

Repeated Games with Many Players*

Takuo Sugaya and Alexander Wolitzky

Stanford GSB and MIT

March 22, 2021

Abstract

Motivated by the problem of sustaining cooperation in large communities with limited information, we analyze sequences of repeated games with imperfect public monitoring where the population size N , discount factor δ , and signal information K (which can measure either the cardinality of the signal space or the mutual information between signals and actions) vary together. We show that if $(1 - \delta)N/K \rightarrow \infty$ then payoffs cannot exceed those consistent with approximately myopic play. If instead $(1 - \delta)N \log(N)/K \rightarrow 0$ then a folk theorem holds under *random auditing*, where each player's action is monitored with the same probability in every period. Thus, up to $\log(N)$ slack, the prospects for cooperation are determined by the ratio of the discount rate $r \approx 1 - \delta$ and the per-capita information K/N , and there is no benefit of monitoring different players' actions "jointly." If attention is restricted to strongly symmetric equilibria, cooperation is possible only under much more severe parameter restrictions.

Keywords: repeated games, large population, random auditing, mutual information

JEL codes: C72, C73

*For helpful comments, we thank Drew Fudenberg, Stephen Morris, Yuliy Sannikov, Satoru Takahashi, and seminar participants at Indiana.

1 Introduction

Repeated games are used to model long-term relationships in a variety of settings, many of which involve a large number of players. Examples include voluntary public-goods provision (Pecorino, 1999; Miguel and Gugerty, 2005), community resource management (Ostrom, 1990; Ellickson, 1991), and informal risk-sharing (Kocherlakota, 1996; Ligon, Thomas, and Worrall, 2002). However, the standard analysis of repeated games with patient players (e.g., Fudenberg, Levine, and Maskin, 1994; henceforth “FLM”) fixes all parameters of the game except the discount factor δ and considers the limit as $\delta \rightarrow 1$. This approach does not capture situations where, while players are patient (i.e., $\delta \approx 1$), they are not necessarily patient “compared to the population size N ” (i.e., $(1 - \delta)N$ may or may not be close to 0). In addition, since standard results are based on statistical identifiability conditions that hold “generically” regardless of the number of players, they also do not capture the possibility that more information may be required to support cooperation in larger groups.

In this paper, we extend the standard analysis of repeated games with imperfect public monitoring by considering sequences of games where the population size N , discount factor δ , and “signal information” K vary together, subject to uniform upper bounds on the magnitude of the players’ stage-game payoffs and the number of actions available to each player, and a uniform lower bound on “noise,” which we formalize as a constraint that each player takes each available action with positive probability. In our model, the signal information K can measure either the cardinality of the public signal space (in logs) or the mutual information between the public signal and the action profile (viewed as a random variable whose distribution depends on the public history); our results are the same either way.

Up to a little slack, we find that the scope for cooperation is determined by the ratio of the discount rate $r \approx (1 - \delta)$ and the per-capita signal information K/N . Thus, cooperation in very large groups requires very patient players or very informative signals, or both. When rN/K is small, so cooperation is possible, cooperation can be supported under a simple *random auditing* scheme, where each player is monitored with the same probability in every period. If instead monitoring and incentive-provision are implemented “collectively,” cooperation is impossible unless r is exponentially small in N . Since this condition seems unlikely

to be satisfied in most applications, we interpret this result as an impossibility theorem for large-group cooperation under collective monitoring, and view the totality of our results as endorsing the importance of individual monitoring in large groups. For example, our results suggest that it is more feasible to sustain voluntary public-goods provision in a large group by randomly monitoring a few agents' contributions rather than by monitoring the total level of public goods provided.

More precisely, we first show that if $(1 - \delta)N/K \rightarrow \infty$ along a sequence of repeated games satisfying our assumptions, then eventually (along the sequence) payoffs in any perfect public equilibrium (PPE) cannot significantly exceed those consistent with approximately myopic play. As we discuss in detail below, this finding improves on earlier results by Green (1980), Sabourian (1990), Fudenberg, Levine, and Pesendorfer (1998), and Al-Najjar and Smorodinsky (2000, 2001), which establish conditions under which play in repeated games becomes approximately myopic in the limit as $N \rightarrow \infty$ for fixed δ and K .¹ Our result comes from combining probabilistic arguments of the kind used in these papers with dynamic programming and geometric arguments of the kind used by FLM. Roughly speaking, a measure of the “average influence” of the players' actions on the public signal can be bounded by $\sqrt{K/N}$, while the “average movement” of the players' continuation payoffs along a tangent hyperplane at some extreme point of the equilibrium payoff set can be bounded by $1/\sqrt{1 - \delta}$, and the “average incentive” that can be provided to the players at an extreme point of the equilibrium payoff set without sacrificing efficiency is bounded by the product of these terms, or $\sqrt{K/((1 - \delta)N)}$. Since this product goes to 0 when $(1 - \delta)N/K \rightarrow \infty$, payoffs cannot exceed those consistent with approximately myopic play along such a sequence of repeated games. Significantly, this negative result imposes no restriction on the monitoring structures considered along the sequence, except that the signal information (i.e., signal-space cardinality or mutual information) equals K .

We then provide a near-converse to this negative result: if $(1 - \delta)N \log(N)/K \rightarrow 0$ then a folk theorem holds when monitoring takes the form of random auditing, where a constant number of players (proportional to K) are publicly selected, uniformly at random, at the

¹As we will see, the arguments in some of these papers can be adapted to give results that apply when N , δ , and K vary together; however, these results are substantially weaker than ours.

end of each period, and the public signal reveals only these players’ actions. Thus, outside of the “small” range of parameters where $(1 - \delta) N/K \not\rightarrow \infty$ and $(1 - \delta) N \log(N)/K \not\rightarrow 0$ (so that $\delta \rightarrow 1$ at least as fast as $N/K \rightarrow \infty$ but not faster than $N \log(N)/K \rightarrow \infty$), random auditing supports cooperation whenever cooperation can be supported under any monitoring structure with the same signal information K : that is, random auditing is an approximately optimal monitoring structure in large populations. In particular, there is little benefit of “joint monitoring” schemes, where the public signal cannot be decomposed into conditionally independent signals of distinct players’ actions.²

Our main result thus establishes that a simple form of “individual monitoring”—random auditing—is approximately optimal outside a small parameter range. We complement this result by investigating the prospects for cooperation in strongly symmetric equilibrium (SSE) in symmetric games. SSE are equilibria where play is symmetric at every history. These equilibria capture an extreme form of joint monitoring, where everyone is punished together after “bad” aggregate outcomes. In symmetric games, the gap between the range of parameters where cooperation is supportable in general PPE as opposed to SSE is thus a measure of the advantage of individual monitoring over “fully collectivized” monitoring.

We show that this gap is large: if there exists $\rho > 0$ such that $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$ along a sequence of repeated games satisfying our assumptions, then eventually payoffs in any SSE cannot exceed those consistent with approximately myopic play.³ Hence, for the “large” set of parameters where $(1 - \delta) N \log(N)/K \rightarrow 0$ but $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$ for some $\rho > 0$, a folk theorem holds for PPE under random auditing, but essentially no long-term incentives can be provided in any SSE. Intuitively, an optimal SSE provides incentive through a statistical “tail test,” where the players are collectively punished if the number of players who take the desired action falls below a threshold, and we show that such a tail test can provide incentives for N players to take the desired action while maintaining a constant “false positive rate” only if the size of the penalty when the test is failed can be taken to be exponentially large in N . Since the maximum penalty size in a repeated game

²Note our careful wording: random auditing is approximately optimal, but joint monitoring schemes may also be approximately optimal, and might even slightly outperform random auditing. Whether this is so is an open question.

³We also establish a near-converse: if there exists $\rho > 0$ such that $(1 - \delta) \exp(N^{1+\rho}) \rightarrow 0$ then a large set of symmetric payoff vectors can be supported in SSE.

is proportional to $1/(1 - \delta)$, it follows that $1/(1 - \delta)$ must be exponentially large in N to support cooperation in a SSE (in contrast to our main result, which says that $1/(1 - \delta)$ must be only approximately linear in N to support cooperation in a PPE).

The paper proceeds as follows: following a brief discussion of related literature, Section 2 presents the model, Section 3 presents our main results for PPE in general games, Section 4 presents our supporting results for SSE in symmetric games, and Section 5 concludes.

1.1 Related Literature

The most closely related papers are those on PPE in repeated games with imperfect public monitoring and a fixed population size (e.g., Abreu, Pearce, and Stacchetti, 1990; henceforth, “APS”; Fudenberg and Levine, 1994; FLM) and those on justifying myopic play in large-population repeated games (Green, 1980; Sabourian, 1990; Fudenberg, Levine, and Pesendorfer, 1998; Al-Najjar and Smorodinsky, 2000, 2001; Pai, Roth, and Ullman, 2014). Our results combine ideas from these two branches of the literature to characterize how the population size, the discount factor, and the informativeness of the public signal interact to determine the prospects for cooperation.⁴

We believe this paper is the first in the repeated games literature to measure signal informativeness by the mutual information between the signal and the action profile, and to investigate maximal equilibrium payoffs subject to a constraint on informativeness.⁵ Our results thus concern optimal monitoring structure design in repeated games, although we consider only asymptotic results rather than exact optimality for fixed parameters. In the context of static moral hazard problems, optimal monitoring design subject to information-theoretic constraints was recently studied by Georgiadis and Szentes (2020), Hoffman, Inderst, and Opp (2020), and Li and Yang (2020), while an earlier literature (Maskin and Riley, 1985;

⁴For fixed N , Hörner and Takahashi (2016) consider the rate of convergence in δ of the PPE payoff set to the limit set characterized by Fudenberg and Levine (1994) and Fudenberg, Levine, and Maskin (1994). Their analysis, which concerns rates of convergence in one parameter to a known limit payoff set, is not very closely related to ours, which asks how the relationship between different parameters determines the limit payoff set. Farther afield, there is also some work suggesting that cooperation in repeated games is harder to sustain in larger groups based on evolutionary models (Boyd and Richerson, 1988) and simulations (Bowles and Gintis, 2011; Chapter 5).

⁵Entropy methods have previously been used in the repeated games literature to study automaton strategies (Neyman and Okada, 1999, 2000), communication (Gossner, Hernández, and Neyman, 2006), and reputation effects (Gossner, 2011; Ekmekci, Gossner, and Wilson, 2011; Faingold, 2020).

Khalil and Lawarree, 1995; Lewis and Sappington, 1995) studied the choice between monitoring inputs and outputs. Random auditing, which we find to be approximately optimal, also arises in static, costly state-verification models (Reinganum and Wilde, 1985; Border and Sobel, 1987; Mookherjee and Png, 1989).

Our exercise of varying N and K together with δ relates to the literature on repeated games with frequent actions, where the information structure varies together with δ in a specific, parametric manner (Abreu, Milgrom, and Pearce, 1991; Fudenberg and Levine, 2007, 2009; Sannikov and Skrzypacz, 2007, 2010; Rahman, 2014). The most closely-related results here are Sannikov and Skrzypacz’s (2007) theorem on the impossibility of collusion in duopoly with frequent actions and Brownian noise, as well as a similar result by Fudenberg and Levine (2007). These results are related to our anti-folk theorem for SSE, as we explain in Section 4, but are not very related to our results for PPE.

Finally, in earlier work (Sugaya and Wolitzky, 2021) we studied the relation among N , K , and δ in a repeated random-matching game with private monitoring and incomplete information, where each player is “bad” with some probability. While there are superficial similarities in some of the results, the logic of that model is very different: society has enough information to determine which players are bad after a single period of play, but this information is decentralized, so the question is whether the players can aggregate information fast enough to ensure that it pays to be good. In contrast, the current paper involves complete information and public signals with i.i.d. noise, so the issue is the precision of public monitoring rather than the speed of information aggregation.

2 Model

Stage Games. A stage game $G = (I, \mathcal{A}, \hat{u})$ consists of a finite set of players $I = \{1, \dots, N\}$, a finite action set \mathcal{A}_i for each player $i \in I$, and a payoff function $\hat{u}_i : \mathcal{A} \rightarrow \mathbb{R}$ for each $i \in I$ (where $\mathcal{A} = \prod_{i \in I} \mathcal{A}_i$ and $\hat{u}(a) = (\hat{u}_i(a))_{i \in I}$ for $a \in \mathcal{A}$). Let $M_i = |\mathcal{A}_i|$. We assume that there is some independent noise in the implementation of players’ actions, so that, whenever a player i “intends” to take an action $a_i \in \mathcal{A}_i$, the realized distribution of her action is $(1 - \varepsilon) a_i + (\varepsilon / (M_i - 1)) \sum_{a'_i \in \mathcal{A}_i} a'_i$, where $\varepsilon > 0$ is the *noise level*. (The noise level could

easily be allowed to vary across players and actions, at the cost of slightly more complicated notation.) Note that, by appropriately choosing a distribution of intended actions $\alpha_i \in \Delta \mathcal{A}_i$, player i can implement any action distribution $\hat{\alpha}_i \in \Delta \mathcal{A}_i$ satisfying $\hat{\alpha}_i(\hat{a}_i) \geq \varepsilon / (M_i - 1)$ for all $\hat{a}_i \in \mathcal{A}_i$.⁶ Given a stage game G and noise level ε , we define the *expected payoff function* $u_i : \mathcal{A} \rightarrow \mathbb{R}$ for each $i \in I$ by letting $u_i(a)$ equal player i 's expected payoff, with payoff function \hat{u}_i and noise ε , at intended action profile a . We assume throughout that $\varepsilon \leq 1/2$ and $M_i \geq 2$ for all $i \in I$.

Sets of Payoffs. Given a stage game G and noise level ε (which induce expected payoff functions $(u_i)_{i \in I}$), we define the sets of feasible payoffs, feasible and individually rational payoffs, and payoffs consistent with “almost-myopic” play. The feasible payoff set is $\mathcal{V} = \text{co} \{ \mathbf{v} \in \mathbb{R}^N : \mathbf{v} = u(a) \text{ for some } a \in \mathcal{A} \}$, where co denotes convex hull. Player i 's minmax payoff is $\underline{v}_i = \min_{\alpha_{-i} \in \prod_{j \neq i} \Delta \mathcal{A}_j} \max_{a_i \in \mathcal{A}_i} u_i(a_i, \alpha_{-i})$, where we extend payoff functions from pure to mixed actions as usual. The feasible and individually rational payoff set is $\mathcal{V}^* = \{ \mathbf{v} \in \mathcal{V} : v_i \geq \underline{v}_i \text{ for all } i \in I \}$. For each $\mathbf{v} \in \mathbb{R}^N$ and $\eta > 0$, let $\mathcal{C}_{\mathbf{v}}(\eta) = \prod_{i \in I} [v_i - \eta, v_i + \eta]$, and let $\mathcal{C}(\eta) = \{ \mathbf{v} \in \mathbb{R}^N : \mathcal{C}_{\mathbf{v}}(\eta) \subseteq \mathcal{V}^* \}$. That is, $\mathcal{C}(\eta)$ is the set of points \mathbf{v} such that the N -dimensional cube with center \mathbf{v} and side-length 2η lies entirely within \mathcal{V}^* .

Given a mixed action profile $\alpha \in \Delta^* \mathcal{A} := \prod_i \Delta \mathcal{A}_i$, player i 's maximum static deviation gain at α is $d_i(\alpha) = \max_{a_i \in \mathcal{A}_i} u_i(a_i, \alpha_{-i}) - u_i(\alpha)$. For each $\eta > 0$, define

$$\begin{aligned} \mathcal{A}(\eta) &= \left\{ \alpha \in \Delta^* \mathcal{A} : \frac{1}{N} \sum_{i \in I} d_i(\alpha) \leq \eta \right\}, \\ \mathcal{M}(\eta) &= \text{co} \{ \mathbf{v} \in \mathbb{R}^N : \mathbf{v} = u(\alpha) \text{ for some } \alpha \in \mathcal{A}(\eta) \}, \quad \text{and} \\ \bar{\mathcal{M}}(\eta) &= \{ \mathbf{v} \in \mathbb{R}^N : \|\mathbf{v} - \mathbf{v}'\| \leq \eta \text{ for some } \mathbf{v}' \in \mathcal{M}(\eta) \}, \end{aligned}$$

where $\|\cdot\|$ is the Euclidean norm in \mathbb{R}^N . That is, $\mathcal{A}(\eta)$ is the set of mixed actions where the per-player average deviation gain is no more than η (i.e., the set of mixed actions consistent with almost-myopic play), $\mathcal{M}(\eta)$ is the convex hull of the set of payoff vectors attained by mixed actions in $\mathcal{A}(\eta)$ (i.e., the set of payoff vectors consistent with almost-myopic play), and $\bar{\mathcal{M}}(\eta)$ is the set of payoff vectors within Euclidean distance η of a point in $\mathcal{M}(\eta)$.

⁶To do so, she sets $\alpha_i(a_i) = (\hat{\alpha}_i(a_i) - \varepsilon / (M_i - 1)) / (1 - \varepsilon M_i / (M_i - 1))$ for all $a_i \in \mathcal{A}_i$.

To preview, our anti-folk theorem will provide conditions under which every equilibrium payoff vector is contained in $\bar{\mathcal{M}}(\eta)$, while our folk theorem will provide conditions under which every payoff vector in $\mathcal{C}(\eta)$ arises in an equilibrium. In the Online Appendix, we consider a canonical public-goods game where each player chooses *Contribute* or *Don't Contribute*, and a player's payoff is the fraction of players who contribute less a constant $c \in (0, 1)$ (independent of N) if she contributes herself; in this game, we show that for every $v \in (0, 1 - c)$, there exists $\eta > 0$ such that the symmetric payoff vector where all players receive payoff v lies in $\mathcal{C}(\eta)$ for all N .

