Repeated Games with Many Players*

Takuo Sugaya — Stanford GSB
Alexander Wolitzky — MIT

July 19, 2023

Abstract

Motivated by the problem of sustaining cooperation in large communities with limited information, we analyze the relationship between population size, discounting, and monitoring in repeated games with individual-level noise. We identify the ratio of the discount rate and the *per-capita channel capacity* of the monitoring structure as a key determinant of the possibility of cooperation. If this ratio is large, cooperation is impossible: all repeated-game Nash equilibrium payoffs are consistent with approximately myopic play. Conversely, if this ratio is small and the monitoring structure is given by *random monitoring*, where each player is monitored with the same probability in every period, cooperation is possible: a folk theorem holds. Moreover, if attention is restricted to *linear perfect public equilibria* (which model collective incentive-provision), cooperation is possible only under much more severe parameter restrictions.

Keywords: repeated games, large populations, individual-level noise, χ^2-divergence, mutual information, channel capacity, folk theorem, random monitoring, linear equilibrium, penalty contract

JEL codes: C72, C73

*For helpful comments, we thank Tommaso Denti, Drew Fudenberg, Stephen Morris, Yuliy Sannikov, Satoru Takahashi, and many seminar participants. Wolitzky acknowledges financial support from the NSF.

Two neighbours may agree to drain a meadow which they possess in common; because it is easy for them to know each other's mind; and each must perceive that the immediate consequence of his failing in his part is the abandoning of the whole project. But it is very difficult, and indeed impossible, that a thousand persons should agree in any such action; it being difficult for them to concert so complicated a design, and still more difficult for them to execute it; while each seeks pretext to free himself of the trouble and expense, and would lay the whole burden on others.

—David Hume, *A Treatise of Human Nature*

1 Introduction

Hume's intuition notwithstanding, large groups of individuals often have a remarkable capacity for cooperation, even in the absence of external contractual enforcement (Ostrom, 1990; Ellickson, 1991; Seabright, 2004). Cooperation in large groups usually seems to rely on accurate monitoring of individual agents' actions, together with sanctions that narrowly target deviators. These are key features of the community resource management settings documented by Ostrom (1990), as well as the local public goods provision environment studied by Miguel and Gugerty (2005), who in a development context found that parents who fell behind on their school fees and other voluntary contributions faced social sanctions.[1] Large cartels appear to operate on similar principles. For example, the Federation of Quebec Maple Syrup Producers—a government-sanctioned cartel that organizes more than 7,000 producers, accounting for over 90% of Canadian maple syrup production—strictly monitors its members' sales, and producers who violate its rules regularly have their sugar shacks searched and their syrup impounded, and can also face fines, legal action, and ultimately the seizure of their farms (Kuitenbrouwer, 2016; Edmiston and Hamilton, 2018). In contrast, we are not aware of any evidence that individual maple syrup producers—or the parents studied by Miguel and Gugerty, or the farmers, fishers, and herders studied by Ostrom—are motivated by the fear of starting a price war or other general breakdown of cooperation.

The principle that large-group cooperation requires precise monitoring and personalized

[1] Similar effects have also been found in the context of group lending (Karlan, 2007; Feigenberg, Field, and Pande, 2013).

sanctions seems like common sense, but it is not reflected in current repeated game models. The standard analysis of repeated games with patient players (e.g., Fudenberg, Levine, and Maskin, 1994; henceforth FLM) fixes all parameters of the game except the discount factor δ and considers the limit as $\delta \to 1$. This approach does not capture situations where, while players are patient ($\delta \approx 1$), they are not necessarily patient in comparison to the population size N (so $(1-\delta) N$ may or may not be close to 0). In addition, since standard results are based on statistical identification conditions that hold generically regardless of the number of players, they also do not capture the possibility that more information may be required to support cooperation in larger groups. Finally, since there is typically a vast multiplicity of cooperative equilibria in the $\delta \to 1$ limit, standard results also say little about what kind of strategies must be used to support large-group cooperation: for example, whether it is better to rely on personalized sanctions (e.g., fines) or collective ones (e.g., price wars).

This paper extends the standard analysis of repeated games with imperfect monitoring by letting the population size, discount factor, stage game, and monitoring structure all vary together. These aspects of the repeated game can vary in a general manner: we assume only a uniform upper bound on the magnitude of the players' stage-game payoffs and a uniform lower bound on the amount of "individual-level noise." Our main results provide necessary and sufficient conditions for cooperation as a function of N, δ, and a measure of the "informativeness" of the monitoring structure. We also show that cooperation is possible only under much more restrictive conditions if society exclusively relies on collective sanctions, such as price wars (or, a la Hume, "the abandoning of the whole project"). In sum, we show that large-group cooperation requires a lot of patience and/or a lot of information, and cannot be based on collective sanctions for reasonable parameter values.

We now preview our main ideas and results. We model individual-level noise by assuming that each player i's action a_i stochastically determines an *individual-level outcome* x_i, independently across players, and that the distribution of observed signals y (the *outcome monitoring structure*) depends on the action profile $a = (a_i)$ only through the outcome profile $x = (x_i)$. This setup follows earlier work by Fudenberg, Levine, and Pesendorfer (1998; henceforth FLP) and al-Najjar and Smorodinsky (2000, 2001; henceforth a-NS). We find that a useful measure of the informativeness of the outcome monitoring structure is its *channel capacity*, C. This is a standard measure in information theory, which in our context is defined as the maximum expected reduction in uncertainty (entropy) about the outcome profile x

that results from observing the signal y. Channel capacity obeys the elementary inequality $C \leq \log |Y|$, where Y is the set of possible signal realizations. Due to this inequality, using channel capacity permits more general results as compared to measuring informativeness by the number of possible signal realizations (as FLP and a-NS do). At the same time, channel capacity is convenient to work with, as it lets us apply tools from information theory such as Pinsker's inequality and the chain rule for mutual information, which play key roles in our analysis.

Our first result (Theorem 1) is that if $(1-\delta) N/C$—the ratio of the discount rate $1-\delta$ and the per-capita channel capacity C/N—is large, then cooperation is impossible: all repeated game Nash equilibrium payoffs are consistent with approximately myopic play. This result builds on a general necessary condition for cooperation in repeated games that we establish in a companion paper (Sugaya and Wolitzky, 2023a; henceforth SW). Compared to that result, the key difference is that here we consider games with individual-level noise, which allows a connection between the main information measure in SW (the χ^2-*divergence* of the signal distribution following a deviation from the equilibrium signal distribution) and the channel capacity of the outcome monitoring structure.

Our second result (Theorem 2) provides a partial converse to Theorem 1 for a specific monitoring structure: *random monitoring*, where in each period a certain number M out of the N players are chosen at random and their outcomes are perfectly revealed, while nothing is learned about the other players' outcomes. Under random monitoring, channel capacity is proportional to the number of monitored players M, and we show that if $(1-\delta) N \log(N)/M$ is small then cooperation is possible: a large set of payoffs arise as perfect equilibria in the repeated game, i.e., a folk theorem holds. This result implies that the condition on δ, N, and C in Theorem 1 is tight up to $\log(N)$ slack. Moreover, while random monitoring is admittedly special, in Appendix B we generalize Theorem 2 to a similar result that holds for any *public, product-structure monitoring* (Theorem 4).

Our final result (Theorem 3) considers the implications of restricting society to "collective" sanctions and rewards under public monitoring. We formalize this restriction by focusing on *linear perfect public equilibria*, where all on-equilibrium-path continuation payoff vectors lie on a line in $\mathbb{R}^N$. When the stage game is symmetric and the line in question is the 45° line, linear equilibria reduce to *strongly symmetric equilibria*, which are a standard model of collusion through the threat of price wars (Green and Porter, 1984; Abreu, Pearce, and Stacchetti, 1986; Athey, Bagwell, and Sanchirico, 2004). We show that if there exists $\rho > 0$ such that $(1-\delta)\exp\left(N^{1-\rho}\right)$ is large, then all equilibrium payoffs are consistent with approximately myopic play. Since this condition holds even if $N \to \infty$ much slower than $\delta \to 1$, we interpret this result as a near-impossibility theorem for large-group cooperation based on collective incentives.[2]

1.1 Related Literature

Prior research on repeated games has established folk theorems in the $\delta \to 1$ limit for fixed N, as well as "anti-folk" theorems in the $N \to \infty$ limit for fixed δ, but has not considered the case where N and δ vary together.

The closest paper is our companion work, SW. That paper establishes general necessary and sufficient conditions for cooperation in repeated games as a function of discounting and monitoring. Relative to SW, the current paper introduces two features that are specific to large-population games: individual-level noise and the possibility that N varies together with discounting and monitoring. Individual-level noise is crucial for our anti-folk theorem (Theorem 1), while letting N vary with discounting and monitoring is the key novelty in our folk theorems (Theorems 2 and 4).

The most relevant folk theorems are due to FLM, Kandori and Matsushima (1998), and SW. However, these papers fix the stage game while taking $\delta \to 1$ (and also letting monitoring vary, in the case of SW), and their proof approach does not easily extend to the case where N and δ vary together. Our proof of Theorems 2 and 4 takes a different approach, which is based on "block strategies" as in Matsushima (2004) and Hörner and Olszewski (2006), and involves a novel application of some large deviations bounds.

Other than that in SW, the most relevant anti-folk theorems are those of FLP, a-NS, Pai, Roth, and Ullman (2014), and Awaya and Krishna (2016, 2019). Following earlier work by Green (1980) and Sabourian (1990), these papers establish conditions under which play is approximately myopic as $N \to \infty$ for fixed δ.[3] These conditions can be adapted to the case

[2] It is well-known that strongly symmetric equilibria are typically less efficient than general perfect public equilibria (see, e.g., FLM). Our result is instead that the relationship between N and δ required for *any* non-trivial incentive provision differs dramatically between strongly symmetric/linear equilibria and general ones.

[3] Awaya and Krishna instead establish conditions under which cheap talk is valuable. Green and Sabourian's papers impose a continuity condition on the mapping from action distributions to signal distributions. Continuity is implied by FLP/a-NS's individual noise assumption.

where N, δ, and monitoring vary together, but the results so obtained are weaker than ours, and are not tight up to log terms. The key difference is that these results rely on bounds on the strength of players' incentives that have a worse order in $1-\delta$ than that given in SW. In sum, the earlier literature established anti-folk theorems as $N \to \infty$ for fixed δ, while our paper tightly (up to log terms) characterizes the tradeoff among N, δ, and monitoring.[4]

Since the monitoring structure varies with δ in our model, we also relate to repeated games with frequent actions, where the monitoring structure varies with δ in a particular, parametric manner (e.g., Abreu, Milgrom, and Pearce, 1991; Fudenberg and Levine, 2007, 2009; Sannikov and Skrzypacz, 2007, 2010). The most relevant results here are Sannikov and Skrzypacz's (2007) theorem on the impossibility of collusion with frequent actions and Brownian noise, as well as a related result by Fudenberg and Levine (2007). These results relate to our anti-folk theorem for linear equilibrium, as we explain in Section 5.[5]

Finally, in Sugaya and Wolitzky (2021) we studied the relationship among N, δ, and monitoring in repeated random-matching games with private monitoring and incomplete information, where each player is "bad" (i.e., a *Defect* commitment type) with some probability. In that model, society has enough information to determine which players are bad after a single period of play, but this information is disaggregated, and supporting cooperation requires sufficiently quick information diffusion. In contrast, in the current paper there is complete information and monitoring can be public, so the analysis concerns monitoring precision (the "amount" of information available to society) rather than the speed of information diffusion (the "distribution" of information). In general, whether the key obstacle to cooperation is that societal information is *insufficient* or *disaggregated* distinguishes "large-population repeated games," such as FLP, a-NS, and the current paper, from "community enforcement" models, such as Kandori (1992), Ellison (1994), and our earlier paper.

2 Model

We consider a general model of repeated games with individual-level noise and imperfect monitoring.

[4] Farther afield, there is also work suggesting that repeated-game cooperation is harder to sustain in larger groups based on evolutionary models (e.g., Boyd and Richerson, 1988), simulations (e.g., Bowles and Gintis, 2011; Chapter 4), and experiments (e.g., Camera, Casari, and Bigoni, 2013).

[5] Another somewhat related question is the rate of convergence of the equilibrium payoff set as $\delta \to 1$ (Hörner and Takahashi, 2016; Sugaya and Wolitzky, 2023b).

Stage Games. A stage game $G = (I, A, u)$ consists of a finite set of players $I = \{1, \ldots, N\}$, a finite product set of actions $A = \times_{i \in I} A_i$, and a payoff function $u_i : A \to \mathbb{R}$ for each $i \in I$. The interpretation is that $u_i(a)$ is player i's expected payoff at action profile a. We denote the range of player i's payoff function by $\bar{u}_i = \max_{a,a'} u_i(a) - u_i(a')$.

Noise. There is a finite product set of *individual outcomes* $X = \times_{i \in I} X_i$ and a row-stochastic *noise matrix* $\pi^i \in [0,1]^{A_i \times X_i}$ for each player i such that, when action profile $a \in A$ is played, outcome profile $x \in X$ is realized with probability $\pi_{a,x} = \prod_i \pi^i_{a_i,x_i}$. We call the pair (X, π) a *noise structure*. Let $\underline{\pi}^i = \min_{a_i,x_i} \pi^i_{a_i,x_i}$ and assume that $\underline{\pi}^i > 0$ for each i: we call this assumption *individual-level noise*. The point of this setup is that signals will depend on a only through x.[6]

For a natural example of a noise structure, suppose that there is some independent noise in the execution of the players' actions, so that a_i is player i's intended action and x_i is her realized action. In this case, $X = A$, and $\pi_{a_i,a_i'}$ is the probability that player i "trembles" to a_i' when she intends to take a_i. We refer to this example as *noisy actions*.

Monitoring. An *outcome monitoring structure* (Y, q) consists of a finite product set of signal profiles $Y = \times_{i \in I} Y^i$ and a family of conditional probability distributions $q(y|x)$. The distribution of signal profiles thus depends only on the realized outcome profile. The outcome monitoring structure (Y, q) is a primitive object in our model: we are interested in properties of (Y, q) that (together with the other model primitives) are necessary or sufficient for supporting cooperative outcomes.

Given an outcome monitoring structure (Y, q), we denote the probability of signal profile y at action profile a by $p(y|a) = \sum_x \pi_{a,x} q(y|x)$. We refer to the pair (Y, p) as the *action monitoring structure induced by* (X, π, Y, q). The action monitoring structure (Y, p) is a derived object in our model: it plays an important role in our analysis, but we will avoid imposing assumptions directly on (Y, p), and instead consider the implications of properties of the noise structure (X, π) and the outcome monitoring structure (Y, q) for (Y, p).[7]

[6] Since we will assume that players do not observe their own payoffs in addition to their signals, it is natural to require that players' realized payoffs are determined by their signals, and hence depend on a only through x. However, this assumption is not necessary for our analysis.

[7] Our companion paper, SW, allows general monitoring structures and directly considers properties of the action monitoring structure (Y, p). The current paper imposes the additional structure that (Y, p) factors into a noise structure (X, π) and an outcome monitoring structure (Y, q). This additional structure lets us formulate the individual-level noise assumption.

Let $\bar{Y} \subseteq Y$ denote the set of signal profiles y such that there exists an outcome profile x at which $q(y|x) > 0$. Since $\underline{\pi}^i > 0$ for each i, at any action profile a, signals are supported on $\bar{Y}$: $p(y|a) > 0$ iff $y \in \bar{Y}$.

Mutual Information and Channel Capacity. Given a distribution of outcomes $\xi \in \Delta(X)$, a standard measure of the informativeness of a signal y about the realized outcome x is the *mutual information* between x and y, defined as

$$\mathbf{I}(\xi) = \sum_{x \in X, y \in \bar{Y}} \xi(x)\, q(y|x) \log\left(\frac{q(y|x)}{\sum_{x' \in X} \xi(x')\, q(y|x')}\right).^{8}$$

Mutual information measures the expected reduction in uncertainty (entropy) about x that results from observing y. The mutual information between x and y is an endogenous object in our model, as it depends on the distribution ξ of x, which in turn is determined by the players' actions, a. Next, denote the set of outcome distributions ξ that can arise for some action distribution α under noise structure (X, π) by

$$\vartheta = \left\{\xi \in \Delta(X) : \exists \alpha \in \Delta(A) \text{ such that } \zeta(x) = \sum_{a \in A} \alpha(a)\, \pi_{a,x} \text{ for all } x \in X\right\}.$$

Finally, define the *channel capacity* of the tuple (X, π, Y, q) as

$$C = \max_{\xi \in \vartheta} \mathbf{I}(\xi).$$

Channel capacity is an exogenous measure of the informativeness of y about x, as it is defined as a function of only the noise structure (X, π) and the outcome monitoring structure (Y, q).[9] Note that C is no greater than the entropy of the signal y, which in turn is at most $\log|Y|$ (Theorem 2.6.3 of Cover and Thomas, 2006; henceforth CT). Channel capacity plays a central role in information theory, because it is the maximum rate at which information can be transmitted over a noisy channel (*Shannon's channel coding theorem*, CT Theorem 7.7.1). Our analysis does not use this theorem; we only use channel capacity as an exogenous upper bound on mutual information. In turn, mutual information arises in our analysis because

[8] In this paper, all logarithms are base e.

[9] We define C as the maximum of $\mathbf{I}(\zeta)$ over $\zeta \in \vartheta$ rather than $\zeta \in \Delta(X)$, because only $\zeta \in \vartheta$ can ever arise. This definition makes Theorem 1 (the only result stated in terms of C) stronger than it would be if we instead took the maximum over all $\zeta \in \Delta(X)$.

it obeys useful properties, in particular the *chain rule* (CT, Theorem 2.5.2) and *Pinsker's inequality* (CT, Lemma 11.6.1). These properties play key roles in the proof of Theorem 1.[10]

Some Special Monitoring Structures. Some of our results will assume that monitoring is public and has a product structure. A monitoring structure (Y, q) is *public* if all players observe the same signal: $y^i = y^j$ for all $i, j \in I, y \in Y$. In this case, we ease notation by identifying the public signal with any one player's signal. A public monitoring structure (Y, q) has a *product structure* if there exists sets $(Y_i)_{i\in I}$ and a family of conditional distributions $(q_i (y_i|x_i))_{i\in I,,y_i\in Y_i,x_i\in X_i}$ such that $Y = \prod_i Y_i$ and $q(y|x) = \prod_i q_i (y_i|x_i)$ for all y, x: that is, the public signal y consists of conditionally independent signals of each player's individual outcome.[11] Note that if (Y, q) is public and has a product structure, then so does (Y, p), meaning that there exists a family of conditional distributions $(p_i (y_i|a_i))_{i\in I,,y_i\in Y_i,a_i\in A_i}$ (given by $p_i (y_i|a_i) = \sum_{x_i} \pi^i_{a_i,x_i} q_i (y_i|x_i)$) such that $p(y|a) = \prod_i p_i (y_i|a_i)$ for all y, a.

A particular public, product monitoring structure is *random monitoring.* Under random monitoring, at the end of every period a certain number M of players are selected uniformly at random, and the public signal perfectly reveals their identities and their realized individual outcomes. That is, under random monitoring of M players, $Y_i = X_i \cup \{\emptyset\}$ for all i, and

$$q_i (y_i|x_i) = \begin{cases} \frac{M}{N} & \text{if } y_i = x_i, \\ 0 & \text{if } y_i \in X_i \backslash \{x_i\}, \\ 1 - \frac{M}{N} & \text{if } y_i = \emptyset. \end{cases}$$

Note that the channel capacity of random monitoring is no more than $M \log (\max_i |X_i|)$.

Repeated Games. A *repeated game with individual-level noise* $\Gamma = (G, X, \pi, Y, q, \delta)$ is described by a stage game, a noise structure, an outcome monitoring structure, and a discount factor $\delta \in [0, 1)$. In each period $t = 1, 2, \ldots$, (i) the players observe the outcome of a public randomizing device z_t drawn from the uniform distribution over $[0, 1]$, (ii) the players take actions a, (iii) the outcome x is drawn according to $\pi_{a,x}$, (iv) the signal y is

[10]We are not the first to recognize the value of entropy methods in repeated games. These methods have previously been used to study issues including complexity and bounded recall (Neyman and Okada, 1999, 2000; Hellman and Peretz, 2020), communication (Gossner, Hernández, and Neyman, 2006), and reputation effects (Gossner, 2011; Ekmekci, Gossner, and Wilson, 2011; Faingold, 2020). However, other than sharing a reliance on entropy methods, our results and proofs are not very related to these papers'.

