A Few Bad Apples Spoil the Barrel: An Anti-Folk Theorem for Anonymous Repeated Games with Incomplete Information*

Takuo Sugaya Alexander Wolitzky

August 3, 2020

Abstract

We study anonymous repeated games where players may be "commitment types" who always take the same action. We establish a stark anti-folk theorem: if the distribution of the number of commitment types satisfies a smoothness condition and the game has a "pairwise dominant" action, this action is almost always taken. This implies that cooperation is impossible in the repeated prisoner's dilemma with anonymous random matching. We also bound equilibrium payoffs for general games. Our bound implies that industry profits converge to zero in linear-demand Cournot oligopoly as the number of firms increases.

Keywords: anonymous games, repeated games, incomplete information, community enforcement

JEL codes: C72, C73, D82

*Sugaya: Graduate School of Business, Stanford University, tsugaya@stanford.edu. Wolitzky: Department of Economics, MIT, wolitzky@mit.edu. For helpful comments, we thank Daron Acemoglu, Glenn Ellison, Jeff Ely, Drew Fudenberg, Yuval Heller, Matt Jackson, George Mailath, Stephen Morris, Satoru Takahashi, Omer Tamuz, the anonymous referees, and seminar particiants at Princeton, Stanford, and Wharton. We thank Eitan Sapiro-Gheiler for careful proofreading. Wolitzky acknowledges financial support from the NSF and the Sloan Foundation.

1 Introduction

The folk theorem of repeated games asserts that a group of rational players, however large, can cooperate if they are sufficiently patient and have enough information about each other's past behavior.[1] But it seems more realistic to assume that, if a group is large enough, it probably contains some irrational agents. For example, Kandori (1992) and Ellison (1994) show that rational players can support cooperation in the prisoner's dilemma with anonymous random matching by relying on *contagion strategies*: whenever a player sees anyone defect, she starts defecting against everyone. But Ellison also notes (p. 578)

> 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.

In this paper, we show that the folk theorem fails for large groups when players are anonymous—so a player's payoff depends only on her own action and the number of opponents taking each action—and may be "commitment types" who always take the same action. For example, in the prisoner's dilemma with anonymous random matching, population size N, and discount factor δ, cooperation is impossible when N is large, even if $(1-\delta)N$ is small. Similarly, in linear-demand Cournot oligopoly, industry profits converge to zero as $N \to \infty$, even if $(1-\delta)N \to 0$.

The key assumption behind our results is that the distribution of the number of commitment types is "smooth": roughly speaking, for every number $n < N$, the probability that n players are commitment types is close to the probability that $n+1$ players are commitment types. For instance, this assumption is satisfied if each player is a commitment type with independent probability z, for any fixed $z \in (0,1)$, and $N \to \infty$.

To see why the folk theorem fails with a smooth distribution of commitment types, observe that, if a rational type deviates from her equilibrium strategy by instead following the strategy of a commitment type, and if the number of "true" commitment types is n, then the population distribution of actions is exactly what it would have been if the rational player had not deviated and the number of commitment types had been $n+1$. Smoothness thus implies that a single deviation from the rational-type strategy to the commitment-type strategy has a small impact on the population distribution of actions. Therefore, the commitment type's action cannot perform

[1] The literature on the folk theorem is enormous. For a textbook treatment, see Mailath and Samuelson (2006).

much better than the rational type's equilibrium strategy against the equilibrium action distribution (a fact we formalize in Lemma 1). Finally, in many games this fact implies that the rational type's equilibrium strategy almost always prescribes the commitment type action, which yields an anti-folk theorem. For example, if the game is the prisoner's dilemma and the commitment type's action is defect, the rational type's equilibrium strategy must almost always defect.

More precisely, we consider anonymous repeated games with one rational type and one commitment type.[2] We first establish Lemma 1: the commitment-type strategy cannot yield a much higher payoff than the rational-type equilibrium strategy, where the size of the gap depends on the smoothness of the distribution of the number of commitment types. We then consider games with a "pairwise dominant" action a^*—meaning that, whenever one player takes action a^* and another player takes a different action a, the player taking a^* obtains a strictly higher payoff than the player taking a—and assume that commitment types take this action. For instance, defection is pairwise dominant in the prisoner's dilemma. Our main result (Theorem 1) shows that, as $N \to \infty$, the pairwise dominant action is almost always taken in every Nash equilibrium. We then briefly consider implications of Lemma 1 for games without a pairwise dominant action, showing in particular that industry profits converge to zero in linear-demand Cournot oligopoly as the number of firms increases.

The paper concludes by discussing implications of our approach beyond anonymous repeated games. One such implication is an elementary proof of a version of Mailath and Postlewaite's (1990) impossibility theorem for public good provision in large populations, which unlike existing proofs allows types to be correlated.

1.1 Related Literature

This paper relates to several branches of literature. Most directly, we contribute to the literature on repeated games with anonymous random matching by showing that the existence of cooperative equilibria in such models is not robust to introducing a smooth distribution of commitment types.[3]

There are two related strands of literature on anti-folk theorems. First, several papers following Green (1980) and Sabourian (1990) consider large-population, complete-information games where the impact of each player's action on the aggregate signal distribution is small.[4] These papers give

[2] Our results can be extended to allow multiple commitment types at the cost of additional notation. We discuss this extension in Section 6.

[3] Without allowing commitment types, Kandori (1992) and Ellison (1994) showed that mutual cooperation is supportable in the prisoner's dilemma, and Deb, Sugaya, and Wolitzky (2020) established a general folk theorem.

[4] See also Levine and Pesendorfer (1995), Fudenberg, Levine, and Pesendorfer (1998), Al-Najjar and Smorodinsky

conditions such that, for fixed δ, all Nash equilibria of the repeated game converge to static Nash equilibria as $N \to \infty$; significantly, they do not give anti-folk theorems in the sense of persistent inefficiency as $\delta \to 1$. In contrast, with incomplete information we allow arbitrarily informative signals of actions (e.g., perfect monitoring) and give conditions such that convergence to static Nash equilibrium obtains uniformly in δ.

Second, like our paper, the "reputation" literature shows that introducing a small amount of incomplete information in repeated games can lead to anti-folk theorems (Mailath and Samuelson, 2006). Reputation models typically consider a small number of long-run players (often only one), and are thus far from the large anonymous games we consider. In particular, the key argument that bounds a rational player's payoff in this literature (due to Fudenberg and Levine, 1989) is that, if a rational player follows the strategy of a commitment type, this eventually causes her opponents to start taking favorable actions in response. In contrast, the key argument that bounds a rational player's payoff in our large anonymous games is that, if a rational player follows the strategy of a commitment type, this has only a small effect on the distribution of her opponents' actions.

Finally, a literature closer to mechanism design considers measures of the pivotality or influence of a player's type on an aggregate outcome, and gives conditions under which most players' influence must be small in large populations. For instance, al-Najjar and Smorodinsky (2000) show that, with independent types, players' average influence on a bounded, real-valued aggregate outcome goes to zero as $N \to \infty$. This result is distinct from our condition that the distribution of the number of agents with a specific type does not vary much with a particular player's type. And this distinction makes a difference: al-Najjar and Smorodinsky (2001) apply their notion of influence to continuation payoffs in repeated games to derive a Green-Sabourian-type result that depends on the order of limits between N and δ, while our results are uniform in δ.[5]

2 Model

A symmetric N-player stage game with action set A and payoff function $u : A \to \mathbb{R}$ is *anonymous* if, for any $i \in I = \{1, \ldots, N\}$, any permutation π on $I \backslash \{i\}$, and any action profile $\mathbf{a} = (a_j)_{j \in I} \in A$, we have $u_i(a_1, \ldots, a_{i-1}, a_i, a_{i+1}, \ldots, a_N) = u_i\left(a_{\pi(1)}, \ldots, a_{\pi(i-1)}, a_i, a_{\pi(i+1)}, \ldots, a_{\pi(N)}\right)$: that is, a player's payoff depends only on her own action and the number of opponents taking each action. Fix a finite, anonymous stage game, and normalize the range of u to lie in $[0, 1]$. Throughout

(2001), Pai, Roth, and Ullman (2017), and Awaya and Krishna (2016, 2019).

[5] Another way to appreciate the difference is to note that we allow correlated types.

the paper, whenever we write a player's payoff as $u(\mathbf{a})$ (without an i-subscript on u), the first element of $\mathbf{a}$ refers to the player's own action and the remaining elements refer to the opponents' actions, which by anonymity can be ordered arbitrarily: thus, $u(a_i, a_{-i})$ is player i's payoff when she takes action $a_i \in A$ and her opponents take actions $a_{-i} \in A^{(N-1)}$. In contrast, $u_i(\mathbf{a}) = u_i(a_1, \ldots, a_{i-1}, a_i, a_{i+1}, \ldots, a_N)$.

The stage game is played repeatedly in periods $t = 1, 2, \ldots$. After taking an action in period t, each player i observes a signal $y_{i,t}$ drawn from a probability distribution that depends on the history of past actions and signals $\left((\mathbf{a}_\tau, \mathbf{y}_\tau)_{\tau=1}^{t-1}, \mathbf{a}_t\right)$, where $\mathbf{y}_\tau = (y_{1,\tau}, \ldots, y_{N,\tau})$ and $\mathbf{a}_\tau = (a_{1,\tau}, \ldots, a_{N,\tau})$. A *history* for player i at the beginning of period t is thus $h_i^t = (a_{i,\tau}, y_{i,\tau})_{\tau=1}^{t-1}$, with $h_i^1 = \emptyset$. A *strategy* σ_i for player i maps histories h_i^t to $\Delta(A)$ for each t.

Each player i has a type $\theta_i \in \{R, B\}$, where R is the *rational* type and B is the *bad* (commitment) type.[6] Rational types maximize expected discounted payoffs with discount factor $\delta \in [0, 1)$. Bad types always play a particular action $a^* \in A$; we call this strategy *Always* a^*. A *strategy profile* $\sigma = (\sigma_i)_i$ specifies the strategy σ_i that each player i follows *when she is rational*; of course, when she is bad, she plays *Always* a^*.

There is a common prior p on the set of players' types $\{R, B\}^N$, which we assume is symmetric: for every permutation π on I and every type profile $(\theta_1, \ldots, \theta_N)$, $p(\theta_1, \ldots, \theta_N) = p\left(\theta_{\pi(1)}, \ldots, \theta_{\pi(N)}\right)$. The repeated game is thus parameterized by the tuple $\Gamma = (N, A, u, \delta, a^*, p)$. We denote the probability that a player is bad by $z = \sum_{\theta:\theta_i=B} p(\theta)$.

Given a strategy profile σ, denote player i's expected discounted per-period payoff conditional on type profile θ by $U_i(\theta) \in [0, 1]$, and denote player i's expected payoff by $U_i = \sum_\theta p(\theta) U_i(\theta) \in [0, 1]$. Since this expectation includes the possibility that $\theta_i = B$, we implicitly assume that bad types have the same utility function as rational types.[7]

Our results concern the set of equilibrium values of per-capita utilitarian social welfare, $\sum_i U_i/N$. This set is only expanded by letting the players access a public randomization device. Since the stage game and prior are symmetric, when public randomization is available any social welfare level attainable by an asymmetric strategy profile is also attained by the symmetric profile where public randomization is first used to randomly permute the players' strategies. We therefore allow public randomization and restrict attention to symmetric strategy profiles (and often drop the i subscript from $U_i(\theta)$ and U_i).

[6] We discuss the case with multiple commitment types in Section 6.

[7] One could alternatively consider player i's expected utility conditional on the event $\theta_i = R$. This would involve a little more notation while giving essentially the same results.

Note that we allow $\delta = 0$, so our results apply equally to one-shot games.

3 Preliminaries

3.1 Lower Bound for Rational Players' Payoffs

We first show that the bad type's action cannot perform much better than the rational type's equilibrium strategy against the equilibrium action distribution (Lemma 1). This *lower* bound for rational players' payoffs will later drive our anti-folk theorem, which in particular implies an *upper* bound for payoffs.

For $n \in \{0, 1, ..., N\}$, let $\mathcal{B}_n$ denote the event that the realized number of bad types is n, and let p_n denote the probability of this event. Conditional on the event that a given player i is rational ($\theta_i = R$), let q_n denote the probability that n out of the remaining $N - 1$ players are bad. Since symmetry implies that

$$\Pr\left(\mathcal{B}_n \wedge \theta_i = R\right) = \frac{N-n}{N} p_n,$$

we see that q_n is given by

$$q_n = \Pr\left(\mathcal{B}_n | \theta_i = R\right) = \frac{\Pr\left(\mathcal{B}_n \wedge \theta_i = R\right)}{\Pr\left(\theta_i = R\right)} = \frac{N-n}{N} \frac{p_n}{1-z}.$$

We also let $q_N = 0$ by convention. Next, conditional on the event that a given player is rational, denote the probability that $n - 1$ out of the remaining $N - 1$ players are bad by

$$q_n^- = q_{n-1} \text{ for } n \in \{1, \ldots, N\},$$

with $q_0^- = 0$ by convention. Note that, given the convention that $q_N = q_0^- = 0$, $q = (q_n)_{n=0}^N$ and $q^- = (q_n^-)_{n=0}^N$ are both probability distributions on $\{0, \ldots, N\}$. Denote the total variation distance between these probability distributions by

$$\Delta_{q,q^-} = \max_{\mathcal{N} \subset \{0,\ldots,N\}} \left| \sum_{n \in \mathcal{N}} \left(q_n - q_n^-\right) \right|. \tag{1}$$

Note that if a rational player deviates by playing *Always* a^* instead of her equilibrium strategy and the realized number of bad types is $n - 1$, then the population distribution of actions is the same as it would be if this player had not deviated and the realized number of bad types were

n. Thus, from the perspective of a rational player (and assuming that the equilibrium strategy of rational players is something other than *Always* a^*), q_n is the probability that n players in the population play *Always* a^* when she follows her equilibrium strategy, and q_n^- is the probability that n players in the population play *Always* a^* when she deviates to *Always* a^*. The distance between the distributions q and q^-, Δ_{q,q^-}, is therefore a measure of the detectability of a deviation by a rational player from her equilibrium strategy to *Always* a^*.

Fix a symmetric Nash equilibrium σ. For $n \in \{0, 1, ..., N\}$, let u_n denote the expected payoff of a random player in the population when there are n bad types, given by

$$u_n = \mathbb{E}\left[U_i\left(\theta\right)|\mathcal{B}_n\right].$$

Let u_n^R denote a rational player's expected payoff when there are n bad types, given by

$$u_n^R = \mathbb{E}\left[U_i\left(\theta\right)|\theta_i = R, \mathcal{B}_n\right].$$

Let u_n^B denote a bad player's expected payoff when there are n bad types, given by

$$u_n^B = \mathbb{E}\left[U_i\left(\theta\right)|\theta_i = B, \mathcal{B}_n\right].$$

Note that, for each n, we have

$$u_n = \frac{N-n}{N}u_n^R + \frac{n}{N}u_n^B.$$

Moreover,

$$U = \sum_{n=0}^{N} p_n u_n.$$

We let $u_0^B = 1$ by convention. This convention makes the following lemma, which gives the desired lower bound on rational players' payoffs, as strong as possible.

Lemma 1 *For any anonymous game and any symmetric Nash equilibrium, the following bounds apply:*

1. ***Rational player payoff bound:*** $\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta_{q,q^-}$.

2. ***Social welfare bound:*** $U \geq \sum_{n=0}^{N} p_n u_n^B - (1-z)\,\Delta_{q,q^-}$.

Proof. In any Nash equilibrium, a rational player must prefer her equilibrium strategy to deviating to *Always* a^*. A rational player's equilibrium payoff is $\sum_{n=0}^{N-1} q_n u_n^R$. If a rational player instead plays *Always* a^*, then for each realized number of "true" bad types n, she receives the same payoff as that received in equilibrium by a bad type when the true number of bad types is $n+1$. Thus, her expected payoff from such a deviation is $\sum_{n=0}^{N-1} q_n u_{n+1}^B$. Now note that

$$\begin{aligned}
\sum_{n=0}^{N-1} q_n u_{n+1}^B &= \sum_{n=0}^{N-1} q_n u_n^B + \sum_{n=0}^{N-1} q_n u_{n+1}^B - \sum_{n=0}^{N-1} q_n u_n^B \\
&= \sum_{n=0}^{N-1} q_n u_n^B + \sum_{n=1}^{N} q_n^- u_n^B - \sum_{n=0}^{N-1} q_n u_n^B \\
&= \sum_{n=0}^{N-1} q_n u_n^B + \sum_{n=0}^{N} q_n^- u_n^B - \sum_{n=0}^{N} q_n u_n^B \\
&= \sum_{n=0}^{N-1} q_n u_n^B - \sum_{n=0}^{N} \left(q_n - q_n^-\right) u_n^B \\
&\geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta_{q,q^-}. \qquad (2)
\end{aligned}$$

Here, the third equality follows because $q_0^- = q_N = 0$, and the final inequality follows because (recalling that $u_n^B \in [0,1]$)

$$\sum_{n=0}^{N} \left(q_n - q_n^-\right) u_n^B \leq \sum_{n: q_n \geq q_n^-} \left(q_n - q_n^-\right) u_n^B \leq \sum_{n: q_n \geq q_n^-} \left(q_n - q_n^-\right) = \Delta_{q,q^-}.$$

This establishes the rational player payoff bound.

To derive the social welfare bound, let $r_n = \Pr\left(\mathcal{B}_n | \theta_N = B\right)$, and note that

$$p_n = \begin{cases} (1-z)\, q_0 & \text{if } n = 0, \\ (1-z)\, q_n + z r_n & \text{if } 1 \leq n \leq N-1, \\ z r_N & \text{if } n = N. \end{cases}$$

Putting this together with the definition of U and the rational player payoff bound, we obtain

$$\begin{aligned}
U &= (1-z)\sum_{n=0}^{N-1} q_n u_n^R + z\sum_{n=1}^{N} r_n u_n^B \\
&\geq (1-z)\left(\sum_{n=0}^{N-1} q_n u_n^B - \Delta_{q,q^-}\right) + z\sum_{n=0}^{N} r_n u_n^B \\
&= \left((1-z)\,q_0 + \left(\sum_{n=1}^{N-1}(1-z)\,q_n + z r_n\right) + z r_N\right) u_n^B - (1-z)\,\Delta_{q,q^-} \\
&= \sum_{n=0}^{N} p_n u_n^B - (1-z)\,\Delta_{q,q^-}.
\end{aligned}$$

■

3.2 Smooth Type Distributions

The payoff bounds established in Lemma 1 are most significant when Δ_{q,q^-} is small. We say that a sequence $(N,p)_N$ with $N \to \infty$ (where, for each N, p is a symmetric prior on $\{R,B\}^N$) has a *smooth distribution of bad types* if

$$\lim_{N\to\infty} \Delta_{q,q^-} = 0.$$

Similarly, a sequence of games indexed by N, $(\Gamma)_N$, has a *smooth distribution of bad types* if this true of $(N,p)_N$. We now discuss when a sequence $(N,p)_N$ has a smooth distribution of bad types.

Suppose p is log-concave: $\frac{p_n}{p_{n-1}} \geq \frac{p_{n+1}}{p_n}$ for all $n \in \{1,\ldots,N-1\}$.[8] Then the maximum in (1) is attained by a set $\mathcal{N}$ that takes a "threshold" form $\mathcal{N} = \{n^*,\ldots,N\}$ for some threshold $n^* \in \{0,\ldots,N\}$. This yields

$$\Delta_{q,q^-} = q_{n^*-1} = \frac{N-n^*+1}{N}\frac{p_{n^*-1}}{1-z}.$$

Therefore, when p is log-concave, the sequence $(N,p)_N$ has an smooth distribution of bad types if and only if $\max_{n\leq N-1} \frac{N-n}{N} p_n \to 0$.

This is a mild condition. For example, when the players' types $(\theta_i)_{i\in I}$ are independent, p is log-concave,[9] and $\max_{n\leq N-1} \frac{N-n}{N} p_n \to 0$ whenever z remains bounded away from 0 and 1 as $N \to \infty$. More generally, a log-concave distribution of bad types is smooth if the probability that the number of bad types takes on any particular value n converges to zero as $N \to \infty$.

[8] See Bagnoli and Bergstrom (2005) for a survey of log-concave probability distributions, with many examples.

[9] Here $p_n = \binom{N}{n}\varepsilon^n(1-\varepsilon)^{N-n}$, and hence $\frac{p_n}{p_{n-1}} = \frac{N-n+1}{n}\frac{\varepsilon}{1-\varepsilon}$, which is decreasing in n.

For an example where the distribution of bad types is *not* smooth, consider independent types with zN held constant at some $\bar{n} \in \mathbb{N}$ as $N \to \infty$, so the distribution of the number of bad types converges to a Poisson distribution with parameter $\bar{n}$. Then

$$\Delta_{q,q^-} = q_{\bar{n}} = \frac{N-\bar{n}}{N}\binom{N}{\bar{n}}\left(\frac{\bar{n}}{N}\right)^{\bar{n}}\left(\frac{N-\bar{n}}{N}\right)^{N-\bar{n}-1}.$$

For instance, if $\bar{n}=1$—on average, there is exactly one bad player in the population—then

$$\Delta_{q,q^-} = \left(\frac{N-1}{N}\right)^{N-1} \sim \frac{1}{e}.$$

Thus, Lemma 1 can provide a meaningful bound even if on average there is only a single bad player. If instead $\bar{n}$ is sufficiently large, then Stirling's approximation gives

$$\Delta_{q,q^-} = \frac{N-\bar{n}}{N}\binom{N}{\bar{n}}\left(\frac{\bar{n}}{N}\right)^{\bar{n}}\left(\frac{N-\bar{n}}{N}\right)^{N-\bar{n}-1} \sim \frac{1}{\sqrt{2\pi\bar{n}}}.$$

Thus, Lemma 1 can provide a very tight bound even if the expected number of bad types stays finite as $N \to \infty$.

If p is not log-concave, then Δ_{q,q^-} need not converge to 0 even if $\max_{n \le N-1} \frac{N-n}{N} p_n \to 0$. For example, $\Delta_{q,q^-} = 1$ if the number of bad types is known in advance to be even. In this (rather artificial) case, the conclusion of Lemma 1 is vacuous.

4 Anti-Folk Theorem for Games with a Pairwise Dominant Action

We say that an action $a^* \in A$ is *pairwise dominant* if there exists a positive number $c > 0$ such that, for any action $a \neq a^*$, if player i takes a^*, player j takes a, and the remaining players take any actions $a_{-ij} \in A^{(N-2)}$, then player i's payoff exceeds player j's by at least c: that is,

$$u(a^*, a, a_{-ij}) - u(a, a^*, a_{-ij}) > c \text{ for all } a\,(\neq a^*) \in A, a_{-ij} \in A^{(N-2)}.$$

To interpret this definition, note that if the impact of a single opponent's action on a player's payoff is small, then $u(a, a^*, a_{-ij}) \approx u(a, a, a_{-ij})$, so the definition of a pairwise dominant action reduces to that of a dominant action, with c equal to the minimum payoff gain from taking a^* rather than another action. For example, this equivalence holds in large-population anonymous random matching games, where $u(a_i, a_{-i}) = (1/(N-1)) \sum_{j \neq i} \hat{u}(a_i, a_j)$ for some function $\hat{u} : A^2 \to \mathbb{R}$

fixed independent of N. More generally, a dominant action is also pairwise dominant if it imposes a negative externality on other players, in that $u(a, a^*, a_{-ij}) \leq u(a, a, a_{-ij})$ for all a, a_{-ij}; and a non-dominant action can be pairwise dominant only if it imposes a sufficiently large negative externality.

