Repeated Games with Many Players*

Takuo Sugaya
Stanford GSB

Alexander Wolitzky
MIT

August 15, 2022

Abstract

Motivated by the problem of sustaining cooperation in large communities with limited information, we analyze the relationship between the population size, the discount factor, and the monitoring structure 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, all repeated-game Nash equilibrium payoffs are consistent with approximately myopic play. We provide a near-converse to this result under public, product-structure monitoring. For example, if the above ratio is small, a folk theorem holds under *random auditing*, where each player is monitored with the same probability in every period. If instead attention is restricted to *linear equilibria* (a generalization of strongly symmetric equilibria), 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 auditing, linear equilibrium

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.

1 Introduction

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 high-quality 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. Large cartels seem 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 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 identifiability 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 individual sanctions (e.g., fines) or collective ones (e.g., price wars).

In this paper, we extend 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 quite general manner: we assume only a uniform upper bound on the magnitude of the players' stage-game payoffs and a uniform lower bound on *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. 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 preview our main ideas and results. We model individual-level noise by assuming that each player's action a_i stochastically determines an *individual-level outcome* x_i, independently across players, and that the distribution of signals $y = (y_i)$ (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 (1996; 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 the current context is defined as the maximum expected reduction in uncertainty (entropy) about the outcome profile x that results from observing the signal profile y. (The elements of y can be distinct, as monitoring need not be public.) Channel capacity obeys the elementary inequality $C \leq \log |Y|$, where Y is the set of possible signal realizations. Due to this inequality, our results based on channel capacity are more general than they would be if we simply measured informativeness by the cardinality of the set of signal realizations. At the same time, channel capacity is convenient to work with, as it lets us use results 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 $(1 - \delta) N/C$—that is, the ratio of the discount rate $1 - \delta$ to the per-capita channel capacity C/N—is large, then all repeated-game Nash equilibrium payoffs are consistent with almost-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, 2022a; 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 key information measure in SW (the χ^2-*divergence* between the signal distribution at equilibrium and that following a deviation) and channel capacity.

Our second result (Theorem 2) provides a near-converse to the first, under some additional structure on monitoring. We show that for games with *public, product-structure monitoring* that satisfies a condition that we call η-*individual identifiability*, for $\eta > 0$ such that $(1-\delta)\log(N)/\eta$ is small, a large set of payoffs arise as perfect equilibria in the repeated game.[1] A simple example of public, product-structure monitoring that satisfies η-individual identifiability is η-*random auditing*, where in each period ηN players are chosen at random and and perfectly monitored, while nothing is revealed about the other players' actions. Since the channel capacity of η-random auditing is a constant multiple of ηN, this example describes a class of monitoring structures where a folk theorem holds whenever $(1-\delta) N \log(N)/C$ is small. This shows that the condition of our first theorem is tight up to $\log(N)$ slack.

Our final result (Theorem 3) considers the implications of restricting society to "collective" sanctions and rewards. We formalize this restriction by focusing on *linear 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 $(1-\delta)\exp(N^{1-\rho})$ is large, then all equilibrium payoffs are consistent with almost-myopic play. Since this condition holds even if $N \to \infty$ much slower than $\delta \to 1$, we interpret this result as an 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. Our main results on the three-way relationship among

[1] The definitions of public, product-structure monitoring and individual identifiability follow FLM. The η-identifiability condition is a version of FLM's individual full rank with η "slack."

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

N, δ, and monitoring also do not have close antecedents in the literature.

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. The current paper introduces two features that are specific to large-population games: individual-level noise (the key feature underlying Theorem 1), and the possibility that N varies together with discounting and monitoring (the key feature underlying Theorem 2).

The most relevant folk theorems are due to FLM, Kandori and Matsushima (1998, henceforth KM), and SW. The proof approach in these papers does not easily extend to the case where N and δ vary together. Our proof of Theorem 2 thus take 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 becomes approximately myopic as $N \to \infty$ for fixed δ.[3] These conditions can be adapted to the case 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 the literature on 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 related results here are Sannikov and Skrzypacz's (2007) theorem on the impossibility of collusion in duopoly with frequent actions and Brownian noise, as well as a similar result by Fudenberg and Levine (2007). These results relate to our anti-folk theorem for linear equilibrium, as we explain in

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

[4] Farther afield, there is also work suggesting that repeated-game cooperation is harder to sustain in larger groups based on evolutionary models (Boyd and Richerson, 1988) and simulations (Bowles and Gintis, 2011; Chapter 5).

Section 5.[5]

Finally, in earlier work (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" 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, so the question is whether information diffuses quickly enough to ensure that it pays to be good. 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 friction is that societal information is *insufficient* or *disaggregated* distinguishes "large-population repeated game" models, 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.

Stage Games. A stage game $G = (I, A, u)$ consists of a finite set of players $I = \{1, \dots, N\}$, 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. 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 *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. One natural interpretation of the individual outcomes is that there is some independent noise in the execution of the players' actions, so that a_i is player i's intended

[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, 2022b).

action and x_i is her realized action—in this case, $X = A$. Also, since we will assume that players do not observe their own payoffs in addition to 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.

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

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

$$\mathbf{I}(\xi) = \sum_{x \in X, y \in \bar{Y}} \xi(x) q(y|x) \log \frac{q(y|x)}{\sum_{x' \in X} \xi(x') q(y|x')}.$$[6]

Mutual information measures the expected reduction in uncertainty (entropy) about x that results from observing y. This is an endogenous object, as it depends on the prior distribution ξ of x. The *channel capacity* of (Y, q) is defined as

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

This is an exogenous measure of the informativeness of y about x. 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 use channel capacity only as an upper bound on mutual information. In turn, mutual information arises in our analysis because it obeys useful properties, in particular the

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

chain rule (CT, Theorem 2.5.2) and *Pinsker's inequality* (CT, Lemma 11.6.1): see Lemma 3 in Section 3.3.[7]

Our folk theorem (Theorem 2) 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 simply denote the public signal by y. 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.[8] 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.

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 drawn according to $q(y|x)$, and (v) each player i observes y^i.[9] A *history* h_i^t for player i at the beginning of period t thus takes the form $h_i^t = \left((z_{t'}, a_{i,t'}, y^i_{t'})_{t'=1}^{t-1}, z_t \right)$, while a *strategy* σ_i for player i maps histories h_i^t to distributions over actions $a_{i,t}$. A *repeated game outcome* $\mu \in (A \times X \times Y)^\infty$ (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.

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

[7] Entropy methods have previously been used to study various issues in repeated games, 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, these papers are not very related to ours.

[8] 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).

[9] The analysis and results are unchanged if player i also observes her own individual outcome x_i.

$\min\{|A_i|, |X_i|\} \leq 1/\underline{\pi}$ for all i. We also say that Γ is $\bar{u}$-*bounded* if the range of stage-game payoffs is bounded above by $\bar{u}$ (but noise is not necessarily bounded).

Our anti-folk theorem (Theorem 1) will apply not only for any Nash equilibrium in Γ, but also for any Nash equilibrium in an associated repeated game Γ^B, which we call the *blind game*. The blind game (which we introduced in SW) is a variant of Γ where (i) the players have access to a neutral mediator, (ii) at the beginning of each period, the mediator privately recommends an action $r_i \in A_i$ to each player i, and (iii) at the end of each period, the mediator observes the signal y (drawn according to $p\left((y^i)_i \mid (a_i)_i\right)$), while the players observe nothing. Players remember their own past actions, and the mediator does not observe the players' actions. Thus, a history for player i in the repeated game Γ^B takes the form $h_i^t = \left((r_{i,t'}, a_{i,t'})_{t'=1}^{t-1}, r_{i,t}\right)$, while a history for the mediator takes the form $h_0^t = \left((r_{i,t'})_i, (y_{t'}^i)_i\right)_{t'=1}^{t-1}$. By standard arguments, any repeated game outcome μ that is induced by a Nash equilibrium in Γ is also induced by a Nash equilibrium in Γ^B.

We can also defined the *mediated game* Γ^M, which differs from Γ^B in that each player i continues to observe y_i at the end of a period. Thus, Γ^M is obtained from Γ by introducing a mediator who observes y, and we can then obtain Γ^B from Γ^M by removing the signal components y_i from the players' information sets. The set of Nash equilibrium outcomes in Γ^M is larger than that in Γ, but smaller than that in Γ^B.

Manipulations. 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 recommended action 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

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

and player i's *maximum gain* at $\alpha \in \Delta(A)$ is

$$g_i(\alpha) = \max_{s_i : A_i \to \Delta A_i} g_i(s_i, \alpha).$$

Sets of 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 $\sum_n \beta_n = 1$.[10] For each $v \in \mathbb{R}^N$ and $\varepsilon > 0$, let $B_v(\varepsilon) = \prod_i [v_i - \varepsilon, v_i + \varepsilon]$, and let $B(\varepsilon) = \{v \in \mathbb{R}^N : B_v(\varepsilon) \subseteq F^*\}$. 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, for any $\varepsilon > 0$, the set of *action distributions consistent with ε-myopic play* is

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

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

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

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 $g_i(\alpha)$ is convex as the maximum of affine functions, and hence $A(\varepsilon)$ and $M(\varepsilon)$ are convex sets.

