1992, sec. 5), Okuno-Fujiwara and Postlewaite (1995), and Clark, Fudenberg, and Wolitzky (2021), among others.

are frequent. The idea is simply to view the overall repeated game as a collection of $N(N-1)/2$ bilateral relationships, one for each pair of players, and use grim trigger strategies within each bilateral relationship.

Let $F = \text{co}\{(0,0),(1,1),(1+G,-L),(-L,1+G)\}$ denote the convex hull of the feasible payoff set in the two-player prisoner's dilemma. Let $F^{\eta} = \{(v_1,v_2) \in F : v_1, v_2 \geq \eta\}$ denote the set of feasible payoffs where each player receives payoff at least $\eta > 0$. Given parameters (N, δ), let $E \subseteq \mathbb{R}^N$ denote the set of rational player sequential equilibrium payoff vectors; that is, $(v_i)_i \in E$ if there exists a sequential equilibrium where $U_i = v_i$ for all $i \in I$.

PROPOSITION 1. Fix a constant $\eta > 0$ and a sequence $(N, \delta)_l$ indexed by $l \in \mathbb{N}$ satisfying $\lim_{l \to \infty} (1 - \delta_l)N_l = 0$. For each $l \in \mathbb{N}$ and each $i, j \in I_l$ with $i \neq j$, fix $(v_{i,j}, v_{j,i}) \in F^{\eta}$. There exists $\bar{l} > 0$ such that, for all $l > \bar{l}$, in the game with parameters (N_l, δ_l) there is a payoff vector $v \in E_l$ satisfying

$$\left| \left(\frac{1}{N_l - 1} \sum_{j \in I \setminus \{i\}} (1 - \varepsilon)v_{i,j} \right) - v_i \right| < \varepsilon\eta \text{ for all } i \in I_l.$$

Proof. See appendix A, section C. QED

The proof of proposition 1 uses bilateral grim trigger strategies to show that any profile of bilaterally feasible and strictly individually rational payoff pairs is sustainable in sequential equilibrium when bilateral interactions are frequent. We conjecture that an even larger set of payoffs can be supported using more complex strategies. For example, player 1 may be willing to accept a negative present value payoff in her relationship with player 2 if she is compensated with a positive payoff in her relationship with player 3. In principle, such payoff vectors can be supported by having players occasionally communicate implicitly via actions.[18] We do not pursue such a result here.

One can also consider the case where $\delta \to 1$ and $N \to \infty$ at the same rate, so the bilateral interaction frequency stays constant. Here partial cooperation is possible: the maximum equilibrium value of U can exceed zero, but a folk theorem typically does not hold (as follows from applying the proof of theorem 1 when $(1 - \delta)N$ is constant but large).

IV. Cooperation with Communication

We now show that if players can exchange cheap talk messages with their partners before taking actions, cooperation is possible whenever $(1 - \delta) \log N \to 0$. We assume that the set of possible messages is finite but can be taken arbitrarily large relative to the population size, N. Rational types communicate strategically to maximize their expected utility. As for

---

[18] Of course, such communication would have to be incentivized.

bad types, we require that their communication strategy does not affect the distribution of opposing actions they face in any period. This implies that if we endowed bad types with utility functions that depend on the sequence of opposing actions they face, the strategies we construct would form an equilibrium for any such functions (although our analysis does not involve specifying utilities for bad types).[19]

We prove versions of our result for two solution concepts: robust η-Nash equilibrium and robust (exact) sequential equilibrium, where in both cases the word "robust" refers to the property that a bad type's communication strategy does not affect the distribution of actions she faces in any period. In this section, a strategy σ_i for player i specifies, as a function of her type and history, both a distribution over messages to send to her partner and a distribution of actions, where the specified action is always D when the player's type is bad.

DEFINITION 1.   For any $\eta > 0$, a strategy profile $\sigma = (\sigma_i)_{i \in I}$ is a robust η-Nash equilibrium if the following conditions hold for each $i \in I$:

1. If player i is rational ($x_i = 0$), any deviation by player i improves her expected per-period payoff by at most η: for any strategy $\tilde{\sigma}_i$,

$$\mathbb{E}^\sigma\left[(1 - \delta)\sum_{t=1}^\infty \delta^{t-1} u(a_{i,t}, a_{\mu_{i,t},t})\big|x_i = 0\right]$$
$$\geq \mathbb{E}^{\tilde{\sigma}_i, \sigma_{-i}}\left[(1 - \delta)\sum_{t=1}^\infty \delta^{t-1} u(a_{i,t}, a_{\mu_{i,t},t})\big|x_i = 0\right] - \eta.$$

2. If player i is bad ($x_i = 1$), her communication strategy does not affect the distribution of actions she faces in any period: for any strategy $\tilde{\sigma}_i$ (which always takes D when $x_i = 1$),

$$\Pr^\sigma(a_{\mu_{i,t},t} = C|x_i = 1) = \Pr^{\tilde{\sigma}_i, \sigma_{-i}}(a_{\mu_{i,t},t} = C|x_i = 1).$$

This is a permissive version of approximate equilibrium, because (rational) players' gains from deviating are required to be small only in ex ante terms. This permissiveness allows a simple equilibrium construction, which we subsequently complement with a more complicated construction for exact sequential equilibrium.

THEOREM 2.   Fix a sequence $(N, \delta)_l$ satisfying $(1 - \delta) \log N \to 0$. With cheap talk, for any $\eta > 0$ there exists $\bar{l} > 0$ such that, for every $l \geq \bar{l}$, in the game with parameters (N_l, δ_l) there is a robust η-Nash equilibrium in which rational players always cooperate with each other along the equilibrium path of play.

Our proof of this result (in app. A, sec. D) involves strategies that seem fairly realistic. Each player keeps track of a blacklist of players (other than

[19] An even more demanding equilibrium concept would require that the rational types' strategies remain optimal against any communication strategies for the bad types. Our constructions do not satisfy this stronger requirement.

herself) who she believes have ever previously played D against a rational opponent. Every period, players communicate their blacklists to their partners and then update their own blacklists before taking actions. Players take C against opponents who are not on their blacklists and take D against opponents on their blacklists.

To see that these strategies form an η-Nash equilibrium whenever $(1 - \delta) \log N \to 0$, note that since the only consequence of communication is that players stop cooperating with opponents on their blacklists, a player is always indifferent between reporting any two blacklists to her partner (recalling that a player never appears on her own blacklist). Moreover, in equilibrium all blacklisted players are indeed bad, so it is optimal to take D against players on one's blacklist. It remains to show that it is approximately optimal to take C against players who are not on one's blacklist (on the equilibrium path).

To show this, observe that if a player defects against a rational opponent, she is added to his blacklist, and her blacklisted status then spreads through the population exponentially quickly, regardless of her own future behavior. We formalize this observation with a lemma on diffusion processes. Consider random matching among N agents. For each $i \in I$, agent i's state at time t is $s_{i,t} = (s_{i,j,t})_{j \in I} \in \{0, 1\}^N$. When agents i and i' match at time t, the jth component of each of their states updates to

$$s_{i,j,t+1} = s_{i',j,t+1} = \max\{s_{i,j,t}, s_{i',j,t}\} \text{ if } j \notin \{i, i'\},$$

$$s_{i,j,t+1} = s_{i,j,t} \text{ and } s_{i',j,t+1} = s_{i',j,t} \qquad \text{if } j \in \{i, i'\}.$$

An interpretation is that $s_{i,j,t} = 1$ means that agent i knows a rumor about agent j at time t, and agents share all the rumors they know with their partners, except for rumors concerning themselves. The following lemma says that if agents i and $i + 1$ initially know a rumor about i (for each i), then the probability that all N agents know all N rumors by time T converges to one exponentially in T.

LEMMA 4. Consider the above diffusion process with $s_{i,i,1} = s_{i,i+1,1} = 1$ and $s_{i,j,1} = 0$ for all i and $j \notin \{i, i + 1\}$. There exist constants $c > 0$ and $Z > 0$ (independent of N) such that, for all $T > Z \log N$,

$$\Pr(s_{i,j,T} = 1 \text{ for all } i, j \in I) \geq 1 - \exp\left(-\frac{cT}{\log N}\right).$$

Frieze and Grimmett (1985) prove a similar result when each agent shares a rumor with a randomly selected receiver, rather than having players meet in pairs as in the current model.[20] Since pairwise matching

yields a different stochastic process for the number of informed players, we provide a complete proof in appendix B, section B.1 (app. B is available online). The basic idea is the same as in Frieze and Grimmett (1985). So long as most players are uninformed, informed players are unlikely to meet each other, so the number of informed players grows exponentially. Then, once most players are informed, uninformed players are unlikely to meet each other, so the number of uninformed players shrinks exponentially. The spread of each rumor thus approximates a logistic curve, as in standard diffusion models.