Repeated Games. In a repeated game, a stage game is played repeatedly in periods $t = 1, 2, \dots$. Before taking actions in period t , the players observe the outcome of a public randomization device z_t drawn from the uniform distribution over $[0, 1]$. After taking actions in period t , the players observe a public signal $y_t \in \mathcal{Y}$ (where the set of possible signal realizations \mathcal{Y} is assumed to be finite) drawn from a probability distribution $p(\cdot | \hat{a}_t) \in \Delta Y$ that depends only on the realized actions in the current period, \hat{a}_t . The pair (\mathcal{Y}, p) is the *monitoring structure*. Players discount payoffs with a common discount factor $\delta \in (0, 1)$.

A repeated game is thus described by a tuple $\Gamma = (G, \varepsilon, \delta, \mathcal{Y}, p)$. We consider sequences of repeated games indexed by $l \in \mathbb{N}$. Thus, $I, \mathcal{A}, \hat{u}, \varepsilon, \delta, \mathcal{Y}$, and p (as well as the parameter K introduced later on, which measures signal information) all implicitly depend on l , although we usually suppress this dependence to simplify notation. Throughout, we restrict attention to sequences $(\Gamma)^l$ with uniformly bounded payoffs, a uniformly bounded number of actions for each player, and uniformly bounded noise: there exists $\bar{u} \in \mathbb{R}_+$ such that $2|u_i^l(a)| \leq \bar{u}$ for all $l \in \mathbb{N}$, $i \in I^l$, and $a \in A^l$; there exists $M \in \mathbb{N}$ such that $M_i^l \leq M$ for all $l \in \mathbb{N}$ and $i \in I^l$; and there exists $\underline{\varepsilon} > 0$ such that $\varepsilon^l \geq \underline{\varepsilon}$ for all $l \in \mathbb{N}$. Note that the difference between any two stage games payoffs is bounded by \bar{u} : for example, $d_i^l(\alpha) \leq \bar{u}$ for all $l \in \mathbb{N}$, $i \in I^l$ and $\alpha \in \Delta^* A^l$.

Histories, Strategies, and Equilibria. A *history* h_i^t for player i at the beginning of period t takes the form $h_i^t = ((z_\tau, a_{i,\tau}, \hat{a}_{i,\tau}, y_\tau)_{\tau=1}^{t-1}, z_t)$, with $h_i^1 = z_1$. Thus, players observe their own intended and realized actions (though obviously a player has no reason to condition her play on her past intended actions). A *strategy* σ_i for player i maps histories

h_i^t to distributions of intended actions $\Delta\mathcal{A}_i$, for each t . Let $\mathbf{v}^\sigma \in \mathcal{V}$ denote the (expected, discounted) payoff vector induced by strategy profile σ .

A *public history* h^t at the beginning of period t takes the form $h^t = ((z_\tau, y_\tau)_{\tau=1}^{t-1}, z_t)$. A strategy σ_i for player i is *public* if it depends on h_i^t only through its public component h^t . A *perfect public equilibrium* (PPE) is a profile of public strategies that, beginning at any period t and any public history h^t , forms a Nash equilibrium from that period on.⁷ Let $E(\Gamma) \subseteq \mathbb{R}^N$ denote the set of PPE payoff vectors in repeated game Γ : that is, the set of vectors $\mathbf{v} \in \mathbb{R}^N$ such that $\mathbf{v} = \mathbf{v}^\sigma$ for some PPE σ .

All of our results concern PPE rather than more general sequential equilibria. Restricting attention to PPE in repeated games with public monitoring is a common practice, which is usually justified by the fact that such equilibria have a tractable recursive structure and are permissive enough to yield a folk theorem under appropriate statistical conditions (as shown by FLM). In the current context, there is an additional reason to restrict attention to PPE, which is subtle but is crucial for our approach: we are interested in how much information society needs to support cooperation, and in a PPE “society’s information” is naturally measured by the informativeness of the public signal Y about the realized actions \hat{A} . If instead players used more complex strategies that additionally depend on their own past actions (so-called “private strategies,” as in Kandori and Obara (2006)), each player would have different beliefs about the history of play, and it would not be obvious how to measure the total amount of information in the system.⁸

Mutual Information, Signal-Space Cardinality, and Random Auditing. A standard measure of the informativeness of Y about \hat{A} is the *mutual information* between these random variables. Given a joint probability distribution on $\mathcal{Y} \times \mathcal{A}$ (the set of possible real-

⁷As usual, this definition allows players to consider deviations to arbitrary, non-public strategies; but such deviations are irrelevant because, whenever a player’s opponents use public strategies, she has a public strategy that is a best response.

⁸One can imagine different ways of measuring “decentralized societal information.” We offer one suggestion in the Conclusion.

izations of Y and \hat{A}), this quantity is defined as

$$\mathbf{I}(Y, \hat{A}) = \sum_{y \in \mathcal{Y}} \sum_{\hat{a} \in \mathcal{A}} \Pr(y, \hat{a}) \log \frac{\Pr(y, \hat{a})}{\Pr(y) \Pr(\hat{a})}.$$
⁹

Mutual information measures the expected reduction in uncertainty (entropy) about \hat{A} that results from observing Y .¹⁰ This is an endogenous object in our model, as it depends on the distribution of \hat{A} (and hence on the players' strategies) in addition to $p(\cdot|\hat{a})$. Given a mixed action profile $\alpha \in \Delta^* \mathcal{A}$, we write $\mathbf{I}(Y, \hat{A}|\alpha)$ for the mutual information between Y and \hat{A} under the joint distribution on $\mathcal{Y} \times \mathcal{A}$ that results when the players' intended actions are distributed according to α and then, given the resulting realized actions \hat{a} , the signal is distributed according to $p(\cdot|\hat{a})$. Note that, for any $\alpha \in \Delta^* \mathcal{A}$, the entropy of \hat{A} given α is of order N , due to the independent noise in implementing the players' actions.¹¹

Given a stage game G , a noise level ε , and a monitoring structure (\mathcal{Y}, p) , for any number $K \geq 0$ let $\Delta_K^* \mathcal{A} \subseteq \Delta^* \mathcal{A}$ denote the set of mixed action profiles α that satisfy $\mathbf{I}(Y, \hat{A}|\alpha) \leq K$. That is, $\alpha \in \Delta_K^* \mathcal{A}$ if, when the players' intended actions are distributed according to α , in expectation observing y reduces uncertainty about the players' realized actions by at most K . Given a repeated game $\Gamma = (G, \varepsilon, \delta, \mathcal{Y}, p)$ and a number K , let $E(\Gamma, K) \subseteq E(\Gamma)$ denote the set of vectors $\mathbf{v} \in \mathbb{R}^N$ such that $\mathbf{v} = \mathbf{v}^\sigma$ for some PPE σ satisfying $\sigma(h^t) \in \Delta_K^* \mathcal{A}$ for every public history h^t (or equivalently, $\mathbf{I}(Y, \hat{A}|\sigma(h^t)) \leq K$ for every public history h^t). That is, $E(\Gamma, K)$ is the set of payoff vectors that are attained in a PPE where, at any period and any history, observing y reduces uncertainty about the realized current-period actions by at most K . Note that, letting $\Delta_\infty^* \mathcal{A} = \Delta^* \mathcal{A}$, we have $E(\Gamma, \infty) = E(\Gamma)$.

Our anti-folk theorem (Theorem 1.1) gives conditions under which the set $E(\Gamma, K)$ is “small” under any monitoring structure. For deriving an anti-folk theorem (a negative result), it is advantageous that imposing an upper bound on mutual information places few restriction on the monitoring structure. However, readers who prefer assumptions about the

⁹In this paper, all logarithms are base e .

¹⁰And also the expected reduction in uncertainty about Y that would result from observing \hat{A} , since $\mathbf{I}(Y, \hat{A}) = \mathbf{I}(\hat{A}, Y)$. Of course, our players observe Y but not \hat{A} .

¹¹For example, if $M_i = M$ for all $i \in I$, then the entropy of \hat{A} attains its minimum value of $N((1 - \varepsilon) \log(1/(1 - \varepsilon)) + \varepsilon \log((M - 1)/\varepsilon))$ when all players take pure strategies, and it attains its maximum value of $N \log M$ when all players mix uniformly over all actions.

monitoring structure that are independent of the players' strategies may note that $\mathbf{I}(Y, \hat{A})$ is always bounded by the entropy of Y , which in turn is bounded by $\log |\mathcal{Y}|$ (Theorem 2.6.3 of Cover and Thomas, 2006; henceforth "CT"), so that $\Delta_K^* \mathcal{A} = \Delta^* \mathcal{A}$ whenever $\log |\mathcal{Y}| \leq K$. Therefore, the same conditions which guarantee that $E(\Gamma, K)$ is small under any monitoring structure also guarantee that $E(\Gamma)$ (i.e., the entire PPE payoff set) is small under any monitoring structure where $\log |\mathcal{Y}| \leq K$ (see Corollary 2.1). Our anti-folk theorem thus improves on earlier results that impose bounds on $|\mathcal{Y}|$ (Fudenberg, Levine, and Pesendorfer, 1998; al-Najjar and Smorodinsky, 2000, 2001), both because bounding mutual information is less restrictive than bounding $|\mathcal{Y}|$, and because (as we will see) our results are also stronger for a given bound on $|\mathcal{Y}|$.

In contrast, our folk theorem (Theorem 1.2) gives conditions under which the set $E(\Gamma, K)$ is "large" whenever the monitoring structure is given by *random auditing*, where $\lfloor K/\log M \rfloor$ players are selected uniformly at random and the public signal reveals their identities and their realized actions. For deriving a folk theorem (a positive result), limiting attention to a restrictive class of monitoring structures again yields a stronger result. Under random auditing, $\log |\mathcal{Y}| > K$ because the signal reveals the monitored players' identities as well as their actions; however, because the monitored players are selected uniformly at random, it remains true that $\Delta_K^* \mathcal{A} = \Delta^* \mathcal{A}$, and hence that $E(\Gamma, K) = E(\Gamma)$. To see this formally, write $\mathcal{Y} = (\mathcal{Y}_1, \mathcal{Y}_2)$, where \mathcal{Y}_1 consists of all subsets $S \subseteq I$ with $|S| = \lfloor K/\log M \rfloor$, and \mathcal{Y}_2 consists of all vectors of realized actions of the form $(\hat{a}_i)_{i \in S}$ for $S \subseteq I$ with $|S| = \lfloor K/\log M \rfloor$. Under random auditing, $y_1 \in \mathcal{Y}_1$ is chosen uniformly at random, and $y_2 = (\hat{a}_i)_{i \in y_1}$, so

$$\mathbf{I}(Y, \hat{A}) = \sum_{y_1 \in \mathcal{Y}_1} \frac{1}{|\mathcal{Y}_1|} \sum_{(\hat{a}_i)_{i \in y_1}} \Pr((\hat{a}_i)_{i \in y_1}) \log \frac{1}{\Pr((\hat{a}_i)_{i \in y_1})} \leq \log \left(\max_{y_1 \in \mathcal{Y}_1} |\mathcal{A}_{i \in y_1}| \right) \leq K.$$

3 Perfect Public Equilibria in General Games

This section presents our main result, Theorem 1, which gives an anti-folk theorem under any monitoring structure when $(1 - \delta)N/K \rightarrow \infty$ and a folk theorem under random auditing when $(1 - \delta)N \log(N)/K \rightarrow 0$. Our anti-folk theorem combines probabilistic arguments with repeated-game analysis in the style of APS and FLM. To understand this result and

its relation to the literature, it is helpful to start with a “warm-up” anti-folk theorem that isolates the consequences of pivotality or “influence” considerations. The main anti-folk theorem combines these considerations with an APS/FLM-type analysis, and the folk theorem (Theorem 1.2) shows that the resulting condition on N , δ , and K is tight up to $\log(N)$ slack.

3.1 “Warm-Up” Anti-Folk Theorem

Our first result says that if $(1 - \delta) \sqrt{N/K} \rightarrow \infty$ along a sequence of repeated games, then eventually (along the sequence), in every PPE where the mutual information between the signal and the action profile never exceeds K , play is approximately myopic at every history.

Proposition 1 *Consider a sequence of repeated games Γ and numbers K satisfying $(1 - \delta) \sqrt{N/K} \rightarrow \infty$. For any $\eta > 0$, there exists $\bar{l} > 0$ such that, for every $l \geq \bar{l}$, every PPE σ in Γ^l such that $\sigma(h^t) \in \Delta_K^* \mathcal{A}$ for every period t and history h^t , and every history h^t , we have $\sigma(h^t) \in \mathcal{A}(\eta)$.*

Since $\Delta_K^* \mathcal{A} = \Delta^* \mathcal{A}$ whenever $\log |\mathcal{Y}| \leq K$, the following corollary is immediate.

Corollary 1 *Consider a sequence of repeated games Γ and numbers K satisfying $\log |\mathcal{Y}| \leq K$ and $(1 - \delta) \sqrt{N/K} \rightarrow \infty$. For any $\eta > 0$, there exists $\bar{l} > 0$ such that, for every $l \geq \bar{l}$, every PPE σ in Γ^l , and every history h^t , we have $\sigma(h^t) \in \mathcal{A}(\eta)$.*

The proof of Proposition 1 (as well as the proof of our main anti-folk theorem, Theorem 1.1) relies on the notion of a player being “constrained” by noise to take a particular action. Recall that if player i ’s intended action is a_i , then her realized action is $\hat{a}_i = a_i$ with probability $1 - \varepsilon$ and is $\hat{a}_i = a'_i$ for each $a'_i \neq a_i$ with probability $\varepsilon / (M_i - 1)$. Define the event that player i is *constrained* as follows: if player i ’s intended action is a_i , (i) with probability $1 - \varepsilon M_i / (M_i - 1)$, she is *unconstrained*, and then her realized action is $\hat{a}_i = a_i$ with conditional probability 1, and (ii) for each $a'_i \in \mathcal{A}_i$ (including $a'_i = a_i$), she is *constrained to play a'_i* with probability $\varepsilon / (M_i - 1)$, and in this case her realized action is $\hat{a}_i = a'_i$ with conditional probability $1 - \varepsilon$ and is each $\hat{a}_i \neq a'_i$ with conditional probability $\varepsilon / (M_i - 1)$. With this notation, the total probability that a player’s realized action equals her intended

action a_i is decomposed as

$$\underbrace{1 - \frac{M_i}{M_i - 1}\varepsilon}_{\Pr(\text{unconstrained})} + \underbrace{\frac{\varepsilon}{M_i - 1}}_{\Pr(\text{constrained to play } a_i)} (1 - \varepsilon) + \underbrace{\frac{(M_i - 1)\varepsilon}{M_i - 1}}_{\Pr(\text{constrained to play some } a'_i \neq a_i)} \left(\frac{\varepsilon}{M_i - 1} \right) = 1 - \varepsilon,$$

and the $\varepsilon/(M_i - 1)$ probability that her realized action equals each $a'_i \neq a_i$ can be decomposed similarly.

We will make use of the following simple lemma. In what follows, $\Pr(\hat{a}|\alpha)$ is the probability that the realized action profile equals $\hat{a} \in \mathcal{A}$ when the (possibly mixed) intended action profile equals $\alpha \in \Delta^* \mathcal{A}$.

Lemma 1 *Fix a period t and an action profile $\tilde{a} \in \mathcal{A}$. For each player i , let ω_i denote the indicator function for the event that i is constrained to play \tilde{a}_i in period t .*

1. *For each player i , action $a_i \in \mathcal{A}_i$, opposing action profile $\alpha_{-i} \in \prod_{j \neq i} \mathcal{A}_j$, and realized action profile $\hat{a} \in \mathcal{A}$, we have*

$$\Pr(\hat{a}|a_i, \alpha_{-i}) - \Pr(\hat{a}|\tilde{a}_i, \alpha_{-i}) = \left(1 - \frac{\varepsilon}{M_i - 1}\right) (\Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 0) - \Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 1)). \quad (1)$$

2. *We have*

$$\mathbf{I}(Y; (\omega_i)_{i \in I}) \leq \mathbf{I}(Y; \hat{A}). \quad (2)$$

Proof. For part 1, recall that $\Pr(\omega_i = 1) = \varepsilon/(M_i - 1)$. By construction, we have

$$\Pr(\hat{a}|a_i, \alpha_{-i}) = \left(1 - \frac{\varepsilon}{M_i - 1}\right) \Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 0) + \frac{\varepsilon}{M_i - 1} \Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 1).$$

Since the distribution of realized actions is the same when player i deviates to \tilde{a}_i as when she is constrained to play \tilde{a}_i , we have $\Pr(\hat{a}|\tilde{a}_i, \alpha_{-i}) = \Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 1)$. Hence,

$$\begin{aligned} \Pr(\hat{a}|a_i, \alpha_{-i}) - \Pr(\hat{a}|\tilde{a}_i, \alpha_{-i}) &= \Pr(\hat{a}|a_i, \alpha_{-i}) - \Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 1) \\ &= \left(1 - \frac{\varepsilon}{M_i - 1}\right) (\Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 0) - \Pr(\hat{a}|a_i, \alpha_{-i}, \omega_i = 1)). \end{aligned}$$

Part 2 is a consequence of the data-processing inequality (CT, Theorem 2.8.1), since Y and $(\omega_i)_{i \in I}$ are independent conditional on \hat{A} . ■

Proposition 1 now follows fairly easily from Lemma 3 of Sugaya and Wolitzky (2021), which bounds the “average influence” of N players’ types on the public signal. In the proof of Proposition 1, a player’s “type” will be the indicator function for the event that she is constrained to play a particular action in a given period. (The notion of a player’s type and the way the lemma is used were different in our earlier paper.¹²) We restate the lemma here in a slightly different form.¹³ Let $(x)_+ = \max\{x, 0\}$.

