

Repeated Games with Many Players*

Takuo Sugaya

Alexander Wolitzky

Stanford GSB

MIT

April 10, 2026

Abstract

Motivated by the problem of sustaining cooperation in large groups with limited information, we analyze the relationship between group size, monitoring precision, and discounting in repeated games with independent, player-level noise. The viability of cooperation under independent noise is linked to the *per-capita channel capacity* of the stage game monitoring structure. We show that cooperation is impossible if the per-capita channel capacity is much smaller than the discount rate. A folk theorem under a novel identification condition provides a near converse. If attention is restricted to *co-linear equilibria* (a generalization of strongly symmetric equilibria), cooperation is possible only under much more severe parameter restrictions.

Keywords: repeated games, large population, imperfect monitoring, noise, mutual information, channel capacity, folk theorem

JEL codes: C72, C73, D86

*For helpful comments, we thank Tommaso Denti, Drew Fudenberg, Stephen Morris, Andy Postlewaite, Yuliy Sannikov, Satoru Takahashi, and many seminar participants.

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

—David Hume, *A Treatise of Human Nature*

1 Introduction

Hume's intuition notwithstanding, large groups of individuals often have a remarkable capacity for cooperation, even in the absence of external contractual enforcement (Ostrom, 1990; Ellickson, 1991; Seabright, 2004). Cooperation in large groups typically relies on accurate monitoring of individual agents' actions, together with sanctions that narrowly target deviators. For example, these are key features of the community resource management problems documented by Ostrom (1990), the local public goods problems studied by Miguel and Gugerty (2005), and the group lending settings studied by Karlan (2007) and Feigenberg, Field, and Pande (2013). Large cartels seem to operate similarly. For instance, 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—monitors its members' sales, and producers who violate its rules can 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, it does not seem that individual maple syrup producers—or the farmers, fishers, and herders studied by Ostrom, or the villagers in the development economics studies cited above—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 targeted 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 the parameters of the game other than the discount

factor δ and considers the limit as $\delta \rightarrow 1$. 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).¹ Also, standard results are based on statistical identification conditions that hold generically regardless of the number of players, so they do not capture the possibility that more information may be required to support cooperation in larger groups. Finally, there is typically a vast multiplicity of cooperative equilibria in the $\delta \rightarrow 1$ limit, so 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 targeted sanctions (e.g., fines) or collective ones (e.g., price wars; or, in Hume’s example, “the abandoning of the whole project”).

This paper extends the study of repeated games by investigating the relationship between the discount factor, the number of players, and the monitoring structure. Rather than focusing on the $\delta \rightarrow 1$ limit, we let all these features vary flexibly, assuming only a uniform upper bound on the range of the stage game payoffs and a uniform lower bound on the amount of independent, player-level “noise.” Our main results provide necessary and (somewhat stronger) sufficient conditions for cooperation as a function of discounting, group size, and a measure of societal information or monitoring precision. We also establish severe obstacles to cooperation under collective incentives. In sum, we show that large-group cooperation requires a high level of patience and/or information, and that it cannot be based on collective incentives for reasonable parameter values.

We now preview our model and results. We model independent, player-level noise by assuming that, in each period of the repeated game, each player i ’s action a_i stochastically determines an *individual outcome* x_i , independently across players, and that the distribution of the public signal y depends on the action profile $a = (a_i)$ only through the outcome profile $x = (x_i)$.² As the following example illustrates, absent noise there may be no tradeoff between discounting, group size, and monitoring precision.

Example 1. Suppose N players repeatedly play a prisoner’s dilemma with a binary

¹Comparing $1 - \delta$ and N has a concrete meaning in some games. For example, in repeated games with random matching (where a player meets each opponent with probability $1/(N - 1)$) with a fixed real-time discount rate r and meetings every Δ units of time (so the discount factor is $\delta = e^{-r\Delta}$), the expected time between meetings for each pair of players is approximately $(1 - \delta)N$, so $(1 - \delta)N \rightarrow 0$ means each pair of players meets frequently, while $(1 - \delta)N \rightarrow \infty$ means each pair of players meets infrequently.

²This setup follows prior work such as Fudenberg, Levine, and Pesendorfer (1998) and al-Najjar and Smorodinsky (2000, 2001).

public signal $y_t \in \{0, 1\}$ in each period, where $y_t = 0$ if *every* player cooperates in period t , and $y_t = 1$ if *any* player defects in period t . A player’s stage game payoff is the fraction of players who cooperate, less a constant (independent of N) if she cooperates herself. In this game, each player’s action is pivotal in determining y_t when the others cooperate, so the range of values for the discount factor for which mutual cooperation is a sequential equilibrium outcome is independent of N . Thus, a single “bit” of information in each period (the binary signal y_t) can form the basis for cooperation in an arbitrarily large group of players in a repeated game where the range of stage game payoffs, the cost of cooperation, and the discount factor are all fixed independent of N .

Now introduce noise. For example, let each player “tremble” in her choice of action with probability π , independently across players and periods, with π fixed independent of N . Assume that the distribution of the public signal depends only on the players’ realized actions, not their intended actions. Then a single bit of information in each period can no longer motivate cooperation by a large group of players for a fixed discount factor, because the probability that a single player is pivotal for the realization of a vector of T binary signals, for any fixed T , goes to zero as $N \rightarrow \infty$, and signal realizations in the distant future have only a small impact on the players’ payoffs. Moreover, a novel implication of our results is that, for any fixed δ , the number of bits of information (e.g., the log of the number of possible stage game signal realizations) that is required to support cooperation is not only increasing in N , but in fact is proportional to N .

In general, we model the “amount of information” available to society as the *channel capacity*, C , of the conditional signal distribution $q(Y|X)$.³ Channel capacity is a standard measure in information theory, which in our context is defined as the maximum *mutual information* $I(X;Y)$ between the profile of individual outcomes X and the signal Y , for any distribution of action profiles: that is, the expected reduction in uncertainty about the outcome profile X that results from observing the signal Y . Channel capacity is a convenient measure of information in games with independent noise, because it provides a bound for the average influence of the players’ actions a on the distribution of the signal Y . In addition, channel capacity is bounded by the entropy of the signal Y , which in turn is bounded by

³We use capital letters for random variables and lower-case letters for their realizations, so X and Y denote the (random) outcome profile and signal.

the log of the number of possible signal realizations.⁴ Our bound on equilibrium incentives in terms of channel capacity thus immediately implies a bound in terms of the number of possible signal realizations. Hence, our results based on channel capacity improve on prior results based on the number of possible signal realizations (in particular, results of Fudenberg, Levine, and Pesendorfer (1998) and al-Najjar and Smorodinsky (2000, 2001), which we discuss below).

We obtain three results on the relationship between information, discounting, and group size in repeated games with independent, player-level noise.

First, if $(1 - \delta)N/C$ —the ratio of the discount rate $1 - \delta$ and the per-capita channel capacity C/N —is large, then cooperation is impossible: all repeated game Nash equilibrium outcomes are consistent with approximately myopic (static optimal) play.⁵ This shows that large-group cooperation requires a high level of patience and/or information. We prove this result by combining inequalities for mutual information in games with independent noise with bounds on the strength of players’ equilibrium incentives in repeated games, which we developed in a companion paper (Sugaya and Wolitzky 2023, henceforth SW).

Second, this result is tight up to a factor of $\log N$. In particular, under a *random auditing* monitoring structure, where each player’s individual outcome is publicly observed with independent probability C/N , a folk theorem holds if $(1 - \delta)N(\log N)/C \rightarrow 0$. We prove this result as a corollary of a more general folk theorem for repeated games with product structure monitoring, where the monitoring structure, discount factor, and stage game all vary simultaneously.

Third, we contrast these results with the situation under collective incentive-provision. We model collective incentives by focusing on equilibria where the players’ equilibrium continuation payoffs are co-linear. When the stage game is symmetric and the continuation payoff vectors lie on the 45° line, co-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 cooperation in a co-linear equilibrium is impossible, unless $(1 - \delta)^{-1}$ is exponentially large relative to N .⁶ Practically speaking, this is an impossibility theorem for large-group

⁴See, e.g., Cover and Thomas (2006, Theorem 2.6.4).

⁵Throughout, we refer to δ as the *discount factor* and $1 - \delta$ as the *discount rate*. Our notion of “approximately myopic play” is that the average static deviation gain across players is small.

⁶It is well-known that strongly symmetric equilibria are typically less efficient than general perfect public equilibria in repeated games. We instead show that the relationship between N and δ required for *any*

cooperation under collective incentives. The intuition is that optimal co-linear incentives take the form of a *tail test*, where the players are all punished if the number of “good” outcomes x_i falls short of a threshold n^* . For such a test, the ratio of the probability that one player’s action is pivotal for the tail test and the probability that the test is failed converges to zero as $N \rightarrow \infty$, unless these probabilities are both exponentially small in N . But a tail test where the pivot probability is exponentially small provides only small incentives, unless the size of the punishment—which is proportional to $(1 - \delta)^{-1}$ —is exponentially large.

Related Literature. Prior research on repeated games has established folk theorems in the $\delta \rightarrow 1$ limit for fixed N , as well as impossibility theorems for cooperation in the $N \rightarrow \infty$ limit for fixed δ , but has not studied the relationship between δ and N required to support cooperation. The closest paper is our companion work, SW. That paper establishes necessary and sufficient conditions for cooperation in repeated games as a function of discounting and monitoring precision. Relative to SW, the current paper introduces two features specific to large-population games: independent noise and the three-way relationship between group size, discounting, and monitoring. Independent noise is crucial for all of our results, while letting N vary together with discounting and monitoring is the key novelty in our folk theorem (Theorem 2).⁷

Other than those in SW, the most relevant necessary conditions for cooperation are those of Fudenberg, Levine, and Pesendorfer (1998), al-Najjar and Smorodinsky (2000, 2001), Pai, Roth, and Ullman (2016), and Awaya and Krishna (2016, 2019). Following earlier work by Green (1980) and Sabourian (1990), these papers establish conditions under which equilibrium play in a repeated game is approximately myopic as $N \rightarrow \infty$ for fixed δ .⁸ These conditions can be adapted to the case where N , δ , and monitoring vary together, but the implied conditions under which payoffs are consistent with approximately myopic play are weaker than ours (and are not tight up to log terms). As we explain in Section 3, the key difference is that prior results rely on bounding the strength of players’ incentives period-by-period, rather than bounding the strength of incentives over the entire game, which results in

non-trivial incentive provision differs dramatically between strongly symmetric equilibria and general ones.

⁷SW was split off from an earlier version of the current paper. SW contains the results from the original paper that do not rely on independent noise or letting N vary together with discounting and monitoring, while the current paper contains the results that do rely on these features.

⁸Awaya and Krishna focus on conditions under which cheap talk is valuable. Green and Sabourian’s papers impose a continuity condition on the mapping from action distributions to signal distributions. Continuity is implied by independent noise.

a higher-order bound in the discount rate than that given in SW ($(1 - \delta)^{-1}$ vs. $(1 - \delta)^{-1/2}$). In sum, prior work has established impossibility theorems for cooperation as $N \rightarrow \infty$ for fixed δ , while our paper tightly (up to log terms) characterizes the tradeoff between N , δ , and monitoring that is required for supporting cooperation.⁹

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

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

Entropy methods have been used in repeated games to study issues including complexity and bounded recall (Neyman and Okada, 1999, 2000; Hellman and Peretz, 2020), communication (Gossner, Hernández, and Neyman, 2006), and reputation effects (Gossner, 2011; Ekmekci, Gossner, and Wilson, 2012; Faingold, 2020). However, other than the shared reliance on entropy methods, these papers are not very related to ours.

We also relate to papers on optimal monitoring design, although we consider only asymptotic results rather than exact optimality for fixed parameters. In static moral hazard problems, optimal monitoring design subject to information-theoretic constraints has been studied by Georgiadis and Szentes (2020), Li and Yang (2020), and Hoffmann, Inderst, and Opp (2021). Random auditing, which we find to be approximately optimal, also arises in costly state-verification models (Reinganum and Wilde, 1985; Border and Sobel, 1987; Mookherjee and Png, 1989).

Finally, in Sugaya and Wolitzky (2021) we studied the relationship between N , δ , and

⁹Farther 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), simulations (Bowles and Gintis, 2011; Chapter 4), and experiments (Camera, Casari, and Bigoni, 2013).

information in repeated random-matching games with private monitoring and incomplete information, where each player is “bad” (a *Defect* commitment type) with positive probability. In that model, society has enough information to determine which players are bad after a single period of play, but this information is disaggregated, and supporting cooperation requires sufficiently quick information diffusion. In contrast, the current paper has complete information and public monitoring, 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 obstacle to cooperation is that society’s information is *insufficient* or *disaggregated* distinguishes large-population repeated game models, such as Fudenberg, Levine, and Pesendorfer (1998), al-Najjar and Smorodinsky (2001), and the current paper, from community enforcement models, such as Kandori (1992), Ellison (1994), and our earlier paper.

2 Repeated Games with Independent Noise

We consider a general model of repeated games with independent, player-level noise.

Stage Games. A stage game $G = (I, \mathcal{A}, u)$ consists of a finite set of players $I = \{1, \dots, N\}$, a finite product set of actions $\mathcal{A} = \times_{i \in I} \mathcal{A}_i$, and a payoff function $u_i : \mathcal{A} \rightarrow \mathbb{R}$ for each $i \in I$. We assume that $|\mathcal{A}_i| \geq 2$ for all i , and denote the range of player i ’s payoff function by $\bar{u}_i = \max_{a, a'} u_i(a) - u_i(a')$. Given \bar{u}_i , by adding a constant, without affecting incentives, we can assume that $|u_i(a)| \leq \bar{u}_i/2$. For any $\bar{u} > 0$, we say that payoffs are \bar{u} -bounded if $\bar{u}_i \leq \bar{u}$ for all i .

Payoff Sets. The feasible payoff set is $F = \text{co} \{ \{u(a)\}_{a \in \mathcal{A}} \} \subseteq \mathbb{R}^N$ (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 $\alpha_n \in \Delta(\mathcal{A})$ 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 any $\varepsilon > 0$, let $B(\varepsilon)$ denote the set of payoff vectors v such that the cube with center v and side-length 2ε lies entirely within F^* : that is, $B(\varepsilon) = \{v \in \mathbb{R}^N : B_v(\varepsilon) \subseteq F^*\}$, where $B_v(\varepsilon) = \times_{i \in I} [v_i - \varepsilon, v_i + \varepsilon]$. We compute

¹⁰Here and throughout, we linearly extend payoff functions to mixed actions. In this paper, “Nash equilibrium” always allows mixed strategies.

$B(\varepsilon)$ in a public goods game in Appendix B. Our folk theorem (Theorem 2) will provide conditions under which all payoff vectors in $B(\varepsilon)$ are attained as repeated game equilibria.