By lemma 4, a player who takes D against a rational opponent is very likely to find herself completely excluded from cooperation within $O(\log N)$ periods. Hence, if $(1 - \delta) \log N \approx 0$, taking D against a rational opponent is unprofitable (at an on-path history where one has not yet been blacklisted). Moreover, lemma 4 also implies that all bad types are very likely to be blacklisted within $O(\log N)$ periods. Hence, if $(1 - \delta) \log N \approx 0$, a rational player's payoff loss from cooperating with not-yet-blacklisted bad types is small. The prescribed strategies thus form an η-Nash equilibrium.

Several challenges arise in attempting to transform this η-Nash equilibrium into an (exact) sequential equilibrium. The most serious concerns the low-probability event where a player learns that the fraction of rational types in the population is actually much smaller than $1 - \varepsilon$. In the extreme, suppose that player 1 observes (and/or is told about) a large number of defecting players and eventually comes to believe that player 2 is the only other rational player in the population. When player 1 subsequently meets player 2, if $(1 - \delta)N \approx \infty$, she should play D against him even if he is not on her blacklist, because players 1 and 2 now effectively find themselves in a two-player repeated game with discount factor $\delta^{N-1} \approx 0$ (since they meet on average once every $N - 1$ periods). Moreover, this problem cannot easily be avoided by specifying that players take D if they learn that there are few other rational players: under such strategies a player must assess whether her opponents believe that there are few rational players, whether they believe that their opponents believe this, and so on, and the equilibrium can easily unravel.

We therefore need a more sophisticated approach to construct an exact Nash or sequential equilibrium. The basic idea is to prescribe cooperation only after a player learns through communication that a large enough fraction of the population is rational, while preventing unraveling by excusing players who defect at "erroneous" histories where they failed to learn that there are enough rational types. To identify when a player's history was erroneous in this sense, her opponents must aggregate their information about her history, which by lemma 4 can be achieved in $O(\log N)$ periods with high probability. Note that a given player's opponents can collectively identify her history through their own past observations; moreover, by making their continuation payoffs independent of

whether she is rewarded or punished, they can be induced to communicate this information honestly.[21] This proof approach lets us support a wide range of payoffs as sequential equilibria. However, not surprisingly, the strategies used in the proof (in app. B, sec. B.2) are much more complicated than those used to prove theorem 2.

To state this more general theorem, first fix N and δ and denote the (random) set of rational players by $\theta^* \subseteq I$. For each θ^*, let $F(\theta^*) \subseteq \mathbb{R}^{|\theta^*|}$ denote the set of payoff profiles for the rational players that are feasible when rational types always take D against bad types. That is, letting $\mathbf{a}_i : \{-i\} \to \{C, D\}$ specify an action for player i as a function of her opponent's identity, player i's expected payoff as a function of $\mathbf{a} = (\mathbf{a}_j)_{j \in I}$ equals $\hat{u}_i(\mathbf{a}) = (1/(N-1))\Sigma_j u_i(\mathbf{a}_i(j), \mathbf{a}_j(i))$. We then define $F(\theta^*) = \text{co}(\{\hat{u}_i(\mathbf{a})\}_{i \in \theta^*, \mathbf{a} \in \mathbf{A}(\theta^*)}) \subseteq \mathbb{R}^{|\theta^*|}$, where $\mathbf{A}(\theta^*) = \{\mathbf{a} : \mathbf{a}_i(j) = D \text{ if } i \notin \theta^* \text{ or } j \notin \theta^*\}$. Let $F^*(\theta^*) = F(\theta^*) \cap \mathbb{R}_+^{|\theta^*|}$ denote the set of individually rational payoff profiles in $F(\theta^*)$. Note that $F^*(\theta^*)$ implicitly depends on N but not on δ. Finally, for any $\theta^* \in \{0, 1\}^N$, $\alpha \in (0, 1 - \varepsilon)$, and $\eta \in (0, 1)$, we define the set $F^{\alpha, \eta}(\theta^*) \subseteq F^*(\theta^*)$ as follows:

1. If $|\theta^*| \geq \alpha N$, then

$$F^{\alpha, \eta}(\theta^*)$$
$$= \left\{ v^{\theta^*} \in \mathbb{R}_+^{|\theta^*|} : \prod_{i \in \theta^*} \left[v_i^{\theta^*} - \eta, v_i^{\theta^*} + \eta \right] \subseteq F^*(\theta^*) \text{ for all } i \in \theta^* \right\}.$$

2. If $|\theta^*| < \alpha N$, then $F^{\alpha, \eta}(\theta^*)$ is the one-element set consisting of the zero vector in $\mathbb{R}^{|\theta^*|}$.

Intuitively, if $|\theta^*| \geq \alpha N$, then $F^{\alpha, \eta}(\theta^*)$ is the set of payoff profiles for the rational players that are feasible and individually rational with η slack, while if $|\theta^*| < \alpha N$, then $F^{\alpha, \eta}(\theta^*)$ consists of the payoff vector that results from mutual defection. It is not hard to show that $F^{\alpha, \eta}(\theta^*)$ is nonempty whenever $\eta < \alpha/2$.[22]

THEOREM 3. Fix a sequence $(N, \delta)_l$ satisfying $(1 - \delta) \log N \to 0$, and fix any $\alpha \in (0, 1 - \varepsilon)$, $\eta \in (0, \alpha/2)$, and $\gamma > 0$. With cheap talk, there exists $\bar{l} > 0$ such that, for any $l > \bar{l}$ and any $(v^{\theta^*})_{\theta^*}$ satisfying $v^{\theta^*} \in F^{\alpha, \eta}(\theta^*)$ for all $\theta^* \subseteq I_l$, in the game with parameters (N_l, δ_l) there is a robust sequential equilibrium σ satisfying

[21] This approach to incentivizing communication was introduced by Compte (1998) and Kandori and Matsushima (1998) in the context of general repeated games with private monitoring.

[22] This is true by definition when $|\theta^*| < \alpha N$, so suppose that $|\theta^*| \geq \alpha N$. Since a matched pair of rational players can obtain any payoff pair in the square $\text{co}\{(0, 0), (0, 1), (1, 0), (1, 1)\}$, we have $(w_i)_{i \in \theta^*} \in F^*(\theta^*)$ for any $(w_i)_{i \in \theta^*}$ satisfying $w_i \in [0, p^{\theta^*}]$ for each $i \in \theta^*$, where $p^{\theta^*} := (|\theta^*| - 1)/(N - 1)$ gives the probability that a rational player meets another rational player. Since $p^{\theta^*} > |\theta^*|/N \geq \alpha > 2\eta$, we have $\prod_{i \in \theta^*}[\alpha/2 - \eta, \alpha/2 + \eta] \subset F^*(\theta^*)$, and hence $(\alpha/2, ..., \alpha/2) \in F^{\alpha, \eta}(\theta^*)$.

$$\left| \mathbb{E}^{\sigma} \left[(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u(a_{i,t}, a_{\mu_{i,t},t}) | \theta^* \right] - v_i^{\theta^*} \right| < \gamma \text{ for all } \theta^* \subseteq I_l \text{ and } i \in \theta^*.$$

Proof.   See appendix B, section B.2. QED

The following proposition (which is similar to proposition 3 of Kandori 1992) shows that the $(1 - \delta) \log N \to 0$ sufficient condition in theorem 3 is nearly the best possible.

PROPOSITION 2.   Fix a sequence $(N, \delta)_l$ satisfying $(1 - \delta)^{1+\rho} \log N \to \infty$ for some $\rho > 0$. There exists $\bar{l} > 0$ such that, for any $l > \bar{l}$, in the game with parameters (N_l, δ_l) the unique Nash equilibrium is *Always Defect*.