[11]Our notation is thus that Y^i denotes the set of possible signals observed by player i (for any monitoring structure), while Y_i denotes the set of public signals of player i's individual outcome (for public, product structure monitoring).

drawn according to $q(y|x)$, and (v) each player i observes y^i.[12] A *history* h_i^t for player i at the beginning of period t thus takes the form $h_i^t = \left((z_{t'}, a_{i,t'}, y_{t'}^i)_{t'=1}^{t-1}, z_t \right)$, while a *strategy* σ_i for player i maps histories h_i^t to distributions over actions $a_{i,t}$. A *repeated game outcome* $\mu \in \Delta((A \times X \times Y)^\infty)$ (not to be confused with a single profile of individual outcomes x) is a distribution over infinite paths of actions, individual outcomes, and signals. Players maximize discounted expected payoffs with discount factor δ.

Under public monitoring, the *public history* h^t at the beginning of period t takes the form $h^t = \left((z_{t'}, y_{t'})_{t'=1}^{t-1}, z_t \right)$. A strategy σ_i for player i is *public* if it depends on player i's history h_i^t only through its public component h^t. A *perfect public equilibrium* (PPE) is a profile of public strategies that, beginning at any period t and any public history h^t, forms a Nash equilibrium from that period on.[13] The set of PPE payoff vectors is denoted by $E \subseteq \mathbb{R}^N$.

For any $\bar{u} > 0$ and $\underline{\pi} > 0$, we say that a repeated game Γ is $(\bar{u}, \underline{\pi})$*-bounded* if the range of stage-game payoffs is bounded above by $\bar{u}$ and individual-level noise is bounded below by $\underline{\pi}$: that is, if $\bar{u}_i \leq \bar{u}$ and $\underline{\pi}_i \geq \underline{\pi}$ for all i. Note that if Γ is $(\bar{u}, \underline{\pi})$-bounded then $|X_i| \leq 1/\underline{\pi}$ for all i.

Target Payoffs. Finally, we define some relevant sets of payoff vectors. The feasible payoff set is $F = \text{co}\left\{ \{u(a)\}_{a \in A} \right\} \subseteq \mathbb{R}^N$ (where co denotes convex hull). Let $F^* \subseteq F$ denote the set of payoff vectors that weakly Pareto-dominate a payoff vector which is a convex combination of static Nash payoffs: that is, $v \in F^*$ if $v \in F$ and there exists a collection of static Nash equilibria (α_n) and non-negative weights (β_n) such that $v \geq \sum_n \beta_n u(\alpha_n)$ and $\sum_n \beta_n = 1$.[14] For each $v \in \mathbb{R}^N$ and $\varepsilon > 0$, let $B_v(\varepsilon) = \prod_i [v_i - \varepsilon, v_i + \varepsilon]$ and let $B(\varepsilon) = \left\{ v \in \mathbb{R}^N : B_v(\varepsilon) \subseteq F^* \right\}$. That is, $B(\varepsilon)$ is the set of payoff vectors $v \in \mathbb{R}^N$ such that the cube with center v and side-length 2ε lies entirely within F^*.

Next, a *manipulation* for a player i is a mapping $s_i : A_i \to \Delta(A_i)$. The interpretation is that when player i is "supposed to play" a_i, she instead plays $s_i(a_i)$. Player i's *gain* from manipulation s_i at a (possibly correlated) action profile distribution $\alpha \in \Delta(A)$ is

$$g_i(s_i, \alpha) = \sum_a \alpha(a) \left(u_i(s_i(a_i), a_{-i}) - u_i(a) \right).$$

[12] Our analysis and results are unchanged if each player i also observes her own individual outcome x_i.

[13] As usual, this definition allows players to consider deviations to arbitrary, non-public strategies; however, such deviations are irrelevant because, whenever a player's opponents use public strategies, she has a public strategy as a best response.

[14] Here and throughout, we linearly extend payoff functions to mixed actions.

Player i's *maximum gain* at $\alpha \in \Delta(A)$ is $\bar{g}_i(\alpha) = \max_{s_i : A_i \to \Delta(A_i)} g_i(s_i, \alpha)$. For any $\varepsilon > 0$, the set of *action distributions consistent with ε-myopic play* is

$$A(\varepsilon) = \left\{ \alpha \in \Delta(A) : \frac{1}{N} \sum_i \bar{g}_i(\alpha) \leq \varepsilon \right\},$$

and the set of *payoff vectors consistent with ε-myopic play* is

$$W(\varepsilon) = \left\{ v \in \mathbb{R}^N : v = u(\alpha) \text{ for some } \alpha \in A(\varepsilon) \right\}.$$

That is, an action distribution α is consistent with ε-myopic play if the per-player average deviation gain at α is less than ε. Note that $\bar{g}_i(\alpha)$ is convex as the maximum of affine functions, and hence $A(\varepsilon)$ and $W(\varepsilon)$ are convex sets.

Our anti-folk theorem will provide conditions under which all repeated game equilibrium payoff vectors are contained in the set $W(\varepsilon)$, while our folk theorem will provide conditions under which all payoff vectors in the set $B(\varepsilon)$ can be attained as repeated game equilibria. These results are interesting insofar as $W(\varepsilon)$ is "small" and $B(\varepsilon)$ is "large." As a check that $B(\varepsilon)$ is indeed reasonably large, in Appendix A we consider a canonical public-goods game where each player chooses *Contribute* or *Don't Contribute*, and a player's payoff is the fraction of players who contribute less a constant $c \in (0, 1)$ (independent of N) if she contributes herself. In this game, we show that for every $v \in (0, 1 - c)$ there exists $\varepsilon > 0$ such that the symmetric payoff vector where all players receive payoff v lies in $B(\varepsilon)$, for all N. We discuss $W(\varepsilon)$ in Section 6, following our main results.

3 Non-Cooperation under Insufficient Monitoring

3.1 Anti-Folk Theorem

Our first result is that whenever per-capita channel capacity is much smaller than the discount rate, all Nash equilibrium payoff vectors are consistent with approximately myopic play. Cooperation in large groups thus requires a lot of information or a low discount rate.

Theorem 1 *Fix any $\bar{u} > 0$ and $\underline{\pi} > 0$. For any $\varepsilon > 0$, there exists $k > 0$ such that, in any*

$(\bar{u}, \underline{\pi})$-bounded repeated game where

$$\frac{(1-\delta)\,N}{C} > k,$$

all Nash equilibrium payoff vectors are consistent with ε-myopic play.

We emphasize that Theorem 1 holds for all monitoring structures (including private monitoring) and all Nash equilibria (so no equilibrium refinement is needed).

When N is large, the implied necessary condition for cooperation—that $(1-\delta)\,N/C$ is not too large—is easier to satisfy in some classes of repeated games than in others. For example, if the space of possible signal realizations Y is fixed independently of N, then, since $C \leq \log |Y|$, the necessary condition implies that δ must converge to 1 at least as fast as $N \to \infty$, which is a restrictive condition. This negative conclusion applies for traditional applications of repeated games with public monitoring where the signal space is fixed independent of N, such as when the public signal is the market price facing Cournot competitors, the level of pollution in a common water source, the output of team production, or some other aggregate statistic.

However, in other types of games C naturally scales linearly with N, so that $(1-\delta)\,N/C$ is small whenever players are patient (regardless of the population size). In repeated games with random matching (Kandori, 1992; Ellison, 1994; Deb, Sugaya, and Wolitzky, 2020), players match in pairs each period and $y_t^i = a_{m(i,t),t}$, where $m(i,t) \in I \setminus \{i\}$ denotes player i's period-t partner. In these games, $C = N \log |A_i|$, so per-capita channel capacity is independent of N. Intuitively, in random matching games each player gets a distinct signal of the overall action profile, so the total amount of information available to society is proportional to the population size. Channel capacity also scales linearly with N in public-monitoring games where the public signal is a vector that includes a distinct signal of each player's action, as in the ratings systems used by websites like eBay and AirBnB. In general, C/N may be constant in games where players are monitored "separately," rather than being monitored jointly through an aggregate statistic.

Remark 1 *In applications like Cournot competition, pollution, or team production, the signal space may be modeled as a continuum, in which case the constraint $C \leq \log |Y|$ is vacuous. However, our results extend to the case where Y is a compact metric space and there exists another compact metric space Z and a function $f^N : X^N \to Z$ (which can vary* *with N) such that the signal distribution admits a conditional density of the form $q_{Y|Z}(y|z)$, where Y, Z, and $q_{Y|Z}$ are fixed independent of N. (For example, in Cournot competition z is industry output and y is the market price, which depends on z and a noise term with variance fixed independent of N.) In this case,*

$$C = \max_{\xi \in \vartheta} \int_{y \in \bar{Y}} \sum_{x \in X} \xi(x)\, q_{Y|Z}\left(y|f^N(x)\right) \log\left(\frac{q_{Y|Z}\left(y|f^N(x)\right)}{\sum_{x' \in X} \xi(x')\, q_{Y|Z}\left(y|f^N(x')\right)}\right),$$

which is bounded by

$$\bar{C} = \max_{q_Z \in \Delta(Z)} \int_{y \in \bar{Y}} \int_{z \in Z} q_Z(z)\, q_{Y|Z}(y|z) \log\left(\frac{q_{Y|Z}(y|z)}{\int_{z' \in Z} q_Z(z')\, q_{Y|Z}(y|z')}\right).$$

Since $\bar{C}$ is independent of N, it follows that C is bounded independent of N.

Remark 2 *Prior results by FLP, a-NS, and Pai, Roth, and Ullman (2014) establish anti-folk theorems as $N \to \infty$ for fixed δ. If we fix a noise structure and a product monitoring structure and let N and δ vary together (with the same noise and monitoring structure for each player), the arguments in these papers could be used to show that cooperation is impossible if $(1-\delta)^2 N \to \infty$. In contrast, Theorem 1 implies the stronger result that cooperation is impossible if $(1-\delta) N \to \infty$. As we will explain, the improvement comes from applying results from SW together with entropy methods. Moreover, Theorem 2 will imply that the relationship between $1-\delta$ and N in Theorem 1 is tight up to* log *terms.*

The remainder of this section proves Theorem 1. The proof proceeds in three steps, each of which is fairly straightforward given the right definitions and prior results. First, we define a measure of the *detectability* of a manipulation under the induced action monitoring structure (Y, p), and we show that every Nash equilibrium payoff vector is attained by an action distribution where the average deviation gain is bounded in terms of the ratio of the average detectability and the discount rate (Lemma 1). This lemma is an extension of Theorem 1 of SW. Second, we show that with individual-level noise, detectability for each player is bounded by the mutual information between that player's individual outcome x_i and the signal profile y (Lemma 2). This lemma follows from the Cauchy-Schwarz and Pinsker inequalities. Third, we show that with individual-level noise, the average across players i of the mutual information between x_i and y is bounded by the channel capacity C (Lemma 3).

This lemma follows from the chain rule for mutual information and the definition of channel capacity. Combining the three lemmas delivers the theorem.

3.2 Bounding Incentives by Detectability

Our first lemma requires some terminology. First, define the *detectability* of a manipulation s_i at an action profile distribution α as

$$\chi_i^2\left(s_i,\alpha\right)=\sum_{a\in A,y\in\bar{Y}}\alpha\left(a\right)p\left(y|a\right)\left(\frac{p\left(y|a\right)-p\left(y|s_i\left(a_i\right),a_{-i}\right)}{p\left(y|a\right)}\right)^2. \tag{1}$$

Note that detectability is a function of the induced action monitoring structure (Y,p). Indeed, when $\alpha\left(a\right)=1$ for some action profile a, detectability is the χ^2-*divergence* of the manipulated signal distribution $p\left(\cdot|s_i\left(a_i\right),a_{-i}\right)$ from the prescribed distribution $p\left(\cdot|a\right)$. The χ^2-divergence is a standard measure of statistical distance. Note that it is well-defined because p has full support on $\bar{Y}$.

Second, denote the variance of player i's payoff under an action profile distribution $\alpha\in\Delta\left(A\right)$ by $V_i\left(\alpha\right)=\mathrm{Var}_{a\sim\alpha}\left(u_i\left(a\right)\right)$. For any set of players $J\subseteq I$, action profile distribution $\alpha\in\Delta\left(A\right)$, and profile of manipulations $s_J=\left(s_i\right)_{i\in J}$ for players $i\in J$, we also define "group average" versions of the deviation gain g_i, detectability χ_i^2, and payoff variance V_i, by

$$g_J\left(s_J,\alpha\right)=\frac{1}{|J|}\sum_{i\in J}g_i\left(s_i,\alpha\right),\quad \chi_J^2\left(s_J,\alpha\right)=\frac{1}{|J|}\sum_{i\in J}\chi_i^2\left(s_i,\alpha\right),\quad\text{and}\quad V_J\left(\alpha\right)=\frac{1}{|J|}\sum_{i\in J}V_i\left(\alpha\right).$$

Third, given a repeated game outcome $\mu\in\Delta\left(\left(A\times X\times Y\right)^\infty\right)$, let $\alpha_t^\mu\in\Delta\left(A\right)$ denote the marginal distribution of period-t action profiles under μ, and define $\alpha^\mu\in\Delta\left(A\right)$, the *occupation measure over action profiles induced by* μ, by

$$\alpha^\mu\left(a\right)=\left(1-\delta\right)\sum_{t=1}^\infty\delta^{t-1}\alpha_t^\mu\left(a\right)\qquad\text{for each } a\in A.$$

Note that the payoff vector under repeated game outcome μ equals

$$\left(1-\delta\right)\sum_{t=1}^\infty\delta^{t-1}\sum_{a\in A}\alpha_t^\mu\left(a\right)u\left(a\right)=\sum_{a\in A}\left(1-\delta\right)\sum_{t=1}^\infty\delta^{t-1}\alpha_t^\mu\left(a\right)u\left(a\right)=\sum_{a\in A}\alpha^\mu\left(a\right)u\left(a\right)=u\left(\alpha^\mu\right). \tag{2}$$

The occupation measure is thus a sufficient statistic for the players' payoffs.

Now we can state our first lemma, which bounds the players' gains from manipulations at an equilibrium outcome as a function of the discount factor, the detectability of the manipulations, and the on-path variance of the players' payoffs, where deviation gain, detectability, and variance are all evaluated at the equilibrium occupation measure.

Lemma 1 *For any Nash equilibrium outcome μ, any set of players J, and any profile of manipulations s_J, we have*

$$g_J\left(s_J,\alpha^\mu\right) \leq \sqrt{\frac{\delta}{1-\delta}\chi_J^2\left(s_J,\alpha^\mu\right)V_J\left(\alpha^\mu\right)}. \tag{3}$$

In particular, any Nash equilibrium payoff vector is consistent with ε-myopic play, where

$$\varepsilon = \max_{s_I,a}\sqrt{\frac{\delta}{1-\delta}\chi_I^2\left(s_I,a\right)\bar{u}^2}.$$

Proof. The special case of Lemma 1 where J is required to be a singleton is Theorem 1 of SW.[15] The result for general J follows as a corollary, because if $g_i\left(s_i,\alpha^\mu\right) \leq \sqrt{\left(\delta/\left(1-\delta\right)\right)\chi_i^2\left(s_i,\alpha^\mu\right)V_i\left(\alpha^\mu\right)}$ for each $i \in I$, then by Cauchy-Schwarz, for any $J \subseteq I$,

$$\begin{aligned} g_J\left(s_J,\alpha^\mu\right) &= \frac{1}{|J|}\sum_{i\in J} g_i\left(s_i,\alpha^\mu\right) \leq \frac{1}{|J|}\sum_{i\in J}\sqrt{\frac{\delta}{1-\delta}\chi_i^2\left(s_i,\alpha^\mu\right)V_i\left(\alpha^\mu\right)} \\ &\leq \frac{1}{|J|}\sqrt{\frac{\delta}{1-\delta}\sum_{i\in J}\chi_i^2\left(s_i,\alpha^\mu\right)\sum_{i\in J}V_i\left(\alpha^\mu\right)} = \sqrt{\frac{\delta}{1-\delta}\chi_J^2\left(s_J,\alpha^\mu\right)V_J\left(\alpha^\mu\right)}. \end{aligned}$$

■

The logic behind the singleton case of Lemma 1 is discussed at length in SW. Briefly, the bound, (3), comes from considering a player's incentive to manipulate according to s_i in each period, and decomposing the variance of continuation payoffs across periods. An important feature is that (3) is of order $(1-\delta)^{-1/2}$, whereas simpler bounds that consider incentives in a single period only are naturally of order $(1-\delta)^{-1}$. This difference is a key reason why Theorem 1 gives a tighter relationship between $1-\delta$ and N, as compared to prior results. But it is not the whole story, because Lemma 1 relates $1-\delta$ and the maximum average

[15] Theorem 1 of SW is actually more general, in that it applies for any Nash equilibrium outcome in a "blind" variant of the repeated game Γ, which has more equilibria than Γ itself. Theorem 1 in the current paper likewise extends to the "blind game." SW also do not assume individual-level noise.

detectability $\max_{s_I,a} \chi_I^2(s_I, a)$, which in general is not tightly related to N. However, as we now show, average detectability can be bounded by per-capita channel capacity in games with individual-level noise.

3.3 Bounding Detectability by Channel Capacity

We start with an intermediate lemma, which follows from standard inequalities.

Lemma 2 *For any player i, any manipulation s_i, and any action profile a, we have*

$$\chi_i^2(s_i, a) \leq \frac{4\mathbf{I}^a(x_i; y)}{\underline{\pi}^2}, \tag{4}$$

where $\mathbf{I}^a(x_i; y)$ denotes the mutual information between x_i and y when action profile a is played.

Proof. For any $a \in A$, let $\Pr^a$ denote the resulting probability distribution over (X, Y). For any $x_i \in X_i$, $y \in Y$, and $a \in A$, we have $\Pr^a(x_i, y) = \pi_{a_i,x_i} \Pr^a(y|x_i) = p(y|a) \Pr^a(x_i|y)$. Hence, since $\pi_{a_i,x_i} \geq \underline{\pi}$, we have

$$(\Pr^a(y|x_i) - p(y|a))^2 = \left(\frac{p(y|a)}{\pi_{a_i,x_i}}(\Pr^a(x_i|y) - \pi_{a_i,x_i})\right)^2 \leq \left(\frac{p(y|a)}{\underline{\pi}}(\Pr^a(x_i|y) - \pi_{a_i,x_i})\right)^2. \tag{5}$$

For any player i, manipulation s_i, and action profile a, we thus have

$$\begin{aligned}
\chi_i^2(s_i, a) &= \sum_{y(\in \bar{Y})} \frac{(p(y|a) - p(y|s_i(a_i), a_{-i}))^2}{p(y|a)} = \sum_y \frac{\left(\sum_{x_i}\left(\pi_{a_i,x_i} - \pi_{s_i(a_i),x_i}\right)\Pr^a(y|x_i)\right)^2}{p(y|a)} \\
&= \sum_y \frac{\left(\sum_{x_i}\left(\pi_{a_i,x_i} - \pi_{s_i(a_i),x_i}\right)(\Pr^a(y|x_i) - p(y|a))\right)^2}{p(y|a)} \\
&\leq \sum_{x_i}\left(\pi_{a_i,x_i} - \pi_{s_i(a_i),x_i}\right)^2 \sum_y \frac{\sum_{x_i}(\Pr^a(y|x_i) - p(y|a))^2}{p(y|a)} \\
&\leq \frac{2}{\underline{\pi}^2}\sum_y p(y|a)\sum_{x_i}(\Pr^a(x_i|y) - \pi_{a_i,x_i})^2 \leq \frac{2}{\underline{\pi}^2}\sum_y p(y|a)\left(\sum_{x_i}|\Pr^a(x_i|y) - \pi_{a_i,x_i}|\right)^2 \\
&\leq \frac{4}{\underline{\pi}^2}\sum_y p(y|a)\sum_{x_i}\Pr^a(x_i|y)\log\left(\frac{\Pr^a(x_i|y)}{\pi_{a_i,x_i}}\right) = \frac{4\mathbf{I}(x_i; y|a)}{\underline{\pi}^2},
\end{aligned}$$

where the first inequality follows by Cauchy-Schwarz, the second follows by (5) and $\sum_{x_i}\left(\pi_{a_i,x_i}-\pi_{a_i',x_i}\right)\leq 2$, the third is immediate, and the fourth follows by Pinsker's inequality (CT, Lemma 11.6.1). ■

We can now use Lemma 2, the chain rule for mutual information, and the independence of individual-level noise to bound average detectability by per-capita channel capacity.