In contrast, our impossibility theorem (Theorem 1) will provide conditions under which all repeated game equilibria are “ ε -myopic.” To define this, let a *manipulation* for a player i be a mapping $s_i : \mathcal{A}_i \rightarrow \Delta(\mathcal{A}_i)$. The interpretation is that when player i is “supposed” to play a_i , she instead plays $s_i(a_i)$. Player i ’s *gain* from manipulation s_i at an action profile distribution $\alpha \in \Delta(\mathcal{A})$ is

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

Player i ’s *maximum gain* at $\alpha \in \Delta(\mathcal{A})$ is $\bar{g}_i(\alpha) = \max_{s_i: \mathcal{A}_i \rightarrow \Delta(\mathcal{A}_i)} g_i(s_i, \alpha)$. For any $\varepsilon \geq 0$, the set of ε -myopic action distributions is

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

and the set of ε -myopic payoff vectors is

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

Note that $\mathcal{A}(\varepsilon)$ and $V(\varepsilon)$ are convex polytopes, and that $\mathcal{A}(0)$ is the set of static correlated equilibria, with $V(0)$ the corresponding payoff set. Since $\mathcal{A}(\varepsilon) \rightarrow \mathcal{A}(0)$ and $V(\varepsilon) \rightarrow V(0)$ as $\varepsilon \rightarrow 0$, $\mathcal{A}(\varepsilon)$ and $V(\varepsilon)$ approximate the sets of static correlated equilibria and the corresponding payoff set as $\varepsilon \rightarrow 0$. In general, an action distribution α is ε -myopic if the per-player average deviation gain at α is less than ε . If the game is symmetric and α is a symmetric distribution, this implies that all players have deviation gains smaller than ε . Otherwise, it allows a small fraction of players to have large gains. In Appendix A, we provide some results comparing $V(\varepsilon)$ with the (smaller) set of payoff vectors that are consistent with *all* players having small deviation gains (i.e., the set of stage game ε -correlated equilibria).

Noise. We assume the presence of independent, player-level noise. Formally, there is a finite product set of *outcome profiles* $\mathcal{X} = \times_{i \in I} \mathcal{X}_i$, where \mathcal{X}_i is the set of *individual outcomes* for player i . When action profile $a \in \mathcal{A}$ is played, the outcome profile $x \in \mathcal{X}$ is drawn from a product distribution $\pi(\cdot|a) = \times_i \pi_i(\cdot|a_i)$, where $\pi_i(\cdot|a_i) \in \Delta(\mathcal{X}_i)$. We call the pair (\mathcal{X}, π)




Figure 1: The noise structure (\mathcal{X}, π) and the outcome monitoring structure (\mathcal{Y}, q) jointly determine the action monitoring structure (\mathcal{Y}, p) .

a *noise structure*. Let $\underline{\pi}_i = \min_{a_i, x_i} \pi_i(x_i|a_i)$ and assume that $\min_i \underline{\pi}_i > 0$. For any $\underline{\pi} > 0$, we say that noise is $\underline{\pi}$ -*bounded* if $\underline{\pi}_i \geq \underline{\pi}$ for all i . Note that if noise is $\underline{\pi}$ -bounded then $|\mathcal{X}_i| \leq 1/\underline{\pi}$ for all i . We assume that $|\mathcal{X}_i| \geq 2$ for at least one player i , which implies that noise can be $\underline{\pi}$ -bounded only for $\underline{\pi} \leq 1/2$.

A simple example of a noise structure arises when there is independent noise in the execution of the players' actions, so that a_i is player i 's intended action and x_i is her realized action. In this case, $\mathcal{X} = \mathcal{A}$ and $\pi_i(a'_i|a_i)$ is the probability that player i “trembles” to a'_i when she intends to take a_i . Another example is a moral hazard in teams problem, where $\mathcal{A}_i \subseteq [\underline{\pi}, 1 - \underline{\pi}]$ is a set of “effort levels” corresponding to the probability of “success” on an individual task, $\mathcal{X}_i = \{Success, Failure\}$, and $\pi_i(Success|a_i) = a_i$.

Monitoring. An *outcome monitoring structure* (\mathcal{Y}, q) consists of a finite set of possible signal realizations \mathcal{Y} and a conditional probability distribution $q(\cdot|x) \in \Delta(\mathcal{Y})$ for each outcome profile x . The signal distribution thus depends only on the outcome profile and not directly on the action profile. In other words, if we view the action profile, the outcome profile, and the signal as random variables A , X , and Y , they form a Markov chain $A \rightarrow X \rightarrow Y$.

Given an outcome monitoring structure (\mathcal{Y}, q) , we denote the probability of signal profile y at action profile a by $p(y|a) = \sum_x \pi(x|a) q(y|x)$. We refer to the pair (\mathcal{Y}, p) as the *action monitoring structure* induced by $(\mathcal{X}, \pi, \mathcal{Y}, q)$. Without loss, we assume that for every $y \in \mathcal{Y}$, there exists $x \in \mathcal{X}$ such that $q(y|x) > 0$. Since $\underline{\pi}_i > 0$ for each i , this implies that p has full support: $p(y|a) > 0$ for all a, y . We also linearly extend p to mixed actions: for $\alpha \in \Delta(A)$, $p(y|\alpha) = \sum_a \alpha(a) p(y|a)$.

Figure 1 summarizes the relationship between the noise structure (\mathcal{X}, π) , the outcome monitoring structure (\mathcal{Y}, q) , and the action monitoring structure (\mathcal{Y}, p) .

Finally, for any action profile $a \in \mathcal{A}$, let $\varphi^a \in \Delta(\mathcal{X} \times \mathcal{Y})$ denote the joint distribution

on $\mathcal{X} \times \mathcal{Y}$ when a is played, so that X has distribution $\pi(\cdot|a)$, and, conditional on each realization x , Y has distribution $q(\cdot|x)$.

Repeated Games. A repeated game with independent noise $\Gamma = (I, \mathcal{A}, u, \mathcal{X}, \pi, \mathcal{Y}, q, \delta)$ is described by a stage game (I, \mathcal{A}, u) , a noise structure (\mathcal{X}, π) , an outcome monitoring structure (\mathcal{Y}, q) , and a discount factor $\delta \in [0, 1)$. In each period $t = 1, 2, \dots$, (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 from distribution $\pi(\cdot|a)$, and (iv) the signal y is drawn from distribution $q(\cdot|x)$ and is publicly observed.¹¹ A history h_i^t for player i at the beginning of period t thus takes the form $h_i^t = ((z_{t'}, a_{i,t'}, y_{t'})_{t'=1}^{t-1}, z_t)$. A strategy σ_i for player i maps histories h_i^t to distributions over actions $a_{i,t}$. A strategy σ_i is *public* if it depends on h_i^t only through the *public history* $h^t = ((z_{t'}, y_{t'})_{t'=1}^{t-1}, z_t)$. A *Nash equilibrium* is a strategy profile where each player's strategy maximizes her discounted expected payoff. 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. Let $E \subseteq \mathbb{R}^N$ denote the set of PPE payoff vectors .

Repeated Game Outcomes and Occupation Measures. A repeated game outcome $\mu \in \Delta((\mathcal{A} \times \mathcal{X} \times \mathcal{Y})^\infty)$ is a distribution over infinite paths of actions, individual outcomes, and signals. Each strategy profile σ induces a unique outcome μ . In turn, each outcome μ induces a marginal distribution over period t action profiles $\alpha_t^\mu \in \Delta(\mathcal{A})$, as well as an *occupation measure* over action profiles, defined as

$$\alpha^\mu = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \alpha_t^\mu.$$

The occupation measure describes the “discounted expected fraction of periods” where each action profile is played in the course of the repeated game. Intuitively, it captures how the stage game is played “on average.” Note that the players’ payoffs are determined by the occupation measure, as

$$(1 - \delta) \sum_t \delta^{t-1} \sum_a \alpha_t^\mu(a) u(a) = \sum_a (1 - \delta) \sum_t \delta^{t-1} \alpha_t^\mu(a) u(a) = \sum_a \alpha^\mu(a) u(a) = u(\alpha^\mu).$$

¹¹It is natural to require that players’ realized payoffs depend only on their own actions and the signal. However, this assumption is not necessary for our analysis.

We say that a repeated game strategy profile σ is ε -myopic if the corresponding occupation measure $\alpha^\mu \in \Delta(\mathcal{A})$, viewed as a correlated action profile in the stage game, is ε -myopic: that is, if $\alpha^\mu \in \mathcal{A}(\varepsilon)$. In this case, the players' repeated game payoffs are also ε -myopic: the repeated game payoff vector is $u(\alpha^\mu) \in V(\varepsilon)$.

3 Necessary Conditions for Cooperation

This section develops our first main result: cooperation is impossible if the per-capita channel capacity of the outcome monitoring structure (\mathcal{Y}, q) is much smaller than the discount rate $1-\delta$. Before stating the result, we define channel capacity and establish some of its properties.

3.1 Information Theory Preliminaries

Mutual Information and Channel Capacity. Given a distribution of outcomes $\xi \in \Delta(\mathcal{X})$, a standard measure of the informativeness of the signal Y about the outcome X is the *mutual information* between the random variables X and Y , defined as

$$I^\xi(X; Y) = \sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \xi(x) q(y|x) \log \left(\frac{q(y|x)}{\sum_{x' \in \mathcal{X}} \xi(x') q(y|x')} \right).^{12}$$

Mutual information measures the expected reduction in uncertainty (entropy) about X that results from observing Y . The mutual information between X and Y is an endogenous object in our model, as it depends on the distribution ξ of X , which in turn is determined by the players' actions, a .

We denote the set of outcome distributions ξ that can arise for some action distribution α under noise structure (\mathcal{X}, π) by

$$\Xi = \left\{ \xi \in \Delta(\mathcal{X}) : \exists \alpha \in \Delta(\mathcal{A}) \text{ such that } \xi(x) = \sum_{a \in \mathcal{A}} \alpha(a) \pi(x|a) \text{ for all } x \in \mathcal{X} \right\}.$$

¹²In this paper, all logarithms are base e . For a joint distribution $\varphi \in \Delta(\mathcal{X} \times \mathcal{Y})$ with marginals $\varphi_{\mathcal{X}} \in \Delta(\mathcal{X})$ and $\varphi_{\mathcal{Y}} \in \Delta(\mathcal{Y})$ and conditionals $\varphi_{\mathcal{X}|Y} \in \Delta(\mathcal{X})$ and $\varphi_{\mathcal{Y}|X} \in \Delta(\mathcal{Y})$, the definition of mutual information is usually written as $\sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \varphi(x,y) \log \left(\frac{\varphi(x,y)}{\varphi_{\mathcal{X}}(x)\varphi_{\mathcal{Y}}(y)} \right)$. This equals $\sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \varphi_{\mathcal{X}}(x) \varphi_{\mathcal{Y}|X}(y|x) \log \left(\frac{\varphi_{\mathcal{Y}|X}(y|x)}{\sum_{x' \in \mathcal{X}} \varphi_{\mathcal{X}}(x') \varphi_{\mathcal{Y}|X}(y|x')} \right)$, which is the form of the definition used above.

Finally, define the *channel capacity* of the tuple $(\mathcal{X}, \pi, \mathcal{Y}, q)$ as

$$C = \max_{\xi \in \Xi} I^\xi(X; Y).$$

Channel capacity is an exogenous measure of the informativeness of Y about X , as it is determined by the noise structure (\mathcal{X}, π) and the outcome monitoring structure (\mathcal{Y}, q) . Channel capacity plays a central role in information theory as the maximum rate at which information can be transmitted over a noisy channel (*Shannon's channel coding theorem*; Cover and Thomas, 2006, Theorem 7.7.1). Our analysis does not use this theorem; we only use channel capacity as an exogenous upper bound on mutual information. In turn, mutual information is a useful measure for our analysis because it satisfies two key inequalities, which we now discuss.

Mutual Information Inequalities. The first mutual information inequality we use relates $I^{\pi(\cdot|a)}(X_i; Y)$, the mutual information between player i 's individual outcome X_i and the signal Y , and a measure of the “influence” of player i 's action on the distribution of the signal Y . In our setting, the most useful measure of the “influence” of i 's action is the χ^2 -divergence of $p(\cdot|\alpha_i, a_{-i})$ from $p(\cdot|a)$, which is defined as the variance (with respect to $p(\cdot|a)$) of the likelihood ratio difference $1 - p(y|\alpha_i, a_{-i})/p(y|a)$: that is,

$$\chi^2(p(\cdot|\alpha_i, a_{-i}) || p(\cdot|a)) := \sum_y \frac{(p(y|a) - p(y|\alpha_i, a_{-i}))^2}{p(y|a)}.$$

We show that, with full-support noise, χ^2 -divergence can be bounded in terms of mutual information.

Lemma 1 *If noise is $\underline{\pi}$ -bounded, then for any $a \in \mathcal{A}$, $i \in I$, and $\alpha_i \in \Delta(\mathcal{A}_i)$, we have*

$$\chi^2(p(\cdot|\alpha_i, a_{-i}) || p(\cdot|a)) \leq \kappa(\underline{\pi})^2 I^{\pi(\cdot|a)}(X_i; Y), \quad (1)$$

where $\kappa(\underline{\pi}) = \sqrt{2}(1 - 2\underline{\pi})/\underline{\pi}$.

The logic is that, since player i 's action affects the signal Y only through the outcome X_i , if a deviation from a_i to a'_i has a large effect on the distribution of Y , then Y must provide a large amount of information about X_i . Lemma 1 is related to standard *f-divergence*

inequalities (Sason and Verdú, 2016), which relate common measures of the difference between probability distributions such as the total variation distance, χ^2 -divergence, and KL-divergence, but it differs from standard results because of the $A \rightarrow X \rightarrow Y$ Markov chain structure of our model and the full-support noise assumption.

The point of bounding χ^2 -divergence in terms of mutual information is that mutual information also obeys a second inequality, which says that mutual information is sub-additive across players under independent noise.

Lemma 2 *For any action profile $a \in \mathcal{A}$, we have*

$$\sum_i I^{\pi(\cdot|a)}(X_i; Y) \leq I^{\pi(\cdot|a)}(X; Y) \leq C, \quad (2)$$

The first inequality in (2) is where we use the assumption that $(X_i)_{i \in I}$ are independent conditional on a . (The second inequality, $I^{\pi(\cdot|a)}(X; Y) \leq C$, is immediate from the definition of C .) The logic is that if $\sum_i I^{\pi(\cdot|a)}(X_i; Y) > I^{\pi(\cdot|a)}(X; Y)$ then there is some redundancy in the information that Y provides about the different outcomes X_i , which is impossible when $(X_i)_{i \in I}$ are conditionally independent. Note that inequality (2) can be strict: for example, if X_1 and X_2 are independent *Bernoulli* (1/2) variables and Y is the parity of their sum, then $I(X_1; Y) = I(X_2; Y) = 0$ but $I((X_1, X_2); Y) > 0$.

Combining Lemmas 1 and 2 (and dividing by N) yields the inequality

$$\frac{1}{N} \sum_i \chi^2(p(\cdot|a'_i, a_{-i}) || p(\cdot|a)) \leq \kappa(\underline{\pi})^2 \frac{C}{N}.$$

We thus obtain a bound for the *average* influence of a player's action on the signal distribution (measured by χ^2 -divergence) in terms of the *per-capita channel capacity* C/N . This will be a key step in the proof of Theorem 1.

3.2 An Impossibility Theorem for Large-Group Cooperation

We are now ready to state our first main result.

Theorem 1 *In any repeated game with N players, channel capacity C , $\underline{\pi}$ -bounded noise,*

and \bar{u} -bounded payoffs, every Nash equilibrium is ε -myopic, for

$$\varepsilon = \sqrt{\frac{\delta}{1-\delta} \frac{C}{N}} \kappa(\underline{\pi}) \bar{u}, \quad (3)$$

where $\kappa(\underline{\pi}) = \sqrt{2}(1 - 2\underline{\pi})/\underline{\pi}$.