Our anti-folk theorem will provide conditions under which all repeated-game equilibrium payoff vectors are contained in the set $M(\varepsilon)$, while our folk theorem will provide conditions under which all payoff vectors in the set $B(\varepsilon)$ arise in repeated-game equilibria. These results are interesting insofar as $M(\varepsilon)$ is "small" and $B(\varepsilon)$ is "large." As a check that $B(\varepsilon)$ is reasonably large, in Appendix A.1 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 $M(\varepsilon)$ in Section 6, following our main results.

3 Necessary Conditions for Cooperation

3.1 Anti-Folk Theorem

Our first result is that whenever per-capita channel capacity is much smaller than the discount rate, payoffs are consistent with almost-myopic play.

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

Theorem 1 *Fix any $\bar{u} > 0$ and $\underline{\pi} > 0$. For any $\varepsilon > 0$, there exists $k > 0$ such that, for any $(\bar{u}, \underline{\pi})$-bounded repeated game Γ such that $(1-\delta) N/C > k$, any Nash equilibrium payoff vector v in Γ (or, moreover, in Γ^B) is consistent with ε-myopic play.*

Theorem 1 implies that cooperation in large groups requires a large amount of information or a high degree of patience.

When N is large, the implied necessary condition for cooperation—that $(1-\delta) N/C$ is not too large—is restrictive for some classes of repeated games but not others. First, 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 restrictive. This negative conclusion applies for traditional applications of repeated games with public monitoring where the signal space is fixed independently 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. 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 $m(i,t) \in I \backslash \{i\}$ 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 scale linearly with N in public-monitoring games where public information 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; and the market price depends on z and noise,* *where the "amount of noise" is independent of N.) In this case,*

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

which is bounded by

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

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

The proof of Theorem 1 proceeds in three steps, each of which is fairly straightforward once we introduce the right definitions and apply the right 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 in Γ^B is attained by an action distribution where the average deviation gain is bounded by the ratio of the average detectability and the discount rate (Lemma 1). This lemma is an immediate extension of the main result in 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, again under individual-level noise, the average across players i of the mutual information between x_i and y is bounded by 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 more terminology.

First, we define the *detectability* of manipulation s_i at action profile a as

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

and define the *maximum detectability* of manipulation s_i as

$$\chi_i^2(s_i) = \max_a \chi_i^2(s_i, a).$$

Our detectability measure is the χ^2-*divergence* between the probability distributions $p(\cdot|a)$ and $p(\cdot|s_i(a_i), a_{-i})$. The χ^2-divergence is a standard measure of statistical distance. Note that it is well-defined because p has full support on $\bar{Y}$.[11]

Second, denote the variance of player i's payoff under an action profile distribution $\alpha \in \Delta(A)$ by $V_i(\alpha) = \mathrm{Var}_{a\sim\alpha}(u_i(a))$. For any subset of players $J \subseteq I$, action profile distribution $\alpha \in \Delta(A)$, and profile of manipulations $s_J = (s_i)_{i\in J}$ 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

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

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

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

Note that the payoff vector under repeated game outcome μ equals

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

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

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

[11] The detectability measure is the same as in SW, but there we did not assume individual-level noise.

Lemma 1 *For any Nash equilibrium outcome μ in Γ^B, any subset of players J, and any profile of manipulations s_J, we have*

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

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

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

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

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

■

The logic behind the singleton case of Lemma 1 is discussed at length in SW. Briefly, the bound comes from decomposing the variance of a player's continuation payoff: if manipulating is unprofitable, then the signal must vary significantly with the action and the continuation payoff must vary significantly with the signal, and this payoff variation must be delivered relatively quickly due to discounting.

3.3 Bounding Detectability by Channel Capacity

We now show that, with individual-level noise, average detectability under (Y,p) can be bounded by the per-capita channel capacity of (Y,q).

We start with an intermediate lemma.

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

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

where $\mathbf{I}(\cdot;\cdot|\cdot)$ *denotes condition mutual information.*

Lemma 2 follows from standard inequalities, including Cauchy-Schwarz and Pinsker. Pinsker's inequality is where mutual information enters the analysis.

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

Lemma 3 *For any subset of players* $J \subseteq I$ *and any profile of manipulations* s_J*, we have*

$$\chi_J^2(s_J) \leq \frac{4C}{\underline{\pi}^2 |J|}. \tag{5}$$

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

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

For any fixed $\bar{u}, \underline{\pi} > 0$, taking $(1-\delta) N/C$ 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 monitoring structure, 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 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, Lemma 3 shows that with individual noise, only "informative" signals can support strong incentives.

4 Sufficient Conditions for Cooperation

4.1 Folk Theorem

Our second result is a folk theorem for repeated games with public, product-structure monitoring, which implies that the relationship among N, δ, and C in Theorem 1 is tight for these games (up to $\log N$ slack).

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.[12] Denote the set of PPE payoff vectors by $E \subseteq \mathbb{R}^N$.

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

$$\sum_{y_i : p_i(y_i|a_i) \geq \eta} p_i(y_i|a_i) \left(\frac{p_i(y_i|\alpha_i) - p_i(y_i|a_i)}{p_i(y_i|a_i)} \right)^2 \geq \eta \quad \text{for all } i \in I, a_i \in A_i, \alpha_i \in \Delta(A_i \backslash \{a_i\}). \tag{6}$$

This condition is a variant of FLM's individual full rank condition and KM's 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, from the perspective of an observer who ignores signals that occur with probability less than η under a_i. Intuitively, this requires that deviations from a_i are detectable, and that in addition "detectability" does not come entirely from a few very rare signal realizations. This assumption will ensure that player i can be incentivized through rewards whose variance and maximum absolute value are both of order $(1-\delta)/\eta$.[13]

For example, suppose that the monitoring structure is given by *η-random auditing*, where in every period ηN players (i.e., fraction η of the population) are selected uniformly at random, and the public signal perfectly reveals their identities and their realized ("noisy")

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

[13] If (6) were weakened by taking the sum over all y_i (rather than only y_i such that $q_i(y_i|a_i) \geq \eta$), player i could be incentivized by rewards with variance $O((1-\delta)/\eta)$, but not necessarily with maximum absolute value $O((1-\delta)/\eta)$. Our proof of the folk theorem requires controlling both the variance and absolute value of players' rewards, so we need the stronger condition.

actions: that is, $X_i = A_i$, $Y_i = X_i \cup \{\emptyset\}$, and

$$q_i(y_i|x_i) = \begin{cases} \eta & \text{if } y_i = x_i, \\ 0 & \text{if } y_i \in X_i \backslash \{x_i\}, \\ 1-\eta & \text{if } y_i = \emptyset, \end{cases} \quad \text{so that } p_i(y_i|a_i) = \begin{cases} \eta \pi_{a_i,x_i} & \text{if } y_i \in X_i, \\ 1-\eta & \text{if } y_i = \emptyset. \end{cases}$$

For simplicity, suppose also that $\pi_{a_i,a_i'} = \underline{\pi} < 1/(|A_i|+1)$ for all $i, a_i \neq a_i'$. Since $p_i(y_i|x_i) \geq \eta\underline{\pi}$ for all $x_i, y_i \neq \emptyset$, we then have

$$\begin{aligned} \sum_{y_i : p_i(y_i|a_i) \geq \eta\underline{\pi}} p_i(y_i|a_i) \left(\frac{p_i(y_i|\alpha_i) - p_i(y_i|a_i)}{p_i(y_i|a_i)} \right)^2 &\geq \frac{\left(p_i(a_i|a_i) - \max_{a_i' \neq a_i} p_i(a_i|a_i')\right)^2}{p_i(a_i|a_i)} \\ &= \frac{\left(\eta\left(1 - \left(|A_i|-1\right)\underline{\pi}\right) - \eta\underline{\pi}\right)^2}{\eta\left(1 - \left(|A_i|-1\right)\underline{\pi}\right)} \geq \eta\underline{\pi}^2, \end{aligned}$$

so random auditing satisfies $\eta\underline{\pi}^2$-individual identifiability.

Theorem 2 *Fix any $\bar{u} > 0$. For any $\varepsilon > 0$, there exists $k > 0$ such that, for any $\bar{u}$-bounded repeated game Γ with public, product-structure monitoring satisfying η-individual identifiability and $(1-\delta)\log(N)/\eta < k$, we have $B(\varepsilon) \subseteq E$.*

Theorem 2 implies that the relationship among N, δ, and C in Theorem 1 is tight up to $\log N$ slack. To see this, consider η-random auditing, which satisfies $\eta\underline{\pi}^2$-individual identifiability. Since $|X_i| \leq 1/\underline{\pi}$ for each i, η-random auditing has a channel capacity of at most $\eta N \log(1/\underline{\pi})$. Therefore, under η-random auditing Theorem 1 implies that payoffs in any Nash equilibrium are consistent with almost-myopic play if $(1-\delta)/\eta \to \infty$, while Theorem 2 implies that a folk theorem holds in PPE if $(1-\delta)\log(N)/\eta \to 0$.

For any repeated game Γ (possibly with private monitoring), one can define the corresponding *public game* Γ^P, where all players observe the entire signal vector y at the end of each period. Note that Γ^P is a repeated game with public monitoring, and that every PPE payoff in Γ^P is a Nash equilibrium payoff in Γ^M (because the mediator in Γ^M can replicate any public strategy profile in Γ^P, and the players in Γ^M have less information, and hence fewer possible deviations, than in Γ^P). Thus, Theorem 2 implies a folk theorem for mediated repeated games with arbitrary monitoring, in the sense that, for any repeated game Γ, if the corresponding public game Γ^P satisfies the conditions of Theorem 2, then every payoff vector in $B(\varepsilon)$ is an equilibrium payoff in the corresponding mediated game Γ^M.