Lemma 3 *For any set of players J, any profile of manipulations s_J, and any action profile a, we have*

$$\chi_J^2\left(s_J,a\right)\leq\frac{4C}{\underline{\pi}^2\left|J\right|}. \tag{6}$$

Proof. By Lemma 2, for any set of players $J\subseteq I$, any profile of manipulations s_J, and any action profile $a\in A$, we have

$$\chi_J^2\left(s_J,a\right)=\frac{1}{|J|}\sum_{i\in J}\chi_i^2\left(s_i,a\right)\leq\frac{4}{\underline{\pi}^2\left|J\right|}\sum_{i\in J}\mathbf{I}\left(x_i;y|a\right)=\frac{4}{\underline{\pi}^2\left|J\right|}\mathbf{I}\left(x_J;y|a\right),$$

where the last equality follows by the chain rule for mutual information (CT, Theorem 2.5.2), because $\left(x_i\right)_{i\in J}$ are independent conditional on a. The proof is completed by showing that $\mathbf{I}\left(x_J;y|a\right)\leq C$. This holds because $\mathbf{I}\left(x_J;y|a\right)=\mathbf{I}\left(x;y|a\right)-\mathbf{I}\left(x_{I\setminus J};y|a,x_J\right)\leq\mathbf{I}\left(x;y|a\right)\leq C$, where the equality follows by the chain rule, the first inequality follows because mutual information is non-negative (CT, Theorem 2.6.3), and the second inequality follows by the definition of channel capacity, because the distribution of x given a lies in ϑ . ■

Theorem 1 now follows immediately from Lemmas 1 and 3.

Proof of Theorem 1. By Lemmas 1 and 3, all repeated game Nash equilibrium payoff vectors are consistent with ε-myopic play, where

$$\varepsilon=\sqrt{\frac{\delta}{1-\delta}\times\frac{4C}{\underline{\pi}^2N}\times\bar{u}^2}.$$

For any fixed $\bar{u},\underline{\pi}>0$, taking $\left(1-\delta\right)N/C$ sufficiently large makes ε as small as desired. ■

Without individual-level noise, detectability under (Y,p) cannot be bounded by the channel capacity of (Y,q), and Theorem 1 fails. For example, suppose that the stage game is an N-player prisoner's dilemma with a binary public signal y, where $y=0$ if *every* player cooperates, and $y=1$ if *any* player defects. Obviously, mutual cooperation is a repeated game equilibrium outcome for a moderate discount factor, independent of N: under grim trigger strategies where the signal $y = 1$ triggers permanent mutual defection, each player's incentives in this game are the same as in a 2-player prisoner's dilemma with perfect monitoring. This observation is consistent with Lemma 1 because detectability is infinite in this example: when the other players cooperate, a deviation to defection is perfectly detectable. However, channel capacity in this example is $\log 2$, so detectability is infinitely greater than channel capacity. Thus, without individual noise, a monitoring structure can detect deviations (and support strong incentives) even if it not very "informative" in terms of channel capacity. In contrast, Theorem 1 shows that with individual noise, only informative signals can support strong incentives.

4 Cooperation under Random Monitoring

4.1 Folk Theorem

Our second result is a folk theorem for repeated games with random monitoring, which lets the discount factor, the noise structure, the number of monitored players M, and the stage game (including the number of players) vary simultaneously. This result implies that the relationship among N, δ, and C in Theorem 1 is tight (up to a $\log(N)$ term).

We require that individual-level noise is not too extreme. Specifically, define the *maximum detectability* of a noise structure (X, π) as

$$\Delta = \sup\left\{\tilde{\Delta} : \sum_{x_i:\pi_{a_i,x_i}\geq\tilde{\Delta}} \pi_{a_i,x_i}\left(\frac{\pi_{a_i,x_i} - \pi_{a'_i,x_i}}{\pi_{a_i,x_i}}\right)^2 \geq \tilde{\Delta} \quad \text{for all } i \in I, a_i \neq a'_i \in A_i\right\}.$$

This quantity is equal to the maximum detectability $\max_{i,s_i,\alpha} \chi^2_i(s_i, \alpha)$ of the action monitoring structure (Y, p) induced by the noise structure (X, π) together with perfect monitoring of outcomes (i.e., $q(y|x) = \mathbf{1}\{y = x\}$), when we ignore outcomes that occur with probability less than Δ.

For example, with noisy actions (i.e., $X = A$), maximum detectability satisfies

$$\Delta > \min_{i,a_i\neq a'_i} \pi_{a_i,a_i} - 2\pi_{a'_i,a_i},$$

and is thus close to 1 when the "tremble probability" $1-\pi_{a_i,a_i}$ is close to 0 for all i and a_i.[16]

Our folk theorem for random monitoring is as follows.

Theorem 2 *Fix any $\bar{u}>0$ and $\Delta>0$. For any $\varepsilon>0$, there exists $k>0$ such that, in any $(\bar{u},0)$-bounded repeated game with random monitoring of M players and a noise structure with maximum detectability Δ, where*

$$\frac{(1-\delta)\,N\log(N)}{M\Delta}<k,$$

we have $B(\varepsilon)\subseteq E$.

Theorem 2 implies that the relationship among N, δ, and C in Theorem 1 is tight up to a $\log(N)$ term. To see this, note that in a $(\bar{u},\underline{\pi})$-bounded game, random monitoring of M players has a channel capacity of at most $M\log(1/\underline{\pi})$. Thus, under random monitoring of M players with a noise structure with any fixed maximum detectability $\Delta>0$, Theorem 1 implies that all Nash equilibrium payoff vectors are consistent with approximately myopic play if $(1-\delta)\,N/M\to\infty$, while Theorem 2 implies that a perfect folk theorem holds if $(1-\delta)\,N\log(N)\,/M\to 0$.

In Appendix B, we generalize Theorem 2 from random monitoring to arbitrary public, product-structure monitoring. This more general result (Theorem 4) is no harder to prove than Theorem 2, but it is less tightly connected to Theorem 1 because it relies on statistical conditions which are imposed directly on the action monitoring structure (Y,p). For this reason, we defer Theorem 4 to the appendix.

4.2 Overview of the Proof of Theorem 2

Theorem 2 (as well as its generalized version, Theorem 4) is a folk theorem for PPE in repeated games with public monitoring.[17] The standard proof approach, following FLM and Kandori and Matsushima (1998), relies on transferring continuation payoffs among the players along hyperplanes that are tangent to the boundary of the PPE payoff set. Unfortunately,

[16] This follows because $\sum_{x_i:\pi_{a_i,x_i}\geq\pi_{a_i,a_i}-2\pi_{a_i',a_i}}\pi_{a_i,x_i}\left(\frac{\pi_{a_i,x_i}-\pi_{a_i',x_i}}{\pi_{a_i,x_i}}\right)^2\geq\frac{\left(\pi_{a_i,a_i}-\pi_{a_i',a_i}\right)^2}{\pi_{a_i,a_i}}\geq\pi_{a_i,a_i}-2\pi_{a_i',a_i}$.

[17] Specifically, it is a "Nash threat" folk theorem, as F^* is the set of payoffs that Pareto-dominate a convex combination of static Nash equilibria. To extend this result to a "minmax threat" theorem, players must be made indifferent among all actions in the support of a mixed strategy that minmaxes an opponent. This requires a stronger identifiability condition, similar to Kandori and Matsushima's assumption (A1).

this approach encounters difficulties when N and δ vary simultaneously. The problem is that when N is large, changing each player's continuation payoff by a small amount can result in a large overall movement in the continuation payoff vector. Mathematically, FLM's proof relies on the equivalence of the L^1 norm and the Euclidean norm in $\mathbb{R}^N$. Since this equivalence is not uniform in N, their proof does not apply when N and δ vary simultaneously.[18]

Our proof (which is sketched in Appendix C, with details deferred to the online appendix) is instead based on the "block strategy" approach introduced by Matsushima (2004) and Hörner and Olszewski (2006) in the context of repeated games with private monitoring. We view the repeated game as a sequence of T-period blocks, where T is a number proportional to $1/(1-\delta)$. At the beginning of each block, a target payoff vector is determined by public randomization, and with high probability the players take actions throughout the block that deliver the target payoff. Players accrue promised continuation payoff adjustments whenever they are monitored, and the distribution of target payoffs in the next block is set to deliver the promised adjustments. To provide incentives, the required payoff adjustment when a player is monitored is of order N/M, the inverse of the monitoring probability. By the law of large numbers, when $T \gg N/M$, with high probability the total adjustment that a given player accrues over a T-period block is much smaller than $T \sim 1/(1-\delta)$, and is thus small enough that it can be delivered by appropriately specifying the distribution of target payoffs at the start of the next block.

The main difficulty in the proof is caused by the low-probability event that a player accrues an unusually large total adjustment over a block, so that at some point there is no room to provide additional incentives. In this case, the player can no longer be incentivized to take a non-myopic best response, and all players' behavior in the current block must change. Hence, if any player's payoff adjustment is "abnormal," all players' payoffs in the block may be far from the target equilibrium payoffs.

The proof ensures that rare payoff-adjustment abnormalities do not compromise either ex

[18] With random monitoring of M players, the per-period movement in each player's continuation payoff required to provide incentives is of order $(1-\delta)N/M$, so the movement of the continuation payoff vector in $\mathbb{R}^N$ is $O\left((1-\delta)N^{3/2}/M\right)$. For any ball $B \subseteq F^*$, consider the problem of generating the point $v = \operatorname{argmax}_{w \in B} w_1$ using continuation payoffs drawn from B. To satisfy promise keeping, player 1's continuation payoff must be within distance $O(1-\delta)$ of v, so the largest possible movement along a translated tangent hyperplane is $O\left(\sqrt{1-\delta}\right)$. FLM's proof approach thus requires that $(1-\delta)N^{3/2}/M \ll \sqrt{1-\delta}$, or equivalently $(1-\delta)N^3/M^2 \ll 1$, while we assume only $(1-\delta)N\log(N)/M \ll 1$. Hence, while the conditions for Theorem 4 are tight up to $\log(N)$ slack, FLM's approach would instead require slack $N^2/M \geq N$. On the other hand, in SW, we extend FLM's proof to give a folk theorem where discounting and monitoring vary simultaneously for a fixed stage game. There, FLM's approach works because N is fixed.

ante efficiency or the players' incentives. Efficiency is preserved if the block-length T is large enough that the probability that *any* player's payoff adjustment is abnormal is small. Since the per-period payoff adjustment for each player is $O(N/M)$ and the length of a block is $O(1/(1-\delta))$, standard concentration bounds imply that the probability that a given player's payoff adjustment is abnormal is $\exp(-O(M/((1-\delta)N)))$. Hence, by the union bound, the probability that any player's adjustment is abnormal is at most $N\exp(-O(M/((1-\delta)N)))$, which converges to 0 when $(1-\delta)N\log(N)/M \to 0$. This step accounts for the $\log(N)$ gap between Theorem 1 and 2.

Finally, since all players' payoffs are affected whenever any player's payoff adjustment becomes abnormal, incentives would be threatened if a player's action influenced the probability that other players' adjustments become abnormal. We avoid this problem by letting each player's adjustment depend only on the public signals of her own actions. Such a separation of payoff adjustments across players is possible under product structure monitoring. We do not know if Theorems 2 and 4 can be extended to non-product structure monitoring without introducing qualitatively larger (i.e., polynomial) slack.

5 Non-Cooperation under Collective Sanctions

We now consider an arbitrary public monitoring structure and a restricted class of equilibria—*linear perfect public equilibria*—which model collective incentive provision. We will show that cooperation is possible in this class of equilibria only if the discount rate is extremely small relative to the population size. We view this result as a near-impossibility theorem for large-group cooperation under collective sanctions. Put more colorfully, the result formalizes Hume's intuition that large groups cannot support cooperation by threatening "the abandoning of the whole project."

We say that a PPE is *linear* if all continuation payoff vectors lie on a line: for each player $i \neq 1$, there exists a constant $b_i \in \mathbb{R}$ such that, for all public histories h, h', we have $w_i(h') - w_i(h) = b_i(w_1(h') - w_1(h))$, where $w_i(h)$ denotes player i's equilibrium continuation payoff at history h. Relabeling the players if necessary, we can take $|b_i| \leq 1$ for all i without loss. Note that if $b_i \geq 0$ for all i then the players' preferences over histories are all aligned; while if $b_i < 0$ for some i then the players can be divided into two groups with opposite preferences. This notion of linear equilibrium generalizes *strongly symmetric equilibrium*

(SSE) in symmetric games, where $b_i = 1$ for all i.

Our result for linear PPE is as follows.

Theorem 3 *Fix any $\bar{u} > 0$ and $\underline{\pi} > 0$. For any $\varepsilon > 0$ and $\rho > 0$, there exists $k > 0$ such that, in any $(\bar{u}, \underline{\pi})$-bounded repeated game with public monitoring where*

$$(1-\delta)\exp\left(N^{1-\rho}\right) > k,$$

all linear PPE payoff vectors are consistent with ε-myopic play.

Theorem 3 differs from Theorem 1 in the required relationship between N and δ, and also in that Theorem 3 holds for any outcome monitoring precision. Intuitively, this is because optimal linear PPE take a bang-bang form even when the realized outcome profile is perfectly observed, so a binary signal that indicates which of two extreme continuation payoff vectors should be implemented is as effective as any more informative signal.

The proof of Theorem 3 is deferred to the online appendix. To see the main idea, consider the case where the game is symmetric and $b_i = 1$ for all i, so linear equilibria are SSE. Suppose also that we are in the noisy action case ($X = A$) with binary actions and symmetric noise, so that $|A_i| = 2$ and $\pi_{a_i,a_i} = 1 - \underline{\pi}$, $\pi_{a_i,a_i'} = \underline{\pi}$ for each $a_i \neq a_i'$. Finally, suppose we wish to enforce a symmetric pure action profile $\vec{a}_0 = (a_0, \ldots, a_0)$, where $\bar{g}_i(\vec{a}_0) = \nu$. By standard arguments, an optimal SSE takes the form of a "tail test," where the players are all punished if the number n of players for whom $x_i = a_0$ falls below a threshold n^*. Due to individual-level noise, when N is large the distribution of n is approximately normal, with mean $(1-\underline{\pi})N$ and standard deviation $\sqrt{\underline{\pi}(1-\underline{\pi})N}$. Denote the threshold z-score of a tail test with threshold n^* by $z^* = (n^* - (1-\underline{\pi})N)/\sqrt{\underline{\pi}(1-\underline{\pi})N}$, let ϕ and Φ denote the standard normal pdf and cdf, and let $x \in [0, \bar{u}/(1-\delta)]$ denote the size of the penalty when the tail test is failed. We then must have

$$\frac{\phi(z^*)}{\sqrt{\underline{\pi}(1-\underline{\pi})N}}x \geq \nu \quad \text{and} \quad \Phi(z^*)x \leq \bar{u},$$

where the first inequality is incentive compatibility, and the second inequality says that the expected penalty cannot exceed the stage-game payoff range. Dividing the first inequality by the second, we obtain

$$\frac{\phi(z^*)}{\Phi(z^*)} \geq \frac{\nu\sqrt{\underline{\pi}(1-\underline{\pi})N}}{\bar{u}}.$$

The left-hand side of this inequality is the Mills ratio of the standard normal distribution, which is approximately equal to $|z^*|$ when $z^* < 0$. Hence, $|z^*|$ must increase at least linearly with $\sqrt{N}$. But since $\phi(z^*)$ decreases exponentially with $|z^*|$, and hence exponentially with N, Theorem 3 now follows from incentive compatibility, which implies that the product of $\phi(z^*)/\sqrt{\underline{\pi}(1-\underline{\pi})N}$ and $\bar{u}/(1-\delta)$ must exceed ν.[19]

The analysis of tail tests as optimal incentive contracts under normal noise goes back to Mirrlees (1975). The logic of Theorem 3 shows that the size of the penalty in a Mirrleesian tail test must increase exponentially with the variance of the noise.[20] Theorem 3 is related to Proposition 1 of Sannikov and Skrzypacz (2007), which is an anti-folk theorem for SSE in a two-player repeated game where actions are observed with additive, normally distributed noise, with variance proportional to $(1-\delta)^{-1}$.[21] As a tail test is optimal in their setting, the reasoning just given implies that incentives can be provided only if $(1-\delta)^{-1}$ increases exponentially with the variance of the noise. Since in their model $(1-\delta)^{-1}$ increases with variance only linearly, they likewise obtain an anti-folk theorem. Similarly, Proposition 2 of Fudenberg and Levine (2007) is an anti-folk theorem in a game with one patient player and a myopic opponent, where the patient player's action is observed with additive, normal noise, with variance proportional to $(1-\delta)^{-\rho}$ for some $\rho > 0$; and their Proposition 3 is a folk theorem when the variance is constant in δ. Theorem 3 suggests that their anti-folk theorem extends whenever variance asymptotically dominates $\left(\log(1-\delta)^{-1}\right)^{1/(1-\rho)}$ for some $\rho > 0$, while their folk theorem extends whenever variance is asymptotically dominated by $\left(\log(1-\delta)^{-1}\right)^{1/(1+\rho)}$ for some $\rho > 0$.

---

[19] Conversely, if π_{a_i,a_i} is sufficiently large for each a_i and $(1-\delta)\exp\left(N^{1+\rho}\right) \to 0$ for some $\rho > 0$, then a folk theorem holds for linear equilibria. Intuitively, a target action profile a can now be enforced by a tail test where the players are all punished only if $x_i \neq a_i$ for *every* player i.

[20] We are not aware of a reference to this point in the literature.

[21] Their interpretation is that the players change their actions every Δ units of time, where $\delta = e^{-r\Delta}$ for fixed $r > 0$ and variance is inversely proportional to Δ, for example as a consequence of observing the average increments of a Brownian process.

6 Discussion

6.1 How Large is $W(\varepsilon)$?

Recall that Theorem 1 gives conditions under which all equilibrium payoffs lie in the set

$$W(\varepsilon) = \left\{ v \in \mathbb{R}^N : v = u(\alpha) \text{ for some } \alpha \text{ such that } \frac{1}{N}\sum_i \bar{g}_i(\alpha) \leq \varepsilon \right\}.$$

Payoffs in $W(\varepsilon)$ are attained by action distributions where the per-player average deviation gain is less than ε; however, a few players can have large deviation gains. A more standard notion of "ε-myopic play" is that *all* players' deviations gains are less than ε. The corresponding payoff vectors are the *static ε-correlated equilibrium payoffs*, given by

$$CE(\varepsilon) = \left\{ v \in \mathbb{R}^N : v = u(\alpha) \text{ for some } \alpha \text{ such that } \bar{g}_i(\alpha) \leq \varepsilon \text{ for all } i \right\}.$$

We now compare the sets $W(\varepsilon)$ and $CE(\varepsilon)$. We first give an example where $W(\varepsilon)$ and $CE(\varepsilon)$ are very different (and $W(\varepsilon)$ cannot be replaced by $CE(\varepsilon)$ in Theorem 1). We then give a condition under which maximum per-capita utilitarian welfare $\sum_i v_i/N$ is "similar" in $W(\varepsilon)$ and $CE(c\sqrt{\varepsilon})$, for a constant c. Intuitively, $W(\varepsilon)$ and $CE(\varepsilon)$ can be very different if incentive constraints bind for only a few players, and these players' actions have large effects on others' payoffs; while maximum utilitarian welfare in $W(\varepsilon)$ and $CE(c\sqrt{\varepsilon})$ are similar if each player's action has a small effect on every opponent's payoff.

For an example where $W(\varepsilon)$ and $CE(\varepsilon)$ differ, consider a "product choice" game where player 1 is a seller who chooses high or low quality (H or L), and the other $N-1$ players are buyers who choose whether to buy or not (B or D). If the seller takes $a_1 \in \{H, L\}$ and a buyer i takes $a_i \in \{B, D\}$, this buyer's payoff is given by

$$\begin{array}{ccc} & B & D \\ H & 1 & 0 \\ L & -1 & 0 \end{array},$$

while the seller's payoff is given by

$$\frac{2k}{N} - \mathbf{1}\{a_1 = H\},$$

where $k \in \{0, 1, \ldots, N\}$ is the number of buyers who take B. Suppose also that $X = A$ and $\underline{\pi}^i = \underline{\pi} \in (0, 1/3)$ for all i. Note that this game is $(3, \underline{\pi})$-bounded.