In particular, for any $\underline{\pi} > 0$, $\bar{u} > 0$, and $\varepsilon > 0$; and any sequence of repeated games $(\Gamma)^k$ with $\underline{\pi}$ -bounded noise and \bar{u} -bounded payoffs such that $(1 - \delta)N/C \rightarrow \infty$ (where δ , N , and C depend on k); there exists \bar{k} such that, for every $k \geq \bar{k}$, all Nash equilibria in game Γ^k are ε -myopic.

Theorem 1 is a counterpoint to the folk theorem. While the folk theorem gives conditions under which the equilibrium payoff set is “large” in the limit where $\delta \rightarrow 1$ while the other parameters of the game are held fixed, Theorem 1 shows that supporting non-myopic payoffs requires not only that the discount rate $1 - \delta$ is small in absolute terms, but also that it is not much larger than the per-capita channel capacity C/N . This conclusion holds for any sequence of repeated games satisfying a uniform lower bound on noise and a uniform upper bound on the range of the stage game payoffs.

Theorem 1 can be compared to prior results by Fudenberg, Levine, and Pesendorfer (1998), al-Najjar and Smorodinsky (2000, 2001), and Pai, Roth, and Ullman (2014). These papers measure information by the number of possible signal realizations $|\mathcal{Y}|$ (rather than channel capacity), and establish impossibility results for cooperation when $N \rightarrow \infty$ for fixed δ and $|\mathcal{Y}|$. When N , δ , and $|\mathcal{Y}|$ vary together, arguments similar to the ones in these papers could be used to show that cooperation is impossible (i.e., all repeated game Nash equilibrium occupation measures are ε -myopic) if $(1 - \delta)^2 N / \log |\mathcal{Y}| \rightarrow \infty$.¹³ Compared to this result, Theorem 1 is qualitatively stronger in two ways: it replaces $\log |\mathcal{Y}|$ with $C \leq \log |\mathcal{Y}|$, and it replaces $(1 - \delta)^2$ with $1 - \delta$. The first of these improvements comes from the mutual information inequalities noted above. The second improvement comes from applying Theorem 1 of SW, which bounds the strength of equilibrium incentives in repeated games by a factor of $(1 - \delta)^{-1/2}$, rather than the naïve bound of $(1 - \delta)^{-1}$. This improved bound relies on focusing on incentives at the occupation measure, rather than considering

¹³Fudenberg, Levine, and Pesendorfer (1998) and al-Najjar and Smorodinsky (2001) focus on strategies that condition only on public signals, yielding the stronger conclusion that cooperation collapses to static ε -Nash equilibria, rather than “ ε -correlated equilibria” as in the present paper. A similar restriction would likewise let us strengthen the conclusion of Theorem 1 to a version of ε -Nash equilibrium.

incentives history-by-history as in the prior literature.¹⁴

We can also compare Theorem 1 to an immediate implication of Theorem 1 of SW, namely that Theorem 1 also holds when (3) is replaced by

$$\varepsilon = \frac{1}{N} \sum_i \sqrt{\frac{\delta}{1-\delta} \max_{a \in A} \chi^2(p(\cdot | s_i(a_i), a_{-i}) || p(\cdot | a))} \bar{u}. \quad (4)$$

This alternative result does not involve mutual information, and it also does not require a lower bound on noise. However, it is inadequate for the current paper’s objectives, for two reasons. First, we wish to bound incentives in terms of an single aggregate measure of societal information. This is achieved in (3)—where ε is bounded in terms of C —but not in (4)—where ε is bounded in terms of $(1/N) \sum_i \sqrt{\max_a \chi^2(p(\cdot | s_i(a_i), a_{-i}) || p(\cdot | a))}$, which is the average of N distinct, player-specific information measures. Second, we wish to generalize and strengthen the prior results of Fudenberg, Levine, and Pesendorfer (1998), al-Najjar and Smorodinsky (2000, 2001), and Pai, Roth, and Ullman (2014). We have already explained how this is achieved by the bound in (3). However, it is not achieved by the bound in (4), as absent noise there is no general relationship between $(1/N) \sum_i \sqrt{\max_a \chi^2(p(\cdot | s_i(a_i), a_{-i}) || p(\cdot | a))}$ and the cardinality of the signal space $|\mathcal{Y}|$.

While we have assumed that the signal y is publicly observed for simplicity, Theorem 1 also holds for repeated games with private monitoring. Indeed, the same result holds for the *blind repeated game*, where in each period the signal y is observed only by a mediator (rather than being directly observed by the players themselves), who then privately recommends actions to the players. Theorem 1 thus depends only on the precision of the signal y (measured by channel capacity), and not on how information about y is distributed among the players.¹⁵

¹⁴For this reason, while we improve the $(1-\delta)$ -order of the bounds in the literature for equilibrium payoffs and occupation measures, we do not improve the $(1-\delta)$ -order for bounds on payoffs or play in any single given period. An intuition for Theorem 1 of SW can be seen by considering a strategy where an agent’s performance is reviewed every T periods, with continuation play determined by the outcome of the review. For the agent to put weight independent of δ on the outcome of the review, a review must occur every $O((1-\delta)^{-1})$ periods. This implies that the standard deviation of the count of each signal realization over the course of the review is $O((1-\delta)^{-1/2})$, and hence the probability that a single signal is pivotal for the review is $O((1-\delta)^{1/2})$. Since the gain from deviating in a single period is $O(1-\delta)$, the agent’s “incentive strength” is $O((1-\delta)^{1/2} / (1-\delta)) = O((1-\delta)^{-1/2})$.

¹⁵See SW for more on blind games.

Proof of Theorem 1. Theorem 1 of SW implies that, for any Nash equilibrium outcome μ , any player i , and any manipulation s_i , we have

$$g_i(s_i, \alpha^\mu) \leq \sqrt{\frac{\delta}{1-\delta} \sum_a \alpha^\mu(a) \chi^2(p(\cdot|s_i(a_i), a_{-i}) || p(\cdot|a)) \bar{u}}.$$

Hence, by Lemma 1,

$$g_i(s_i, \alpha^\mu) \leq \sqrt{\frac{\delta}{1-\delta} \sum_a \alpha^\mu(a) I^{\pi(\cdot|a)}(X_i; Y) \kappa(\underline{\pi}) \bar{u}}.$$

Taking the maximum over manipulations s_i and averaging across players gives

$$\begin{aligned} \frac{1}{N} \sum_i \bar{g}_i(\alpha) &\leq \frac{1}{N} \sum_i \sqrt{\frac{\delta}{1-\delta} \sum_a \alpha^\mu(a) I^{\pi(\cdot|a)}(X_i; Y) \kappa(\underline{\pi}) \bar{u}} \\ &\leq \sqrt{\frac{1}{N} \sum_i \frac{\delta}{1-\delta} \sum_a \alpha^\mu(a) I^{\pi(\cdot|a)}(X_i; Y) \kappa(\underline{\pi}) \bar{u}} \\ &= \sqrt{\frac{\delta}{1-\delta} \frac{1}{N} \sum_a \alpha^\mu(a) \sum_i I^{\pi(\cdot|a)}(X_i; Y) \kappa(\underline{\pi}) \bar{u}} \\ &\leq \sqrt{\frac{\delta}{1-\delta} \frac{C}{N} \kappa(\underline{\pi}) \bar{u}}, \end{aligned}$$

where the second inequality is by Jensen and the third is by Lemma 2. This establishes Theorem 1. ■

In large groups, the necessary condition for cooperation implied by Theorem 1—that $(1-\delta)N/C$ is not too large—is easier to satisfy in some classes of repeated games than in others. For example, if the space of possible signal realizations \mathcal{Y} is fixed independent of N , then, since $C \leq \log |\mathcal{Y}|$, the necessary condition implies that $(1-\delta)^{-1}$ must be at least proportional to N , which is a restrictive condition in large groups. This negative conclusion applies for traditional applications of repeated games with public monitoring where the signal space is fixed independent of N , such as when the public signal is the market price facing Cournot competitors, the level of exploitation of a common-pool resource, the output of team production, or some other aggregate statistic.

However, in other settings C naturally scales linearly with N , so $(1-\delta)N/C$ is small whenever players are patient—regardless of group size. In repeated games with random

matching (Kandori, 1992; Ellison, 1994; Deb, Sugaya, and Wolitzky, 2020), players match in pairs each period, and each player observes her partner’s action. In these games, $C = N \log |\mathcal{A}_i|$, so per-capita channel capacity is independent of N . Intuitively, in a random matching game each player gets a distinct signal of the overall action profile, so the total amount of information available to society is proportional to N . Similarly, channel capacity scales linearly with N in public monitoring games where the public signal is a vector that includes a distinct signal of each player’s action, as in the ratings systems used by online platforms like eBay and AirBnB (Dellarocas, 2003; Tadelis, 2016). In general, C/N can be expected to be roughly independent of the population size in settings where players are monitored “separately,” rather than being monitored jointly through an aggregate statistic.

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

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

which is bounded by

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

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

Remark 2 *Theorem 1 also extends to games where noise is independent across “groups” of players, rather than individuals. For example, consider a repeated game with random matching, where the actions (a_i, a_j) of matched partners i and j generate an outcome $x_{i,j}$ with probability $\pi_{i,j}(x_{i,j}|a_i, a_j)$ satisfying $\min_{a_i, a_j, x_{i,j}} \pi_{i,j}(x_{i,j}|a_i, a_j) \geq \underline{\pi}$, independently across matches. Then Lemma 2 holds with $2C$ in place of C , because the sum $\sum_i I^{\pi(\cdot|a)}(X_{i,m(i)}; Y)$*

(where $m(i)$ denotes i 's partner) can be split into two sums of the mutual information of independent random variables, where each sum is bounded by C as in Lemma 2. Theorem 1 then likewise holds with $2C$ in place of C . More generally, if the players interact in disjoint groups of size K each period with independent noise across groups, then Lemma 2 and Theorem 1 hold with KC in place of C . Note that the $K = N$ case entails dropping the assumption of independent noise entirely, but then Theorem 1 holds with $\varepsilon = \sqrt{(\delta/(1-\delta)) C \kappa(\underline{\pi}) \bar{u}}$, which does not depend on N .

4 Sufficient Conditions for Cooperation

This section establishes a folk theorem for repeated games with public, product structure monitoring, where the discount factor, monitoring structure, and stage game (including the number of players N) vary simultaneously. The theorem will imply that the relationship between N , δ , and C in Theorem 1 is tight up to a $\log(N)$ factor.

Our folk theorem allows independent noise, but does not require it: this section does not require that $\min_i \underline{\pi}_i > 0$. We do however require that monitoring has a *product structure*: there exist sets $(\mathcal{Y}_i)_{i \in I}$ and a family of conditional distributions $(q_i(y_i|x_i))_{i,y_i,x_i}$ such that $\mathcal{Y} = \prod_i \mathcal{Y}_i$ 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. Note that if (\mathcal{Y}, q) has a product structure, then so does the action monitoring structure (\mathcal{Y}, p) , meaning that there exists a family of conditional distributions $(p_i(y_i|a_i))_{i,y_i,a_i}$ (given by $p_i(y_i|a_i) = \sum_{x_i} \pi_i(x_i|a_i) q_i(y_i|x_i)$) such that $p(y|a) = \prod_i p_i(y_i|a_i)$ for all y, a .

We also need an identification condition. For any $\eta \in (0, 1)$, we say that the action monitoring structure (\mathcal{Y}, p) satisfies *η -individual identifiability* if

$$\sum_{y_i: p_i(y_i|a_i) \geq \eta} \frac{(p_i(y_i|a_i) - p_i(y_i|\alpha_i))^2}{p_i(y_i|a_i)} \geq \eta \quad \text{for all } i \in I, a_i \in \mathcal{A}_i, \alpha_i \in \Delta(\mathcal{A}_i \setminus \{a_i\}). \quad (5)$$

This condition is a variant of FLM's individual full rank condition and Kandori and Matsushima's (1998) assumption (A2''). It says that the influence on the signal distribution (measured by χ^2 -divergence) of a deviation from a_i to any mixed action α_i supported on $\mathcal{A}_i \setminus \{a_i\}$ is at least η , ignoring signals that occur with probability less than η under a_i . Intuitively, this requires that deviations from a_i are sufficiently detectable, and that in addition

detection does not rest on very rare signal realizations. This assumption will ensure that players can be motivated by rewards whose variance and maximum absolute value are both of order $(1 - \delta)/\eta$.¹⁶

We establish the following folk theorem. (Recall that $B(\varepsilon)$ is the “target” set of payoff vectors defined in Section 2, and E is the set of PPE payoff vectors.)

Theorem 2 *For any $\bar{u} > 0$ and $\varepsilon > 0$; and any sequence of repeated games $(\Gamma)^k$ with \bar{u} -bounded payoffs and product structure monitoring satisfying η -individual identifiability such that $(1 - \delta) \log(N)/\eta \rightarrow 0$ (where δ , N , and η depend on k); there exists \bar{k} such that, for every $k \geq \bar{k}$, we have $B(\varepsilon) \subseteq E$.*

To see why Theorem 2 implies that Theorem 1 is tight up to a $\log(N)$ factor, consider a game where $\mathcal{X} = \mathcal{A}$ with uniform noise, so that $\pi_i(a'_i|a_i) = \underline{\pi}$ for all $i, a_i, a'_i \neq a_i$, and assume that $\underline{\pi} < (\max_i |\mathcal{A}_i| + 1)^{-1}$. Suppose that the outcome monitoring structure (\mathcal{Y}, q) is given by η -random auditing, where in every period the public signal perfectly reveals each player’s identity and realized individual outcome with probability η . That is, under η -random auditing, $\mathcal{Y}_i = \mathcal{X}_i \cup \{\emptyset\}$ for all i , and

$$q_i(y_i|x_i) = \begin{cases} \eta & \text{if } y_i = x_i, \\ 0 & \text{if } y_i \in \mathcal{X}_i \setminus \{x_i\}, \\ 1 - \eta & \text{if } y_i = \emptyset, \end{cases} \quad \text{so that} \quad p_i(y_i|a_i) = \begin{cases} \eta\pi_i(y_i|a_i) & \text{if } y_i \in \mathcal{X}_i, \\ 1 - \eta & \text{if } y_i = \emptyset. \end{cases}$$

Note that the channel capacity under η -random auditing is at most $\eta N \log(\max_i |\mathcal{A}_i|)$. In addition, η -random auditing satisfies $\eta\underline{\pi}$ -individual identifiability, because, any i, a_i, α_i , we have

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

¹⁶If (5) were relaxed by taking the sum over all y_i (rather than only y_i such that $p_i(y_i|a_i) \geq \eta$), player i could be motivated by rewards with variance $O((1 - \delta)/\eta)$, but not necessarily with maximum absolute value $O((1 - \delta)/\eta)$. Our analysis requires controlling both the variance and absolute value of players’ rewards, so we need the stronger condition. We also note that the current definition of η -individual identifiability coincides with $\sqrt{\eta}$ -individual identifiability in the terminology in SW.

where the last inequality uses $\underline{\pi} < (\max_i |\mathcal{A}_i| + 1)^{-1}$. Thus, by Theorem 2, η -random auditing is a monitoring structure with channel capacity at most $C = \eta N \log(\max_i |\mathcal{A}_i|)$, under which a folk theorem holds whenever $(1 - \delta) N \log(N) / C \rightarrow 0$. Therefore, Theorem 1’s conclusion that play is ε -myopic if $(1 - \delta) N / C \rightarrow \infty$ can be improved by at most a $\log(N)$ factor.¹⁷

The assumption that payoffs are uniformly bounded plays a different role in Theorems 1 and 2. Theorem 1 requires bounded payoffs to bound the variation in players’ continuation payoffs. Theorem 2 requires bounded payoffs to bound players’ one-shot deviation gains. For example, the conclusion of Theorem 2 does not hold for repeated Bertrand competition where the size of the market (and hence the gain from undercutting one’s rivals to win the entire market) is proportional to the number of firms.¹⁸

We now discuss the proof of Theorem 2. Theorem 2 is a folk theorem for repeated games with public monitoring.¹⁹ The standard proof, following FLM and Kandori and Matsushima (1998), relies on continuation payoffs transfers along hyperplanes tangent to the boundary of the PPE payoff set. Unfortunately, this approach encounters difficulties when N and δ vary simultaneously. The problem is that when N is large, changing each player’s continuation payoff by a small amount can result in a large overall movement in the continuation payoff vector. Mathematically, FLM’s proof relies on the equivalence of the L^1 norm and the Euclidean norm in \mathbb{R}^N . Since this equivalence is not uniform in N , their proof does not apply when N and δ vary simultaneously.²⁰

Our proof of Theorem 2 is instead based on the “block strategy” approach introduced by Matsushima (2004) and Hörner and Olszewski (2006) in the context of repeated games

¹⁷Note that Theorem 1 holds verbatim if C is taken to be an upper bound for channel capacity rather than its exact value, because the theorem’s conclusion is stronger when C is smaller.