In this section, we assume that the action a^* played by bad types is pairwise dominant.[10] Denote social welfare when everyone takes the pairwise dominant action by $U^* = u(a^*, ..., a^*) \in [0, 1]$. We also let $b > 0$ denote the greatest impact on social welfare that can result from a player switching from a^* to another action, given by

$$b = \sup_{a_i \in A, a_{-i} \in A^{(N-1)}} \left| \sum_{j=1}^{N} \left(u_j(a_1, \ldots, a_{i-1}, a_i, a_{i+1}, \ldots, a_N) - u_j(a_1, \ldots, a_{i-1}, a_i^*, a_{i+1}, \ldots, a_N) \right) \right|.$$

Example 1 (Prisoner's Dilemma with Anonymous Random Matching) *Suppose that in each period players match in pairs (uniformly at random and independently across periods) to play the prisoner's dilemma stage game*

	C	D
C	$\frac{1+L}{1+G+L}, \frac{1+L}{1+G+L}$	$0, 1$
D	$1, 0$	$\frac{L}{1+G+L}, \frac{L}{1+G+L}$

(3)

where $G, L > 0$. This is an anonymous N-player game, where a player's stage-game payoff when she takes action $a \in \{C, D\}$, m out of her $(N-1)$ opponents take action D, and her remaining $(N-m-1)$ opponents take action C, is given by

$$\frac{N-m-1}{N-1} u(a, C) + \frac{m}{N-1} u(a, D).$$

In this game, action D is pairwise dominant, with

$$\begin{aligned} c &= \frac{1}{1+G+L} \left(\min\{G, L\} + \frac{1}{N-1} (1 + \max\{G, L\}) \right), \\ b &= \frac{1 + |G-L|}{1+G+L}. \end{aligned}$$

Our main result is the following anti-folk theorem for anonymous repeated games with a pairwise dominant action.

[10] We thus consider games that have a pairwise dominant action. Clearly, any game has at most one such action.

Theorem 1 *For any anonymous repeated game Γ with a pairwise dominant action, in any Nash equilibrium social welfare U satisfies*

$$|U - U^*| \leq (1-z)\, b \frac{1+c}{c} \Delta_{q,q^-}. \tag{4}$$

In particular, for any sequence $(\Gamma)_N$ of anonymous repeated games with a pairwise dominant action satisfying $\liminf_{N\to\infty} c_N > 0$ *and* $\limsup_{N\to\infty} b_N < \infty$ *and a smooth distribution of bad types, and any corresponding sequence of Nash equilibrium social welfare levels $(U)_N$, we have*

$$\lim_{N\to\infty} |U_N - U_N^*| = 0. \tag{5}$$

For example, consider the repeated prisoner's dilemma with anonymous random matching where bad types always defect. If we fix the payoff parameters G and L and vary N and δ, then along any sequence with a smooth distribution of bad types, social welfare converges to the payoff from mutual defection, $L/(1+G+L)$. Crucially, this conclusion does not depend on how δ varies along the sequence: for instance, it applies even if $(1-\delta)N \to 0$.

To see the intuition for Theorem 1, note that as in Lemma 1, bad players' expected discounted equilibrium payoffs cannot be much greater than rational players'. However, since a^* is pairwise dominant, bad players' payoffs exceed rational players' by at least c multiplied by the expected discounted frequency with which rational players take actions other than a^*. Therefore, this frequency must be small—that is, rational players must almost always take a^*. Finally, when bad players always take a^* and rational players almost always take a^*, social welfare is close to U^*.[11]

Proof. Fix a symmetric Nash equilibrium σ. For $n \in \{1, \ldots, N-1\}$, let γ_n denote the "expected discounted frequency" with which a rational player takes an action other than a^* when there are n bad types, given by

$$\gamma_n = (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} \sum_{h_i^t} \Pr^\sigma \left(h_i^t | \theta_i = R, \mathcal{B}_n\right) \left(1 - \sigma_{i,t}\left(h_i^t\right)[a^*]\right).$$

We note that

$$u_n^B \geq u_n^R + \gamma_n c \quad \text{for all } n \in \{1, \ldots, N-1\}. \tag{6}$$

[11] More precisely, we show that rational players take actions other than a^* with frequency at most $((1+c)/c)\,\Delta_{q,q^-}$. Since the probability that a given player is rational is $1-z$ and the payoff impact of switching a single player's action from a^* to another action is at most b, we see that $|U - U^*|$ is at most $(1-z)\, b\, ((1+c)/c)\, \Delta_{q,q^-}$.

This inequality holds because, in any period where a rational player takes an action other than a^*, a bad player's payoff exceeds her payoff by at least c. So if a rational player takes an action other than a^* in period t with probability

$$\gamma_{n,t} = \sum_{h_i^t} \Pr^\sigma \left(h_i^t | \theta_i = R, \mathcal{B}_n\right) \left(1 - \sigma_{i,t}\left(h_i^t\right)[a^*]\right),$$

a bad player's expected period-t payoff exceeds hers by at least $\gamma_{n,t}c$, and taking a discounted sum over periods implies that a bad player's repeated game payoff exceeds a rational player's by at least $(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\gamma_{n,t}c = \gamma_n c$.

Combining the rational player payoff bound from Lemma 1 with (6), and recalling that $u_0^B = 1$ by convention, we obtain

$$\Delta_{q,q^-} \geq \sum_{n=0}^{N-1} q_n \left(u_n^B - u_n^R\right) \geq \sum_{n=1}^{N-1} q_n \left(u_n^B - u_n^R\right) \geq \sum_{n=1}^{N-1} q_n \gamma_n c.$$

Now define $\gamma = \sum_{n=1}^{N-1} q_n \gamma_n$. Since $q_0 = q_0 - q_0^- \leq \Delta_{q,q'}$, we have

$$\gamma = q_0\gamma_0 + \sum_{n=1}^{N} q_n\gamma_n \leq \Delta_{q,q^-} + \frac{1}{c}\Delta_{q,q^-} = \frac{1+c}{c}\Delta_{q,q^-}.$$

Finally, the ex ante probability that a given player takes an action other than a^* in period t equals $(1-z)\sum_{n=0}^{N-1} q_n\gamma_{n,t}$. Hence, (per-capita) expected social welfare in period t differs from U^* by at most $(1-z)\,b\sum_{n=0}^{N-1} q_n\gamma_{n,t}$. Therefore, ex ante expected welfare differs from U^* by at most

$$(1-z)\,b\,(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\sum_{n=0}^{N-1} q_n\gamma_{n,t} = (1-z)\,b\gamma \leq (1-z)\,b\frac{1+c}{c}\Delta_{q,q^-}.$$

This yields (4), and taking $\Delta_{q,q^-} \to 0$ yields (5) ■

Let us clarify the difference between dominant and pairwise dominant actions under the conditions on payoffs required by Theorem 1. If an action a^* is strictly dominant, the minimum payoff gain from taking a^* rather than another action is bounded away from 0 as $N \to \infty$, and the total externality b imposed by switching from a^* to another action is bounded as $N \to \infty$, then a^* is pairwise dominant for sufficiently large N, and thus Theorem 1 implies that a^* is almost always played. But the converse is false: even when $N \to \infty$ while c and b remain bounded (so the payoff conditions of Theorem 1 are satisfied), a pairwise dominant action does not need to be dominant.

Therefore, Theorem 1 applies in some games without a strictly dominant action.

For example, suppose that $A = \left\{a^0, a^1, a^*\right\}$ and payoffs are given by

$$u\left(a_i, a_{-i}\right) = \begin{cases} 1/2 & \text{if } a_i = a^*, \\ 1 & \text{if } a_i = a^1 \text{ and } a_j = a^0 \text{ for all } j \neq i, \\ 0 & \text{otherwise.} \end{cases}$$

Note that a^* is pairwise dominant with $c = 1/2$, since whenever some player takes a^* and another player takes a different action, the first player's payoff is 1/2 and the second player's payoff is 0; however, a^* is not dominant, because $a_i = a^1$ is the unique best response when $a_j = a^0$ for all $j \neq i$. Moreover, b is also equal to 1/2: switching a player's action from a^* to a^1 changes her own payoff by 1/2 without affecting anyone else's payoff; and switching her action from a^* to a^0 decreases her own payoff by 1/2, while either affecting no one else's payoff or increasing a single other player's payoff by 1. Thus, assuming a smooth distribution of bad types, Theorem 1 implies that a^* is almost always played.

However, when $N \to \infty$ a pairwise dominant action a^* does satisfy the weaker condition that, for any other action a, a^* is a strictly better-response than a against any mixture of a's and a^*'s.

Proposition 1 *Fix a sequence of stage games $(N, A, u)_N$ with a pairwise dominant action a^* satisfying $\liminf_{N\to\infty} c_N > 0$ and $\limsup_{N\to\infty} b_N < \infty$. There exists $\bar{N}$ such that, for all $N > \bar{N}$, all $a \in A$, and all $M \in \{0, \ldots, N-1\}$, we have*

$$u\left(a^*, (a)^M, (a^*)^{N-M-1}\right) > u\left(a, (a)^M, (a^*)^{N-M-1}\right),$$

where $(a)^M$ is the vector of M a's and $(a^*)^{N-M-1}$ is the vector of $(N-M-1)$ a^*'s.

Proof. The proof follows easily from the definitions of c and b and is thus omitted. ■

We also note that Theorem 1 generalizes to stochastic games where the payoff function u depends on the profile of players' histories $h^t = \left(h_i^t\right)_{i\in I}$. Specifically, letting $u^t\left(a; h^t\right)$ denote the period-t stage-game payoff at action profile a and history profile h^t, assume that

$$u^t\left(a^*, a, a_{-ij}; h^t\right) - u^t\left(a, a^*, a_{-ij}; h^t\right) > c \text{ for all } a\,(\neq a^*) \in A, a_{-ij} \in A^{(N-2)}, t, h^t, \tag{7}$$

and let

$$b=\sup_{a_i\in A,a_{-i}\in A^{(N-1)},t,h^t}\left|\sum_{j=1}^{N}\left(u_j^t\left(a_1,\ldots,a_{i-1},a_i,a_{i+1},\ldots,a_N;h^t\right)-u_j^t\left(a_1,\ldots,a_{i-1},a_i^*,a_{i+1},\ldots,a_N;h^t\right)\right)\right|.$$

With these values for c and b, Theorem 1 holds verbatim for stochastic games, by the same proof.

Example 2 (Non-Uniform Matching) *Consider again the prisoner's dilemma with anonymous random matching and payoff matrix (3), but now allow the matching process to be non-uniform, non-stationary, and history-dependent. Specifically, assume that, given history profile* $h^t=\left(h_i^t\right)_i$ *at the beginning of period* t*, players* i *and* j *meet in period* t *with probability* $\psi_{ij}\left(h^t\right)$*. Continue to assume that the matching process is symmetric across players ex ante: for any permutation* π *on* I*, we have* $\psi_{ij}\left(\left(h_k^t\right)_k\right)=\psi_{\pi(i)\pi(j)}\left(\left(h_{\pi(k)}^t\right)_k\right)$ *for all* $\left(i,j,t,h^t\right)$*. In order for the action* D *to remain pairwise dominant, we must assume that the meeting probabilities between any two players cannot be too unequal: otherwise, some player who takes* D *could receive a lower payoff than another player who takes* C*, if the latter player is much more likely to meet a partner who takes* C*. In particular, letting*

$$R:=\frac{\sup_{i,j,t,h^t}\psi_{ij}\left(h^t\right)}{\inf_{i,j,t,h^t}\psi_{ij}\left(h^t\right)},$$

we assume that

$$R^2<1+G.$$

This assumption implies that action D *remains pairwise dominant, so Theorem 1 holds with the appropriate choice of* $c>0$*.*[12]

(To see why D *is pairwise dominant whenever* $R^2<1+G$*, suppose player* i *takes* D*, player* j *takes* C*, and fraction* α_{-ij} *of the remaining players take* C*. Note that, for any set of players* $\mathcal{S}\subset I$*, any* h^t*-measurable events* $\mathcal{E}$ *and* $\mathcal{E}'$*, and any* $k\in I$*, we have*

$$1/R\leq\frac{\Pr\left(\mu_k\left(t\right)\in\mathcal{S}|\mathcal{E}\right)}{\Pr\left(\mu_k\left(t\right)\in\mathcal{S}|\mathcal{E}'\right)}\leq R.$$

Hence, by Bayes' rule, the probability that player i *meets an opponent who takes* C *is at least* α_{-ij}/R*, and the probability that player* j *meets an opponent who takes* C *is at most* $\min\left\{R\alpha_{-ij},1\right\}$*. So the*

[12] And with the same value of b as in the uniform random matching case: $b=\left(1+|G-L|\right)/\left(1+G+L\right)$.

difference between player i's payoff and player j's payoff is at least

$$\frac{\alpha_{-ij}}{R}+\left(1-\frac{\alpha_{-ij}}{R}\right)\frac{L}{1+G+L}-\min\left\{R\alpha_{-ij},1\right\}\frac{1+L}{1+G+L}.$$

This expression is strictly positive whenever $R^2<1+G$.[13]*)*

5 General Games

The payoff bounds established in Lemma 1 can also be useful in games without a pairwise dominant action. As an example of such an analysis, in this section we derive an implication of Lemma 1 involving the concave closure of the payoff function, and use it to show that industry profits in linear Cournot oligopoly converge to zero as the number of firms increases.

Fixing a strategy profile σ, let $\boldsymbol{\alpha}_t(\sigma)\in\Delta\left(A^N\right)$ denote the resulting distribution over action profiles in period t. Let $\boldsymbol{\alpha}(\sigma)=(1-\delta)\sum_{t=1}^{\infty}\delta^t\boldsymbol{\alpha}_t(\sigma)\in\Delta\left(A^N\right)$. That is, for each action profile $\mathbf{a}\in A^N$, $\boldsymbol{\alpha}(\sigma)$ is the "discounted frequency" with which $\mathbf{a}$ is played under σ.

Next, let $U(\mathbf{a})=(1/N)\sum_{i=1}^{N}u_i(\mathbf{a})$ denote social welfare at action profile $\mathbf{a}\in A^N$, and let $U(\boldsymbol{\alpha})=\sum_{\mathbf{a}\in A^N}\boldsymbol{\alpha}[\mathbf{a}]U(\mathbf{a})$ denote expected social welfare at action profile distribution $\boldsymbol{\alpha}\in\Delta\left(A^N\right)$. Let $\bar{U}:\Delta\left(A^N\right)\to[0,1]$ denote the *concavification* of U: that is, the smallest concave function $\bar{U}$ that satisfies $\bar{U}(\boldsymbol{\alpha})\geq U(\boldsymbol{\alpha})$ for all $\boldsymbol{\alpha}\in\Delta\left(A^N\right)$. Finally, let $\underline{u}:\Delta\left(A^{N-1}\right)\to[0,1]$ denote the *convexification* of the function $u(a^*,\cdot):\Delta\left(A^{N-1}\right)\to[0,1]$ given by $u(a^*,\boldsymbol{\alpha}_{-i})=\sum_{a_{-i}\in A^{(N-1)}}\boldsymbol{\alpha}_{-i}[a_{-i}]u(a^*,a_{-i})$: that is, the greatest convex function $\underline{u}$ that satisfies $\underline{u}(a^*,\boldsymbol{\alpha}_{-i})\leq u(a^*,\boldsymbol{\alpha}_{-i})$ for all $\boldsymbol{\alpha}_{-i}\in\Delta\left(A^{N-1}\right)$.

Proposition 2 *For any anonymous game and any symmetric Nash equilibrium σ, we have*

$$\bar{U}(\boldsymbol{\alpha}(\sigma))\geq\underline{u}(\boldsymbol{\alpha}_{-i}(\sigma))-(1-z)\Delta_{q,q^-}.$$

Proposition 2 follows from the social welfare bound of Lemma 1, because (as the proof shows) $\bar{U}(\boldsymbol{\alpha}(\sigma))\geq U$ and $\underline{u}(\boldsymbol{\alpha}(\sigma))\leq\sum_{n=0}^{N}p_n u_n^B$.

Proof. Let $\boldsymbol{\alpha}_t=\boldsymbol{\alpha}_t(\sigma)$ and $\boldsymbol{\alpha}=\boldsymbol{\alpha}(\sigma)$. We have

$$\bar{U}(\boldsymbol{\alpha})=\bar{U}\left((1-\delta)\sum_t\delta^t\boldsymbol{\alpha}_t\right)\geq(1-\delta)\sum_t\delta^t\bar{U}(\boldsymbol{\alpha}_t)\geq(1-\delta)\sum_t\delta^tU(\boldsymbol{\alpha}_t)=U,$$

[13] This follows from straightforward algebra, considering separately the cases where $R\alpha_{-ij}<1$ and $R\alpha_{-ij}\geq1$.

where the first inequality follows because $\bar{U}$ is concave and the second follows because $\bar{U}$ is everywhere greater than U. Similarly, we have

$$\begin{aligned}\underline{u}\left(\boldsymbol{\alpha}_{-i}\right) &= \underline{u}\left(a^*, (1-\delta)\sum_t \delta^t \boldsymbol{\alpha}_{-i,t}\right) \\ &\leq (1-\delta)\sum_t \delta^t \underline{u}\left(a^*, \boldsymbol{\alpha}_{-i,t}\right) \leq (1-\delta)\sum_t \delta^t u\left(a^*, \boldsymbol{\alpha}_{-i,t}\right) \leq \sum_{n=0}^{N} p_n u_n^B,\end{aligned}$$

where the first inequality follows because $\underline{u}$ is convex, the second follows because $\underline{u}$ is everywhere less than $u\left(a^*, \cdot\right)$, and the third follows because of our convention that $u_0^B = 1$. The result follows from combining these inequalities with the social welfare bound of Lemma 1. ■

Example 3 (Linear Cournot oligopoly) *Suppose that in every period each of N firms produces quantity* $a_i \geq 0$ *and payoffs are given by* $u_i\left(a_i, a_{-i}\right) = \max\left\{1 - \sum_{j=1}^{N} a_j, 0\right\} a_i$.[14] *We assume that bad types always take the static Nash equilibrium action, so* $a^* = 1/\left(N+1\right)$. *(An interpretation is that bad types are not aware that the other firms are trying to collude, and thus expect the static equilibrium to be played.) We also assume that the distribution of the number of bad types is smooth. Under these assumptions, expected industry profits* $\sum_{i=1}^{N} U_i$ *converge to zero along any sequence of Nash equilibria as* $N \to \infty$ *(regardless of how* δ *may vary with* N*, including the case where* $\left(1-\delta\right) N \to 0$*).*[15]

To see this, first note that $u_n^B \in [0, 1/\left(N+1\right)]$ *for each* n. *This implies that* Δ_{q,q^-} *can be replaced by* $\Delta_{q,q^-}/\left(N+1\right)$ *in the statements of Lemma 1 and Proposition 2.*[16]. *Moreover, it is without loss to restrict attention to symmetric equilibria in which industry output* $\sum_i a_i$ *is bounded by* 1 *with probability* 1.[17] *Note that the proof of Proposition 2 involves only on-path action profiles, so the resulting payoff bound continues to apply if we restrict attention to output profiles where* $\sum_i a_i \leq 1$ *when defining* $\bar{U}$ *and* $\underline{u}$. *With this restriction, industry profits are concave in* $\sum_i a_i$, *and hence in* $\mathbf{a}$.

[14] We thus normalize marginal costs to zero, and assume that prices are bounded by zero to keep payoffs bounded. The assumption that the kink in the demand curve is located at marginal cost is not essential. Also, as our proofs have covered only finite games (although they extend to the case where A is a compact metric space and u is continuous), the set of feasible quantity levels may be taken to be finite.

[15] In this example we consider total payoffs rather than per-capita payoffs, since with a fixed demand curve the maximum *feasible* per-firm profits go to zero as $N \to \infty$.

[16] In particular, the last step of the derivation of (2) uses the fact that $u_n^B \in [0,1]$ to conclude that $\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta_{q,q^-}$; if the upper bound for u_n^B is replaced by $1/\left(N+1\right)$, the corresponding conclusion is $\sum_{n=0}^{N-1} q_n u_n^R \geq \sum_{n=0}^{N-1} q_n u_n^B - \Delta_{q,q^-}/\left(N+1\right)$.

[17] This follows by considering the relaxed problem where the only available strategies are the rational-type equilibrium strategy and the bad-type strategy. In this relaxed problem, if industry output ever exceeds 1, the output of rational types can be reduced so that industry output exactly equals 1 without affecting anyone's payoff from following either strategy.

Therefore, $\bar{U}(\boldsymbol{\alpha}) = U(\boldsymbol{\alpha})$ *for all* $\boldsymbol{\alpha} \in \Delta\left(A^N\right)$. *Similarly,* $u(a^*, a_{-i}) = \left(1 - a^* - \sum_{j \neq i} a_j\right) a^*$, *which is linear in* $\sum_{j \neq i} a_j$, *and hence in* a_{-i}. *Therefore,* $\underline{u}(\boldsymbol{\alpha}_{-i}) = u(a^*, \boldsymbol{\alpha}_{-i})$ *for all* $\boldsymbol{\alpha}_{-i} \in \Delta\left(A^{N-1}\right)$. *Letting* $\boldsymbol{\alpha} = \boldsymbol{\alpha}(\sigma) \in \Delta\left(A^N\right)$, *Proposition 2 now implies that*

$$U(\boldsymbol{\alpha}) \geq u(a^*, \boldsymbol{\alpha}_{-i}) - \Delta_{q,q^-} / (N+1).$$

To complete the proof, let $\bar{a} = \mathbb{E}^{\boldsymbol{\alpha}}\left[\sum_{i=1}^N a_i\right]$. *Since* $\left(1 - \sum_{i=1}^N a_i\right) \sum_{i=1}^N a_i$ *is concave in* $\sum_{i=1}^N a_i$, *we have*

$$U(\boldsymbol{\alpha}) \leq \frac{1}{N}(1 - \bar{a})\bar{a},$$

Similarly, since $\mathbb{E}^{\boldsymbol{\alpha}}\left[\sum_{j \neq i} a_j\right] = \frac{N-1}{N}\bar{a}$ *by symmetry, we have*

$$u(a^*, \boldsymbol{\alpha}_{-i}) = \left(1 - \frac{N-1}{N}\bar{a} - \frac{1}{N+1}\right)\frac{1}{N+1}.$$

We conclude that

$$\begin{aligned} \left(1 - \frac{N-1}{N}\bar{a} - \frac{1}{N+1}\right)\frac{1}{N+1} - \frac{1}{N}(1-\bar{a})\bar{a} &\leq \frac{1}{N+1}\Delta_{q,q^-}, \text{ or} \\ 1 - \frac{N-1}{N}\bar{a} - \frac{1}{N+1} - \frac{N+1}{N}(1-\bar{a})\bar{a} &\leq \Delta_{q,q^-}. \end{aligned}$$

This inequality implies that if $\lim_{N\to\infty} \Delta_{q,q^-} = 0$ *then* $\lim_{N\to\infty} \bar{a} = 1$ *as well. (Otherwise, the left-hand side would not converge to* 0.*) Finally, industry profits equal* $NU(\boldsymbol{\alpha}) \leq (1-\bar{a})\bar{a}$, *which converges to zero whenever* $\lim_{N\to\infty} \bar{a} = 1$.

6 Discussion

We conclude by assessing the prospects for extending our results to settings with multiple types, to mechanism design problems, and to non-anonymous games.

6.1 Multiple Commitment Types

Our results extend straightforwardly to the case with one rational type and K commitment types, each of whom is committed to an arbitrary repeated game strategy σ^k. In this case, let the vector $\mathbf{n} \in \{0, ..., N\}^K$ count the realized number of players of each commitment type, and let $q_{\mathbf{n}}$ be the probability of $\mathbf{n}$ conditional of the event that a given player is rational. For each commitment strategy σ^k, let $q_{\mathbf{n}}^{\sigma^k}$ denote the probability that the the realized number of players of each commitment type differs from $\mathbf{n}$ in that one fewer player is committed to σ^k. Note that, if a single rational player deviates by playing σ^k, then $q_{\mathbf{n}}^{\sigma^k}$ is the probability that number of players who play each commitment strategy is given by $\mathbf{n}$. Let $\Delta_{q,q^{\sigma^k}} = \max_{\mathcal{N}\subset\{0,\ldots,N\}^K} \left|\sum_{\mathbf{n}\in\mathcal{N}} \left(q_{\mathbf{n}} - q_{\mathbf{n}}^{\sigma^k}\right)\right|$. A straightforward extension of Lemma 1 then implies that $\sum_{\mathbf{n}} q_{\mathbf{n}} u_{\mathbf{n}}^R \geq \sum_{\mathbf{n}} q_{\mathbf{n}} u_{\mathbf{n}}^{\sigma^k} - \Delta_{q,q^{\sigma^k}}$ for every commitment strategy σ^k, where the sum is taken over all vectors $\mathbf{n}$, and $u_{\mathbf{n}}^{\sigma^k}$ is the expected utility of the σ^k commitment type conditional on vector $\mathbf{n}$. Finally, in games with a pairwise dominant action a^*, Theorem 1 holds with $\Delta_{q,q^{Always\ a^*}}$ in place of Δ_{q,q^-} (by the same proof), with the slight modification that an additional $+bz$ term must be added to the right-hand sides of equations (4) and (5) to reflect the fact that commitment types other than the *Always* a^* type can take actions besides a^*.