Lemma 2 *Let $\omega_1, \dots, \omega_N$ be i.i.d. binary random variables with $\Pr(\omega_i = 1) = \varepsilon \leq 1/2$ for all $i \in I$, and let Y be a random variable satisfying $\mathbf{I}(Y; (\omega_i)_{i \in I}) \leq K$. Then*

$$\frac{1}{N} \sum_{i \in I} \sum_{y \in \mathcal{Y}} (\Pr(y|\omega_i = 0) - \Pr(y|\omega_i = 1))_+ \leq \sqrt{\frac{K}{\varepsilon N}}. \quad (3)$$

Lemma 2 follows from Pinsker’s inequality (CT, Lemma 11.6.1) and the Cauchy-Schwarz inequality, together with some elementary manipulations. A simple intuition can be given in the special case where $\mathcal{Y} = \{0, 1\}^K$ (in which case $\mathbf{I}(Y; (\omega_i)_{i \in I}) \leq K$ for any joint distribution of Y and $(\omega_i)_{i \in I}$). In this case, average influence (the left-hand side of (3)) is maximized by dividing the N players into K equal-sized groups, and specifying that the n^{th} component of the signal Y takes value 1 iff $\omega_i = 1$ for at least $\varepsilon N/K$ of the players in the n^{th} group: under this “majority rule” scheme, each player is pivotal with probability approximately $\sqrt{K/(\varepsilon N)}$ (the right-hand side of (3)). The results of Fudenberg, Levine and Pesendorfer (1998) and Al-Najjar and Smorodinsky (2000, 2001) are based on a similar bound for average influence in the $K = 1$ case, which can be easily be extended by induction to give a bound of $K/\sqrt{\varepsilon N}$ in general (which is weaker than (3)). This bound can be used to

¹²There, the influence-bounding lemma was used to analyze players’ incentives to follow a different payoff type’s strategy for the entire game, so a player’s “type” was their payoff type, and the “signal” was the infinite sequence of the player’s partners’ histories over the course of the game. Here, we analyze different incentive constraints in every period, so a player’s “type” reflects whether they are hit by noise in the current period, and the “signal” is the current public signal. The argument in the earlier paper is thus quite different from Proposition 1 in the current paper, and it is even more distant from our main anti-folk theorem, Theorem 1.

¹³The statement in our earlier paper assumed that $\mathcal{Y} = \{0, 1\}^K$; however, the proof requires only that $\mathbf{I}(Y; (\omega_i)_{i \in I}) \leq K$.

prove Proposition 1 under the stronger assumption that $(1 - \delta) \sqrt{N}/K \rightarrow \infty$, by the same reasoning as that by which (3) is used to prove Proposition 1 under the stated assumption.¹⁴

Proof of Proposition 1. Fix a PPE σ such that $\sigma(h^t) \in \Delta_K^* \mathcal{A}$ for every period t and history h^t . Since σ is a PPE, for every player i , period t , history h^t , and action $a_i \in \mathcal{A}_i$, we have

$$\begin{aligned} & (1 - \delta) u_i(\sigma(h^t)) + \delta \sum_{y \in \mathcal{Y}} \Pr(y|\sigma(h^t)) U_{i,t+1}^\sigma(h^t, y) \\ & \geq (1 - \delta) u_i(a_i, \sigma_{-i}(h^t)) + \delta \sum_{y \in \mathcal{Y}} \Pr(y|a_i, \sigma_{-i}(h^t)) U_{i,t+1}^\sigma(h^t, y), \end{aligned}$$

where $U_{i,t+1}^\sigma(h^t, y)$ denotes player i 's continuation payoff from period $t + 1$ at history (h^t, y) under strategy profile σ . Since $u_i(a_i, \sigma_{-i}(h^t)) - u_i(\sigma(h^t)) \leq d_i(\sigma(h^t))$ and $|U_{i,t+1}^\sigma(h^{t+1}) - U_{i,t+1}^\sigma(\tilde{h}^{t+1})| \leq \bar{u}$ for any two histories h^{t+1} and \tilde{h}^{t+1} , this inequality implies that

$$d_i(\sigma(h^t)) \leq \frac{\delta}{1 - \delta} \sum_{y \in \mathcal{Y}} (\Pr(y|\sigma(h^t)) - \Pr(y|a_i, \sigma_{-i}(h^t)))_+ \bar{u},$$

Summing over players i , for any action profile $a \in \mathcal{A}$, we have

$$\sum_{i \in I} d_i(\sigma(h^t)) \leq \frac{\delta}{1 - \delta} \sum_{i \in I} \sum_{y \in \mathcal{Y}} (\Pr(y|\sigma(h^t)) - \Pr(y|a_i, \sigma_{-i}(h^t)))_+ \bar{u}. \quad (4)$$

Now, fixing the action profile a , let ω_i denote the indicator function for the event that i is constrained to take a_i in period t . Since the signal distribution depends only on the realized actions, for each $y \in \mathcal{Y}$, (1) implies that

$$\begin{aligned} (\Pr(y|\sigma(h^t)) - \Pr(y|a_i, \sigma_{-i}(h^t)))_+ &= \left(1 - \frac{\varepsilon}{M - 1}\right) (\Pr^\sigma(y|\omega_i = 0) - \Pr^\sigma(y|\omega_i = 1))_+ \\ &\leq (\Pr^\sigma(y|\omega_i = 0) - \Pr^\sigma(y|\omega_i = 1))_+ \end{aligned} \quad (5)$$

¹⁴Earlier results by Green (1980) and Sabourian (1990) directly assume that the map from distributions of strategies to distributions of outcomes is continuous. This approach not allow quantitative comparisons among N , δ , and K .

Since $\sigma(h^t) \in \Delta_K^* \mathcal{A}$, by (2) we have $\mathbf{I}(Y; (\omega_i)_{i \in I}) \leq \mathbf{I}(Y; \hat{A}) \leq K$, and hence by (3) we have

$$\frac{1}{N} \sum_{i \in I} \sum_{y \in \mathcal{Y}} (\Pr^\sigma(y|\omega_i = 0) - \Pr^\sigma(y|\omega_i = 1))_+ \leq \sqrt{\frac{K}{\varepsilon N}}. \quad (6)$$

By (4), (5), and (6), we have

$$\frac{1}{N} \sum_{i \in I} d_i(\sigma(h^t)) \leq \frac{\delta}{1 - \delta} \sqrt{\frac{K}{\varepsilon N}}.$$

Since $(1 - \delta) \sqrt{N/K} \rightarrow \infty$, for any $\eta > 0$ there exists \bar{l} such that, for every $l \geq \bar{l}$,

$$\frac{\delta}{1 - \delta} \sqrt{\frac{K}{\varepsilon N}} \leq \eta,$$

and hence $\sigma(h^t) \in \mathcal{A}(\eta)$. ■

3.2 Main Results: Folk and Anti-Folk Theorem

We are now ready to state our main result.

Theorem 1 *Consider a sequence of repeated games Γ and numbers K .*

1. (Anti-Folk Theorem) *Suppose that $(1 - \delta) N/K \rightarrow \infty$. For any $\eta > 0$, there exists \bar{l} such that, for every $l \geq \bar{l}$, $E(\Gamma^l, K^l) \subseteq \bar{\mathcal{M}}^l(\eta)$.*
2. (Folk Theorem) *Suppose that (\mathcal{Y}, p) is random auditing, $1 \leq \lfloor K/\log M \rfloor \leq N$, and $(1 - \delta) N \log(N)/K \rightarrow 0$. For any $\eta > 0$, there exists \bar{l} such that, for every $l \geq \bar{l}$, $\mathcal{C}^l(\eta) \subseteq E(\Gamma^l, K^l)$.*

Since $\Delta_K^* \mathcal{A} = \Delta^* \mathcal{A}$ (and hence $E(\Gamma, K) = E(\Gamma)$) whenever $\log |\mathcal{Y}| \leq K$, as well as whenever (\mathcal{Y}, p) is random auditing of $\lfloor K/\log M \rfloor$ players, the following corollary is immediate.

Corollary 2 *Consider a sequence of repeated games Γ and numbers K .*

1. (Anti-Folk Theorem) *Suppose that $\log |\mathcal{Y}| \leq K$ and $(1 - \delta) N/K \rightarrow \infty$. For any $\eta > 0$, there exists \bar{l} such that, for every $l \geq \bar{l}$, $E(\Gamma^l) \subseteq \bar{\mathcal{M}}^l(\eta)$.*

2. (Folk Theorem) *Suppose that (\mathcal{Y}, p) is random auditing, $1 \leq \lfloor K/\log M \rfloor \leq N$, and $(1 - \delta)N \log(N)/K \rightarrow 0$. For any $\eta > 0$, there exists \bar{l} such that, for every $l \geq \bar{l}$, $\mathcal{C}^l(\eta) \subseteq E(\Gamma^l)$.*

We defer the proof of Theorem 1 to the Appendix. Here we explain the main ideas, including why our sufficient condition for the folk theorem requires $\log(N)$ slack.

The proof of the anti-folk theorem starts by assuming that $E \not\subseteq \bar{\mathcal{M}}(\eta)$ (where we suppress the dependence of $E(\Gamma, K)$ on (Γ, K) , and suppress the dependence of all variables on l), and ultimately concludes that $(1 - \delta)N/K$ cannot be too large. The first step in the proof (Lemma 3) is a purely geometric result. It shows that if $E \not\subseteq \bar{\mathcal{M}}(\eta)$ then there exists a payoff vector \mathbf{v}^* on boundary of E and a ball B^* that contains E and is tangent to E at \mathbf{v}^* such that $\mathbf{v}^* \notin \mathcal{M}(\eta)$ and the radius of B^* is “sufficiently small.” This result is not trivial, because the curvature of E at an arbitrary boundary point $\mathbf{v} \notin \bar{\mathcal{M}}(\eta)$ may be very small (or even zero). Nonetheless, we show that the existence of a boundary point $\mathbf{v} \notin \bar{\mathcal{M}}(\eta)$, together with the fact that E and $\mathcal{M}(\eta)$ are convex sets that lie in an N -dimensional cube with side-length \bar{u} , implies that there exists a boundary point outside $\mathcal{M}(\eta)$ (though not necessarily outside $\bar{\mathcal{M}}(\eta)$) where the curvature of E is large enough that a tangent ball of sufficiently small radius exists.

Next, since \mathbf{v}^* is on the boundary of E but outside $\mathcal{M}(\eta)$, the per-player average static deviation gain at an equilibrium action profile that generates \mathbf{v}^* (together with continuation payoffs drawn from E) is at least η , and therefore the product of a measure of the per-player average influence on the public signal and the average change in continuation payoffs must exceed η . A probabilistic argument (Lemma 5), which is similar to Lemma 2 and is likewise based on Pinsker’s inequality, bounds the average influence measure by \sqrt{K}/N .¹⁵ The average continuation payoff movement can be decomposed into the average movement in the direction normal to the boundary of E at \mathbf{v}^* and the direction tangent to the boundary of E at \mathbf{v}^* . We show that both of these terms are of order at most $\sqrt{N/(1 - \delta)}$. For the normal direction, this follows from the standard self-generation and promise-keeping constraints. For the tangent direction, it follows because the curvature of E at \mathbf{v}^* is sufficiently large: that

¹⁵The difference from the $\sqrt{K/N}$ bound of Lemma 2 comes because Lemma 5 involves the L^2 -norm, while Lemma 2 involves the L^1 -norm.

is, the radius on the enclosing tangent ball B^* is “sufficiently small,” where this is defined precisely so that an $O\left(\sqrt{N/(1-\delta)}\right)$ bound obtains. Finally, since the product of the influence bound, \sqrt{K}/N , and the continuation payoff bound in either the normal or tangent direction, a term of order $\sqrt{N/(1-\delta)}$, is of order $\sqrt{K/((1-\delta)N)}$; and we have seen that this product must exceed η ; we can conclude that $(1-\delta)N/K$ cannot be too large.

The proof of the folk theorem is constructive. We view the repeated game as a sequence of T -period blocks of periods, where T is a number proportional to $1/(1-\delta)$. At the beginning of each block, a target payoff vector is determined by public randomization, and with high probability the players take actions throughout the block that deliver the target payoff. Players are randomly audited throughout the block, and the distribution of target payoffs in the next block is determined by the players’ actions in the periods where they were audited, in such a way so as to provide incentives for correct play. Since each player is audited with probability $\lfloor K/\log M \rfloor / N$ in every period, the adjustment to a player’s continuation payoff in the event that she is audited must be of order N/K to provide incentives. So long as each player is audited a small number of times in every block and $1/(1-\delta)$ is much greater than N/K , the total required adjustment to continuation payoffs is small enough that it can be delivered by appropriately specifying the distribution of target payoffs at the start of the next block.¹⁶

The main difficulty in the proof is that the number of times each player is audited in a block is random and potentially very large. With small probability, a player is audited so many times in a block that the distribution of target payoffs for the next block cannot be modified any further. In this case, she can no longer be incentivized to take a non-myopic best response, and all players’ behavior in the current block must change. Thus, if any player is monitored an “abnormal” number of times in a block, all players’ payoffs in that block may be far from the target equilibrium payoffs. Therefore, to prove the folk theorem, we need the length of a block to be long enough that the probability that any player is monitored an abnormal number of times is small. Since each of the N players is

¹⁶The reader may wonder why we cannot use the same proof approach as FLM, which does not involve dividing the game into blocks. FLM’s proof relies the equivalence of the Euclidean norm in \mathbb{R}^N and the L^1 norm. This equivalence is not uniform in N , so their proof does not directly apply when N and δ vary together. We do not know if it is possible to give an alternate proof of Theorem 1.2 that is closer to FLM’s proof.

monitored with probability of order K/N and the length of a block is of order $1/(1-\delta)$, this condition turns out to require that $(1-\delta)N \log(N)/K$ is small. The extra $\log(N)$ term in this expression arises for a similar reason as in the well-known “coupon collector’s problem” in probability theory: if in each period K balls are drawn with replacement from an urn containing N balls, then the expected number of periods before all N balls have been drawn at least once is approximately $N \log(N)/K$; and this is also the number of periods before it becomes unlikely that the realized number of times any ball is drawn differs significantly from its expectation.¹⁷

Let us note a surprising aspect of Theorem 1. The theorem shows that random auditing is approximately optimal, in that it allows a folk theorem whenever non-myopic incentives can be provided (except for the small set of parameters where $(1-\delta)N/K \not\rightarrow \infty$ but $(1-\delta)N \log(N)/K \rightarrow 0$). But random auditing does not maximize pivot probability or influence: average influence under random auditing is of order K/N , while average influence under joint monitoring schemes of the kind discussed in the previous subsection is of a higher order, $\sqrt{K/N}$. Thus, random auditing provides “approximately maximal incentives” but not “approximately maximal influence.” The explanation for this puzzle is that, under the self-generation constraint that continuation payoffs must be drawn from E , there is a tradeoff among the magnitudes of different players’ continuation payoff movements: starting from a point near the boundary of E at which E exhibits significant curvature, continuation payoffs remain in E if a few players’ payoffs move a lot or many players’ payoffs move a little, but not if many players’ payoffs move a lot. In particular, near a boundary point where E is locally like a ball of radius $O(\sqrt{N})$, each player’s payoff can move only about $1/\sqrt{N}$ times as far when everyone’s payoffs move at once. This \sqrt{N} reduction in incentive power exactly offsets the \sqrt{N} gain in influence from monitoring everyone at once. Hence, by monitoring a few players each period but adjusting their continuation payoffs by a large amount, random auditing provides as strong incentives as does any joint monitoring scheme that monitors more players but adjusts their continuation payoffs by less.

Finally, while we have emphasized the case where $N \rightarrow \infty$ along the sequence of repeated games Γ , none of our analysis requires this assumption. Note that our definition of

¹⁷See, e.g., https://en.wikipedia.org/wiki/Coupon_collector%27s_problem.

random auditing assumes that a fixed, integer number of players $\lfloor K/\log M \rfloor$ are monitored each period, but the proof of Theorem 1.2 applies equally when every player is monitored with independent probability $K/\log M$ each period. With this alternate notion of random auditing, the assumption that $\lfloor K/\log M \rfloor \geq 1$ can be dropped, and the proof of Theorem 1.2 shows that, if N is bounded by some \bar{N} and every player is monitored with probability $K/\log M \leq 1$ each period, a folk theorem holds whenever $(1 - \delta)/K \rightarrow 0$. Conversely, for any N , Theorem 1.1 establishes an anti-folk theorem when the mutual information between the signal and the action profile is bounded by K and $(1 - \delta)/K \rightarrow \infty$.

4 Strongly Symmetric Equilibria in Symmetric Games

We now turn to folk and anti-folk theorems for SSE in symmetric games. Our main motivation for studying this restricted class of equilibria is to clarify the gap between the scope for cooperation with individual monitoring and incentives (which was characterized by our main result, Theorem 1) and the scope for cooperation with fully collectivized monitoring and incentives (which is captured by SSE). One also might have expected SSE to perform well in our setting, because collective monitoring maximizes average influence as discussed in Section 3.1; however, this intuition turns out to be wrong, because we will see that optimal SSE are not those that maximize average influence. Finally, some influential papers on collusion have restricted attention to SSE (Green and Porter, 1984; Abreu, Pearce, and Stacchetti, 1986; Athey, Bagwell, and Sanchirico, 2004), so our results may be relevant for analyzing collusion among many firms in industries where this restriction is well-motivated.

A stage game is *symmetric* if $\mathcal{A}_i = \mathcal{A}_j$ for all i, j and, for every i and every permutation ρ of I , we have $u_i(a_1, \dots, a_N) = u_{\rho(i)}(a_{\rho(1)}, \dots, a_{\rho(N)})$. Denote the common action space by \mathcal{A}_0 . A public strategy profile $\sigma = (\sigma_i)_{i \in I}$ in a symmetric game is *strongly symmetric* if $\sigma_i(h^t) = \sigma_j(h^t)$ for every $i \neq j$ and every public history h^t . In this case, we slightly abuse notation by writing $\sigma : h^t \rightarrow \Delta \mathcal{A}_0$. A *strongly symmetric equilibrium* (SSE) is a PPE in strong symmetric strategies.