In this game, for any $\varepsilon > 0$, when N is sufficiently large, we have $(H, B, \ldots, B) \in A(\varepsilon)$, and hence $(1, 1, \ldots, 1) \in W(\varepsilon)$. This follows because the per-player average deviation gain at action profile $(H, B, \ldots, B)$ equals $1/N$: the seller has a deviation gain of 1, while each buyer has a deviation gain of 0. Thus, Theorem 1 does not preclude $(1, 1, \ldots, 1)$ (or any convex combination of $(1, 1, \ldots, 1)$ and $(0, 0, \ldots, 0)$) as an equilibrium payoff vector, even when $(1 - \delta) N/C$ is very large. This is reassuring, because the monitoring structure given by perfect monitoring of the seller's realized action (i.e., $Y = \{H, L\}$, $q(y|x) = \mathbf{1}\{y = x_1\}$) has channel capacity $\log 2$ and supports the payoff vector

$$\left(\frac{1 - 3\underline{\pi}}{1 - 2\underline{\pi}}, \frac{1 - 3\underline{\pi}}{1 - 2\underline{\pi}}, \ldots, \frac{1 - 3\underline{\pi}}{1 - 2\underline{\pi}}\right), \quad \text{for all } \delta > \frac{1}{2 - 3\underline{\pi}} \text{ and all } N \geq 2.$$[22]

In contrast, the greatest symmetric payoff vector in $CE(\varepsilon)$ is $(\varepsilon, \varepsilon, \ldots, \varepsilon)$, because the seller's deviation gain equals the probability that she takes H.

Intuitively, even though the efficient action profile $(H, B, \ldots, B)$ is not a static ε-correlated equilibrium, it can be supported as a repeated game equilibrium with "not very informative" monitoring. The reason is that only one player (the seller) is tempted to deviate at the efficient action profile, so monitoring one player suffices to support this action profile regardless of the population size (the number of buyers).

Next, for any $d \in (0, \bar{u})$, say that *per-capita externalities are bounded by* d if $\left|u_i\left(a_j', a_{-j}\right) - u_i(a)\right| \leq d/N$ for all $i \neq j, a_j', a$. For example, in a repeated random matching game, d can be taken as the maximum impact of a player's action on her partner's payoff, which is independent of N. In contrast, in the product choice game, per-capita externalities cannot be bounded uniformly in N, because the seller exerts an externality of 2 on each buyer who purchases.

In games with bounded per-capita externalities, maximum per-capita utilitarian welfare in $W(\varepsilon)$ and $CE\left(\sqrt{8d\varepsilon}\right)$ are "similar."

Proposition 1 *Assume that per-capita externalities are bounded by d. Then, for any $\varepsilon \in$*

[22] This is a standard calcuation, which results from considering "forgiving trigger strategies" that prescribe Nash reversion with probability ϕ when $y = L$. The smallest value of ϕ that induces the seller to take H is given by $\phi = (1 - \delta) / (\delta - 3\delta\underline{\pi})$, and substituting this into the value recursion $v = (1 - \delta)(1) + \delta(1 - \underline{\pi}\phi) v$ yields $v = (1 - 3\underline{\pi}) / (1 - 2\underline{\pi})$.

$(0, d)$ *and any* $v \in W(\varepsilon)$, *there exists* $v' \in CE\left(\sqrt{8d\varepsilon}\right)$ *such that*

$$\frac{1}{N}\left|\sum_{i \in I}(v_i - v_i')\right| \leq \sqrt{\frac{2\varepsilon}{d}}\bar{u}.$$

6.2 Conclusion

This paper has developed a theory of large-group cooperation based on repeated games with individual-level noise where the population size, discount factor, stage game, and monitoring structure all vary together in a flexible manner. Our main results establish necessary and sufficient conditions for cooperation, which identify the ratio of the discount rate and the *per-capita channel capacity* of the outcome monitoring structure as a key statistic. For a class of monitoring structures, our necessary and sufficient conditions coincide up to $\log(N)$ slack. We also show that cooperation in a *linear equilibrium* is possible only under much more stringent conditions. This last result demonstrates a sense in which large-group cooperation must rely on personalized sanctions.

Our results raise several questions for future theoretical and applied research. On the theory side, this paper has focused on insufficient monitoring precision as an obstacle to large-group cooperation. In reality, insufficient precision coexists with other obstacles to cooperation, such as monitoring being decentralized (as in community enforcement models) and the possibility that some players may be irrational or fail to understand the equilibrium being played (as in our earlier work, Sugaya and Wolitzky, 2020, 2021). Combining these features may help develop a richer and more realistic perspective on the prospects for large-group cooperation.

As for applied work, more systematic empirical or experimental evidence on the determinants of cooperation in large-population repeated games would be valuable.[23] In particular, our results predict that, while either personalized or collective sanctions can work in small groups, personalized sanctions are much more effective in large groups. It would be interesting to test this prediction.

[23] Camera and Casari (2009) and Duffy and Ochs (2009), among others, run experiments on repeated games with random matching and private monitoring, i.e., community enforcement. As explained in the introduction, community enforcement raises additional issues beyond the ones we focus on, which arise even under public monitoring. Camera, Casari, and Bigoni (2013) include a treatment with public monitoring (without individual-level noise), where they find that larger groups cooperate less.

Appendix

A The Set $B(\varepsilon)$ in A Public-Goods Game

Consider the public-goods game where each player chooses *Contribute* or *Don't Contribute*, and a player's payoff is the fraction of players who contribute less a constant $c \in (0,1)$ (independent of N) if she contributes herself. Fix any $v \in (0, 1-c)$, let $v = (v, \ldots, v) \in \mathbb{R}^N$, and let $\varepsilon = cv(1-c-v)/4 > 0$. We show that $B_v(\varepsilon) \subseteq F$ for all N. Since no one contributing is a Nash equilibrium with 0 payoffs, this implies that $B_v(\varepsilon) \subseteq F^*$, and hence $v \in B(\varepsilon)$, for all N.

Fix any N. Since the game is symmetric, to show that $B_v(\varepsilon) \subseteq F$ it suffices to show that, for any number $n \in \{0, \ldots, N\}$, there exists a feasible payoff vector where n "favored" players receive payoffs no less than $v + \varepsilon$, and the remaining $N - n$ "disfavored" players receive payoffs no more than $v - \varepsilon$. First, consider the mixed action profile α^1 where favored players contribute with probability $\frac{v+\varepsilon}{1-c}$ and disfavored players always contribute. At this profile, favored players receive payoff $f(n) := \frac{n}{N}\frac{v+\varepsilon}{1-c} + \left(1 - \frac{n}{N}\right)(1) - c\frac{v+\varepsilon}{1-c}$, while disfavored players receive payoff $g(n) := \frac{n}{N}\frac{v+\varepsilon}{1-c} + \left(1 - \frac{n}{N}\right)(1) - c$. Now consider the mixed action profile α^2 where favored players contribute with probability $\frac{(v+\varepsilon)^2}{(1-c)f(n)}$ and disfavored players contribute with probability $\frac{v+\varepsilon}{f(n)}$. Note that each player's payoff at profile α^2 equals her payoff at profile α^1 multiplied by $\frac{v+\varepsilon}{f(n)}$. Therefore, at profile α^2, favored players receive payoff $f(n)\frac{v+\varepsilon}{f(n)} = v + \varepsilon$, while disfavored players receive payoff

$$
\begin{aligned}
g(n)\frac{v+\varepsilon}{f(n)} &= \left(f(n) - \left(1 - \frac{v+\varepsilon}{1-c}\right)c\right)\frac{v+\varepsilon}{f(n)} \\
&\leq v + \varepsilon - \left(1 - \frac{v+\varepsilon}{1-c}\right)c(v+\varepsilon) \quad \text{(since } f(n) \leq 1\text{)} \\
&\leq v - \varepsilon \quad \text{(since } \varepsilon = cv(1-c-v)/4\text{)}.
\end{aligned}
$$

B A More General Folk Theorem

For any $\eta > 0$, we say that a public action monitoring structure (Y, p) satisfies *η-individual identifiability* if

$$\sum_{y_i : p_i(y_i|a_i) \geq \eta^2} p_i(y_i|a_i) \left(\frac{p_i(y_i|a_i) - p_i(y_i|\alpha_i)}{p_i(y_i|a_i)} \right)^2 \geq \eta^2 \quad \text{for all } i \in I, a_i \in A_i, \alpha_i \in \Delta(A_i \backslash \{a_i\}). \tag{7}$$

This condition is a variant of FLM's individual full rank condition and Kandori and Matsushima's (1998) assumption (A2"). It says that the detectability of a deviation from a_i to any mixed action α_i supported on $A_i \backslash \{a_i\}$ is at least η^2, from the perspective of an observer who ignores signals that occur with probability less than η^2 under a_i. Intuitively, this requires that deviations from a_i are detectable, and that in addition detection does not rest on very rare signal realizations. This assumption will ensure that players can be incentivized through rewards whose variance and maximum absolute value are both of order $(1-\delta)/\eta^2$.[24]

Our general folk theorem is as follows.

Theorem 4 *Fix any $\bar{u} > 0$. For any $\varepsilon > 0$, there exists $k > 0$ such that, for any $(\bar{u}, 0)$-bounded repeated game with public, product-structure monitoring satisfying η-individual identifiability and*

$$\frac{(1-\delta)\log(N)}{\eta^2} < k, \tag{8}$$

we have $B(\varepsilon) \subseteq E$.

To prove Theorem 2 from Theorem 4, it suffices to show that random monitoring of M/N players with a noise structure with Δ detectability satisfies $\sqrt{\Delta M/N}$-individual identifiability. To see this, note that, under random monitoring of M/N players, we have

$$p_i(y_i|a_i) = \begin{cases} \frac{M}{N}\pi_{a_i,y_i} & \text{if } y_i \in X_i, \\ 1 - \frac{M}{N} & \text{if } y_i = \emptyset. \end{cases}$$

[24] If (7) were weakened by taking the sum over all y_i (rather than only y_i such that $p_i(y_i|a_i) \geq \eta^2$), player i could be incentivized by rewards with variance $O\left((1-\delta)/\eta^2\right)$, but not necessarily with maximum absolute value $O\left((1-\delta)/\eta^2\right)$. Our analysis requires controlling both the variance and absolute value of players' rewards, so we need the stronger condition.

We then have

$$\sum_{y_i : p_i(y_i|a_i) \geq \Delta M/N} p_i\left(y_i|a_i\right)\left(\frac{p_i\left(y_i|\alpha_i\right) - p_i\left(y_i|a_i\right)}{p_i\left(y_i|a_i\right)}\right)^2 = \frac{M}{N}\sum_{x_i : \pi_{a_i,x_i} \geq \Delta} \pi_{a_i,x_i}\left(\frac{\pi_{a_i,x_i} - \pi_{a_i',x_i}}{\pi_{a_i,x_i}}\right)^2 \geq \frac{\Delta M}{N}.$$

Hence, random monitoring of M/N players with a noise structure with Δ detectability satisfies $\sqrt{\Delta M/N}$-individual identifiability.

C Sketch of the Proof of Theorem 4

Fix any $v \in B(\varepsilon)$. We show that, for sufficiently large δ, the cube $B_v(\varepsilon/2)$ is self-generating. Since $B(\varepsilon)$ is compact, this implies that, for sufficiently large δ, $B(\varepsilon)$ is self-generating, and hence $B(\varepsilon) \subseteq E$.

Since $B_v(\varepsilon/2)$ is a cube, for each extreme point $v^* \in B_v(\varepsilon/2)$, there exists $\zeta \in \{-1,1\}^N$ such that $v_i^* \in \operatorname{argmax}_{w \in B_v(\varepsilon/2)} \zeta_i w_i$ for all i. To self-generate $B_v(\varepsilon/2)$, it is sufficient that, for each $\zeta \in \{-1,1\}^N$ and v^* satisfying $v_i^* \in \operatorname{argmax}_{w \in B_v(\varepsilon/2)} \zeta_i w_i$ for all i, we can find a number $T \in \mathbb{N}$, a T-period strategy σ, and a history-contingent continuation payoff $w\left(h^{T+1}\right)$ such that the following three conditions hold:

Promise Keeping $v_i^* = (1-\delta)\sum_{t=1}^T \delta^{t-1}\mathbb{E}^\sigma\left[u_i\left(a_t\right)\right] + \delta^T \mathbb{E}^\sigma\left[w_i\left(h^{T+1}\right)\right]$ for all i.

Incentive Compatibility $\tilde{\sigma}_i = \sigma_i$ is optimal in the T-period repeated game with objective $\mathbb{E}^{\tilde{\sigma}_i,\sigma_{-i}}\left[(1-\delta)\sum_{t=1}^T \delta^{t-1}u_i(a) + \delta^T w_i\left(h^{T+1}\right)\right]$, for all i.

Self Generation $w\left(h^{T+1}\right) \in B_v(\varepsilon/2)$ for all h^{T+1}.

Since $B_v(\varepsilon/2)$ is the cube with center v and side-length ε, and $v_i^* \in \operatorname{argmax}_{w \in B_v(\varepsilon/2)} \zeta_i w_i$ for all i, we have $w\left(h^{T+1}\right) \in B_v(\varepsilon/2)$ iff $\zeta_i\left(w_i\left(h^{T+1}\right) - v_i\right) \in [-\varepsilon, 0]$ for all i. Thus, defining $\psi_i\left(h^{T+1}\right) = \left(\delta^T/(1-\delta)\right)\left(w_i\left(h^{T+1}\right) - v_i^*\right)$, we can rewrite the above conditions as

Promise Keeping $v_i = \frac{1-\delta}{1-\delta^T}\mathbb{E}^\sigma\left[\sum_{t=1}^T \delta^{t-1}u\left(a_t\right) + \psi_i\left(h^{T+1}\right)\right]$ for all i.

Incentive Compatibility $\tilde{\sigma}_i = \sigma_i$ is optimal in the T-period repeated game with objective $\mathbb{E}^{\tilde{\sigma}_i,\sigma_{-i}}\left[\sum_{t=1}^T \delta^{t-1}u(a) + \psi_i\left(h^{T+1}\right)|\sigma_i', \sigma_{-i}\right]$, for all i.

Self Generation $-\frac{\delta^T}{1-\delta}\varepsilon \leq \zeta_i\psi_i\left(h^{T+1}\right) \leq 0$ for all i, h^{T+1}. Moreover, since $\lim_{\delta\to 1} -\frac{\delta^T}{1-\delta}\varepsilon = -\infty$, it suffices to require that $\zeta_i\psi_i\left(h^{T+1}\right) \leq 0$ for all i, h^{T+1}.

Fix ζ and v^*, and take $T = O\left((1-\delta)^{-1}\right)$. We construct a T-period strategy σ and a "reward function" $\psi_i\left(h^{T+1}\right)$ that satisfy the above conditions.

By (7), for each recommendation r_i, there exists $f_{i,r_i}(y_i)$ such that (i) augmenting player i's utility by $f_{i,r_i}(y_i)$ incentivizes her to take r_i, (ii) the expectation of $f_{i,r_i}(y_i)$ when player i takes r_i equals 0, and (iii) the variance of $f_{i,r_i}(y_i)$ is of order η^2. Indeed, these properties are achieved by taking $f_{i,r_i}(y_i)$ proportional to the likelihood ratio difference $\min_{\alpha_i \in \Delta(A_i \setminus \{a_i\})} (p_i(y_i|a_i) - p_i(y_i|\alpha_i))/p_i(y_i|a_i)$. (See Lemma 4.)

Since $v \in B(\varepsilon)$ and $v^* \in B_v(\varepsilon/2)$, there exists $\bar{\alpha} \in \Delta(A)$ such that $\zeta_i(u_i(\bar{\alpha}) - v_i^*) = \varepsilon/2$. Suppose that the recommendation profile r is drawn according to $\bar{\alpha}$ by public randomization (and players follow their recommendations), and define the reward function $\tilde{\psi}_i\left(h^{T+1}\right) = \sum_t \delta^{t-1} f_{i,r_{i,t}}(y_{i,t}) - \zeta_i \frac{1-\delta^T}{1-\delta}\frac{\varepsilon}{2}$. We call $\tilde{\psi}_i\left(h^{T+1}\right)$ the "base reward." We show that this strategy and reward function satisfy promise keeping and incentive compatibility, and also satisfy self generation with high probability. We then show how to modify the strategy and reward function to ensure that self generation is always satisfied.

Since $f_{i,r_i}(y_i)$ has 0 mean, promise keeping is immediate:

$$v_i = \frac{1-\delta}{1-\delta^T}\mathbb{E}^{\sigma}\left[\sum_{t=1}^{T}\delta^{t-1}u_i(a) + \tilde{\psi}_i\left(h^{T+1}\right)\right] = u_i(\bar{\alpha}) - \zeta_i\frac{\varepsilon}{2} = v_i^*.$$

Next, incentive compatibility holds because

$$\mathbb{E}^{\tilde{\sigma}_i,\sigma_{-i}}\left[\sum_{t=1}^{T}\delta^{t-1}u_i(a_t) + \tilde{\psi}_i\left(h^{T+1}\right)\right] = \mathbb{E}^{\tilde{\sigma}_i,\sigma_{-i}}\left[\sum_{t=1}^{T}\delta^{t-1}\left(u_i(a_t) + f_{i,r_{i,t}}(y_{i,t})\right)\right] - \frac{1-\delta^T}{1-\delta}\zeta_i\frac{\varepsilon}{2},$$

so the augmented per-period payoff is $u_i(a) + f_{i,r_{i,t}}(y_{i,t})$. Moreover, since the variance of f_{i,r_i} is $O(\eta^2)$ and T is $O\left((1-\delta)^{-1}\right)$, by a standard concentration inequality, the self generation constraint $\zeta_i\tilde{\psi}_i\left(h^{T+1}\right) \leq 0$ holds for all i with probability at least

$$N\exp\left(-\frac{\frac{1-\delta^T}{1-\delta}\zeta_i\frac{\varepsilon}{2}}{\sqrt{T\eta^2}}\right) \approx \exp\left(-\sqrt{\frac{(1-\delta)\log N}{\eta^2}}\right).$$

Therefore, by (8), self generation holds with high probability when k is small. (See Lemmas 5 and 7.)

We now modify the strategy and reward to satisfy self-generation at every history. To this end, define a stopping time as the first period τ such that

$$\zeta_i \sum_{t=1}^{\tau} \delta^{t-1} f_{i,r_{i,t}}\left(y_{i,t}\right) > \bar{f}, \tag{9}$$

where $\bar{f}$ is a positive constant less than $\left(\left(1-\delta^T\right)/\left(1-\delta\right)\right)\varepsilon/2$. That is, in (the random) period τ, for a player, the base reward $\tilde{\psi}_i\left(h^{T+1}\right)$ becomes abnormal. If no such period arises, define $\tau = T$. By the same concentration argument as above, abnormality does not happen to any player's base reward (that is, $\tau = T$) with high probability: in particular,

$$\Pr\left(\tau < T\right) \approx \exp\left(-\sqrt{\frac{\left(1-\delta\right)\log N}{\eta^2}}\right). \tag{10}$$

We now define the modified strategy.

If $\tau = T$, then in every period r is drawn according to $\bar{\alpha}$ and the reward equals $\tilde{\psi}_i\left(h^{T+1}\right)$.

If $\tau < T$, then let I^* be the set of players whose base reward satisfies (9). For each $i \in I^*$, we add or subtract a constant from the rewards of players $-i$ to satisfy self generation. Since monitoring has a product structure, players $-i$ cannot control the realization of player i's reward. Thus, this addition or subtraction does not affect incentives.

If I^* is a singleton, $I^* = \{i\}$, then player i starts taking a static best response. Meanwhile, players $-i$ take r_{-i} drawn from $\bar{\alpha}$ if $\zeta_i = 1$, and take static Nash actions $\left(\alpha_j^{NE}\right)_{j\neq i}$ if $\zeta_i = -1$. Let $u_i\left(\zeta_i\right)$ be player i's resulting instantaneous payoff. Since $v^* \in F^*$, we have $\zeta_i\left(u_i\left(\zeta_i\right) - u_i\left(\bar{\alpha}\right)\right) \geq 0$. Hence, if player i's period t reward is fixed at $u_i\left(\bar{\alpha}\right) - u_i\left(\zeta_i\right)$, self generation is satisfied, and player i's period t augmented payoff equals $u_i\left(\bar{\alpha}\right)$. If instead $|I^*| \geq 2$, then all players' subsequent rewards equal 0.

Since $\tau = T$ with high probability by (10), expected payoffs under the modified strategy and reward are close to v. Further adjusting the rewards by a small constant thus achieves promise keeping. Moreover, self generation now holds by construction. Finally, for any period $t > \tau$, incentive compatibility holds, because either a player's reward is fixed and she is supposed to take a static best response, or she is incentivized by the base reward function.