Although our results extend to the case with multiple commitment types, we have focused on the case with a single commitment type because calculating $\Delta_{q,q^{\sigma^k}}$ (the total variation distance between two K-dimensional probability distributions) is much simpler when $K = 1$. However, in two leading special cases computing $\Delta_{q,q^{\sigma^k}}$ for arbitrary K is not much harder than it is when $K = 1$. The first is when types are independent across players: in this case, to compute $\Delta_{q,q^{\sigma^k}}$ we need only keep track of the number of σ^k commitment types, as in the $K = 1$ case. The second is when types can be correlated but, conditional on the event that a given set of players are not rational, their specific commitment types are determined independently: in this case, to compute $\Delta_{q,q^{\sigma^k}}$ we need only keep track of the number of σ^k commitment types and the number of rational types, and thus calculate the distance between two 2-dimensional distributions.

6.2 Multiple Rational Types, Incentive Compatibility, and Mechanism Design

Lemma 1 can also be extended to settings with multiple rational types. Suppose there are K rational types and (for simplicity) no commitment types. Let $\mathbf{n} \in \{0, ..., N\}^K$ count the realized number of players of each type. Fixing a pair of types $\left(\theta_i, \hat{\theta}_i\right)$ and conditioning on the event that a given player has type θ_i, let $q_{\mathbf{n}}^{\theta_i}$ be the probability of $\mathbf{n}$, let $q_{\mathbf{n}}^{\theta_i,\hat{\theta}_i}$ denote the probability that the realized number of players of each type differs from $\mathbf{n}$ in that one fewer player has type $\hat{\theta}_i$ and one more player has type θ_i, and let $\Delta_{q^{\theta_i},q^{\theta_i,\hat{\theta}_i}} = \max_{\mathcal{N}\subset\{0,\ldots,N\}^K} \left|\sum_{\mathbf{n}\in\mathcal{N}} \left(q_{\mathbf{n}}^{\theta_i} - q_{\mathbf{n}}^{\theta_i,\hat{\theta}_i}\right)\right|$. Next, fixing a symmetric Nash equilibrium, let $u_{\mathbf{n}}^{\theta_i}$ denote the equilibrium expected utility of a type θ_i player conditional on the vector $\mathbf{n}$. Let $u_{\mathbf{n}}^{\theta_i,\hat{\theta}_i}$ denote the expected utility that a type θ_i player receives from following the equilibrium strategy of a type $\hat{\theta}_i$ player, conditional on $\mathbf{n}$. (Put differently, $u_{\mathbf{n}}^{\theta_i,\hat{\theta}_i}$

is the expected utility according to the type θ_i utility function of the equilibrium outcome obtained by type $\hat{\theta}_i$ conditional on $\mathbf{n}$.) Lemma 1 then implies that $\sum_{\mathbf{n}} q_{\mathbf{n}}^{\theta_i} u_{\mathbf{n}}^{\theta_i} \geq \sum_{\mathbf{n}} q_{\mathbf{n}}^{\theta_i} u_{\mathbf{n}}^{\theta_i,\hat{\theta}_i} - \Delta_{q^{\theta_i},q^{\theta_i,\hat{\theta}_i}}$.

Of course, we have simply shown that this inequality is one implication of incentive compatibility: the fact that a type θ_i player prefers to follow her own equilibrium strategy rather than the equilibrium strategy of type $\hat{\theta}_i$. But it may be a useful implication in some mechanism design problems. For example, this form of Lemma 1 can be used to give an elementary proof of a version of Mailath and Postlewaite's (1990) impossibility theorem for large-population public good provision, assuming that the type space is discrete and the prior is symmetric and satisfies $\lim_{N\to\infty} \Delta_{q^{\theta_i},q^{\theta_i,\hat{\theta}_i}} = 0$ for all $\left(\theta_i, \hat{\theta}_i\right)$ but is not necessarily independent. (In contrast, Mailath and Postlewaite's proof requires independence, as do all other proofs of their result that we aware of, such as that of al-Najjar and Smorodinsky (2000).[18])

The proof can be easily sketched: Recall that a type θ_i player receives utility $\theta_i y - t_i$, where $y \in \{0, 1\}$ is the public provision level and t_i is the player's payment. Thus, $\sum_{\mathbf{n}} q_{\mathbf{n}}^{\theta_i} \left(u_{\mathbf{n}}^{\theta_i,\hat{\theta}_i} - u_{\mathbf{n}}^{\theta_i}\right)$ equals the difference between the expected payment of a type θ_i player and a type $\hat{\theta}_i$ player, according to the beliefs of a type θ_i player. Taking $\hat{\theta}_i$ to be the lowest type and normalizing this type to zero, individual rationality for type $\hat{\theta}_i$ implies that this type makes non-positive payments, so $\sum_{\mathbf{n}} q_{\mathbf{n}}^{\theta_i} \left(u_{\mathbf{n}}^{\theta_i,0} - u_{\mathbf{n}}^{\theta_i}\right)$ weakly exceeds the expected payment of a type θ_i player according to her own beliefs, $t(\theta_i)$. Thus, for each type θ_i, Lemma 1 implies that $t(\theta_i) \leq \Delta_{q^{\theta_i},q^{\theta_i,0}}$.[19] Since $\lim_{N\to\infty} \Delta_{q^{\theta_i},q^{\theta_i,0}} = 0$ for each θ_i, we have $\lim_{N\to\infty} t(\theta_i) = 0$ for each θ_i. Taking an expectation over θ_i and applying the law of iterated expectations then implies that ex ante expected per-capita payments converge to zero as $N \to \infty$. Therefore, if the per-capita cost of providing the good is positive and constant in N (as Mailath and Postlewaite assume), the probability that it is provided also converges to zero as $N \to \infty$.

[18] Independence (combined with full support) is much stronger than our condition that $\lim_{N\to\infty} \Delta_{q^{\theta_i},q^{\theta_i,\hat{\theta}_i}} = 0$ for all $\left(\theta_i, \hat{\theta}_i\right)$. For example, our condition is satisfied whenever types are conditionally independent with full support (given some common random variable). On the other hand, Mailath and Postlewaite allow continuous types and an asymmetric prior. We can allow continuous types if the prior satisfies an appropriate continuity condition, but symmetry is crucial for our approach. Also, while Mailath and Postlewaite's proof requires independence, in Appendix 2 of their paper they present an example with correlated types, the logic of which is similar to that of our result.

[19] Recall that Lemma 1 assumes bounded utilities, which in the current context implies bounded transfers. This explains why mechanisms like those of Crémer and McLean (1988) are not effective. Note that Mailath and Postlewaite's Appendix 2 example similarly assumes bounded transfers.

6.3 Non-Anonymous Games

Our analysis depends critically on the anonymity assumption: a player's payoff is a function of her own action and the number of opponents taking each action. This assumption is what makes the distribution of the number of bad types so important. To apply our approach in more general games, one would need to find some statistic ω of the vector of players' types $(\theta_i)_{i \in I}$ with the properties that (1) a player's payoff depends only on her own action and type and ω, and (2) the distribution of ω is not very responsive to a change in a single player's type. In anonymous games, ω is the number of players of each type. In games with multiple populations (e.g., buyers and sellers) where players are anonymous within each population (as in the "semi-anonymous" games studied by Kalai, 2004), ω could be taken to be the number of players of each type in each population. Whether there are other interesting classes of games where such a statistic can be found is an open question.

An important example to which our results do *not* directly extend is the repeated prisoner's dilemma with non-anonymous random matching (where players observe their partners' identities before taking actions). In this game, for any N and any distribution of the number of bad types, it is straightforward to support cooperation among rational types when δ is sufficiently high: simply prescribe "bilateral grim trigger strategies," where each player views herself as playing a separate 2-player repeated games with each opponent and plays grim trigger in all of them.

However, some trace of the negative conclusion of Theorem 1 does survive in non-anonymous random matching games. Bilateral grim trigger strategies are robust to allowing bad types but support cooperation only if $(1-\delta) N \to 0$, and are thus ineffective in very large populations. In contrast, Kandori (1992) and Ellison (1994) showed that contagion strategies support cooperation whenever $(1-\delta) \log N \to 0$; however, such strategies are not robust to bad types. In a companion paper (Sugaya and Wolitzky, 2020), we generalize this observation to show that introducing bad types logarithmically decreases the maximum population size for which cooperation is sustainable in the repeated prisoner's dilemma with non-anonymous random matching.

References

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

[2] Al-Najjar, Nabil I., and Rann Smorodinsky. "Large Nonanonymous Repeated Games." *Games and Economic Behavior* 37.1 (2001): 26-39.

[3] Awaya, Yu, and Vijay Krishna. “On Communication and Collusion.” *American Economic Review* 106.2 (2016): 285-315.

[4] Awaya, Yu, and Vijay Krishna. “Communication and Cooperation in Repeated Games.” *Theoretical Economics* 14.2 (2019): 513-553.

[5] Bagnoli, Mark, and Ted Bergstrom. “Log-Concave Probability and its Applications.” *Economic Theory* 26.2 (2005): 445-469.

[6] Crémer, Jacques, and Richard P. McLean. “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions.” *Econometrica* 56.6 (1988): 1247-1257.

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

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

[9] Fudenberg, Drew, and David K. Levine. “Reputation and Equilibrium Selection in Games with a Patient Player.” *Econometrica* 57.4 (1989): 759-778

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

[11] Green, Edward J. “Noncooperative Price Taking in Large Dynamic Markets.” *Journal of Economic Theory* 22.2 (1980): 155-182.

[12] Kalai, Ehud. “Large Robust Games.” *Econometrica* 72.6 (2004): 1631-1665.

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

[14] Levine, David K., and Wolfgang Pesendorfer. “When are Agents Negligible?” *American Economic Review* 85.5 (1995): 1160-1170.

[15] Mailath, George J., and Andrew Postlewaite. “Asymmetric Information Bargaining Problems with Many Agents.” *Review of Economic Studies* 57.3 (1990): 351-367.

[16] Mailath, George J., and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships.* Oxford University Press, 2006.

[17] Pai, Mallesh M., Aaron Roth, and Jonathan Ullman. “An Antifolk Theorem for Large Repeated Games.” *ACM Transactions on Economics and Computation (TEAC)* 5.2 (2017): 10.

[18] Sabourian, Hamid. “Anonymous Repeated Games with a Large Number of Players and Random Outcomes.” *Journal of Economic Theory* 51.1 (1990): 92-110.

[19] Sugaya, Takuo, and Alexander Wolitzky. “Communication and Community Enforcement.” *Working paper* (2020).

Communication and Community Enforcement*

Takuo Sugaya and Alexander Wolitzky

Stanford GSB and MIT

September 3, 2020

Abstract

We study the repeated prisoner's dilemma with random matching, a canonical model of community enforcement with decentralized information. We assume (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 assumptions preclude cooperation in the absence of explicit communication, but that introducing within-match cheap talk communication restores cooperation. Thus, communication is essential for community enforcement.

Keywords: communication, community enforcement, repeated games, incomplete information

JEL codes: C72, C73, D83

*For helpful comments, we thank Daron Acemoglu, Glenn Ellison, Jeff Ely, Drew Fudenberg, Yuval Heller, Matt Jackson, George Mailath, Stephen Morris, Satoru Takahashi, Omer Tamuz, and seminar particiants at Princeton, Stanford, and Wharton. Wolitzky acknowledges financial support from the NSF and the Sloan Foundation.

"He that filches from me my good name
Robs me of that which not enriches him
And makes me poor indeed."
—*Othello*, III.3.

1 Introduction

Everyday experience and a wealth of evidence from across the social sciences indicate that communication about the reputation of third parties—gossip—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 gossip tomorrow and ostracism the day after keeps us on good behavior today.

While gossip's role in supporting 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 gossip 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 gossip and individualized punishment rather than contagion-like strategies? Are these features truly essential for supporting cooperation, or are they merely quirks of the particular cooperative equilibrium in which we happen to find ourselves?

This paper establishes the essentiality of gossip 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—like 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 remem-

[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.

ber) 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: players can then treat the overall repeated game as a collection of two-player 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 1), 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 $(1-\delta) N \to 0$—that is, whenever bilateral interactions become frequent.[2]

In stark contract, our first main result (Theorem 1) shows that average payoffs converge to the mutual defection payoff along any sequence where $(1-\delta) N \to \infty$—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 $(1-\delta) N \to \infty$, 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 towards 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 can convey only one bit of information each period. We show that $O(N)$ bits are required to provide significant information about N players'

[2]To see why $(1-\delta) N \to 0$ corresponds to frequent bilateral interactions, suppose 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 interact $1/(\Delta \times (N-1)) \approx r/((1-\delta) N)$ times per unit of real time on average, $(1-\delta) N \to 0$ means that each pair of players interact frequently, while $(1-\delta) N \to \infty$ means that each of pair of players rarely interact.

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

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 $(1-\delta)\log N \to 0$. 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 each player keeps track of a "blacklist" of opponents whom she believes have ever defected against a rational player, players share their blacklists with each other prior to taking actions, and 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 are 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, we show that cheap talk enables cooperation with infrequent bilateral interactions if and only if the number of possible messages that a player can send grows exponentially in N—that is, if and only if the message set is rich enough to include binary summaries of the individual reputations of a positive fraction of the population. Thus, not only does the particular equilibrium we construct rely on "gossip" about specific players' past behavior, but any cooperative equilibrium must involve the communication of similarly rich information.

1.1 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 cooperation in a repeated random matching game, and 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).[3] 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 bad types, which motivates the current paper's focus on individualized punishment.[4]

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 through making the observation that a partner defected in the past informative of his being a bad type (in Heller and Mohlin). These papers are therefore less closely related to ours.[5]

In many papers on the folk theorem in repeated games, implicit communication through actions is just as effective as explicit communication through cheap talk. For example, this is the case in Hörner and Olszewski's (2006) folk theorem with almost perfect monitoring and 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.[6]

Several papers on community enforcement and repeated games on networks fix $\delta < 1$ and show

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

[4]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," (p. 578).

[5]One result closer in spirit to ours is Heller and Mohlin's Theorem 1, which shows that cooperation is impossible in the "offensive" (submodular) PD 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.

[6]Deb (2020) establishes a folk theorem for anonymous random matching games with explicit communication. Here communication has the distinctive role of serving to relax anonymity, as players can identify each other via endogenous "names."

that cooperation is easier to support when the news that a defection occurred spreads more quickly (Raub and Weesie, 1990; Ahn and Suominen, 2001; Dixit, 2003; Lippert and Spagnolo, 2011; Ali and Miller, 2013; Wolitzky, 2013; Balmaceda and 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).[7] 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, 2019). While we of course also insist that communication is incentive compatible, Theorem 3 shows that this constraint is ultimately not binding in our model.

2 Model

A set $I = \{1, \ldots, N\}$ of N players interact in discrete time, $t = 1, 2, \ldots$, 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.[8] 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. We assume that rational and bad players have the same payoff function, so a player's expected payoff equals $1 - \varepsilon$ times her expected payoff when she is rational (and plays optimally) plus ε times her expected payoff when she is bad (and is constrained to always play D).

[7] 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.

[8] We discuss generalizations to multiple "commitment types" and to correlated types in Section 5.1.

This assumption is mostly just for convenience; we discuss later on where it has bite.

Matching is *non-anonymous.* That is, at the beginning of each period t, every player i observes the identity 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 $h_i^t = \left(\left(\mu_{i,\tau}, a_{i,\tau}, a_{\mu_{i,\tau},\tau} \right)_{\tau=1}^{t-1}, \mu_{i,t} \right)$, with $h_i^1 = \mu_{i,1}$. In Section 4, we augment the game by allowing pre-play cheap talk communication within each match. The description of a history for player i will then also include 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 $\sigma_i \left(h_i^t \right)$ 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 by U_i (recalling that this value is a weighted average of her payoff when rational and when bad). Denote the population average payoff by $U = \sum_i U_i / N$. Note that, since the minmax payoff is 0 and the maximum sum of stage game payoff is 2, $U \in [0, 1]$ 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 $(1 - \delta) \log N \to \infty$, even with within-match cheap talk.[9] 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 $(1 - \delta) \log N \to 0$.

3 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 $U \to 0$.*

[9]We establish the former result in Section 3.2 and the latter (which is essentially the same as Proposition 3 of Kandori (1992)) in Section 4.

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

We prove Theorem 1 in Section 3.1, deferring some details to the appendix. We then present a partial converse—a folk theorem when $(1-\delta)N \to 0$—in Section 3.2.

3.1 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 property holds even when we consider strategy profiles $\sigma = (\sigma_i)_i$ 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 $\mu^t = (\mu_{i,\tau})_{i=1,\tau\leq t}^N$ denote the first t periods of the match realization, let $h_i^t = \left(a_{i,\tau}, a_{\mu_{i,\tau},\tau}\right)_{\tau=1}^{t-1}$ 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 $\left(h_i^t, \mu^t\right)$ to a mixed action. That is, $\sigma_i\left(h_i^t, \mu^t\right)$ is the (possibly mixed) action taken by player i in period t at history $\left(h_i^t, \mu^t\right)$. (Note that μ^t includes the identity of i's period t partner.) Fix such a strategy profile $\sigma = (\sigma_i)_i$.

Let 0_i (resp., 1_i) denote the event that player i is rational (resp., bad). Finally, for $x_i \in \{0_i, 1_i\}$ and $x_j \in \{0_j, 1_j\}$, let $\Pr\left(h_i^t, h_j^t | x_i, x_j, \mu^t\right)$ 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 outcome distribution 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*). The condition that player i prefers σ_i to *Always Defect* is thus precisely the "incentive compatibility" condition that i's expected payoff when rational exceeds her expected payoff when bad: this condition appears in the appendix as equation (3). The first step of the proof of Theorem 1 puts this incentive compatibility constraint in a more convenient form and averages over players $i \in I$. (Proofs of lemmas are deferred to the appendix.)

Lemma 1 *If each player's expected payoff is higher when she is rational than when she is bad (that is, if equation (3) in the appendix holds for all $i \in I$) then*

$$\begin{aligned} & (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_i \frac{1}{N}\sum_{h_i^t}\Pr\left(h_i^t|0_i,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right)\min\{G,L\} \\ \leq \quad & (1-\varepsilon)(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{i,j\neq i}\frac{1}{N(N-1)}\sum_{h_j^t}\left(\begin{array}{c}\Pr\left(h_j^t|0_i,0_j,\mu^t\right) \\ -\Pr\left(h_j^t|1_i,0_j,\mu^t\right)\end{array}\right)_+ (1+G)\,. \end{aligned} \tag{1}$$

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 benefit from averted punishment goes to 0 if $(1-\delta)N \to \infty$. Roughly speaking, this amounts to showing that the (expected, discounted, average) impact 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* 0*.*

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

$$U \leq 2(1-\varepsilon)(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_i \frac{1}{N}\sum_{h_i^t}\Pr\left(h_i^t|0_i,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right),$$

as the population payoff would equal the right-hand side of this inequality if total within-match payoffs increased by 2 whenever a player takes C rather than D, which is an upper bound by our assumption that $G < 1+L$. Since Lemmas 1 and 2 imply that the right-hand side of this inequality goes to 0 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 impact of N players' types on a k-dimensional binary random variable is non-negligible only if $k/N \nrightarrow 0$.[10]

[10] We thank Omer Tamuz for providing a proof of this lemma.

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*

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

Let us provide some intuition and a proof sketch for Lemma 3. The $k = 1$ case is the relatively familiar result that the average impact 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 impact" of $\sum_{i=1}^{N} \sum_{s \in \{0,1\}} (\Pr(s|0_i) - \Pr(s|1_i))_+ \approx \sum_{i=1}^{N} 1/\sqrt{N} = \sqrt{N}$.[11] The lemma asserts that in general the total impact 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 in each of them: each voter is then pivotal with probability approximately $1/\sqrt{N/k}$, and summing over voters gives a total impact 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 impact of a vote.

The proof of Lemma 3 proceeds as follows. First, we use Pinsker's inequality and some manipulations to show that the impact of i's vote, $\sum_{s \in \{0,1\}^k} (\Pr(s|0_i) - \Pr(s|1_i))_+$, is at most $\sqrt{I_i/\underline{\varepsilon}}$, where I_i denotes the mutual information between S and X_i. Elementary properties of mutual information, together with independence of the X_i's, imply that $\sum_i I_i$ is at most the entropy of S, which is at most k since S is a k-dimensional binary random variable. Therefore, the sum of the squared impacts is at most $\sum_i I_i/\underline{\varepsilon} \leq k/\underline{\varepsilon}$, and hence the sum of the impacts is at most $\sqrt{kN/\underline{\varepsilon}}$ by the $\ell_1 - \ell_2$ norm inequality.[12]

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 $\delta \in (0,1)$.[13]

[11] See for example Lemma A of Fudenberg, Levine, and Pesendorfer (1998) and Theorem 2 of Al-Najjar and Smorodinsky (2000).

[12] One could also try to prove Lemma 3 by induction on k. This approach easily gives a bound of order $k\sqrt{N}$. Such an approach is used in recent work 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.

[13] One way to prove this result is to show by differentiation that $\sum_{t=1}^{\infty} \delta^t t^{-\alpha}$ is log-convex in α, and then conclude that $\sum_{t=1}^{\infty} \delta^t t^{-1/2} \leq \sqrt{\left(\sum_{t=1}^{\infty} \delta^t\right)\left(\sum_{t=1}^{\infty} \delta^t t^{-1}\right)} = \sqrt{\left(\frac{\delta}{1-\delta}\right)\left(\frac{\delta}{(1-\delta)^2}\right)} = \frac{\delta}{(1-\delta)^{3/2}}$.

Now, since h_j^t 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

$$
\begin{aligned}
& (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{i,j\neq i}\frac{1}{N(N-1)}\sum_{h_j^t}\left(\Pr\left(h_j^t|0_i,0_j,\mu^t\right)-\Pr\left(h_j^t|1_i,0_j,\mu^t\right)\right)_+ \\
= \; & (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j}\frac{1}{N(N-1)}\sum_{i\neq j}\sum_{h_j^t}\left(\Pr\left(h_j^t|0_i,0_j,\mu^t\right)-\Pr\left(h_j^t|1_i,0_j,\mu^t\right)\right)_+ \\
\leq \; & (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\frac{N}{N(N-1)}\sqrt{\frac{2(t-1)(N-1)}{\underline{\varepsilon}}} \qquad \text{(by (2))} \\
= \; & \sqrt{\frac{2}{\underline{\varepsilon}}}\frac{1}{\sqrt{N-1}}(1-\delta)\sum_{t=1}^{\infty}\delta^{t}\sqrt{t} \\
\leq \; & \sqrt{\frac{2}{\underline{\varepsilon}}}\frac{1}{\sqrt{N-1}}\frac{1}{\sqrt{1-\delta}}.
\end{aligned}
$$

Hence, if $(1-\delta)N \to \infty$ then the right-hand side of (1) goes to 0.

3.2 A Converse: Cooperation with Frequent Bilateral Interactions

We now provide a partial converse to Theorem 1, which shows that the presence of bad types does not hinder cooperation when bilateral interactions are frequent. The idea is to simply 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 $E \subset \mathbb{R}^N$ denote the set of rational-player sequential equilibrium payoff vectors: that is, $(v_i)_i \in E$ if there exists a sequential equilibrium where each player i's expected payoff when rational equals v_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 $l > \bar{l}$, there exists a payoff vector $v \in E_l$ satisfying*

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

Proof. See Appendix A.3. ■

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 by a positive payoff in her relationship with player 3. In principle, such payoffs vectors can be supported by having players occasionally communicate implicitly via actions.[14] We do not pursue such a result here.

One could also consider the case where $\delta \to 1$ and $N \to \infty$ at the same rate, so that the bilateral interaction frequency stays constant. Here partial cooperation is possible: the maximum equilibrium value of U can exceed 0, but a folk theorem typically does not hold (as follows from applying the proof of Theorem 1 in the case where $(1-\delta) N$ is constant but large).

4 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 all players (both rational types and bad types) communicate strategically to maximize their expected utility, and we assume that the set of possible messages is finite but can be taken arbitrarily large relative to the population size N.