In a symmetric game, given a mixed action $\alpha_0 \in \Delta^* \mathcal{A}_0$, denote the symmetric action profile where all players take α_0 by $\bar{\alpha}_0 \in \Delta^* \mathcal{A}_0$, denote the profile where player i takes

action $a \in \mathcal{A}_0$ and everyone else takes α_0 by $(a, \bar{\alpha}_0^{-i})$, and denote player i 's payoff at this profile by $u(a, \bar{\alpha}_0^{-i})$. Denote the greatest and smallest symmetric stage game payoffs by $\bar{v} = \max_{a_0 \in \mathcal{A}_0} u(\bar{a}_0)$ and $\underline{v} = \min_{a_0 \in \mathcal{A}_0} u(\bar{a}_0)$, respectively. Denote the greatest and smallest symmetric stage game Nash equilibrium payoffs by \bar{v}^N and \underline{v}^N , respectively. Let $\bar{d} = \min_{a_0: u(\bar{a}_0) = \bar{v}} \max_{a \in \mathcal{A}_0} u(a, \bar{a}_0^{-i}) - \bar{v}$ and $\underline{d} = \min_{a_0: u(\bar{a}_0) = \underline{v}} \max_{a \in \mathcal{A}_0} u(a, \bar{a}_0^{-i}) - \underline{v}$. Given any $\alpha_0 \in \Delta \mathcal{A}_0$, let $d(\alpha_0) = \max_{a \in \mathcal{A}_0} u(a, \bar{\alpha}_0^{-i}) - u(\bar{\alpha}_0)$. Given any $\eta > 0$, let $A_0(\eta) = \{\alpha_0 \in \Delta \mathcal{A}_0 : d(\alpha_0) \leq \eta\}$. Let $\bar{v}(\eta) = \max_{\alpha_0 \in A_0(\eta)} u(\bar{\alpha}_0)$ and let $\underline{v}(\eta) = \min_{\alpha_0 \in A_0(\eta)} u(\bar{\alpha}_0)$. Denote the set of SSE payoffs in repeated game Γ by $E^S(\Gamma)$: note that this set is an interval due to public randomization.

The following is our folk/anti-folk theorem for SSE. Note that the signal information K does not play a role here, except that the folk theorem requires $|\mathcal{Y}| \geq 2$. This is because, since continuation payoffs like on the 45° line in an SSE, there is no advantage to allowing more than two signal realizations. Also, since the folk theorem is the less important part of the result, we content ourselves with a “Nash-threat” result for simplicity; this could be strengthened to a minmax-threat folk theorem if the monitoring structure is rich enough to track deviations from the minmaxing action profile, in addition to deviations from the target action profile (this may require $|\mathcal{Y}| > 2$).

Theorem 2 *Consider a sequence (G, ε, δ) . For a given monitoring structure (\mathcal{Y}, p) , let $\Gamma = (G, \varepsilon, \delta, \mathcal{Y}, p)$.*

1. (Anti-Folk Theorem) *Suppose that there exists $\rho > 0$ such that $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$. For any sequence of monitoring structures (\mathcal{Y}, p) and any $\eta > 0$, there exists \bar{l} such that, for every $l \geq \bar{l}$, $E^S(\Gamma) \subseteq [\underline{v}(\eta), \bar{v}(\eta)]$.*
2. (Folk Theorem) *Suppose that there exists $\rho > 0$ such that $(1 - \delta) \exp(N^{1+\rho}) \rightarrow 0$.*
 - (a) *Suppose that $\limsup \frac{\varepsilon \bar{d}}{1-2\varepsilon} < \liminf \frac{\bar{v} - \underline{v}^N}{2}$. There exists a sequence of monitoring structures (\mathcal{Y}, p) with $|\mathcal{Y}| = 2$ such that the following holds: for any $\eta \in (\limsup \frac{(M_0-1)\varepsilon d}{M_0-1-2\varepsilon}, \liminf \frac{\bar{v} - \underline{v}^N}{2})$, there exists \bar{l} such that, for every $l \geq \bar{l}$, $[\underline{v}^N, \bar{v} - \eta] \subseteq E^S(\Gamma)$.*
 - (b) *Suppose that $\limsup \frac{(1-\varepsilon)d}{1-2\varepsilon} < \liminf \frac{\bar{v}^N - \underline{v}}{2}$. There exists a sequence of monitoring structures (\mathcal{Y}, p) with $|\mathcal{Y}| = 2$ such that the following holds: for any*

$\eta \in (\limsup \frac{(M_0-1)(1-\varepsilon)d}{M_0-1-2\varepsilon}, \liminf \frac{\bar{v}^N - v}{2}),$ there exists \bar{l} such that, for every $l \geq \bar{l}$, $[v + \eta, \bar{v}^N] \subseteq E^S(\Gamma)$.

Since the condition that $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$ for some $\rho > 0$ is very permissive, we view Theorem 2 primarily as an impossibility theorem for large-group cooperation under collective monitoring.

The proof of Theorem 2 is deferred to the Online Appendix. To understand the result, suppose we wish to enforce a symmetric pure action profile (a, \dots, a) , where $d(a) = \eta$. By standard arguments, an optimal SSE takes the form of a “tail test,” where the players are all punished if the number of players n who take action a falls below a threshold n^* . Due to the i.i.d. noise ε , the distribution of n is approximately normal when N is large, with mean $(1 - \varepsilon)N$ and standard deviation $\sqrt{\varepsilon(1 - \varepsilon)N}$. Denote the threshold z -score by $z^* = (n^* - (1 - \varepsilon)N) / \sqrt{\varepsilon(1 - \varepsilon)N}$, let ϕ and Φ denote the standard normal pdf and cdf, and let $x \in [0, \bar{u} / (1 - \delta)]$ denote the size of the penalty when the tail test is failed. We have

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

where the first inequality is incentive-compatibility, and the second says that the expected penalty cannot exceed the maximum difference between any two stage-game payoffs. Dividing the first inequality by the second, we obtain

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

The left-hand side of this inequality is the statistical score of a normal tail test, which is approximately equal to $-z^*$ for $z^* < 0$. Hence, z^* must decrease at least linearly with \sqrt{N} . Since $\phi(z^*)$ decreases exponentially with z^* , and hence exponentially with N , Theorem 2.1 now follows from the above incentive-compatibility condition, which implies that the product of $\phi(z^*) / \sqrt{\varepsilon(1 - \varepsilon)N}$ and $\bar{u} / (1 - \delta)$ must exceed η . Conversely, the proof of Theorem 2.2 is very simple: if $(1 - \delta) \exp(N^{1+\rho}) \rightarrow 0$ for some $\rho > 0$, the desired action profile (a, \dots, a) can be enforced by a tail test with $n^* = 0$, so that the players are collectively punished only if *no one* takes action a .

Theorem 2.1 is related to Proposition 1 of Sannikov and Skrzypacz (2007), which is an anti-folk theorem for SSE in a two-player repeated game where actions are observed with additive, normally distributed noise, with variance proportional to $1/(1-\delta)$. (The interpretation is that players change their actions every Δ units of time, where $\delta = e^{-r\Delta}$ for fixed $r > 0$ and variance is inversely proportional to Δ , for example as a consequence of observing the average increments of a Brownian process.) As a tail test is also optimal in their setting, the reasoning just given implies that incentives can be provided only if $1/(1-\delta)$ increases exponentially with the variance of the noise. Since in their model $1/(1-\delta)$ increases with the variance only linearly, they likewise obtain an anti-folk theorem. Similarly, Proposition 2 of Fudenberg and Levine (2007) is an anti-folk theorem in a game with one long-run player and series of short-run players, where the long-run player's action is observed with additive, normal noise with variance proportional to $1/(1-\delta)^\rho$ for some $\rho > 0$; and their Proposition 3 is a folk theorem when the variance is constant in δ . Theorem 2 suggests that their anti-folk theorem extends whenever the variance asymptotically dominates $(\log 1/(1-\delta))^{1/(1-\rho)}$ for some $\rho > 0$, while their folk theorem extends whenever the variance is asymptotically dominated by $(\log 1/(1-\delta))^{1/(1+\rho)}$ for some $\rho > 0$.¹⁸

5 Conclusion

To help better understand whether and how large communities with limited information can support cooperation, we have considered sequences of repeated games with imperfect public monitoring where the population size N , discount factor δ , and signal information K vary together. Analyzing this problem involves combining probabilistic and information-theoretic ideas with repeated game theory based on dynamic programming and geometry. Our main result is that, except for a “small” range of parameters not covered by our results, the prospects for cooperation are determined by the ratio of the discount rate $r \approx 1 - \delta$ and the per-capita information K/N : a folk theorem holds when the ratio rN/K is small,

¹⁸More broadly, the analysis of tail tests as optimal incentive contracts under normal noise goes back to Mirrlees (1975). The logic of Theorem 2.1 shows that the size of the penalty in a Mirrleesian tail test must increase exponentially with the variance of the noise. We are not aware of references to this point in the literature.

and an anti-folk theorem holds when it is large. Whenever the ratio is small, cooperation can be supported by a simple random auditing scheme, which is somewhat surprising because random auditing does not maximize the players' pivot probability or "influence." If instead monitoring and incentives are carried out "collectively," which we model by restricting attention to strongly symmetric equilibria, cooperation is possible only under much more severe parameter restrictions.

In reality, information in large communities is typically decentralized: different agents have different pieces of information. Extending our model to account for decentralized information would require developing a measure of "society's information" for this case. One possibility (not the only one) would be to modify our model by assuming that the signal Y is observed only by a mediator or "social planner," who then privately recommends an action to each player. It would be interesting to see if such a model yields results similar to those in the public monitoring case.

A Appendix: Proof of Theorem 1

A.1 Proof of Theorem 1.1 (Anti-Folk Theorem)

We suppress the dependence of $E(\Gamma, K)$ on (Γ, K) . We first show that if $E \not\subseteq \bar{\mathcal{M}}(\eta)$ then there exists a ball B^* in \mathbb{R}^N that (i) is tangent to E at a point \mathbf{v}^* that lies outside of $\mathcal{M}(\eta)$, (ii) contains E , and (iii) has a sufficiently small radius r .

For $x \in \mathbb{R}$, let $\text{sign}(x) = 1$ if $x > 0$, 0 if $x = 0$, and -1 if $x < 0$. Given a vector $\mathbf{v} \in \mathbb{R}^N$, let $\text{sign}(\mathbf{v}) = (\text{sign}(v_1), \dots, \text{sign}(v_N)) \in \{-1, 0, 1\}^N$. Note that $\mathbf{v} \cdot \text{sign}(\mathbf{v}) > 0$ for any $\mathbf{v} \neq \mathbf{0}$. For a set $V \subseteq \mathbb{R}^N$, let ∂V denote the relative boundary of V . Denote the set of unit vectors ("directions") in \mathbb{R}^N by $\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^N : \sum_{i \in I} |\lambda_i|^2 = 1\}$. A vector is *normal* to a surface at a point if it is orthogonal to a tangent hyperplane.

Lemma 3 *Suppose that $E \not\subseteq \bar{\mathcal{M}}(\eta)$. There exists a ball B^* of radius r and a payoff vector $\mathbf{v}^* \in \partial E \cap \partial B^* \setminus \mathcal{M}(\eta)$ such that $E \subseteq B^*$ and, letting $\boldsymbol{\lambda}^* \in \Lambda$ denote a normal vector of B^* at \mathbf{v}^* , we have*

$$r\boldsymbol{\lambda}^* \cdot \text{sign}(\boldsymbol{\lambda}^*) \leq \left(\frac{2\bar{u}^2}{\eta} + \bar{u} \right) N. \quad (7)$$




Figure 1: The left panel illustrates how we fix λ , d , and $\hat{\mathbf{v}}$. The center panel illustrates how we fix h , λ^* , α^* , $\tilde{\mathbf{v}}$, \mathbf{v}^* , $\bar{\mathbf{v}}$, $B(0)$, and B^* . The bottom-right panel reproduces the purple triangle from the center panel.

Proof. Fix

$$\lambda \in \arg \max_{\lambda' \in \Lambda} \left(\max_{\mathbf{v}' \in E} \lambda' \cdot \mathbf{v}' - \max_{\mathbf{v}' \in \bar{\mathcal{M}}(\eta)} \lambda' \cdot \mathbf{v}' \right).$$

Since $E \not\subseteq \bar{\mathcal{M}}(\eta)$, we have $\max_{\mathbf{v}' \in E} \lambda \cdot \mathbf{v}' - \max_{\mathbf{v}' \in \bar{\mathcal{M}}(\eta)} \lambda \cdot \mathbf{v}' > 0$. Let

$$d = \max_{\mathbf{v}' \in E} \lambda \cdot \mathbf{v}' - \max_{\mathbf{v}' \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v}' \quad \text{and} \quad h = \max_{\mathbf{v}' \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v}'. \quad (8)$$

Note that $d \leq \bar{u}\sqrt{N}$, because $E, \mathcal{M}(\eta) \subseteq \mathcal{V}$ and the Euclidean distance between any two points in \mathcal{V} is at most $\bar{u}\sqrt{N}$. It is also easily shown that $d \geq \eta\lambda \cdot \text{sign}(\lambda)$. To see this, note that $\max_{\mathbf{v}' \in \bar{\mathcal{M}}(\eta)} \lambda \cdot \mathbf{v}' = \max_{\mathbf{v}' \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v}' + \eta\lambda \cdot \text{sign}(\lambda)$.¹⁹ Hence,

$$d = \max_{\mathbf{v}' \in E} \lambda \cdot \mathbf{v}' - \max_{\mathbf{v}' \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v}' = \max_{\mathbf{v}' \in E} \lambda \cdot \mathbf{v}' - \left(\max_{\mathbf{v}' \in \bar{\mathcal{M}}(\eta)} \lambda \cdot \mathbf{v}' - \eta\lambda \cdot \text{sign}(\lambda) \right) \geq \eta\lambda \cdot \text{sign}(\lambda). \quad (9)$$

Finally, fix $\hat{\mathbf{v}} \in \arg \max_{\mathbf{v}' \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v}'$. The specification of λ , d , and $\hat{\mathbf{v}}$ is illustrated in the left panel of Figure 1.

¹⁹Proof: For each $\mathbf{v} \in \bar{\mathcal{M}}(\eta)$, there exists $\mathbf{v}' \in \mathcal{M}(\eta)$ such that $\max_i |v_i - v'_i| \leq \eta$, and hence $\lambda \cdot \mathbf{v} \leq \lambda \cdot \mathbf{v}' + \eta\lambda \cdot \text{sign}(\lambda)$. Therefore, $\max_{\mathbf{v} \in \bar{\mathcal{M}}(\eta)} \lambda \cdot \mathbf{v} \leq \max_{\mathbf{v} \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v} + \eta\lambda \cdot \text{sign}(\lambda)$. Conversely, for each $\mathbf{v} \in \mathcal{M}(\eta)$ we have $\mathbf{v} + \eta\text{sign}(\lambda) \in \bar{\mathcal{M}}(\eta)$, and hence $\max_{\mathbf{v} \in \bar{\mathcal{M}}(\eta)} \lambda \cdot \mathbf{v} \geq \lambda \cdot (\mathbf{v} + \eta \cdot \text{sign}(\lambda)) = \max_{\mathbf{v} \in \mathcal{M}(\eta)} \lambda \cdot \mathbf{v} + \eta\lambda \cdot \text{sign}(\lambda)$.

For any $\alpha \in \mathbb{R}_+$, let $B(\alpha)$ be the N -dimensional closed ball with center $\hat{\mathbf{v}} - \alpha\boldsymbol{\lambda}$ and radius

$$r = \frac{\bar{u}^2 N + d^2}{d}. \quad (10)$$

Since $d \leq \bar{u}\sqrt{N}$, we have $r \geq \bar{u}\sqrt{N}$, and therefore (since $\hat{\mathbf{v}}$, the center of $B(0)$, is in \mathcal{V}) $B(0) \supseteq \mathcal{V} \supseteq E \cup \mathcal{M}(\eta)$. Note also that $dr \leq 2\bar{u}^2 N$. Let

$$\alpha^* = \min \{ \alpha \in \mathbb{R}_+ : \partial E \cap \partial B(\alpha) \neq \emptyset \},$$

and let $B^* = B(\alpha^*)$. That is, B^* is given by translating $B(0)$ in direction $-\boldsymbol{\lambda}$ until the translated ball becomes tangent to E . Let $\mathbf{v}^* \in \partial E \cap \partial B^*$ (a tangency point between E and B^*), let $\boldsymbol{\lambda}^*$ be the normal vector of B^* at \mathbf{v}^* , and let $\mathbf{o}^* = \hat{\mathbf{v}} - \alpha^*\boldsymbol{\lambda}$ (the center of B^*). The specification of $\boldsymbol{\lambda}^*$, \mathbf{v}^* , $B(0)$, B^* , and α^* are illustrated in the center panel of Figure 1.

As $E \subseteq B^*$ by construction, it remains to show that $\mathbf{v}^* \notin \mathcal{M}(\eta)$ and that (7) holds.