Proof.   The maximum number of players who can possibly learn about a single player's deviation within t periods is $\min\{2^t, N\}$. Hence, the difference in a player's total continuation payoff when she conforms rather than deviates is at most

$$\sum_{t=1}^{\infty} \delta^t \min\left\{ \frac{2^t}{N}, 1 \right\} (1 + G + L),$$

where here $1 + G + L$ represents the maximum difference between any two stage-game payoffs. Note that for any $\eta \in (0, 1)$,

$$
\begin{aligned}
\sum_{t=1}^{\infty} \delta^t \min\left\{ \frac{2^t}{N}, 1 \right\} &= \sum_{t=1}^{\lfloor \eta \log_2 N \rfloor} \delta^t \frac{2^t}{N} + \sum_{t=\lfloor \eta \log_2 N \rfloor + 1}^{\infty} \delta^t \min\left\{ \frac{2^t}{N}, 1 \right\} \\
&\leq \sum_{t=1}^{\lfloor \eta \log_2 N \rfloor} \frac{2^t}{N} + \frac{\delta^{\lfloor \eta \log_2 N \rfloor}}{1 - \delta} \qquad\qquad (4) \\
&\leq \frac{\lfloor \eta \log_2 N \rfloor \times N^{\eta}}{N} + \frac{\exp(-(1 - \delta) \lfloor \eta \log_2 N \rfloor)}{1 - \delta},
\end{aligned}
$$

where the last inequality follows because the first term in the second line is the sum of $\lfloor \eta \log_2 N \rfloor$ terms that are each less than $2^{\lfloor \eta \log_2 N \rfloor} \leq N^{\eta}$, and $\delta^{\lfloor \eta \log_2 N \rfloor} = \exp(\lfloor \eta \log_2 N \rfloor \log \delta) \leq \exp(-(1 - \delta) \lfloor \eta \log_2 N \rfloor)$. Note that (4) goes to zero whenever there exists $\rho > 0$ such that $(1 - \delta)^{1+\rho} \log N \to \infty$. Thus, if $(1 - \delta)^{1+\rho} \log N \to \infty$ for some $\rho > 0$, the unique Nash equilibrium is *Always Defect*. QED

The remainder of this section previews the equilibrium construction for theorem 3 (sec. IV.A) and extends the model to decouple the frequencies at which players interact and communicate (sec. IV.B).

A.   *Sketch of the Equilibrium Construction for Theorem 3*

The proof of theorem 3 proceeds by constructing a block belief–free equilibrium. Block belief–free equilibria were introduced by Hörner

and Olszewski (2006) in the context of repeated games with almost-perfect monitoring, and they were extended to anonymous random matching games by Deb, Sugaya, and Wolitzky (2020) and to ex post equilibria in games with incomplete information by Sugaya and Yamamoto (2020). The current proof combines elements from these three papers. The main novelty is that since cooperation is impossible in the rare event that there are few rational types, we must keep track of players' beliefs about the number of rational types. In particular, the equilibrium cannot be ex post with respect to the set of rational types. On the other hand, the availability of cheap talk makes providing incentives for truthful communication easier, as compared with the case where communication can be executed only through payoff-relevant actions.

The proof shows that strategies of the following form give a sequential equilibrium. In the very first period of the repeated game, all rational players are supposed to cooperate. Given the realized period 1 action profile, we let $\theta \subseteq I$ denote the set of players who cooperated in period 1. Thus, θ is always a subset of θ^*, and in equilibrium θ equals θ^*. As we will see, all players will eventually abandon cooperation in the event that $|\theta| < \alpha N$. Period 1 thus plays a distinguished role in the equilibrium construction.

Following period 1, the repeated game is viewed as an infinite sequence of finite blocks of T^{**} consecutive periods, where T^{**} is a large number specified in the proof. At the beginning of each block, each player i selects her state profile $(x_i^\theta)_{\theta \subseteq I} \in (\{G, B\})_{\theta \subseteq I}$ for the block, which specifies a state $x_i^\theta \in \{G, B\}$ for each possible realization of θ, the set of players who cooperated in period 1.[23] Intuitively, even if at some point in the game player i comes to believe with probability one that the set of players who cooperated in period 1 was θ, she continues to entertain the possibility that the set of period 1 cooperators was actually some $\theta' \neq \theta$, and she keeps track of a state $x_i^{\theta'} \in \{G, B\}$ for each possible set θ'.

The interpretation of player i's state is as follows: as in Hörner and Olszewski (2006), Deb, Sugaya, and Wolitzky (2020), and Sugaya and Yamamoto (2020), player i can be viewed as the arbiter of player $i + 1$'s payoff, meaning that player $i + 1$'s equilibrium continuation payoff is high when player i is in the good state G, and player $i + 1$'s equilibrium continuation payoff is low when player i is in the bad state B. Specifically, $x_i^\theta = G$ means that, if in the coming block the players reach agreement that the set of period 1 cooperators was θ, then player i prescribes a high continuation payoff for player $i + 1$ (which is delivered both by player i cooperating with player $i + 1$ herself and by player i instructing other players to cooperate with player $i + 1$); similarly, $x_i^\theta = B$ means that, if agreement is reached

---

[23] These states have nothing to do with the states of the diffusion process analyzed in lemma 4.

that the set of period 1 cooperators was θ, then player i prescribes a low continuation payoff for player $i + 1$ (and thus defects against player $i + 1$ herself while also instructing others to defect against player $i + 1$).[24] While it might seem unnecessary for the players to form a belief about θ anew in every block (since, of course, the true value of θ is determined once and for all by the period 1 action profile), this approach conveniently preserves the equilibrium's recursive structure.[25]

The defining feature of a block belief–free equilibrium is that, for each i and θ, all players other than player $i + 1$ (including player i herself) are indifferent as to whether player i selects state $x_i^\theta = G$ and $x_i^\theta = B$, while player $i + 1$ is better off when player i selects $x_i^\theta = G$. Player i can thus be prescribed to randomize between $x_i^\theta = G$ and $x_i^\theta = B$ with a probability depending on her history in the previous block, so as to provide incentives for player $i + 1$ in the previous block. We now describe how a block is constructed so as to ensure that players are incentivized to both take the prescribed actions and communicate honestly.

Each block is divided into several subblocks. In the first portion of the block, players defect while communicating about who cooperated in period 1 so as to reach agreement about θ, as well as communicating their state profiles $(x_i^\theta)_{\theta \subseteq I}$.[26] Then, in each of K main subblocks (which together comprise the vast majority of the block and hence determine the equilibrium payoffs), for multiple periods players take their prescribed actions (which depend on their states, the period 1 history, and the history within the block so far), and players then communicate about their observations within the block so far. Importantly, if the agreed-on state θ satisfies $|\theta| < \alpha N$, then all players are prescribed defection in the main subblocks.

Communication is always executed through a protocol that facilitates truth telling. In essence, when player i meets player j, she reports her past direct observations to him (i.e., her past actions and her past opponents' identities and actions), and also, for each third party $k \notin \{i, j\}$, she tells him all information that she has previously learned via a chain of players that excludes player k. Players thus "tag" each piece of information with

---

[24] "Agreement" here means that all players learn the same set of period 1 cooperations, according to a communication protocol described below. The players can reach agreement without it being common knowledge that they have done so.

[25] This aspect of the construction is facilitated by specifying that "trembles" are much more likely in earlier blocks. Thus, if the players' communications in a prior block indicated that the set of period 1 cooperators is θ, while their communications in the current block indicate that it is $\theta' \neq \theta$, all players believe that the communications in the earlier block were erroneous and proceed in the current block as if the true set were θ'.

[26] Technically, the first portion of the block is divided into three subblocks: one to reach agreement about θ, one to reach agreement about the state profiles, and one to ensure that, if communication in the first two subblocks was unsuccessful (e.g., if a player deviated or if the matching process took on an unlikely realization that prevented all players from meeting and thus reaching agreement), this fact becomes known to all players with high probability.

the identities of the players who have previously conveyed it; for example, one piece of information might be, "I heard from Alice that she heard from Bob that Carol defected against David in period 5." Since tagging occurs at each step, so long as players other than player k have not deviated, player j can trust any information he receives that is tagged as coming from a chain that excludes k. As a consequence, a player cannot unilaterally affect others' inferences about any variable other than her direct observations. Moreover, since a player's direct observations in any period t are also observed by her period t partner, a player cannot unilaterally affect others' inferences about these variables either.