Remark 1 *This setup relies on the assumption that bad types have the same utility function as rational types. An interpretation is that bad types differ from rational types only in that they are technologically unable to cooperate, or can cooperate only at a prohibitive cost. The import of this assumption is that, like rational players, bad players use communication to try to get their partners to cooperate with them—it is not actually important that rational and bad types have exactly the same preferences (for example, the same discount factor). While this assumption closes the model in a simple manner, there are other reasonable ways of specifying bad types' communication strategies. For instance, if we specified that bad types always send the same message, or if we could otherwise freely specify bad types' communication strategies, this would only make our results easier to prove. On the other hand, our results would be harder to prove if we required that communication*

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

among rational types was robust to any specification of bad types' communication strategies, or if we assumed that bad types' objective is to minimize rational types' payoffs, and we do not pursue these more difficult results here. Thus, broadly speaking, the interpretation of bad types in this section is that they are uncooperative and selfish, but not malevolent; and we study such types because selfishness is both more tractable and (we believe) more realistic than malevolence.

We prove two versions of our result. First, we show that for any $\eta > 0$ there exists an η-Nash equilibrium in which rational types always cooperate with each other on path.

Theorem 2 *Fix a sequence* $(N, \delta)_l$ *satisfying* $\lim_{l \to \infty} (1 - \delta_l) \log N_l = 0$*. With cheap talk, for every* $\eta > 0$ *there exists* $\bar{l} > 0$ *such that, for each* $l \geq \bar{l}$*, there exists an* η*-Nash equilibrium in which rational players always cooperate with each other along the equilibrium path of play.*

This result is relatively simple, and its proof (in Appendix A.4) involves strategies that seem quite realistic. Each player keeps track of a "blacklist" of players whom she believes have ever previously played D against a rational opponent. Every period, all players communicate their blacklists to their partners 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 $(1 - \delta) \log N \to 0$, first note 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. Formally, we rely on the following lemma.

Lemma 4 *Consider uniform random matching among* N *agents. Suppose that agent 1 knows a "rumor" in period 1, and in every period all agents other than agent 2 who know the rumor share it with their partners; agent 2, meanwhile, never shares the rumor. Then, letting* $T = Z \log_2 N$*, there exist constants* $c > 0$ *and* $\bar{Z} > 0$ *(independent of* N*) such that, for all* $Z > \bar{Z}$*, the probability that everyone knows the rumor at time* T *is at least* $1 - \exp(-cZ)$*.*

Moreover, suppose there are N *different rumors, where initially agent* i *knows rumor* i*, and agent* $i + 1$ *shares all rumors except rumor* i*. Then, letting* $T = Z \log_2 N$*, there exist (different) constants* $c > 0$ *and* $\bar{Z} > 0$ *(again independent of* N*) such that, for all* $Z > \bar{Z}$*, the probability that everyone knows* all N *rumors at time* T *is at least* $1 - \exp(-cZ)$*.*

Frieze and Grimmett (1985) prove a similar result in the related model where, every period, each informed player shares the rumor with a receiver selected uniformly at random from the population—rather than having players meet in pairs, as in the current model.[15] Since pairwise matching yields a different stochastic process for the number of informed players, we provide a complete proof in Online Appendix B.1. The basic idea, though, is the same as in Frieze and Grimmett. 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, *un*informed players are unlikely to meet each other, so the number of uninformed players shrinks exponentially.

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 $(1-\delta)\log N \approx 0$, deviating to D against a rational opponent is unprofitable.

However, other deviations from this strategy profile may be (slightly) profitable, which is why it is only an η-Nash equilibrium. First, a standard problem is that a player who punishes a deviant rational opponent gets blacklisted herself, so (off path) players do not have incentives to punish deviators. This problem, though, is easily addressed by modifying the criterion for getting on the blacklist. One simple fix would be specifying that a player gets blacklisted only if she plays D against an opponent who simultaneously plays C against her.

A far more serious problem arises in the low-probability event that a player learns that the fraction of rational players in the population is actually much smaller than $1-\varepsilon$. In the extreme, suppose player 1 witnesses (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. Then, when player 1 meets player 2, if $(1-\delta)N \approx \infty$ she should play D against him even if he is not on her blacklist—this conclusion follows 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 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

[15] Frieze and Grimmett also do not consider the possibility that a single agent refuses to spread the rumor. While we need to take this feature into account (since we cannot rely on a deviant player to "self-incriminate"), it has little effect on the proof of Lemma 4.

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, they can be induced to communicate this information honestly by making their own continuation payoffs independent of whether she player is rewarded or punished.[16] This proof approach not only let us construct an exact sequential equilibrium, it also lets us also support a wider range of payoffs. But, not surprisingly, the strategies used in the proof (in Online Appendix 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 $\theta^* \subset I$. For each θ^*, let $F(\theta^*)$ denote the set of feasible payoff profiles where players outside θ^* always play D. 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 $\hat{u}_i(\mathbf{a}) = \frac{1}{N-1} \sum_j u_i(\mathbf{a}_i(j), \mathbf{a}_j(i))$. We then define $F(\theta^*) = \text{co}(\{\hat{u}(\mathbf{a})\}_{\mathbf{a} \in \mathbf{A}(\theta^*)}) \subset \mathbb{R}^N$, where $\mathbf{A}(\theta^*) = \{\mathbf{a} : \mathbf{a}_j(k) = D \ \forall j \notin \theta^*, k \neq j\}$. Let $F^*(\theta^*) = F(\theta^*) \cap \mathbb{R}_+^N$ denote the set of feasible and individually rational payoffs. Note that $F^*(\theta^*)$ implicitly depends on N, but not on δ.

Now fix a sequence $(N, \delta)_l$. For any $\alpha \in (0, 1-\varepsilon)$ and any $\eta \in (0, 1)$, we define $F^{\alpha,\eta} \subset \mathbb{R}_+^N$ as the set of payoff profiles $v \in \mathbb{R}_+^N$ such that there exists $\boldsymbol{v} \in \mathbb{R}_+^{N|\Theta|}$ with $\Theta = 2^I$ satisfying the following three conditions:

1. Letting $v^{\theta^*} \in \mathbb{R}_+^N$ denote the θ^*-component of $\boldsymbol{v}$, we have $v = \sum_{\theta^*} \binom{N}{|\theta^*|} (1-\varepsilon)^{|\theta^*|} \varepsilon^{N-|\theta^*|} v^{\theta^*}$.
2. For each θ^* satisfying $|\theta^*| \geq \alpha N$, we have $B^\eta(v^{\theta^*}) \subset F^*(\theta^*)$, where B^η denotes the ball of radius η. In contrast, for each θ^* satisfying $|\theta^*| < \alpha N$, we have $v^{\theta^*} = 0$.
3. For each $i \in I$ and each $\theta^*, \theta^{*\prime}$ satisfying (i) $|\theta^*|, |\theta^{*\prime}| \geq \alpha N$, (ii) $\theta^* \ni i$, and (iii) $\theta^{*\prime} \not\ni i$, we have $v_i^{\theta^*} - v_i^{\theta^{*\prime}} \geq \eta$.

Intuitively, $F^{\alpha,\eta}$ is the set of feasible and strictly individually rational expected payoffs such that no cooperation occurs when $|\theta^*| < \alpha N$ and each player's expected payoff is strictly greater when she is rational than when she is bad, where all strict constraints hold with η slack.[17] Note

[16] This approach to incentivizing communiation was introduced by Compte (1998) and Kandori and Matsushima (1998) in the context of repeated games with public monitoring.

[17] Strict individual rationality implies that for each θ^* all players receive strictly positive expected payoffs, including bad players. We could instead require bad players to receive payoff 0. In this case, the definition of $F(\theta^*)$ can be

that $F^{\alpha,\eta}$ implicitly depends on N (but not δ). In addition, if $\alpha = 0$ then $F^{\alpha,\eta} \supset F^{\eta}$, where F^{η} is the payoff set obtained in Proposition 1. Together with continuity in α, this inclusion implies that $F^{\alpha,\eta}$ is non-empty for all sufficiently small α and all $\eta \in (0,1)$. Let E^* denote the set of sequential equilibrium payoff profiles, which implicitly depends on N and δ.[18]

Theorem 3 *Fix a sequence $(N,\delta)_l$ satisfying $\lim_{l\to\infty}(1-\delta_l)\log N_l = 0$, and fix any $\alpha \in (0, 1-\varepsilon)$ and $\eta \in (0,1)$. With cheap talk, $F^{\alpha,\eta} \subseteq E^*$ for all sufficiently large l.*

Proof. See Appendix B.2. ■

The $(1-\delta)\log N \to 0$ sufficient condition in Theorem 3 is nearly the best possible: the maximum number of players who could possibly learn about a deviation by player i within t periods is 2^t. Thus, for each $\eta \in (0,1)$, the "cost" to player i from deviating is at most

$$\begin{aligned}\sum_{t=1}^{\infty}\delta^t \min\left\{\frac{2^t}{N},1\right\}(1+G+L) &\leq \left(\sum_{t=1}^{\eta\log N}\frac{2^t}{N}+\delta^{\eta\log N}\right)(1+G+L)\\ &\leq \left(\frac{\eta\log N\times N^{\eta}}{N}+\exp\left(-\eta\left(1-\delta\right)\log N\right)\right)(1+G+L),\end{aligned}$$

which goes to 0 whenever $(1-\delta)\log N \to \infty$. Thus, if $(1-\delta_l)\log N_l \to \infty$ then for sufficiently large l the unique Nash equilibrium is *Always Defect*.[19]

Before discussing the proof of Theorem 3, we clarify that the divergent conclusions of Theorems 1 and 3 depend on the assumption that the set of available cheap talk messages is arbitrarily large. Suppose that the set of possible messages each player can send to her partner is required to be a finite set M with cardinality $|M| = \kappa$, so the game is now parameterized by (N,δ,κ).

Theorem 1' *With cheap talk from message sets of cardinality κ, for any sequence of parameters (N,δ,κ) satisfying $\sqrt{(1-\delta)N/\log\kappa} \to \infty$ and any corresponding sequence of Nash equilibrium population payoffs (U), we have $U \to 0$.*

Thus, for cheap talk to support cooperation when bilateral interactions are infrequent (more precisely, when $(1-\delta)^{\beta}N \to \infty$ for any $\beta > 1$), the cardinality of the message set must increase exponentially in N—that is, the message set must be rich enough for players to be able to convey

modified by assuming that rational players always take D when matched with bad players. The proof of Theorem 3 then goes through as written, except that equation (19) in Online Appendix B.2 is imposed only for players $i \in \theta$.

[18] The set E^* differs from the set E defined in Section 3.2, in that the latter set contains vectors of players' expected payoffs conditional on being rational, while E^* contains vectors of unconditional expected payoffs.

[19] This observation is essentially the same as Proposition 3 of Kandori (1992).

binary summaries of the individual reputations of a positive fraction of the population. This observation does not exactly prove that conversation about individual players' reputation ("gossip") is a necessary feature of any cooperative equilibrium, but it does show that any cooperative equilibrium must rely on a communication system that is rich enough to accommodate gossip.[20]

Proof of Theorem 1'. The proof is the same as that of Theorem 1, except for two steps. First, instead of considering a deviation by the rational type of player i from her equilibrium strategy to *Always Defect*, we must consider the deviation to playing *Always Defect* together with sending cheap talk messages as she were the bad type of player i. The requirement that such a deviation is unprofitable yields equation (3) (in the appendix). Second, player j's history h_j^t may now be viewed as a vector of $2\left(1+\lceil\log_2\kappa\rceil\right)(t-1)$ binary random variables, rather than $2(t-1)$ binary random variables as in the model without cheap talk. Replacing $2(t-1)$ with $2\left(1+\lceil\log_2\kappa\rceil\right)(t-1)$ in the last step of the proof of Theorem 1 implies that a sufficient condition for $U\to 0$ is

$$\sqrt{\frac{\underline{\varepsilon}\,(1-\delta)\,N}{2\left(1+\lceil\log_2\kappa\rceil\right)}}\to\infty.$$

This conclusion holds whenever $\sqrt{(1-\delta)\,N/\log\kappa}\to\infty$. ■

4.1 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 were extended to community enforcement 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 much easier relative to the case where communication can be executed only through payoff-relevant actions.

Specifically, the proof shows that strategies of the following form give a sequential equilibrium:

[20] In reality, agents may sometimes implicitly communicate through the fine details of their actions rather than literally engaging in conversation. For example, Cramton and Schwartz (2000) report that bidders in the 1990s FCC spectrum auctions exchanged information via the trailing digits of their bids. Interpreting κ as the maximum number of actions (e.g., bids) taken along the equilibrium path, the logic of Theorem 1' applies equally to such situations.

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 $\theta \subset I$ 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* $\left(x_i^\theta\right)_{\theta \subset I} \in (\{G, B\})_{\theta \subset I}$ 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. Intuitively, even if at some point in the game player i comes to believe with probability 1 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 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. In particular, $x_i^\theta = G$ means that, if in the coming block the players reach consensus 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 also by player i "instructing" other players to cooperate with player $i+1$); similarly, $x_i^\theta = B$ means that, if consensus is reached 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 $i+1$). While it might seem unnecessary for the players to form a consensus 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.[21]

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

[21] 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'.

$x_i^\theta = G$ 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 sub-blocks. In the first portion of the block, players defect while communicating about who cooperated in period 1 so as to reach consensus about θ, as well as communicating their state profiles $\left(x_i^\theta\right)_{\theta \subset I}$.[22] Then, in each of K "main sub-blocks" (which together comprise the vast majority of the block, and hence determine the equilibrium payoffs), for many periods players take their prescribed pure actions (which depend deterministically 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 consensus state θ satisfies $|\theta| < \alpha N$ then all players are prescribed defection in the main sub-blocks.

Communication is always executed through a protocol that is intended to facilitate truthtelling. 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 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 i 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 a reduced continuation payoff at the beginning of the next block), this strategy ensures incentive compatibility at on-path histories.

[22] Technically, the first portion of the block is divided into three sub-blocks: one to reach consensus about θ, one to reach consensus about the state profiles, and one to ensure that, if communication in the first two sub-blocks 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 consensus), this fact becomes known to all players with high probability.

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 such players to communicate player j's history honestly.

Finally, at the end of the block, each player i learns the block history of player $i+1$ 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 $i+1$.

5 Discussion

5.1 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.

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 an arbitrary repeated game strategy. Theorem 1 extends immediately to this more general setting, regardless of the value of ε', by the same proof. 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.[23] 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 replace period 1 in the proof of Theorem 3.) But we have not verified this conjecture.

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

[23] 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'.

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.[24] In contrast, the proofs of Theorems 2 and 3 can easily accommodate correlated types.

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, in the sense of fixing N and then taking $\varepsilon \to 0$. If we instead first fix ε and then take $N \to \infty$, contagion strategies breaks down, but also our proof of Theorem 1 does not apply.[25] 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.

5.2 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 gossip is essential for supporting cooperation in large populations. We believe our model and results provide a more realistic perspective on large group cooperation than earlier analyses which 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 with 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 allowing communication devices that are

[24] Specifically, independence implies that $\sum_i I_i \leq k$, where I_i is the mutual information between S and X_i (see equation (5) in the appendix). As the X_i's become correlated, the upper bound of k increases towards N, which increases the upper bound in (2).

[25] 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.

more powerful than plain cheap talk.

Richer communication devices must also be introduced to support cooperation in extremely large (e.g., continuum) populations, as well as to model real-world informational institutions such as 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 (2020). 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?).[26]

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.

A Appendix: Omitted Proofs

The condition in the proof of Theorem 1 that player i's expected payoff is higher when she is rational than when she is bad is

$$\begin{aligned}
&\underbrace{(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\begin{pmatrix}(1-\varepsilon)\Pr\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)u\left(\sigma_i(h_i^t,\mu^t),\sigma_j(h_j^t,\mu^t)\right)\\+\varepsilon\Pr\left(h_i^t,h_j^t|0_i,1_j,\mu^t\right)u\left(\sigma_i(h_i^t,\mu^t),D\right)\end{pmatrix}}_{i\text{'s expected payoff when rational and playing }\sigma_i}\\
\geq\ &\underbrace{(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\begin{pmatrix}(1-\varepsilon)\Pr\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)u\left(D,\sigma_j(h_j^t,\mu^t)\right)\\+\varepsilon\Pr\left(h_i^t,h_j^t|1_i,1_j,\mu^t\right)u\left(D,D\right)\end{pmatrix}}_{i\text{'s expected payoff when rational and playing \textit{Always Defect} (which equals }i\text{'s expected payoff when bad)}},\qquad(3)
\end{aligned}$$

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

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

A.1 Proof of Lemma 1

Since $u(D,D)=0$ and $u(C,D)=-L$, (3) is equivalent to

$$\begin{aligned}
& (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\begin{pmatrix}\Pr\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)u\left(\sigma_i(h_i^t,\mu^t),\sigma_j(h_j^t,\mu^t)\right)\\ -\frac{\varepsilon}{1-\varepsilon}\Pr\left(h_i^t|0_i,1_j,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right)L\end{pmatrix}\\
\geq\ & (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\Pr\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)u\left(D,\sigma_j(h_j^t,\mu^t)\right).
\end{aligned}$$

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

$$\begin{aligned}
& (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\begin{pmatrix}\Pr\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)\begin{pmatrix}u\left(\sigma_i(h_i^t,\mu^t),\sigma_j(h_j^t,\mu^t)\right)\\ -u\left(D,\sigma_j(h_j^t,\mu^t)\right)\end{pmatrix}\\ -\frac{\varepsilon}{1-\varepsilon}\Pr\left(h_i^t|0_i,1_j,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right)L\end{pmatrix}\\
\geq\ & (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\Pr\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)-\Pr\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)\right)u\left(D,\sigma_j(h_j^t,\mu^t)\right).
\end{aligned}$$

Since $u(C,a)-u(D,a)\leq-\min\{G,L\}$ and $u(D,a)\in\{0,1+G\}$ for each $a\in\{C,D\}$, a weaker necessary condition is

$$\begin{aligned}
& (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\begin{pmatrix}\Pr\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right)\\ +\frac{\varepsilon}{1-\varepsilon}\Pr\left(h_i^t|0_i,1_j,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right)\end{pmatrix}\min\{G,L\}\\
\leq\ & (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_i^t,h_j^t}\left(\Pr\left(h_i^t,h_j^t|0_i,0_j,\mu^t\right)-\Pr\left(h_i^t,h_j^t|1_i,0_j,\mu^t\right)\right)_+(1+G),
\end{aligned}$$

or equivalently

$$\begin{aligned}
& (1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{h_i^t}\Pr\left(h_i^t|0_i,\mu^t\right)\Pr\left(\sigma_i(h_i^t,\mu^t)=C\right)\min\{G,L\}\\
\leq\ & (1-\varepsilon)(1-\delta)\sum_t \delta^{t-1}\sum_{\mu^t}\Pr\left(\mu^t\right)\sum_{j\neq i}\frac{1}{N-1}\sum_{h_j^t}\left(\Pr\left(h_j^t|0_i,0_j,\mu^t\right)-\Pr\left(h_j^t|1_i,0_j,\mu^t\right)\right)_+(1+G).
\end{aligned}$$

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

A.2 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

$$\begin{aligned} H(A) &= -\sum_{a\in\mathcal{A}} \Pr(A=a)\log_2(\Pr(A=a)), \\ H(A|B) &= -\sum_{a\in\mathcal{A}}\sum_{b\in\mathcal{B}} \Pr(A=a,B=b)\log_2(\Pr(A=a|B=b)), \\ I(A;B) &= \sum_{a\in\mathcal{A}}\sum_{b\in\mathcal{B}} \Pr(A=a,B=b)\log_2\left(\frac{\Pr(A=a,B=b)}{\Pr(A=a)\Pr(B=b)}\right) \\ &= H(A)-H(A|B)=H(B)-H(B|A)=I(B;A). \end{aligned}$$

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

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 $H(S)\le k$. Denote the "impact" of X_i on S by

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

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

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

Note that

$$\begin{aligned} M_i(S) &= \sum_{s\in\{0,1\}^k}\left(\Pr(S=s|X_i=0)-\Pr(S=s)-(\Pr(S=s|X_i=1)-\Pr(S=s))\right)_+ \\ &= \sum_{x\in\{0,1\}}\sum_{s\in\{0,1\}^k}\left(\Pr(S=s|X_i=x)-\Pr(S=s)\right)_+. \end{aligned}$$

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,

$$\begin{aligned}
\sum_{s\in\{0,1\}^k} \left(\Pr\left(S=s|X_i=0\right)-\Pr\left(S=s\right)\right)_+ &\leq \sqrt{\frac{1}{2}D_{KL}\left(P^0||P\right)} \qquad \text{and} \\
\sum_{s\in\{0,1\}^k} \left(\Pr\left(S=s|X_i=1\right)-\Pr\left(S=s\right)\right)_+ &\leq \sqrt{\frac{1}{2}D_{KL}\left(P^1||P\right)},
\end{aligned}$$

where $D_{KL}\left(\cdot||\cdot\right)$ denotes Kullback-Leibler divergence. Note that

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

and

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

Hence, we have

$$\begin{aligned}
M_i\left(S\right) &\leq \sqrt{\frac{1}{2}D_{KL}\left(P^0||P\right)}+\sqrt{\frac{1}{2}D_{KL}\left(P^1||P\right)} \\
&= \sqrt{\frac{1}{2\left(1-\varepsilon\right)}\sum_{s\in\{0,1\}^k} \Pr\left(S=s,X_i=0\right)\log_2\left(\frac{\Pr\left(S=s,X_i=0\right)}{\Pr\left(S=s\right)\Pr\left(X_i=0\right)}\right)} \\
&\quad +\sqrt{\frac{1}{2\varepsilon}\sum_{s\in\{0,1\}^k} \Pr\left(S=s,X_i=1\right)\log_2\left(\frac{\Pr\left(S=s,X_i=0\right)}{\Pr\left(S=s\right)\Pr\left(X_i=1\right)}\right)} \\
&\leq \sqrt{\frac{1}{2\underline{\varepsilon}}}\left(\begin{array}{c} \sqrt{\sum_{s\in\{0,1\}^k}\Pr\left(S=s,X_i=0\right)\log_2\left(\frac{\Pr(S=s,X_i=0)}{\Pr(S=s)\Pr(X_i=0)}\right)} \\ +\sqrt{\sum_{s\in\{0,1\}^k}\Pr\left(S=s,X_i=1\right)\log_2\left(\frac{\Pr(S=s,X_i=0)}{\Pr(S=s)\Pr(X_i=1)}\right)} \end{array}\right) \\
&\leq \sqrt{\frac{1}{\underline{\varepsilon}}}\sqrt{\sum_{x\in\{0,1\}}\sum_{s\in\{0,1\}^k} \Pr\left(S=s,X_i=x\right)\log_2\left(\frac{\Pr\left(S=s,X_i=x\right)}{\Pr\left(S=s\right)\Pr\left(X_i=x\right)}\right)} \\
&= \sqrt{\frac{I\left(S;X_i\right)}{\underline{\varepsilon}}}, \qquad (4)
\end{aligned}$$

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

We next show that

$$\sum_i I(S; X_i) \leq k. \tag{5}$$

To see this, first note that

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

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

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

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

$$\sum_i (H(X_i) - I(S; X_i)) \geq \sum_i H(X_i) - k \Leftrightarrow \sum_i I(S; X_i) \leq k.$$

Combining (4) and (5), we have

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

or

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

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

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

A.3 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 $\left(a_t^{i,j}\right)_{t\in\mathbb{N}}$.

Suppose each player i conditions her behavior against each player $j \neq i$ only on the history of outcomes in past (i,j) matches, and in particular follows $\left(a_t^{i,j}\right)_{t\in\mathbb{N}}$ 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 $i \neq j$, we have

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

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

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

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 $\left(a_t^{i,j}\right)_{t\in\mathbb{N}}$ until j deviates from this path, and then obtains payoff 0 forever. Suppose the first deviation by j from action path $\left(a_t^{i,j}\right)_{t\in\mathbb{N}}$ occurs in period t. Then i's payoff against j is at least $\left(1-\delta^t\right)u_{i,j}^{<t}+\delta^t(1-\delta)(-L)$, where $u_{i,j}^{<t}$ is i's average payoff from the first $t-1$ periods of action path $\left(a_t^{i,j}\right)_{t\in\mathbb{N}}$. Note that $u_{i,j}^{<t}$ satisfies

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

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

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

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

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

Moreover, i's payoff against j is non-positive, 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.