We first show that $\boldsymbol{\lambda} \cdot \mathbf{v}^* > h$, which implies that $\mathbf{v}^* \notin \mathcal{M}(\eta)$. Let

$$\tilde{\mathbf{v}} = \hat{\mathbf{v}} + (r - \alpha^*)\boldsymbol{\lambda}. \quad (11)$$

Since B^* has center $\hat{\mathbf{v}} - \alpha^*\boldsymbol{\lambda}$ and radius r , we have $\tilde{\mathbf{v}} \in \partial B^*$, and $\boldsymbol{\lambda}$ is normal to B^* at $\tilde{\mathbf{v}}$ (as illustrated in the center panel of Figure 1). Hence,

$$\boldsymbol{\lambda} \cdot \tilde{\mathbf{v}} = \max_{\mathbf{v}' \in B^*} \boldsymbol{\lambda} \cdot \mathbf{v}'. \quad (12)$$

Now decompose the vector $\tilde{\mathbf{v}} - \mathbf{v}^*$ into the direction $\boldsymbol{\lambda}$ and an orthogonal direction $\boldsymbol{\lambda}^\perp \in \Lambda$, so that

$$\tilde{\mathbf{v}} - \mathbf{v}^* = \beta\boldsymbol{\lambda} + \gamma\boldsymbol{\lambda}^\perp \text{ for some } \beta, \gamma \in \mathbb{R}^N. \quad (13)$$

(This decomposition is illustrated by the purple triangle in the center and bottom-right panels of Figure 1). We derive bounds for β and γ .

First, since $E \subseteq B^*$ and (12), we have $\boldsymbol{\lambda} \cdot \tilde{\mathbf{v}} = \max_{\mathbf{v}' \in B^*} \boldsymbol{\lambda} \cdot \mathbf{v}' \geq \max_{\mathbf{v}' \in E} \boldsymbol{\lambda} \cdot \mathbf{v}' = d + h$, and therefore

$$\beta = \boldsymbol{\lambda} \cdot (\tilde{\mathbf{v}} - \mathbf{v}^*) \geq d + h - \boldsymbol{\lambda} \cdot \mathbf{v}^*. \quad (14)$$

Second, we note that $\beta \leq r$. To see this, let $\bar{\mathbf{v}} = \mathbf{v}^* + \alpha^* \boldsymbol{\lambda}$. Since $\mathbf{v}^* \in \partial B^*$ and B^* is given by translating $B(0)$ by α^* in direction $-\boldsymbol{\lambda}$, we have $\bar{\mathbf{v}} \in \partial B(0)$. Hence,

$$\hat{\mathbf{v}} - \mathbf{v}^* = (\beta - r + \alpha^*) \boldsymbol{\lambda} + \gamma \boldsymbol{\lambda}^\perp \text{ by (11) and (13), and } \hat{\mathbf{v}} - \bar{\mathbf{v}} = (\beta - r) \boldsymbol{\lambda} + \gamma \boldsymbol{\lambda}^\perp. \quad (15)$$

Since $\mathbf{v}^* \in B(0)$, $\bar{\mathbf{v}} \in \partial B(0)$, and $\hat{\mathbf{v}}$ is the center of $B(0)$, we have $\|\hat{\mathbf{v}} - \mathbf{v}^*\| \leq r = \|\hat{\mathbf{v}} - \bar{\mathbf{v}}\|$. This inequality implies that $|\beta - r + \alpha^*| \leq |\beta - r|$. Since $\alpha^* \geq 0$ by construction, it follows that $\beta \leq r$.

Third, since $\hat{\mathbf{v}}, \mathbf{v}^* \in \mathcal{V}$ satisfy (15), we have

$$\bar{u}\sqrt{N} \geq \|\hat{\mathbf{v}} - \mathbf{v}^*\| \geq |\gamma|. \quad (16)$$

Now we use these bounds to show that $\boldsymbol{\lambda} \cdot \mathbf{v}^* > h$. From (13) and the fact that $\boldsymbol{\lambda}$ is normal to B^* at $\tilde{\mathbf{v}}$, applying the Pythagorean identity to the orange triangle in the center panel of Figure 1, we have

$$\|\mathbf{v}^* - \mathbf{o}^*\|^2 = \|\tilde{\mathbf{v}} - \mathbf{o}^* - \beta \boldsymbol{\lambda}\|^2 + \|\gamma \boldsymbol{\lambda}^\perp\|^2.$$

Note that

$$\begin{aligned} \|\tilde{\mathbf{v}} - \mathbf{o}^* - \beta \boldsymbol{\lambda}\| &= r - \beta \quad (\text{since } \tilde{\mathbf{v}} - \mathbf{o}^* \text{ is parallel to } \boldsymbol{\lambda} \text{ and } \|\tilde{\mathbf{v}} - \mathbf{o}^*\| = r \geq \beta) \\ &\geq r - (d + h) + \boldsymbol{\lambda} \cdot \mathbf{v}^* \quad (\text{by (14)}). \end{aligned}$$

Together with the facts that $\|\mathbf{v}^* - \mathbf{o}^*\| = r$ (because $\mathbf{v}^* \in \partial B^*$) and $\|\gamma \boldsymbol{\lambda}^\perp\| = |\gamma| \leq \bar{u}\sqrt{N}$ (by (16)), we have

$$r^2 \leq (r - (d + h) + \boldsymbol{\lambda} \cdot \mathbf{v}^*)^2 + \bar{u}^2 N.$$

Since $r \geq \beta \geq d + h - \boldsymbol{\lambda} \cdot \mathbf{v}^*$, if $\boldsymbol{\lambda} \cdot \mathbf{v}^* \leq h$ then this inequality would require that

$$r^2 \leq (r - d)^2 + \bar{u}^2 N, \text{ or } r \leq \frac{\bar{u}^2 N + d^2}{2d},$$

which is false under our definition of r , (10). Hence, $\boldsymbol{\lambda} \cdot \mathbf{v}^* > h$.

Next, we establish (7). Since $\boldsymbol{\lambda}^*$ is normal to B^* at \mathbf{v}^* we have

$$r\boldsymbol{\lambda}^* = \mathbf{v}^* - \mathbf{o}^* = \tilde{\mathbf{v}} - \beta\boldsymbol{\lambda} - \gamma\boldsymbol{\lambda}^\perp - \mathbf{o}^* \quad (\text{by (13)}).$$

Since $\boldsymbol{\lambda}$ is normal to B^* at $\tilde{\mathbf{v}}$, we have $r\boldsymbol{\lambda} = \tilde{\mathbf{v}} - \mathbf{o}^*$. Hence, $r\boldsymbol{\lambda}^* = (r - \beta)\boldsymbol{\lambda} - \gamma\boldsymbol{\lambda}^\perp$. We have

$$\begin{aligned} r\boldsymbol{\lambda}^* \cdot \text{sign}(\boldsymbol{\lambda}^*) &= ((r - \beta)\boldsymbol{\lambda} - \gamma\boldsymbol{\lambda}^\perp) \cdot \text{sign}((r - \beta)\boldsymbol{\lambda} - \gamma\boldsymbol{\lambda}^\perp) \\ &= \sum_i |(r - \beta)\lambda_i - \gamma\lambda_i^\perp| \\ &\leq \sum_i (|(r - \beta)\lambda_i| + |\gamma\lambda_i^\perp|) \\ &= |r - \beta| \boldsymbol{\lambda} \cdot \text{sign}(\boldsymbol{\lambda}) + |\gamma| \boldsymbol{\lambda}^\perp \cdot \text{sign}(\boldsymbol{\lambda}^\perp) \\ &\leq r\boldsymbol{\lambda} \cdot \text{sign}(\boldsymbol{\lambda}) + \bar{u}\sqrt{N}\boldsymbol{\lambda}^\perp \cdot \text{sign}(\boldsymbol{\lambda}^\perp) \quad (\text{by } 0 \leq \beta \leq r \text{ and (16)}) \\ &\leq \frac{r}{\eta}d + \bar{u}N \quad (\text{by (9) and } \boldsymbol{\lambda}^\perp \cdot \text{sign}(\boldsymbol{\lambda}^\perp) \leq \sqrt{N}) \\ &\leq \left(\frac{2\bar{u}^2}{\eta} + \bar{u} \right) N \quad (\text{by } dr \leq 2\bar{u}^2N). \end{aligned}$$

■

Recall that E is the set of payoff vectors that are attained in a PPE satisfying $\sigma(h^t) \in \Delta_K^* \mathcal{A}$ for all t, h^t . By standard arguments (e.g., Proposition 7.3.2 of Mailath and Samuelson, 2006), the set E is self-generating: that is, for any $\mathbf{v} \in E$, there exist $\alpha \in \Delta_K^* \mathcal{A}$ and $(\mathbf{w}(y))_y$ such that

Promise-Keeping. $v_i = (1 - \delta)u_i(\alpha) + \delta\mathbb{E}[w_i(Y) | \alpha]$ for all $i \in I$,

Incentive-Compatibility. $\text{supp}(\alpha_i) \subseteq \text{argmax}_{a'_i} (1 - \delta)u_i(a'_i, \alpha_{-i}) + \delta\mathbb{E}[w_i(Y) | a'_i, \alpha_{-i}]$ for all $i \in I$,

Self-Generation. $\mathbf{w}(y) \in E$ for all $y \in \mathcal{Y}$.

Take $\mathbf{v} = \mathbf{v}^*$, and take any corresponding α and $(\mathbf{w}(y))_y$. Since $\mathbf{v}^* \in \partial B^* \setminus \mathcal{M}(\eta)$ and $w(y) \in E \subseteq B^*$ for every y , we have $\mathbf{u}(\alpha) \notin B^*$. Moreover, since $\mathcal{M}(\eta) \subseteq B^*$, this implies that $\alpha \notin \mathcal{A}(\eta)$: that is, $\sum_i d_i(\alpha) > N\eta$.




Figure 2: Decomposition of $\mathbf{x}(y)$ and $\mathbf{x}(y')$ for two signals y and y' , where y satisfies $|\frac{1-\delta}{\delta}\beta(y)| \leq r$, and y' satisfies $|\frac{1-\delta}{\delta}\beta(y')| \geq r$.

By promise-keeping and incentive-compatibility, defining

$$\mathbf{x}(y) = \frac{\delta}{1-\delta} (\mathbf{w}(y) - \mathbf{v}^*), \text{ we have} \quad (17)$$

$$v_i = u_i(\alpha) + \mathbb{E}[x_i(Y) | \alpha] \text{ for all } i \in I, \text{ and} \quad (18)$$

$$\text{supp}(\alpha_i) \subseteq \underset{a'_i}{\text{argmax}} u_i(a'_i, \alpha_{-i}) + \mathbb{E}[x_i(Y) | a'_i, \alpha_{-i}] \text{ for all } i \in I. \quad (19)$$

For each $y \in \mathcal{Y}$, decompose the vector $\mathbf{x}(y)$ into the normal direction $\boldsymbol{\lambda}^*$ and an orthogonal direction $\boldsymbol{\lambda}^\perp(y) \in \Lambda$:

$$\mathbf{x}(y) = \beta(y) \boldsymbol{\lambda}^* + \gamma(y) \boldsymbol{\lambda}^\perp(y) \text{ for some } \beta(y), \gamma(y) \in \mathbb{R}. \quad (20)$$

(This decomposition is illustrated in Figure 2.) Note that since $\mathbf{v}^* = \underset{\mathbf{v} \in B^*}{\text{argmax}} \boldsymbol{\lambda}^* \cdot \mathbf{v}$ and $\mathbf{w}(Y) \in E \subseteq B^*$, we have $\beta(y) \leq 0$ for all y .

We can rewrite (19) as, for each a'_i ,

$$\mathbb{E}[x_i(Y) | \alpha] - \mathbb{E}[x_i(Y) | a'_i, \alpha_{-i}] \geq u_i(a'_i, \alpha_{-i}) - u_i(\alpha). \quad (21)$$

Since the signal distribution depends only on realized actions, by (1) the left-hand side of

(21) equals

$$\begin{aligned} & \left(1 - \frac{\varepsilon}{M_i - 1}\right) (\mathbb{E}[x_i(Y) | \alpha, \omega_i = 0] - \mathbb{E}[x_i(Y) | \alpha, \omega_i = 1]) \\ &= \left(1 - \frac{\varepsilon}{M_i - 1}\right) \sum_y (\Pr(y | \omega_i = 0) - \Pr(y | \omega_i = 1)) x_i(y). \end{aligned}$$

Recalling that $d_i(\alpha) = \max_{a'_i} u_i(a'_i, \alpha_{-i}) - u_i(\alpha)$ and taking $a'_i \in \operatorname{argmax}_{a'_i} u_i(a'_i, \alpha_{-i})$, we have

$$\left(1 - \frac{\varepsilon}{M_i - 1}\right) \sum_y (\Pr(y | \omega_i = 0) - \Pr(y | \omega_i = 1)) x_i(y) \geq d_i(\alpha).$$

Together with (20) and $\sum_i d_i(\alpha) > N\eta$, this implies

$$\begin{aligned} \eta &\leq \frac{1}{N} \sum_i \sum_y |\Pr(y | \omega_i = 0) - \Pr(y | \omega_i = 1)| |\beta(y) \lambda_i^*| \\ &\quad + \frac{1}{N} \sum_i \sum_y |\Pr(y | \omega_i = 0) - \Pr(y | \omega_i = 1)| |\gamma(y) \lambda_i^\perp(y)|. \end{aligned} \quad (22)$$

The following lemma implies that if $(1 - \delta)N/K \rightarrow \infty$ then eventually (22) is violated, and therefore eventually $E \subseteq \bar{\mathcal{M}}(\eta)$. This completes the proof.

Lemma 4 *If $(1 - \delta)N/K \rightarrow \infty$ then*

$$\frac{1}{N} \sum_i \sum_y |\Pr(y | \omega_i = 0) - \Pr(y | \omega_i = 1)| |\beta(y) \lambda_i^*| \rightarrow 0 \quad (23)$$

and

$$\frac{1}{N} \sum_i \sum_y |\Pr(y | \omega_i = 0) - \Pr(y | \omega_i = 1)| |\gamma(y) \lambda_i^\perp(y)| \rightarrow 0. \quad (24)$$

We begin with a lemma concerning mutual information.

Lemma 5 *Fix a mixed action profile $\alpha \in \Delta^* \mathcal{A}$ and a monitoring structure (\mathcal{Y}, p) , and let $(\Pr(y, \omega))_{y \in \mathcal{Y}, \omega \in \{0,1\}^N}$ denote the resulting joint probability distribution on $y \in \mathcal{Y}$ and $\omega = (\omega_1, \dots, \omega_N) \in \{0,1\}^N$. For each $y \in \mathcal{Y}$, define*

$$P_y = \sqrt{\sum_i (\Pr(y | \omega_i = 1) - \Pr(y | \omega_i = 0))^2}.$$

Then

$$P_y \leq \frac{M}{\varepsilon} \sqrt{2 \Pr(y) \sum_i \sum_{\omega_i \in \{0,1\}} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)}}. \quad (25)$$

Proof. Recall that, for each i , $\Pr(\omega_i = 0) \geq \Pr(\omega_i = 1) = \varepsilon / (M_i - 1) \geq \varepsilon / M$. Note that

$$\begin{aligned} |\Pr(y|\omega_i = 1) - \Pr(y|\omega_i = 0)| &= \sum_{\omega_i \in \{0,1\}} \left| \frac{\Pr(y, \omega_i) - \Pr(y) \Pr(\omega_i)}{\Pr(\omega_i)} \right| \\ &\leq \frac{M}{\varepsilon} \sum_{\omega_i \in \{0,1\}} |\Pr(y, \omega_i) - \Pr(y) \Pr(\omega_i)|. \end{aligned}$$

Define $D_i(y) = \sum_{\omega_i \in \{0,1\}} |\Pr(y, \omega_i) - \Pr(y) \Pr(\omega_i)| = \Pr(y) \sum_{\omega_i \in \{0,1\}} |\Pr(\omega_i|y) - \Pr(\omega_i)|$.

By Pinsker's inequality, we have

$$\begin{aligned} \left(\sum_{\omega_i \in \{0,1\}} |\Pr(\omega_i|y) - \Pr(\omega_i)| \right)^2 &\leq 2 \sum_{\omega_i \in \{0,1\}} \Pr(\omega_i|y) \ln \frac{\Pr(\omega_i|y)}{\Pr(\omega_i)} \\ &= \frac{2}{\Pr(y)} \sum_{\omega_i \in \{0,1\}} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)}. \end{aligned}$$

Hence,

$$(D_i(y))^2 \leq \Pr(y)^2 \left(\sum_{\omega_i \in \{0,1\}} |\Pr(\omega_i|y) - \Pr(\omega_i)| \right)^2 \leq 2 \Pr(y) \sum_{\omega_i \in \{0,1\}} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)}.$$

Finally, since

$$(P_y)^2 = \sum_i (\Pr(y|\omega_i = 1) - \Pr(y|\omega_i = 0))^2 \leq \left(\frac{M}{\varepsilon} \right)^2 \sum_i (D_i(y))^2,$$

we have

$$P_y \leq \frac{M}{\varepsilon} \sqrt{2 \Pr(y) \sum_i \sum_{\omega_i \in \{0,1\}} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)}}.$$

■

Finally, we prove Lemma 4.