Theorem 2 is a "Nash threat" folk theorem, as F^* is the set of payoffs that Pareto-dominate a static Nash equilibrium. 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).[14]

4.2 Remarks on the Proof

Theorem 2 is a folk theorem for PPE in repeated games with public monitoring. The standard proof approach, following FLM and KM, relies on transferring continuation payoffs among the players along hyperplanes that are tangent to the boundary of the PPE payoff set. Unfortunately, this approach encounters difficulties when N and δ vary together. 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, which makes self-generation difficult to satisfy. Mathematically, FLM's proof relies 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 together.

To see the problem in more detail, note that under η-individual identifiability, the per-period movement in each player's continuation payoff required to provide incentives is of order $(1-\delta)/\eta$, so the movement of the continuation payoff vector in $\mathbb{R}^N$ is $O\left(\sqrt{N}(1-\delta)/\eta\right)$. Fix a ball B contained in F^*, and consider the problem of generating the point $v = \operatorname{argmax}_{w \in B} w_1$—the point in B that maximizes player 1's payoff—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 greatest movement along a tangent hyperplane is $O\left(\sqrt{1-\delta}\right)$. FLM's proof approach thus requires that $\sqrt{N}(1-\delta)/\eta \ll \sqrt{1-\delta}$, or equivalently $(1-\delta)N/\eta^2 \ll 1$, while we assume only $(1-\delta)\log(N)/\eta \ll 1$. Therefore, while the conditions for Theorem 2 are tight up to $\log N$ slack, FLM's approach would instead require slack N.[15]

[14] Specifically, consider the $|Y_i| \times |A_i|$ matrix Q_i whose (y_i, a_i) entry is $q_i(y_i|a_i)$. We assume this matrix has full column rank, and hence there exists a $|A_i| \times |Y_i|$ matrix P_i^{-1} such that $P_i^{-1}P_i$ is the identity matrix. A sufficient condition for the minmax threat folk theorem is that the absolute value of each entry of P_i^{-1} is no more than $1/\eta$ (and $(1-\delta)\log(N)/\eta \to 0$).

[15] In SW, we extend FLM's proof to characterize the tradeoff between discounting and monitoring for an arbitrary fixed stage game. In that paper, FLM's proof approach works, because N is held fixed as discounting and monitoring vary.

Our proof of Theorem 2 (in Appendix A.4) 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 of periods, 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 throughout the block based on the public signals of their actions, and the distribution of target payoffs in the next block is set so as to deliver the promised adjustments. By η-individual identifiability, the required adjustment to each player's continuation payoff in every period is $O(1/\eta)$. By the law of large numbers, when $T \gg 1/\eta$, with high probability the total adjustment that a given player accrues over a T-period block is of order less than T, 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 the target payoff for the next block cannot be modified any further. 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. Thus, if any player's payoff adjustment is "abnormal," all players' payoffs in the block may be far from the target equilibrium payoffs.

To prove the theorem, we must ensure that rare payoff-adjustment abnormalities do not compromise either ex ante efficiency or the players' incentives. Efficiency is preserved if the blocks are long enough that the probability that *any* player's payoff adjustment is abnormal is small. Since the per-period payoff adjustment for each player is $O(1/\eta)$ 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 $\exp(-O(\eta/(1-\delta)))$. Hence, by union bound, the probability that any player's adjustment is abnormal is at most $N \exp(-O(\eta/(1-\delta)))$, which converges to 0 when $(1-\delta)\log(N)/\eta \to 0$. This step in the proof accounts for the $\log(N)$ slack.

Finally, since all players' payoffs are affected when any player's payoff adjustment becomes abnormal, incentives would be threatened if one player's action affected the probability that another player's adjustment becomes abnormal. We avoid this problem by letting each player's adjustment depend only on the signals of her own actions. Such a separation of payoff adjustments across players is possible because we assume product structure monitoring. We do not know if Theorem 2 can be extended to non-product structure monitoring without introducing polynomial slack in N.[16]

5 Comparison with Linear Equilibria

We say that a Nash equilibrium 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 on-path complete histories of play 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.[17] Relabeling the players if necessary, we can take $|b_i| \leq 1$ for all i without loss. This notion of linear equilibrium generalizes that of a *linear public perfect equilibrium* in a game with public monitoring, where continuation payoff vectors at all public histories lie on a line, as well as that of a *strongly symmetric equilibrium* (SSE) in a symmetric game with public monitoring, where the line in question is additionally required to be the 45° line.

Linear equilibria are of interest because they model collective incentive provision. If $b_i \geq 0$ for all i, all players have the same preferences over histories, and are thus all rewarded or punished together. If $b_i < 0$ for some i, the players can be divided into two groups, where each group is rewarded when the other is punished.

Our final result is that cooperation in a linear equilibrium is possible only under extremely restrictive conditions on N and δ. We view this result as essentially an impossibility theorem for large-group cooperation under collective incentives.

Theorem 3 *Fix any $\bar{u} > 0$ and $\underline{\pi} > 0$. For any $\varepsilon > 0$ and $\rho > 0$, there exists $k > 0$ such that, for any $(\bar{u}, \underline{\pi})$-bounded repeated game Γ such that $(1-\delta)\exp(N^{1-\rho}) > k$, any linear equilibrium payoff vector v in Γ (or, moreover, in Γ^B) is consistent with ε-myopic play.*

Theorem 3 differs from Theorem 1 not only in the required relationship between N and δ, but also in that the conclusion of Theorem 3 applies no matter how informative the

[16] As noted above, we conjecture that the approach of FLM and KM—which requires only a pairwise identfiability condition, not a product structure—yields a folk theorem if $(1-\delta)N/\eta^2 \to 0$.

[17] A complete history of play h at the beginning of period t takes the form $\left((z_{t'}, y_{t'}, a_{t'})_{t'=1}^{t-1}, z_t\right)$.

outcome monitoring structure is. Intuitively, this is because optimal linear equilibria have 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 played is as effective as any more informative signal.

Theorem 3 is proved in Appendix A.5. 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 we wish to enforce a symmetric pure action profile $\vec{a}_0 = (a_0, \ldots, a_0)$, where $g_i(\vec{a}_0) = \eta$, and suppose for simplicity that $|A_i| = 2$, $X = A$, and $\pi_{a_i,x_i} = 1 - \underline{\pi}$ whenever $a_i = x_i$, and $\pi_{a_i,x_i} = \underline{\pi}$ otherwise. 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, the distribution of n is approximately normal when N is large, 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

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

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

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

The left-hand size of this inequality is the statistical *score* of a normal tail test, which is approximately equal to $|z^*|$. 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 η.[18]

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

[18] 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*.

tail test must increase exponentially with the variance of the noise.[19] 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 $1/(1-\delta)$. (The 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.) As a tail test is also optimal in their setting, the reasoning just given implies that incentives can be provided only if $1/(1-\delta)$ increases exponentially with the variance of the noise. Since in their model $1/(1-\delta)$ 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 $1/(1-\delta)^{\rho}$ 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 the variance asymptotically dominates $(\log 1/(1-\delta))^{1/(1-\rho)}$ for some $\rho > 0$, while their folk theorem extends whenever the variance is asymptotically dominated by $(\log 1/(1-\delta))^{1/(1+\rho)}$ for some $\rho > 0$.

6 Discussion

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

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

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

Payoff vectors in $M(\varepsilon)$ are consistent with ε-myopic play, in the sense that they 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 at an action distribution $\alpha \in A(\varepsilon)$. A more standard notion of "ε-myopia" is that *all* players' deviations gains are less than ε. The set of payoff vectors consistent with this notion is the set of *static ε-correlated equilibrium payoffs*,

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

given by

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

We can illustrate the set $M(\varepsilon)$ by comparing it with $CE(\varepsilon)$. We give a simple example where $M(\varepsilon)$ and $CE(\varepsilon)$ are very different (and $M(\varepsilon)$ cannot be replaced by $CE(\varepsilon)$ in Theorem 1), and then give a general bound on the distance between $M(\varepsilon)$ and $CE(c\varepsilon)$ for any sufficiently large constant c. Intuitively, $M(\varepsilon)$ and $CE(\varepsilon)$ can be very different if moral hazard binds for only a few players, and these players' actions have large effects on others' payoffs; while $M(\varepsilon)$ and $CE(c\varepsilon)$ are similar if each player's action has a small effect on every opponent's payoff.

For an example where $M(\varepsilon)$ 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

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

and 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 $(1, 1, \ldots, 1) \in M(\varepsilon)$. 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, and 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 action (i.e., $Y = \{H, L\}$, $q(y|x) = \mathbf{1}\{y = x_1\}$) has channel capacity $\log 2$ and supports the payoff vector

$$\left(\frac{1-3\underline{\pi}}{1-2\underline{\pi}}, \frac{1-3\underline{\pi}}{1-2\underline{\pi}}, \ldots, \frac{1-3\underline{\pi}}{1-2\underline{\pi}}\right), \quad \text{for all } \delta > \frac{1}{2-3\underline{\pi}} \text{ and all } N \geq 2.$$[20]

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

We can also bound the distance between $M(\varepsilon)$ and $CE(c\varepsilon)$, for any sufficiently large constant c.