¹⁸For repeated Bertrand competition with a fixed market size, Theorem 2 holds vacuously as $B(\varepsilon) = \emptyset$ for sufficiently large N . In contrast, the public goods game in Appendix B is an example with uniformly bounded payoffs where $B(\varepsilon)$ is “large” for all N .

¹⁹Specifically, it is a “Nash threat” folk theorem, as F^* is the set of payoffs that Pareto-dominate a convex combination of static Nash equilibria. To extend this result to a “minmax threat” theorem, players must be made indifferent among all actions in the support of a mixed strategy that minmaxes an opponent. This requires a stronger identifiability condition, similar to Kandori and Matsushima’s assumption (A1).

²⁰To see the problem in more detail, η -individual identifiability implies that the movement in each player’s per-period continuation payoff required to provide incentives is of order $(1 - \delta) / \eta$, so the movement of the continuation payoff vector in the L^2 norm is $O(\sqrt{N}(1 - \delta) / \eta)$. Fix a ball B contained in V^* , 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 . Since player 1’s continuation payoff must be within $O(1 - \delta)$ distance of v , the greatest movement along a tangent hyperplane is $O(\sqrt{1 - \delta})$. FLM’s proof approach thus requires $\sqrt{N}(1 - \delta) / \eta \ll \sqrt{1 - \delta}$, or $(1 - \delta) N / \eta^2 \ll 1$, while we assume only $(1 - \delta) \log(N) / \eta \ll 1$. Thus, while the conditions for Theorem 2 are tight up to $\log(N)$ slack, FLM’s approach requires slack N .

with private monitoring. We view the repeated game as a sequence of T -period blocks of periods, where T is a number proportional to $(1 - \delta)^{-1}$. 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 future rewards 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 rewards. By η -individual identifiability, incentives can be provided with per-period rewards of maximum size $O(\eta^{-1})$, and the rewards can be normalized to have zero mean. Therefore, by the law of large numbers, when $T \gg \eta^{-1}$, with high probability the total reward that a 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 reward over a block, so that at some point the target payoff for the next block cannot be further incremented. In this case, the player can no longer be motivated to take a non-myopic best response, and all players' behavior in the current block must change. Thus, if any player's reward is "abnormal," all players' payoffs in that block may be far from the target equilibrium payoffs.

To prove the theorem, we must ensure that abnormal rewards do not compromise either ex ante efficiency or the players' incentives. Efficiency is preserved if the blocks length T is large enough that the probability that *any* player's total reward is abnormal is small. Since the per-period rewards have size $O(\eta^{-1})$ and the length of a block is $O((1 - \delta)^{-1})$, standard concentration bounds imply that the probability that a given player's total reward is abnormal is $\exp(-O(\eta/(1 - \delta)))$. Hence, by the union bound, the probability that any player's total reward is abnormal is at most $N \exp(-O(\eta/(1 - \delta)))$, which converges to 0 when $(1 - \delta) \log(N)/\eta \rightarrow 0$. This step in the proof accounts for the $\log(N)$ slack.

Finally, since all players' payoffs are affected whenever any player's reward becomes abnormal, incentives would be threatened if one player's action affected the probability that another player's reward becomes abnormal. We avoid this problem by letting each player's reward depend only on the signals of her own actions. This separation of rewards 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.²¹

5 Co-Linear Equilibria

We now consider a restricted class of equilibria—*co-linear equilibrium*—which model collective incentive-provision in repeated games. We will show that cooperation is possible in this class of equilibria only if the discount rate is exponentially small relative to the population size. We view this as an impossibility theorem for cooperation under collective incentives for any “reasonably large” group. More colorfully, the result can be seen as a formalization of Hume’s intuition that large groups cannot support cooperation by threatening “the abandoning of the whole project.”

Formally, a *co-linear equilibrium* is a PPE where the players’ continuation payoffs at all public histories are co-linear: for each player $i \neq 1$, there exists a constant $b_i \in \mathbb{R}$ such that, for all public histories h, h' , we have $w_i(h') - w_i(h) = b_i(w_1(h') - w_1(h))$, where $w_i(h)$ denotes player i ’s equilibrium continuation payoff at history h . Relabeling the players if necessary, we can take $|b_i| \leq 1$ for all i without loss. Note that if $b_i \geq 0$ for all i then the players’ preferences over histories are all aligned; while if $b_i < 0$ for some i then the players can be divided into two groups with opposing preferences. Note also that the notion of co-linear equilibrium generalizes *strongly symmetric equilibrium* (SSE) in symmetric games, where $b_i = 1$ for all i .²²

Our result for co-linear equilibria is as follows.

Theorem 3 *For any $\underline{\pi} > 0$, $\bar{u} > 0$, $\varepsilon > 0$, and $\rho > 0$; and any sequence of repeated games (Γ^k) with $\underline{\pi}$ -bounded noise and \bar{u} -bounded payoffs such that $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$ (where δ and N depend on k); there exists $\bar{k} > 0$ such that, for every $k > \bar{k}$, all co-linear equilibria in game Γ^k are ε -myopic.*

Theorem 3 shows that collective incentives are ineffective unless the discount rate is exponentially small relative to N . Comparing Theorem 2 with Theorem 3, we see that

²¹As noted above, we conjecture that the approach of FLM and Kandori and Matsushima yields a folk theorem if $(1 - \delta) N/\eta^2 \rightarrow 0$. Their approach requires only pairwise identifiability rather than a product structure, so we conjecture that the product structure can be relaxed to pairwise identifiability if $(1 - \delta) N/\eta^2 \rightarrow 0$. We do not know if such a relaxation is possible under the weaker hypothesis of Theorem 2.

²²Thus, co-linear equilibrium generalizes SSE in two ways: to asymmetric games, and to an arbitrary linear relationship among the players’ continuation payoffs.

targeted incentives are much more effective than collective ones when the discount rate is not exponentially small. Notably, this result holds even when information is scarce (e.g., C is relatively small), so that precisely monitoring all players is infeasible.

Theorem 3 differs from Theorem 1 in the required relationship between N and δ , and also in that Theorem 3 holds for any outcome monitoring precision, so the channel capacity C does not show up in the statement. The intuition for why the theorem does not involve C is that optimal co-linear equilibria take a bang-bang form even when the realized outcome profile is perfectly observed, so a binary signal that indicates which of two extreme continuation payoff vectors should be implemented is as effective as any more informative signal.

It is well-known that, with noisy monitoring, SSE are inefficient for any discount factor (e.g., Mailath and Samuelson, 2006, Proposition 8.2.1). In contrast, Theorem 3 shows that the relationship between N and δ required to provide *any* non-trivial equilibrium incentives is dramatically different between SSE (or, more generally, co-linear equilibria) and arbitrary equilibria.

To see the logic of Theorem 3, consider the case where the game is symmetric and $b_i = 1$ for all i , so co-linear equilibria are SSE. Suppose also that $\mathcal{X} = \mathcal{A}$ with binary actions and symmetric noise, so that $|\mathcal{A}_i| = 2$ and $\pi(a_i|a_i) = 1 - \underline{\pi}$, $\pi(a'_i|a_i) = \underline{\pi}$ for each $a_i \neq a'_i$. Finally, suppose we wish to enforce a symmetric pure action profile $\vec{a}_0 = (a_0, \dots, a_0)$, where $\bar{g}_i(\vec{a}_0) = \nu$. By standard arguments, it can be shown that 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^* .²³ Due to independent noise, when N is large the distribution of n is approximately normal, with mean $(1 - \underline{\pi})N$ and standard deviation $\sqrt{\underline{\pi}(1 - \underline{\pi})N}$. Now, 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 $\tau \in [0, \bar{u} / (1 - \delta)]$ denote the size of the penalty when the tail test is failed. We then must have

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

where the first inequality is the incentive compatibility condition that the product of the pivot probability $\phi(z^*) / \sqrt{\underline{\pi}(1 - \underline{\pi})N}$ and the penalty size τ must exceed the gain from

²³The analysis of tail tests as optimal incentive contracts under normal noise goes back to Mirrlees (1975). The logic of Theorem 3 shows that the size of the penalty in a Mirrleesian tail test must increase exponentially with the variance of the noise.

deviating ν , and the second inequality is the promise-keeping condition that the expected penalty cannot exceed the stage-game payoff range. Dividing the first inequality by the second, we obtain

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

The left-hand side of this inequality is the *Mills ratio* of the standard normal distribution, which is approximately equal to $|z^*|$ when $z^* < 0$. Hence, to satisfying incentive compatibility and promise-keeping, $|z^*|$ 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 $\phi(z^*)/\sqrt{\pi(1-\pi)N}$ and $\bar{u}/(1-\delta)$ (the upper bound for τ) must exceed ν .

Intuitively, the weakness of co-linear equilibrium is that the probability that a single player's action is pivotal for a tail test is of order $\phi(z^*)/\sqrt{N}$, while the probability that the test is failed is $\Phi(z^*)$, and the former is much smaller than the latter unless z^* is much less than zero, which in turn is consistent with equilibrium incentives only if $(1-\delta)^{-1}$ is exponentially large.

We also note a converse to Theorem 3: if π_{a_i, a_i} is sufficiently large for each a_i and $(1-\delta)\exp(N^{1+\rho}) \rightarrow 0$ for some $\rho > 0$, then a folk theorem holds for co-linear equilibria. Intuitively, a target action profile a can now be enforced by a tail test where the players are all punished only if $x_i \neq a_i$ for *every* player i .

Theorem 3 is related to Proposition 1 of Sannikov and Skrzypacz (2007), which is an impossibility theorem for SSE in a two-player repeated game where actions are observed with additive, normally distributed noise, with variance proportional to $(1-\delta)^{-1}$.²⁴ As a tail test is optimal in their setting, the proof of Theorem 3 implies that non-vanishing incentives can be provided only if $(1-\delta)^{-1}$ increases exponentially with the variance of the noise. Since in their model $(1-\delta)^{-1}$ increases with variance only linearly, they likewise obtain an impossibility result. Similarly, Proposition 2 of Fudenberg and Levine (2007) is an impossibility theorem in a game with one patient player and a myopic opponent, where the patient player's action is observed with additive, normal noise, with variance proportional to $(1-\delta)^{-\rho}$ for some $\rho > 0$; and their Proposition 3 is a folk theorem when the variance

²⁴The 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 increments of a Brownian process.

is constant in δ . Theorem 3 suggest that their impossibility theorem extends whenever variance asymptotically dominates $(-\log(1-\delta))^{1/(1-\rho)}$ for some $\rho > 0$, while their folk theorem extends whenever variance is asymptotically dominated by $(-\log(1-\delta))^{1/(1+\rho)}$ for some $\rho > 0$.

6 Conclusion

This paper has developed a theory of large-group cooperation in repeated games. Our key assumption is that monitoring is imperfect and respects a degree of independence across players. Our main results establish necessary and (somewhat stronger) sufficient conditions for cooperation in terms of the number of players, the discount factor, and the *per-capita channel capacity* of the monitoring structure. We also show that cooperation in a *co-linear equilibrium*, where the players' rewards are co-linear, is possible only under much more stringent conditions. This result demonstrates a sense in which large-group cooperation must rely on targeted sanctions. Notably, this result holds even when information is scarce, so that precisely monitoring all players is infeasible.

Our results raise several questions for future theoretical and applied research. On the theory side, this paper has focused on insufficient monitoring precision as an obstacle to large-group cooperation. In reality, noisy monitoring coexists with other obstacles to cooperation, such as decentralized monitoring (as in community enforcement models) and the possibility that a small fraction of players may be irrational or fail to understand the equilibrium being played (as in, e.g., Sugaya and Wolitzky 2020, 2021). Combining these features may lead to a richer and more realistic perspective on the determinants of large-group cooperation. We also believe it could be interesting to explore the implications of independent noise and limited monitoring precision for organizational design, for example the design of managerial hierarchies. Finally, another open question is whether some version of our results survives under an appropriate relaxation of independent noise.

As for applied work, more systematic empirical or experimental evidence on the determinants of large-group cooperation under imperfect monitoring would be valuable.²⁵ For

²⁵Camera and Casari (2009) and Duffy and Ochs (2009), among others, run experiments on repeated games with random matching and private monitoring, i.e., community enforcement. Community enforcement raises additional issues beyond the ones we focus on, which arise even under public monitoring. Camera, Casari, and Bigoni (2013) include a treatment with public monitoring, and find that larger groups cooperate less.

example, a novel prediction of our paper is that targeted sanctions are much more effective than collective ones in large groups, even when the total amount of available information about agents' performance is small. It would be interesting to test this prediction.

Appendix

A Comparison of $V(\varepsilon)$ and ε -Correlated Equilibria

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

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

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

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

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

For an example where $V(\varepsilon)$ 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

$$\mathbf{1}\{a_i = B\}(-1 + 2 \times \mathbf{1}\{a_1 = H\}),$$

while the seller’s payoff is given by

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

where $k \in \{0, 1, \dots, N - 1\}$ is the number of buyers who take B . Suppose also that the players tremble with independent, uniform noise $\underline{\pi} \in (0, 1/3)$. Note that in this game the payoff range is bounded by 3 and noise is bounded by $\underline{\pi}$.