Proof of Lemma 4. We have

$$\begin{aligned}
& \frac{1}{N} \sum_i \sum_y |\Pr(y|\omega_i = 0) - \Pr(y|\omega_i = 1)| |\beta(y) \lambda_i^*| \\
& \leq \frac{1}{N} \sum_y |\beta(y)| \sqrt{\sum_i (\Pr(y|\omega_i = 0) - \Pr(y|\omega_i = 1))^2} \quad (\text{by Cauchy-Schwarz and } \lambda^* \in \Lambda) \\
& \leq \frac{\sqrt{2}M}{\underline{\varepsilon}N} \sum_y \sqrt{\Pr(y) \beta(y)^2} \sqrt{\sum_i \sum_{\omega_i} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)}} \quad (\text{by (25)}) \\
& \leq \frac{\sqrt{2}M}{\underline{\varepsilon}N} \sqrt{\sum_y \Pr(y) \beta(y)^2} \sqrt{\sum_i \sum_y \sum_{\omega_i} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)}} \quad (\text{by Cauchy-Schwarz}).
\end{aligned} \tag{26}$$

Note that

$$\sum_y \sum_{\omega_i} \Pr(y, \omega_i) \ln \frac{\Pr(y, \omega_i)}{\Pr(y) \Pr(\omega_i)} = \mathbf{I}(Y, \omega_i),$$

the mutual information between Y and ω_i . We have

$$\begin{aligned}
\sum_i \mathbf{I}(Y, \omega_i) & \leq \sum_i \mathbf{I}(Y, \hat{A}_i) \quad (\text{by the data-processing inequality}) \\
& = \mathbf{I}(Y, \hat{A}) \quad (\text{by independence of } (\hat{a}_i)_i \text{ conditional on } \alpha) \\
& \leq K \quad (\text{by } \alpha \in \Delta_K^* \mathcal{A}).
\end{aligned}$$

Together with (26), to establish (23) it suffices to show that

$$\frac{\sqrt{K}}{N} \sqrt{\sum_y \Pr(y) \beta(y)^2} \rightarrow 0. \tag{27}$$

Similarly, to establish (24) it suffices to show that

$$\frac{\sqrt{K}}{N} \sqrt{\sum_y \Pr(y) \gamma(y)^2} \rightarrow 0. \tag{28}$$

We first establish (27). Recall that

$$\begin{aligned}\mathbf{x}(y) &= \frac{\delta}{1-\delta} (\mathbf{w}(y) - \mathbf{v}^*) = \beta(y) \boldsymbol{\lambda}^* + \gamma(y) \boldsymbol{\lambda}^\perp(y) && \text{by (17) and (20), and} \\ \mathbf{v}^* &= \mathbf{u}(\alpha) + \mathbb{E}[\mathbf{x}(Y) | \alpha] && \text{by (18)}.\end{aligned}$$

Since $\boldsymbol{\lambda}^* \cdot \boldsymbol{\lambda}^* = 1$, $\boldsymbol{\lambda}^* \cdot \boldsymbol{\lambda}^\perp(y) = 0$, $\mathbf{u}(\alpha) \in \mathcal{V}$, $\mathbf{v}^* \in \mathcal{V}$, and $\beta(y) \leq 0$, we have

$$\mathbb{E}[\beta(y)] = \boldsymbol{\lambda}^* \cdot (\mathbf{v}^* - \mathbf{u}(\alpha)) \in [-\bar{u}\sqrt{N}, 0]. \quad (29)$$

Moreover, by the Pythagorean theorem, for each y we have

$$\|\mathbf{x}(y)\|^2 = \beta(y)^2 + \gamma(y)^2 \Rightarrow |\beta(y)| \leq \frac{\delta}{1-\delta} \|\mathbf{w}(y) - \mathbf{v}^*\| \leq \frac{\bar{u}\sqrt{N}}{1-\delta}.$$

Note that the distribution of $\beta(y)$ given by $\beta(y) = -\bar{u}\sqrt{N}/(1-\delta)$ with probability $1-\delta$ and $\beta(y) = 0$ with probability δ is a mean-preserving spread of any distribution of $\beta(y)$ that satisfies $\mathbb{E}[\beta(y)] = -\bar{u}\sqrt{N}$ and $\beta(y) \in [-\bar{u}\sqrt{N}/(1-\delta), 0]$ for all y . Hence, since $(\cdot)^2$ is convex, Jensen's inequality implies

$$\sum_y \Pr(y) \beta(y)^2 \leq (1-\delta) \left(-\frac{\bar{u}\sqrt{N}}{1-\delta} \right)^2 + \delta(0)^2 = \bar{u}^2 \frac{N}{1-\delta}.$$

Therefore, we have

$$\frac{\sqrt{K}}{N} \sqrt{\sum_y \Pr(y) \beta(y)^2} \leq \frac{\sqrt{K}}{N} \sqrt{\bar{u}^2 \frac{N}{1-\delta}} = \bar{u} \sqrt{\frac{K}{N(1-\delta)}},$$

which converges to 0 whenever $(1-\delta)N/K \rightarrow \infty$. This establishes (27).

We next establish (28). Applying the Pythagorean theorem to the two red triangles in Figure 2, we have

$$\left(\frac{1-\delta}{\delta} \gamma(y) \right)^2 \leq \begin{cases} r^2 - (r - |\frac{1-\delta}{\delta} \beta(y)|)^2 & \text{if } |\frac{1-\delta}{\delta} \beta(y)| \leq r \\ r^2 - (|\frac{1-\delta}{\delta} \beta(y)| - r)^2 & \text{if } |\frac{1-\delta}{\delta} \beta(y)| \geq r \end{cases} \leq 4r \frac{1-\delta}{\delta} |\beta(y)|.$$

Hence,

$$\sum_y \Pr(y) \gamma(y)^2 \leq 4 \frac{\delta}{1-\delta} \sum_y P(y) r |\beta(y)| \leq 4 \frac{\delta}{1-\delta} r \left| \sum_y \Pr(y) \beta(y) \right|, \quad (30)$$

where the last inequality follows because $\beta(y) \leq 0$ for all y . By (29), we have

$$\left| \sum_y \Pr(y) \beta(y) \right| = \boldsymbol{\lambda}^* \cdot (\mathbf{u}(\alpha) - \mathbf{v}^*) \leq \bar{u} \boldsymbol{\lambda}^* \cdot \text{sign}(\boldsymbol{\lambda}^*).$$

Together with (30) and (7), we have

$$\sum_y \Pr(y) \gamma(y)^2 \leq 4 \frac{\delta}{1-\delta} r \bar{u} \boldsymbol{\lambda}^* \cdot \text{sign}(\boldsymbol{\lambda}^*) \leq 4 \frac{\delta}{1-\delta} \bar{u} \left(\frac{2\bar{u}^2}{\eta} + \bar{u} \right) N.$$

Therefore,

$$\frac{\sqrt{K}}{N} \sqrt{\sum_y P(y) \gamma(y)^2} \leq \frac{\sqrt{K}}{N} \sqrt{4 \frac{\delta}{1-\delta} \bar{u} \left(\frac{2\bar{u}^2}{\eta} + \bar{u} \right) N} = \sqrt{4\delta\bar{u} \left(\frac{2\bar{u}^2}{\eta} + \bar{u} \right) \frac{K}{(1-\delta)N}},$$

which converges to 0 whenever $(1-\delta)N/K \rightarrow \infty$. This establishes (28). ■

A.2 Proof of Theorem 1.2 (Folk Theorem)

A.2.1 Preliminaries

Fix any $\eta > 0$. Recall that $\{i \in y_{1,t}\}$ denotes the event that player i is monitored in period t . Note that

$$\mathbb{E}[\mathbf{1}\{i \in y_{1,t}\}] = \frac{1}{N} \left\lfloor \frac{K}{\lceil \log M \rceil} \right\rfloor, \quad \text{and} \quad \text{Var}(\mathbf{1}\{i \in y_{1,t}\}) = \frac{1}{N} \left\lfloor \frac{K}{\log M} \right\rfloor \left(1 - \frac{1}{N} \left\lfloor \frac{K}{\log M} \right\rfloor \right).$$

To simplify notation, let $\mu = \lfloor K/\log M \rfloor / N$ and let $\nu = (\lfloor K/\log M \rfloor / N) (1 - \lfloor K/\log M \rfloor / N)$. Note that these variables vary with l along a sequence of repeated games, and $\mu \geq 1/N$ because $\lfloor K/\log M \rfloor \geq 1$. For any $T \in \mathbb{N}$ and $\bar{Y} \in \mathbb{R}_+$, Bennett's inequality (Bennett, 1962)

implies that

$$\Pr \left(\sum_{t=1}^T \mathbf{1} \{i \in y_{i,t}\} \geq T\mu + \bar{Y} \right) \leq \exp \left(-T\nu h \left(\frac{\bar{Y}}{T\nu} \right) \right), \quad (31)$$

where $h(x) = (1+x) \log(1+x) - x$ is a function that will appear throughout the proof.

We begin with a lemma that will be used to choose T and \bar{Y} (as a function of l) so that the bound in (31) is sufficiently small and some other inequalities used in the proof are also satisfied.

Lemma 6 *There exists \bar{l} such that, for every $l \geq \bar{l}$, there exist $T \in \mathbb{N}$ and $\bar{Y} \in \mathbb{R}$ that satisfy the following three inequalities:*

$$300\bar{u}N \exp \left(-T\nu h \left(\frac{\bar{Y}}{9\bar{u}T\nu} \right) \right) \leq \eta, \quad (32)$$

$$32\bar{u} \frac{1-\delta}{1-\delta^T} \frac{\bar{Y}}{\mu} \leq \eta, \quad (33)$$

$$8\bar{u} \left(\frac{1-\delta^T}{\delta^T} + \frac{1-\delta}{\delta^T} \left(T + \frac{\bar{Y}}{\mu} \right) \right) \leq \eta. \quad (34)$$

Proof. Fix $\rho > 0$ sufficiently large so that

$$\frac{\log M}{27\bar{u}\rho} < \frac{1}{2}, \text{ and} \quad (35)$$

$$\frac{32\bar{u} \log M}{\rho} < \eta. \quad (36)$$

Next, fix $\beta > 0$ sufficiently small so that

$$\frac{32\bar{u}\beta \log M}{1 - e^{-\rho\beta}} < \eta. \quad (37)$$

$$24\bar{u} (e^{\rho\beta} - 1) < \eta, \quad (38)$$

$$24\bar{u}\rho\beta e^{-\rho\beta} < \eta, \text{ and} \quad (39)$$

$$24\bar{u}\beta e^{\rho\beta} \log M < \eta. \quad (40)$$

Such β exists because, since $\lim_{\beta \rightarrow 0} \beta / (1 - e^{-\rho\beta}) = 1/\rho$, (36) implies that $\lim_{\beta \rightarrow 0} 32\bar{u}\beta \log(M) / (1 - e^{-\rho\beta}) < \eta$. Finally, let $T = \lfloor \rho\beta / (1 - \delta) \rfloor$ and let $\bar{Y} = (\beta / (1 - \delta)) (K/N)$. (Thus, ρ and β are independent of l , while T and \bar{Y} depend on l .) With T and \bar{Y} so-defined as

functions of l , we show that there exists \bar{l} such that conditions (32)–(34) hold for all $l \geq \bar{l}$.

Condition (32): Note that, for any $x > 0$,

$$h(x) \geq \frac{x^2}{2} - \frac{x^3}{3}. \quad (41)$$

This follows because $h(0) = 0$ and, for any $x > 0$,

$$\frac{d}{dx}(h(x)) = \log(1+x) \geq \frac{x}{1+x} \geq x(1-x) = \frac{d}{dx}\left(\frac{x^2}{2} - \frac{x^3}{3}\right).$$

Note also that, for any x ,

$$\frac{d}{dx}\left(\frac{h(x)}{x}\right) = \frac{1}{x} - \frac{\log(1+x)}{x^2} \geq \frac{1}{x} - \frac{x}{x^2} = 0. \quad (42)$$

Hence, letting

$$x = \frac{\log M}{9\bar{u}\rho} = \frac{\bar{Y}(1-\delta)}{9\bar{u}\rho\beta\mu} \text{ and } x' = \frac{\bar{Y}}{9\bar{u}T\nu},$$

and noting that $0 < x < x'$ (since $\frac{\rho\beta}{1-\delta} \geq T$ and $\mu > \nu$), we have

$$\begin{aligned} N \exp\left(-T\nu h\left(\frac{Y}{9\bar{u}T\nu}\right)\right) &= N \exp\left(-\frac{\bar{Y}}{9\bar{u}} \frac{1}{x'} h(x')\right) \\ &\leq N \exp\left(-\frac{\bar{Y}}{9\bar{u}} \frac{1}{x} h(x)\right) \quad (\text{by (42)}) \\ &\leq N \exp\left(-\frac{\bar{Y}}{9\bar{u}} \left(\frac{x}{2} - \frac{x^2}{3}\right)\right) \quad (\text{by (41)}) \\ &= N \exp\left(-\frac{\beta K}{(1-\delta)N} \frac{1}{9\bar{u}} \left(\frac{1}{2} \frac{\log M}{9\bar{u}\rho} - \frac{1}{3} \left(\frac{\log M}{9\bar{u}\rho}\right)^2\right)\right) \\ &= \exp\left(-\frac{\beta\psi K}{(1-\delta)N} + \log N\right), \end{aligned}$$

where

$$\psi := \frac{\log M}{81\bar{u}^2\rho} \left(\frac{1}{2} - \frac{\log M}{27\bar{u}\rho}\right) > 0 \quad (\text{by (35)}).$$

Since $(1-\delta) \frac{N \log N}{K} \rightarrow 0$, there exists \bar{l}_1 such that, for all $l \geq \bar{l}_1$, we have

$$\exp\left(-\frac{\beta\psi K}{(1-\delta)N} + \log N\right) \leq \frac{\eta}{300\bar{u}},$$

and thus (32) holds.

Condition (33): Since $\bar{Y} = \frac{\beta\mu \log M}{1-\delta}$ and $\lim_{l \rightarrow \infty} \delta^T = \lim_{\delta \rightarrow 1} \delta^{\frac{\rho\beta}{1-\delta}} = e^{-\rho\beta}$, we have

$$\lim_{l \rightarrow \infty} 32\bar{u} \frac{1-\delta}{1-\delta^T} \frac{\bar{Y}}{\mu} = \lim_{l \rightarrow \infty} \frac{32\bar{u}\beta \log M}{1-\delta^T} = \frac{32\bar{u}\beta \log M}{1-e^{-\rho\beta}} < \eta \quad (\text{by (37)}).$$

Hence, there exists \bar{l}_2 such that, for all $l \geq \bar{l}_2$, (33) holds.

Condition (34): First,

$$\lim_{l \rightarrow \infty} 8\bar{u} \frac{1-\delta^T}{\delta^T} = \lim_{\delta \rightarrow 1} 8\bar{u} \frac{1-\delta^{\frac{\rho\beta}{1-\delta}}}{\delta^{\frac{\rho\beta}{1-\delta}}} = 8\bar{u} (e^{\rho\beta} - 1) < \frac{\eta}{3} \quad (\text{by (38)}).$$

Second,

$$\lim_{l \rightarrow \infty} 8\bar{u} \frac{1-\delta}{\delta^T} T = \lim_{\delta \rightarrow 1} 8\bar{u} \frac{\rho\beta}{\delta^{\frac{\rho\beta}{1-\delta}}} = 8\bar{u} \frac{\rho\beta}{e^{\rho\beta}} < \frac{\eta}{3} \quad (\text{by (39)}).$$

Third,

$$\lim_{l \rightarrow \infty} 8\bar{u} \frac{1-\delta}{\delta^T} \frac{\bar{Y}}{\mu} = 8\bar{u} e^{\rho\beta} \beta \log M < \frac{\eta}{3} \quad (\text{by (40)}).$$

In total, the $l \rightarrow \infty$ limit of the left-hand side of (34) is strictly less than η . Hence, there exists \bar{l}_3 such that, for all $l \geq \bar{l}_3$, (34) holds.

The proof is now completed by taking $\bar{l} = \max\{\bar{l}_1, \bar{l}_2, \bar{l}_3\}$. ■

A.2.2 Equilibrium Construction

Fix any l , T , and \bar{Y} that satisfy (32)–(34), as well any $v \in \mathcal{C}(\eta)$. For each extreme point v^* of $\mathcal{C}_v(\eta/2)$, we construct a PPE in a T -period, finitely repeated game augmented with continuation values drawn from $\mathcal{C}_v(\eta/2)$ that generates payoff vector v^* . By standard arguments, this implies that $\mathcal{C}_v(\eta/2) \subseteq E(\Gamma)$, and hence that $v \in E(\Gamma)$.²⁰ Finally, since $v \in \mathcal{C}(\eta)$ was chosen arbitrarily, it follows that $\mathcal{C}(\eta) \subset E(\Gamma)$.

More precisely, for each $\zeta \in \{-1, 1\}^N$ and $v^* = \operatorname{argmax}_{v \in \mathcal{C}_v(\eta/2)} \zeta \cdot v$, we construct a strategy profile σ in a T -period, finitely repeated game (which we call a *block strategy profile*) together with a continuation value function $w : H^{T+1} \rightarrow \mathbb{R}^N$ that satisfy

²⁰In particular, at each history h^{T+1} that marks the end of a block, public randomization is 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 $w(h^{T+1})$.

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

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

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

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

Promise-Keeping.

$$v_i^* = \frac{1 - \delta}{1 - \delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \pi_i (h^{T+1}) \right] \text{ for all } i. \quad (43)$$

Incentive-Compatibility.

$$\sigma_i \in \operatorname{argmax}_{\tilde{\sigma}_i} \mathbb{E}^{\tilde{\sigma}_i, \sigma_{-i}} \left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \pi_i (h^{T+1}) \right] \text{ for all } i. \quad (44)$$

Self-Generation.

$$\zeta_i \frac{1 - \delta}{\delta^T} \pi_i (h^{T+1}) \in [-\eta, 0] \text{ for all } i \text{ and } h^{T+1}. \quad (45)$$

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

First, fix a correlated action profile $\bar{\alpha} \in \Delta \mathcal{A}$ such that, for each i ,

$$u_i (\bar{\alpha}) = v_i^* + \zeta_i \frac{\eta}{2}. \quad (46)$$

Such an $\bar{\alpha}$ exists because $v^* \in \mathcal{C}_v (\eta/2)$ and $\mathcal{C}_v (\eta) \subseteq F^*$. Also, fix an arbitrary static Nash equilibrium $\alpha^{NE} \in \Delta^* \mathcal{A} = \prod_i \Delta \mathcal{A}_i$.

Next, for each $i \in I$ and $\alpha \in \Delta^* \mathcal{A}$, we define functions $f_{i,\alpha} : \mathcal{A}_i \rightarrow \mathbb{R}$ and $g_{i,\alpha} : \mathcal{A}_i \rightarrow \mathbb{R}$, which will later be used in constructing σ and π .