The communication protocol thus ensures that each player cannot prevent her opponents from aggregating information about her own behavior. Together with the threat of punishment (which takes the form of both blacklisting within the block and reduced continuation payoffs at the beginning of the next block), this strategy ensures incentive compatibility at on-path histories. To ensure incentive compatibility at off-path histories, other players reward player j for punishing player i in period t if and only if they confirm through a chain that excludes player j that player j's period t history was one where he was prescribed to punish player i. Since player j's continuation payoff is determined solely by the state of player $j - 1$, which does not affect the payoffs of players other than j, it is optimal for player j's opponents to communicate her history honestly.

Finally, at the end of the block, each player i learns the player $i + 1$'s history in the block through a chain that excludes $i + 1$. She then uses this information to adjust her state mixing probability at the beginning of the next block, so as to deliver the promised continuation payoff to player $i + 1$.

B. Decoupling Interaction Frequency and Communication Frequency

Suppose that rather than matching every period to communicate and interact (i.e., play the prisoner's dilemma) with the same partner, players instead form "communication matches" (where they exchange cheap talk messages but do not play the prisoner's dilemma) every Δ_M units of time and "interaction matches" (where they play the prisoner's dilemma but do not exchange messages) every Δ_A units of time, where players rematch uniformly at random in each meeting. The proofs of theorems 2 and 3 do not rely on the assumption that players communicate and interact with the same partner each period, so if $\Delta_M = \Delta_A = \Delta$ and we fix a real-time discount rate r, we recover these results for any sequence (N, Δ) satisfying $\Delta \log N \to 0$ (because $1 - \delta = 1 - e^{-r\Delta} \approx r\Delta$). More interestingly, if $\Delta_M \neq \Delta_A$, it is straightforward to show that the conclusions of theorems 2 and 3 hold for any

sequence (N, Δ_M, Δ_A) satisfying $\Delta_A \to 0$ and $\Delta_M \log N \to 0$. That is, what these results really require is that players interact frequently (so that the discounting between interaction matches, $1 - e^{-r\Delta_A}$, is small) and that information spreads through the population quickly when players share all their information in every communication match. Since information spreads exponentially in the number of communication matches (by lemma 4), the latter condition requires only that $\Delta_M \log N \to 0$. We explain this point further in appendix B, section B.3.

Putting our results together, we can characterize the prospects for cooperation for a wide range of combinations of the parameters N, Δ_M, and Δ_A. In particular, a folk theorem holds if either

1. $\Delta_A N \to 0$ (regardless of Δ_M; this is a version of proposition 1) or
2. $\Delta_A \to 0$ and $\Delta_M \log N \to 0$ (by modifying theorem 2 or 3 as just discussed).

Conversely, population average payoffs converge to zero in any Nash equilibrium if either

1. $\Delta_A \to \infty$ (regardless of Δ_M; this is obvious),
2. $(\min\{\Delta_A, \Delta_M\})^{1+\rho} \log N \to \infty$ for some $\rho > 0$ (this is a version of proposition 2), or
3. $\Delta_A N \to \infty$ and $\Delta_A N / \Delta_M \to 0$ (this is an extension of theorem 1—which directly applies when $\Delta_A N \to \infty$ and $\Delta_M = \infty$—to the case where Δ_M is large but finite; we omit the proof).

Thus, if we focus on the case where $\Delta_A / \Delta_M \nrightarrow 0$, so the communication frequency is not much less than the interaction frequency (as seems realistic), we obtain either a folk theorem or an impossibility result whenever Δ_A converges to either zero or ∞, $\Delta_A N$ converges to either zero or ∞, and either $\Delta_M \log N \to 0$ or $\Delta_M^{1+\rho} \log N \to \infty$ for some $\rho > 0$.

V. Discussion

A. *Possible Extensions*

We discuss the prospects for extending our model in some technical directions, deferring a broader discussion of future research to the next subsection.

1. Multiple Commitment Types

While the simple "bad types" we consider seem natural and realistic, there is little reason to rule out additional behavioral types that are

committed to strategies other than *Always Defect*. Let us continue to assume that each player is bad with probability ε, while introducing a probability ε' that each player may be committed to some other (arbitrary) repeated game strategy. Theorem 1 extends immediately to this more general setting, with the modified conclusion that $\lim \inf U \le \varepsilon'(1 + G)$ (to account for the fact that the ε' commitment types may cooperate). Theorem 2 also extends, provided that ε' is sufficiently small: if instead ε' is large, then the other commitment type strategies could provide incentives that overturn those in our construction—for example, by rewarding players for defecting against rational opponents.[27] Finally, we conjecture that theorem 3 extends if ε' is sufficiently small and in addition the set of commitment types has the property that there exists a pure strategy for the rational types and a finite time T such that, when the rational types follow this strategy, each commitment type takes a different action than the rational types at some time $t < T$. (Under this property, the first T periods may collectively replace period 1 in the proof of theorem 3.) But we have not verified this conjecture.

2.   Correlated Types

Our analysis of anonymous games in Sugaya and Wolitzky (2020) allows players' types to be correlated, so long as the distribution of the number of bad types satisfies a smoothness condition. In the present paper, independence is used critically in lemma 3. As correlation is introduced, the exponent on N in equation (2) increases and the required condition on δ and N in theorem 1 becomes more stringent, eventually becoming impossible to satisfy when types are perfectly correlated.[28] In contrast, the proofs of theorems 2 and 3 can easily accommodate correlated types.

3.   Independent Noise

An interesting open question is how theorem 1 might extend with i.i.d. noise, where each player is forced to play D with independent probability ε in every period, rather than with probability ε being forced to play D in all periods. Ellison (1994, proposition 2) shows that contagion strategies (which require only $(1 - \delta) \log N \to 0$) are robust to i.i.d. noise if $\varepsilon \to 0$ for a fixed N. If instead $N \to \infty$ for a fixed $\varepsilon > 0$, contagion strategies

---

[27] However, if each commitment type takes a deterministic sequence of actions and messages (rather than responding to its opponents' behavior), then theorem 2 holds for any ε'.

[28] Specifically, independence implies that $\Sigma_i I(S; X_i) \le k$, where $I(S; X_i)$ represents the mutual information between S and X_i (see eq. [A3]). As the X_i's become correlated, the upper bound of k increases toward N, which increases the upper bound in (2).

break down, but also our proof of theorem 1 does not apply.[29] In this case, determining the critical discount factor for supporting cooperation as a function of N seems to require a more intricate analysis of sequential rationality constraints.

B.   Conclusion

This paper has analyzed community enforcement in the presence of "bad types" who never cooperate. We established two main results. First, without explicit communication, community enforcement is ineffective, in that cooperation is sustainable only if bilateral interactions are frequent. Second, introducing ordinary conversation (cheap talk) between matched partners enables cooperation with infrequent bilateral interaction, so long as the population size is not exponentially greater than $1/(1 - \delta)$. Together, these results show that communication is essential for supporting cooperation in large populations. We believe that our model and results provide a more realistic perspective on large-group cooperation than earlier analyses that focused on anonymous agents and collective punishment.

There are a few promising ways in which the theory could be brought even closer to reality. First, while we have shown that communication enables cooperation under realistic-seeming strategies that are approximately optimal (theorem 2), our proof for exact sequential equilibrium (theorem 3) relies on much more complicated strategies that should not be taken literally as a description of real-world behavior. A natural next question is whether and how cooperation can be supported in sequential equilibrium using simpler strategies, perhaps while allowing communication devices that are more powerful than plain cheap talk.

Richer communication devices must also be introduced to support cooperation in populations where N is exponentially greater than $1/(1 - \delta)$ (or where $N = \infty$), as well as to model real-world informational institutions such as fiat money, credit bureaus, and online ratings systems. Recent papers on this topic include Heller and Mohlin (2018), Bhaskar and Thomas (2019), and Clark, Fudenberg, and Wolitzky (2021). Individuals' incentives to provide information to such institutions remain relatively poorly understood, as does these institutions' robustness to dishonest or malicious reporting (e.g., what happens if some agents are "bad communication types," in addition to the "bad action types" we considered?).[30]

---

[29] With i.i.d. noise, the probability that a player is forced to play D for k consecutive periods is ε^k. We could thus apply lemma 3 to this model with $\underline{\varepsilon}^k$ in place of $\underline{\varepsilon}$. But this bound is too loose to yield the conclusion of theorem 1.