In this game, for any $\varepsilon > 0$, when N is sufficiently large, we have $(H, B, \dots, B) \in \mathcal{A}(\varepsilon)$, and hence $(1, 1, \dots, 1) \in V(\varepsilon)$. This follows because the per-player average deviation gain at action profile (H, B, \dots, B) equals $1/N$: the seller has a deviation gain of 1, while each buyer has a deviation gain of 0. Thus, Theorem 1 does not preclude $(1, 1, \dots, 1)$ (or any convex combination of $(1, 1, \dots, 1)$ and $(0, 0, \dots, 0)$) as an equilibrium payoff vector, even when $(1 - \delta)N/C$ is very large. This is reassuring, because the monitoring structure given by perfect monitoring of the seller's realized action (i.e., $\mathcal{Y} = \{H, L\}$, $q(y|x) = \mathbf{1}\{y = x_1\}$) has channel capacity $\log 2$ and supports the payoff vector $((1 - 3\underline{\pi}) / (1 - 2\underline{\pi}), \dots, (1 - 3\underline{\pi}) / (1 - 2\underline{\pi}))$ for all $\delta \geq 1 / (2 - 3\underline{\pi})$ and all $N \geq 2$.²⁶ In contrast, the greatest symmetric payoff vector in $CE(\varepsilon)$ is $(\varepsilon, \varepsilon, \dots, \varepsilon)$, because the seller's deviation gain equals the probability that she takes H .

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

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

In games with bounded per-capita externalities, any level of per-capita utilitarian welfare that is attainable in $V(\varepsilon)$ can also be approximated in $CE(\sqrt{8d\varepsilon})$.

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

²⁶This is a standard calculation, which results from considering “tolerant 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 $v = (1 - 3\underline{\pi}) / (1 - 2\underline{\pi})$.

$(0, 2d)$ and any $v \in V(\varepsilon)$, there exists $v' \in CE(\sqrt{8d\varepsilon})$ such that

$$\left| \frac{1}{N} \sum_i (v_i - v'_i) \right| \leq \sqrt{\frac{2\varepsilon}{d}} \bar{u}.$$

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

B 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 $v = (v, \dots, v) \in \mathbb{R}^N$, and let $\varepsilon = cv(1 - c - v)/4 > 0$. We show that $B_v(\varepsilon) \subseteq F$ for all N . Since no one contributing is a Nash equilibrium with 0 payoffs, this implies that $B_v(\varepsilon) \subseteq F^*$, and hence $v \in B(\varepsilon)$, for all N .

Fix any N . Since the game is symmetric, to show that $B_v(\varepsilon) \subseteq F$ it suffices to show that, for any number $n \in \{0, \dots, N\}$, there exists a feasible payoff vector where n “favored” players receive payoffs no less than $v + \varepsilon$, and the remaining $N - n$ “disfavored” players receive payoffs no more than $v - \varepsilon$. First, consider the mixed action profile α^1 where favored

players contribute with probability $\frac{v+\varepsilon}{1-c}$ and disfavored players always contribute. At this profile, favored players receive payoff $f(n) := \frac{n}{N} \frac{v+\varepsilon}{1-c} + (1 - \frac{n}{N})(1) - c \frac{v+\varepsilon}{1-c}$, while disfavored players receive payoff $g(n) := \frac{n}{N} \frac{v+\varepsilon}{1-c} + (1 - \frac{n}{N})(1) - c$. Now consider the mixed action profile α^2 where favored players contribute with probability $\frac{(v+\varepsilon)^2}{(1-c)f(n)}$ and disfavored players contribute with probability $\frac{v+\varepsilon}{f(n)}$. Note that each player's payoff at profile α^2 equals her payoff at profile α^1 multiplied by $\frac{v+\varepsilon}{f(n)}$. Therefore, at profile α^2 , favored players receive payoff $f(n) \frac{v+\varepsilon}{f(n)} = v + \varepsilon$, while disfavored players receive payoff

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

C Proof of Lemma 1

Since χ^2 -divergence is a convex function, it suffices to consider the case where α_i is degenerate on some $a'_i \in \mathcal{A}_i$. For any $x_i \in \mathcal{X}_i$ and $y \in \mathcal{Y}$, we have $\varphi^a(x_i, y) = \pi_i(x_i|a_i) \varphi^a(y|x_i) = p(y|a) \varphi^a(x_i|y)$. Hence, since $\pi_i(x_i|a_i) \geq \underline{\pi}$, we have

$$(\varphi^a(y|x_i) - p(y|a))^2 = \left(\frac{p(y|a)}{\pi_i(x_i|a_i)} (\varphi^a(x_i|y) - \pi_i(x_i|a_i)) \right)^2 \leq \left(\frac{p(y|a)}{\underline{\pi}} (\varphi^a(x_i|y) - \pi_i(x_i|a_i)) \right)^2. \quad (6)$$

Now, for any a , i , and a'_i , we have

$$\begin{aligned}
\chi^2(p(\cdot|a'_i, a_{-i}) || p(\cdot|a)) &= \sum_y \frac{(\sum_{x_i} (\pi_i(x_i|a_i) - \pi_i(x_i|a'_i)) \varphi^a(y|x_i))^2}{p(y|a)} \\
&= \sum_y \frac{(\sum_{x_i} (\pi_i(x_i|a_i) - \pi_i(x_i|a'_i)) (\varphi^a(y|x_i) - p(y|a)))^2}{p(y|a)} \\
&\leq \sum_{x_i} (\pi_i(x_i|a_i) - \pi_i(x_i|a'_i))^2 \sum_y \frac{\sum_{x_i} (\varphi^a(y|x_i) - p(y|a))^2}{p(y|a)} \\
&\leq \frac{2(1-2\underline{\pi})^2}{\underline{\pi}^2} \sum_y p(y|a) \sum_{x_i} (\varphi^a(x_i|y) - \pi_i(x_i|a_i))^2 \\
&\leq \frac{(1-2\underline{\pi})^2}{\underline{\pi}^2} \sum_y p(y|a) \left(\sum_{x_i} |\varphi^a(x_i|y) - \pi_i(x_i|a_i)| \right)^2 \\
&\leq \frac{2(1-2\underline{\pi})^2}{\underline{\pi}^2} \sum_y p(y|a) \sum_{x_i} \varphi^a(x_i|y) \log \left(\frac{\varphi^a(x_i|y)}{\pi_i(x_i|a_i)} \right) \\
&= \kappa(\underline{\pi})^2 I^{\pi(\cdot|a)}(X_i; Y),
\end{aligned}$$

where the first inequality follows by Cauchy-Schwarz; the second follows by (6) and

$$\begin{aligned}
\sum_{x_i} (\pi_i(x_i|a_i) - \pi_i(x_i|a'_i))^2 &\leq (\pi_i(\mathcal{X}^+|a_i) - \pi_i(\mathcal{X}^+|a'_i))^2 \\
&\quad + (\pi_i(\mathcal{X} \setminus \mathcal{X}^+|a_i) - \pi_i(\mathcal{X} \setminus \mathcal{X}^+|a'_i))^2 \leq 2(1-2\underline{\pi})^2;
\end{aligned}$$

the third follows by the $L^1 - L^2$ norm inequality (using $\sum_{x_i} \varphi^a(x_i|y) - \pi_i(x_i|a_i) = 0$); and the fourth follows by *Pinsker's inequality* (CT, Lemma 11.1.1), which states that for any two probability distributions ζ and ζ' on a finite set \mathcal{Z} , we have $(\sum_z |\zeta(z) - \zeta'(z)|)^2 \leq 2 \sum_z \zeta(z) \log(\zeta(z)/\zeta'(z))$.

D Proof of Lemma 2

We recall some basic concepts from information theory (see, e.g., Chapter 2 of Cover and Thomas, 2006). For any discrete random variable Z with distribution ζ , its *entropy* is $H(Z) = -\sum_z \zeta(z) \log \zeta(z)$. For any pair of discrete random variables (Z, Z') with joint

distribution ζ , the mutual information $I(Z; Z')$ satisfies

$$I(Z; Z') = \sum_{z, z'} \zeta(z, z') \log \left(\frac{\zeta(z, z')}{\zeta(z)\zeta(z')} \right) = H(Z) - H(Z|Z'),$$

where the *conditional entropy* $H(Z|Z')$ is $H(Z|Z') = -\sum_{z, z'} \zeta(z, z') \log \zeta(z|z')$. We also recall the *independence bound on entropy* (Cover and Thomas, Theorem 2.6.6): if $Z = (Z_1, \dots, Z_N)$ then $H(Z) \leq \sum_i H(Z_i)$, with equality if and only if the Z_i are independent.

We now prove inequality (2). Suppressing the superscript $\pi(\cdot|a)$, we have

$$\begin{aligned} \sum_i I(X_i; Y) &= \sum_i (H(X_i) - H(X_i|Y)) \\ &= \sum_i H(X_i) - \sum_i H(X_i|Y) \leq H(X) - H(X|Y) = I(X; Y), \end{aligned}$$

where the inequality follows because, by the independence bound on entropy and independence of the X_i , we have $H(X) = \sum_i H(X_i)$ and $H(X|Y) \leq \sum_i H(X_i|Y)$. Finally, $I(X; Y) \leq C$ by definition of channel capacity.

E Proof of Theorem 2

E.1 Preliminaries

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

Lemma 3 *Under η -individual identifiability, for each $i \in I$ and $r_i \in \mathcal{A}_i$ there exists a function $f_{i, r_i} : \mathcal{Y}_i \rightarrow \mathbb{R}$ such that*

$$\mathbb{E}[f_{i, r_i}(y_i) | r_i] - \mathbb{E}[f_{i, r_i}(y_i) | a_i] \geq \bar{u} \quad \text{for all } a_i \neq r_i, \quad (7)$$

$$\mathbb{E}[f_{i, r_i}(y_i) | r_i] = 0, \quad (8)$$

$$\text{Var}(f_{i, r_i}(y_i) | r_i) \leq \bar{u}^2/\eta, \quad \text{and} \quad (9)$$

$$|f_{i, r_i}(y_i)| \leq 2\bar{u}/\eta \quad \text{for all } y_i. \quad (10)$$

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

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

Note that (5) is equivalent to $d(p_i(r_i), \text{co}(P_i(r_i))) \geq \sqrt{\eta}$ for all $i \in I, r_i \in \mathcal{A}_i$, where $d(\cdot, \cdot)$ denotes Euclidean distance in $\mathbb{R}^{|\mathcal{Y}_i^*|}$. Hence, by the separating hyperplane theorem, there exists $x = (x(y_i))_{y_i \in \mathcal{Y}_i^*} \in \mathbb{R}^{|\mathcal{Y}_i^*|}$ such that $\|x\| = 1$ and $(p_i(r_i) - p) \cdot x \geq \sqrt{\eta}$ for all $p \in P_i(r_i)$. By definition of p_i and P_i , this implies that $\sum_{y_i \in \mathcal{Y}_i^*} \left(\frac{p_i(y_i|r_i) - p_i(y_i|a_i)}{\sqrt{p_i(y_i|r_i)}} \right) x(y_i) \geq \sqrt{\eta}$ for all $a_i \neq r_i$. Now define

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

Clearly, conditions (7) and (8) hold. Moreover, since $\mathbb{E}[f_{i,r_i}(y_i)|r_i] = 0$ and the term $\sum_{\tilde{y}_i \in \mathcal{Y}_i^*} \sqrt{p_i(\tilde{y}_i|r_i)} x(\tilde{y}_i)$ is independent of y_i , we have

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

and hence (9) holds. Finally, (10) holds since, for each $y_i \in \mathcal{Y}_i^*$,

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

■

Now fix $i \in I$ and $r_i \in \mathcal{A}_i$, and suppose that $y_{i,t} \sim p_i(\cdot|r_i)$ for each period $t \in \mathbb{N}$, independently across periods (which would be the case in the repeated game if r_i were taken in every period). By (9), for any $T \in \mathbb{N}$, we have

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

Together with (8) and (10), Bernstein's inequality now implies that, for any $T \in \mathbb{N}$ and $\bar{f} \in \mathbb{R}_+$, we have

$$\Pr \left(\sum_{t=1}^T \delta^{t-1} f_{i,r_i}(y_{i,t}) \geq \bar{f} \right) \leq \exp \left(- \frac{\bar{f}^2 \eta}{2 \left(\frac{1-\delta^{2T}}{1-\delta^2} \bar{u}^2 + \frac{2}{3} \bar{f} \bar{u} \right)} \right). \quad (11)$$

Our second lemma fixes T and \bar{f} so that the bound in (11) is sufficiently small, and some other conditions used in the proof also hold.

Lemma 4 *There exists $\varkappa > 0$ such that, whenever $(1 - \delta) \log(N) / \eta < \varkappa$, there exist $T \in \mathbb{N}$ and $\bar{f} \in \mathbb{R}$ that satisfy the following three inequalities:*

$$60\bar{u}N \exp \left(- \frac{\left(\frac{\bar{f}}{3}\right)^2 \eta}{2 \left(\frac{1-\delta^{2T}}{1-\delta^2} \bar{u}^2 + \frac{2}{3} \frac{\bar{f}}{3} \bar{u} \right)} \right) \leq \varepsilon, \quad (12)$$

$$8 \frac{1-\delta}{1-\delta^T} \left(\bar{f} + \frac{2\bar{u}}{\eta} \right) \leq \varepsilon, \quad (13)$$

$$4\bar{u} \frac{1-\delta^T}{\delta^T} + \frac{1-\delta}{\delta^T} \left(\bar{f} + \frac{2\bar{u}}{\eta} \right) \leq \varepsilon. \quad (14)$$

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

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

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

$$\frac{4}{9} \sqrt{36 \left(\log \left(\frac{60\bar{u}}{\varepsilon} \right) + \log(N) \right) \frac{1-\delta}{1-\delta^T} \frac{1}{\eta}} \leq 1 \quad \text{and} \quad (15)$$

$$8\bar{u} \left(\sqrt{36 \left(\log \left(\frac{60\bar{u}}{\varepsilon} \right) + \log(N) \right) \frac{1-\delta}{1-\delta^T} \frac{1}{\eta}} + \frac{1-\delta}{1-\delta^T} \frac{2}{\eta} \right) \leq \varepsilon. \quad (16)$$

It now follows from straightforward algebra (provided in Appendix E.4) that (12)–(14) hold.

■

E.2 Equilibrium Construction

Fix any T and \bar{f} that satisfy (12)–(14), 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 $v \in E(\Gamma)$.²⁷ 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 $w : H^{T+1} \rightarrow \mathbb{R}^N$ such that, letting $\psi_i(h^{T+1}) = \frac{\delta^T}{1-\delta} (w_i(h^{T+1}) - v_i^*)$, we have

$$\text{Promise Keeping: } v_i^* = \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \psi_i(h^{T+1}) \right] \quad \text{for all } i, \quad (17)$$

$$\text{Incentive Compatibility: } \sigma_i \in \operatorname{argmax}_{\tilde{\sigma}_i} \mathbb{E}^{\tilde{\sigma}_i, \sigma_{-i}} \left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \psi_i(h^{T+1}) \right] \quad \text{for all } i, \quad (18)$$

$$\text{Self Generation: } \zeta_i \psi_i(h^{T+1}) \in \left[-\frac{\delta^T}{1-\delta} \varepsilon, 0 \right] \quad \text{for all } i \text{ and } h^{T+1}. \quad (19)$$

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

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

$$u_i(\bar{\alpha}) = v_i^* + \zeta_i \varepsilon / 2 \quad \text{for all } i, \quad (20)$$

and fix a probability distribution over static Nash equilibria $\alpha^{NE} \in \Delta(\prod_i \Delta(\mathcal{A}_i))$ such that $u_i(\alpha^{NE}) \leq v_i^* - \varepsilon/2$ 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 $t \in \{1, \dots, T\}$, we define a *state* $\theta_{i,t} \in \{0, 1\}$ for player i in period t . The states are determined by the public history, and so are common knowledge among the players. We first specify players' prescribed actions as a function of the state, and then specify the state as a function

²⁷Specifically, 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 $w(h^{T+1})$.

of the public history.