Proposition 1 *Fix $b > 0$ such that $\left|u_i\left(a_j', a_{-j}\right) - u_i(a)\right| \leq b/N$ for all $i \neq j, a_j', a$. Then, for any $v \in M(\varepsilon)$ and any $c \geq \sqrt{4b/\varepsilon}$, there exists $v' \in CE(c\varepsilon)$ such that*

$$\frac{1}{N}\left|\sum_{i \in I} v_i - \sum_{i \in I} v_i'\right| \leq \frac{\bar{u}}{c}.$$

For example, in a repeated random matching game, b is the impact of a player's action on her partner's payoff, which is independent of N. Note that for the smallest value of c permitted by Proposition 1, we have $c\varepsilon \approx \sqrt{4b\varepsilon}$, so when ε is small every $v' \in CE(c\varepsilon)$ is an approximate static correlated equilibrium, albeit with an approximating factor of order $\sqrt{\varepsilon}$ rather than ε.

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 monitoring structure as a key statistic. For a specific class of monitoring structures (*random auditing*), 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 formalizes a sense in which large-group cooperation must rely on personalized sanctions.

[20] 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 $\phi = (1-\delta)/(\delta - 3\delta\pi)$, and substituting this into the value recursion $v = (1-\delta)(1) + \delta(1-\pi\phi)v$ yields $v = (1-3\underline{\pi})/(1-2\underline{\pi})$.

We hope our results raise questions for future theoretical and applied research. On the theory side, we believe that connections between repeated games, mechanism design, and information theory remain under-explored. As for applied work, it would be interesting to see more systematic empirical or experimental evidence on the determinants of large-group cooperation, for example on the relative efficacy of individual and collective sanctions in populations of different sizes.

A Appendix

A.1 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.

To see this, 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 $v - \varepsilon$. Fix such an n, and let $\psi = n/N$.

Consider the mixed action profile α^1 where favored players *Contribute* with probability $(v+\varepsilon)/(1-c) \in (0,1)$ and disfavored players always *Contribute*. At this profile, favored players receive payoff

$$f(\psi) := \psi \frac{v+\varepsilon}{1-c} + (1-\psi)(1) - c\frac{v+\varepsilon}{1-c},$$

while disfavored players receive payoff

$$g(\psi) := \psi \frac{v+\varepsilon}{1-c} + (1-\psi)(1) - c.$$

Note that $f'(\psi) < 0$, so $f(\psi) \geq f(1) = v + \varepsilon$.

With $f(\psi)$ so defined, consider the mixed action profile α^2 where favored players *Contribute* with probability $(v+\varepsilon)^2/((1-c) f(\psi)) \in (0,1)$ and disfavored players *Contribute* with probability $(v+\varepsilon)/f(\psi) \in (0,1)$. Note that each player's payoff at profile α^2 equals her payoff at profile α^1 multiplied by $(v+\varepsilon)/f(\psi)$. Therefore, at profile α^2, favored players receive payoff

$$f(\psi)\frac{v+\varepsilon}{f(\psi)} = v + \varepsilon,$$

while disfavored players receive payoff

$$\begin{aligned} g(\psi)\frac{v+\varepsilon}{f(\psi)} &= \left(f(\psi)-\left(1-\frac{v+\varepsilon}{1-c}\right)c\right)\frac{v+\varepsilon}{f(\psi)} \\ &\leq v+\varepsilon-\left(1-\frac{v+\varepsilon}{1-c}\right)c(v+\varepsilon) \quad \text{(since } f(\psi)\leq 1\text{)} \\ &\leq v-\varepsilon, \end{aligned}$$

where the last inequality follows from $\varepsilon = cv(1-c-v)/4$ and straightforward algebra.

A.2 Proof of Lemma 2

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

$$(\Pr(y|a,x_i)-p(y|a))^2 = \left(\frac{p(y|a)}{\pi_{a_i,x_i}}\left(\Pr(x_i|a,y)-\pi_{a_i,x_i}\right)\right)^2 \leq \left(\frac{p(y|a)}{\underline{\pi}}\left(\Pr(x_i|a,y)-\pi_{a_i,x_i}\right)\right)^2. \tag{7}$$

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

$$\begin{aligned}
\chi_i^2(s_i,a) &= \sum_{y\in\bar{Y}} \frac{1}{p(y|a)} \left(p(y|s_i(a_i),a_{-i}) - p(y|a)\right)^2 \\
&= \sum_y \frac{1}{p(y|a)} \left(\sum_{x_i} \left(\pi_{s_i(a_i),x_i} - \pi_{a_i,x_i}\right) \Pr(y|a,x_i) \right)^2 \\
&= \sum_y \frac{1}{p(y|a)} \left(\sum_{x_i} \left(\pi_{s_i(a_i),x_i} - \pi_{a_i,x_i}\right) \left(\Pr(y|a,x_i) - p(y|a)\right) \right)^2 \\
&\leq \sum_{x_i} \left(\pi_{s_i(a_i),x_i} - \pi_{a_i,x_i}\right)^2 \sum_y \frac{1}{p(y|a)} \sum_{x_i} \left(\Pr(y|a,x_i) - p(y|a)\right)^2 \\
&\leq \frac{2}{\underline{\pi}^2} \sum_y p(y|a) \sum_{x_i} \left(\Pr(x_i|a,y) - \pi_{a_i,x_i}\right)^2 \\
&\leq \frac{2}{\underline{\pi}^2} \sum_y p(y|a) \left(\sum_{x_i} \left|\Pr(x_i|a,y) - \pi_{a_i,x_i}\right| \right)^2 \\
&\leq \frac{4}{\underline{\pi}^2} \sum_y p(y|a) \sum_{x_i} \Pr(x_i|a,y) \log \frac{\Pr(x_i|a,y)}{\pi_{a_i,x_i}} \\
&= \frac{4}{\underline{\pi}^2} \sum_{x_i,y} \Pr(x_i,y|a) \log \frac{\Pr(x_i|a,y)}{\pi_{a_i,x_i}} = \frac{4\mathbf{I}(x_i;y|a)}{\underline{\pi}^2}.
\end{aligned}$$

where the first inequality follows by Cauchy-Schwarz, the second follows by (7) and $\sum_{x_i} \left(\pi_{a_i,x_i} - \pi_{a_i,x_i}\right)^2 \leq 2$, the third is immediate, and the fourth follows by Pinsker's inequality (CT, Lemma 11.6.1).

A.3 Proof of Lemma 3

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

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

where the last equality follows by the chain rule for mutual information (CT, Theorem 2.5.2), because $(x_i)_{i\in J}$ are independent conditional on a.

Next, note that

$$\begin{aligned}\mathbf{I}(x_J; y|a) &= \mathbf{I}(x; y|a) - \mathbf{I}\left(x_{I\setminus J}; y|a, x_J\right) \le \mathbf{I}(x; y|a) = \mathbf{I}(x, a; y) - \mathbf{I}(a; y) \\ &= \mathbf{I}(x; y) + \mathbf{I}(a; y|x) - \mathbf{I}(a; y) = \mathbf{I}(x; y) + 0 - \mathbf{I}(a; y) \le \mathbf{I}(x; y) \le C,\end{aligned}$$

where the first equality follows by the chain rule, the first inequality follows because mutual information is non-negative, the second and third equalities again follow by the chain rule, the fourth equality follows because a and y are independent conditional on x, the second inequality again follows by non-negativity, and the last inequality follows by the definition of channel capacity.

In total, we have

$$\chi_J^2(s_J, a) \le \frac{4C}{\underline{\pi}^2 |J|} \quad \text{for all } a \in A,$$

and (5) follows as $\chi_J^2(s_J|a) = \max_a \chi_J^2(s_J, a)$.

A.4 Proof of Theorem 2

A.4.1 Preliminaries

Fix any $\varepsilon > 0$. If $\varepsilon \ge \bar{u}/2$ 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*

$$\begin{aligned}\mathbb{E}[f_{i,r_i}(y_i)|r_i] - \mathbb{E}[f_{i,r_i}(y_i)|a_i] &\ge \bar{u} \quad \textit{for all } a_i \ne r_i, &(8)\\ \mathbb{E}[f_{i,r_i}(y_i)|r_i] &= 0, &(9)\\ \mathrm{Var}(f_{i,r_i}(y_i)|r_i) &\le \bar{u}^2/\eta, \quad \textit{and} &(10)\\ |f_{i,r_i}(y_i)| &\le 2\bar{u}/\eta \quad \textit{for all } y_i. &(11)\end{aligned}$$

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

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

Note that (6) is equivalent to

$$d\left(p_i(r_i; Y_i^*), \mathrm{co}\left(P_i(r_i; Y_i^*)\right)\right) \geq \sqrt{\eta} \quad \text{for all } i \in I, r_i \in A_i,$$

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 $(p_i(r_i; Y_i^*) - p) \cdot x \geq \sqrt{\eta}$ for all $p \in P_i(r_i; Y_i^*)$. By definition of p_i and P_i, this implies that $\sum_{y_i \in Y_i^*} (p_i(y_i|r_i) - p_i(y_i|a_i))\, x(y_i) \geq \sqrt{p_i(y_i|r_i)\,\eta}$ for all $a_i \neq r_i$. Now define

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

Clearly, conditions (8) and (9) hold. Moreover, since $\mathbb{E}[f_{i,r_i}(y_i)|r_i] = 0$ 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

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

and hence (10) holds. Finally, (11) holds since, for each $y_i \in Y_i^*$,

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

■

Now fix $i \in I$ and $r_i \in A_i$, and suppose that $y_{i,t} \sim p_i(\cdot|r_i)$ for each period $t \in \mathbb{N}$, independently across periods (as would be the case in the repeated game if r_i were taken in every period). By (10), for any $T \in \mathbb{N}$, we have

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

Together with (9) and (11), 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

$$\Pr\left(\sum_{t=1}^{T} \delta^{t-1} f_{i,r_i}(y_{i,t}) \geq \bar{f}\right) \leq \exp\left(-\frac{\bar{f}^2 \eta}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2 + \frac{2}{3}\bar{f}\bar{u}\right)}\right). \tag{12}$$

Our second lemma fixes T and $\bar{f}$ so that the bound in (12) is sufficiently small, and some other conditions used in the proof also hold.