A.4 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" $I_{i,t}^D \subset I$. Let $I_{i,1}^D = \emptyset$ 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 i is rational, let S denote the set of bad players, and suppose that $|S| = n$. Fix any $T, Z \in \mathbb{N}$ 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 $1 - n\left(\frac{n-1}{N-1}\right)^T$. 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 $1 - n\left(\frac{N-n}{N-1}\right)^T - \exp(-cZ)$, starting from period $T + Z\log_2 N$ player i obtains payoff 1 when she meets a rational type and obtains payoff 0 when she meets a bad type, for an expected payoff of $\frac{N-1-n}{N-1}$. For the first $T + Z\log_2 N$ periods, and with probability at most $n\left(\frac{N-n}{N-1}\right)^T + \exp(-cZ)$ 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

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

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

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

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,

$$\begin{aligned} & \sum_n p_n \frac{N-1-n}{N-1} - \sum_n p_n \min_{T\in\mathbb{N}, Z>\bar{Z}} \left\{ \left(1-\delta^{T+Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^T + \exp(-cZ) \right\} (1+L) \\ \geq \ & \sum_n p_n \frac{N-1-n}{N-1} - \frac{2}{3}\eta. \end{aligned} \tag{6}$$

First, fix some $\hat{\alpha} \in (0, 1-\varepsilon)$, and fix some $Z > \bar{Z}$ such $\exp(-cZ) \leq \frac{\eta}{4(1+L)}$. 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 $\frac{1}{12}\frac{1}{1+N+\exp(-cZ)}\frac{\eta}{1+L}$. Since $\left(1-\delta^{T+Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^T + \exp(-cZ) \leq 1+N+\exp(-cZ)$ for each N, $n \leq N$, and $T \in \mathbb{N}$, to establish (6) 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, $\exp(-cZ)(1+L) \leq \frac{1}{4}\eta$. 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

$$\left(\left(1-\delta^{T+Z\log_2 N}\right) + n\left(\frac{n-1}{N-1}\right)^T \right)(1+L) \leq \frac{1}{3}\eta. \tag{7}$$

To establish (7), for a given value of l, let $T \in \mathbb{N}$ be the smallest integer such that $N(1-\hat{\alpha})^T \leq \frac{\eta}{4(1+L)}$. 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

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

Hence, to establish (7), it suffices to show that, for sufficiently large l, we have

$$\left(1-\delta^{T+Z\log_2 N}\right)(1+L) \leq \frac{1}{12}\eta. \tag{8}$$

Finally, we have

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

and $(1-\delta)\log_2 N \to 0$ as $l \to \infty$ by hypothesis. This establishes (8) (and hence (6)), as desired.

Bad type equilibrium payoff. Here we take the trivial bound that, when player i is bad, her equilibrium payoff is non-negative.

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 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 $(j_\tau, j_{\tau+1}, \ldots, j_{t-1})$ such that $j_\tau=\mu_{i,\tau}$, $j\in\{j_\tau,\ldots,j_{t-1}\}$, $i\notin\{j_\tau,\ldots,j_{t-1}\}$, and $j_{t'+1}=\mu_{j_{t'},t'+1}$ for each $t'\in\{\tau,\ldots,t-2\}$. 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 at most

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

First fixing Z such that $\exp(-cZ) \leq \frac{\eta}{3(1+G)}$ 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 (6), 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 0 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 non-positive, her payoff under any unilateral deviation is at most

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

As we have seen, this converges to 0 as $l \to \infty$. Hence, the equilibrium strategy is η-optimal.

References

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

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

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

[4] Ali, S. Nageeb, and David A. Miller. "Communication and Cooperation in Markets." *Working Paper* (2020).

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

[6] Awaya, Yu, and Vijay Krishna. "On Communication and Collusion." *American Economic Review* 106.2 (2016): 285-315.

[7] Awaya, Yu, and Vijay Krishna. "Communication and Cooperation in Repeated Games." *Theoretical Economics* 14.2 (2019): 513-553.

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

[9] Barron, Daniel, and Yingni Guo. "The Use and Misuse of Coordinated Punishments." *Working Paper* (2019).

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

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

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

[13] Clark, Daniel, Drew Fudenberg, and Alexander Wolitzky. "Record-Keeping and Cooperation in Large Societies." *Working Paper* (2020).

[14] Compte, Olivier. "Communication in Repeated Games with Imperfect Private Monitoring." *Econometrica* (1998): 597-626.

[15] Cramton, Peter, and Jesse A. Schwartz. "Collusive Bidding: Lessons from the FCC Spectrum Auctions." *Journal of Regulatory Economics* 17.3 (2000): 229-252.

[16] Deb, Joyee. "Cooperation and Community Responsibility," *Journal of Political Economy* 128.5 (2020): 1976-2009.

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

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

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

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

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

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

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

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

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

[26] Fudenberg, Drew, and Eric Maskin. "On the Dispensability of Public Randomization in Discounted Repeated Games." *Journal of Economic Theory* 53.2 (1991): 428-438.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[42] Sugaya, Takuo, and Yuichi Yamamoto. "Common Learning and Cooperation in Repeated Games." *Theoretical Economics* 15.3 (2020): 1175-1219.

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

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

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

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

B Online Appendix

B.1 Proof of Lemma 4

For each $i \neq j$, let $N_t^{-i}(j)$ denote the (random) number of players in the set $-i = I \backslash \{i\}$ who, by period t, have met a player in $-i$ who met a player in $-i$ who... met player j. We wish to show that there exists a constant $c > 0$ such that $\Pr\left(N_T^{-i}(j) = N - 1 \; \forall i, j\right) \geq 1 - \exp(-cZ)$. The idea of the proof is to show that, with high probability, $\min_{i,j} N_t^{-i}(j)$ grows exponentially in t until it reaches a constant fraction of N, and that subsequently $N - \min_{i,j} N_t^{-i}(j)$ shrinks exponentially.

We first show that $\min_{i,j} N_t^{-i}(j)$ grows exponentially until it reaches $\frac{2}{3}N$.

Lemma 5 *There exists $\bar{\gamma} \in \left(0, \frac{1}{2}\right]$ such that, for every N and $n \leq \frac{2}{3}N$,*

$$\begin{aligned}
&\Pr\left(\min_{i,j} N_{t+1}^{-i}(j) \leq (1+\bar{\gamma}) \min_{i,j} N_t^{-i}(j) | \min_{i,j} N_t^{-i}(j) = n\right) \\
\leq \;& N(N-1) \frac{e}{2\pi \bar{\gamma}^{\frac{1}{2}} (1-\bar{\gamma})^{\frac{1}{2}}} \left(\frac{1}{2}\right)^{\frac{\bar{\gamma} n}{2}}.
\end{aligned}$$

Proof. By monotonicity in the number of informed players and symmetry, it suffices to prove that, for each particular $i \neq j$,

$$\Pr\left(N_{t+1}^{-i}(j) \leq (1+\bar{\gamma}) N_t^{-i}(j) | N_t^{-i}(j) = n\right) \leq \frac{e}{2\pi \bar{\gamma}^{\frac{1}{2}} (1-\bar{\gamma})^{\frac{1}{2}}} \left(\frac{1}{2}\right)^{\frac{\bar{\gamma} n}{2}}.$$

Fixing $i \neq j$, and suppressing i and j in the notation, let I_t be the set of players who received player j's message through a path excluding i by period t: thus, $|I_t| = n$. Note that, for each number $n' \leq N - n$ with the same parity as n, $\Pr(N_{t+1} = n + n' | N_t = n)$ is at most

$$\begin{aligned}
&\underbrace{\binom{n}{n'}}_{\substack{\text{who in } I_t \text{ meets} \\ \text{players in } I \backslash (I_t \cup \{i\})}} \times \underbrace{\frac{n-1}{N-1}}_{\substack{\text{"first" player in } I_t \\ \text{meets someone in } I_t}} \times \underbrace{\frac{n-3}{N-3}}_{\substack{\text{"second" (remaining) player in } I_t \\ \text{meets some (remaining) player in } I_t}} \times \cdots \times \frac{n'+1}{N-n+n'+1} \\
= \;& \binom{n}{n'} \prod_{k=1}^{\frac{n-n'}{2}} \frac{n-2k+1}{N-2k+1}.
\end{aligned}$$

This expression is an upper bound, as we neglect the probability that the players in I_t who are selected to meet someone in $I \backslash (I_t \cup \{i\})$ actually do so. Similarly, for each n' with the opposite parity as n, $\Pr(N_{t+1} = n + n' | N_t = n)$ is at most

$$\begin{aligned}
&\underbrace{n}_{\text{who in } I_t \text{ meets } i} \times \underbrace{\frac{1}{N-1}}_{\text{prob. of meeting } i} \times \underbrace{\binom{n-1}{n'}}_{\substack{\text{who in } I_t \text{ meets} \\ \text{players in } I \backslash (I_t \cup \{i\})}} \times \underbrace{\frac{n-2}{N-3} \times \ldots \times \frac{n'+1}{N-n+n'}}_{\substack{\text{remaining players in } I_t \\ \text{match with each other}}} \\
= \;& \binom{n-1}{n'} \prod_{k=1}^{\frac{n-n'+1}{2}} \frac{n-2k+2}{N-2k+1} \leq \binom{n}{n'} \prod_{k=1}^{\frac{n-n'+1}{2}} \frac{n-2k+2}{N-2k+1}.
\end{aligned}$$

For any $\gamma \in (0, \frac{1}{2}]$, if $n' \leq \gamma n$, Stirling's formula gives

$$\binom{n}{n'} \leq \frac{en^{n+\frac{1}{2}}e^{-n}}{2\pi(\gamma n)^{\gamma n+\frac{1}{2}}e^{-\gamma n}((1-\gamma)n)^{(1-\gamma)n+\frac{1}{2}}e^{-(1-\gamma)n}} \leq \frac{e}{2\pi(\gamma)^{\gamma n+\frac{1}{2}}(1-\gamma)^{(1-\gamma)n+\frac{1}{2}}}.$$

We also have

$$\prod_{k=1}^{\frac{n-n'}{2}} \frac{n-2k+1}{N-2k+1} \leq \left(\frac{n-1}{N-1}\right)^{\frac{n-n'}{2}} \leq \left(\frac{n}{N-1}\right)^{\frac{(1-\gamma)n}{2}}, \text{ and}$$

$$\prod_{k=1}^{\frac{n-n'+1}{2}} \frac{n-2k+2}{N-2k+1} \leq \left(\frac{n}{N-1}\right)^{\frac{n-n'+1}{2}} \leq \left(\frac{n}{N-1}\right)^{\frac{(1-\gamma)n}{2}}.$$

Therefore, for any $\gamma \in (0, \frac{1}{2}]$ and $n' \leq \gamma n$, we have

$$\Pr\left(N_{t+1} = n+n' | N_t = n\right) \leq \frac{e}{2\pi(\gamma)^{\gamma n+\frac{1}{2}}(1-\gamma)^{(1-\gamma)n+\frac{1}{2}}}\left(\frac{n}{N-1}\right)^{\frac{(1-\gamma)n}{2}},$$

and hence

$$\begin{aligned}
\Pr\left(N_{t+1} \leq n+\gamma n | N_t = n\right) &\leq \frac{e(\gamma n+1)}{2\pi(\gamma)^{\gamma n+\frac{1}{2}}(1-\gamma)^{(1-\gamma)n+\frac{1}{2}}}\left(\frac{n}{N-1}\right)^{\frac{(1-\gamma)n}{2}} \\
&= \frac{e}{2\pi\gamma^{\frac{1}{2}}(1-\gamma)^{\frac{1}{2}}}\left(\frac{(\gamma n+1)^{\frac{2}{\gamma n}}}{\gamma^2(1-\gamma)^{2\frac{1-\gamma}{\gamma}}}\left(\frac{n}{N-1}\right)^{\frac{1-\gamma}{\gamma}}\right)^{\frac{\gamma n}{2}} \\
&\leq \frac{e}{2\pi(\gamma)^{\frac{1}{2}}(1-\gamma)^{\frac{1}{2}}}\left(\frac{e^2}{\gamma^2(1-\gamma)^{2\frac{1-\gamma}{\gamma}}}\left(\frac{2}{3}\frac{N}{N-1}\right)^{\frac{1-\gamma}{\gamma}}\right)^{\frac{\gamma n}{2}}. \qquad (9)
\end{aligned}$$

Fix $\bar{\gamma} \in \left(0, \frac{1}{2}\right]$ such that

$$\frac{e^2}{\bar{\gamma}^2(1-\bar{\gamma})^{2\frac{1-\bar{\gamma}}{\bar{\gamma}}}}\left(\frac{8}{9}\right)^{\frac{1-\bar{\gamma}}{\bar{\gamma}}} < \frac{1}{2}.$$

Such a $\bar{\gamma}$ exists as the left-hand side of this inequality goes to 0 as $\bar{\gamma} \to 0$. Since $N \geq 4$, we have $\frac{2}{3}\frac{N}{N-1} \leq \frac{8}{9}$. Hence, substituting $\gamma = \bar{\gamma}$ in (9), we have, for every N and $n \leq \frac{2}{3}N$,

$$\Pr\left(N_{t+1} \leq n+\bar{\gamma} n | N_t = n\right) \leq \frac{e}{2\pi(\bar{\gamma})^{\frac{1}{2}}(1-\bar{\gamma})^{\frac{1}{2}}}\left(\frac{1}{2}\right)^{\frac{\bar{\gamma} n}{2}},$$

as desired. ■

Fix $\bar{\gamma}$ satisfying the conditions of Lemma 5. Let $n^*(N)$ satisfy

$$N(N-1)\frac{e}{2\pi\bar{\gamma}^{\frac{1}{2}}(1-\bar{\gamma})^{\frac{1}{2}}}\left(\frac{1}{2}\right)^{\frac{\bar{\gamma} n^*(N)}{2}} = \frac{1}{4}.$$

Note that

$$n^*(N) = \hat{c}\left(\log_2 N + \log_2 (N-1)\right),$$

where $\hat{c} > 0$ is a constant independent of N. The following lemma is an immediate consequence of Lemma 5.

Lemma 6 *For every n satisfying $n^*(N) \leq n \leq \frac{2}{3}N$,*

$$\Pr\left(\min_{i,j} N_{t+1}^{-i}(j) \leq (1+\bar{\gamma}) \min_{i,j} N_t^{-i}(j) \,|\, \min_{i,j} N_t^{-i}(j) = n\right) \leq \frac{1}{4}.$$

We now consider the case where $n \leq n^*(N)$, considering first the subcase where $n \geq 12$.

Lemma 7 *There exists $\bar{N}_1$ such that, for every $N \geq \bar{N}_1$ and n satisfying $12 \leq n \leq n^*(N)$,*

$$\Pr\left(\min_{i,j} N_{t+1}^{-i}(j) \leq \frac{3}{2} \min_{i,j} N_t^{-i}(j) \,|\, \min_{i,j} N_t^{-i}(j) = n\right) \leq \frac{1}{4}.$$

Proof. Taking $\gamma = \frac{1}{2}$ in (9), we have

$$\begin{aligned} & \Pr\left(\min_{i,j} N_{t+1}^{-i}(j) \leq \frac{3}{2} \min_{i,j} N_{t+1}^{-i} | \min_{i,j} N_t^{-i}(j) = n\right) \\ \leq \quad & N(N-1)\frac{e}{2\pi\gamma^{\frac{1}{2}}(1-\gamma)^{\frac{1}{2}}}\left(\frac{e^2}{\gamma^2(1-\gamma)^{2\frac{1-\gamma}{\gamma}}}\left(\frac{n}{N-1}\right)^{\frac{1-\gamma}{\gamma}}\right)^{\frac{\gamma n}{2}} \\ = \quad & N(N-1)\frac{e}{2\pi\left(\frac{1}{2}\right)^{\frac{1}{2}}\left(\frac{1}{2}\right)^{\frac{1}{2}}}\left(\frac{e^2}{\left(\frac{1}{2}\right)^2\left(\frac{1}{2}\right)^2}\frac{n}{N-1}\right)^{\frac{n}{4}}. \end{aligned}$$

Since $12 \leq n \leq \hat{c}\left(\log_2 N + \log_2 (N-1)\right)$, this is at most

$$N(N-1)\frac{e}{\pi}\left(16e^2\frac{\hat{c}\left(\log_2 N + \log_2 (N-1)\right)}{N-1}\right)^3,$$

which is less than $\frac{1}{4}$ for sufficiently large N. ■

The next lemma addresses the subcase with fewer than 12 informed players.

Lemma 8 *There exists $\bar{N}_2$ such that, for every $N \geq \bar{N}_2$,*

$$\Pr\left(\min_{i,j} N_{t+6}^{-i}(j) \leq 12 | \min_{i,j} N_t^{-i}(j) \geq 1\right) \leq \frac{1}{4}.$$

Proof. Fix $i \neq j$, and suppose $N_{t+6}^{-i}(j) \leq 12$. Since $\min_{i,j} N_t^{-i}(j) \geq 1$ and $12 < 2^4$, this is possible only if $N_{t'+1}^{-i}(j) = 2N_{t'}^{-i}(j)$ for at most 3 out of the 6 periods $t' \in \{t+1, \ldots, t+6\}$. That is, in at least 3 out of these 6 periods, some player in $I_{t'}^{-i}(j)$ must meet someone in $I_{t'}^{-i}(j) \cup \{i\}$. Since by hypothesis $N_{t'}^{-i}(j) \leq 12$ for each such period t', the probability of this event is at most $\begin{pmatrix} 6 \\ 3 \end{pmatrix} \times 12 \times \left(\frac{12}{N-1}\right)^3$. Hence,

$$\Pr\left(\min_{i,j} N_{t+6}^{-i}(j) \leq 12 | \min_{i,j} N_t^{-i}(j) \geq 1\right) \leq N(N-1)\frac{20 \times 12^4}{(N-1)^3},$$

which is less than $\frac{1}{4}$ for sufficiently large N. ■

In total, since $\bar{\gamma} \leq \frac{1}{2}$, we have the following lemma:

Lemma 9 *For every $N \geq \max\left\{\bar{N}_1, \bar{N}_2\right\}$,*

1. *For any $\left(N_t^{-i}(j)\right)_{i,j}$ such that $\min_{i,j} N_t^{-i}(j) \geq 1$, we have*

$$\Pr\left(\min_{i,j} N_{t+6}^{-i}(j) \leq 12 | \left(N_t^{-i}(j)\right)_{i,j}\right) \leq \frac{1}{4}.$$

2. *For any $\left(N_t^{-i}(j)\right)_{i,j}$ such that $\min_{i,j} N_t^{-i}(j) = n$ satisfies $12 \leq n \leq \frac{2}{3}N$, we have*

$$\Pr\left(\min_{i,j} N_{t+1}^{-i}(j) \leq (1+\bar{\gamma}) \min_{i,j} N_t^{-i} | \left(N_t^{-i}(j)\right)_{i,j}\right) \leq \frac{1}{4}.$$

We now provide a symmetric bound for the case where $N_t^{-i}(j)$ is "large" for each $i \neq j$. Let $M_t^{-i}(j) = N - 1 - N_t^{-i}(j)$ be the number of players $-i$ who have not yet received player j's message through a path excluding i; and let $J_t^{-i}(j)$ be the set of such players.

Lemma 10 *There exists $\bar{N}_3$ such that, for each $N \geq \bar{N}_3$,*

1. *For any $\left(M_t^{-i}(j)\right)_{i,j}$ such that $\max_{i,j} M_t^{-i}(j) \leq 12$, we have*

$$\Pr\left(\max_{i,j} M_{t+6}^{-i}(j) > 0 | \left(M_t^{-i}(j)\right)_{i,j}\right) \leq \frac{1}{4}.$$

2. *For any $\left(M_t^{-i}(j)\right)_{i,j}$ such that $\max_{i,j} M_t^{-i}(j) = n$ satisfies $12 \leq n \leq \frac{1}{3}N$, we have*

$$\Pr\left(\max_{i,j} M_{t+1}^{-i}(j) \geq (1-\bar{\gamma}) \max_{i,j} M_t^{-i}(j) \,|\, \max_{i,j} M_t^{-i}(j) = n\right) \leq \frac{1}{4}.$$

Proof. Lemmas 5–9 provide an upper bound for the probability that fraction $\bar{\gamma}$ of players in $I_t^{-i}(j)$ do not meet players outside of $I_t^{-i}(j) \cup \{i\}$. The current lemma provides an upper bound for the probability that fraction $\bar{\gamma}$ of players in $J_t^{-i}(j)$ do not meet players outside of $J_t^{-i}(j) \cup \{i\}$. The argument is symmetric. ■

We now combine Lemmas 9 and 10 to prove Lemma 4. We first assume $N \geq \max\{\bar{N}_1, \bar{N}_2, \bar{N}_3\}$.

We have the following properties. First, if $\min_{i,j} N_t^{-i}(j) < 12$, then $\min_{i,j} N_{t+6}^{-i}(j) \geq 12$ with probability at least $\frac{3}{4}$. Second, if $12 \leq \min_{i,j} N_t^{-i}(j) \leq \frac{2}{3}N$, then $\min_{i,j} N_{t+1}^{-i}(j) \geq (1+\bar{\gamma}) \min_{i,j} N_t^{-i}(j)$ with probability at least $\frac{3}{4}$. (And note that $\log_{(1+\bar{\gamma})} \frac{2}{3}N$ "increases" by a factor of $(1+\bar{\gamma})$ suffice to raise $\min_{i,j} N_t^{-i}(j)$ to $\frac{2}{3}N$.) Third, if $\frac{2}{3}N \leq \min_{i,j} N_t^{-i}(j) \leq N - 13$—or equivalently $12 \leq \max_{i,j} M_t^{-i}(j) \leq \frac{1}{3}N$—then $\max_{i,j} M_{t+1}^{-i}(j) \leq (1-\bar{\gamma}) \max_{i,j} M_t^{-i}(j)$ with probability at least $\frac{3}{4}$. (Note that $\log_{(1-\bar{\gamma})} 3\frac{1}{N}$ "decreases" suffice to reduce $\max_{i,j} M_t^{-i}(j)$ to 12.) Finally, if $\max_{i,j} M_t^{-i}(j) \leq 12$, then $\min_{i,j} N_{t+6}^{-i}(j) = N - 1$ (equivalently $\max_{i,j} M_t^{-i}(j) = 0$) with probability at least $\frac{3}{4}$.

Combining these properties, we see that $\Pr\left(\min_{i,j} N_T^{-i}(j) = N - 1\right)$ is lower-bounded by the probability that, out of $T/6$ Bernoulli random variables with parameter $\frac{3}{4}$, the realizations of at least $2 + \log_{(1+\bar{\gamma})} \frac{2}{3}N + \log_{(1-\bar{\gamma})} 3\frac{1}{N}$ of them equal 1. By Hoeffding's inequality, this probability is at least

$$1 - \exp\left(-2\left(\frac{3}{4} - \frac{2 + \log_{(1+\bar{\gamma})} \frac{2}{3}N + \log_{(1-\bar{\gamma})} 3\frac{1}{N}}{\frac{T}{6}}\right)^2 \frac{T}{6}\right).$$

If $T = Z \log_2 N$, then

$$\begin{aligned} \frac{2 + \log_{(1+\bar{\gamma})} \frac{2}{3}N + \log_{(1-\bar{\gamma})} 3\frac{1}{N}}{\frac{Z \log_2 N}{6}} &< \frac{2 + (\log_2 N)\left(\frac{1}{\log_2(1+\bar{\gamma})} - \frac{1}{\log_2(1-\bar{\gamma})}\right)}{\frac{Z \log_2 N}{6}} \\ &< \frac{6}{Z}\left(2 + \frac{1}{\log_2(1+\bar{\gamma})} - \frac{1}{\log_2(1-\bar{\gamma})}\right). \end{aligned}$$

Hence, there exists $\bar{Z}_1 > 0$ such that if $Z > \bar{Z}_1$ then, for all $N \geq \max\{\bar{N}_1, \bar{N}_2, \bar{N}_3\}$, we have

$$\frac{2 + \log_{(1+\bar{\gamma})} \frac{2}{3}N + \log_{(1-\bar{\gamma})} 3\frac{1}{N}}{\frac{Z \log_2 N}{6}} < \frac{1}{4},$$

and hence

$$\Pr\left(\min_{i,j} N_T^{-i}(j) = N - 1\right) \geq 1 - \exp\left(-2\left(\frac{1}{2}\right)^2 \frac{Z \log_2 N}{6}\right) \geq 1 - \exp\left(-\frac{1}{12}Z\right).$$

Finally, for the case $N < \max\{\bar{N}_1, \bar{N}_2, \bar{N}_3\}$, Hoeffding's inequality implies

$$\Pr\left(\min_{i,j} N_T^{-i}(j) = N - 1\right) \geq 1 - N(N-1)\exp\left(-2\left(\frac{1}{N-1}\right)^2 T\right).$$

Hence, there exist $c_1 > 0$ and $\bar{T} > 0$ such that, for all $N < \max\{\bar{N}_1, \bar{N}_2, \bar{N}_3\}$ and $T > \bar{T}$, we have

$$\Pr\left(\min_{i,j} N_T^{-i}(j) = N - 1\right) \geq 1 - \exp(-c_1 T).$$

Taking $c = \min\left\{\frac{1}{12}, c_1\right\}$ and $\bar{Z} = \max\left\{\bar{Z}_1, \bar{T}\right\}$ completes the proof.