Prescribed Equilibrium Actions: For each period t , let $r_t \in \mathcal{A}$ be a pure action profile which is drawn by public randomization at the start of period t from the distribution $\bar{\alpha} \in \Delta(\mathcal{A})$ fixed in (20), and let $\varrho_t^{NE} \in \prod_i \Delta(\mathcal{A}_i)$ be a mixed action profile which is drawn by public randomization at the start of period t from the distribution α^{NE} . The prescribed equilibrium actions are defined as follows.

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

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

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

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

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

States: At the start of each period t , conditional on the public randomization draw of $r_t \in \mathcal{A}$ described above, an additional (“fictitious”) random variable $\tilde{y}_t \in \mathcal{Y}$ 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 $f_{i,r_{i,t}} : \mathcal{Y}_i \rightarrow \mathbb{R}$ be defined as in Lemma 3, and let

$$f_{i,t} = \begin{cases} f_{i,r_{i,t}}(y_{i,t}) & \text{if } S_{i,t} = 0, \\ f_{i,r_{i,t}}(\tilde{y}_{i,t}) & \text{if } S_{i,t} = NE, \\ 0 & \text{if } S_{i,t} = BR. \end{cases} \quad (21)$$

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 .²⁸ Later in the proof, $f_{i,t}$ will be a component of the “reward” earned by player i in period t , which will be reflected in player i ’s end-of-block continuation payoff function $\psi : H^{T+1} \rightarrow \mathbb{R}$.

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

$$\theta_{i,t} = \mathbf{1} \left\{ \exists t' \leq t : \left| \sum_{t''=1}^{t'-1} \delta^{t''-1} f_{i,t''} \right| \geq \bar{f} \right\}. \quad (22)$$

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 5 *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 .*

²⁸Intuitively, introducing the variable \tilde{y}_t , rather than simply using $y_{i,t}$ everywhere in (21), ensures that the distribution of $f_{i,t}$ does not depend on player i ’s opponents’ strategies.

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

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

■

Continuation Value Function: We now construct the continuation value function $\psi : H^{T+1} \rightarrow \mathbb{R}^N$. For each player i and end-of-block history h^{T+1} , player i 's continuation value $\psi_i(h^{T+1})$ will be defined as the sum of T “rewards” $\psi_{i,t}$, where $t = 1, \dots, T$, and a constant term c_i that does not depend on h^{T+1} .

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

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

$$\psi_{i,t} = \delta^{t-1} f_{i,r_{i,t}}(y_{i,t}). \quad (23)$$

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

$$\psi_{i,t} = \delta^{t-1} (u_i(\bar{\alpha}) - u_i(\alpha_t^*)). \quad (24)$$

3. Otherwise,

$$\psi_{i,t} = \delta^{t-1} (-\zeta_i \bar{u} - u_i(\alpha_t^*) + \mathbf{1}\{S_{i,t} = 0\} f_{i,r_{i,t}}(y_{i,t})). \quad (25)$$

The constant c_i is defined as

$$c_i = -\mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(\mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 0 \right\} u_i(\bar{\alpha}) - \mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 1 \right\} \zeta_i \bar{u} \right) \right] + \frac{1 - \delta^T}{1 - \delta} v_i^*. \quad (26)$$

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

$$|c_i| \leq 2\bar{u} \frac{1 - \delta^T}{1 - \delta}. \quad (27)$$

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

$$\psi_i(h^{T+1}) = c_i + \sum_{t=1}^T \psi_{i,t}. \quad (28)$$

E.3 Verification of the Equilibrium Conditions

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

Lemma 6 *We have*

$$\Pr \left(\max_{i \in I, t \in \{1, \dots, T\}} \theta_{i,t} = 0 \right) \geq 1 - \frac{\varepsilon}{20\bar{u}}. \quad (29)$$

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

$$\Pr \left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \delta^{t'-1} f_{i,t'} \right| \geq \bar{f} \right) \leq \frac{\varepsilon}{20\bar{u}N}. \quad (30)$$

To see this, let $\tilde{f}_{i,t} = f_{i,r_{i,t}}(\tilde{y}_{i,t})$. Note that the variables $(\tilde{f}_{i,t})_{t=1}^T$ are independent (unlike the variables $(f_{i,t})_{t=1}^T$). Since $(\tilde{f}_{i,t'})_{t'=1}^t$ 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

$$\Pr \left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \delta^{t'-1} f_{i,t'} \right| \geq \bar{f} \right) \leq \Pr \left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \delta^{t'-1} \tilde{f}_{i,t'} \right| \geq \bar{f} \right). \quad (31)$$

Since $(\tilde{f}_{i,t})_{t=1}^T$ are independent, Etemadi's inequality implies that

$$\Pr \left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \delta^{t'-1} \tilde{f}_{i,t'} \right| \geq \bar{f} \right) \leq 3 \max_{t \in \{1, \dots, T\}} \Pr \left(\left| \sum_{t'=1}^t \delta^{t'-1} \tilde{f}_{i,t'} \right| \geq \frac{\bar{f}}{3} \right). \quad (32)$$

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 (10), $\mathbb{E}[x_{i,t}] = 0$ by (8), and

$$\text{Var} \left(\sum_{t'=1}^t x_{i,t'} \right) = \sum_{t'=1}^t \text{Var}(x_{i,t'}) \leq \frac{1 - \delta^T}{1 - \delta} \frac{\bar{u}^2}{\eta} \quad \text{by (9)}.$$

Therefore, by Bernstein's inequality ((11), which again applies because $(\tilde{f}_{i,t})_{t=1}^T$ are independent) and (12), we have, for each $t \leq T$,

$$\Pr \left(\left| \sum_{t'=1}^t \delta^{t'-1} \tilde{f}_{i,t'} \right| \geq \frac{\bar{f}}{3} \right) \leq \frac{\varepsilon}{60\bar{u}N}. \quad (33)$$

Finally, (31), (32), and (33) together imply (30). ■

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

Lemma 7 *For each player i and block strategy profile σ , incentive compatibility holds (i.e., (18) is satisfied) if and only if*

$$\text{supp } \sigma_i(h^t) \subseteq \underset{a_{i,t} \in \mathcal{A}_i}{\text{argmax}} \mathbb{E}^{\sigma^{-i}} [\delta^{t-1} u_{i,t} + \psi_{i,t} | h^t, a_{i,t}] \quad \text{for all } t \text{ and } h^t. \quad (34)$$

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

$$\mathbb{E}^\sigma [\delta^{t'-1} u_{i,t'} + \psi_{i,t'} | h^t] = \mathbb{E}^\sigma \left[\delta^{t'-1} \left(\mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t'} = 0 \right\} u_i(\bar{\alpha}) - \mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t'} = 1 \right\} \zeta_i \bar{u} \right) | h^t \right]. \quad (35)$$

We now verify that (34) holds. Fix a player i , period t , and history h^t . We consider several cases, which parallel the definition of the reward $\psi_{i,t}$.

1. If $\theta_{j,t} = 0$ for all $j \in I$, recall that the equilibrium action profile is the r_t that is prescribed by public randomization z_t . For each action $a_i \neq r_{i,t}$, by (7) and (23), and recalling that $\bar{u} \geq \max_a u_i(a) - \min_a u_i(a)$, we have

$$\begin{aligned} & \mathbb{E}^{\sigma^{-i}} [\delta^{t-1} u_{i,t} + \psi_{i,t} | h^t, z_t, a_{i,t} = r_{i,t}] - \mathbb{E}^{\sigma^{-i}} [\delta^{t-1} u_{i,t} + \psi_{i,t} | h^t, z_t, a_{i,t} = a_i] \\ &= \delta^{t-1} (\mathbb{E} [u_i(r_t) + f_{i,r_{i,t}}(y_{i,t}) | a_{i,t} = r_{i,t}] - \mathbb{E} [u_i(a_i, r_{-i,t}) + f_{i,r_{i,t}}(y_{i,t}) | a_{i,t} = a_i]) \\ &\geq 0, \quad \text{so (34) holds.} \end{aligned}$$

2. If $\theta_{i,t} = 1$ and $\theta_{j,t} = 0$ for all $j \neq i$, then the reward $\psi_{i,t}$ specified by (24) does not depend on $y_{i,t}$. Hence, (34) reduces to the condition that every action in $\text{supp } \sigma_i(h^t)$ is a static best responses to $\sigma_{-i}(h^t)$. This conditions holds for the prescribed action profile, $(r'_i \in BR_i(r_{-i,t}), r_{-i,t})$ or $\varrho_{i,t}^{NE}$.
3. Otherwise: (a) If $S_{i,t} = 0$, then (34) holds because it holds in Case 1 above and (23) and (25) differ only by a constant independent of $y_{i,t}$. (b) If $S_{i,t} \neq 0$, then either $\theta_{j,t} = \theta_{j',t} = 1$ for distinct players j, j' , or there exists a unique player $j \neq i$ with $\theta_{j,t} = 1$, and for this player we have $\zeta_j = -1$. In both cases, ϱ_t^{NE} is prescribed. Since the reward $\psi_{i,t}$ specified by (25) does not depend on $y_{i,t}$, (34) reduces to the condition that every action in $\text{supp } \sigma_i(h^t)$ is a static best responses to $\sigma_{-i}(h^t)$, which holds for the prescribed action profile ϱ_t^{NE} .

Promise Keeping: This essentially holds by construction: we have

$$\begin{aligned}
& \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \psi_i(h^{T+1}) \right] \\
&= \frac{1-\delta}{1-\delta^T} \left(\mathbb{E}^\sigma \left[\sum_{t=1}^T (\delta^{t-1} u_{i,t} + \psi_{i,t}) \right] + c_i \right) \quad (\text{by (28)}) \\
&= \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^T \delta^{t-1} \left(\mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 0 \right\} u_i(\bar{\alpha}) - \mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 1 \right\} \zeta_i \bar{u} \right) + c_i \right] \quad (\text{by (35)}) \\
&= v_i^* \quad (\text{by (26)}), \text{ so (17) holds.}
\end{aligned}$$

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

Lemma 8 *For every end-of-block history h^{T+1} , we have*

$$\zeta_i \sum_{t=1}^T \psi_{i,t} \leq \bar{f} + \frac{2\bar{u}}{\eta} \quad \text{and} \quad (36)$$

$$\left| \sum_{t=1}^T \psi_{i,t} \right| \leq \bar{f} + \frac{2\bar{u}}{\eta} + 2\bar{u} \frac{1-\delta^T}{1-\delta}. \quad (37)$$

In addition,

$$\zeta_i c_i \leq -\frac{1-\delta^T}{1-\delta} \frac{\varepsilon}{8}. \quad (38)$$

To establish self generation ((19)), it suffices to show that, for each h^{T+1} , $\zeta_i \psi_i(h^{T+1}) \leq 0$ and $|\psi_i(h^{T+1})| \leq (\delta^T / (1 - \delta)) \varepsilon$. This now follows because

$$\begin{aligned}
\zeta_i \psi_i(h^{T+1}) &= \zeta_i \left(c_i + \sum_{t=1}^T \psi_{i,t} \right) \leq -\frac{1 - \delta^T}{1 - \delta} \frac{\varepsilon}{8} + \bar{f} + 2\bar{u}/\eta \quad (\text{by (36) and (38)}) \\
&\leq \frac{1 - \delta^T}{8(1 - \delta)} \left(-\varepsilon + 8 \left(\frac{1 - \delta}{1 - \delta^T} \right) (\bar{f} + 2\bar{u}/\eta) \right) \leq 0 \quad (\text{by (13)}), \quad \text{and} \\
|\psi_i(h^{T+1})| &\leq |c_i| + \left| \sum_{t=1}^T \psi_{i,t} \right| \\
&\leq 4\bar{u} \frac{1 - \delta^T}{1 - \delta} + \bar{f} + 2\bar{u}/\eta \quad (\text{by (27) and (37)}) \\
&= \frac{1 - \delta^T}{1 - \delta} 4\bar{u} + \bar{f} + 2\bar{u}/\eta \leq \frac{\delta^T}{1 - \delta} \varepsilon \quad (\text{by (14)}),
\end{aligned}$$

which completes the proof.

E.4 Omitted Details for the Proof of Lemma 4

We show that, with the stated definitions of T and \bar{f} , (15) and (16) imply (12)–(14). First, note that

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

Hence,

$$\begin{aligned}
\frac{2\bar{f}(1 - \delta^2)}{9\bar{u}(1 - \delta^{2T})} &< \frac{4}{9\bar{u}} \frac{1 - \delta}{1 - \delta^T} \sqrt{36 \left(\log \left(\frac{60\bar{u}}{\varepsilon} \right) + \log(N) \right) \frac{1 - \delta^T}{1 - \delta} \frac{\bar{u}^2}{\eta}} \\
&= \frac{4}{9} \sqrt{36 \left(\log \left(\frac{60\bar{u}}{\varepsilon} \right) + \log(N) \right) \frac{1 - \delta}{1 - \delta^T} \frac{1}{\eta}} \leq 1 \quad (\text{by (15)}).
\end{aligned}$$

Therefore,

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

Moreover,

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

Hence, we have

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

This establishes (12).

Next, we have

$$8 \frac{1-\delta}{1-\delta^T} \left(\bar{f} + \frac{2\bar{u}}{\eta} \right) = 8\bar{u} \left(\sqrt{36 \left(\log \left(\frac{60\bar{u}}{\varepsilon} \right) + \log(N) \right) \frac{1-\delta}{1-\delta^T} \frac{1}{\eta}} + \frac{1-\delta}{1-\delta^T} \frac{2}{\eta} \right) \leq \varepsilon \quad (\text{by (16)}). \quad (39)$$

This establishes (13).

Finally, by (39) and $8\bar{u}(1-\delta^T)/\delta^T \leq \varepsilon$, we have

$$4\bar{u} \frac{1-\delta^T}{\delta^T} + \frac{1-\delta}{\delta^T} \left(\bar{f} + \frac{2\bar{u}}{\eta} \right) = 4\bar{u} \frac{1-\delta^T}{\delta^T} + \frac{1-\delta^T}{\delta^T} \frac{1-\delta}{1-\delta^T} \left(\bar{f} + \frac{2\bar{u}}{\eta} \right) \leq 4\frac{\varepsilon}{8} + \frac{\varepsilon}{88} \leq \varepsilon.$$

This establishes (14).

E.5 Proof of Lemma 7

We show that player i has a profitable one-shot deviation from σ_i at some history h^t if and only if (34) is violated at h^t . To see this, we first calculate player i 's continuation payoff under σ from period $t+1$ onward (net of the constant c_i and the rewards already accrued $\sum_{t'=1}^t \psi_{i,t'}$). For each $t' \geq t+1$, there are several cases to consider.