[30] There is, however, an interesting empirical literature on these issues in the context of online ratings systems, which is surveyed in sec. 5 of Tadelis (2016).

Finally, in reality large cooperative groups may not be well approximated by the canonical uniform random matching model studied here. Introducing incomplete information (e.g., bad types) into more structured population models—such as models with voluntary separation, assortative matching, or network structure—is another interesting direction for future research.

Appendix A

Omitted Proofs

The condition in the proof of theorem 1 that player i (when rational) prefers strategy σ_i to *Always Defect* is given by

$$
\underbrace{(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\begin{array}{c}(1-\varepsilon)\Pr\!\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)u\!\left(\sigma_i(h_i^t,\mu^t),\sigma_j(h_j^t,\mu^t)\right)\\[4pt]+\varepsilon\Pr\!\left(h_i^t,h_j^t|0_i,1_j,\mu^t\right)u\!\left(\sigma_i(h_i^t,\mu^t),D\right)\end{array}\right)}_{i\text{'s expected payoff when rational and playing }\sigma_i}
$$

$$
\geq \underbrace{(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\begin{array}{c}(1-\varepsilon)\Pr\!\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)u\!\left(D,\sigma_j(h_j^t,\mu^t)\right)\\[4pt]+\varepsilon\Pr\!\left(h_i^t,h_j^t|1_i,1_j,\mu^t\right)u\!\left(D,D\right)\end{array}\right)}_{i\text{'s expected payoff when rational and playing }Always\ Defect},
$$

(A1)

where $u(\cdot,\cdot)$ represents the stage-game payoff function, extended to mixed actions in the usual manner.

A. Proof of Lemma 1

Since $u(D,D)=0$ and $u(C,D)=-L$, (A1) is equivalent to

$$
(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\begin{array}{c}\Pr\!\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)u\!\left(\sigma_i(h_i^t,\mu^t),\sigma_j(h_j^t,\mu^t)\right)\\[4pt]-\dfrac{\varepsilon}{1-\varepsilon}\Pr\!\left(h_i^t|0_i,1_j,\mu^t\right)\Pr(\sigma_i(h_i^t,\mu^t)=C)L\end{array}\right)
$$

$$
\geq (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\Pr\!\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)u\!\left(D,\sigma_j(h_j^t,\mu^t)\right).
$$

Subtracting a like term from both sides, this inequality may be rewritten as

$$
(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\begin{array}{c}\Pr\!\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)\left(\begin{array}{c}u\!\left(\sigma_i(h_i^t,\mu^t),\sigma_j(h_j^t,\mu^t)\right)\\-u\!\left(D,\sigma_j(h_j^t,\mu^t)\right)\end{array}\right)\\[8pt]-\dfrac{\varepsilon}{1-\varepsilon}\Pr\!\left(h_i^t|0_i,1_j,\mu^t\right)\Pr(\sigma_i(h_i^t,\mu^t)=C)L\end{array}\right)
$$

$$
\geq (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\Pr\!\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)-\Pr\!\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)\right)u\!\left(D,\sigma_j(h_j^t,\mu^t)\right)
$$

$$
= (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_j^t}\left(\Pr\!\left(h_j^t|1_i,0_j,\mu^t\right)-\Pr\!\left(h_j^t|0_i,0_j,\mu^t\right)\right)u\!\left(D,\sigma_j(h_j^t,\mu^t)\right).
$$

Since $u(C, a) - u(D, a) \leq -\min\{G, L\}$ and $u(D, a) \in \{0, 1 + G\}$ for each $a \in \{C, D\}$, this inequality implies that

$$(1 - \delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j \neq i}\frac{1}{N-1}\sum_{h_i^t, h_j^t}\Pr\left(h_i^t, h_j^t | 0_i, 0_j, \mu^t\right)\Pr(\sigma_i(h_i^t, \mu^t) = C)\min\{G, L\}$$

$$\geq (1 - \delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j \neq i}\frac{1}{N-1}\sum_{h_j^t}\left(\Pr\left(h_j^t | 0_i, 0_j, \mu^t\right) - \Pr\left(h_j^t | 1_i, 0_j, \mu^t\right)\right)_-(1 + G),$$

or equivalently,

$$(1 - \delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{h_i^t}\Pr(h_i^t | 0_i, \mu^t)\Pr(\sigma_i(h_i^t, \mu^t) = C)\min\{G, L\}$$

$$\leq (1 - \varepsilon)(1 - \delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr(\mu^t)\sum_{j \neq i}\frac{1}{N-1}\sum_{h_j^t}\left(\Pr\left(h_j^t | 0_i, 0_j, \mu^t\right) - \Pr\left(h_j^t | 1_i, 0_j, \mu^t\right)\right)_+(1 + G).$$

Summing this inequality over i and dividing by N yields (1).

B.   Proof of Lemma 3

For random variables A and B taking values in sets $\mathcal{A}$ and $\mathcal{B}$, we denote entropy by $H(A)$, conditional entropy by $H(A|B)$, and mutual information by $I(A; B)$. We have

$$H(A) = -\sum_{a \in \mathcal{A}}\Pr(A = a)\log_2(\Pr(A = a)),$$

$$H(A|B) = -\sum_{a \in \mathcal{A}}\sum_{b \in \mathcal{B}}\Pr(A = a, B = b)\log_2(\Pr(A = a | B = b)),$$

$$I(A; B) = \sum_{a \in \mathcal{A}}\sum_{b \in \mathcal{B}}\Pr(A = a, B = b)\log_2\left(\frac{\Pr(A = a, B = b)}{\Pr(A = a)\Pr(B = b)}\right)$$

$$= H(A) - H(A|B) = H(B) - H(B|A) = I(B; A).$$

Recall that, for any random variables $A_1, \ldots, A_n, B$, we have $H(A_1, \ldots, A_n|B) \leq \sum_{i=1}^n H(A_i|B)$.[31]

Let $X = (X_i)_{i=1}^N$ be a collection of i.i.d. binary random variables with $\Pr(X_i = 1) = \varepsilon$, and let S be a k-dimensional binary random variable defined on the same probability space. Recall that $H(X) = \sum_i H(X_i)$ (by independence) and $H(S) \leq k$. Denote the "influence" of X_i on S by

$$M_i(S) = \sum_{s \in \{0,1\}^k}(\Pr(S = s | X_i = 0) - \Pr(S = s | X_i = 1))_+.$$

Letting $\underline{\varepsilon} = \min\{\varepsilon, 1 - \varepsilon\}$, we wish to show that

$$\sum_{i=1}^N M_i(S) \leq \sqrt{\frac{\log(2)kN}{\underline{\varepsilon}}}.$$

---

[31] This and other basic facts about entropy and mutual information used in the proof can be found in, e.g., Cover and Thomas (2006).

Note that

$$M_i(S) = \sum_{s \in \{0,1\}^k} (\Pr(S = s|X_i = 0) - \Pr(S = s) - (\Pr(S = s|X_i = 1) - \Pr(S = s)))_+$$

$$= \sum_{x \in \{0,1\}} \sum_{s \in \{0,1\}^k} (\Pr(S = s|X_i = x) - \Pr(S = s))_+.$$

Let $P(s) = \Pr(S = s)$, $P^0(s) = \Pr(S = s|X_i = 0)$, and $P^1(s) = \Pr(S = s|X_i = 1)$. By Pinsker's inequality,

$$\sum_{s \in \{0,1\}^k} (\Pr(S = s|X_i = 0) - \Pr(S = s))_+ \leq \sqrt{\frac{\log(2)}{2} D_{KL}(P^0||P)} \quad \text{and}$$

$$\sum_{s \in \{0,1\}^k} (\Pr(S = s|X_i = 1) - \Pr(S = s))_+ \leq \sqrt{\frac{\log(2)}{2} D_{KL}(P^1||P)},$$

where $D_{KL}(\cdot||\cdot)$ denotes Kullback-Leibler divergence (measured in bits). Note that

$$D_{KL}(P^0||P) = \sum_{s \in \{0,1\}^k} \Pr(S = s|X_i = 0)\log_2 \frac{\Pr(S = s|X_i = 0)}{\Pr(S = s)}$$

$$= \frac{1}{1 - \varepsilon} \sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = 0)\log_2\left(\frac{\Pr(S = s, X_i = 0)}{\Pr(S = s)\Pr(X_i = 0)}\right),$$

and

$$D_{KL}(P^1||P) = \frac{1}{\varepsilon} \sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = 1)\log_2\left(\frac{\Pr(S = s, X_i = 0)}{\Pr(S = s)\Pr(X_i = 1)}\right).$$