B.2 Proof of Theorem 3

We first prove the following theorem:

Theorem 4 *Fix a sequence $(N, \delta)_l$ and fix any $\alpha \in (0, 1-\varepsilon)$ and $\eta \in (0,1)$. With cheap talk, if $\lim_l (1-\delta_l) \log N_l = 0$ then, for any $v \in F^{\alpha,\eta}$, we have $v \in E^*$ for all sufficiently large l.*

In Section B.2.10, we extend this result to show that $F^{\alpha,\eta} \subseteq E^*$ for all sufficiently large l.

To prove Theorem 4, we first describe a protocol for the community to circulate messages. This protocol has the feature that, with high probability, the number of periods it takes for everyone to learn the message is on the order of $\log N$; moreover, no single player can prevent the rest of the community from learning. We then use this protocol as a building block in the construction of a block belief-free equilibrium.

B.2.1 Protocol for Players to Circulate Message m

Suppose each player i wishes to disseminate a message m_i throughout the community, where each m_i is an element of some finite set M_i. We say that *players circulate message* $m = (m_i)_i$ *for* T *periods* if the players obey the following protocol for T periods:

In each period $t \in \{1, \dots, T\}$, all players take action D, while sending cheap-talk messages. Each player j has a "state"

$$\left(\zeta_{j,t}^{I,-i}, \zeta_{j,t}^{M,-i}\right)_{i \neq j} \subset \times_{i \neq j} \left(I \times \left(\times_{n \neq i} M_n\right)\right).$$

Intuitively, $\zeta_{j,t}^{I,-i}$ is the set of players k whose message player j has heard (directly or indirectly) via a path that excludes i, and $\zeta_{j,t}^{M,-i}|_k \subset M_k$ is the set of messages reported to j as having been sent by k via a path that excludes i.[27] Note that $\zeta_{j,t}^{M,-i}$ does not include player i's message since it is infeasible to share player i's message via a path that excludes i.

Formally, for each player j and $i \neq j$, $\left(\zeta_{j,1}^{I,-i}, \zeta_{j,1}^{M,-i}\right) = (\{j\}, (\emptyset, \dots, \emptyset, \{m_j\}, \emptyset, \dots, \emptyset))$. In each period t, given $\left(\zeta_{j,t}^{I,-i}, \zeta_{j,t}^{M,-i}\right)_{i \neq j}$, if player j meets player k, player j sends message $\left(\zeta_{j,t}^{I,-i}, \zeta_{j,t}^{M,-i}\right)_{i \notin \{j,k\}}$. That is, player j passes all of his information to player k, except for the "$-k$" information being circulated by players $-k$. Given his opponent's message $\left(\hat{\zeta}_{k,t}^{I,-i}, \hat{\zeta}_{k,t}^{M,-i}\right)_{i \notin \{j,k\}}$, for each $i \notin \{j,k\}$, player j's next-period state is given by $\zeta_{j,t+1}^{I,-i} = \zeta_{j,t}^{I,-i} \cup \hat{\zeta}_{k,t}^{I,-i}$ and $\zeta_{j,t+1}^{M,-i}|_n = \zeta_{j,t}^{M,-i}|_n \cup \hat{\zeta}_{k,t}^{M,-i}|_n$ for all $n \neq i$ (recall that $n \neq i$ since $\zeta_{j,t}^{M,-i} \subset \times_{n \neq i} M_n$). For each $i \in \{j,k\}$, let $\left(\zeta_{j,t+1}^{I,-i}, \zeta_{j,t+1}^{M,-i}\right) = \left(\zeta_{j,t}^{I,-i}, \zeta_{j,t}^{M,-i}\right)$. That is, for each player $n \neq i$, player j adds $\hat{\zeta}_{k,t}^{M,-i}|_n$ to the set of messages reported to him as having been sent by n (note that $k \neq i$ by definition: only player $k \neq i$ hears a message via a path that excludes i). (Throughout, we use hatted variables to denote messages.)

At the end of period T, for each $i \neq j$, if $\zeta_{j,T}^{I,-i} = -i$ and $\left|\zeta_{j,T}^{M,-i}|_n\right| = 1$ for each $n \neq i$, we say player j *infers* message $m_{-i}(j) \in \times_{k \neq i} M_k$, where $m_{-i}(j)|_n$ is equal to the unique element of $\zeta_{j,T}^{M,-i}|_n$, for each n. Otherwise, we say player j infers $m_{-i}(j) = \texttt{error}$.[28] We also say the match realization is *erroneous* if there exists disjoint players $i \neq j \neq k \neq i$ such that, by period T, player i has not met a player in $-k$ who met a player in $-k$ who... met player j. Otherwise, the match is *regular*.

Note that, if all players follow the protocol, then at the end of period T either the match is erroneous or $m_{-i}(j) = m_{-i}$ for all $i \neq j$. Moreover, if $T = Z \log_2 N$, by Lemma 4 the probability that the match is erroneous decreases exponentially in Z. We thus have

Lemma 11 *Let* $T = Z \log_2 N$. *There exist* $c > 0$ *and* $\bar{Z} > 0$ *such that, for all* $Z > \bar{Z}$ *and all* l, *we have*

$$\Pr\left(m_{-i}(j) = m_{-i} \ \forall i \neq j\right) \geq 1 - \exp\left(-cZ\right).$$

Note also that whether or not the event $\{m_{-i}(j) = m_{-i}\}$ obtains is independent of player i's behavior.

[27] For a vector $x \in X^{N-1}$ and $k \in \{1, \dots, N-1\}$, we denote the k^{th} coordinate of x by $x|_k$.

[28] Note that it is possible that $m_{-i}(j)|_n \neq m_{-i'}(j)|_n$ for some $i \neq i' \neq n \neq i$. Intuitively, $m_{-i}(j) = \texttt{error}$ means that j fails to infer a message through a chain of players excluding i, but not necessarily that the messages she infers through all chains are mutually consistent.

B.2.2 Period 1

The very first period of the repeated game plays a special role in our construction. We denote this period by 1^* rather than 1, to clarify that this is the first period of the infinitely repeated game, rather than the first period of a block. In period 1^*, every normal player is supposed to play C. Given the outcome of period 1^*, let θ denote the set of players who took $a_{i,1^*} = C$ as prescribed. (Note that $\theta \subset \theta^*$, as all committed players take D, and some rational players may also take D as the result of a deviation.) In our construction, only players in θ will cooperate with each other. The strategies we construct will take θ as a persistent "state variable," and we denote the set of possible states θ by $\Theta = 2^I$. Note that each player i's period-1^* history, $h_{i,1^*} = \left(\mu_{i,1^*}, a_{i,1^*}, a_{\mu_{i,1^*},1^*}\right)$, is directly informative of θ; for this reason, players' period-1^* histories will play a distinguished role in our construction.[29]

B.2.3 Block Belief-Free Structure

We now describe the general structure of our construction (following period 1^*) and present the corresponding equilibrium conditions.

Block Strategies. We view the repeated game from period 2 on as an infinite sequence of T^{**}-period blocks, where T^{**} is a number to be specified. At the beginning of every block, each player i selects a "strategy state" $x_i^\theta \in \{G, B\}$ for each $\theta \in \Theta$ from a full support probability distribution. Given the vector $\boldsymbol{x}_i = \left(x_i^\theta\right)_{\theta\in\Theta}$ and player i's period-1^* history $h_{i,1^*}$, player i plays a behavioral strategy $\sigma_i^*\left(\boldsymbol{x}_i, h_{i,1^*}\right)$ (her *block strategy*) within the block. That is, in every period $t = 1, \ldots, T^{**}$ of the block, $\sigma_i^*\left(\boldsymbol{x}_i, h_{i,1^*}\right)$ specifies a probability distribution over cheap talk messages and actions as a function of player i's *block history* $h_i^t = (\left(\mu_{i,\tau}, m_{i,\tau}, m_{\mu_{i,\tau},\tau}, a_{i,\tau}, a_{\mu_{i,\tau},\tau}\right)_{\tau=1}^{t-1}, \mu_{i,t})$. Denote player i's strategy set in the T^{**}-period game by Σ_i.

Players are prescribed to play C in period 1^* and subsequently use the same strategy in each block. Thus, a player's entire repeated-game strategy can be summarized by a single block strategy, together with a policy for selecting the strategy state $\boldsymbol{x}_i$ at the start of each block.

Continuation Payoffs. Conditional on the persistent state being equal to θ, player i's equilibrium continuation payoff at the end of a block is a function only of player $(i-1)$'s state x_{i-1}^θ and history $h_{i-1}^{T^{**}}$ in the previous block. (Adopt here the convention that player-names are mod N, so player $(1-1)$ is player N.) Denote this continuation payoff by $w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}})$.

Thus, player $(i-1)$ is the "arbiter" of player i's payoff, in that player $(i-1)$'s choice of her strategy state $\boldsymbol{x}_{i-1}$ determines player i's equilibrium continuation payoff in each state θ. This feature is typical of block belief-free constructions, such as those in Hörner and Olszewski (2006), Deb, Sugaya, and Wolitzky (2019), and Sugaya and Yamamoto (2019).

Beliefs. Players' belief systems $(\beta_i)_{i\in I}$ are specified as a function of the block strategy profile σ. Intuitively, players believe that trembles in the current block are much less likely than trembles in previous blocks, but that, within the current block, trembles in later periods are much more likely than trembles in earlier periods. This has two important implications. First, if a player reaches a history that can be explained by some past opponents' play that does not involve any deviations within the current block, she believes with probability 1 that no one deviated within the current block. Second, if a player reaches a history that cannot be explained without appealing to deviations within the current block, but can be explained by supposing that the only within-block deviation was made by her current opponent in the current period, then she believes with

[29] We omit messages $\left(m_{i,1^*}, m_{\mu_{i,1^*},1^*}\right)$ in the description of $h_{i,1^*}$, as there is no communication in period 1^* in our construction.

probability 1 that this is indeed what occurred.

To construct the belief system, first note that N and T^{**} determine the number of possible block history profiles $\left(h_i^t\right)_{i\in I,t\leq T^{**}}$.[30] Denote this number by $\tilde{c}$. Beliefs are derived from Bayes' rule along a sequence of completely mixed strategy profiles $\left(\sigma^l\right)_{l\in\mathbb{N}}$, in which each player i "trembles" uniformly over all messages and actions with probability $(1/l)^{\tilde{c}(T^{**}b-t)}$ in period $t\in\{1,\ldots,T^{**}\}$ of block b. As $l\to\infty$, the resulting beliefs display the properties discussed above.

Equilibrium Conditions. Fix $\alpha\in(0,1-\varepsilon)$, $\eta\in(0,1)$, and a target payoff $\tilde{v}\in F^{\alpha,\eta}$. Let $\tilde{v}^{\theta^*}$ be the associated value given θ^*. Let p^0 and p^1 denote, respectively, the probability that a given pair of players are both rational, and the probability that exactly one of them is rational. Define $\left(v^{\theta^*}\right)_{\theta^*}$ such that, for each i,

$$v_i^{\theta^*}=\begin{cases}\tilde{v}_i^{\theta^*}-\frac{1-\delta}{\delta}\frac{p^0+p^1\left(\frac{1+G-L}{2}\right)}{\Pr(|\theta^*|\geq\alpha N)} & \text{if } |\theta^*|\geq\alpha N,\\ 0 & \text{otherwise,}\end{cases}\tag{10}$$

and let $v=\sum_{\theta^*}p\left(I\setminus\theta^*\right)v^{\theta^*}$. In order to show $\tilde{v}\in E^*$, it suffices to show that, for sufficiently large l,

$$\left((1-\delta)\left(p^0+p^1\left(\frac{1+G-L}{2}\right)\right)+\delta v_i\right)_i\in E^*.\tag{11}$$

(Note that the left-hand side of this expression is player i's expected payoff when rational players play C in period 1 and receive continuation payoff $\left(v_i^{\theta^*}\right)_{\theta^*}$ starting in period 2.) This follows because

$$\begin{aligned}
&(1-\delta)\left(p^0+p^1\left(\frac{1+G-L}{2}\right)\right)+\delta v_i\\
=\;&(1-\delta)\left(p^0+p^1\left(\frac{1+G-L}{2}\right)\right)\\
&+\delta\left(\sum_{\theta^*:|\theta^*|\geq\alpha N}p\left(I\setminus\theta^*\right)\left(\tilde{v}_i^{\theta^*}-\frac{1-\delta}{\delta}\frac{p^0+p^1\left(\frac{1+G-L}{2}\right)}{\Pr\left(|\theta^*|\geq\alpha N\right)}\right)+\Pr\left(|\theta^*|<\alpha N\right)(0)\right)\\
=\;&\tilde{v}.
\end{aligned}$$

Suppose that l is large enough so that

$$\left|\frac{1-\delta}{\delta}\frac{p^0+p^1\left(\frac{1+G-L}{2}\right)}{\Pr\left(|\theta^*|\geq\alpha N\right)}\right|\leq\frac{\eta}{2}.$$

(This holds for large l, since $\delta\to1$ and $\Pr\left(|\theta^*|\geq\alpha N\right)\to1$ as $l\to\infty$.) Then $\tilde{v}\in F^{\alpha,\eta}$ implies that v satisfies the following conditions:

1. For each θ^* satisfying $|\theta^*|\geq\alpha N$, we have $B^{\frac{\eta}{2}}\left(v^{\theta^*}\right)\subset F^*\left(\theta^*\right)$. In contrast, for each θ^* satisfying $|\theta^*|<\alpha N$, we have $v^{\theta^*}=0$.

2. For each $i\in I$ and each $\theta^*,\theta^{*\prime}$ satisfying (i) $|\theta^*|,\left|\theta^{*\prime}\right|\geq\alpha N$, (ii) $\theta^*\ni i$, and (iii) $\theta^{*\prime}\not\ni i$, we have
$$v_i^{\theta^*}-v_i^{\theta^{*\prime}}\geq\frac{\eta}{2}.\tag{12}$$

[30] The size of the message sets $|M_{i,t}|$ used in the construction will be explicitly determined as a function of N and T^{**} in the course of the proof.

We now provide a sufficient condition to establish (11).

Fix a block length $T^{**} \in \mathbb{N}$, a block strategy profile $\left(\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)\right)_{i\in I, \boldsymbol{x}_i \in \{G,B\}^{|\Theta|}, h_i^{1^*} \in H_i^{1^*}}$, target payoffs (conditional on both θ and $\boldsymbol{x}$) $\left(v_i^\theta(x_{i-1}^\theta)\right)_{i\in I, \theta\in\Theta, x_{i-1}^\theta \in \{G,B\}}$, and continuation payoffs $\left(w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})\right)_{i\in I,\theta\in\Theta, x_{i-1}^\theta\in\{G,B\}, h_{i-1}^{T^{**}+1}\in H_{i-1}^{T^{**}+1}}$. We will show that the following set of conditions is sufficient for $\sum_{\theta^*} p\left(I \setminus \theta^*\right) v^{\theta^*} \in E^*$. In what follows, $\mathbb{E}^\sigma\left[\cdot|\cdot\right]$ denotes conditional expectation under block strategy profile σ, with the corresponding belief system defined above given σ. We also write $\tilde{h}_i^{b,t} \in \tilde{H}_i^{b,t}$ for a generic history in period t of block b of the infinitely repeated game, and write $h_i^t \in H_i^t$ for a generic block history in period t of a block. (Thus, $\tilde{h}_i^{b,t}$ records the outcomes of $(b-1)T^{**} + t - 1$ periods of play, while h_i^t records the outcomes of $t-1$ periods.) Finally, we write $\tilde{h}_i^{b,0} \in \tilde{H}_i^{b,0}$ for a generic repeated game history at the beginning of block b, before the determination of the first match in the block.

1. *[Sequential Rationality]* For each $\boldsymbol{x} \in \{G,B\}^{N|\Theta|}$, each $i \in I$, each $h_i^{1^*} \in H_i^{1^*}$, each $t \in \{1,\dots,T^{**}\}$, each $b \in \mathbb{N}$, and each $\tilde{h}_i^{b,t} \in \tilde{H}_i^{b,t}$, $\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)$ is a maximizer (over $\sigma_i \in \Sigma_i$) of

$$\sum_{h_{-i}^{1^*}\in H_{-i}^{1^*}} \beta_i\left(h_{-i}^{1^*}|\boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right) \mathbb{E}^{\left(\sigma_i, \sigma_{-i}^*\left(\boldsymbol{x}_{-i}, h_{-i}^{1^*}\right)\right)}\left[\begin{array}{c} (1-\delta)\sum_{\tau=1}^{T^{**}} \delta^{\tau-1}\hat{u}_i\left(\mathbf{a}_\tau\right) \\ +\delta^{T^{**}} w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) \end{array} \middle| \boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right].$$

   (Here, the sum $\sum_{\tau=1}^{T^{**}}$ is taken over all periods in the current block b, where the current period $t \in \{(b-1)T^{**}+2, \dots bT^{**}+1\}$ is some period in block b. Note also that sequential rationality is imposed "ex post" over vectors $\boldsymbol{x}_{-i} \in \{G,B\}^{(N-1)|\Theta|}$. This is the defining feature of a block belief-free construction. However, optimality with respect to $h_{-i}^{1^*}$ is demanded only in expectation, not ex post.)

2. *[Promise Keeping]* For each $\theta \in \Theta$, $i \in I$, $x_{i-1}^\theta \in \{G,B\}^N$, $b \in \mathbb{N}$, and $\tilde{h}^{b,0} \in \tilde{H}^{b,0}$,

$$v_i^\theta(x_{i-1}^\theta) = \mathbb{E}^{\sigma^*\left(\boldsymbol{x}, h^{1^*}\right)}\left[(1-\delta)\sum_{t=1}^{T^{**}} \delta^{t-1}\hat{u}_i\left(\mathbf{a}_t\right) + \delta^{T^{**}} w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})|\tilde{h}^{b,0}, \theta\right].$$

   (Note that player i's continuation payoff $v_i^\theta(x_{i-1}^\theta)$ is allowed to depend on $\tilde{h}^b$ only through θ.)

3. *[Self-Generation]* For each $\theta \in \Theta$, $i \in I$, we have either (i) $w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) \in (v_i^\theta(B), v_i^\theta(G))$ for each $x_{i-1}^\theta \in \{G,B\}$ and $h_{i-1}^{T^{**}+1} \in H_{i-1}^{T^{**}+1}$, or (ii) $w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) = v_i^\theta(B) = v_i^\theta(G)$ for each $x_{i-1}^\theta \in \{G,B\}$ and $h_{i-1}^{T^{**}+1} \in H_{i-1}^{T^{**}+1}$.

4. *[Feasibility]* For each $\theta \in \Theta$ and $i \in I$, we have either $v_i^\theta \in (v_i^\theta(B), v_i^\theta(G))$ or $v_i^\theta = v_i^\theta(B) = v_i^\theta(G)$.

   (This implies that, by appropriately randomizing her strategy state x_{i-1}^θ in the first block, player $(i-1)$ can deliver the target payoff v_i^θ to player i. Moreover, this randomization has full support.)

5. *[Incentive to take C in period 1^*]* For each $i \in I$,

$$\delta\sum_{\theta\in\Theta} \Pr\left(\theta|a_{i,1^*} = C, i\in\theta^*\right) v_i^\theta > (1-\delta)\max\{G,L\} + \delta\sum_{\theta\in\Theta}\Pr\left(\theta|a_{i,1^*} = D, i \in \theta^*\right) v_i^\theta.$$

Defining $\pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) := \frac{\delta^{T^{**}}}{1-\delta}\left(w_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) - v_i^\theta(x_{i-1}^\theta)\right)$, we can rewrite these conditions as follows:

1. *[Sequential Rationality]* For each $\boldsymbol{x} \in \{G,B\}^{N|\Theta|}$, $i \in I$, $h_i^{1^*} \in H_i^{1^*}$, $t \in \{1,\ldots,T^{**}\}$, $b \in \mathbb{N}$, and $\tilde{h}_i^{b,t} \in \tilde{H}_i^{b,t}$, $\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)$ maximizes

$$\sum_{h_{-i}^{1^*}\in H_{-i}^{1^*}} \beta_i\left(h_{-i}^{1^*}|\boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right)\mathbb{E}^{\left(\sigma_i,\sigma_{-i}^*\left(\boldsymbol{x}_{-i},h_{-i}^{1^*}\right)\right)}\left[\sum_{\tau=1}^{T^{**}}\delta^{\tau-1}\hat{u}_i\left(\mathbf{a}_\tau\right) + \pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})|\boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right]. \tag{13}$$

2. *[Promise Keeping]* For each $\theta \in \Theta$, $i \in I$, $x_{i-1}^\theta \in \{G,B\}^N$, $b \in \mathbb{N}$, and $\tilde{h}^{b,0} \in \tilde{H}^{b,0}$,

$$v_i^\theta(x_{i-1}^\theta) = \mathbb{E}^{\sigma^*(\boldsymbol{x})}\left[\frac{1-\delta}{1-\delta^{T^{**}}}\sum_{t=1}^{T^{**}}\delta^{t-1}\hat{u}_i\left(\mathbf{a}_t\right) + \pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})|\tilde{h}^{b,0},\theta\right]. \tag{14}$$

3. *[Self-Generation]* For each $\theta \in \Theta$ and $i \in I$, either (i)

$$\text{sign}\left(x_{i-1}^\theta\right)\pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) > 0 \text{ and } \left|\frac{1-\delta}{\delta^{T^{**}}}\pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})\right| < v_i^\theta(G) - v_i^\theta(B) \tag{15}$$

for each $x_{i-1}^\theta \in \{G,B\}$ and $h_{i-1}^{T^{**}+1} \in H_{i-1}^{T^{**}+1}$, or (ii) $\pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) = 0$ and $v_i^\theta(G) = v_i^\theta(B)$ for each $x_{i-1}^\theta \in \{G,B\}$ and $h_{i-1}^{T^{**}+1} \in H_{i-1}^{T^{**}+1}$, where $\text{sign}\left(x_{i-1}^\theta\right) := 1_{\{x_{i-1}^\theta = B\}} - 1_{\{x_{i-1}^\theta = G\}}$.

4. *[Feasibility]* For each $\theta \in \Theta$ and $i \in I$,

$$v_i^\theta \in (v_i^\theta(B), v_i^\theta(G)) \text{ or } v_i^\theta = v_i^\theta(B) = v_i^\theta(G). \tag{16}$$

5. *[Incentive to take C in period 1^*]* For each $i \in I$,

$$\delta\sum_{\theta\in\Theta}\Pr\left(\theta|a_{i,1^*} = C, i \in \theta^*\right)v_i^\theta > (1-\delta)\max\{G,L\} + \delta\sum_{\theta\in\Theta}\Pr\left(\theta|a_{i,1^*} = D, i \in \theta^*\right)v_i^\theta. \tag{17}$$

Lemma 12 *For all $\boldsymbol{v} \in \mathbb{R}^{N|\Theta|}$ and $\delta \in [0,1)$, if there exist $T^{**} \in \mathbb{N}$, $\left(\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)\right)_{i\in I, \boldsymbol{x}_i\in\{G,B\}^{|\Theta|}, h_i^{1^*}\in H_i^{1^*}}$, $\left(v_i^\theta(x_{i-1}^\theta)\right)_{i\in I,\theta\in\Theta,x_{i-1}^\theta\in\{G,B\}}$, and $\left(\pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})\right)_{i\in I,\theta\in\Theta,x_{i-1}^\theta\in\{G,B\},h_{i-1}^{T^{**}+1}\in H_{i-1}^{T^{**}+1}}$ such that Conditions (13)–(17) are satisfied, then*

$$\left((1-\delta)\left(p^0 + p^1\left(\frac{1+G-L}{2}\right)\right) + \delta\sum_{\theta^*}p\left(I\setminus\theta^*\right)v_i^{\theta^*}\right)_i \in E^*. \tag{18}$$

Proof. Conditions (14) and (15) imply that payoffs $\left(v_i^\theta(x_{i-1}^\theta)\right)_{i\in I,\theta\in\Theta,x_{i-1}^\theta\in\{G,B\}}$ can be delivered at the beginning of each block with full support state transition probabilities, and Condition (16) then implies that, by appropriately randomizing over $\left(x_{i-1}^\theta\right)_{i\in I,\theta\in\Theta}$ before the first block (i.e., before period 2 of the repeated game), the target expected payoff vector $\boldsymbol{v}$ can be delivered. This is as in, for example, Hörner and Olszewski (2006). Condition (13) is then a more stringent version of the resulting sequential rationality constraint, as it imposes sequential rationality for each realization of $\boldsymbol{x}_{-i}$, rather than only in expectation. Thus, Conditions (13)–(16) imply that the strategies $\left(\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)\right)_{i\in I, \boldsymbol{x}_i\in\{G,B\}^{|\Theta|}, h_i^{1^*}\in H_i^{1^*}}$ are sequentially rational and deliver continuation payoffs $\boldsymbol{v}$ starting from the second period of the repeated game. Given this, (17) implies that it is optimal for rational players to take C in period 1^*. Finally, the resulting ex ante expected payoffs are given by (18). ■

To prove Theorem 3, it thus suffices to show that, for any $\tilde{v} \in F^{\alpha,\eta}$ and v defined by (10), for sufficiently large l there exist $T^{**} \in \mathbb{N}$, $\left(\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)\right)_{i\in I, \boldsymbol{x}_i\in\{G,B\}^{|\Theta|}, h_i^{1^*}\in H_i^{1^*}}$, $\left(v_i^\theta(x_{i-1}^\theta)\right)_{i\in I,\theta\in\Theta, x_{i-1}^\theta\in\{G,B\}}$, and $\left(\pi_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})\right)_{i\in I,\theta\in\Theta, x_{i-1}^\theta\in\{G,B\}, h_{i-1}^{T^{**}+1}\in H_{i-1}^{T^{**}+1}}$ such that Conditions (13)–(17) are satisfied.