To complete the proof, it remains to establish incentive compatibility for periods $t \leq \tau$. For $t \leq \tau$, player i's augmented period t payoff is $u_i\left(a_i, r_{-i}\right) + f_{i,r_i}\left(y_i\right))$. Thus, to show that it is optimal for player i to follow her recommendation, it suffices to show that she cannot gain by manipulating the stopping time τ.

Since monitoring has a product structure, player i cannot influence others' rewards. Player i also cannot improve her augmented period t payoff by manipulating her own reward, because both $u_i(r)+f_{i,r_{i,t}}(y_{i,t}))$ and $u_i(\zeta_i)+u_i(\bar{\alpha})-u_i(\zeta_i))$ equal $u(\bar{\alpha})$ regardless of whether $t \leq \tau$ or $t > \tau$. However, there is one potential benefit from manipulation: once τ realizes with $I^* = \{i\}$, the chance of a constant being added or subtracted from player i's reward vanishes, but if τ first realizes with $I^* \neq \{i\}$, this addition or subtraction occurs. To prevent this adjustment from affecting player i's incentive, a "fictitious" recommendation $\tilde{r}_t$ is drawn according to $\bar{\alpha}$, and a fictitious signal $\tilde{y}$ is drawn according to $p(\tilde{y}|\tilde{r})$, and the base rewards are updated according to the fictitious recommendations and signals even when $t > \tau$. (See (25) for the definition of the fictitious recommendations and signals.) If player $j \neq i$'s fictitious base reward satisfies (9), we add or subtract a constant from player i's reward. (See (26) for the definition of the event that induces this addition or subtraction. Note also that this fictitious update of player j's base reward is used solely to satisfy player i's incentives and does not affect player j's reward.) Given this modification, player i does not have an incentive to manipulate her own reward to manipulate the distribution of τ (see Lemma 6), and hence incentive compatibility holds (Lemma 8).

D Proof of Proposition 1

We establish the stronger conclusion that, for any $v \in W(\varepsilon)$ and any $c \geq \sqrt{8d/\varepsilon}$, there exists $v' \in CE(c\varepsilon)$ such that $\left|\sum_{i\in I}(v_i - v'_i)\right|/N \leq 4\bar{u}/c$. (The stated conclusion follows by taking $c = \sqrt{8d/\varepsilon}$.) Fix $\varepsilon \in (0,d)$ and $\alpha \in A(\varepsilon)$. Let $J = \{i : \bar{g}_i(\alpha) > c\varepsilon/2\}$, and note that $|J| \leq 2N/c$. Let $\tilde{\alpha} \in \Delta(A)$ be an action distribution that has the same marginal on $A_{I\setminus J}$ as α and that satisfies $\bar{g}_i(\tilde{\alpha}) \leq c\varepsilon$ for all $i \in J$: for example, take a Nash equilibrium in the game among the players in J, where the action distribution among the players in $I\setminus J$ is held fixed. Since $\left|u_i(a'_j, a_{-j}) - u_i(a)\right| \leq d/N$ for all $i \neq j, a'_j, a$, and the actions of at most $2N/c$ players differ between $\tilde{\alpha}$ and α, we have $\bar{g}_i(\tilde{\alpha}) \leq \bar{g}_i(\alpha) + 4d/c$ for each $i \in I\setminus J$. Since $\bar{g}_i(\alpha) \leq c\varepsilon/2$ (as $i \in I\setminus J$) and $4d/c \leq c\varepsilon/2$ (as $c \geq \sqrt{8d/\varepsilon}$), we have $\bar{g}_i(\tilde{\alpha}) \leq c\varepsilon$. Since we assumed that $\bar{g}_i(\tilde{\alpha}) \leq c\varepsilon$ for all $i \in J$, we have $\bar{g}_i(\tilde{\alpha}) \leq c\varepsilon$ for all $i \in I$, and hence $u(\tilde{\alpha}) \in CE(c\varepsilon)$. Finally, since the actions of at most $2N/c$ players differ between $\tilde{\alpha}$ and α, we have $|u_i(\tilde{\alpha}) - u_i(\alpha)| \leq 2d/c \leq 2\bar{u}/c$ for all $i \in I\setminus J$, and by definition of $\bar{u}$ we have $|u_i(\tilde{\alpha}) - u_i(\alpha)| \leq \bar{u}$ for all $i \in J$. Since $c > 2$ and $|J| \leq 2N/c$, we have

$\left|\sum_{i\in I}\left(u_i\left(\tilde{\alpha}\right)-u_i\left(\alpha\right)\right)\right|\leq\left(N-2N/c\right)2\bar{u}/c+\left(2N/c\right)\bar{u}\leq 4N\bar{u}/c.$

References

[1] Abreu, Dilip, David Pearce, and Ennio Stacchetti. "Optimal Cartel Equilibria with Imperfect Monitoring." *Journal of Economic Theory* 39.1 (1986): 251-269.

[2] Abreu, Dilip, David Pearce, and Ennio Stacchetti. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58.5 (1990): 1041-1063.

[3] Abreu, Dilip, Paul Milgrom, and David Pearce. "Information and Timing in Repeated Partnerships." *Econometrica* 59.6 (1991): 1713-1733.

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

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

[6] Athey, Susan, Kyle Bagwell, and Chris Sanchirico. "Collusion and Price Rigidity." *Review of Economic Studies* 71.2 (2004): 317-349.

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

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

[9] Billingsley, Patrick. *Probability and Measure, 3rd ed.* Wiley (1995).

[10] Boucheron, Stéphane, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence.* Oxford University Press, 2013.

[11] Bowles, Samuel, and Herbert Gintis. *A Cooperative Species: Human Reciprocity and its Evolution.* Princeton University Press, 2011.

[12] Boyd, Robert, and Peter J. Richerson. "The Evolution of Reciprocity in Sizable Groups." *Journal of Theoretical Biology* 132.3 (1988): 337-356.

[13] Camera, Gabriele, and Marco Casari. "Cooperation among Strangers under the Shadow of the Future." *American Economic Review* 99.3 (2009): 979-1005.

[14] Camera, Gabriele, Marco Casari, and Maria Bigoni. "Money and Trust among Strangers." *Proceedings of the National Academy of Science*s 110.37 (2013): 14889-14893.

[15] Cover, Thomas M., and Joy A. Thomas. *Elements of Information Theory, 2nd ed.* Wiley (2006).

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

[17] Duffy, John, and Jack Ochs. "Cooperative Behavior and the Frequency of Social Interaction." *Games and Economic Behavior* 66.2 (2009): 785-812.

[18] Edmiston, Jake, and Graeme Hamilton, "The Last Days of Quebec's Maple Syrup Rebellion," *National Post*, April 6, 2018.

[19] Ekmekci, Mehmet, Olivier Gossner, and Andrea Wilson. "Impermanent Types and Permanent Reputations." *Journal of Economic Theory* 147.1 (2012): 162-178.

[20] Ellickson, Robert C. *Order without Law: How Neighbors Settle Disputes.* Harvard University Press, 1991.

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

[22] Faingold, Eduardo. "Reputation and the Flow of Information in Repeated Games." *Econometrica* 88.4 (2020): 1697-1723.

[23] Feigenberg, Benjamin, Erica Field, and Rohini Pande. "The Economic Returns to Social Interaction: Experimental Evidence from Microfinance." *Review of Economic Studies* 80.4 (2013): 1459-1483.

[24] Forges, Francoise. "An Approach to Communication Equilibria." *Econometrica* 54.6 (1986): 1375-1385.

[25] Fudenberg, Drew, and David K. Levine. "Efficiency and Observability with Long-Run and Short-Run Players." J*ournal of Economic Theory* 62.1 (1994): 103-135.

[26] Fudenberg, Drew, and David K. Levine. "Continuous Time Limits of Repeated Games with Imperfect Public Monitoring." *Review of Economic Dynamics* 10.2 (2007): 173-192.

[27] Fudenberg, Drew, and David K. Levine. "Repeated Games with Frequent Signals." *Quarterly Journal of Economics* 124.1 (2009): 233-265.

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

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

[30] Gossner, Olivier. "Simple Bounds on the Value of a Reputation." *Econometrica* 79.5 (2011): 1627-1641.

[31] Gossner, Olivier, Penelope Hernández, and Abraham Neyman. "Optimal Use of Communication Resources." *Econometrica* 74.6 (2006): 1603-1636.

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

[33] Green, Edward J., and Robert H. Porter. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52.1 (1984): 87-100.

[34] Hellman, Ziv, and Ron Peretz. "A survey on entropy and economic behaviour." *Entropy* 22.2 (2020): 157.

[35] Holmström, Bengt. "Moral Hazard and Observability." *Bell Journal of Economics* 10.1 (1979): 74-91.

[36] Hörner, Johannes, and Wojciech Olszewski. "The Folk Theorem for Games with Private Almost-Perfect Monitoring" *Econometrica* 74.7 (2006): 1499-1544.

[37] Hörner, Johannes, and Satoru Takahashi. "How Fast do Equilibrium Payoff Sets Converge in Repeated Games?" *Journal of Economic Theory* 165 (2016): 332-359.

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

[39] Kandori, Michihiro. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory* 102.1 (2002): 1-15.

[40] Kandori, Michihiro, and Hitoshi Matsushima. "Private Observation, Communication and Collusion." *Econometrica* 66.3 (1998): 627-652.

[41] Karlan, Dean S. "Social Connections and Group Banking." *Economic Journal* 117.517 (2007): F52-F84.

[42] Kuitenbrouwer, "'It's Crazy, Isn't It': Quebec's Maple Syrup Rebels Face Ruin as Cartel Crushes Dissent," *Financial Post*, December 5, 2016.

[43] Matsushima, Hitoshi. "Repeated Games with Private Monitoring: Two Players." *Econometrica* 72.3 (2004): 823-852.

[44] Miguel, Edward, and Mary Kay Gugerty. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89.11-12 (2005): 2325-2368.

[45] Mirrlees, James A. "The Theory of Moral Hazard and Unobservable Behaviour: Part I." *Working Paper* (1975) (published in *Review of Economic Studies* 66.1 (1999): 3-21).

[46] Neyman, Abraham, and Daijiro Okada. "Strategic Entropy and Complexity in Repeated Games." *Games and Economic Behavior* 29.1-2 (1999): 191-223.

[47] Neyman, Abraham, and Daijiro Okada. "Repeated Games with Bounded Entropy." *Games and Economic Behavior* 30.2 (2000): 228-247.

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

[49] Pai, Mallesh M., Aaron Roth, and Jonathan Ullman. "An Antifolk Theorem for Large Repeated Games." *ACM Transactions on Economics and Computation (TEAC)* 5.2 (2016): 1-20.

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

[51] Sannikov, Yuliy, and Andrzej Skrzypacz. "Impossibility of Collusion under Imperfect Monitoring with Flexible Production." *American Economic Review* 97.5 (2007): 1794-1823.

[52] Sannikov, Yuliy, and Andrzej Skrzypacz. "The Role of Information in Repeated Games with Frequent Actions." *Econometrica* 78.3 (2010): 847-882.

[53] Seabright, Paul. *The Company of Strangers.* Princeton University Press, 2004.

[54] Sugaya, Takuo, and Alexander Wolitzky. "A Few Bad Apples Spoil the Barrel: An Antifolk Theorem for Anonymous Repeated Games with Incomplete Information." *American Economic Review* 110.12 (2020): 3817-35.

[55] Sugaya, Takuo, and Alexander Wolitzky, "Communication and Community Enforcement." *Journal of Political Economy* 129.9 (2021): 2595-2628.

[56] Sugaya, Takuo, and Alexander Wolitzky, "Monitoring versus Discounting in Repeated Games." *Econometrica*, Forthcoming (2023a).

[57] Sugaya, Takuo, and Alexander Wolitzky, "Performance Feedback in Long-Run Relationships: A Rate of Convergence Approach." *Working Paper* (2023b).

Online Appendix

E Proof of Theorem 4

E.1 Preliminaries

Fix any $\varepsilon > 0$. If $\varepsilon \geq \bar{u}/2$ then $B(\varepsilon) = \emptyset$ and the conclusion of the theorem is trivial, so assume without loss that $\varepsilon < \bar{u}/2$. We begin with two preliminary lemmas. First, for each $i \in I$ and $r_i \in A_i$, we define a function $f_{i,r_i} : Y_i \to \mathbb{R}$ that will later be used to specify player i's continuation payoff as a function of y_i.

Lemma 4 *Under η-individual identifiability, for each $i \in I$ and $r_i \in A_i$ there exists a function $f_{i,r_i} : Y_i \to \mathbb{R}$ such that*

$$\begin{aligned} \mathbb{E}\left[f_{i,r_i}(y_i)\,|r_i\right] - \mathbb{E}\left[f_{i,r_i}(y_i)\,|a_i\right] &\geq \bar{u} \quad \text{for all } a_i \neq r_i, & (11)\\ \mathbb{E}\left[f_{i,r_i}(y_i)\,|r_i\right] &= 0, & (12)\\ \operatorname{Var}\left(f_{i,r_i}(y_i)\,|r_i\right) &\leq \bar{u}^2/\eta^2, \quad \text{and} & (13)\\ \left|f_{i,r_i}(y_i)\right| &\leq 2\bar{u}/\eta^2 \quad \text{for all } y_i. & (14) \end{aligned}$$

Proof. Fix i and r_i. Let $Y_i^* = \{y_i : p_i(y_i, r_i) \geq \eta^2\}$, and let

$$p_i(r_i) = \left(\sqrt{p_i(y_i|r_i)}\right)_{y_i \in Y_i^*} \quad \text{and} \quad P_i(r_i) = \bigcup_{a_i \neq r_i} \left(\frac{p_i(y_i|a_i)}{\sqrt{p_i(y_i|r_i)}}\right)_{y_i \in Y_i^*}.$$

Note that (7) is equivalent to $d(p_i(r_i), \operatorname{co}(P_i(r_i))) \geq \eta$ for all $i \in I, r_i \in A_i$, where $d(\cdot,\cdot)$ denotes Euclidean distance in $\mathbb{R}^{|Y_i^*|}$. Hence, by the separating hyperplane theorem, there exists $x = (x(y_i))_{y_i \in Y_i^*} \in \mathbb{R}^{|Y_i^*|}$ such that $\|x\| = 1$ and $(p_i(r_i) - p) \cdot x \geq \eta$ for all $p \in P_i(r_i)$. By definition of p_i and P_i, this implies that $\sum_{y_i \in Y_i^*} (p_i(y_i|r_i) - p_i(y_i|a_i))\, x(y_i) \geq \eta\sqrt{p_i(y_i|r_i)}$ for all $a_i \neq r_i$. Now define

$$\begin{aligned} f_{i,r_i}(y_i) &= \frac{\bar{u}}{\eta}\left(\frac{x(y_i)}{\sqrt{p_i(y_i|r_i)}} - \sum_{\tilde{y}_i \in Y_i} \frac{p(\tilde{y}_i|r_i)}{\sqrt{p_i(\tilde{y}_i|r_i)}} x_i(\tilde{y}_i)\right) \quad \text{for all } y_i \in Y_i^*, \quad \text{and}\\ f_{i,r_i}(y_i) &= 0 \quad \text{for all } y_i \notin Y_i^*. \end{aligned}$$

Clearly, conditions (11) and (12) hold. Moreover, since $\mathbb{E}\left[f_{i,r_i}(y_i)\,|r_i\right] = 0$ and the term $\sum_{\tilde{y}_i \in Y_i} \sqrt{p(\tilde{y}_i|r_i)}\, x_i(\tilde{y}_i)$ is independent of y_i, we have

$$\operatorname{Var}\left(f_{i,r_i}(y_i)\,|r_i\right) = \mathbb{E}\left[\frac{\bar{u}^2 x(y_i)^2}{\eta^2 p_i(y_i|r_i)}\right] - \mathbb{E}\left[\frac{\bar{u} x_i(y_i)}{\eta\sqrt{p_i(y_i|r_i)}}\right]^2 \leq \frac{\bar{u}^2}{\eta^2} \sum_{y_i \in Y_i^*} x(y_i)^2 \leq \frac{\bar{u}^2}{\eta^2},$$

and hence (13) holds. Finally, (14) holds since, for each $y_i \in Y_i^*$,

$$\left|f_{i,r_i}\left(y_i\right)\right| \leq \frac{\bar{u}}{\eta}\left(\frac{\left|x\left(y_i\right)\right| + \sum_{\tilde{y}_i \in Y_i^*} p\left(\tilde{y}_i|r_i\right)\left|x_i\left(\tilde{y}_i\right)\right|}{\sqrt{p_i\left(y_i|r_i\right)}}\right) \leq \frac{\bar{u}}{\eta^2}\left(1 + \sum_{\tilde{y}_i \in Y_i^*} p\left(\tilde{y}_i|r_i\right)\right) \leq \frac{2\bar{u}}{\eta^2}.$$

■

Now fix $i \in I$ and $r_i \in A_i$, and suppose that $y_{i,t} \sim p_i\left(\cdot|r_i\right)$ for each period $t \in \mathbb{N}$, independently across periods (which would be the case in the repeated game if r_i were taken in every period). By (13), for any $T \in \mathbb{N}$, we have

$$\mathrm{Var}\left(\sum_{t=1}^{T} \delta^{t-1} f_{i,r_i}\left(y_{i,t}\right)\right) = \sum_{t=1}^{T} \delta^{2(t-1)} \mathrm{Var}\left(f_{i,r_i}\left(y_{i,t}\right)\right) \leq \frac{1-\delta^{2T}}{1-\delta^2}\frac{\bar{u}^2}{\eta^2}.$$

Together with (12) and (14), Bernstein's inequality (Boucheron, Lugosi, and Massart, 2013, Theorem 2.10) now implies that, for any $T \in \mathbb{N}$ and $\bar{f} \in \mathbb{R}_+$, we have

$$\Pr\left(\sum_{t=1}^{T} \delta^{t-1} f_{i,r_i}\left(y_{i,t}\right) \geq \bar{f}\right) \leq \exp\left(-\frac{\bar{f}^2\eta^2}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2 + \frac{2}{3}\bar{f}\bar{u}\right)}\right). \tag{15}$$

Our second lemma fixes T and $\bar{f}$ so that the bound in (15) is sufficiently small, and some other conditions used in the proof also hold.

Lemma 5 *There exists $k > 0$ such that, whenever $(1-\delta)\log(N)/\eta^2 < k$, there exist $T \in \mathbb{N}$ and $\bar{f} \in \mathbb{R}$ that satisfy the following three inequalities:*

$$60\bar{u}N\exp\left(-\frac{\left(\frac{\bar{f}}{3}\right)^2\eta^2}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2 + \frac{2}{3}\frac{\bar{f}}{3}\bar{u}\right)}\right) \leq \varepsilon, \tag{16}$$

$$8\frac{1-\delta}{1-\delta^T}\left(\bar{f} + \frac{2\bar{u}}{\eta^2}\right) \leq \varepsilon, \tag{17}$$

$$4\bar{u}\frac{1-\delta^T}{\delta^T} + \frac{1-\delta}{\delta^T}\left(\bar{f} + \frac{2\bar{u}}{\eta^2}\right) \leq \varepsilon. \tag{18}$$

Proof. Let T be the largest integer such that $8\bar{u}\left(1-\delta^T\right)/\delta^T \leq \varepsilon$, and let

$$\bar{f} = \sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log\left(N\right)\frac{1-\delta^T}{1-\delta}\frac{\bar{u}^2}{\eta^2}}.$$

Note that if $(1-\delta)\log(N)/\eta^2 \to 0$ then $1-\delta^T \to \varepsilon/(\varepsilon+8\bar{u})$, and hence $(1-\delta)\log(N)/\left(\eta^2\left(1-\delta^T\right)\right) \to 0$. Therefore, there exists $k>0$ such that, whenever $(1-\delta)\log(N)/\eta^2 < k$, we have

$$\frac{4}{9}\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta^2}} \leq 1 \quad \text{and} \tag{19}$$

$$8\bar{u}\left(\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta^2}}+\frac{1-\delta}{1-\delta^T}\frac{2}{\eta^2}\right) \leq \varepsilon. \tag{20}$$

It now follows from straightforward algebra (provided in Appendix E.4) that (16)–(18) hold for every $k \geq \bar{k}$. ■

E.2 Equilibrium Construction

Fix any k, T, and $\bar{f}$ that satisfy (16)–(18), as well any $v \in B(\varepsilon)$. For each extreme point v^* of $B_v(\varepsilon/2)$, we construct a PPE in a T-period, finitely repeated game augmented with continuation values drawn from $B_v(\varepsilon/2)$ that generates payoff vector v^*. By standard arguments, this implies that $B_v(\varepsilon/2) \subseteq E(\Gamma)$, and hence that $v \in E(\Gamma)$.[25] Since $v \in B(\varepsilon)$ was chosen arbitrarily, it follows that $B(\varepsilon) \subseteq E(\Gamma)$.