Hence, we have

$$M_i(S) \leq \sqrt{\frac{\log(2)}{2} D_{KL}(P^0||P)} + \sqrt{\frac{\log(2)}{2} D_{KL}(P^1||P)}$$

$$= \sqrt{\frac{\log(2)}{2(1 - \varepsilon)} \sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = 0)\log_2\left(\frac{\Pr(S = s, X_i = 0)}{\Pr(S = s)\Pr(X_i = 0)}\right)}$$

$$+ \sqrt{\frac{\log(2)}{2\varepsilon} \sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = 1)\log_2\left(\frac{\Pr(S = s, X_i = 0)}{\Pr(S = s)\Pr(X_i = 1)}\right)}$$

$$\leq \sqrt{\frac{\log(2)}{2\underline{\varepsilon}}} \left(\sqrt{\sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = 0)\log_2\left(\frac{\Pr(S = s, X_i = 0)}{\Pr(S = s)\Pr(X_i = 0)}\right)} \right.$$
$$\left. + \sqrt{\sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = 1)\log_2\left(\frac{\Pr(S = s, X_i = 0)}{\Pr(S = s)\Pr(X_i = 1)}\right)} \right) \quad \text{(A2)}$$

$$\leq \sqrt{\frac{\log(2)}{\underline{\varepsilon}}} \sqrt{\sum_{x \in \{0,1\}} \sum_{s \in \{0,1\}^k} \Pr(S = s, X_i = x)\log_2\left(\frac{\Pr(S = s, X_i = x)}{\Pr(S = s)\Pr(X_i = x)}\right)}$$

$$= \sqrt{\frac{\log(2)I(S; X_i)}{\underline{\varepsilon}}},$$

where the last inequality follows because $D_{KL}(\cdot||\cdot)$ is nonnegative and $\sqrt{a} + \sqrt{b} \le \sqrt{2}\sqrt{a+b}$ for nonnegative a, b by the $\ell_1 - \ell_2$ norm inequality, and the last equality is the definition of mutual information.

We next show that

$$\sum_i I(S; X_i) \le k. \tag{A3}$$

To see this, first note that

$$k \ge H(S) \ge I(S; X) = H(X) - H(X|S) = \sum_i H(X_i) - H(X|S),$$

where the last equality follows from independence of $(X_i)_i$. Hence, $H(X|S) \ge \sum_i H(X_i) - k$, and therefore

$$\sum_i H(X_i|S) \ge H(X|S) \ge \sum_i H(X_i) - k.$$

Since $H(X_i|S) = H(X_i) - I(S; X_i)$, we have

$$\sum_i (H(X_i) - I(S; X_i)) \ge \sum_i H(X_i) - k \Leftrightarrow \sum_i I(S; X_i) \le k.$$

Combining (A2) and (A3), we have

$$\sum_i M_i(S)^2 \le \log(2)\sum_i \frac{I(S; X_i)}{\varepsilon} \le \frac{\log(2)k}{\varepsilon},$$

or

$$\sqrt{\sum_i M_i(S)^2} \le \sqrt{\frac{\log(2)k}{\varepsilon}}.$$

Finally, by the $\ell_1 - \ell_2$ norm inequality, in an N-dimensional space $|x|_1 \le \sqrt{N}|x|_2$. Hence,

$$\sum_i M_i(S) \le \sqrt{N}\sqrt{\sum_i M_i(S)^2} \le \sqrt{\frac{\log(2)kN}{\varepsilon}}.$$

C. Proof of Proposition 1

By lemma 2 of Fudenberg and Maskin (1991), there exists $\bar{\delta} < 1$ such that, for all $(v_{i,j}, v_{j,i}) \in F^\eta$, there exists a sequence of pure action profiles whose discounted average payoffs equal $(v_{i,j}, v_{j,i})$ and whose continuation payoffs starting from any time t are within $\eta/2$ of $(v_{i,j}, v_{j,i})$. Call this action path $(a_t^{i,j})_{t \in N}$.

Suppose that each player i conditions her behavior against each player $j \ne i$ only on the history of outcomes in past (i, j) matches, and in particular follows $(a_t^{i,j})_{t \in N}$ if this path has been followed so far in the (i, j) matches, and otherwise reverts to D in these matches forever. By construction, this strategy profile is a sequential equilibrium if, for all $i \ne j$, we have

$$(1 - \delta) \max\{G, L\} \ \leq \ \frac{\delta}{N - 1}(1 - \varepsilon)\left(v_{i,j} - \frac{\eta}{2}\right).$$

Since $v_{i,j} - (\eta/2) \geq \eta/2$ for all $i \neq j$ by hypothesis, a sufficient condition for this profile to be a sequential equilibrium is $\delta \geq 1/2$ and

$$(1 - \delta)N \ \leq \ \frac{\eta(1 - \varepsilon)}{4 \max\{G, L\}}.$$

If $\lim_l (1 - \delta)N = 0$, there exists $\bar{l} > 0$ such that this inequality is satisfied for all $l > \bar{l}$.

It remains to show that, for l sufficiently high, each player i's expected payoff (when rational) in the resulting sequential equilibrium satisfies

$$v_i \in \left[\frac{1}{N - 1}\sum_{j \neq i}\big((1 - \varepsilon)v_{i,j} - \varepsilon\eta\big), \frac{1}{N - 1}\sum_{j \neq i}(1 - \varepsilon)v_{i,j}\right].$$

When player j is rational, player i obtains payoff $v_{i,j}$ against player j. When player j is bad, i obtains the payoff from action path $(a_t^{i,j})_{t \in N}$ until j deviates from this path and then obtains payoff zero forever. Suppose that the first deviation by j from action path $(a_t^{i,j})_{t \in N}$ occurs in period t. Then i's payoff against j is at least $(1 - \delta^t)u_{i,j}^{<t} + \delta^t(1 - \delta)(-L)$, where $u_{i,j}^{<t}$ is i's average payoff from the first $t - 1$ periods of action path $(a_t^{i,j})_{t \in N}$. Note that $u_{i,j}^{<t}$ satisfies

$$(1 - \delta^t)u_{i,j}^{<t} + \delta^t u_{i,j}^{\geq t} \ = \ v_{i,j},$$

where $u_{i,j}^{\geq t}$ represents i's average payoff starting from period t under action path $(a_t^{i,j})_{t \in N}$, and $u_{i,j}^{\geq t} \leq v_{i,j} + \eta/2$. Hence,

$$(1 - \delta^t)u_{i,j}^{<t} \geq (1 - \delta^t)v_{i,j} - \delta^t\frac{\eta}{2} \geq -\frac{\eta}{2}.$$

Therefore, for δ sufficiently high that $(1 - \delta)L \leq \eta/2$, i's payoff against j is at least

$$(1 - \delta^t)u_{i,j}^{<t} + \delta^t(1 - \delta)(-L) \geq -\frac{\eta}{2} - \frac{\eta}{2} = -\eta.$$

Moreover, i's payoff against j is nonpositive, since j always defects. Hence, i's expected payoff against j is at least $(1 - \varepsilon)v_{i,j} - \varepsilon\eta$ and at most $(1 - \varepsilon)v_{i,j}$. Averaging over $j \neq i$ yields the desired bounds for v_i.

D. Proof of Theorem 2

Consider the following strategies, which do not depend on l.

Equilibrium strategies.—Each player i enters each period t with a blacklist $I_{i,t}^D \subseteq I$. Let $I_{i,1}^D = \varnothing$ for each i.

In period t, player i truthfully reports $I_{i,t}$ to her period t opponent $\mu_{i,t}$ (whether or not i is rational). When rational, i then takes action C if $\mu_{i,t} \notin I_{i,t}^D$ and takes D if $\mu_{i,t} \in I_{i,t}^D$. Bad types always take D.