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

$$60\bar{u}N \exp\left(-\frac{\left(\frac{\bar{f}}{3}\right)^2 \eta}{2\left(\frac{1-\delta^{2T}}{1-\delta^2}\bar{u}^2 + \frac{2}{3}\frac{\bar{f}}{3}\bar{u}\right)}\right) \leq \varepsilon, \tag{13}$$

$$8\frac{1-\delta}{1-\delta^T}\left(\bar{f} + \frac{2\bar{u}}{\eta}\right) \leq \varepsilon, \tag{14}$$

$$4\bar{u}\frac{1-\delta^T}{\delta^T} + \frac{1-\delta}{\delta^T}\left(\bar{f} + \frac{2\bar{u}}{\eta}\right) \leq \varepsilon. \tag{15}$$

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

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

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

$$\frac{4}{9}\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta}} \leq 1 \quad \text{and} \tag{16}$$

$$8\bar{u}\left(\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta}} + \frac{1-\delta}{1-\delta^T}\frac{2}{\eta}\right) \leq \varepsilon. \tag{17}$$

It now follows from straightforward algebra (provided in Appendix A.4.4) that (13)–(15) hold for every $k \geq \bar{k}$. ■

A.4.2 Equilibrium Construction

Fix any k, T, and $\bar{f}$ that satisfy (13)–(15), 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 $v \in E(\Gamma)$.[21] 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 $w : H^{T+1} \to \mathbb{R}^N$ that satisfy

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

Incentive Compatibility. $\sigma_i \in \operatorname{argmax}_{\tilde{\sigma}_i} \mathbb{E}^{\tilde{\sigma}_i, \sigma_{-i}}\left[(1-\delta)\sum_{t=1}^T \delta^{t-1} u_{i,t} + \delta^T w_i\left(h^{T+1}\right)\right]$ for all $i \in I$.

Self Generation. $w\left(h^{T+1}\right) \in B_v(\varepsilon/2)$ for all h^{T+1}. (Note that, since $B_v(\varepsilon/2)$ is cube with side-length ε and $v^* = \operatorname{argmax}_{v \in B_v(\varepsilon/2)} \zeta \cdot v$, this is equivalent to $\zeta_i\left(w_i\left(h^{T+1}\right) - v_i^*\right) \in [-\varepsilon, 0]$ for all i and h^{T+1}.)

Defining $\pi_i\left(h^{T+1}\right) = \left(\delta^T / (1-\delta)\right)\left(w_i\left(h^{T+1}\right) - v_i^*\right)$, these conditions can be rewritten as

Promise Keeping.

$$v_i^* = \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma\left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \pi_i\left(h^{T+1}\right)\right] \qquad \text{for all } i. \tag{18}$$

Incentive Compatibility.

$$\sigma_i \in \operatorname*{argmax}_{\tilde{\sigma}_i} \mathbb{E}^{\tilde{\sigma}_i, \sigma_{-i}}\left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \pi_i\left(h^{T+1}\right)\right] \qquad \text{for all } i. \tag{19}$$

[21] 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 $w\left(h^{T+1}\right)$.

Self Generation.

$$\zeta_i \frac{1-\delta}{\delta^T} \pi_i \left(h^{T+1}\right) \in [-\varepsilon, 0] \quad \text{for all } i \text{ and } h^{T+1}. \tag{20}$$

Fix $\zeta \in \{-1,1\}^N$ and $v^* = \text{argmax}_{v \in B_v(\varepsilon/2)} \zeta \cdot v$. We construct a block strategy profile σ and continuation value function π 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

$$u_i(\bar{\alpha}) = v_i^* + \zeta_i \varepsilon/2 \quad \text{for all } i, \tag{21}$$

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 $B_v(\varepsilon) \subseteq F^*$.

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, which will determine player i's prescribed equilibrium action 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 (21), and let $r_t^{NE} \in A$ be a mixed action profile which is drawn by public randomization at the start of period t from the distribution α^{NE}.[22] 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 r_t^{NE} if $\zeta_i = -1$, where $BR_i(r_{-i}) = \text{argmax}_{a_i \in A_i} u_i(a_i, r_{-i})$ 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 r_t^{NE}.

[22] Technically, the public randomization device Z_t is always a uniform $[0,1]$ random variable. Throughout the proof, whenever we say that a certain variable is "drawn by public randomization," we mean that its realization is determined by public randomization, independently of the other variables in the construction. Since we define only a finite number B of such variables, this can be done by, for example, specifying that if $n = b \bmod B$ then the n^{th} digit of z is used to encode the realization of the b^{th} such variable we define.

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. The most subtle case is when there is a unique player i such that $\theta_{i,t} = 1$. Intuitively, this case will correspond to situations where the signals of player i's actions are "abnormal," which later in the proof will imply that her continuation payoffs cannot be adjusted further without violating the self-generation constraint. In this case, player i starts taking static best responses. Moreover, if $\zeta_i = -1$—so that player i's continuation payoff is already "low"—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 $r_{i,t}^{NE}$.

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 $r_{-i,t}^{NE}$, depending on $(\theta_{j,t})_{j\neq i}$.)

States: At the start of each period t, conditional on the public randomization draw of $r_t \in A$ described above, an additional 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

$$f_{i,t} = \begin{cases} f_{i,r_{i,t}}\left(y_{i,t}\right) & \text{if } S_{i,t} = 0, \\ f_{i,r_{i,t}}\left(\tilde{y}_{i,t}\right) & \text{if } S_{i,t} = NE, \\ 0 & \text{if } S_{i,t} = BR. \end{cases} \tag{22}$$

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 $\tilde{y}_t$.[23] 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 $\pi : H^{T+1} \to \mathbb{R}$.

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

$$\theta_{i,t} = \mathbf{1}\left\{\exists t' \leq t : \left|\sum_{t''=1}^{t'-1} \delta^{t''-1} f_{i,t''}\right| \geq \bar{f}\right\}. \tag{23}$$

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 ac-

[23] Intuitively, introducing the variable $\tilde{y}_t$, rather than simply using $y_{i,t}$ everywhere in (22), ensures that the distribution of $f_{i,t}$ does not depend on player i's opponents' strategies.

cording to $p_j\left(y_{j,t}|r_{j,t}\right)$, 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 $p_j\left(\tilde{y}_{j,t}|r_{j,t}\right)$, 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, $\Pr\left(\left(f_{j,t}\right)_{j\in J}|h^t,z_t,a_{i,t}\right)=\prod_{j\neq i}\Pr\left(f_{j,t}|S_{j,t},r_{j,t}\right)$, which is independent of $a_{i,t}$ as desired. ■

Continuation Value Function: We now construct the continuation value function $\pi:H^{T+1}\to\mathbb{R}^N$. For each player i and end-of-block history h^{T+1}, player i's continuation value $\pi_i\left(h^{T+1}\right)$ will be defined as the sum of T "rewards" $\pi_{i,t}$, where $t=1,\ldots,T$, and a constant term c_i that does not depend on h^{T+1}.

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

1. If $\theta_{j,t}=0$ for all $j\in I$, then

$$\pi_{i,t}=\delta^{t-1}\left(v_i^*+\zeta_i\varepsilon/4-u_i\left(\alpha_t^*\right)+f_{i,r_{i,t}}\left(y_{i,t}\right)\right). \tag{24}$$

2. If $\theta_{i,t}=1$ and $\theta_{j,t}=0$ for all $j\neq i$,

$$\pi_{i,t}=\delta^{t-1}\left(v_i^*+\zeta_i\varepsilon/4-u_i\left(\alpha_t^*\right)\right). \tag{25}$$

3. Otherwise,

$$\pi_{i,t}=\delta^{t-1}\left(-\zeta_i\bar{u}-u\left(\alpha_t^*\right)+\mathbf{1}\left\{S_{i,t}=0\right\}f_{i,r_{i,t}}\left(y_{i,t}\right)\right). \tag{26}$$

The constant c_i is defined as

$$c_i=-\mathbb{E}\left[\sum_{t=1}^{T}\delta^{t-1}\left(\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t}=0\right\}\left(v_i^*+\zeta_i\varepsilon/4\right)-\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t}=1\right\}\zeta_i\bar{u}\right)\right]+\frac{1-\delta^T}{1-\delta}v_i^*. \tag{27}$$

Note that, since $v_i^*+\zeta_i\varepsilon/4$ and v_i^* are both feasible payoffs, we have

$$|c_i|\leq 2\bar{u}\frac{1-\delta^T}{1-\delta}. \tag{28}$$

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

$$\pi_i\left(h^{T+1}\right)=c_i+\sum_{t=1}^{T}\pi_{i,t}. \tag{29}$$

A.4.3 Verification of the Equilibrium Conditions

We now verify that σ and π 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*

$$\Pr\left(\max_{i\in I, t\in\{1,\dots,T\}} \theta_{i,t} = 0\right) \geq 1 - \frac{\varepsilon}{20\bar{u}}. \tag{30}$$