Specifically, for each $\zeta \in \{-1,1\}^N$ and $v^* = \operatorname{argmax}_{v\in B_v(\varepsilon/2)} \zeta \cdot v$, we construct a public strategy profile σ in a T-period, finitely repeated game (which we call a *block strategy profile*) together with a continuation value function $w: H^{T+1} \to \mathbb{R}^N$ such that, letting $\psi_i\left(h^{T+1}\right) = \frac{\delta^T}{1-\delta}\left(w_i\left(h^{T+1}\right)-v_i^*\right)$, we have

$$\textit{Promise Keeping:} \quad v_i^* = \frac{1-\delta}{1-\delta^T}\mathbb{E}^{\sigma}\left[\sum_{t=1}^{T}\delta^{t-1}u_{i,t}+\psi_i\left(h^{T+1}\right)\right] \quad \text{for all } i, \tag{21}$$

$$\textit{Incentive Compatibility:} \quad \sigma_i \in \operatorname*{argmax}_{\tilde{\sigma}_i}\mathbb{E}^{\tilde{\sigma}_i,\sigma_{-i}}\left[\sum_{t=1}^{T}\delta^{t-1}u_{i,t}+\psi_i\left(h^{T+1}\right)\right] \quad \text{for all } i, \tag{22}$$

$$\textit{Self Generation:} \quad \zeta_i\psi_i\left(h^{T+1}\right) \in \left[-\frac{\delta^T}{1-\delta}\varepsilon, 0\right] \quad \text{for all } i \text{ and } h^{T+1}. \tag{23}$$

Fix $\zeta \in \{-1,1\}^N$ and $v^* = \operatorname{argmax}_{v\in B_v(\varepsilon/2)} \zeta \cdot v$. We construct a block strategy profile σ and continuation value function ψ which, in the next subsection, we show satisfy these three conditions. This will complete the proof of the theorem.

First, fix a correlated action profile $\bar{\alpha} \in \Delta(A)$ such that

$$u_i(\bar{\alpha}) = v_i^* + \zeta_i\varepsilon/2 \quad \text{for all } i, \tag{24}$$

and fix a probability distribution over static Nash equilibria $\alpha^{NE} \in \Delta\left(\prod_i \Delta(A_i)\right)$ such that $u_i\left(\alpha^{NE}\right) \leq v_i^* - \varepsilon/2$ for all i. Such $\bar{\alpha}$ and α^{NE} exist because $v^* \in B_v(\varepsilon/2)$ and $B_v(\varepsilon) \subseteq F^*$.

[25] Specifically, at each history h^{T+1} that marks the end of a block, public randomization can be used to select an extreme point v^* to be targeted in the following block, with probabilities chosen so that the expected payoff $\mathbb{E}[v^*]$ equals the promised continuation value $w\left(h^{T+1}\right)$.

We now construct the block strategy profile σ. For each player $i \in I$ and period $t \in \{1, \ldots, T\}$, we define a *state* $\theta_{i,t} \in \{0, 1\}$ for player i in period t. The states are determined by the public history, and so are common knowledge among the players. We first specify players' prescribed actions as a function of the state, and then specify the state as a function of the public history.

Prescribed Equilibrium Actions: For each period t, let $r_t \in A$ be a pure action profile which is drawn by public randomization at the start of period t from the distribution $\bar{\alpha} \in \Delta(A)$ fixed in (24), and let $\varrho_t^{NE} \in \prod_i \Delta(A_i)$ be a mixed action profile which is drawn by public randomization at the start of period t from the distribution α^{NE}. The prescribed equilibrium actions are defined as follows.

1. If $\theta_{i,t} = 0$ for all $i \in I$, the players take $a_t = r_t$.
2. If there is a unique player i such that $\theta_{i,t} = 1$, the players take $a_t = (r'_i, r_{-i,t})$ for some $r'_i \in BR_i(r_{-i,t})$ if $\zeta_i = 1$, and they take ϱ_t^{NE} if $\zeta_i = -1$, where $BR_i(r_{-i}) = \operatorname{argmax}_{a_i \in A_i} u_i(a_i, r_{-i})$ is the set of i's best responses to r_{-i}.
3. If there is more than one player i such that $\theta_{i,t} = 1$, the players take ϱ_t^{NE}.

Let $\alpha_t^* \in \prod_i \Delta(A_i)$ denote the distribution of prescribed equilibrium actions, prior to public randomization z_t.

(It may be helpful to informally summarize the prescribed actions. So long as $\theta_{i,t} = 0$ for all players, the players take actions drawn from the target action distribution $\bar{\alpha}$. If $\theta_{i,t} = 1$ for multiple players, the inefficient Nash equilibrium distribution α^{NE} is played. If $\theta_{i,t} = 1$ for a unique player i, player i starts taking static best responses; moreover, if $\zeta_i = -1$ then α^{NE} is played.)

It will be useful to introduce the following additional state variable $S_{i,t}$, which summarizes player i's prescribed action as a function of $(\theta_{j,t})_{j \in I}$:

1. $S_{i,t} = 0$ if $\theta_{j,t} = 0$ for all $j \in I$, or if there exists a unique player $j \neq i$ such that $\theta_{j,t} = 1$, and for this player we have $\zeta_j = 1$. In this case, player i is prescribed to take $a_{i,t} = r_{i,t}$.
2. $S_{i,t} = NE$ if $\theta_{i,t} = 0$ and either (i) there exists a unique player j such that $\theta_{j,t} = 1$, and for this player we have $\zeta_j = -1$, or (ii) there are two distinct players j, j' such that $\theta_{j,t} = \theta_{j',t} = 1$. In this case, player i is prescribed to take $\varrho_{i,t}^{NE}$.
3. $S_{i,t} = BR$ if $\theta_{i,t} = 1$. In this case, player i is prescribed to best respond to her opponents' actions (which equal either $r_{-i,t}$ or $\varrho_{-i,t}^{NE}$, depending on ζ_i and $(\theta_{j,t})_{j \neq i}$.)

States: At the start of each period t, conditional on the public randomization draw of $r_t \in A$ described above, an additional ("fictitious") random variable $\tilde{y}_t \in Y$ is also drawn by public randomization, with distribution $p(\tilde{y}_t | r_t)$. That is, the distribution of the public randomization draw $\tilde{y}_t$ conditional on the draw r_t is the same as the distribution of the realized public signal profile $\tilde{y}_t$ at action profile r_t; however, the distribution of $\tilde{y}_t$ depends only on the public randomization draw r_t and not on the players' actions. For each player i and period t, let $f_{i,r_{i,t}} : Y_i \to \mathbb{R}$ be defined as in Lemma 4, and let

$$f_{i,t} = \begin{cases} f_{i,r_{i,t}}(y_{i,t}) & \text{if } S_{i,t} = 0, \\ f_{i,r_{i,t}}(\tilde{y}_{i,t}) & \text{if } S_{i,t} = NE, \\ 0 & \text{if } S_{i,t} = BR. \end{cases} \tag{25}$$

Thus, the value of $f_{i,t}$ depends on the state $(\theta_{n,t})_{n \in I}$, the target action profile r_t (which is drawn from distribution $\bar{\alpha}$ as described above), the public signal y_t, and the additional variable $\tilde{y}_t$.[26] Later in the proof, $f_{i,t}$ will be a component of the "reward" earned by player i in period t, which will be reflected in player i's end-of-block continuation payoff function $\psi : H^{T+1} \to \mathbb{R}$.

We can finally define $\theta_{i,t}$ as

$$\theta_{i,t} = \mathbf{1}\left\{\exists t' \leq t : \left|\sum_{t''=1}^{t'-1} \delta^{t''-1} f_{i,t''}\right| \geq \bar{f}\right\}. \tag{26}$$

That is, $\theta_{i,t}$ is the indicator function for the event that the magnitude of the component of player i's reward captured by $(f_{i,t''})_{t''=1}^{t'-1}$ exceeds $\bar{f}$ at any time $t' \leq t$.

This completes the definition of the equilibrium block strategy profile σ. Before proceeding further, we note that a unilateral deviation from σ by any player i does not affect the distribution of the state vector $\left((\theta_{j,t})_{j \neq i}\right)_{t=1}^{T}$. (However, such a deviation does affect the distribution of $(\theta_{i,t})_{t=1}^{T}$.)

Lemma 6 *For any player i and block strategy $\tilde{\sigma}_i$, the distribution of the random vector $\left((\theta_{j,t})_{j \neq i}\right)_{t=1}^{T}$ is the same under block strategy profile $(\tilde{\sigma}_i, \sigma_{-i})$ as under block strategy profile σ.*

Proof. Since $\theta_{j,t} = 1$ implies $\theta_{j,t+1} = 1$, it suffices to show that, for each t, each $J \subseteq I \backslash \{i\}$, each h^t such that $J = \{j \in I \backslash \{i\} : \theta_{j,t} = 0\}$, and each z_t, the probability $\Pr\left((\theta_{j,t+1})_{j \in J} | h^t, z_t, a_{i,t}\right)$ is independent of $a_{i,t}$. Since $\theta_{j,t+1}$ is determined by h^t and $f_{j,t}$, it is enough to show that $\Pr\left((f_{j,t})_{j \in J} | h^t, z_t, a_{i,t}\right)$ is independent of $a_{i,t}$.

Recall that $S_{j,t}$ is determined by h^t, and that if $j \in J$ (that is, $\theta_{j,t} = 0$) then $S_{j,t} \in \{0, NE\}$. If $S_{j,t} = 0$ then player j takes $r_{j,t}$, which is determined by z_t, $y_{j,t}$ is distributed according to $p_j(y_{j,t}|r_{j,t})$, and $f_{j,t}$ is determined by $y_{j,t}$, independently across players conditional on z_t. If $S_{j,t} = NE$ then $\tilde{y}_{j,t}$ is distributed according to $p_j(\tilde{y}_{j,t}|r_{j,t})$, where $r_{j,t}$ is determined by z_t, and $f_{j,t}$ is determined by $\tilde{y}_{j,t}$, independently across players conditional on z_t. Thus, $\Pr\left((f_{j,t})_{j \in J} | h^t, z_t, a_{i,t}\right) = \prod_{j \neq i} \Pr(f_{j,t}|S_{j,t}, r_{j,t})$, which is independent of $a_{i,t}$ as desired. ■

Continuation Value Function: We now construct the continuation value function $\psi : H^{T+1} \to \mathbb{R}^N$. For each player i and end-of-block history h^{T+1}, player i's continuation

[26] Intuitively, introducing the variable $\tilde{y}_t$, rather than simply using $y_{i,t}$ everywhere in (25), ensures that the distribution of $f_{i,t}$ does not depend on player i's opponents' strategies.

value $\psi_i\left(h^{T+1}\right)$ will be defined as the sum of T "rewards" $\psi_{i,t}$, where $t=1,\ldots,T$, and a constant term c_i that does not depend on h^{T+1}.

The rewards $\psi_{i,t}$ are defined as follows:

1. If $\theta_{j,t}=0$ for all $j\in I$, then
$$\psi_{i,t}=\delta^{t-1}f_{i,r_{i,t}}\left(y_{i,t}\right). \tag{27}$$

2. If $\theta_{i,t}=1$ and $\theta_{j,t}=0$ for all $j\neq i$, then
$$\psi_{i,t}=\delta^{t-1}\left(u_i\left(\bar{\alpha}\right)-u_i\left(\alpha_t^*\right)\right). \tag{28}$$

3. Otherwise,
$$\psi_{i,t}=\delta^{t-1}\left(-\zeta_i\bar{u}-u_i\left(\alpha_t^*\right)+\mathbf{1}\left\{S_{i,t}=0\right\}f_{i,r_{i,t}}\left(y_{i,t}\right)\right). \tag{29}$$

The constant c_i is defined as

$$c_i=-\mathbb{E}\left[\sum_{t=1}^{T}\delta^{t-1}\left(\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t}=0\right\}u_i\left(\bar{\alpha}\right)-\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t}=1\right\}\zeta_i\bar{u}\right)\right]+\frac{1-\delta^T}{1-\delta}v_i^*. \tag{30}$$

Note that, since $u_i\left(\bar{\alpha}\right)$ and v_i^* are both feasible payoffs, we have

$$\left|c_i\right|\leq 2\bar{u}\frac{1-\delta^T}{1-\delta}. \tag{31}$$

Finally, for each i and h^{T+1}, player i's continuation value at end-of-block history h^{T+1} is defined as

$$\psi_i\left(h^{T+1}\right)=c_i+\sum_{t=1}^{T}\psi_{i,t}. \tag{32}$$

E.3 Verification of the Equilibrium Conditions

We now verify that σ and ψ satisfy promise keeping, incentive compatibility, and self generation. We first show that $\theta_{i,t}=0$ for all i and t with high probability, and then verify the three desired conditions in turn.

Lemma 7 *We have*

$$\Pr\left(\max_{i\in I,t\in\{1,\ldots,T\}}\theta_{i,t}=0\right)\geq 1-\frac{\varepsilon}{20\bar{u}}. \tag{33}$$

Proof. By union bound, it suffices to show that, for each i, $\Pr\left(\max_{t\in\{1,\ldots,T\}}\theta_{i,t}=1\right)\leq\varepsilon/20\bar{u}N$, or equivalently

$$\Pr\left(\max_{t\in\{1,\ldots,T\}}\left|\sum_{t'=1}^{t}\delta^{t-1}f_{i,t'}\right|\geq\bar{f}\right)\leq\frac{\varepsilon}{20\bar{u}N}. \tag{34}$$

To see this, let $\tilde{f}_{i,t} = f_{i,r_{i,t}}(\tilde{y}_{i,t})$. Note that the variables $\left(\tilde{f}_{i,t}\right)_{t=1}^{T}$ are independent (unlike the variables $(f_{i,t})_{t=1}^{T}$). Since $\left(\tilde{f}_{i,t'}\right)_{t'=1}^{t}$ and $(f_{i,t'})_{t'=1}^{t}$ have the same distribution if $S_{i,t} \neq BR$, while $f_{i,t} = 0$ if $S_{i,t} = BR$, we have

$$\Pr\left(\max_{t\in\{1,\ldots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} f_{i,t'}\right| \geq \bar{f}\right) \leq \Pr\left(\max_{t\in\{1,\ldots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} \tilde{f}_{i,t'}\right| \geq \bar{f}\right). \tag{35}$$

Since $\left(\tilde{f}_{i,t}\right)_{t=1}^{T}$ are independent, Etemadi's inequality (Billingsley, 1995; Theorem 22.5) implies that

$$\Pr\left(\max_{t\in\{1,\ldots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} \tilde{f}_{i,t'}\right| \geq \bar{f}\right) \leq 3 \max_{t\in\{1,\ldots,T\}} \Pr\left(\left|\sum_{t'=1}^{t} \delta^{t-1} \tilde{f}_{i,t'}\right| \geq \frac{\bar{f}}{3}\right). \tag{36}$$

Letting $x_{i,t} = \delta^{t-1}\tilde{f}_{i,t}$, note that $|x_{i,t}| \leq 2\bar{u}/\eta^2$ with probability 1 by (14), $\mathbb{E}[x_{i,t}] = 0$ by (12), and

$$\operatorname{Var}\left(\sum_{t'=1}^{t} x_{i,t'}\right) = \sum_{t'=1}^{t} \operatorname{Var}(x_{i,t'}) \leq \sum_{t'=1}^{T} \operatorname{Var}(x_{i,t'}) = \frac{1-\delta^T}{1-\delta}\frac{\bar{u}^2}{\eta^2} \quad \text{by (13).}$$

Therefore, by Bernstein's inequality ((15), which again applies because $\left(\tilde{f}_{i,t}\right)_{t=1}^{T}$ are independent) and (16), we have, for each $t \leq T$,

$$\Pr\left(\left|\sum_{t'=1}^{t} \delta^{t'-1} \tilde{f}_{i,t'}\right| \geq \frac{\bar{f}}{3}\right) \leq \frac{\varepsilon}{60\bar{u}N}. \tag{37}$$

Finally, (35), (36), and (37) together imply (34). ■

Incentive Compatibility: We use the following lemma (proof in Appendix E.5).

Lemma 8 *For each player i and block strategy profile σ, incentive compatibility holds (i.e., (22) is satisfied) if and only if*

$$\operatorname{supp}\sigma_i\left(h^t\right) \subseteq \operatorname*{argmax}_{a_{i,t}\in A_i} \mathbb{E}^{\sigma_{-i}}\left[\delta^{t-1}u_{i,t} + \psi_{i,t}|h^t, a_{i,t}\right] \quad \textit{for all } t \textit{ and } h^t. \tag{38}$$

In addition, for all $t \leq t'$ and h^t, we have

$$\mathbb{E}^{\sigma}\left[\delta^{t'-1}u_{i,t} + \psi_{i,t'}|h^t\right] = \mathbb{E}^{\sigma}\left[\delta^{t'-1}\left(\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t'} = 0\right\} u_i(\bar{\alpha}) - \mathbf{1}\left\{\max_{j\neq i}\theta_{j,t'} = 1\right\}\zeta_i\bar{u}\right)|h^t\right]. \tag{39}$$

We now verify that (38) holds. Fix a player i, period t, and history h^t. We consider several cases, which parallel the definition of the reward $\psi_{i,t}$.

1. If $\theta_{j,t} = 0$ for all $j \in I$, recall that the equilibrium action profile is the r_t that is prescribed by public randomization z_t. For each action $a_i \neq r_{i,t}$, by (11) and (27), and

recalling that $\bar{u} \geq \max_a u_i(a) - \min_a u_i(a)$, we have

$$\begin{aligned}
& \mathbb{E}^{\sigma_{-i}}\left[\delta^{t-1}u_{i,t} + \psi_{i,t}|h^t, z_t, a_{i,t} = r_{i,t}\right] - \mathbb{E}^{\sigma_{-i}}\left[\delta^{t-1}u_{i,t} + \psi_{i,t}|h^t, z_t, a_{i,t} = a_i\right] \\
= \; & \delta^{t-1}\left(\mathbb{E}\left[u_i(r_t) + f_{i,r_{i,t}}(y_{i,t})\,|a_{i,t} = r_{i,t}\right] - \mathbb{E}\left[u_i(a_i, r_{-i,t}) + f_{i,r_{i,t}}(y_{i,t})\,|a_{i,t} = a_i\right]\right) \\
\leq \; & 0, \quad \text{so (38) holds.}
\end{aligned}$$

2. If $\theta_{i,t} = 1$ and $\theta_{j,t} = 0$ for all $j \neq i$, then the reward $\psi_{i,t}$ specified by (28) does not depend on $y_{i,t}$. Hence, (38) reduces to the condition that every action in $\operatorname{supp}\sigma_i(h^t)$ is a static best responses to $\sigma_{-i}(h^t)$. This conditions holds for the prescribed action profile, $(r'_i \in BR_i(r_{-i,t}), r_{-i,t})$ or $\varrho_{i,t}^{NE}$.

3. Otherwise: (a) If $S_{i,t} = 0$, then (38) holds because it holds in Case 1 above and (27) and (29) differ only by a constant independent of $y_{i,t}$. (b) If $S_{i,t} \neq 0$, then either $\theta_{j,t} = \theta_{j',t} = 1$ for distinct players j, j', or there exists a unique player $j \neq i$ with $\theta_{j,t} = 1$, and for this player we have $\zeta_j = -1$. In both cases, ϱ_t^{NE} is prescribed. Since the reward $\psi_{i,t}$ specified by (29) does not depend on $y_{i,t}$, (38) reduces to the condition that every action in $\operatorname{supp}\sigma_i(h^t)$ is a static best responses to $\sigma_{-i}(h^t)$, which holds for the prescribed action profile ϱ_t^{NE}.

Promise Keeping: This essentially holds by construction: we have

$$\begin{aligned}
& \frac{1-\delta}{1-\delta^T}\mathbb{E}^{\sigma}\left[\sum_{t=1}^{T}\delta^{t-1}u_{i,t} + \psi_i\left(h^{T+1}\right)\right] \\
= \; & \frac{1-\delta}{1-\delta^T}\left(\mathbb{E}^{\sigma}\left[\sum_{t=1}^{T}\left(\delta^{t-1}u_{i,t} + \psi_{i,t}\right)\right] + c_i\right) \quad \text{(by (32))} \\
= \; & \frac{1-\delta}{1-\delta^T}\mathbb{E}^{\sigma}\left[\sum_{t=1}^{T}\delta^{t-1}\left(\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t} = 0\right\}u_i(\bar{\alpha}) - \mathbf{1}\left\{\max_{j\neq i}\theta_{j,t} = 1\right\}\zeta_i\bar{u}\right) + c_i\right] \quad \text{(by (39))} \\
= \; & v_i^* \quad \text{(by (30)), so (21) holds.}
\end{aligned}$$

Self Generation: We use the following lemma (proof in Appendix E.6).