Denote the report of player i's opponent by $\hat{I}_{\mu_{i,t},t}^D$. At the end of period t, i's blacklist updates to

$$I_{i,t+1}^{D} = \begin{cases} I_{i,t}^{D} \cup \hat{I}_{\mu_{i,t},t}^{D} & \text{if } \mu_{i,t} \text{ played } C \text{ or } i \text{ is bad,} \\ I_{i,t}^{D} \cup \hat{I}_{\mu_{i,t},t}^{D} \cup \{\mu_{i,t}\} & \text{if } \mu_{i,t} \text{ played } D \text{ and } i \text{ is rational.} \end{cases}$$

Fix $\eta > 0$. We prove that, for sufficiently large l, these strategies form an η-Nash equilibrium. To do so, we (1) compute lower bounds on the equilibrium payoffs of rational and bad types, (2) compute upper bounds on the payoffs of rational and bad types from any unilateral deviation, and (3) show that the latter cannot exceed the former by more than η.

Rational type equilibrium payoff.—Suppose that i is rational, let S denote the set of bad players, and suppose that $|S| = n$. Fix any $T, Z \in N$ with $Z > \bar{Z}$ (with $\bar{Z}$ defined as in the statement of lemma 4). The probability that every bad player meets a rational player at least once by period T is at least $1 - n((n-1)/(N-1))^{T}$. Conditional on this event, by lemma 4, $I_{i,T+Z\log_2 N}^{D} = S$ with probability at least $1 - \exp(-cZ)$. Hence, with probability at least $1 - n((N-n)/(N-1))^{T} - \exp(-cZ)$, starting from period $T + Z\log_2 N$ player i obtains payoff one when she meets a rational type and obtains payoff zero when she meets a bad type, for an expected payoff of $(N-1-n)/(N-1)$. For the first $T + Z\log_2 N$ periods, and with probability at most $n((N-n)/(N-1))^{T} + \exp(-cZ)$ for the rest of the game, player i's payoff is at least $-L$. In total, rational player i's equilibrium expected payoff, conditional on the event $|S| = n$, is at least

$$\frac{N-1-n}{N-1} - \min_{T \in N, Z > \bar{Z}} \left\{ \left(1 - \delta^{T+Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^{T} + \exp(-cZ) \right\}(1 + L).$$

Taking the expectation with respect to n, rational player i's equilibrium unconditional expected payoff is at least

$$\sum_{n} p_n \frac{N-1-n}{N-1} - \sum_{n} p_n \min_{T \in N, Z > \bar{Z}} \left\{ \left(1 - \delta^{T+Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^{T} \right.$$
$$\left. + \exp(-cZ) \right\}(1 + L),$$

where $p_n = \binom{N}{n} \varepsilon^{n}(1 - \varepsilon)^{N-n}$ denotes the probability that there are n bad types.

We will show that, for sufficiently large l,

$$\sum_{n} p_n \frac{N-1-n}{N-1} - \sum_{n} p_n \min_{T \in N, Z > \bar{Z}} \left\{ \left(1 - \delta^{T+Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^{T} \right.$$
$$\left. + \exp(-cZ) \right\}(1 + L) \tag{A4}$$

$$\geq \sum_{n} p_n \frac{N-1-n}{N-1} - \frac{2}{3}\eta.$$

First, fix some $\hat{\alpha} \in (0, 1 - \varepsilon)$, and fix some $Z > \bar{Z}$ such that $\exp(-cZ) \leq \eta/(4(1 + L))$. Here $\hat{\alpha}$ and Z are independent of l. By the central limit theorem, for sufficiently large N (or l), the probability that there are more than $(1 - \hat{\alpha})N$ bad types, $\sum_{n \geq (1-\hat{\alpha})N} p_n$, is less than $(1/12)(1/(1 + N + \exp(-cZ)))(\eta/(1 + L))$. Since $(1 - \delta^{T+Z\log_2 N}) + n((n-1)/(N-1))^{T} + \exp(-cZ) \leq 1 + N + \exp(-cZ)$

for each N, $n \leq N$, and $T \in N$, to establish (A4) it suffices to show that, for sufficiently large l, there exists T such that, for each $n \leq (1 - \hat{\alpha})N$, we have

$$\left(\left(1 - \delta^{T + Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^T + \exp(-cZ) \right)(1 + L) \leq \frac{7}{12}\eta.$$

By definition of Z, $\exp(-cZ)(1 + L) \leq (1/4)\eta$. Hence, it remains to show that, for sufficiently large l, there exists T such that, for each $n \leq (1 - \hat{\alpha})N$, we have

$$\left(\left(1 - \delta^{T + Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^T \right)(1 + L) \leq \frac{1}{3}\eta. \tag{A5}$$

To establish (A5), for a given value of l, let $T \in N$ be the smallest integer such that $N(1 - \hat{\alpha})^T \leq \eta/(4(1 + L))$. Note that $T \leq \hat{c}\log_2 N$ for some constant $\hat{c}$. Now, for all $n \leq (1 - \hat{\alpha})N$, we have

$$\left(\frac{n-1}{N-1}\right)^T \leq (1 - \hat{\alpha})N\left(\frac{(1 - \hat{\alpha})N - 1}{N-1}\right)^T \leq (1 - \hat{\alpha})N(1 - \hat{\alpha})^T \leq \frac{\eta}{4(1 + L)}.$$

Hence, to establish (A5), it suffices to show that, for sufficiently large l, we have

$$\left(1 - \delta^{T + Z\log_2 N}\right)(1 + L) \leq \frac{1}{12}\eta. \tag{A6}$$

Finally, we have

$$1 - \delta^{T + Z\log_2 N} \leq (1 - \delta)(T + Z\log_2 N) \leq (1 - \delta)(\hat{c} + Z)\log_2 N,$$

and $(1 - \delta)\log_2 N \to 0$ as $l \to \infty$ by hypothesis. This establishes (A6) (and hence [A4]), as desired.

Bad type equilibrium payoff.—Here we take the trivial bound that, when player i is bad, her equilibrium payoff is nonnegative.

Rational type deviation payoff.—We derive an upper bound for player i's payoff under any unilateral deviation. To this end, suppose that player i can observe whether her opponent is rational or bad before acting and always takes D against bad opponents. Moreover, suppose that player i's opponents blacklist her if they learn that she took D against a rational player through a chain of players that excludes player i herself; that is, if player i played D against a rational opponent in period τ, then a rational player j takes D against i in period $t > \tau$ if there exists a sequence of players $(j_\tau, j_{\tau+1}, \dots, j_{t-1})$ such that $j_\tau = \mu_{i,\tau}$, $j \in \{j_\tau, \dots, j_{t-1}\}$, $i \notin \{j_\tau, \dots, j_{t-1}\}$, and $j_{t'+1} = \mu_{j_{t'}, t'+1}$ for each $t' \in \{\tau, \dots, t - 2\}$. By lemma 4, if player i takes D against a rational player in period τ, then, for every $Z > \bar{Z}$, with probability $1 - \exp(-cZ)$ everyone takes D against player i starting from period $Z\log_2 N$. Hence, player i's expected payoff is at most

$$\sum_{n=0}^{N-1} p_n \frac{N - 1 - n}{N - 1} + \min_{Z > \bar{Z}}\left\{\left(1 - \delta^{Z\log_2 N}\right) + \exp(-cZ)\right\}(1 + G).$$

First fixing Z such that $\exp(-cZ) \leq \eta/(3(1 + G))$ and then taking $l \to \infty$, we see that $1 - \delta^{Z\log_2 N} \leq (1 - \delta)Z\log_2 N \to 0$, so for sufficiently large l this is at most

$$\sum_{n=0}^{N-1} p_n \frac{N - 1 - n}{N - 1} + \frac{1}{3}\eta.$$

Comparing this upper bound with the lower bound (A4), we see that the equilibrium strategy is η-optimal.

Bad type deviation payoff.—Since player i always takes D when bad, if she meets a rational player for the first time in period τ, for every $Z > \bar{Z}$, her continuation payoff starting from period $\tau + Z\log_2 N$ is zero with probability at least $1 - \exp(-cZ)$. (As in the case where player i is rational, this holds regardless of player i's own behavior following period τ.) Since player i's payoff against bad opponents is nonpositive, her payoff under any unilateral deviation is at most

$$\delta^\tau \min_{Z > \bar{Z}} \left\{ \left(1 - \delta^{Z\log_2 N}\right) + \exp(-cZ)(1 + G) \right\}.$$

As we have seen, this converges to zero as $l \to \infty$. Hence, the equilibrium strategy is η-optimal.