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

$$\Pr\left(\max_{t\in\{1,\dots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} f_{i,t'}\right| \geq \bar{f}\right) \leq \frac{\varepsilon}{20\bar{u}N}. \tag{31}$$

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

$$\Pr\left(\max_{t\in\{1,\dots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} f_{i,t'}\right| \geq \bar{f}\right) \leq \Pr\left(\max_{t\in\{1,\dots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} \tilde{f}_{i,t'}\right| \geq \bar{f}\right). \tag{32}$$

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

$$\Pr\left(\max_{t\in\{1,\dots,T\}} \left|\sum_{t'=1}^{t} \delta^{t-1} \tilde{f}_{i,t'}\right| \geq \bar{f}\right) \leq 3 \max_{t\in\{1,\dots,T\}} \Pr\left(\left|\sum_{t'=1}^{t} \delta^{t-1} \tilde{f}_{i,t'}\right| \geq \bar{f}/3\right). \tag{33}$$

Letting $x_{i,t} = \delta^{t-1}\tilde{f}_{i,t}$, note that $|x_{i,t}| \leq 2\bar{u}/\eta$ with probability 1 by (11), $\mathbb{E}[x_{i,t}] = 0$ by (9), and

$$\mathrm{Var}\left(\sum_{t'=1}^{t} x_{i,t'}\right) = \sum_{t'=1}^{t} \mathrm{Var}(x_{i,t'}) \leq \sum_{t'=1}^{T} \mathrm{Var}(x_{i,t'}) = \frac{1-\delta^T}{1-\delta}\frac{\bar{u}^2}{\eta} \quad \text{by (10).}$$

Therefore, by Bernstein's inequality ((12), which again applies because $\left(\tilde{f}_{i,t}\right)_{t=1}^{T}$ are independent) and (13), we have, for each $t \leq T$,

$$\Pr\left(\left|\sum_{t'=1}^{t} \delta^{t'-1} \tilde{f}_{i,t'}\right| \geq \bar{f}/3\right) \leq \frac{\varepsilon}{60\bar{u}N}. \tag{34}$$

Finally, (32), (33), and (34) together imply (31). ■

Incentive Compatibility: We use the following lemma (proof in Appendix A.4.5).

Lemma 8 *For each player i and block strategy profile σ, incentive compatibility holds (i.e., (19) is satisfied) if and only if*

$$\operatorname{supp} \sigma_i\left(h^t\right) \subseteq \operatorname*{argmax}_{a_{i,t} \in A_i} \mathbb{E}^{\sigma_{-i}}\left[\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, a_{i,t}\right] \quad \textit{for all } t \textit{ and } h^t. \tag{35}$$

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

$$\mathbb{E}^{\sigma}\left[\delta^{t'-1} u_{i,t} + \pi_{i,t'} | h^t\right] = \mathbb{E}^{\sigma}\left[\delta^{t'-1}\left(\mathbf{1}\left\{\max_{j \neq i} \theta_{j,t'} = 0\right\}\left(v_i^* + \zeta_i \varepsilon/4\right) - \mathbf{1}\left\{\max_{j \neq i} \theta_{j,t'} = 1\right\} \zeta_i \bar{u}\right) | h^t\right]. \tag{36}$$

We now verify that (35) holds. Fix a player i, period t, and history h^t. We consider several cases, which parallel the definition of the reward $\pi_{i,t}$.

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 (8) and (24), and recalling that $\bar{u} \geq \max_a u_i(a) - \min_a u_i(a)$, we have

$$\begin{aligned} & \mathbb{E}^{\sigma_{-i}}\left[\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, z_t, a_{i,t} = r_{i,t}\right] - \mathbb{E}^{\sigma_{-i}}\left[\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, z_t, a_{i,t} = a_i\right] \\ = \ & \delta^{t-1}\left(\mathbb{E}\left[u_i(r_t) + f_{i,r_{i,t}}(y_{i,t}) | a_{i,t} = r_{i,t}\right] - \mathbb{E}\left[u_i(a_i, r_{-i,t}) + f_{i,r_{i,t}}(y_{i,t}) | a_{i,t} = a_i\right]\right) \\ \leq \ & 0, \quad \text{so (35) holds.} \end{aligned}$$

2. If $\theta_{i,t} = 1$ and $\theta_{j,t} = 0$ for all $j \neq i$, then the reward $\pi_{i,t}$ specified by (25) does not depend on $y_{i,t}$. Hence, (35) 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 $r_{i,t}^{NE}$.

3. Otherwise,

   (a) If $S_{i,t} = 0$, then (35) holds because it holds in Case 1 above and (24) and (26) 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, r_t^{NE} is prescribed. Since the reward $\pi_{i,t}$ specified by (26) does not depend on $y_{i,t}$, (35) 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 r_t^{NE}.

Promise Keeping: This essentially holds by construction: we have

$$
\begin{aligned}
& \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^{T} \delta^{t-1} u_{i,t} + \pi_i \left(h^{T+1}\right) \right] \\
= \; & \frac{1-\delta}{1-\delta^T} \left(\mathbb{E}^\sigma \left[\sum_{t=1}^{T} \left(\delta^{t-1} u_{i,t} + \pi_{i,t}\right) \right] + c_i \right) \quad \text{(by (29))} \\
= \; & \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} (v_i^* + \zeta_i \varepsilon/4) - \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} \zeta_i \bar{u} \right) + c_i \right] \quad \text{(by (36))} \\
= \; & v_i^* \quad \text{(by (27)), so (18) holds.}
\end{aligned}
$$

Self Generation: We use the following lemma (proof in Appendix A.4.6).

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

$$
\begin{aligned}
\zeta_i \sum_{t=1}^{T} \pi_{i,t} \;&\leq\; \bar{f} + \frac{2\bar{u}}{\eta} \quad \textit{and} & (37)\\
\left| \sum_{t=1}^{T} \pi_{i,t} \right| \;&\leq\; \bar{f} + \frac{2\bar{u}}{\eta} + 2\bar{u}\frac{1-\delta^T}{1-\delta}. & (38)
\end{aligned}
$$

In addition,

$$
\zeta_i c_i \leq -\frac{1-\delta^T}{1-\delta}\frac{\varepsilon}{8}. \tag{39}
$$

To establish self generation ((20)), it suffices to show that, for each h^{T+1}, $\zeta_i \pi_i \left(h^{T+1}\right) \le 0$ and $\left((1-\delta)/\delta^T\right) \left|\pi_i\left(h^{T+1}\right)\right| \le \varepsilon$. This now follows because

$$\begin{aligned}
\zeta_i \pi_i \left(h^{T+1}\right) &= \zeta_i \left(c_i + \sum_{t=1}^{T} \pi_{i,t}\right) \le -\frac{1-\delta^T}{1-\delta}\frac{\varepsilon}{8} + \bar{f} + 2\bar{u}/\eta \quad \text{(by (37) and (39))} \\
&\le \frac{1-\delta^T}{8(1-\delta)}\left(-\varepsilon + 8\left(\frac{1-\delta}{1-\delta^T}\right)\left(\bar{f} + 2\bar{u}/\eta\right)\right) \le 0 \quad \text{(by (14)),} \quad \text{and} \\
\frac{1-\delta}{\delta^T}\left|\pi_i\left(h^{T+1}\right)\right| &\le \frac{1-\delta}{\delta^T}\left(|c_i| + \left|\sum_{t=1}^{T}\pi_{i,t}\right|\right) \\
&\le \frac{1-\delta}{\delta^T}\left(4\bar{u}\frac{1-\delta^T}{1-\delta} + \bar{f} + 2\bar{u}/\eta\right) \quad \text{(by (28) and (38))} \\
&= \frac{1-\delta^T}{\delta^T}4\bar{u} + \frac{1-\delta}{\delta^T}\left(\bar{f} + 2\bar{u}/\eta\right) \le \varepsilon \quad \text{(by (15)),}
\end{aligned}$$

which completes the proof.

A.4.4 Omitted Details for the Proof of Lemma 5

We show that, with the stated definitions of T and $\bar{f}$, (16) and (17) imply (13)–(15). 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,

$$\begin{aligned}
\frac{2\bar{f}\left(1-\delta^2\right)}{9\bar{u}\left(1-\delta^{2T}\right)} &< \frac{4}{9\bar{u}}\frac{1-\delta}{1-\delta^T}\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta^T}{1-\delta}\frac{\bar{u}^2}{\eta}} \\
&= \frac{4}{9}\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log(N)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta}} \le 1 \quad \text{(by (16)).}
\end{aligned}$$

Therefore,

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

Moreover,

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

Hence, we have

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

This establishes (13).

Next, we have

$$8\frac{1-\delta}{1-\delta^T}\left(\bar{f}+\frac{2\bar{u}}{\eta}\right) = 8\bar{u}\left(\sqrt{36\log\left(\frac{60\bar{u}}{\varepsilon}\right)\log\left(N\right)\frac{1-\delta}{1-\delta^T}\frac{1}{\eta}}+\frac{1-\delta}{1-\delta^T}\frac{2}{\eta}\right) \leq \varepsilon \quad \text{(by (17))}. \tag{40}$$

This establishes (14).

Finally, by (40) and $8\bar{u}\left(1-\delta^T\right)/\delta^T \leq \varepsilon$, we have

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

This establishes (15).

A.4.5 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 (35) 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 $\sum_{t'=1}^{t}\pi_{i,t'}$). For each $t' \geq t+1$, there are several cases to consider.