Lemma 9 *For every end-of-block history h^{T+1}, we have*

$$\zeta_i\sum_{t=1}^{T}\psi_{i,t} \leq \bar{f} + \frac{2\bar{u}}{\eta^2} \quad \textit{and} \tag{40}$$

$$\left|\sum_{t=1}^{T}\psi_{i,t}\right| \leq \bar{f} + \frac{2\bar{u}}{\eta^2} + 2\bar{u}\frac{1-\delta^T}{1-\delta}. \tag{41}$$

In addition,

$$\zeta_i c_i \leq -\frac{1-\delta^T}{1-\delta}\frac{\varepsilon}{8}. \tag{42}$$

To establish self generation ((23)), it suffices to show that, for each h^{T+1}, $\zeta_i \psi_i \left(h^{T+1}\right) \leq 0$ and $\left|\psi_i \left(h^{T+1}\right)\right| \leq \left(\delta^T / (1-\delta)\right) \varepsilon$. This now follows because

$$\begin{aligned}
\zeta_i \psi_i \left(h^{T+1}\right) &= \zeta_i \left(c_i + \sum_{t=1}^{T} \psi_{i,t}\right) \leq -\frac{1-\delta^T}{1-\delta} \frac{\varepsilon}{8} + \bar{f} + 2\bar{u}/\eta^2 \quad \text{(by (40) and (42))} \\
&\leq \frac{1-\delta^T}{8(1-\delta)} \left(-\varepsilon + 8\left(\frac{1-\delta}{1-\delta^T}\right)\left(\bar{f} + 2\bar{u}/\eta^2\right)\right) \leq 0 \quad \text{(by (17)),} \quad \text{and} \\
\left|\psi_i \left(h^{T+1}\right)\right| &\leq |c_i| + \left|\sum_{t=1}^{T} \psi_{i,t}\right| \\
&\leq 4\bar{u}\frac{1-\delta^T}{1-\delta} + \bar{f} + 2\bar{u}/\eta^2 \quad \text{(by (31) and (41))} \\
&= \frac{1-\delta^T}{1-\delta} 4\bar{u} + \bar{f} + 2\bar{u}/\eta^2 \leq \frac{\delta^T}{1-\delta}\varepsilon \quad \text{(by (18)),}
\end{aligned}$$

which completes the proof.

E.4 Omitted Details for the Proof of Lemma 5

We show that, with the stated definitions of T and $\bar{f}$, (19) and (20) imply (16)–(18). First, note that

$$\frac{1-\delta^2}{1-\delta^{2T}} = \frac{(1+\delta)(1-\delta)}{\left(1+\delta^T\right)\left(1-\delta^T\right)} < 2\frac{1-\delta}{1-\delta^T}.$$

Hence,

$$\begin{aligned}
\frac{2\bar{f}\left(1-\delta^2\right)}{9\bar{u}\left(1-\delta^{2T}\right)} &< \frac{4}{9\bar{u}}\frac{1-\delta}{1-\delta^T}\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta^T}{1-\delta}\frac{\bar{u}^2}{\eta^2}} \\
&= \frac{4}{9}\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta^2}} \leq 1 \quad \text{(by (19)).}
\end{aligned}$$

Therefore,

$$60\bar{u}N\exp\left(\frac{-\left(\frac{\bar{f}}{3}\right)^2\eta^2}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2 + \frac{2}{3}\frac{\bar{f}}{3}\bar{u}\right)}\right) \leq 60\bar{u}N\exp\left(\frac{-\left(\frac{\bar{f}}{3}\right)^2\eta^2}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2 + \frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2\right)}\right) = 60\bar{u}N\exp\left(\frac{-\bar{f}^2\eta^2}{36\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2}\right).$$

Moreover,

$$\frac{\bar{f}^2\eta^2}{36\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2} = \frac{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta^T}{1-\delta}}{36\frac{1-\delta^{2T}}{1-\delta^2}} = \frac{1+\delta}{1+\delta^T}\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N) \geq \log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N).$$

Hence, we have

$$60\bar{u}N\exp\left(\frac{-\left(\frac{\bar{f}}{3}\right)^2\eta^2}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2+\frac{2}{3}\frac{\bar{f}}{3}\bar{u}\right)}\right)\le 60\bar{u}N\exp\left(-\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log\left(N\right)\right)=\varepsilon.$$

This establishes (16).

Next, we have

$$8\frac{1-\delta}{1-\delta^T}\left(\bar{f}+\frac{2\bar{u}}{\eta^2}\right)=8\bar{u}\left(\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log\left(N\right)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta^2}}+\frac{1-\delta}{1-\delta^T}\frac{2}{\eta^2}\right)\le\varepsilon\quad\text{(by (20))}. \tag{43}$$

This establishes (17).

Finally, by (43) and $8\bar{u}\left(1-\delta^T\right)/\delta^T\le\varepsilon$, we have

$$4\bar{u}\frac{1-\delta^T}{\delta^T}+\frac{1-\delta}{\delta^T}\left(\bar{f}+\frac{2\bar{u}}{\eta^2}\right)=4\bar{u}\frac{1-\delta^T}{\delta^T}+\frac{1-\delta^T}{\delta^T}\frac{1-\delta}{1-\delta^T}\left(\bar{f}+\frac{2\bar{u}}{\eta^2}\right)\le 4\frac{\varepsilon}{8}+\frac{\varepsilon}{8}\frac{\varepsilon}{8}\le\varepsilon.$$

This establishes (18).

E.5 Proof of Lemma 8

We show that player i has a profitable one-shot deviation from σ_i at some history h^t if and only if (38) is violated at h^t. To see this, we first calculate player i's continuation payoff under σ from period $t+1$ onward (net of the constant c_i and the rewards already accrued $\sum_{t'=1}^{t}\psi_{i,t'}$). For each $t'\ge t+1$, there are several cases to consider.

1. If $\theta_{j,t'}=0$ for all j, then by (12) and (27) we have
$\mathbb{E}^\sigma\left[\delta^{t'-1}u_{i,t'}+\psi_{i,t'}|h^{t'}\right]=\delta^{t'-1}\left(u_i\left(\alpha^*_{t'}\right)+\mathbb{E}\left[f_{i,r_{i,t'}}\left(y_{i,t'}\right)|r_{i,t'}\right]\right)=\delta^{t'-1}u_i\left(\bar{\alpha}\right)$.

2. If $\theta_{i,t'}=1$ and $\theta_{j,t'}=0$ for all $j\ne i$, then by (28) we have
$\mathbb{E}^\sigma\left[\delta^{t'-1}u_{i,t'}+\psi_{i,t'}|h^{t'}\right]=\delta^{t'-1}\left(u_i\left(\alpha^*_{t'}\right)+u_i\left(\bar{\alpha}\right)-u_i\left(\alpha^*_{t'}\right)\right)=\delta^{t'-1}u_i\left(\bar{\alpha}\right)$.

3. Otherwise: (a) If $S_{i,t'}=0$, then by (12) and (29) (and recalling that player i's equilibrium action is $r_{i,t'}$ when $S_{i,t'}=0$) we have
$\mathbb{E}^\sigma\left[\delta^{t'-1}u_{i,t'}+\psi_{i,t'}|h^{t'}\right]=\delta^{t'-1}\left(u_i\left(\alpha^*_{t'}\right)-\zeta_i\bar{u}-u\left(\alpha^*_{t'}\right)+\mathbb{E}\left[f_{i,r_{i,t'}}\left(y_{i,t'}\right)|r_{i,t'}\right]\right)=\delta^{t'-1}\left(-\zeta_i\bar{u}\right)$.
(b) If $S_{i,t'}\ne 0$, then by (29) we have
$\mathbb{E}^\sigma\left[\delta^{t'-1}u_{i,t'}+\psi_{i,t'}|h^{t'}\right]=\delta^{t'-1}\left(u_i\left(\alpha^*_{t'}\right)-\zeta_i\bar{u}-u\left(\alpha^*_{t'}\right)\right)=\delta^{t'-1}\left(-\zeta_i\bar{u}\right)$.

In total, (39) holds, and player i's net continuation payoff under σ from period $t+1$ onward equals

$$\mathbb{E}^\sigma\left[\sum_{t'=t+1}^{T}\delta^{t'-1}\left(\mathbf{1}\left\{\max_{j\ne i}\theta_{j,t'}=0\right\}u_i\left(\bar{\alpha}\right)-\mathbf{1}\left\{\max_{j\ne i}\theta_{j,t'}=1\right\}\zeta_i\bar{u}\right)|h^t\right].$$

By Lemma 6, the distribution of $\left((\theta_{n,t'})_{n\neq i}\right)_{t'=t+1}^{T}$ does not depend on player i's period-t action, and hence neither does player i's net continuation payoff under σ from period $t+1$ onward. Therefore, player i's period-t action $a_{i,t}$ maximizes her continuation payoff from period t onward if and only if it maximizes $\mathbb{E}^{\sigma_{-i}}[\delta^{t-1}u_{i,t}+\psi_{i,t}|h^t,a_{i,t}]$.

E.6 Proof of Lemma 9

Define

$$\begin{aligned}\psi_{i,t}^{v} &= \begin{cases}\delta^{t-1}\left(u_i\left(\bar{\alpha}\right)-u_i\left(\alpha_t^*\right)\right) & \text{if } \theta_{j,t}=0 \text{ for all } j\neq i,\\ \delta^{t-1}\left(-\zeta_i\bar{u}-u_i\left(\alpha_t^*\right)\right) & \text{otherwise,}\end{cases} \quad \text{and}\\ \psi_{i,t}^{f} &= \begin{cases}\delta^{t-1}f_{i,a_{i,t}}\left(y_{i,t}\right) & \text{if either } \theta_{j,t}=0 \text{ for all } j \text{ or } S_{i,t}=0,\\ 0 & \text{otherwise.}\end{cases}\end{aligned}$$

Note that, by (27)–(29), we can write $\psi_{i,t}=\psi_{i,t}^{v}+\psi_{i,t}^{f}$. (Note that, if $\theta_{n,t}=0$ for all $n\in I$, we have $\alpha_t^*=\bar{\alpha}$ and hence $\psi_{i,t}^{v}+\psi_{i,t}^{f}=\delta^{t-1}f_{i,a_{i,t}}\left(y_{i,t}\right)$, as specified in (27).) We show that, for every end-of-block history h^{T+1}, we have

$$\zeta_i\sum_{t=1}^{T}\psi_{i,t}^{v} \in \left[-2\bar{u}\frac{1-\delta^T}{1-\delta},0\right] \quad \text{and} \tag{44}$$

$$\left|\zeta_i\sum_{t=1}^{T}\psi_{i,t}^{f}\right| \leq \bar{f}+\frac{2\bar{u}}{\eta^2}. \tag{45}$$

Since $\psi_{i,t}=\psi_{i,t}^{v}+\psi_{i,t}^{f}$, (44) and (45) imply (40) and (41), which proves the first part of the lemma.

For (44), note that, by definition of the prescribed equilibrium actions, if $\theta_{j,t}=0$ for all $j\neq i$, then (i) if $\zeta_i=1$, we have $u_i\left(\alpha_t^*\right)\geq\sum_a\bar{\alpha}\left(a\right)\min\left\{u_i\left(a\right),\max_{a_i'}u_i\left(a_i',a_{-i}\right)\right\}\geq u_i\left(\bar{\alpha}\right)$; and (ii) if $\zeta_i=-1$, we have $u_i\left(\alpha_t^*\right)\leq\max\left\{u_i\left(\bar{\alpha}\right),u_i\left(\alpha^{NE}\right)\right\}=u_i\left(\bar{\alpha}\right)$. In total, we have $\zeta_i\left(u_i\left(\bar{\alpha}\right)-u_i\left(\alpha_t^*\right)\right)\leq 0$. Since obviously $\zeta_i\left(u_i\left(\bar{\alpha}\right)-u_i\left(\alpha_t^*\right)\right)\geq-2\bar{u}$ and $-\bar{u}-\zeta_iu_i\left(\alpha_t^*\right)\geq-2\bar{u}$, we have

$$\zeta_i\psi_{i,t}^{v}=\begin{cases}\delta^{t-1}\zeta_i\left(u_i\left(\bar{\alpha}\right)-u_i\left(\alpha_t^*\right)\right) & \text{if } \theta_{j,t}=0 \text{ for all } j\neq i,\\ \delta^{t-1}\left(-\bar{u}-\zeta_iu_i\left(\alpha_t^*\right)\right) & \text{otherwise}\end{cases} \in \left[-2\bar{u}\delta^{t-1},0\right].$$

For (45), note that $S_{i,t}=0$ implies $\theta_{i,t}=0$, and hence

$$\left|\zeta_i\sum_{t=1}^{T}\psi_{i,t}^{f}\right|\leq\left|\zeta_i\sum_{t=1}^{T}\mathbf{1}\left\{\theta_{i,t}=0\right\}\delta^{t-1}f_{i,a_{i,t}}\left(y_{i,t}\right)\right|.$$

Since $\theta_{i,t+1}=1$ whenever $\left|\sum_{t'=1,..,t}\delta^{t-1}f_{i,a_{i,t}}\left(y_{i,t}\right)\right|\geq\bar{f}$, and in addition $\left|f_{i,a_{i,t}}\left(y_{i,t}\right)\right|\leq 2\bar{u}/\eta^2$ by (14), this inequality implies (45).

For the second part of the lemma, by (30), we have

$$\begin{aligned}
\zeta_i c_i &= \zeta_i \left(-\mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} u_i(\bar{\alpha}) - \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} \zeta_i \bar{u} \right) \right] + \frac{1-\delta^T}{1-\delta} v_i^* \right) \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} \zeta_i \left(v_i^* - u_i(\bar{\alpha}) \right) + \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} \underbrace{\left(\bar{u} + \zeta_i v_i^* \right)}_{\in[0,2\bar{u}]} \right) \right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} \left(\frac{-\varepsilon}{2} \right) + \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} 2\bar{u} \right) \right] \quad \text{by (24)} \\
&\leq -\frac{1-\delta^T}{1-\delta} \left(\left(1 - \frac{\varepsilon}{20\bar{u}} \right) \frac{\varepsilon}{2} + \left(\frac{\varepsilon}{20\bar{u}} \right) 2\bar{u} \right) \quad \text{(by (33))} \\
&\leq -\frac{1-\delta^T}{1-\delta} \frac{\varepsilon}{8} \quad (\text{as } \varepsilon < \bar{u}/2).
\end{aligned}$$

F Proof of Theorem 3

Fix a linear PPE with coefficients $b = (1, b_2, \ldots, b_N)$, where $|b_i| \leq 1$ for all i. Let $I^+ = \{i : b_i \geq 0\}$ and $I^- = \{i : b_i < 0\}$. Define

$$\underline{v}_i = \begin{cases} \inf_h w_i(h) & \text{if } i \in I^+, \\ \sup_h w_i(h) & \text{if } i \in I^-, \end{cases} \quad \text{and} \quad \bar{v}_i = \begin{cases} \sup_h w_i(h) & \text{if } i \in I^+, \\ \inf_h w_i(h) & \text{if } i \in I^-. \end{cases}$$

Since $W(\varepsilon)$ is convex, it suffices to show that $\underline{v}, \bar{v} \in W(\varepsilon)$.

In the following lemma, given $\alpha \in \Delta(A)$ and a function $\omega : A \times Y \to \mathbb{R}$, $\mathbb{E}^\alpha[\omega(r,y)]$ denotes expectation where $r \sim \alpha$ and then $y \sim p(\cdot|r)$, and $\mathbb{E}^{\alpha,a_i'}[\omega(r,y)]$ denotes expectation where $r \sim \alpha$ and then $y \sim p(\cdot|a_i', r_{-i})$.

Lemma 10 *There exist $\alpha \in \Delta(A)$ and $\omega : A \times Y \to \mathbb{R}$ such that*

$$\begin{aligned}
\bar{v} &= \mathbb{E}^\alpha\left[u(r) - b\omega(r,y) \right], \\
\mathbb{E}^\alpha\left[u_i(r) - b_i\omega(r,y) | r_i = a_i \right] &\geq \mathbb{E}^{\alpha,a_i'}\left[u_i(a_i', r_{-i}) - b_i\omega(r,y) | r_i = a_i \right] \quad \textit{for all } i, a_i \in \operatorname{supp} \alpha_i, a_i' \in A_i, \\
\omega(r,y) &\in \left[0, \frac{\delta}{1-\delta}\bar{u} \right] \quad \textit{for all } r, y.
\end{aligned}$$

Moreover, if the constraint $\omega(r,y) \in \left[0, \frac{\delta}{1-\delta}\bar{u}\right]$ is replaced with $\omega(r,y) \in \left[-\frac{\delta}{1-\delta}\bar{u}, 0\right]$, then the same statement holds with $\underline{v}$ in place of $\bar{v}$.

Proof. Let $E = \{(1-\beta)\underline{v} + \beta\bar{v} : \beta \in [0,1]\}$. By standard arguments, E is self-generating: for any $v \in E$, there exist $\alpha \in \Delta(A)$ and $w : A \times Y \to E$ such that

$$\begin{aligned}
v &= \mathbb{E}^\alpha\left[u(r) + \delta w(r,y) \right] \quad \text{and} \\
\mathbb{E}^\alpha\left[u_i(r) + \delta w_i(r,y) | r_i = a_i \right] &\geq \mathbb{E}^{\alpha,a_i'}\left[u_i(a_i', r_{-i}) + \delta w_i(r,y) | r_i = a_i \right] \quad \text{for all } i, a_i \in \operatorname{supp} \alpha_i, a_i' \in A_i.
\end{aligned}$$

Since $v \in E$ and $w(r,y) \in E$ for all r,y, we have $v_i - w_i(r,y) = b_i(v_1 - w_1(r,y))$ for all i,r,y. Since $\bar{v}_1 \geq v_1$ for all $v \in E$, if $v = \bar{v}$ then $w_1(r,y) \leq v_1$ for all r,y. Hence, taking $v = \bar{v} = (1-\delta)u(\alpha) + \delta b\mathbb{E}[w(r,y)|\alpha]$ and defining $\omega(r,y) = \frac{\delta}{1-\delta}(\bar{v}_1 - w_1(r,y)) \in \left[0, \frac{\delta}{1-\delta}\bar{u}\right]$ for all r,y, and letting $\mathbb{E}[\cdot]$ denote expectation where $y \sim p(\cdot|a)$, we have, for all a,r,

$$
\begin{aligned}
u(a) - b\mathbb{E}[\omega(r,y)] &= u(a) - b\mathbb{E}\left[\frac{\delta}{1-\delta}(\bar{v}_1 - w_1(r,y))\right] \\
&= u(a) - \mathbb{E}\left[\frac{\delta}{1-\delta}(\bar{v} - w(r,y))\right] = (1-\delta)u(a) + \delta\mathbb{E}[w(r,y)],
\end{aligned}
$$

and the result follows. Similarly, if $v = \underline{v}$ then $w_1(r,y) \geq v_1$ for all r,y, and the symmetric conclusion holds. ■

Taking α and ω as in Lemma 10, we have, for any player i and manipulation s_i,

$$
\begin{aligned}
g_i(s_i,\alpha) &\leq \sum_{a_i} \alpha_i(a_i)\left(\mathbb{E}^{\alpha,s_i(a_i)}[b_i\omega(r,y)|r_i = a_i] - \mathbb{E}^{\alpha}[b_i\omega(r,y)|r_i = a_i]\right) \\
&\leq \sum_{r} \alpha(r)\max_{a_i}|\mathbb{E}[\omega(r,y)|r,a_i] - \mathbb{E}[\omega(r,y)|r]|,
\end{aligned}
$$

where the second inequality uses $|b_i| \leq 1$. Hence,

$$
\begin{aligned}
\frac{1}{N}\sum_i \bar{g}_i(\alpha) &\leq \frac{1}{N}\sum_i\sum_r \alpha(r)\max_{a_i}|\mathbb{E}[\omega(r,y)|r,a_i] - \mathbb{E}[\omega(r,y)|r]| \\
&\leq \max_{r,a}\frac{1}{N}\sum_i |\mathbb{E}[\omega(y)|a_i,r_{-i}] - \mathbb{E}[\omega(y)|r]|.
\end{aligned}
$$

We conclude that $\sum_i \bar{g}_i(\alpha)/N$ is bounded by the solution to the program

$$
\begin{aligned}
\max_{(Y,p),r,a,\omega} \frac{1}{N}\sum_i |\mathbb{E}[\omega(y)|a_i,r_{-i}] - \mathbb{E}[\omega(y)|r]| \quad &\text{s.t.} \\
\omega(y) \in \left[0, \frac{\delta}{1-\delta}\bar{u}\right] \quad &\text{for all } y, \\
\mathbb{E}[\omega(y)|r] \leq \bar{u}&,
\end{aligned}
$$

where the last line holds because $\mathbb{E}[\omega(y)|r] = u_1(r) - \bar{v}_1 \leq \bar{u}$. The remainder of the proof shows that the value of this program converges to 0 if $(1-\delta)\exp(N^{1-\rho}) \to \infty$ for $\rho > 0$.