References

Ahn, Illtae, and Matti Suominen. 2001. "Word-of-Mouth Communication and Community Enforcement." *Internat. Econ. Rev.* 42 (2): 399–415.

Ali, S. Nageeb, and David A. Miller. 2013. "Enforcing Cooperation in Networked Societies." Working paper, Univ. California, San Diego.

———. 2016. "Ostracism and Forgiveness." *A.E.R.* 106 (8): 2329–48.

———. 2020. "Communication and Cooperation in Markets." Working paper.

Al-Najjar, Nabil I., and Rann Smorodinsky. 2000. "Pivotal Players and the Characterization of Influence." *J. Econ. Theory* 92 (2): 318–42.

Awaya, Yu, and Vijay Krishna. 2016. "On Communication and Collusion." *A.E.R.* 106 (2): 285–315.

———. 2019. "Communication and Cooperation in Repeated Games." *Theoretical Econ.* 14 (2): 513–53.

Balmaceda, Felipe, and Juan F. Escobar. 2017. "Trust in Cohesive Communities." *J. Econ. Theory* 170:289–318.

Barron, Daniel, and Yingni Guo. 2021. "The Use and Misuse of Coordinated Punishments." *Q.J.E.* 136 (1): 471–504.

Baumeister, Roy F., Liqing Zhang, and Kathleen D. Vohs. 2004. "Gossip as Cultural Learning." *Rev. General Psychology* 8 (2): 111–21.

Bhaskar, V., and Caroline Thomas. 2019. "Community Enforcement of Trust with Bounded Memory." *Rev. Econ. Studies* 86 (3): 1010–32.

Bowen, T. Renee, David M. Kreps, and Andrzej Skrzypacz. 2013. "Rules with Discretion and Local Information." *Q.J.E.* 128 (3): 1273–320.

Clark, Daniel, Drew Fudenberg, and Alexander Wolitzky. 2021. "Record-Keeping and Cooperation in Large Societies." *Rev. Econ. Studies*, forthcoming.

Compte, Olivier. 1998. "Communication in Repeated Games with Imperfect Private Monitoring." *Econometrica* 66 (3): 597–626.

Cover, Thomas M., and Joy A. Thomas. 2006. *Elements of Information Theory*, 2nd ed. Wiley Series in Telecommunications and Signal Processing. Hoboken, NJ: Wiley.

Cramton, Peter, and Jesse A. Schwartz. 2000. "Collusive Bidding: Lessons from the FCC Spectrum Auctions." *J. Regulatory Econ.* 17 (3): 229–52.

Deb, Joyee. 2020. "Cooperation and Community Responsibility." *J.P.E.* 128 (5): 1976–2009.

Deb, Joyee, Takuo Sugaya, and Alexander Wolitzky. 2020. "The Folk Theorem in Repeated Games with Anonymous Random Matching." *Econometrica* 88 (3): 917–64.

Dilmé, Francesc. 2016. "Helping Behavior in Large Societies." *Internat. Econ. Rev.* 57 (4): 1261–78.

Dixit, Avinash. 2003. "Trade Expansion and Contract Enforcement." *J.P.E.* 111 (6): 1293–317.

Dunbar, Robin. 2004. "Gossip in Evolutionary Perspective." *Rev. General Psychology* 8 (2): 100–110.

Ellickson, Robert C. 1994. *Order without Law: How Neighbors Settle Disputes.* Cambridge, MA: Harvard Univ. Press.

Ellison, Glenn. 1994. "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching." *Rev. Econ. Studies* 61 (3): 567–88.

Feinberg, Matthew, Robb Willer, and Michael Schultz. 2014. "Gossip and Ostracism Promote Cooperation in Groups." *Psychological Sci.* 25 (3): 656–64.

Frieze, Alan M., and Geoffrey R. Grimmett. 1985. "The Shortest-Path Problem for Graphs with Random Arc-Lengths." *Discrete Appl. Math.* 10 (1): 57–77.

Fudenberg, Drew, David Levine, and Wolfgang Pesendorfer. 1998. "When Are Nonanonymous Players Negligible?" *J. Econ. Theory* 79 (1): 46–71.

Fudenberg, Drew, and Eric Maskin. 1991. "On the Dispensability of Public Randomization in Discounted Repeated Games." *J. Econ. Theory* 53 (2): 428–38.

Ghosh, Parikshit, and Debraj Ray. 1996. "Cooperation in Community Interaction without Information Flows." *Rev. Econ. Studies* 63 (3): 491–519.

Gluckman, Max. 1963. "Papers in Honor of Melville J. Herskovits: Gossip and Scandal." *Current Anthropology* 4 (3): 307–16.

Greif, Avner. 1993. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition." *A.E.R.* 83 (3): 525–48.

Harrington, Jr., Joseph E. 1995. "Cooperation in a One-Shot Prisoners' Dilemma." *Games and Econ. Behavior* 8 (2): 364–77.

Hedetniemi, Sandra M., Stephen T. Hedetniemi, and Arthur L. Liestman. 1988. "A Survey of Gossiping and Broadcasting in Communication Networks." *Networks* 18 (4): 319–49.

Heller, Yuval, and Erik Mohlin. 2018. "Observations on Cooperation." *Rev. Econ. Studies* 85 (4): 2253–82.

Hörner, Johannes, and Wojciech Olszewski. 2006. "The Folk Theorem for Games with Private Almost-Perfect Monitoring." *Econometrica* 74 (6): 1499–544.

Kandori, Michihiro. 1992. "Social Norms and Community Enforcement." *Rev. Econ. Studies* 59 (1): 63–80.

Kandori, Michihiro, and Hitoshi Matsushima. 1998. "Private Observation, Communication and Collusion." *Econometrica* 66 (3): 627–52.

Klein, Daniel B. 1992. "Promise Keeping in the Great Society: A Model of Credit Information Sharing." *Econ. and Politics* 4 (2): 117–36.

Lippert, Steffen, and Giancarlo Spagnolo. 2011. "Networks of Relations and Word-of-Mouth Communication." *Games and Econ. Behavior* 72 (1): 202–17.

Noon, Mike, and Rick Delbridge. 1993. "News from behind My Hand: Gossip in Organizations." *Org. Studies* 14 (1): 23–36.

Okuno-Fujiwara, Masahiro, and Andrew Postlewaite. 1995. "Social Norms and Random Matching Games." *Games and Econ. Behavior* 9 (1): 79–109.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge: Cambridge Univ. Press.

Raub, Werner, and Jeroen Weesie. 1990. "Reputation and Efficiency in Social Interactions: An Example of Network Effects." *American J. Soc.* 96 (3): 626–54.

Shah, Devavrat. 2009. *Gossip Algorithms.* Boston: Now.

Sommerfeld, Ralf D., Hans-Jürgen Krambeck, Dirk Semmann, and Manfred Milinski. 2007. "Gossip as an Alternative for Direct Observation in Games of Indirect Reciprocity." *Proc. Nat. Acad. Sci. USA* 104 (44): 17435–40.

Sugaya, Takuo, and Alexander Wolitzky. 2020. "A Few Bad Apples Spoil the Barrel: An Anti-folk Theorem for Anonymous Repeated Games with Incomplete Information." *A.E.R.* 110 (12): 3817–35.

Sugaya, Takuo, and Yuichi Yamamoto. 2020. "Common Learning and Cooperation in Repeated Games." *Theoretical Econ.* 15 (3): 1175–219.

Tadelis, Steven. 2016. "Reputation and Feedback Systems in Online Platform Markets." *Ann. Rev. Econ.* 8:321–40.

Takahashi, Satoru. 2010. "Community Enforcement When Players Observe Partners' Past Play." *J. Econ. Theory* 145 (1): 42–62.

Wolitzky, Alexander. 2012. "Cooperation with Network Monitoring." *Rev. Econ. Studies* 80 (1): 395–427.

———. 2015. "Communication with Tokens in Repeated Games on Networks." *Theoretical Econ.* 10 (1): 67–101.