1. If $\theta_{j,t'} = 0$ for all j, then by (9) and (24) we have

$$\begin{aligned}\mathbb{E}^\sigma\left[\delta^{t'-1}u_{i,t'}+\pi_{i,t'}|h^{t'}\right] &= \delta^{t'-1}\left(u_i\left(\alpha^*_{t'}\right)+v^*_i+\zeta_i\varepsilon/4-u_i\left(\alpha^*_{t'}\right)+\mathbb{E}\left[f_{i,r_{i,t'}}\left(y_{i,t'}\right)|r_{i,t'}\right]\right)\\ &= \delta^{t'-1}\left(v^*_i+\zeta_i\varepsilon/4\right).\end{aligned}$$

2. If $\theta_{i,t'} = 1$ and $\theta_{j,t'} = 0$ for all $j \neq i$, then by (25) we have

$$\mathbb{E}^{\sigma}\left[\delta^{t'-1}u_{i,t'} + \pi_{i,t'}|h^{t'}\right] = \delta^{t'-1}\left(u_i\left(\alpha_{t'}^*\right) + v_i^* + \zeta_i\varepsilon/4 - u_i\left(\alpha_{t'}^*\right)\right) = \delta^{t'-1}\left(v_i^* + \zeta_i\varepsilon/4\right).$$

3. Otherwise,

   (a) If $S_{i,t'} = 0$, then by (9) and (26) (and recalling that player i's equilibrium action is $r_{i,t'}$ when $S_{i,t'} = 0$) we have

$$\begin{aligned}\mathbb{E}^{\sigma}\left[\delta^{t'-1}u_{i,t'} + \pi_{i,t'}|h^{t'}\right] &= \delta^{t'-1}\left(u_i\left(\alpha_{t'}^*\right) - \zeta_i\bar{u} - u\left(\alpha_{t'}^*\right) + \mathbb{E}\left[f_{i,r_{i,t'}}\left(y_{i,t'}\right)|r_{i,t'}\right]\right) \\ &= \delta^{t'-1}\left(-\zeta_i\bar{u}\right).\end{aligned}$$

   (b) If $S_{i,t'} \neq 0$, then by (26) we have

$$\mathbb{E}^{\sigma}\left[\delta^{t'-1}u_{i,t'} + \pi_{i,t'}|h^{t'}\right] = \delta^{t'-1}\left(u_i\left(\alpha_{t'}^*\right) - \zeta_i\bar{u} - u\left(\alpha_{t'}^*\right)\right) = \delta^{t'-1}\left(-\zeta_i\bar{u}\right)$$

In total, (36) holds, and player i's net continuation payoff under σ from period $t+1$ onward equals

$$\mathbb{E}^{\sigma}\left[\sum_{t'=t+1}^{T}\delta^{t'-1}\left(\mathbf{1}\left\{\max_{j\neq i}\theta_{j,t'} = 0\right\}\left(v_i^* + \zeta_i\varepsilon/4\right) - \mathbf{1}\left\{\max_{j\neq i}\theta_{j,t'} = 1\right\}\zeta_i\bar{u}\right)|h^t\right].$$

By Lemma 6, the distribution of $\left(\left(\theta_{n,t'}\right)_{n\neq i}\right)_{t'=t+1}^{T}$ 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 $\mathbb{E}^{\sigma_{-i}}[\delta^{t-1}u_{i,t} + \pi_{i,t}|h^t, a_{i,t}]$.

A.4.6 Proof of Lemma 9

Define

$$\pi^v_{i,t} = \begin{cases} \delta^{t-1}\left(v^*_i + \zeta_i\varepsilon/4 - u_i\left(\alpha^*_t\right)\right) & \text{if } \theta_{j,t} = 0 \text{ for all } j \neq i, \\ \delta^{t-1}\left(-\zeta_i\bar{u} - u_i\left(\alpha^*_t\right)\right) & \text{otherwise} \end{cases}, \quad \text{and}$$

$$\pi^f_{i,t} = \begin{cases} \delta^{t-1} f_{i,a_{i,t}}\left(y_{i,t}\right) & \text{if either } \theta_{j,t} = 0 \text{ for all } j \text{ or } S_{i,t} = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Note that, by (24)–(26), we can write $\pi_{i,t} = \pi^v_{i,t} + \pi^f_{i,t}$. We show that, for every end-of-block history h^{T+1}, we have

$$\zeta_i \sum_{t=1}^{T} \pi^v_{i,t} \in \left[-2\bar{u}\frac{1-\delta^T}{1-\delta}, 0\right] \quad \text{and} \tag{41}$$

$$\left|\zeta_i \sum_{t=1}^{T} \pi^f_{i,t}\right| \leq \bar{f} + 2\bar{u}/\eta. \tag{42}$$

Since $\pi_{i,t} = \pi^v_{i,t} + \pi^f_{i,t}$, (41) and (42) imply (37) and (38), which proves the first part of the lemma.

For (41), 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 $u_i\left(\alpha^*_t\right) \geq \sum_a \bar{\alpha}\left(a\right) \min\left\{u_i\left(a\right), \max_{a'_i} u_i\left(a'_i, a_{-i}\right)\right\} \geq u_i\left(\bar{\alpha}\right) = v^*_i + \varepsilon/2$, by (21); and (ii) if $\zeta_i = -1$, we have $u_i\left(\alpha^*_t\right) \leq \max\left\{u_i\left(\bar{\alpha}\right), u_i\left(\alpha^{NE}\right)\right\} = u_i\left(\bar{\alpha}\right) = v^*_i - \varepsilon/2$, again by (21). In total, we have $\zeta_i\left(v^*_i + \zeta_i\varepsilon/4 - u_i\left(\alpha^*_t\right)\right) \leq -\varepsilon/4$. Since obviously $\zeta_i\left(v^*_i + \zeta_i\varepsilon/4 - u_i\left(\alpha^*_t\right)\right) \geq -2\bar{u}$ and $-\bar{u} - \zeta_i u_i\left(\alpha^*_t\right) \geq -2\bar{u}$, we have

$$\zeta_i \pi^v_{i,t} = \begin{cases} \delta^{t-1}\zeta_i\left(v^*_i + \zeta_i\varepsilon/4 - u_i\left(\alpha^*_t\right)\right) & \text{if } \theta_{j,t} = 0 \text{ for all } j \neq i \\ \delta^{t-1}\left(-\bar{u} - \zeta_i u_i\left(\alpha^*_t\right)\right) & \text{otherwise} \end{cases} \in \left[-2\bar{u}\delta^{t-1}, 0\right].$$

For (42), note that $S_{i,t} = 0$ implies $\theta_{i,t} = 0$, and hence

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

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 $\left|f_{i,a_{i,t}}\left(y_{i,t}\right)\right| \leq 2\bar{u}/\eta$ by (11), this inequality implies (42).

For the second part of the lemma, by (27), we have

$$
\begin{aligned}
\zeta_i c_i &= \zeta_i \left(-\mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} (v_i^* + \zeta_i \varepsilon/4) - \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} \zeta_i \bar{u} \right) \right] + \frac{1-\delta^T}{1-\delta} v_i^* \right) \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} (-\varepsilon/4) + \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} \underbrace{(\bar{u} + \zeta_i v_i^*)}_{\in[0,2\bar{u}]} \right) \right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \left(\mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 0 \right\} (-\varepsilon/4) + \mathbf{1}\left\{ \max_{j\neq i} \theta_{j,t} = 1 \right\} 2\bar{u} \right) \right] \\
&\leq -\frac{1-\delta^T}{1-\delta} \left(\left(1 - \frac{\varepsilon}{20\bar{u}}\right) \frac{\varepsilon}{4} + \left(\frac{\varepsilon}{20\bar{u}}\right) 2\bar{u} \right) \quad \text{(by (30))} \\
&\leq -\frac{1-\delta^T}{1-\delta} \frac{\varepsilon}{8} \quad (\text{as } \varepsilon < \bar{u}/2).
\end{aligned}
$$

A.5 Proof of Theorem 3

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

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

By standard arguments, for every $\beta \in [0,1]$, there exists a linear equilibrium with the same weights b and expected payoff $v = (1-\beta)\underline{v} + \beta\bar{v}$ such that the set $\{v : \exists h \text{ s.t. } v = u(h)\}$ is closed and there exist histories h and h' such that $w(h) = \underline{v}$ and $w(h') = \bar{v}$. Since $M(\varepsilon)$ is convex, it thus suffices to show that $\underline{v}, \bar{v} \in M(\varepsilon)$.

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*

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

If the constraint $\omega(r,y) \in [0, (\delta/(1-\delta))\bar{u}]$ is replaced with $\omega(r,y) \in [-(\delta/(1-\delta))\bar{u}, 0]$, 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

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

Since $v \in E$ and $w(r,y) \in E$ for all r, y, we have

$$v_i - w_i(r,y) = b_i(v_1 - w_1(r,y)) \quad \text{for all } i \neq 1, r \in A, y \in Y.$$

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 $\omega(r,y) = (\delta/(1-\delta))(\bar{v}_1 - w_1(r,y)) \in [0, (\delta/(1-\delta))\bar{u}]$ 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 argument applies. ■

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

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

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

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

We conclude that $\sum_i g_i(\alpha)/N$ is bounded by the solution to the program

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

where the last constraint follows because $\mathbb{E}[\omega_1(y)|r]=u_1(r)-\bar{v}_1\leq\bar{u}$. 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 maximization is over (r,a,ω). Recall that $p(y|a)=\sum_x \pi_{a,x}q(y|x)$. Note that 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 $\hat{p}=Mp$.) To see this, fix any $(r,a,\hat{\omega})$ that is feasible with signal distribution $\hat{p}$, and define $\omega(r,y)=\sum_{\hat{y}}M(\hat{y}|y)\hat{\omega}(r,\hat{y})$. Then

$$\sum_y p(y|a)\omega(r,y)=\sum_y p(y|a)\sum_{\hat{y}}M(\hat{y}|y)\hat{\omega}(r,\hat{y})=\sum_{\hat{y}}\hat{p}(\hat{y}|a)\hat{\omega}(r,\hat{y}) \text{ for all } a,$$

so (r,a,ω) is feasible with signal distribution p and yields the same value in the sub-program.