1. If $\theta_{j,t'} = 0$ for all j , then by (8) and (23) we have

$$\mathbb{E}^\sigma \left[\delta^{t'-1} u_{i,t'} + \psi_{i,t'} | h^{t'} \right] = \delta^{t'-1} \left(u_i(\alpha_{t'}^*) + \mathbb{E} \left[f_{i,r_{i,t'}}(y_{i,t'}) | r_{i,t'} \right] \right) = \delta^{t'-1} u_i(\bar{\alpha}).$$

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

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

3. Otherwise: (a) If $S_{i,t'} = 0$, then by (8) and (25) (and recalling that player i 's equilibrium action is $r_{i,t'}$ when $S_{i,t'} = 0$) we have

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

(b) If $S_{i,t'} \neq 0$, then by (25) we have

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

In total, (35) holds, and player i 's net continuation payoff under σ from period $t + 1$ onward equals

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

By Lemma 5, the distribution of $\left((\theta_{n,t'})_{n \neq i} \right)_{t'=t+1}^T$ does not depend on player i 's period- t action, and hence neither does player i 's net continuation payoff under σ from period $t + 1$ onward. Therefore, player i 's period- t action $a_{i,t}$ maximizes her continuation payoff from period t onward if and only if it maximizes $\mathbb{E}^{\sigma-i} [\delta^{t-1} u_{i,t} + \psi_{i,t} | h^t, a_{i,t}]$.

E.6 Proof of Lemma 8

Define

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

Note that, by (23)–(25), we can write $\psi_{i,t} = \psi_{i,t}^v + \psi_{i,t}^f$. (Note that, if $\theta_{n,t} = 0$ for all $n \in I$, we have $\alpha_t^* = \bar{\alpha}$ and hence $\psi_{i,t}^v + \psi_{i,t}^f = \delta^{t-1} f_{i,a_{i,t}}(y_{i,t})$, as specified in (23).) We show that, for every end-of-block history h^{T+1} , we have

$$\zeta_i \sum_{t=1}^T \psi_{i,t}^v \in \left[-2\bar{u} \frac{1 - \delta^T}{1 - \delta}, 0 \right] \quad \text{and} \quad (40)$$

$$\left| \zeta_i \sum_{t=1}^T \psi_{i,t}^f \right| \leq \bar{f} + \frac{2\bar{u}}{\eta}. \quad (41)$$

Since $\psi_{i,t} = \psi_{i,t}^v + \psi_{i,t}^f$, (40) and (41) imply (36) and (37), which proves the first part of the lemma.

For (40), 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 $u_i(\alpha_t^*) \geq \sum_a \bar{\alpha}(a) \min \{u_i(a), \max_{a'_i} u_i(a'_i, a_{-i})\} \geq u_i(\bar{\alpha})$; and (ii) if $\zeta_i = -1$, we have $u_i(\alpha_t^*) \leq \max \{u_i(\bar{\alpha}), u_i(\alpha^{NE})\} = u_i(\bar{\alpha})$. In total, we have $\zeta_i(u_i(\bar{\alpha}) - u_i(\alpha_t^*)) \leq 0$. Since obviously $\zeta_i(u_i(\bar{\alpha}) - u_i(\alpha_t^*)) \geq -2\bar{u}$ and $-\bar{u} - \zeta_i u_i(\alpha_t^*) \geq -2\bar{u}$, we have

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

For (41), note that $S_{i,t} = 0$ implies $\theta_{i,t} = 0$, and hence

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

Since $\theta_{i,t+1} = 1$ whenever $\left| \sum_{t'=1, \dots, t} \delta^{t'-1} f_{i,a_{i,t'}}(y_{i,t'}) \right| \geq \bar{f}$, and in addition $|f_{i,a_{i,t}}(y_{i,t})| \leq 2\bar{u}/\eta$ by (10), this inequality implies (41).

For the second part of the lemma, by (26), we have

$$\begin{aligned} \zeta_i c_i &= \zeta_i \left(-\mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(\mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 0 \right\} u_i(\bar{\alpha}) - \mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 1 \right\} \zeta_i \bar{u} \right) \right] + \frac{1 - \delta^T}{1 - \delta} v_i^* \right) \\ &= \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(\mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 0 \right\} \zeta_i (v_i^* - u_i(\bar{\alpha})) + \mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 1 \right\} \underbrace{(\bar{u} + \zeta_i v_i^*)}_{\in [0, 2\bar{u}]} \right) \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(\mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 0 \right\} \left(\frac{-\varepsilon}{2} \right) + \mathbf{1} \left\{ \max_{j \neq i} \theta_{j,t} = 1 \right\} 2\bar{u} \right) \right] \text{ by (20)} \\ &\leq -\frac{1 - \delta^T}{1 - \delta} \left(\left(1 - \frac{\varepsilon}{20\bar{u}} \right) \frac{\varepsilon}{2} - \left(\frac{\varepsilon}{20\bar{u}} \right) 2\bar{u} \right) \text{ (by (29))} \\ &\leq -\frac{1 - \delta^T}{1 - \delta} \frac{\varepsilon}{8} \text{ (as } \varepsilon < \bar{u}/2 \text{).} \end{aligned}$$

F Proof of Theorem 3

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

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

Since $V(\varepsilon)$ is convex, it suffices to show that $\underline{v}, \bar{v} \in V(\varepsilon)$, where $\underline{v} = (\underline{v}_i)_{i \in I}$ and $\bar{v} = (\bar{v}_i)_{i \in I}$.

In the following lemma, for any $\alpha \in \Delta(\mathcal{A})$ and $f : \mathcal{A} \times \mathcal{Y} \rightarrow \mathbb{R}$, $\mathbb{E}^\alpha[f(r, y)]$ denotes expectation where $r \sim \alpha$ and then $y \sim p(\cdot|r)$, and $\mathbb{E}^{\alpha, a'_i}[f(r, y)]$ denotes expectation where $r \sim \alpha$ and then $y \sim p(\cdot|a'_i, r_{-i})$.

Lemma 9 *There exist $\alpha \in \Delta(\mathcal{A})$ and $\tau : \mathcal{A} \times \mathcal{Y} \rightarrow \mathbb{R}$ such that*

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

Moreover, if the constraints $\tau(r, y) \in [0, \frac{\delta}{1-\delta}\bar{u}]$ and $\mathbb{E}^\alpha[\tau(r, y)] \leq \bar{u}$ are replaced by $\tau(r, y) \in [-\frac{\delta}{1-\delta}\bar{u}, 0]$ and $\mathbb{E}^\alpha[\tau(r, y)] \geq -\bar{u}$, 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 $\alpha \in \Delta(\mathcal{A})$ and $w : \mathcal{A} \times \mathcal{Y} \rightarrow E$ such that

$$\begin{aligned} v &= \mathbb{E}^\alpha[(1-\delta)u(r) + \delta w(r, y)] \quad \text{and} \\ \mathbb{E}^\alpha[(1-\delta)u_i(r) + \delta w_i(r, y) | r_i = a_i] &\geq \mathbb{E}^{\alpha, a'_i}[(1-\delta)u_i(a'_i, r_{-i}) + \delta w_i(r, y) | r_i = a_i], \end{aligned}$$

for all $i, a_i \in \text{supp } \alpha_i, a'_i \in \mathcal{A}_i$. Since $v \in E$ and $w(r, y) \in E$ for all r, y , we have $v_i - w_i(r, y) = b_i(v_1 - w_1(r, y))$ for all i, r, y . Since $\bar{v}_1 \geq v_1$ for all $v \in E$, if $v = \bar{v}$ then $w_1(r, y) \leq v_1$ for all r, y . Hence, taking $v = \bar{v} = (1-\delta)u(\alpha) + \delta\mathbb{E}^\alpha[w(r, y)]$ and defining

$\tau(r, y) = \frac{\delta}{1-\delta} (\bar{v}_1 - w_1(r, y))$, we have $\tau(r, y) \in [0, \frac{\delta}{1-\delta} \bar{u}]$ and

$$\begin{aligned} \mathbb{E}^\alpha [\tau(r, y)] &= \frac{\delta}{1-\delta} \mathbb{E}^\alpha [(1-\delta) u_1(\alpha) + \delta \mathbb{E}^\alpha [w_1(r, y)] - w_1(r, y)] \\ &\leq \frac{\delta}{1-\delta} (1-\delta) (u_1(\alpha) - \mathbb{E}^\alpha [w_1(r, y)]) \leq \bar{u}. \end{aligned}$$

Moreover, we have

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

and, for all $i, a_i \in \text{supp } \alpha_i, a'_i \in \mathcal{A}_i$,

$$\begin{aligned} \mathbb{E}^\alpha [(1-\delta) u_i(r) + \delta w_i(r, y) | r_i = a_i] &\geq \mathbb{E}^{\alpha, a'_i} [(1-\delta) u_i(a'_i, r_{-i}) + \delta w_i(r, y) | r_i = a_i] \\ &\iff \\ \mathbb{E}^\alpha \left[u_i(r) + \frac{\delta}{1-\delta} (w_i(r, y) - \bar{v}_i) | r_i = a_i \right] &\geq \mathbb{E}^{\alpha, a'_i} \left[u_i(a'_i, r_{-i}) + \frac{\delta}{1-\delta} (w_i(r, y) - \bar{v}_i) | r_i = a_i \right] \\ &\iff \\ \mathbb{E}^\alpha [u_i(r) - b_i \tau(r, y) | r_i = a_i] &\geq \mathbb{E}^{\alpha, a'_i} [u_i(a'_i, r_{-i}) - b_i \tau(r, y) | r_i = a_i]. \end{aligned}$$

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

Taking α and τ as in Lemma 9 and recalling that $|b_i| \leq 1$ for all i , we see that $\sum_i \bar{g}(\alpha) / N$ is bounded by the value of the program

$$\max_{(\mathcal{Y}, \mathcal{P}), r, a, \tau} \frac{1}{N} \sum_{i \in I} |\mathbb{E}^r [\tau(y)] - \mathbb{E}^{(a_i, r_{-i})} [\tau(y)]| \quad \text{s.t.}$$

$$\tau(y) \in \left[0, \frac{\delta}{1-\delta} \bar{u} \right] \quad \text{for all } y, \quad (42)$$

$$\mathbb{E}^r [\tau(y)] \leq \bar{u}. \quad (43)$$

In turn, this is bounded by the sum of the value of the program

$$\max_{(\mathcal{Y}, p), r, a, \tau} \frac{1}{N} \sum_{i \in I_+} \mathbb{E}^r [\tau(y)] - \mathbb{E}^{(a_i, r_{-i})} [\tau(y)] \quad \text{s.t.} \quad (42), (43), \quad (44)$$

and the sum of the corresponding program with I_- in place of I_+ . To prove the theorem, we show that, for any $\underline{\pi} > 0$ and $\rho > 0$, each of these values converges to 0 along any sequence (N, δ) where $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$. Without loss, it suffices to consider the first program, (44).

We first establish that the solution is a tail test. Let $N_+ = |I_+|$.

Lemma 10 *Program (44) is solved by a tail test, where $\mathcal{Y} = \mathcal{X} = \mathcal{A} = \{0, 1\}^{N_+}$; $q(y|x) = \mathbf{1}\{y = x\}$ for all y, x ; $r_i \neq a_i$ for all i ; $\pi_i(r_i|r_i) = 1 - \underline{\pi}$ and $\pi_i(a_i|r_i) = \underline{\pi}$ for all $i, r_i \neq a_i$; and, letting $n = |\{i : y_i = r_i\}|$,*

$$\tau(y) = \begin{cases} \frac{\delta}{1-\delta} \bar{u} & \text{if } n > n^*, \\ \beta \frac{\delta}{1-\delta} \bar{u} & \text{if } n = n^*, \\ 0 & \text{if } n < n^*, \end{cases}$$

for some $n^* \in \{0, 1, \dots, N_+\}$ and $\beta \in [0, 1]$.

We prove Lemma 10 in the next subsection. Lemma 10 implies that the program becomes

$$\begin{aligned} \max_{n^* \in \{0, 1, \dots, N_+\}, \beta \in [0, 1]} & \frac{N_+}{N} \frac{\delta}{1-\delta} \bar{u} (1 - 2\underline{\pi}) (\beta \Pr(n_{-i} = n^* - 1) + (1 - \beta) \Pr(n_{-i} = n^*)) \\ \text{s.t.} & \beta \Pr(n = n^*) + \Pr(n \geq n^* + 1) \leq \frac{\bar{u}}{\frac{\delta}{1-\delta} \bar{u}}, \end{aligned}$$

where $n_{-i} = |\{j \neq i : y_j = r_j\}|$ and the probabilities are binomial with parameter $1 - \underline{\pi}$. Since the value of the program is maximized when $N = N_+$, we henceforth assume that this is the case. We now show that, for any $\rho > 0$, there exist $c_0, c_1 > 0$ such that, for each N , the value of the program is at most

$$\max \left\{ \frac{\delta}{1-\delta} \bar{u} \exp(-c_0 N^{1-\rho}), c_1 N^{-\rho/2} \bar{u} \right\}.$$

This completes the proof, as if \bar{u} is fixed and $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$ then both terms converge to 0.

We bound the program separately for n^* such that $|n^* - (1 - \underline{\pi})N| > N^{1-\rho/2}$ and n^* such that $|n^* - (1 - \underline{\pi})N| \leq N^{1-\rho/2}$. In the first case, by Hoeffding's inequality, there exists $c_0 > 0$ such that

$$\min \{ \Pr(n_{-i} \geq n^* - 1), \Pr(n_{-i} \leq n^*) \} \leq \exp(-c_0 N^{1-\rho}).$$

Since the value of the program is at most $2 \frac{\delta}{1-\delta} \bar{u} \min \{ \Pr(n_{-i} \geq n^* - 1), \Pr(n_{-i} \leq n^*) \}$, this gives the desired bound when $|n^* - (1 - \underline{\pi})N| > N^{1-\rho/2}$.

For the second case, the value of the program is at most

$$\begin{aligned} & \frac{\delta}{1-\delta} \bar{u} \frac{\beta \Pr(n_{-i} = n^* - 1) + (1-\beta) \Pr(n_{-i} = n^*)}{\beta \Pr(n \geq n^*) + \Pr(n \geq n^* + 1)} \frac{\bar{u}}{\frac{\delta}{1-\delta} \bar{u}} \\ & \leq \frac{(\beta \Pr(n_{-i} = n^* - 1) + (1-\beta) \Pr(n_{-i} = n^*))}{\beta \Pr(n \geq n^*) + (1-\beta) \Pr(n \geq n^* + 1)} \bar{u} \\ & \leq \left(\frac{\Pr(n_{-i} = n^* - 1)}{\Pr(n \geq n^*)} + \frac{\Pr(n_{-i} = n^*)}{\Pr(n \geq n^* + 1)} \right) \bar{u}. \end{aligned}$$

By McKay (1989, Theorem 2), for any $m \geq (1 - \underline{\pi})N$, we have

$$\Pr(n \geq m) \geq \sqrt{N\underline{\pi}(1-\underline{\pi})} \Pr(n_{-i} = m-1) \frac{1 - \Phi\left((m - (1 - \underline{\pi})N) / \sqrt{N\underline{\pi}(1 - \underline{\pi})}\right)}{\phi\left((m - (1 - \underline{\pi})N) / \sqrt{N\underline{\pi}(1 - \underline{\pi})}\right)}.$$

If $(1 - \underline{\pi})N \leq n^* \leq (1 - \underline{\pi})N + N^{1-\rho/2}$, applying this inequality for $m \in \{n^*, n^* + 1\}$, together with the inverse Mills ratio inequality $(1 - \Phi(x)) / \phi(x) \geq 1 / (1 + x)$, we have

$$\begin{aligned} \frac{\Pr(n_{-i} = n^* - 1 | r_{-i})}{\Pr(n \geq n^* | r)} + \frac{\Pr(n_{-i} = n^* | r_{-i})}{\Pr(n \geq n^* + 1 | r)} & \leq 2 \frac{1}{\sqrt{N\underline{\pi}(1-\underline{\pi})}} \left(\frac{n^* - (1 - \underline{\pi})N}{\sqrt{N\underline{\pi}(1-\underline{\pi})}} + 1 \right) \\ & \leq 2 \left(\frac{N^{-\rho/2}}{\sqrt{\underline{\pi}(1-\underline{\pi})}} + \frac{1}{\sqrt{N\underline{\pi}(1-\underline{\pi})}} \right). \end{aligned}$$

Thus, there exists $c_1 > 0$ such that the value of the program is at most $c_1 N^{-\rho/2} \bar{u}$. Symmetrically, the same bound applies when $(1 - \underline{\pi})N - N^{1-\rho/2} \leq n^* \leq (1 - \underline{\pi})N$.