Lemma 7 For each $i \in I$ and $\alpha \in \Delta^* \mathcal{A}$,

1. There exists a function $f_{i,\alpha} : \mathcal{A}_i \rightarrow [-3\bar{u}, 3\bar{u}]$ such that

$$u_i(a_i, \alpha_{-i}) + \mathbb{E}[f_{i,\alpha}(\hat{a}_i) | a_i, \alpha] = u_i(\alpha) \text{ for each } a_i \in \mathcal{A}_i, \quad (47)$$

$$\mathbb{E}[f_{i,\alpha}(\hat{a}_i) | \alpha] = 0. \quad (48)$$

2. There exists a function $g_{i,\alpha} : \mathcal{A}_i \rightarrow [-8\bar{u}, 8\bar{u}]$ such that

$$u_i(a_i, \alpha_{-i}) + \mathbb{E}[g_{i,\alpha}(\hat{a}_i) | a_i, \alpha_{-i}] = -4\zeta_i \bar{u} \text{ for each } a_i \in \mathcal{A}_i, \quad (49)$$

$$\zeta_i g_{i,\alpha}(\hat{a}_i) \leq 0. \quad (50)$$

Proof. For part 1, for each $a_i \in \mathcal{A}_i$, define

$$\Delta_{i,\alpha}(a_i) = u_i(\alpha) - u_i(a_i, \alpha_{-i})$$

and

$$f_{i,\alpha}(a_i) = \frac{1}{1-\varepsilon} \left(\Delta_{i,\alpha}(a_i) - \frac{\varepsilon}{M_i} \sum_{a'_i \in \mathcal{A}_i} \Delta_{i,\alpha}(a'_i) \right).$$

We verify that $f_{i,\alpha}$ satisfies the desired conditions, which is a matter of straightforward algebra. First,

$$|f_{i,\alpha}(a_i)| \leq \frac{1+\varepsilon}{1-\varepsilon} \max_{a_i \in \mathcal{A}_i} |\Delta_{i,\alpha}(a_i)| \leq \frac{1+\varepsilon}{1-\varepsilon} \bar{u} \leq 3\bar{u},$$

where the last inequality holds because $\varepsilon \leq 1/2$. Second,

$$\begin{aligned}
& \mathbb{E}[f_{i,\alpha}(\hat{a}_i) | a_i, \alpha_{-i}] \\
&= (1 - \varepsilon) f_{i,\alpha}(a_i) + \frac{\varepsilon}{M_i} \sum_{a'_i \in \mathcal{A}_i} f_{i,\alpha}(a'_i) \\
&= \left(\Delta_{i,\alpha}(a_i) - \frac{\varepsilon}{M_i} \sum_{a'_i \in \mathcal{A}_i} \Delta_{i,\alpha}(a'_i) \right) + \frac{\varepsilon}{(1 - \varepsilon) M_i} \sum_{a'_i \in \mathcal{A}_i} \left(\Delta_{i,\alpha}(a'_i) - \frac{\varepsilon}{M_i} \sum_{a''_i \in \mathcal{A}_i} \Delta_{i,\alpha}(a''_i) \right) \\
&= \left(\Delta_{i,\alpha}(a_i) - \frac{\varepsilon}{M_i} \sum_{a'_i \in \mathcal{A}_i} \Delta_{i,\alpha}(a'_i) \right) + \frac{\varepsilon}{(1 - \varepsilon) M_i} \left(\sum_{a'_i \in \mathcal{A}_i} \Delta_{i,\alpha}(a'_i) - \varepsilon \sum_{a'_i \in \mathcal{A}_i} \Delta_{i,\alpha}(a'_i) \right) \\
&= \Delta_{i,\alpha}(a_i).
\end{aligned}$$

Finally, $\mathbb{E}[f_{i,\alpha}(\hat{a}_i) | \alpha] = \sum_{a_i} \alpha_i(a_i) \mathbb{E}[f_{i,\alpha}(\hat{a}_i) | a_i, \alpha_{-i}] = \sum_{a_i} \alpha_i(a_i) \Delta_{i,\alpha}(a_i) = u_i(\alpha) - \sum_{a_i} \alpha_i(a_i) u_i(a_i, \alpha_{-i}) = 0$.

For part 2, for each $a_i \in \mathcal{A}_i$, let

$$g_{i,\alpha}(a_i) = f_{i,\alpha}(a_i) - u_i(\alpha) - 4\zeta_i \bar{u}.$$

Since $|f_{i,\alpha}(a_i)| \leq 3\bar{u}$ and $|u_i(\alpha)| \leq \bar{u}$, we have $|g_{i,\alpha}(a_i)| \leq 8\bar{u}$. In addition, (47) implies (49), and (48) (together with $|u_i(\alpha)| \leq \bar{u}$) implies (50). ■

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* $(E_{i,t}, F_{i,t}) \in \{0, 1\}^2$ for player i in period t , which will determine player i 's prescribed equilibrium action in period t . The states are determined by the public history, and so are common knowledge among the players. We first specify players' prescribed actions as a function of the state, and then specify the state as a function of the public history.

Prescribed Equilibrium Actions: The prescribed equilibrium actions are defined as follows. For each period t , let $a_t \in \mathcal{A}$ 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 (46).²¹

²¹Technically, the public randomization device Z_t is always a uniform $[0, 1]$ random variable. Throughout the proof, whenever we say that a certain variable is "drawn by public randomization," we mean that its realization is encoded in the realization of public randomization, independently of the other variables in the construction. Since we define only a finite number B of such variables, this can be done by, for example, specifying that if $n = b \bmod B$ then the n^{th} digit of z is used to encode the realization of the b^{th} such variable we define.

1. If $E_{i,t} = F_{i,t} = 0$ for all $i \in I$, the players take a_t .
2. If $E_{i,t} = 0$ for all i and there is a unique player i such that $F_{i,t} = 1$, the players take $(a'_i, a_{-i,t})$ for some $a'_i \in BR_i(a_{-i,t})$ if $\zeta_i = 1$, and take $(a'_i, \alpha_{-i}^{\min})$ for some $a'_i \in BR_i(\alpha_{-i}^{\min})$ if $\zeta_i = -1$, where $BR_i(\alpha_{-i}) = \operatorname{argmax}_{a_i \in \mathcal{A}_i} u_i(a_i, \alpha_{-i})$ is the set of i 's best responses to α_{-i} , and $\alpha_{-i}^{\min} \in \operatorname{argmin}_{\alpha_{-i} \in \Delta^* \mathcal{A}_{-i}} \max_{a_i \in \mathcal{A}_i} u_i(a_i, \alpha_{-i})$ is a mixed action profile that minmaxes player i .
3. If $E_{i,t} = 1$ for any i , or if $E_{i,t} = 0$ for all i and there is more than one player i such that $F_{i,t} = 1$, the players take α^{NE} .

Given the current state $(E_{i,t}, F_{i,t})_{i \in I}$ and the public randomization realization z_t , let $\alpha_t^* \in \Delta^* \mathcal{A}$ be the prescribed equilibrium action profile defined above. (In particular, $\alpha_t^* \in \left\{ a_t, ((a'_i, a_{-i,t}))_{i \in I, a'_i \in BR_i(a_{-i,t})}, ((a'_i, \alpha_{-i}^{\min}))_{i \in I, a'_i \in BR_i(\alpha_{-i}^{\min})}, \alpha^{NE} \right\}$.)

(It may be helpful to informally summarize the prescribed actions. So long as $E_{i,t} = F_{i,t} = 0$ for all players, the players take actions drawn from the target action distribution $\bar{\alpha}$. If $E_{i,t} = 1$ for any player, or if $F_{i,t} = 1$ for multiple players, the players take the arbitrary static Nash equilibrium α^{NE} . The most subtle case is when $E_{i,t} = 0$ for all i and there is a unique player i such that $F_{i,t} = 1$. Intuitively, this case will correspond to situations where player i 's monitored actions are “abnormal,” which later in the proof will imply that her continuation payoffs cannot be adjusted further without violated the self-generation constraint. In this case, player i starts taking static best responses. Moreover, if $\zeta_i = -1$ —so that player i 's continuation payoff is already “low”—her opponents start minmaxing her.)

States: The first component of player i 's period- t state, $E_{i,t}$, is defined as

$$E_{i,t} = \mathbf{1} \left\{ \sum_{t'=1}^{t-1} \mathbf{1} \{i \in y_{1,t'}\} \geq \mu T + \bar{Y} \right\}. \quad (51)$$

That is, $E_{i,t}$ is the indicator function for the event that the number of times that player i 's action has been monitored by period t exceeds the expected number of times that she is monitored in the entire T -period block by at least \bar{Y} .

The definition of the second component of player i 's period- t state, $F_{i,t}$, is more complicated. At the start of each period t , conditional on the draw of $a_t \in \mathcal{A}$ described above, an

additional random variable $\tilde{a}_t \in \mathcal{A}$ is also drawn by public randomization, such that $\tilde{a}_{i,t} = a_{i,t}$ with probability $1 - \varepsilon$ and $\tilde{a}_{i,t} = a'_{i,t}$ with probability $\varepsilon / (M_i - 1)$ for each $\tilde{a}_{i,t} \neq a'_{i,t}$, independently across players i . That is, the distribution of the public randomization draw \tilde{a}_t conditional on the draw a_t is the same as the distribution of the realized action profile \hat{a}_t when the profile of the players' intended actions is a_t ; however, note that the distribution of \tilde{a}_t depends only on the public randomization draw a_t , and not on the players' intended actions. For each player i and period t , let $f_{i,a_t} : \mathcal{A}_i \rightarrow [-3\bar{u}, 3\bar{u}]$ be defined as in Lemma 7, and let

$$f_{i,t} = \begin{cases} f_{i,a_t}(\hat{a}_{i,t}) & \text{if } E_{j,t} = F_{j,t} = 0 \text{ for all } j \in I, \\ f_{i,a_t}(\tilde{a}_{i,t}) & \text{if } E_{j,t} = 0 \text{ for all } j \in I, F_{i,t} = 0, \text{ and } F_{j,t} = 1 \text{ for some } j \neq i, \\ 0 & \text{if } E_{j,t} = 1 \text{ for some } j \in I \text{ or } F_{i,t} = 1. \end{cases} \quad (52)$$

Thus, the value of $f_{i,t}$ depends on the state $(E_{i,t}, F_{i,t})_{i \in I}$, the target action profile a_t (which is drawn from distribution $\bar{\alpha}$ as described above), the additional variable \tilde{a}_t , and the realized action profile \hat{a}_t . (Intuitively, the reason for introducing the variable \tilde{a}_t , rather than simply taking \hat{a}_t in place of \tilde{a}_t in (52), is that we want to ensure that the distribution of $f_{i,t}$ does not depend on $F_{j,t}$ for $j \neq i$, even though the distribution of $\hat{a}_{i,t}$ does depend on $F_{j,t}$.) Later in the proof, $f_{i,t}$ will be a component of the ‘‘reward’’ earned by player i in period t , which will be reflected in player i 's end-of-block continuation payoff function $\pi : H^{T+1} \rightarrow \mathbb{R}$.

We can finally define $F_{i,t}$ as

$$F_{i,t} = \mathbf{1} \left\{ \exists t' \leq t : \left| \sum_{t''=1}^{t'-1} \mathbf{1} \{i \in y_{1,t''}\} \delta^{t''-1} f_{i,t''} \right| \geq \bar{Y} \right\}. \quad (53)$$

That is, $F_{i,t}$ 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{Y} at any time $t' \leq t$.

This completes the definition of the 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 $\left((E_{j,t})_j, (F_{j,t})_{j \neq i} \right)_{t=1}^T$. (However, such a deviation can affect the distribution of $(F_{i,t})_{t=1}^T$.)

Lemma 8 For any player i and block strategy $\tilde{\sigma}_i$, the distribution of the random vector $\left((E_{j,t})_j, (F_{j,t})_{j \neq i} \right)_{t=1}^T$ is the same under block strategy profile $(\tilde{\sigma}_i, \sigma_{-i})$ as under block strategy profile σ .

Proof. Note first that the distribution of $\left((E_{j,t})_j \right)_{t=1}^T$ does not depend on the players' strategies. Now fix a player i and strategy $\tilde{\sigma}_i$. We argue that the distribution of $\left((E_{j,t})_j, (f_{j,t})_{j \neq i} \right)_{t=1}^T$ is the same under $(\tilde{\sigma}_i, \sigma_{-i})$ as under σ . Since the distribution of $\left((E_{j,t})_j \right)_{t=1}^T$ does not depend on anyone's strategy, and $f_{j,t} = 0$ whenever $E_{j',t} = 1$ for any player j' or $F_{j,t} = 1$, it suffices to show that, for each t , the distribution of $(f_{j,t})_{j \neq i: F_{j,t}=0}$ conditional on the event $\{E_{j,t'} = 0 \forall j, t' \leq t\}$ and the previous realizations $\left((f_{j,t'})_{j \neq i} \right)_{t'=1}^{t-1}$ is the same under $(\tilde{\sigma}_i, \sigma_{-i})$ and σ . This follows because, conditional on $\{E_{j,t'} = 0 \forall j, t' \leq t\}$ and $\left((f_{j,t'})_{j \neq i} \right)_{t'=1}^{t-1}$, for any $j \neq i$ such that $F_{j,t} = 0$, if $F_{j',t} = 0$ for all j' then $f_{j,t} = f_{j,a_t}(\hat{a}_{j,t})$, and $\hat{a}_{j,t} = a_{j,t}$ with probability $1 - \varepsilon$ and $\hat{a}_{j,t} = a'_{j,t}$ with probability $\varepsilon / (M_j - 1)$ for each $a'_{j,t} \neq a_{j,t}$; and if $F_{j',t} = 1$ for some j' then $f_{j,t} = f_{j,a_t}(\tilde{a}_{j,t})$, and $\tilde{a}_{j,t} = a_{j,t}$ with probability $1 - \varepsilon$ and $\tilde{a}_{j,t} = a'_{j,t}$ with probability $\varepsilon / (M_j - 1)$ for each $a'_{j,t} \neq a_{j,t}$. Hence, the distribution of $\left((E_{j,t})_j, (f_{j,t})_{j \neq i} \right)_{t=1}^T$ is the same under $(\tilde{\sigma}_i, \sigma_{-i})$ and σ , and therefore the distribution of $\left((E_{j,t})_j, (F_{j,t})_{j \neq i} \right)_{t=1}^T$ is also the same under these two profiles. ■

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

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

1. If $E_{j,t} = F_{j,t} = 0$ for all j , then

$$\pi_{i,t} = \delta^{t-1} \frac{\mathbf{1}\{i \in y_{1,t}\}}{\mu} \left(v_i^* + \zeta_i \frac{\eta}{4} - u_i(\alpha_t^*) + f_{i,\alpha_t^*}(\hat{a}_{i,t}) \right). \quad (54)$$

2. If $E_{j,t} = 0$ for all j and player i is the unique player such that $F_{i,t} = 1$, then

$$\pi_{i,t} = \delta^{t-1} \left(v_i^* + \zeta_i \frac{\eta}{4} - u_i(\alpha_t^*) \right). \quad (55)$$

3. If $E_{j,t} = 0$ for all j and there exists a player $j \neq i$ such that $F_{j,t} = 1$, then

$$\pi_{i,t} = \delta^{t-1} \frac{\mathbf{1}\{i \in y_{1,t}\}}{\mu} g_{i,\alpha_i^*}(\hat{a}_{i,t}). \quad (56)$$

4. If $E_{j,t} = 1$ for some j , then $\pi_{i,t} = 0$.

The constant c_i is defined as

$$c_i = -\mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left((1 - \max_{j \in I} E_{j,t}) \begin{pmatrix} (1 - \max_{j \neq i} F_{j,t}) (v_i^* + \zeta_i \frac{\eta}{4}) \\ + (\max_{j \neq i} F_{j,t}) (-4\zeta_i \bar{u}) \end{pmatrix} + (\max_{j \in I} E_{j,t}) u_i(\alpha^{NE}) \right) \right] + \frac{1 - \delta^T}{1 - \delta} v_i^*. \quad (57)$$

Note that, since $(v_i^* + \zeta_i \eta/4)$, $u_i(\alpha^{NE})$, and v_i^* are all feasible payoffs, we have

$$|c_i| \leq 5 \frac{1 - \delta^T}{1 - \delta} \bar{u}. \quad (58)$$

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

$$\pi_i(h^{T+1}) = c_i + \sum_{t=1}^T \pi_{i,t}. \quad (59)$$

A.2.3 Verification of the Equilibrium Conditions

We now verify that σ and π satisfy the promise-keeping, incentive-compatibility, and self-generation conditions. We first establish that $E_{i,t} = F_{i,t} = 0$ for all i and t with high probability, and then verify the three desired conditions in turn.