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 the 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

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

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

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

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

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

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

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

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

$$\max_{r,a,\omega} \frac{1}{N} \sum_i \left|\mathbb{E}\left[\omega\left(x\right)|a_i, r_{-i}\right] - \mathbb{E}\left[\omega\left(x\right)|r\right]\right| \quad \text{s.t.} \tag{43}$$

$$\omega\left(x\right) \in \left[0, \frac{\delta}{1-\delta}\bar{u}\right] \qquad \text{for all } x \in A, \tag{44}$$

$$\mathbb{E}\left[\omega\left(x\right)|r\right] \leq \bar{u}, \tag{45}$$

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

$$\max_{a_{-i},r_{-i},\omega_{-i}:A_{-i}\to\mathbb{R}} \frac{1}{N}\sum_{j\neq i}\left|\mathbb{E}\left[\omega_{-i}\left(x_{-i}\right)|a_j,r_{-j}\right]-\mathbb{E}\left[\omega_{-i}\left(x_{-i}\right)|r\right]\right| \quad \text{s.t. (44), (45).}$$

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 $\omega(x) = \omega_{-i}(x_{-i})$ for all x. We thus assume that $a_i \neq r_i$ for all i.

We now show that the value of program (43)–(45) converges to 0, which completes the proof. Note that this value is less than the sum of the values of the two programs

$$\begin{aligned}
&\max_{r,a,\omega}\frac{1}{N}\sum_i\left(\mathbb{E}\left[\omega\left(x\right)|a_i,r_{-i}\right]-\mathbb{E}\left[\omega\left(x\right)|r\right]\right)_+ \quad \text{s.t. (44), (45), and}\\
&\max_{r,a,\omega}\frac{1}{N}\sum_i\left(\mathbb{E}\left[\omega\left(x\right)|r\right]-\mathbb{E}\left[\omega\left(x\right)|a_i,r_{-i}\right]\right)_+ \quad \text{s.t. (44), (45).}
\end{aligned}$$

We show that the value of the first of these programs converges to 0. A symmetric argument shows that the value of the second program also converges to 0, which implies that the value of program (43)–(45) converges to 0 as well, as desired.

Letting $\lambda \geq 0$ denote the multiplier on (45), it is immediate that the solution to the first program takes the form

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

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

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

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

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

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

$$\max_{n^* \in \{0,1,\ldots,N\}, \beta \in [0,1]} \frac{\delta}{1-\delta} \bar{u} (1-2\underline{\pi}) (\beta \Pr(n_{-i} = n^* - 1|r_{-i}) + (1-\beta) \Pr(n_{-i} = n^*|r_{-i})) \tag{46}$$

$$\text{s.t.} \quad \beta \Pr(n = n^*|r) + \Pr(n \geq n^* + 1|r) \leq \frac{1-\delta}{\delta}, \tag{47}$$

where

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

Fix $\rho > 0$ and a sequence, indexed by k, of games with $(1-\delta)\exp(N^{1-\rho}) > k$ and pairs (n^*, β) that satisfy the constraint (47). 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 the objective (46) 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 the objective (46) 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 $|\underline{\pi} - (n^*-1)/(N-1)| > c$. By Hoeffding's inequality (Boucheron, Lugosi, and Massart, 2013, Theorem 2.8),

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

Hence, for every $\tilde{k}$, there is some $k \geq \tilde{k}$ such that the value of (46) is at most

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

Since $(1-\delta)\exp(N^{1-\rho}) \to \infty$, we have $\exp(-2c^2(N-1))/(1-\delta) \to 0$ for all $c > 0$, and hence (46) 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

$$\left| \underline{\pi} - \frac{n^* - 1}{N - 1} \right| \leq c. \tag{48}$$

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*

$$\frac{\beta \Pr(n = n^*|r) + \Pr(n \geq n^* + 1|r)}{\beta \Pr(n_{-i} = n^* - 1|r_{-i}) + (1 - \beta) \Pr(n_{-i} = n^*|r_{-i})} \geq m(1 - \gamma). \tag{49}$$

Proof. Fix $c > 0$ and take k sufficiently large that (48) holds. For any $m \in \mathbb{N}$, we have

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

By (48), for any $\gamma' > 0$, for sufficiently large k we have $(n^* - 1)/(n^* + m) \geq 1 - \gamma'$, and hence

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

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

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

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 $k\geq\tilde{k}$, (49) 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 (46) satisfies

$$\begin{aligned}
&\frac{\delta}{1-\delta}\bar{u}\left(1-2\underline{\pi}\right)\left(\beta\Pr\left(n_{-i}=n^*-1|r_{-i}\right)+\left(1-\beta\right)\Pr\left(n_{-i}=n^*|r_{-i}\right)\right)\\
\leq\;&\bar{u}\left(1-2\underline{\pi}\right)\frac{\beta\Pr\left(n_{-i}=n^*-1|r_{-i}\right)+\left(1-\beta\right)\Pr\left(n_{-i}=n^*|r_{-i}\right)}{\beta\Pr\left(n=n^*|r\right)+\Pr\left(n\geq n^*+1|r\right)}\quad\text{(by (47))}\\
\leq\;&\frac{\bar{u}\left(1-2\underline{\pi}\right)}{m\left(1-\gamma\right)}\quad\text{(by (49)).}
\end{aligned}$$

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.

A.6 Proof of Proposition 1

Fix $\alpha\in A\left(\varepsilon\right)$. Let $J=\left\{i:g_i\left(\alpha\right)>c\varepsilon/2\right\}$, and note that $|J|\leq N/c$. Let $\tilde{\alpha}\in\Delta\left(A\right)$ be any action distribution whose marginal on $A_{I\backslash J}$ coincides with that of α and satisfies $g_i\left(\tilde{\alpha}\right)\leq c\varepsilon$ 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 $I\backslash J$ is held fixed. Since $\left|u_i\left(a'_j,a_{-j}\right)-u_i\left(a\right)\right|\leq b/N$ for all $i\neq j,a'_j,a$, and the actions of at most N/c players differ between $\tilde{\alpha}$ and α, we have $g_i\left(\tilde{\alpha}\right)\leq g_i\left(\alpha\right)+2b/c$ for each $i\in I\backslash J$. Since $g_i\left(\alpha\right)\leq c\varepsilon/2$ (as $i\in I\backslash J$) and $2b/c\leq c\varepsilon/2$ (as $c\geq\sqrt{4b/\varepsilon}$), we have $g_i\left(\tilde{\alpha}\right)\leq c\varepsilon$. Since we assumed that $g_i\left(\tilde{\alpha}\right)\leq c\varepsilon$ for all $i\in J$, we have $g_i\left(\tilde{\alpha}\right)\leq c\varepsilon$ for all $i\in I$, and hence $u\left(\tilde{\alpha}\right)\in CE\left(c\varepsilon\right)$. Finally, since the actions of at most N/c players differ between $\tilde{\alpha}$ and α, we have $\left|u_i\left(\tilde{\alpha}\right)-u_i\left(\alpha\right)\right|\leq b/c\leq\bar{u}/c$ for all $i\in I\backslash J$, and by definition of $\bar{u}$ we have $\left|u_i\left(\tilde{\alpha}\right)-u_i\left(\alpha\right)\right|\leq\bar{u}$ for all $i\in J$. Since $|J|\leq N/c$, we have $\left|\sum_{i\in I}u_i\left(\tilde{\alpha}\right)-\sum_{i\in I}u_i\left(\alpha\right)\right|\leq N\bar{u}/c$.

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

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

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

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

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

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

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

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

[21] 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.

[22] 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.

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

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

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

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

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

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

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

[30] Hébert, Benjamin. "Moral Hazard and the Optimality of Debt." *Review of Economic Studies* 85.4 (2018): 2214-2252.

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

[32] Hoffmann, Florian, Roman Inderst, and Marcus Opp. "Only Time will Tell: A Theory of Deferred Compensation." *Review of Economic Studies* 88.3 (2021): 1253-1278.

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

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

[35] Jewitt, Ian, Ohad Kadan, and Jeroen M. Swinkels. "Moral Hazard with Bounded Payments." *Journal of Economic Theory* 143.1 (2008): 59-82.

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

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

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

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

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

[41] 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.

[42] 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).

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

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

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

[46] 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.

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

[48] Sadzik, Tomasz, and Ennio Stacchetti. "Agency Models with Frequent Actions." *Econometrica* 83.1 (2015): 193-237.

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

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

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

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

[53] Sugaya, Takuo, and Alexander Wolitzky, "Informational Requirements for Cooperation." *Working Paper* (2022a).

[54] Sugaya, Takuo, and Alexander Wolitzky, "Rate of Convergence in Repeated Games: A Universal Speed Limit." *Working Paper* (2022b).