F.1 Proof of Lemma 10

First, consider the sub-program where $(\mathcal{X}, \pi, \mathcal{Y}, q)$ is fixed, so the objective is maximized over (r, a, τ) . By Blackwell's theorem, the value of the sub-program with signal distribution p is greater than that with signal distribution \hat{p} , if \hat{p} is a *garbling* of p . (That is, viewing p and \hat{p} as $|\mathcal{Y}| \times |\mathcal{A}|$ matrices, there is a $|\mathcal{Y}| \times |\mathcal{Y}|$ Markov matrix M such that $\hat{p} = Mp$.) For any noise structure (\mathcal{X}, π) , the action monitoring structure (\mathcal{Y}, p) induced by any outcome monitoring structure (\mathcal{Y}, q) is clearly a garbling of that induced by the outcome monitoring structure where $\mathcal{Y} = \mathcal{X}$ and $q(y|x) = \mathbf{1}\{y = x\}$ for all y, x , so that $p(y|a) = \pi(y|a)$ for all y, a . It is thus without loss to focus on this (\mathcal{Y}, q) .

In addition, if we let $\mathcal{X} = \mathcal{A}$ and, for each $r, a \in \mathcal{A}$ and i , let

$$\bar{\pi}_i(y_i|r_i) = \begin{cases} 1 - \underline{\pi} & \text{if } y_i = r_i, \\ \underline{\pi} & \text{if } y_i = a_i, \\ 0 & \text{otherwise,} \end{cases} \quad \bar{\pi}_i(y_i|a_i) = \begin{cases} 1 - \underline{\pi} & \text{if } y_i = a_i, \\ \underline{\pi} & \text{if } y_i = r_i, \\ 0 & \text{otherwise,} \end{cases}$$

and $\bar{\pi}_i(y_i|\tilde{a}_i) = \mathbf{1}\{y_i = \tilde{a}_i\}$ for $\tilde{a}_i \notin \{a_i, r_i\}$,

then π_i is a garbling of $\bar{\pi}_i$ for each i , and hence π is a garbling of $\bar{\pi}$. To see this, since $\underline{\pi} < 1/2$, the matrix $\bar{\pi}_i$ is invertible, and the matrix inverse $\bar{\pi}_i^{-1}$ is given by

$$\bar{\pi}_i^{-1}(y_i|r_i) = \begin{cases} \frac{1-\underline{\pi}}{1-2\underline{\pi}} & \text{if } y_i = r_i, \\ -\frac{\underline{\pi}}{1-2\underline{\pi}} & \text{if } y_i = a_i, \\ 0 & \text{otherwise,} \end{cases} \quad \bar{\pi}_i^{-1}(y_i|a_i) = \begin{cases} \frac{1-\underline{\pi}}{1-2\underline{\pi}} & \text{if } y_i = a_i, \\ -\frac{\underline{\pi}}{1-2\underline{\pi}} & \text{if } y_i = r_i, \\ 0 & \text{otherwise,} \end{cases}$$

and $\bar{\pi}_i^{-1}(y_i|\tilde{a}_i) = \mathbf{1}\{y_i = \tilde{a}_i\}$ for $\tilde{a}_i \notin \{a_i, r_i\}$.

We can then calculate the matrix $M_i := \pi_i \bar{\pi}_i^{-1}$ as

$$M_i(y_i|\tilde{a}_i) = \begin{cases} \frac{(1-\underline{\pi})\pi_i(y_i|r_i) - \underline{\pi}\pi_i(y_i|a_i)}{1-2\underline{\pi}} & \text{if } \tilde{a}_i = r_i, \\ \frac{(1-\underline{\pi})\pi_i(y_i|a_i) - \underline{\pi}\pi_i(y_i|r_i)}{1-2\underline{\pi}} & \text{if } \tilde{a}_i = a_i, \\ \pi_i(y_i|\tilde{a}_i) & \text{otherwise.} \end{cases}$$

Note that, for each \tilde{a}_i , we have $\sum_{y_i} M_i(y_i|\tilde{a}_i) = 1$ and, since $\pi_i(y_i|\tilde{a}_i) \geq \underline{\pi}$ for all y_i, \tilde{a}_i ,

$$\frac{(1-\underline{\pi})\pi_i(y_i|r_i) - \underline{\pi}\pi_i(y_i|a_i)}{1-2\underline{\pi}} \geq \frac{(1-\underline{\pi})\underline{\pi} - \underline{\pi}(1-\underline{\pi})}{1-2\underline{\pi}} = 0,$$

and similarly $\frac{(1-\pi)\pi_i(y_i|a_i)-\pi\pi_i(y_i|r_i)}{1-2\pi} \geq 0$. Thus, M_i is a Markov matrix satisfying $\pi_i = M_i\bar{\pi}_i$.

It is thus without loss to take $(\mathcal{X}, \pi) = (\mathcal{A}, \bar{\pi})$. The program then simplifies to

$$\begin{aligned} \max_{r,a,\tau} \frac{1}{N} \sum_{i \in I_+} \sum_y \bar{\pi}(y|r) \left(1 - \frac{\bar{\pi}_i(y_i|a_i)}{\bar{\pi}_i(y_i|r_i)}\right) \tau(y) \quad \text{s.t.} \\ \tau(y) \in \left[0, \frac{\delta}{1-\delta} \bar{u}\right] \quad \text{for all } y, \\ \sum_y p(y|r) \tau(y) \leq \bar{u}. \end{aligned}$$

Here it is without loss to take $a_i \neq r_i$ for all i , as if $a_i = r_i$ then the same value can be attained by taking $\tau(y)$ independent of y_i , at which point a_i can then be taken different from r_i without affecting the value. Letting $\lambda \geq 0$ denote the multiplier on $\sum_y p(y|r) \tau(y) \leq \bar{u}$, we see that, for each y , the Lagrangian is increasing in $\tau(y)$ if and only if

$$\frac{1}{N} \sum_i \frac{\bar{\pi}_i(y_i|a_i)}{\bar{\pi}_i(y_i|r_i)} \leq 1 - \lambda.$$

It follows that $\tau(y)$ takes the prescribed form.

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] Border, Kim C., and Joel Sobel. “Samurai Accountant: A Theory of Auditing and Plunder.” *The Review of Economic Studies* 54.4 (1987): 525-540.
- [10] Bowles, Samuel, and Herbert Gintis. *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton University Press, 2011.
- [11] Boyd, Robert, and Peter J. Richerson. “The Evolution of Reciprocity in Sizable Groups.” *Journal of Theoretical Biology* 132.3 (1988): 337-356.
- [12] Camera, Gabriele, and Marco Casari. “Cooperation among Strangers under the Shadow of the Future.” *American Economic Review* 99.3 (2009): 979-1005.
- [13] Camera, Gabriele, Marco Casari, and Maria Bigoni. “Money and Trust among Strangers.” *Proceedings of the National Academy of Sciences* 110.37 (2013): 14889-14893.
- [14] Cover, Thomas M., and Joy A. Thomas. *Elements of Information Theory, 2nd ed.* Wiley (2006).
- [15] Deb, Joyee, Takuo Sugaya, and Alexander Wolitzky. “The Folk Theorem in Repeated Games with Anonymous Random Matching.” *Econometrica* 88.3 (2020): 917-964.
- [16] Dellarocas, Chrysanthos. “The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms.” *Management Science* 49.10 (2003): 1407-1424.
- [17] Duffy, John, and Jack Ochs. “Cooperative Behavior and the Frequency of Social Interaction.” *Games and Economic Behavior* 66.2 (2009): 785-812.
- [18] Edmiston, Jake, and Graeme Hamilton, “The Last Days of Quebec’s Maple Syrup Rebellion,” *National Post*, April 6, 2018.
- [19] Ekmekci, Mehmet, Olivier Gossner, and Andrea Wilson. “Impermanent Types and Permanent Reputations.” *Journal of Economic Theory* 147.1 (2012): 162-178.
- [20] Ellickson, Robert C. *Order without Law: How Neighbors Settle Disputes*. Harvard University Press, 1991.
- [21] Ellison, Glenn. “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching.” *Review of Economic Studies* 61.3 (1994): 567-588.
- [22] Faingold, Eduardo. “Reputation and the Flow of Information in Repeated Games.” *Econometrica* 88.4 (2020): 1697-1723.
- [23] Feigenberg, Benjamin, Erica Field, and Rohini Pande. “The Economic Returns to Social Interaction: Experimental Evidence from Microfinance.” *Review of Economic Studies* 80.4 (2013): 1459-1483.

- [24] Forges, Françoise. “An Approach to Communication Equilibria.” *Econometrica* 54.6 (1986): 1375-1385.
- [25] Fudenberg, Drew, and David K. Levine. “Efficiency and Observability with Long-Run and Short-Run Players.” *Journal of Economic Theory* 62.1 (1994): 103-135.
- [26] Fudenberg, Drew, and David K. Levine. “Continuous Time Limits of Repeated Games with Imperfect Public Monitoring.” *Review of Economic Dynamics* 10.2 (2007): 173-192.
- [27] Fudenberg, Drew, and David K. Levine. “Repeated Games with Frequent Signals.” *Quarterly Journal of Economics* 124.1 (2009): 233-265.
- [28] Fudenberg, Drew, David Levine, and Eric Maskin. “The Folk Theorem with Imperfect Public Information.” *Econometrica* 62.5 (1994): 997-1039.
- [29] Fudenberg, Drew, David Levine, and Wolfgang Pesendorfer. “When are Nonanonymous Players Negligible?” *Journal of Economic Theory* 79.1 (1998): 46-71.
- [30] Georgiadis, George, and Balazs Szentes. “Optimal Monitoring Design.” *Econometrica* 88.5 (2020): 2075-2107.
- [31] Gossner, Olivier. “Simple Bounds on the Value of a Reputation.” *Econometrica* 79.5 (2011): 1627-1641.
- [32] Gossner, Olivier, Penelope Hernández, and Abraham Neyman. “Optimal Use of Communication Resources.” *Econometrica* 74.6 (2006): 1603-1636.
- [33] Green, Edward J. “Noncooperative Price Taking in Large Dynamic Markets.” *Journal of Economic Theory* 22.2 (1980): 155-182.
- [34] Green, Edward J., and Robert H. Porter. “Noncooperative Collusion under Imperfect Price Information.” *Econometrica* 52.1 (1984): 87-100.
- [35] Hellman, Ziv, and Ron Peretz. “A survey on entropy and economic behaviour.” *Entropy* 22.2 (2020): 157.
- [36] Hoffmann, Florian, Roman Inderst, and Marcus Opp. “Only Time Will Tell: A Theory of Deferred Compensation.” *The Review of Economic Studies* 88.3 (2021): 1253-1278.
- [37] Holmström, Bengt. “Moral Hazard in Teams.” *Bell Journal of Economics* (1982): 324-340.
- [38] Hörner, Johannes, and Wojciech Olszewski. “The Folk Theorem for Games with Private Almost-Perfect Monitoring” *Econometrica* 74.7 (2006): 1499-1544.
- [39] Kandori, Michihiro. “Social Norms and Community Enforcement.” *Review of Economic Studies* 59.1 (1992): 63-80.

- [40] Kandori, Michihiro, and Hitoshi Matsushima. "Private Observation, Communication and Collusion." *Econometrica* 66.3 (1998): 627-652.
- [41] Karlan, Dean S. "Social Connections and Group Banking." *Economic Journal* 117.517 (2007): F52-F84.
- [42] Kuitenbrouwer, "'It's Crazy, Isn't It': Quebec's Maple Syrup Rebels Face Ruin as Cartel Crushes Dissent," *Financial Post*, December 5, 2016.
- [43] Li, Anqi, and Ming Yang. "Optimal incentive contract with endogenous monitoring technology." *Theoretical Economics* 15.3 (2020): 1135-1173.
- [44] Mailath, George, and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, 2006.
- [45] Matsushima, Hitoshi. "Repeated Games with Private Monitoring: Two Players." *Econometrica* 72.3 (2004): 823-852.
- [46] McKay, Brendan. "On Littlewood's Estimate for the Binomial Distribution." *Advances in Applied Probability* 21.2 (1989), 475-478.
- [47] 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.
- [48] 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).
- [49] Mookherjee, Dilip, and Ivan Png. "Optimal Auditing, Insurance, and Redistribution." *The Quarterly Journal of Economics* 104.2 (1989): 399-415.
- [50] Neyman, Abraham, and Daijiro Okada. "Strategic Entropy and Complexity in Repeated Games." *Games and Economic Behavior* 29.1-2 (1999): 191-223.
- [51] Neyman, Abraham, and Daijiro Okada. "Repeated Games with Bounded Entropy." *Games and Economic Behavior* 30.2 (2000): 228-247.
- [52] Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
- [53] 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.
- [54] Qian, Yingyi. "Incentives and Loss of Control in an Optimal Hierarchy." *Review of Economic Studies* 61.3 (1994): 527-544.
- [55] Reinganum, Jennifer F., and Louis L. Wilde. "Income tax compliance in a principal-agent framework." *Journal of Public Economics* 26.1 (1985): 1-18.

- [56] Sabourian, Hamid. “Anonymous Repeated Games with a Large Number of Players and Random Outcomes.” *Journal of Economic Theory* 51.1 (1990): 92-110.
- [57] Sannikov, Yuliy, and Andrzej Skrzypacz. “Impossibility of Collusion under Imperfect Monitoring with Flexible Production.” *American Economic Review* 97.5 (2007): 1794-1823.
- [58] Sannikov, Yuliy, and Andrzej Skrzypacz. “The Role of Information in Repeated Games with Frequent Actions.” *Econometrica* 78.3 (2010): 847-882.
- [59] Sason, Igal, and Sergio Verdú. “ f -Divergence Inequalities.” *IEEE Transactions on Information Theory* 62.11 (2016): 5973-6006.
- [60] Seabright, Paul. *The Company of Strangers*. Princeton University Press, 2004.
- [61] Sugaya, Takuo, and Alexander Wolitzky. “A Few Bad Apples Spoil the Barrel: An Anti-folk Theorem for Anonymous Repeated Games with Incomplete Information.” *American Economic Review* 110.12 (2020): 3817-35.
- [62] Sugaya, Takuo, and Alexander Wolitzky, “Communication and Community Enforcement.” *Journal of Political Economy* 129.9 (2021): 2595-2628.
- [63] Sugaya, Takuo, and Alexander Wolitzky, “Monitoring versus Discounting in Repeated Games.” *Econometrica* 91.5 (2023): 1727-1761.
- [64] Tadelis, Steven. “Reputation and Feedback Systems in Online Platform Markets.” *Annual Review of Economics* 8.1 (2016): 321-340.