We first consider the sub-program where (Y,p) is fixed, so the objective is maximized over (r,a,ω). Recall that $p(y|a) = \sum_x \pi_{a,x} q(y|x)$. By Blackwell's theorem, the value of the sub-program with signal distribution p is greater than that with signal distribution $\hat{p}$, if $\hat{p}$ is a *garbling* of p. (That is, there exists a Markov matrix M such that $\hat{p} = Mp$.) Consequently, it is without loss to let $Y = X$ and $q(y|x) = \mathbf{1}\{y = x\}$ for all y,x, so that $p(x|a) = \pi_{a,x}$ for all a, x. Now fix $r, a \in A$, and for each i, define $\bar{X} = A$,

$$\bar{\pi}^i_{a_i,x_i} = \begin{cases} 1-\underline{\pi} & \text{if } x_i = a_i, \\ \underline{\pi} & \text{if } x_i = r_i, \\ 0 & \text{otherwise,} \end{cases} \qquad \bar{\pi}^i_{r_i,x_i} = \begin{cases} 1-\underline{\pi} & \text{if } x_i = r_i, \\ \underline{\pi} & \text{if } x_i = a_i, \\ 0 & \text{otherwise,} \end{cases} \qquad \bar{\pi}^i_{\tilde{a}_i,x_i} = \mathbf{1}\{x_i = \tilde{a}_i\} \text{ for } \tilde{a}_i \notin \{a_i, r_i\},$$

and finally $\bar{\pi}_{\tilde{a},x} = \prod_i \bar{\pi}^i_{\tilde{a}_i,x_i}$ for all $\tilde{a}, x$. The following lemma implies that the value of our program is upper-bounded by that with $X = \bar{X}$ and $\pi = \bar{\pi}$.

Lemma 11 *π is a garbling of $\bar{\pi}$.*

Proof. Since (x_i) are independent conditional on (a_i), it suffices to show that π^i is a garbling of $\bar{\pi}^i$ for each i. Since $\underline{\pi} < 1/2$, the matrix $\bar{\pi}^i$ is invertible, with inverse matrix $\hat{\pi}^i$ given by

$$\hat{\pi}^i_{a_i,\hat{a}_i} = \begin{cases} \frac{1-\underline{\pi}}{1-2\underline{\pi}} & \text{if } \hat{a}_i = a_i, \\ -\frac{\underline{\pi}}{1-2\underline{\pi}} & \text{if } \hat{a}_i = r_i, \\ 0 & \text{otherwise,} \end{cases} \qquad \hat{\pi}^i_{r_i,\hat{a}_i} = \begin{cases} \frac{1-\underline{\pi}}{1-2\underline{\pi}} & \text{if } \hat{a}_i = r_i, \\ -\frac{\underline{\pi}}{1-2\underline{\pi}} & \text{if } \hat{a}_i = a_i, \\ 0 & \text{otherwise,} \end{cases} \qquad \hat{\pi}^i_{\tilde{a}_i,\hat{a}_i} = \mathbf{1}\{\hat{a}_i = \tilde{a}_i\} \text{ for } \tilde{a}_i \notin \{a_i, r_i\}.$$

The matrix $M^i := \pi^i \hat{\pi}^i$ is easily calculated as

$$M^i_{\hat{a}_i,x_i} = \begin{cases} \pi^i_{\hat{a}_i,x_i} \frac{1-\underline{\pi}}{1-2\underline{\pi}} - \left(1 - \pi^i_{\hat{a}_i,x_i}\right) \frac{\underline{\pi}}{1-2\underline{\pi}} & \text{if } \hat{a}_i \in \{a_i, r_i\}, \\ \pi^i_{\hat{a}_i,x_i} & \text{otherwise.} \end{cases}$$

Note that, for $\hat{a}_i \in \{a_i, r_i\}$,

$$\sum_{x_i} M^i_{\hat{a}_i,x_i} = \frac{|A_i| - 1 - \underline{\pi}}{|A_i| - 1 - |A_i|\,\underline{\pi}} - (|A_i| - 1) \frac{\underline{\pi}}{|A_i| - 1 - |A_i|\,\underline{\pi}} = 1,$$

and clearly $\sum_{x_i} M^i_{\hat{a}_i,x_i} = 1$ for $\hat{a}_i \notin \{a_i, r_i\}$. In addition, since $\pi^i_{\hat{a}_i,x_i} \geq \underline{\pi}$ for all $\hat{a}_i, x_i$, we have

$$\frac{\pi^i_{\hat{a}_i,x_i}(1-\underline{\pi}) - \left(1 - \pi^i_{\hat{a}_i,x_i}\right)\underline{\pi}}{1 - 2\underline{\pi}} \geq \frac{\underline{\pi}(1-\underline{\pi}) - (1-\underline{\pi})\underline{\pi}}{|X_i| - 1 - |X_i|\,\underline{\pi}} = 0,$$

and clearly $M^i_{\hat{a}_i,x_i} \leq 1$ for all $\hat{a}_i, x_i$. So M^i is a Markov matrix and $\pi^i = M^i \bar{\pi}^i$, completing the proof. ■

Given Lemma 11, our program simplifies to

$$\max_{r,a,\omega} \frac{1}{N} \sum_i |\mathbb{E}[\omega(x)\,|a_i, r_{-i}] - \mathbb{E}[\omega(x)\,|r]| \quad \text{s.t.} \tag{46}$$

$$\omega(x) \in \left[0, \frac{\delta}{1-\delta}\bar{u}\right] \quad \text{for all } x \in A, \tag{47}$$

$$\mathbb{E}[\omega(x)\,|r] \leq \bar{u}, \tag{48}$$

where x is distributed $\bar{\pi}_{\tilde{a},x}$. Note that, for $\tilde{a} = r$ or $\tilde{a} = (a_i, r_{-i})$ for some i, $\bar{\pi}_{\tilde{a},x} > 0$ iff $x \in \times_i \{a_i, r_i\}$. Note also that it is without loss to take $a_i \neq r_i$ for all i. For, if $a_i = r_i$ then the program becomes

$$\max_{a_{-i},r_{-i},\omega_{-i}:A_{-i}\to\mathbb{R}} \frac{1}{N}\sum_{j\neq i}\left|\mathbb{E}\left[\omega_{-i}\left(x_{-i}\right)|a_j,r_{-j}\right]-\mathbb{E}\left[\omega_{-i}\left(x_{-i}\right)|r\right]\right| \quad \text{s.t. (47), (48).}$$

Any feasible triple $(a_{-i},r_{-i},\omega_{-i})$ in this reduced program can be extended to a feasible triple (a,r,ω) with $a_i\neq r_i$ in the original program which gives the same value, by defining $\omega\left(x\right)=\omega_{-i}\left(x_{-i}\right)$ for all x. We thus assume that $a_i\neq r_i$ for all i.

We now show that the value of program (46)–(48) converges to 0, which completes the proof. Note that this value is less than the sum of the values of the two programs

$$\begin{aligned}
&\max_{r,a,\omega}\frac{1}{N}\sum_i\left(\mathbb{E}\left[\omega\left(x\right)|a_i,r_{-i}\right]-\mathbb{E}\left[\omega\left(x\right)|r\right]\right)_+ \quad \text{s.t. (47), (48), and}\\
&\max_{r,a,\omega}\frac{1}{N}\sum_i\left(\mathbb{E}\left[\omega\left(x\right)|r\right]-\mathbb{E}\left[\omega\left(x\right)|a_i,r_{-i}\right]\right)_+ \quad \text{s.t. (47), (48).}
\end{aligned}$$

We show that the value of the former program converges to 0. A symmetric argument shows that the value of the latter program also converges to 0, which implies that the value of program (46)–(48) converges to 0 as well, as desired.

Letting $\lambda\geq 0$ denote the multiplier on (48), the solution to the first program satisfies

$$\omega\left(x\right)=\begin{cases}\frac{\delta}{1-\delta}\bar{u} & \text{if } \frac{\left(\frac{1}{N}\sum_i\bar{\pi}_{(a_i,r_{-i}),x}\right)-\bar{\pi}_{r,x}}{\bar{\pi}_{r,x}}>\lambda,\\ 0 & \text{if } \frac{\bar{\pi}_{(a_i,r_{-i}),x}-\bar{\pi}_{r,x}}{\bar{\pi}_{r,x}}<\lambda,\end{cases}=\begin{cases}\frac{\delta}{1-\delta}\bar{u} & \text{if } \frac{1}{N}\sum_i\frac{\bar{\pi}_{(a_i,r_{-i}),x}}{\bar{\pi}_{r,x}}>\lambda+1,\\ 0 & \text{if } \frac{1}{N}\sum_i\frac{\bar{\pi}_{(a_i,r_{-i}),x}}{\bar{\pi}_{r,x}}<\lambda+1.\end{cases}$$

For all $x\in\times_i\{a_i,r_i\}$, we have

$$\frac{\bar{\pi}_{(a_i,r_{-i}),x}}{\bar{\pi}_{r,x}}=\begin{cases}\frac{1-\underline{\pi}}{\underline{\pi}} & \text{if } x_i=a_i,\\ \frac{\underline{\pi}}{1-\underline{\pi}} & \text{if } x_i=r_i.\end{cases}$$

Since $\frac{1-\underline{\pi}}{\underline{\pi}}>\frac{\underline{\pi}}{1-\underline{\pi}}$ (as $\underline{\pi}<1/2$), it follows that there exists $n^*\in\{0,1,\ldots,N\}$ and $\beta\in[0,1]$ such that

$$\omega\left(x\right)=\begin{cases}\frac{\delta}{1-\delta}\bar{u} & \text{if } \{i:x_i=a_i\}>n^*,\\ \beta\frac{\delta}{1-\delta}\bar{u} & \text{if } \{i:x_i=a_i\}=n^*,\\ 0 & \text{if } \{i:x_i=a_i\}<n^*.\end{cases}$$

Let $n=|\{i:x_i=a_i\}|$ and let $n_{-i}=|\{j\neq i:x_j=a_j\}|$. Note that, for any n^*,

$$\begin{aligned}
\Pr\left(n=n^*|a_i,r_{-i}\right) &= \left(1-\underline{\pi}\right)\Pr\left(n_{-i}=n^*-1|r_{-i}\right)+\underline{\pi}\Pr\left(n_{-i}=n^*|r_{-i}\right), \quad \text{and}\\
\Pr\left(n=n^*|r\right) &= \underline{\pi}\Pr\left(n_{-i}=n^*-1|r_{-i}\right)+\left(1-\underline{\pi}\right)\Pr\left(n_{-i}=n^*|r_{-i}\right),
\end{aligned}$$

and hence $\Pr\left(n\geq n^*|a_i,r_{-i}\right)-\Pr\left(n\geq n^*|r_{-i}\right)=\left(1-2\underline{\pi}\right)\Pr\left(n_{-i}=n^*-1|r_{-i}\right)$. Therefore, the program becomes

$$\max_{n^*\in\{0,1,\ldots,N\},\beta\in[0,1]} \frac{\delta}{1-\delta}\bar{u}\left(1-2\underline{\pi}\right)\left(\beta\Pr\left(n_{-i}=n^*-1|r_{-i}\right)+\left(1-\beta\right)\Pr\left(n_{-i}=n^*|r_{-i}\right)\right) \quad (49)$$

$$\text{s.t.}\quad \beta\Pr\left(n=n^*|r\right)+\Pr\left(n\geq n^*+1|r\right)\leq\frac{1-\delta}{\delta}, \quad (50)$$

where

$$\Pr\left(n_{-i}=n^*|r_{-i}\right)=\binom{N-1}{n^*}\underline{\pi}^{n^*}\left(1-\underline{\pi}\right)^{N-1-n^*}\quad\text{and}\quad\Pr\left(n=n^*|r\right)=\binom{N}{n^*}\underline{\pi}^{n^*}\left(1-\underline{\pi}\right)^{N-n^*}.$$

Fix $\rho>0$ and a sequence, indexed by k, of games with $(1-\delta)\exp\left(N^{1-\rho}\right)>k$ and pairs (n^*,β) that satisfy (50). Fix $\varepsilon>0$, and suppose toward a contradiction that, for every $\bar{k}$, there is some $k\geq\bar{k}$ such that the value of (49) exceeds ε. Taking a subsequence and relabeling $\bar{k}$ if necessary, this implies that there exists $\bar{k}$ such that, for every $k\geq\bar{k}$, the value of (49) exceeds ε.

We consider two cases and derive a contradiction in each of them.

First, suppose that there exists $c>0$ such that, for every $\tilde{k}$, there is some $k\geq\tilde{k}$ satisfying $\left|\underline{\pi}-\left(n^*-1\right)/\left(N-1\right)\right|>c$. By Hoeffding's inequality (Boucheron, Lugosi, and Massart, 2013, Theorem 2.8),

$$\Pr\left(n_{-i}\geq n^*-1|r_{-i}\right)\leq\exp\left(-2\left(\underline{\pi}-\frac{n^*-1}{N-1}\right)^2\left(N-1\right)\right).$$

Hence, for every $\tilde{k}$, there is some $k\geq\tilde{k}$ such that the value of (49) is at most

$$\frac{\delta}{1-\delta}\bar{u}\left(1-2\underline{\pi}\right)\exp\left(-2\left(\underline{\pi}-\frac{n^*-1}{N-1}\right)^2\left(N-1\right)\right)\leq\frac{\delta}{1-\delta}\bar{u}\left(1-2\underline{\pi}\right)\exp\left(-2c^2\left(N-1\right)\right).$$

Since $(1-\delta)\exp\left(N^{1-\rho}\right)\to\infty$, we have $\exp\left(-2c^2\left(N-1\right)\right)/\left(1-\delta\right)\to0$ for all $c>0$, and hence (49) is less than ε for sufficiently large k, a contradiction.

Second, suppose that for any $c>0$, there exists $\tilde{k}$ such that, for every $k\geq\tilde{k}$, we have

$$\left|\underline{\pi}-\frac{n^*-1}{N-1}\right|\leq c. \quad (51)$$

For this case, we establish a final lemma.

Lemma 12 *For any $m\in\mathbb{N}$ and any $\gamma>0$, there exists $\tilde{k}$ such that, for every $k\geq\tilde{k}$, we have*

$$\frac{\beta\Pr\left(n=n^*|r\right)+\Pr\left(n\geq n^*+1|r\right)}{\beta\Pr\left(n_{-i}=n^*-1|r_{-i}\right)+\left(1-\beta\right)\Pr\left(n_{-i}=n^*|r_{-i}\right)}\geq m\left(1-\gamma\right). \quad (52)$$

Proof. Fix $c > 0$ and take k sufficiently large that (51) holds. For any $m \in \mathbb{N}$, we have

$$\begin{aligned}
\frac{\Pr\left(n \geq n^* + 1 | r\right)}{\Pr\left(n_{-i} = n^* | r_{-i}\right)} &= \sum_{n=n^*+1}^{N} \frac{N\left(1-\underline{\pi}\right)}{N-n^*} \frac{(N-n^*)!n^*!}{(N-n)!n!} \left(\frac{\underline{\pi}}{1-\underline{\pi}}\right)^{n-n^*} \\
&\geq \sum_{n=n^*+1}^{N} \frac{N\left(1-c\right)}{N-1} \left(\frac{N-n^*}{n}\right)^{n-n^*} \left(\frac{n^*-1-c\left(N-1\right)}{N-n^*+c\left(N-1\right)}\right)^{n-n^*} \\
&\geq \sum_{n=n^*+1}^{n^*+m} \left(1-c\right) \left(\frac{N-n^*}{n^*+m} \times \frac{n^*-1-c\left(N-1\right)}{N-n^*+c\left(N-1\right)}\right)^{m} \\
&= m\left(1-c\right) \left(\frac{N-n^*}{n^*+m} \times \frac{n^*-1-c\left(N-1\right)}{N-n^*+c\left(N-1\right)}\right)^{m}.
\end{aligned}$$

By (51), for any $\gamma' > 0$, for sufficiently large k we have $\left(n^*-1\right) / \left(n^*+m\right) \geq 1-\gamma'$, and hence

$$\begin{aligned}
\frac{N-n^*}{n^*+m} \times \frac{n^*-1-c\left(N-1\right)}{N-n^*+c\left(N-1\right)} &\geq \left(1-\gamma'\right) \frac{N-n^*}{n^*-1} \times \frac{n^*-1-c\left(N-1\right)}{N-n^*+c\left(N-1\right)} \\
&= \left(1-\gamma'\right) \frac{1-c\frac{N-1}{n^*-1}}{1+c\frac{N-1}{N-n^*}} \\
&\geq \left(1-\gamma'\right) \frac{1-\frac{c}{\underline{\pi}-c}}{1+\frac{c}{1-\underline{\pi}-c}} = \frac{\left(1-\gamma'\right)\left(\underline{\pi}-2c\right)\left(1-\underline{\pi}-c\right)}{\left(\underline{\pi}-c\right)\left(1-\underline{\pi}\right)},
\end{aligned}$$

which converges to $1-\gamma'$ as $c \to 0$. Hence, for any $\gamma > 0$, there exists $\tilde{k}$ sufficiently large such that, for every $k \geq \tilde{k}$,

$$\frac{\Pr\left(n \geq n^* + 1 | r\right)}{\Pr\left(n_{-i} = n^* | r_{-i}\right)} \geq m\left(1-c\right) \left(\frac{\left(1-\gamma'\right)\left(\underline{\pi}-2c\right)\left(1-\underline{\pi}-c\right)}{\left(\underline{\pi}-c\right)\left(1-\underline{\pi}\right)}\right)^{m} \geq m\left(1-\gamma\right).$$

We therefore have

$$\frac{\beta \Pr\left(n = n^* | r\right) + \Pr\left(n \geq n^* + 1 | r\right)}{\Pr\left(n_{-i} = n^* | r_{-i}\right)} \geq \frac{\Pr\left(n \geq n^* + 1 | r\right)}{\Pr\left(n_{-i} = n^* | r_{-i}\right)} \geq m\left(1-\gamma\right).$$

Similarly, for any m and $\gamma > 0$, there exists $\tilde{k}$ such that, for every $k \geq \tilde{k}$, we have

$$\frac{\beta \Pr\left(n = n^* | r\right) + \Pr\left(n \geq n^* + 1 | r\right)}{\Pr\left(n_{-i} = n^* - 1 | r_{-i}\right)} \geq m\left(1-\gamma\right).$$

Together, these inequalities imply that, for any m and $\gamma > 0$, there exists $\tilde{k}$ such that, for every $k \geq \tilde{k}$, (52) holds. ■

Thus, for any $m \in \mathbb{N}$ and any $\gamma > 0$, there exists $\tilde{k}$ such that, for every $k \geq \tilde{k}$, the value of (49) satisfies

$$
\begin{aligned}
& \frac{\delta}{1-\delta}\bar{u}\left(1-2\underline{\pi}\right)\left(\beta\Pr\left(n_{-i}=n^{*}-1|r_{-i}\right)+\left(1-\beta\right)\Pr\left(n_{-i}=n^{*}|r_{-i}\right)\right) \\
\leq \;& \bar{u}\left(1-2\underline{\pi}\right)\frac{\beta\Pr\left(n_{-i}=n^{*}-1|r_{-i}\right)+\left(1-\beta\right)\Pr\left(n_{-i}=n^{*}|r_{-i}\right)}{\beta\Pr\left(n=n^{*}|r\right)+\Pr\left(n\geq n^{*}+1|r\right)} \quad \text{(by (50))} \\
\leq \;& \frac{\bar{u}\left(1-2\underline{\pi}\right)}{m\left(1-\gamma\right)} \quad \text{(by (52)).}
\end{aligned}
$$

Taking m and γ such that $\bar{u}\left(1-2\underline{\pi}\right)/\left(m\left(1-\gamma\right)\right)<\varepsilon$ gives the desired contradiction.