Condition (17). It is immediate that Condition (17) is satisfied for sufficiently large l. For, taking C gives player i payoff at least

$$(1-\delta)\left(\Pr\left(a_{\mu_{i,1^*},1^*} = C | i \in \theta^*\right)(1) + \Pr\left(a_{\mu_{i,1^*},1^*} = D | i \in \theta^*\right)(-L)\right)$$
$$+\delta\left(\Pr\left(|\theta| \geq \alpha N | i \in \theta\right) \min_{\theta:|\theta|\geq\alpha N, \theta\ni i} v_i^\theta + \left(1 - \Pr\left(|\theta| \geq \alpha N | i \in \theta\right)\right)(0)\right),$$

while taking D gives player i payoff at most

$$(1-\delta)\left(\Pr\left(a_{\mu_{i,1^*},1^*} = C | i \in \theta^*\right)(1+G) + \Pr\left(a_{\mu_{i,1^*},1^*} = D | i \in \theta^*\right)(0)\right) + \delta \max_{\theta:|\theta|\geq\alpha N, \theta\not\ni i} v_i^\theta.$$

Since $\lim_l \Pr\left(|\theta| \geq \alpha N | i \in \theta\right) = 1$, (12) implies (17) for sufficiently large l.

B.2.4 Target Actions

We now define a target (opponent identity-contingent) action profile $\mathbf{a}^{x^\theta}$ for each state $\theta \subset I$.

For θ satisfying $|\theta| < \alpha N$, we define $\mathbf{a}_i^{x^\theta}(j) = D$ for all $x^\theta \in \{G,B\}^N$ and $i \neq j$. That is, all players are prescribed defection. In this case, we define $v_i^\theta(G) = v_i^\theta(B) = 0$. Note that, to satisfy (15), this requires $\pi_i^\theta\left(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}\right) = 0$ for all $x_{i-1}^\theta \in \{G,B\}$ and $h_{i-1}^{T^{**}+1} \in H_{i-1}^{T^{**}+1}$.

For θ satisfying $|\theta| \geq \alpha N$, for each $x^\theta \in \{G,B\}^N$ we define $\mathbf{a}^{x^\theta}$ such that, for each $i \in I$, $u_i\left(\mathbf{a}^{x^\theta}\right) > v_i^\theta$ if $x_{i-1}^\theta = G$, and $u_i\left(\mathbf{a}^{x^\theta}\right) < v_i^\theta$ if $x_{i-1}^\theta = B$.[31] Define $v_i^\theta(G)$, $v_i^\theta(B)$, and $\bar{\varepsilon} > 0$ such that

$$\left(\max_{x^\theta:x_{i-1}^\theta=B} u_i\left(\mathbf{a}^{x^\theta}\right)\right)_+ \leq v_i^\theta(B) + 4\bar{\varepsilon} < v_i^\theta < v_i^\theta(G) - 4\bar{\varepsilon} \leq \min_{x^\theta:x_{i-1}^\theta=G} u_i\left(\mathbf{a}^{x^\theta}\right). \tag{19}$$

Note that such $\bar{\varepsilon}$ exists since $B^{\frac{\eta}{2}}\left(v^\theta\right) \subset F^*(\theta)$. With these definitions, (16) is satisfied.

B.2.5 Structure of the Block

Each block consists of the following sub-blocks: Let

$$K := \left\lceil \frac{\max\{G, L\}}{\bar{\varepsilon}} \right\rceil. \tag{20}$$

[31] As in Hörner and Olszewski (2006) and several subsequent papers, it may actually be necessary for players to cycle through a sequence of distinct action profiles $\mathbf{a}^{x^\theta}$ to achieve average payoffs $u_i > v_i^\theta$ (resp., $u_i < v_i^\theta$) for i such that $x_{i-1}^\theta = G$ (resp., $x_{i-1}^\theta = B$). Accommodating this possibility poses no difficulty for the proof, so we follow Hörner and Olszewski (and others) in assuming that a single action profile suffices.

Now fix Z sufficiently large such that $Z \geq \bar{Z}$ (with c and $\bar{Z}$ given in Lemma 11) and

$$(K+3)\left(\frac{1}{Z}+2\exp(-cZ)\right)\bar{u} \leq \bar{\varepsilon}, \tag{21}$$

where $\bar{u} = 2\max\{L, 1+G\}$. Let $T = Z\log_2 N$. In what follows, recall that players always take action D while circulating information.

1. **1^*-communication sub-block** (the first T periods of the block): Players circulate information about h^{1^*}.
2. **$\boldsymbol{x}$-communication sub-block** (the next T periods): Players circulate information about $\boldsymbol{x}$.
3. **Supplemental round** 0 (the next T periods): Players circulate information about the first two sub-blocks.
4. **Main sub-block** k (there are K main sub-blocks, each lasting for $(1+Z)T$ periods, and each divided into the following two rounds):
   (a) **Main round** k (the first ZT periods of the sub-block): Players take the target actions (and do not send cheap talk messages).
   (b) **Supplemental round** k (the next T periods of the sub-block): Players circulate information about the history up to the end of main round k.

Recall that T^{**} denotes the length of the block, or equivalently the last period of supplemental round K. Let T^* denote the last period of main round K. Note that $T^{**} = (3+K(1+Z))Z\log_2 N$. Since $(1-\delta)\log N \to 0$, we have

$$\limsup_{l\to\infty}(1-\delta)T^* \leq \limsup_{l\to\infty}(1-\delta)T^{**} = 0. \tag{22}$$

B.2.6 Reduction Lemma

We now show that, by communicating their histories during supplemental round K (the last such round in the block) and adjusting continuation payoffs appropriately, the players can effectively cancel the effects of discounting while letting continuation payoffs depend on $\left(x^{\theta}_{-i}, h^{1^*}, h^{T^*+1}\right)$ rather than $\left(x^{\theta}_{i-1}, h^{T^{**}+1}_{i-1}\right)$ (when $|\theta| \geq \alpha N$).[32]

Let $\Sigma_i^{T^*}$ denote the set of i's block strategies up to period T^*. We show that the following conditions are sufficient for (18).

1. *[Sequential Rationality]* For each $\boldsymbol{x} \in \{G,B\}^{N|\Theta|}$, $i \in I$, $h_i^{1^*} \in H_i^{1^*}$, $t \in \{1,\ldots,T^*\}$, $b \in \mathbb{N}$, and $\tilde{h}_i^{b,t} \in \tilde{H}_i^{b,t}$, $\sigma_i^{T^*}\left(\boldsymbol{x}_i, h_i^{1^*}\right)$ maximizes (over $\sigma_i \in \Sigma_i^{T^*}$)

$$\sum_{h_{-i}^{1^*}\in H_{-i}^{1^*}} \beta_i\left(h_{-i}^{1^*}|\boldsymbol{x}_{-i},\tilde{h}_i^{b,t}\right)\mathbb{E}^{\left(\sigma_i,\sigma^*_{-i}\left(\boldsymbol{x}_{-i},h_{-i}^{1^*}\right)\right)}\left[\begin{array}{c} 1_{\{\theta:|\theta|\geq\alpha N\}}\left(\sum_{\tau=1}^{T^*}\hat{u}_i(\mathbf{a}_\tau)+\pi_i^{\theta}(x^{\theta}_{-i},h^{1^*},h^{T^*+1})\right) \\ +1_{\{\theta:|\theta|<\alpha N\}}\sum_{\tau=1}^{T^*}\delta^{\tau-1}\hat{u}_i(\mathbf{a}_\tau) \\ |\boldsymbol{x}_{-i},\tilde{h}_i^{b,t}\end{array}\right]. \tag{23}$$

[32] A similar but more complicated argument appears in Deb, Sugaya, and Wolitzky (2020).

2. *[Promise Keeping]* For each $\theta \in \Theta$, $\boldsymbol{x} \in \{G, B\}^{N|\Theta|}$, $i \in I$, $b \in \mathbb{N}$, and $\tilde{h}^{b,0} \in \tilde{H}^{b,0}$,

$$v_i^\theta(x_{i-1}^\theta) = \mathbb{E}^{\sigma^*\left(\boldsymbol{x}, h^{1^*}\right)} \left[\begin{array}{c} 1_{\{\theta:|\theta|\geq\alpha N\}} \frac{1}{T^*} \left(\sum_{\tau=1}^{T^*} \hat{u}_i \left(\mathbf{a}_\tau\right) + \pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1}) \right) \\ +1_{\{\theta:|\theta|<\alpha N\}} \frac{1-\delta}{1-\delta^{T^*}} \sum_{\tau=1}^{T^*} \delta^{\tau-1} \hat{u}_i \left(\mathbf{a}_\tau\right) \end{array} |\theta, \tilde{h}^{b,0} \right]. \quad (24)$$

3. *[Self-Generation]* For each $\theta \in \Theta$, $i \in I$, $\boldsymbol{x}_{i-1} \in \{G, B\}^{|\Theta|}$, $h^{1^*} \in H^{1^*}$, and $h^{T^*+1} \in H^{T^*+1}$,

$$\text{sign}\left(x_{i-1}^\theta\right) \pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1}) \begin{cases} > 0 & \text{for } \theta \text{ satisfying } |\theta| \geq \alpha N, \\ = 0 & \text{for } \theta \text{ satisfying } |\theta| < \alpha N. \end{cases} \quad (25)$$

and

$$\left|\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1})\right| \begin{cases} \leq 2\bar{u}T^* & \text{for } \theta \text{ satisfying } |\theta| \geq \alpha N, \\ = 0 & \text{for } \theta \text{ satisfying } |\theta| < \alpha N. \end{cases} \quad (26)$$

Lemma 13 *For any sequence $(N, \delta)_l$ such that $(1-\delta)\log N \to 0$, suppose there exists $\bar{l}$ such that, for each $l \geq \bar{l}$ and corresponding $(N, \delta)_l$, there exist $\left(\sigma_i^{T^*}\left(\boldsymbol{x}_i, h_i^{1^*}\right)\right)_{i\in I, \boldsymbol{x}_i \in \{G,B\}^{|\Theta|}, h_i^{1^*}\in H_i^{1^*}}$ and $\left(\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1})\right)_{i\in I, x_{-i}^\theta\in\{G,B\}^{N-1}, h^{1^*}, h^{T^*+1}}$ such that Conditions (23)–(26) are satisfied. Then, for sufficiently large l, (18) holds.*

Proof. We have fixed $v_i^\theta(x_{i-1}^\theta)$. Fix $\sigma_i^{T^*}$ and π_i^θ satisfying (23)–(26). We will construct $\tilde{\sigma}_i^*$, $\tilde{\pi}_i^\theta$, and $\tilde{v}_i^\theta(x_{i-1}^\theta)$ that satisfy (13)–(16).

We extend strategy $\sigma_i^{T^*} \in \Sigma_i^{T^*}$ to a strategy $\tilde{\sigma}_i^* \in \Sigma_i$ by specifying that players circulate message $m = (m_i)_i = (\boldsymbol{x}_i, h_i^{1^*}, h_i^{T^*+1})_i$ in supplemental round K.

Given player $i-1$'s history in supplemental round K, we define $\tilde{\pi}_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1})$ as follows. (i) If $m_{-i}(i-1) = \texttt{error}$, then $\tilde{\pi}_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) = 0$ for each $x_{i-1}^\theta, h_{i-1}^{T^{**}+1}$. (ii) Otherwise, player $i-1$ infers $(h_{-i}^{1^*}(i-1), h_{-i}^{T^*+1}(i-1))$. Since matching is pairwise, there exists a unique $(h^{1^*}(i-1), h^{T^*+1}(i-1))$ that is consistent with $(h_{-i}^{1^*}(i-1), h_{-i}^{T^*+1}(i-1))$. Given $h^{T^*+1}(i-1)$, let $\mathbf{a}_t(i-1)$ be the action in period t. We define

$$\tilde{\pi}_i^\theta(x_{i-1}^\theta, h_{i-1}^{T^{**}+1}) = \begin{cases} \begin{array}{c} \frac{T^*\left(1-\delta^{T^*}\right)\text{sign}\left(x_{i-1}^\theta\right)\bar{u} + \sum_{t=1}^{T^*}\left(1-\delta^{t-1}\right)\hat{u}_i(\mathbf{a}_t(i-1))}{\Pr(m_{-i}(i-1)\neq\texttt{error})} \\ + \frac{\pi_i^\theta(\boldsymbol{x}_{-i}(i-1), h^{1^*}(i-1), h^{T^*+1}(i-1))}{\Pr(m_{-i}(i-1)\neq\texttt{error})} \end{array} & \text{if } |\theta| \geq \alpha N, \\ 0 & \text{if } |\theta| < \alpha N. \end{cases}$$

Finally, we define

$$\tilde{v}_i^\theta(x_{i-1}^\theta) = \begin{cases} \frac{1-\delta}{1-\delta^{T^{**}}} T^* v_i^\theta(x_{i-1}^\theta) + (1-\delta)\, T^* \text{sign}\left(x_{i-1}^\theta\right) \bar{u} & \text{if } |\theta| \geq \alpha N, \\ v_i^\theta(x_{i-1}^\theta) & \text{if } |\theta| < \alpha N. \end{cases}$$

As $l \to \infty$, since $\frac{1-\delta}{1-\delta^{T^{**}}} T^* \to 1$ and $(1-\delta)\, T^* \to 0$, we have $\tilde{v}_i^\theta(x_{i-1}^\theta) \to v_i^\theta(x_{i-1}^\theta)$ uniformly for each θ. It remains to show that $\tilde{\sigma}_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)$, $\tilde{\beta}^*$, $\tilde{\pi}_i^\theta$, and $\tilde{v}_i^\theta(x_{i-1}^\theta)$ satisfy (13)–(16).

Note that player i's payoff depends on the outcome of play in supplemental round K only through her stage game payoffs (which are maximized by taking D) and $m_{-i}(i-1)$. Since player i cannot affect the distribution of $m_{-i}(i-1)$, following $\sigma_i^*\left(\boldsymbol{x}_i, h_i^{1^*}\right)|_{h_i^{T^*+1}}$ is optimal. Given this, by the law of iterated expectation, in period $t \leq T^*$, the expected value of

$$\sum_{\tau=t}^{T^*} \delta^{t-1} \hat{u}_i(\mathbf{a}_\tau) + \tilde{\pi}_i^\theta (x_{i-1}^\theta, h_{i-1}^{T^{**}})$$

given θ, $\boldsymbol{x}$, h^{1^*}, and $\tilde{h}^t$ is equal to

$$1_{\{\theta:|\theta|\geq\alpha N\}} \left(\left(1-\delta^{T^*}\right) T^* \text{sign}\left(x_{i-1}^\theta\right) \bar{u} + \sum_{\tau=t}^{T^*} \hat{u}_i(\mathbf{a}_\tau) + \pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1}) \right) + 1_{\{\theta:|\theta|<\alpha N\}} \sum_{\tau=t}^{T^*} \delta^{\tau-1} \hat{u}_i(\mathbf{a}_\tau).$$

Ignoring the constant $\left(1-\delta^{T^*}\right) T^* \text{sign}\left(x_{i-1}^\theta\right) \bar{u}$, (23) implies (13).

For $|\theta| < \alpha N$, (14)–(16) hold since all the payoffs and rewards are zero regardless of x_{i-1}^θ and $h_{i-1}^{T^{**}+1}$. For $|\theta| \geq \alpha N$, (14) follows from (24) given the definition of $\tilde{v}_i^\theta(x_{i-1}^\theta)$. In addition, (25) and (26) imply (15) for sufficiently large l since $\delta \to 1$ and $v_i^\theta(G) - v_i^\theta(B) \geq 8\bar{\varepsilon}$ by (19). Finally, (19) implies that there is $4\bar{\varepsilon}$ slack between $v_i^\theta(x_{i-1}^\theta)$ and v_i^θ for each θ satisfying $|\theta| \geq \alpha N$. Hence, (16) holds with $\tilde{v}_i^\theta(x_{i-1}^\theta)$ for sufficiently large l. ■

B.2.7 Equilibrium Strategies

We now complete the description of the equilibrium strategies.

It will be useful to define the notion of a "detectable deviation" by player i. As we will see, given player i's period 1^* history $h_i^{1^*}$ and her strategy state $\boldsymbol{x}_i$, her block strategy is pure along the equilibrium path of play. Given $h_i^{1^*}$ and an on-path period t block history h_i^t, we say that a period t message $m_{i,t}$ is a *detectable deviation* if there does not exist a strategy state $\hat{\boldsymbol{x}}_i$ such that $\left(h_i^t, m_{i,t}\right)$ occurs with positive probability given $\left(\hat{\boldsymbol{x}}_i, h_i^{1^*}\right)$; similarly, given a triple $\left(h_i^t, m_{i,t}, m_{\mu_{i,t},t}\right)$, an action $a_{i,t}$ is a *detectable deviation* if there does not exist a strategy state $\hat{\boldsymbol{x}}_i$ such that $\left(h_i^t, m_{i,t}, m_{\mu_{i,t},t}, a_{i,t}\right)$ occurs with positive probability given $\left(\hat{\boldsymbol{x}}_i, h_i^{1^*}\right)$. We say a player *detectably deviates* if she plays a detectable deviation.

1^*-Communication Sub-Block Players circulate message $m = (m_i)_i$, where m_i is the set of players whom player i knows to have taken C in period 1^*: that is, $m_i = \left\{i, \mu_{i,1^*}\right\} \cap \theta$.

Let h_i^{T+1} be player i's history at the end of the sub-block. We define $\theta\left(h_i^{T+1}\right) = \emptyset$ if, for some $j \neq i$, either $\zeta_{i,T}^{I,-j} \neq -j$ (i.e., i does not receive each player's message through a path excluding j) or $m_{-j}(i) = \texttt{error}$ (i.e., i receives inconsistent messages through a path excluding j). We also define $\theta\left(h_i^{T+1}\right) = \emptyset$ if there exist $j \neq j' \neq k \neq j$ such that $m_{-j}(i)|_k \neq m_{-j'}(i)|_k$ (i.e., i receives inconsistent messages through a path excluding j and through a path excluding j'). Otherwise, we define $\theta\left(h_i^{T+1}\right) = \bigcup_{j\neq i} \bigcup_{k\neq j} m_{-j}(i)|_k$ (i.e., $\theta\left(h_i^{T+1}\right)$ is the set of players who i has been told took C in period 1^*).

Lemma 11 immediately implies the following result.

Lemma 14 *Suppose all players follow the protocol. There exist $c > 0$ and $\bar{Z} > 0$ such that, for all $Z > \bar{Z}$ and all l, we have*

$$\Pr\left(\theta(h_i^{T+1}) = \theta \ \forall i\right) \geq 1 - \exp(-cZ).$$

We record two key properties of player i's beliefs about θ. Suppose the current block is block b.

First, for each $t \geq T+1$, $\tilde{h}_i^{b,t} \in \tilde{H}_i^{b,t}$, and $\boldsymbol{x}_{-i} \in \{G,B\}^{|\Theta|(N-1)}$, player i believes that $\theta \supseteq \theta(h_i^{T+1})$:

$$\sum_{\theta \supseteq \theta(h_i^{T+1})} \beta_i \left(\theta | \boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right) = 1. \tag{27}$$

This is trivial if $\theta_i(h_i^{T+1}) = \emptyset$. Otherwise, since trembles in earlier blocks are more likely, $\beta_i \left(\theta = \theta(h_i^{T+1}) | \boldsymbol{x}_{-i}, \tilde{h}_i^{b,T+1}\right) = 1$ (i.e., player i believes that $\theta = \theta\left(h_i^{T+1}\right)$ at the end of the 1*-communication sub-block). Moreover, since trembles are more likely in later periods within the block, player i continues to believe that $\theta = \theta\left(h_i^{T+1}\right)$ for the duration of the block.

Second, for each $t \geq T+1$, $\tilde{h}_i^{b,t} \in \tilde{H}_i^{b,t}$, and $\boldsymbol{x}_{-i} \in \{G,B\}^{|\theta|(N-1)}$, player i believes that $\theta \supseteq \theta(h_j^{T+1})$ for each $j \neq i$:

$$\sum_{\theta \supseteq \theta(h_j^{T+1})} \beta_{i,t} \left(\theta | \boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right) = 1. \tag{28}$$

This holds by similar reasoning. Note that, if player i deviates in the 1*-communication sub-block, this can only switch $\theta(h_j^{T+1})$ from θ to $\emptyset$ (this is because, to have $\theta(h_j^{T+1}) \neq \emptyset$, player j needs to receive messages from every player except for i through the path excluding i; hence, by telling a lie, player i can only create an inconsistency in the messages that player j receives), and hence cannot affect the probability that $\theta \supseteq \theta(h_j^{T+1})$.

$\boldsymbol{x}$-Communication Sub-Block Players circulate message $m = (m_i)_i = (\boldsymbol{x}_i)_i$. Slightly abusing notation, let $h_i^{\leq 0}$ denote player i's history at the end of the $\boldsymbol{x}$-communication sub-block.

If player i infers $m_{-j}(i) = \texttt{error}$ for some $j \neq i$, we define $\boldsymbol{x}(i) = \boldsymbol{B}$, where $\boldsymbol{B} \in \{G,B\}^{N|\Theta|}$ denotes the vector with B in every component. If instead i infers some $m_{-j}(i) \in \times_{n \neq i} M_n$ for each $j \neq i$, then:

1. If there exists $\hat{\boldsymbol{x}}_{-i} \in \{G,B\}^{(N-1)|\Theta|}$ such that $m_{-j}(i)|_n = \hat{\boldsymbol{x}}_{-i}|_n$ for all $j \neq i \neq n \neq j$, we define $\boldsymbol{x}(i) = (\boldsymbol{x}_i, \hat{\boldsymbol{x}}_{-i})$.
2. Otherwise, we define $\boldsymbol{x}(i) = \boldsymbol{B}$.[33]

Finally, we define $x(i) = \boldsymbol{x}(i)^{\theta\left(h_i^{T+1}\right)}$.

Supplemental Round 0 Players circulate message $m = (m_i)_i = \left(h_i^{\leq 0}\right)_i$. Let $h_i^{<1}$ denote player i's history at the end of supplemental round 0.

We define $I^D\left(h_i^{<1}\right) = 1$ if any of the following hold:

1. $\left|(\theta(h_i^{T+1}))\right| < \alpha N$.
2. Player i detectably deviates in either the $\boldsymbol{x}$-communication sub-block or supplemental round 0.
3. $m_{-j}(i) = \texttt{error}$ for some $j \neq i$ in either the $\boldsymbol{x}$-communication sub-block or supplemental round 0.

[33] Note that this can occur even if $m_{-j}(i) \neq \texttt{error}$, as in the situation noted in footnote 28.

Otherwise, for each $j \neq i$, $m_{-j}(i) = \times_{n \neq i} h_n^{\leq 0}$ for some $\times_{n \neq i} h_n^{\leq 0} \in \times_{n \neq i} H_n^{\leq 0}$. If there exists a player $j \neq i$ such that, according to history $\left(h_i^{\leq 0}, m_{-j}(i)\right)$, player j detectably deviated in the 1^*-communication sub-block or the $\boldsymbol{x}$-communication sub-block, then we define $I^D\left(h_i^{<1}\right) = 1$. Otherwise, we define $I^D\left(h_i^{<1}\right) = 0$.

Main Sub-Block k, $k \in \{1, \ldots, K\}$ For each $k \in \{1, \ldots, K\}$, each player i enters sub-block k with state variables $x(i) \in \{G, B\}^N$ and $I^D\left(h_i^{<k}\right) \in \{0, 1\}$. The state variable $x(i)$ was determined at the end of the $\boldsymbol{x}$-communication sub-block, and remains constant throughout the main sub-blocks. The state variable $I^D\left(h_i^{<1}\right)$ was determined at the end of supplemental round 0; the state variable $I^D\left(h_i^{<k}\right)$ may switch from 0 to 1 during some main sub-block, in which case it remains equal to 1 for the duration of the block.

We now define player i's strategy in main sub-block k as a function of $x(i)$ and $I^D\left(h_i^{<k}\right)$, and then specify how $I^D\left(h_i^{<k+1}\right)$ evolves.

Main round actions as a function of $x(i)$ and $I^D\left(h_i^{<k}\right)$: If $I^D\left(h_i^{<k}\right) = 1$, then player i takes D throughout the round. If $I^D\left(h_i^{<k}\right) = 0$, then player i takes $\mathbf{a}_i^{x(i)}$ throughout the round, unless she herself deviates from $\mathbf{a}_i^{x(i)}$ during the round. If such a deviation occurs, she takes D for the rest of the round. Let $h_i^{\leq k}$ denote player i's history at the end of main round k.

Supplemental round communication as a function of $h_i^{\leq k}$: Players circulate message $m = (m_i)_i = \left(h_i^{\leq k}\right)_i$.

Determination of $I^D\left(h_i^{<k+1}\right)$: Set $I^D\left(h_i^{<k+1}\right) = 1$ if any of the following hold:

1. $I^D\left(h_i^{<k}\right) = 1$.
2. Player i detectably deviated during main sub-block k.
3. $m_{-j}(i) = \texttt{error}$ for some $j \neq i$ during supplemental round k.
4. For each $j \neq i$, $m_{-j}(i) = \times_{n \neq i} h_n^{\leq k}$ for some $\times_{n \neq i} h_n^{\leq k} \in \times_{n \neq i} H_n^{\leq k}$, and there exists a player $j \neq i$ such that, according to history $\left(h_i^{\leq k}, m_{-j}(i)\right)$, player j detectably deviated during main round k.

Otherwise, set $I^D\left(h_i^{<k+1}\right) = 0$.

B.2.8 Reward Function

Given the above block strategy profile, we now define the reward function $\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1})$. For θ satisfying $|\theta| < \alpha N$, define $\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1}) = 0$ for all $(x_{-i}^\theta, h^{1^*}, h^{T^*+1})$. This satisfies Conditions (24)–(26); we verify Condition (23) (sequential rationality) in the next subsection. For the remainder of this section, assume $|\theta| \geq \alpha N$.

Given $\left(h^{1^*}, h^{T^*+1}\right)$, we define $\chi_i(h^{1^*}, h^{T^*+1}) = 1$ if there exists a player $j \neq i$ who detectably deviated from the prescribed block strategy (according to h^{T^*+1}) or if the match realization was erroneous in any round in the current block (again, according to h^{T^*+1}). We define $\chi_i(h^{1^*}, h^{T^*+1}) = 0$ otherwise. Lemma 11 immediately implies the following result.

Lemma 15 *Suppose player i's opponents follow the prescribed strategy. For all l and all θ (and regardless of player i's own strategy), we have*

$$\Pr\left(\chi_i(h^{1^*}, h^{T^*+1}) = 1|\theta\right) \leq (3+K)\exp(-cZ).$$

Next, define $I_i^D(h^{1^*}, h^{T^*+1}) = 1$ if player i detectably deviated from the prescribed strategy, and define $I_i^D(h^{1^*}, h^{T^*+1}) = 0$ otherwise. Finally, given $\left(h^{1^*}, h^{T^*+1}\right)$ satisfying $I_i^D(h^{1^*}, h^{T^*+1}) = 0$, define $\hat{\boldsymbol{x}}_i(h^{1^*}, h^{T^*+1})$ to be that value of $\hat{\boldsymbol{x}}_i$ for which the history $\left(h_i^{1^*}, h_i^{T^*+1}\right)$ is consistent with player i taking strategy $\sigma_i^{T^*}(\hat{\boldsymbol{x}}_i, h_i^{1^*})$. Such $\hat{\boldsymbol{x}}_i$ is uniquely determined since player i communicates $\hat{\boldsymbol{x}}_i$ in the $\boldsymbol{x}$-communication sub-block.

Define the function $\pi_i^{\text{cancel}}\left(x_{i-1}^\theta, \boldsymbol{a}\right) : \{G, B\} \times A^N \to [-\bar{u}, \bar{u}]$ such that, for each $\boldsymbol{a} \in A^N$, we have

$$\begin{cases} \hat{u}_i(\boldsymbol{a}) + \pi_i^{\text{cancel}}\left(x_{i-1}^\theta, \boldsymbol{a}\right) = \text{sign}\left(x_{i-1}^\theta\right)\frac{1}{2}\bar{u} \\ \text{sign}\left(x_{i-1}^\theta\right)\pi_i^{\text{cancel}}\left(x_{i-1}^\theta, \boldsymbol{a}\right) \geq 0 \end{cases} \tag{29}$$

Thus, the function $\pi_i^{\text{cancel}}\left(x_{i-1}^\theta, \boldsymbol{a}\right)$ cancels player i's instantaneous utility and leaves player i a negative (resp., positive) payoff when $x_{i-1}^\theta = G$ (resp., B)

If $\chi_i(h^{1^*}, h^{T^*+1}) = 1$, define

$$\hat{\pi}_i^\theta(x_{-i}^\theta, h^{T^*+1}) = \sum_{t=1}^{T^*} \pi_i^{\text{cancel}}\left(x_{i-1}^\theta, \boldsymbol{a}_t\right). \tag{30}$$

If $\chi_i(h^{1^*}, h^{T^*+1}) = 0$, define

$$\hat{\pi}_i^\theta(x_{-i}^\theta, h^{T^*+1}) = \begin{cases} 1_{\left\{I_i^D(h^{1^*}, h^{T^*+1})=0\right\}}\bar{\varepsilon}T^* & \text{if } x_{i-1}^\theta = B, \\ -1_{\left\{I_i^D(h^{1^*}, h^{T^*+1})=1\right\}}2\bar{u}T^* & \text{if } x_{i-1}^\theta = G. \end{cases}$$

That is, if $x_{i-1}^\theta = B$ then player i is rewarded if she follows the prescribed strategy; and if $x_{i-1}^\theta = G$ then she is punished if she detectably deviates.

Let

$$u_i\left(x^\theta, h^{1^*}\right) = \frac{1}{T^*}\mathbb{E}^{\sigma(\boldsymbol{x})}\left[\sum_{\tau=1}^{T^*} \hat{u}_i(\mathbf{a}_\tau) + \hat{\pi}_i^\theta(x_{-i}^\theta, h^{T^*+1}) + \text{sign}\left(x_{i-1}^\theta\right)\bar{\varepsilon}T^* | h^{1^*}, \theta\right]. \tag{31}$$

Note that the right hand side of (31) depends only on x^θ and h^{1^*} since (i) when $\chi_i(h^{1^*}, h^{T^*+1}) = 0$, $\mathbf{a}_\tau = \mathbf{a}^{x^\theta}$ in main rounds, (ii) when $\chi_i(h^{1^*}, h^{T^*+1}) = 1$, (29) implies that player i's payoff from $\hat{u}_i$ and π_i^{cancel} depends only on x_{i-1}^θ in main rounds, (iii) the distribution of $\chi_i(h^{1^*}, h^{T^*+1})$ is determined by the match realization, and (iv) in non-main rounds, players take defection for sure.

Note that

$$
\begin{aligned}
& \left|u_i\left(x^\theta,h^{1^*}\right)-\hat{u}_i\left(\mathbf{a}^{x^\theta}\right)\right| \\
= \ & \Pr\left(\chi_i(h^{1^*},h^{T^*+1})=0|\theta\right)\left|\frac{1}{T^*}\mathbb{E}^{\sigma(\boldsymbol{x})}\left[\left(\begin{array}{c}\sum_{\tau=1}^{T^*}\hat{u}_i\left(\mathbf{a}_\tau\right)\\ +\hat{\pi}_i^\theta(x_{-i}^\theta,h^{T^*+1})\\ +\text{sign}\left(x_{i-1}^\theta\right)\bar{\varepsilon}T^*\end{array}\right)|h^{1^*},\theta,\chi_i(h^{1^*},h^{T^*+1})=0\right]-\hat{u}_i\left(\mathbf{a}^{x^\theta}\right)\right| \\
& +\Pr\left(\chi_i(h^{1^*},h^{T^*+1})=1|\theta\right)\left|\frac{1}{T^*}\mathbb{E}^{\sigma(\boldsymbol{x})}\left[\left(\begin{array}{c}\sum_{\tau=1}^{T^*}\hat{u}_i\left(\mathbf{a}_\tau\right)\\ +\hat{\pi}_i^\theta(x_{-i}^\theta,h^{T^*+1})\\ +\text{sign}\left(x_{i-1}^\theta\right)\bar{\varepsilon}T^*\end{array}\right)|h^{1^*},\theta,\chi_i(h^{1^*},h^{T^*+1})=1\right]-\hat{u}_i\left(\mathbf{a}^{x^\theta}\right)\right| \\
\leq \ & \Pr\left(\chi_i(h^{1^*},h^{T^*+1})=0|\theta\right)\left(\frac{(3+K)\,T}{T^*}\bar{u}+2\bar{\varepsilon}\right)+\Pr\left(\chi_i(h^{1^*},h^{T^*+1})=1|\theta\right)(2\bar{u}+\bar{\varepsilon}) \\
\leq \ & 2\bar{\varepsilon}+\left(\frac{3+K}{Z}+2\Pr\left(\chi_i(h^{1^*},h^{T^*+1})=1|\theta\right)\right)\bar{u} \\
\leq \ & 2\bar{\varepsilon}+(3+K)\left(\frac{1}{Z}+2\exp\left(-cZ\right)\right)\bar{u}\quad\text{(by Lemma 15)} \\
\leq \ & 3\bar{\varepsilon}\quad\text{(by (21))}.
\end{aligned}
\tag{32}
$$

Here the first inequality follows because (i) when $\chi_i(h^{1^*},h^{T^*+1})=0$, $\mathbf{a}_\tau=\mathbf{a}^{x^\theta}$ in main rounds (i.e., in all but $(3+K)\,T$ periods), (ii) the magnitude of $\hat{u}_i\left(\mathbf{a}_\tau\right)$ is bounded by $\frac{1}{2}\bar{u}$, and (iii) on-path (i.e., when $I_i^D(h^{1^*},h^{T^*+1})=0$), the magnitude of $\hat{\pi}_i^\theta(x_{-i}^\theta,h^{T^*+1})$ is bounded by $\bar{\varepsilon}T^*$.

We now define the reward function

$$
\begin{aligned}
& \pi_i^\theta(x_{-i}^\theta,h^{1^*},h^{T^*+1}) \\
= \ & 1_{\left\{I_i^D\left(h^{1*},h^{T^*+1}\right)=0\right\}}\left(v_i^\theta(x_{i-1}^\theta)-u_i\left(\hat{x}_i^\theta(h^{T^*+1}),x_{-i}^\theta,h^{1^*}\right)\right)T^*+\hat{\pi}_i^\theta(x_{-i}^\theta,h^{T^*+1})+\text{sign}\left(x_{i-1}^\theta\right)\bar{\varepsilon}T^*.
\end{aligned}
$$

We verify that, with this reward function, Conditions (23)–(26) are satisfied. This will complete the proof. We first establish Conditions (24)–(26), deferring Condition (23) (sequential rationality) to the next subsection.

Since $I_i^D(h^{1^*},h^{T^*+1})=0$ on path, (31) implies that expected per-period block payoffs given $|\theta|\geq\alpha N$ equal $v_i^\theta(x_{i-1}^\theta)$. Hence, (24) holds.

By (32) and (19), we have

$$
\text{sign}\left(x_{i-1}^\theta\right)\left(v_i^\theta(x_{i-1}^\theta)-u_i\left(\hat{x}_i^\theta(h^{T^*+1}),x_{-i}^\theta,h^{1^*}\right)\right)\geq 0.
$$

Together with (29) and (30), this implies

$$
\text{sign}\left(x_{i-1}^\theta\right)\pi_i^\theta(x_{-i}^\theta,h^{1^*},h^{T^*+1})\geq\bar{\varepsilon}T, \tag{33}
$$

and hence (25).

Moreover, if $I_i^D\left(h^{1^*},h^{T^*+1}\right)=0$ then

$$
\left|\pi_i^\theta(x_{-i}^\theta,h^{1^*},h^{T^*+1})\right|\leq\left|v_i^\theta(x_{i-1}^\theta)-u_i\left(\hat{x}_i^\theta(h^{T^*+1}),x_{-i}^\theta,h^{1^*}\right)\right|T^*+2\bar{\varepsilon}T^*\leq 2\bar{u}T^*;
$$

and if $I_i^D\left(h^{1^*}, h^{T^*+1}\right) = 1$ then

$$\left|\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1})\right| \leq 2\bar{u}T^*.$$

Hence, (26) holds.

B.2.9 Verifying Sequential Rationality (Condition (23))

Given (30) and $|\theta| \geq \alpha N$, if $\chi_i(h_{1^*}, h^{T^*+1}) = 1$, then any action is optimal for player i. Since $\Pr\left(\chi_i(h_{1^*}, h^{T^*+1}) = 1|\sigma_i, \theta\right)$ is independent of σ_i, it is without loss to verify sequential rationality conditional on the event $\{\chi_i(h^{1^*}, h^{T^*+1}) = 0 \vee |\theta| < \alpha N\}$. We thus restrict attention to pairs $\left(\boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right)$ such that $\Pr\left(\{\chi_i(h^{1^*}, h^{T^*+1}) = 0 \vee |\theta| < \alpha N\}|\boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right) > 0$. Note this implies that (27) and (28) hold conditional on the triple $\left(\{\chi_i(h^{1^*}, h^{T^*+1}) = 0 \vee |\theta| < \alpha N\}, \boldsymbol{x}_{-i}, \tilde{h}_i^{b,t}\right)$. We consider separately the cases $|\theta| < \alpha N$ and $\left\{\chi_i(h^{1^*}, h^{T^*+1}) = 0 \wedge |\theta| \geq \alpha N\right\}$.

Conditional on $|\theta| < \alpha N$, by (28), player i believes that $\left|\theta(h_n^{T+1})\right| < \alpha N$ for each $n \neq i$. Hence, she believes that players $-i$ take D throughout the block and $\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1}) = 0$ regardless of her own strategy. It is therefore optimal for player i to take D in each period and send any messages. And this behavior is indeed what is prescribed for player i, since (27) implies that $\left|\theta(h_i^{T+1})\right| < \alpha N$.

It remains to verify sequential rationality conditional on $\left\{\chi_i(h^{1^*}, h^{T^*+1}) = 0 \wedge |\theta| \geq \alpha N\right\}$. We proceed in three steps.

It is optimal to take D and send any message after player i detectably deviates after the 1^*-*communication sub-block.*

Let τ be the first period in which player i detectably deviated. First, suppose that τ is before supplemental round 0. Then, regardless of player i's behavior after period τ, the fact that $\chi_i(h^{1^*}, h^{T^*+1}) = 0$ (and hence matching is regular) implies that players $-i$ will become aware of player i's deviation at the end of supplemental round 0 and will then take D for the rest of the block. Moreover, the reward function is constant:

$$\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1}) = \hat{\pi}_i^\theta(x_{-i}^\theta, h^{T^*+1}) = \begin{cases} 0 & \text{if } x_{i-1}^\theta = B, \\ -2\bar{u}T^* & \text{if } x_{i-1}^\theta = G. \end{cases}$$

Hence, taking D and sending any messages is optimal for player i.

Second, suppose τ is in or after supplemental round 0. Then, regardless of player i's behavior after period τ, players $-i$ take $\mathbf{a}^{x^\theta}$ in the main sub-block and take D in other rounds until next supplemental round; and subsequently (since matching is regular) they will switch to D for the rest of the block. Again, the reward is constant. Hence, taking D and sending any messages is optimal.

It is optimal not to detectably deviate from the equilibrium strategy at on-path histories.

We compare the maximum gain in within-block payoffs from a detectable deviation to the minimum loss in the reward function. Since matching is regular, players $-i$ switch to D starting in the next main round. Hence, the maximum gain in within-block payoffs is at most $Z^2 \log_2 N \times \max\{G, L\}$. In contrast, if $x_{i-1}^\theta = B$, the loss in the reward function from switching $I_i^D\left(h^{1^*}, h^{T^*+1}\right)$ from 0 to 1 is at least $\bar{\varepsilon}T^*$; this comes from the $\hat{\pi}_i^\theta\left(x_{-i}^\theta, h^{T^*+1}\right)$ term in the reward function, noting that the term

$$1_{\left\{I_i^D(h^{1^*}, h^{T^*+1})=0\right\}}\left(v_i^\theta(x_{i-1}^\theta) - u_i\left(\hat{x}_i^\theta(h^{T^*+1}), x_{-i}^\theta, h^{1^*}\right)\right)$$

in the reward $\pi_i^\theta(x_{-i}^\theta, h^{1^*}, h^{T^*+1})$ is non-negative. By (20), $\bar{\varepsilon}T^* \geq Z^2 \log_2 N \times \max\{G, L\}$, so deviating is unprofitable when $x_{i-1}^\theta = B$. If instead $x_{i-1}^\theta = G$, the loss in the reward function from switching $I_i^D\left(h^{1^*},h^{T^*+1}\right)$ from 0 to 1 is at least

$$2\bar{u}T^* - \left|v_i^\theta(x_{i-1}^\theta) - u_i\left(\hat{x}_i^\theta(h^{T^*+1}), x_{-i}^\theta, h^{1^*}\right)\right| T^* \geq \bar{u}T^* \geq Z^2 \log_2 N \times \max\{G, L\},$$

where the first inequality follows because $\left|v_i^\theta(x_{i-1}^\theta) - u_i\left(\hat{x}_i^\theta(h^{T^*+1}), x_{-i}^\theta, h^{1^*}\right)\right| \leq \bar{u}$. In total, for any x_{i-1}^θ, the net deviation gain is negative.

It is optimal to send message $\hat{\boldsymbol{x}}_i = \boldsymbol{x}_i$ in the $\boldsymbol{x}$-communication sub-block.

We show that, for any $\boldsymbol{x}_{-i}$, player i is indifferent among the block strategies $(\sigma_i(\boldsymbol{x}_i))_{\boldsymbol{x}_i}$. Player i's expected payoff conditional on $|\theta| \geq \alpha N$ equals

$$\mathbb{E}\left[\frac{1}{T^*}\mathbb{E}^{\sigma(\boldsymbol{x})}\left[\sum_{\tau=3T+1}^{T^*} \hat{u}_i(\mathbf{a}_\tau) + \hat{\pi}_i^\theta(x_{-i}^\theta, h^{T^*+1}) + \text{sign}\left(x_{i-1}^\theta\right)\bar{\varepsilon}T^*|h^{1^*},\theta\right]|\boldsymbol{x}_{-i}\right] = v_i^\theta(x_{i-1}^\theta).$$

Moreover, her payoff conditional on $\chi_i(h^{1^*},h^{T^*+1}) = 1$ equals $\text{sign}\left(x_{i-1}^\theta\right)\left(\frac{1}{2}+\bar{\varepsilon}T^*\right)$. Since these payoffs depend on $\boldsymbol{x}$ only through x_{i-1}^θ, and additionally $\Pr\left(\chi_i(h^{1^*},h^{T^*+1}) = 1\right)$ is independent of $\boldsymbol{x}$, it follows that player i's expected payoff conditional on $\left\{\chi_i(h^{1^*},h^{T^*+1}) = 0 \wedge |\theta| \geq \alpha N\right\}$ also depends on $\boldsymbol{x}$ only through x_{i-1}^θ. This completes the proof of Theorem 4.

B.2.10 Proof of Theorem 3

Since $F^{\alpha,\eta}$ is compact, the following lemma is sufficient:

Lemma 16 *For any $\tilde{v} \in F^{\alpha,\eta}$, there exist $\rho > 0$ such that $B^\rho(\tilde{v}) \subseteq E^*$ for all sufficiently large l.*

Proof. Fix $\tilde{v} \in F^{\alpha,\eta}$ and the associated state-contingent payoffs $\left(\tilde{v}^\theta\right)_\theta$. Define $\left(v^\theta\right)_\theta$ as in (10), and then fix $\left(v_i^\theta\left(x_{i-1}^\theta\right)\right)_{i\in I,\theta\in\Theta,x_{i-1}^\theta\in\{G,B\}}$ and $\bar{\varepsilon} > 0$ to satisfy (19). The proof of Theorem 4 shows that a convex set $\mathbf{V} \subset \mathbb{R}^{N|\Theta|}$ is self-generating for sufficiently large l if it satisfies

$$\begin{aligned} V_i^\theta &\subseteq \left(v_i^\theta(B) + 4\bar{\varepsilon}, v_i^\theta(G) - 4\bar{\varepsilon}\right) \text{ for } |\theta| \geq \alpha N, \\ V_i^\theta &= \{0\} \text{ for } |\theta| < \alpha N. \end{aligned} \tag{34}$$

Define the set $\mathbf{V}^{\hat{\rho}} \subset \mathbb{R}^{N|\Theta|}$ by

$$\begin{aligned} V_i^{\theta,\hat{\rho}} &= \left[v_i^\theta - \hat{\rho}, v_i^\theta + \hat{\rho}\right] \text{ for } |\theta| \geq \alpha N, \\ V_i^\theta &= \{0\} \text{ for } |\theta| < \alpha N. \end{aligned}$$

Fix $\hat{\rho} > 0$ sufficiently small so that $\mathbf{V}^{\hat{\rho}}$ satisfies (34) and (17) for sufficiently large δ; such $\hat{\rho} > 0$ exists since $V^\theta = \{v^\theta\}$ satisfies (34) and $\left(v^\theta\right)_\theta$ satisfies (17) with strict inequality. Now fix $\bar{l}_1$ sufficiently large so that, for each $l \geq \bar{l}_1$, the set $\mathbf{V}^{\hat{\rho}}$ is self-generating. Since (17) holds, rational players take C in period 1^*. Hence, for each $\left(\hat{v}^\theta\right)_\theta \in \mathbf{V}^{\hat{\rho}}$, there exists a sequential equilibrium such that, for each $\theta \in \Theta$, the resulting payoff when $\theta^* = \theta$ equals $\hat{v}^\theta$.

Let $\hat{V}^{\hat{\rho}} \subset \mathbb{R}^N$ be the set of $\hat{v} \in \mathbb{R}^N$ such that there exists $\left(\hat{v}^\theta\right)_\theta \in \mathbf{V}^{\hat{\rho}}$ satisfying

$$\left((1-\delta)\left(p^0 + p^1\left(\frac{1+G-L}{2}\right)\right) + \delta\sum_{\theta^*} p\left(I\setminus\theta^*\right)\hat{v}_i^{\theta^*}\right)_i = \hat{v}_i.$$

(Note that this set $\hat{V}^{\hat{\rho}}$ depends on l.) Since any payoff vector in $\mathbf{V}^{\hat{\rho}}$ is implementable, so is any expected payoff in $\hat{V}^{\hat{\rho}}$. Since $V_i^{\theta,\hat{\rho}} = \left[v_i^{\theta} - \hat{\rho}, v_i^{\theta} + \hat{\rho}\right]$ for $|\theta| \geq \alpha N$, by taking $\rho > 0$ sufficiently small and $\bar{l}_2$ sufficiently large, we have $B^{\rho}(\tilde{v}) \subseteq \hat{V}^{\hat{\rho}}$ for all $l \geq \bar{l}_2$. For such $\rho > 0$, we have $B^{\rho}(\tilde{v}) \subseteq E^*$ for all $l \geq \max\left\{\bar{l}_1, \bar{l}_2\right\}$, as desired. ■