Lemma 9 *We have*

$$\Pr(E_{i,t} = 0 \text{ for all } i \in I \text{ and } t \in \{1, \dots, T\}) \geq 1 - \frac{\eta}{100\bar{u}}, \text{ and} \quad (60)$$

$$\Pr(F_{i,t} = 0 \text{ for all } i \in I \text{ and } t \in \{1, \dots, T\}) \geq 1 - \frac{\eta}{100\bar{u}}. \quad (61)$$

Proof. To establish (60), it suffices to show that, for each i ,

$$\Pr(E_{i,T} = 1) \leq \frac{1}{N} \frac{\eta}{100\bar{u}}.$$

This holds because, by (31) (i.e., Bennett's inequality), the fact that $h(x)$ is increasing (which is implied by (42)), and (32), we have

$$\Pr(E_{i,T} = 1) \leq \exp\left(-T\nu h\left(\frac{\bar{Y}}{T\nu}\right)\right) \leq \exp\left(-T\nu h\left(\frac{\bar{Y}}{9\bar{u}T\nu}\right)\right) \leq \frac{1}{N} \frac{\eta}{300\bar{u}} \leq \frac{1}{N} \frac{\eta}{100\bar{u}}.$$

To establish (61), it suffices to show that, for each i , $\Pr(\max_t F_{i,t \in \{1, \dots, T\}} = 1) \leq (1/N)(\eta/100\bar{u})$, or

$$\Pr\left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \mathbf{1}\{i \in y_{1,t'}\} \delta^{t-1} f_{i,t'} \right| \geq \bar{Y}\right) \leq \frac{1}{N} \frac{\eta}{100\bar{u}}. \quad (62)$$

To see this, let $\tilde{f}_{i,t} = f_{i,a_t}(\tilde{a}_{i,t})$. 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 $E_{j,t} = 0$ for all $j \in I$ and $F_{i,t} = 0$, while $f_{i,t} = 0$ if $E_{j,t} = 1$ for some $j \in I$ or $F_{i,t} = 1$, we have

$$\Pr\left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \mathbf{1}\{i \in y_{1,t'}\} \delta^{t-1} f_{i,t'} \right| \geq \bar{Y}\right) \leq \Pr\left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \mathbf{1}\{i \in y_{1,t'}\} \delta^{t-1} \tilde{f}_{i,t'} \right| \geq \bar{Y}\right). \quad (63)$$

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

$$\Pr\left(\max_{t \in \{1, \dots, T\}} \left| \sum_{t'=1}^t \mathbf{1}\{i \in y_{1,t'}\} \delta^{t-1} \tilde{f}_{i,t'} \right| \geq \bar{Y}\right) \leq 3 \max_{t \in \{1, \dots, T\}} \Pr\left(\left| \sum_{t'=1}^t \mathbf{1}\{i \in y_{1,t'}\} \delta^{t-1} \tilde{f}_{i,t'} \right| \geq \frac{\bar{Y}}{3}\right). \quad (64)$$

Letting $x_{i,t} = \mathbf{1}\{i \in y_{1,t}\} \delta^{t-1} \tilde{f}_{i,t}$, note that $|x_{i,t}| \leq 3\bar{u}$ with probability 1, $\mathbb{E}[x_{i,t}] = 0$, and $\text{Var}(x_{i,t}) \leq (3\bar{u})^2 \nu$. Therefore, by Bennett's inequality (which again applies because $(\tilde{f}_{i,t})_{t=1}^T$ are independent), the fact that $h(x)/x$ is increasing (by (42)), and (32), we have,

for each $t \leq T$,

$$\begin{aligned} \Pr \left(\left| \sum_{t'=1}^t \mathbf{1}_{\{i \in y_{1,t'}\}} \delta^{t'-1} \tilde{f}_{i,t'} \right| \geq \frac{\bar{Y}}{3} \right) &\leq \exp \left(-t\nu h \left(\frac{\bar{Y}}{9\bar{u}t\nu} \right) \right) \\ &\leq \exp \left(-T\nu h \left(\frac{\bar{Y}}{9\bar{u}T\nu} \right) \right) \leq \frac{1}{N} \frac{\eta}{300\bar{u}}. \end{aligned} \quad (65)$$

Finally, (63), (64), and (65) together imply (62). ■

Incentive-Compatibility: The following lemma simplifies the verification of incentive-compatibility.

Lemma 10 *For each player i and block strategy profile σ , incentive-compatibility holds (i.e., (44) is satisfied) if and only if*

$$\text{supp } \sigma_i(h^t) \subseteq \underset{\tilde{a}_{i,t} \in \mathcal{A}_i}{\text{argmax}} \mathbb{E}^{\sigma^{-i}} [\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, \tilde{a}_{i,t}] \text{ for all } t \text{ and } h^t. \quad (66)$$

Proof. We show that player i has a profitable one-shot deviation from σ_i at some history h^t if and only if (66) is violated at h^t . To see this, we first calculate player i 's continuation payoff under σ from period $t+1$ onwards (net of the constant c_i and the rewards already accrued $\sum_{t'=1}^t \pi_{i,t'}$). For any $t' \geq t+1$, there are four cases, which parallel the definition of the reward $\pi_{i,t}$.

1. If $E_{j,t'} = F_{j,t'} = 0$ for all j , then by Lemma 7 and (54), we have

$$\begin{aligned} \mathbb{E}^\sigma \left[\delta^{t'-1} u_{i,t'} + \pi_{i,t'} | h^{t'} \right] &= \delta^{t'-1} \left(u_i(\alpha_{t'}^*) + v_i^* + \zeta_i \frac{\eta}{4} - u_i(\alpha_{t'}^*) + \mathbb{E} \left[f_{i,\alpha_{t'}^*}(\hat{a}_{i,t'}) | \alpha_{t'}^* \right] \right) \\ &= \delta^{t'-1} \left(v_i^* + \zeta_i \frac{\eta}{4} \right). \end{aligned}$$

2. If $E_{j,t} = 0$ for all j and player i is the unique player such that $F_{i,t} = 1$, then by Lemma 7 and (55) we have

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

3. If $E_{j,t} = 0$ for all j and there exists a player $j \neq i$ such that $F_{j,t} = 1$, then by Lemma

7 and (56), we have

$$\mathbb{E}^\sigma \left[\delta^{t'-1} u_{i,t'} + \pi_{i,t'} | h^{t'} \right] = \delta^{t'-1} \left(u_i(\alpha_{t'}^*) + \mathbb{E} \left[g_{i,\alpha_{t'}^*}(\hat{a}_{i,t'}) | \alpha_{t'}^* \right] \right) = \delta^{t-1} (-4\zeta_i \bar{u}).$$

4. If $E_{j,t} = 1$ for some j , then $\pi_{i,t} = 0$, so we have

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

In total, player i 's continuation payoff under σ from period $t + 1$ onwards equals

$$\mathbb{E}^\sigma \left[\sum_{t'=t+1}^T \delta^{t'-1} \left((1 - \max_j E_{j,t'}) \left((1 - \max_{j \neq i} F_{j,t'}) (v_i^* + \zeta_i \frac{\eta}{4}) + (\max_{j \neq i} F_{j,t'}) (-4\zeta_i \bar{u}) \right) + (\max_j E_{j,t'}) u_i(\alpha^{NE}) \right) \right].$$

By Lemma 8, the distribution of $\left((E_{j,t'})_j, (F_{j,t'})_{j \neq i} \right)_{t'=t+1}^T$ does not depend on player i 's period- t action, and hence neither does player i 's continuation payoff under σ from period $t + 1$ onwards. Therefore, player i 's period- t action $\tilde{a}_{i,t}$ maximizes her continuation payoff from period t onwards if and only if it maximizes $\mathbb{E}^{\sigma_{-i}} [\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, \tilde{a}_{i,t}]$. ■

We now verify (66). Fix a player i , period t , and history h^t . We again consider four cases.

1. If $E_{j,t} = F_{j,t} = 0$ for all j , then for each action $\tilde{a}_{i,t}$, by Lemma 7 and (54) (and recalling that $\alpha_t^* = a_t$ when $E_{j,t} = F_{j,t} = 0$ for all j), we have

$$\begin{aligned} \mathbb{E}^{\sigma_{-i}} [\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, \tilde{a}_{i,t}] &= \delta^{t-1} \left(u_i(\tilde{a}_{i,t}, a_{-i,t}) + v_i^* + \zeta_i \frac{\eta}{4} - u_i(a_t) + \mathbb{E} [f_{i,a_t}(\hat{a}_{i,t}) | \tilde{a}_{i,t}, a_{-i,t}] \right) \\ &= \delta^{t-1} \left(v_i^* + \zeta_i \frac{\eta}{4} - u_i(a_t) \right). \end{aligned}$$

Since this does not depend on $\tilde{a}_{i,t}$, (66) holds.

2. If $E_{j,t} = 0$ for all j and player i is the unique player such that $F_{i,t} = 1$, then the reward $\pi_{i,t}$ specified by (55) does not depend on h_t . Hence, (66) 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, $(a'_i \in BR_i(a_{-i,t}), a_{-i,t})$ or $(a'_i \in BR_j(\alpha_{-i}^{\min}), \alpha_{-i}^{\min})$.

3. If $E_{j,t} = 0$ for all j and there exists a player $j \neq i$ such that $F_{j,t} = 1$, then for each action $\tilde{a}_{i,t}$, by Lemma 7 and (56), we have

$$\mathbb{E}^{\sigma_{-i}} [\delta^{t-1} u_{i,t} + \pi_{i,t} | h^t, \tilde{a}_{i,t}] = \delta^{t-1} (u_i(\tilde{a}_{i,t}, \alpha_{-i,t}^*) + \mathbb{E} [g_{i,\alpha_t^*}(\hat{a}_{i,t}) | \tilde{a}_{i,t}, \alpha_{-i,t}^*]) = \delta^{t-1} (-4\zeta_i \bar{u}).$$

Since this does not depend on $\tilde{a}_{i,t}$, (66) holds.

4. If $E_{j,t} = 1$ for some j , then $\pi_{i,t} = 0$, so (66) 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 condition holds for the prescribed action profile, α^{NE} .

Promise Keeping: By the definition of π_i , we have

$$\begin{aligned} & \frac{1-\delta}{1-\delta^T} \mathbb{E}^\sigma \left[\sum_{t=1}^T \delta^{t-1} u_{i,t} + \pi_i(h^{T+1}) \right] \\ = & \frac{1-\delta}{1-\delta^T} \left(\mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \begin{pmatrix} (1 - \max_j E_{j,t}) \begin{pmatrix} (1 - \max_{j \neq i} F_{j,t}) (v_i^* + \zeta_i \frac{\eta}{4}) \\ + (\max_{j \neq i} F_{j,t}) (-4\zeta_i \bar{u}) \end{pmatrix} \\ + (\max_j E_{j,t}) u_i(\alpha^{NE}) \end{pmatrix} \right] + c_i \right) \\ = & v_i^*. \end{aligned}$$

Self Generation: We first show that $\zeta_i \pi_i(h^{T+1}) \leq 0$ for every end-of-block history h^{T+1} , and then show that $\zeta_i ((1-\delta)/\delta^T) \pi_i(h^{T+1}) \geq -\eta$ for every h^{T+1} .

We first claim that, for every end-of-block history h^{T+1} ,

$$\zeta_i \sum_{t=1}^T \pi_{i,t} \leq \frac{\bar{Y} + 3\bar{u}}{\mu}. \quad (67)$$

To see this, first note that if $E_{j,t} = 1$ for some j or $F_{j,t} = 1$ for some $j \neq i$, then (56) and Lemma 7 imply that $\zeta_i \pi_{i,t} \leq 0$. Similarly, if $E_{j,t} = 0$ for each $j \in I$ and i is the unique player with $F_{i,t} = 1$, then (55) and (46) imply that

$$\zeta_i \pi_{i,t} = \delta^{t-1} \left(v_i^* + \zeta_i \frac{\eta}{4} - u_i(\alpha_t^*) \right) = \delta^{t-1} \left(v_i^* + \zeta_i \frac{\eta}{4} - \left(v_i^* + \zeta_i \frac{\eta}{2} \right) \right) = -\delta^{t-1} \frac{\eta}{4} \leq 0.$$

Hence, (54) implies that

$$\zeta_i \sum_{t=1}^T \pi_{i,t} \leq \zeta_i \sum_{t=1}^T \mathbf{1} \{E_{j,t} = F_{j,t} = 0 \forall j \in I\} \delta^{t-1} \frac{\mathbf{1} \{i \in y_{1,t}\}}{\mu} \left(v_i^* + \zeta_i \frac{\eta}{4} - u_i(a_t) + f_{i,a_t}(\hat{a}_{i,t}) \right).$$

Since $F_{i,t+1} = 1$ whenever $E_{j,t} = F_{j,t} = 0$ for all j and $\left| \sum_{t'=1,\dots,t} \delta^{t-1} \mathbf{1} \{i \in y_{1,t'}\} f_{i,a_{t'}}(\hat{a}_{i,t'}) \right| \geq \bar{Y}$, and in addition $v_i^* + \zeta_i \eta/4 - u_i(a_t) \leq 0$ and $|f_{i,a_t}(\hat{a}_{i,t})| \leq 3\bar{u}$, this inequality implies (67).

Next, by (57), we have

$$\begin{aligned} \zeta_i c_i &= -\zeta_i \left(\mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left((1 - \max_j E_{j,t}) \begin{pmatrix} (1 - \max_{j \neq i} F_{j,t}) (v_i^* + \zeta_i \frac{\eta}{4}) \\ + (\max_{j \neq i} F_{j,t}) (-4\zeta_i \bar{u}) \end{pmatrix} \right) + (\max_j E_{j,t}) u_i(\alpha^{NE}) \right] - \frac{1 - \delta^T}{1 - \delta} v_i^* \right) \\ &= -\zeta_i \left(\mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left((1 - \max_j E_{j,t}) \begin{pmatrix} (1 - \max_{j \neq i} F_{j,t}) (\zeta_i \frac{\eta}{4}) \\ + (\max_{j \neq i} F_{j,t}) (-4\zeta_i \bar{u} - v_i^*) \end{pmatrix} \right) + (\max_j E_{j,t}) (u_i(\alpha^{NE}) - v_i^*) \right] \right) \\ &\leq -\frac{1 - \delta^T}{1 - \delta} \left(\left(1 - \frac{\eta}{100\bar{u}}\right) \frac{\eta}{4} - \frac{\eta}{100\bar{u}} 5\bar{u} \right) \quad (\text{by (60)}) \\ &\leq -\frac{1 - \delta^T}{1 - \delta} \frac{\eta}{8} \end{aligned} \tag{68}$$

Finally, for each h^{T+1} , we have

$$\begin{aligned} \zeta_i \pi_i(h^{T+1}) &= \zeta_i \left(c_i + \sum_{t=1}^T \pi_{i,t} \right) \leq -\frac{1 - \delta^T}{1 - \delta} \frac{\eta}{8} + \frac{\bar{Y} + 3\bar{u}}{\mu} \quad (\text{by (67) and (68)}) \\ &\leq \frac{1 - \delta^T}{8(1 - \delta)} \left(-\eta + 32\bar{u} \frac{1 - \delta}{1 - \delta^T} \frac{\bar{Y}}{\mu} \right) \leq 0 \quad (\text{by } \bar{u} \geq 1 \text{ and (33)}). \end{aligned}$$

We now show that $\zeta_i ((1 - \delta)/\delta^T) \pi_i(h^{T+1}) \geq -\eta$ for every end-of-block history h^{T+1} . Note that $|\pi_{i,t}| \leq 3\bar{u} + (\mathbf{1} \{i \in y_{i,t}\} / \mu) 5\bar{u}$ for all t , and that $\pi_{i,t} = 0$ if $|\{t' \leq t : i \in y_{1,t'}\}| \geq$

$\mu T + \bar{Y}$. Since $|c_i| \leq ((1 - \delta^T) / (1 - \delta)) 5\bar{u}$ (by (58)), we have

$$\begin{aligned} \frac{1 - \delta}{\delta^T} |\pi_i(h^{T+1})| &\leq \frac{1 - \delta}{\delta^T} \left(\frac{1 - \delta^T}{1 - \delta} 5\bar{u} + \sum_{t=1}^T \delta^{t-1} 3\bar{u} + (\mu T + \bar{Y}) \frac{5\bar{u}}{\mu} \right) \\ &= \frac{1 - \delta^T}{\delta^T} 8\bar{u} + \frac{1 - \delta}{\delta^T} \left(T + \frac{\bar{Y}}{\mu} \right) 5\bar{u} \leq \eta \quad (\text{by (34)}). \end{aligned}$$

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 (1990): 1041-1063.
- [3] Al-Najjar, Nabil I., and Rann Smorodinsky. “Pivotal Players and the Characterization of Influence.” *Journal of Economic Theory* 92.2 (2000): 318-342.
- [4] Al-Najjar, Nabil I., and Rann Smorodinsky. “Large Nonanonymous Repeated Games.” *Games and Economic Behavior* 37.1 (2001): 26-39.
- [5] Athey, Susan, Kyle Bagwell, and Chris Sanchirico. “Collusion and Price Rigidity.” *Review of Economic Studies* 71.2 (2004): 317-349.
- [6] Bennett, George. “Probability Inequalities for the Sum of Independent Random Variables.” *Journal of the American Statistical Association* 57.297 (1962): 33-45.
- [7] Billingsley, Patrick. *Probability and Measure, 3rd ed.* Wiley (1995).
- [8] Border, Kim C., and Joel Sobel. “Samurai Accountant: A Theory of Auditing and Plunder.” *Review of Economic Studies* 54.4 (1987): 525-540.
- [9] Bowles, Samuel, and Herbert Gintis. *A Cooperative Species: Human Reciprocity and its Evolution.* Princeton University Press, 2011.
- [10] Boyd, Robert, and Peter J. Richerson. “The Evolution of Reciprocity in Sizable Groups.” *Journal of Theoretical Biology* 132.3 (1988): 337-356.
- [11] Cover, Thomas M., and Joy A. Thomas. *Elements of Information Theory, 2nd ed.* Wiley (2006).
- [12] Ekmekci, Mehmet, Olivier Gossner, and Andrea Wilson. “Impermanent Types and Permanent Reputations.” *Journal of Economic Theory* 147.1 (2012): 162-178.
- [13] Ellickson, Robert C. *Order without Law: How Neighbors Settle Disputes.* Harvard University Press, 1991.

- [14] Faingold, Eduardo. “Reputation and the Flow of Information in Repeated Games.” *Econometrica* 88.4 (2020): 1697-1723.
- [15] 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.
- [16] 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.
- [17] Fudenberg, Drew, and David K. Levine. “Repeated Games with Frequent Signals.” *Quarterly Journal of Economics* 124.1 (2009): 233-265.
- [18] Fudenberg, Drew, David Levine, and Eric Maskin. “The Folk Theorem with Imperfect Public Information.” *Econometrica* 62.5 (1994): 997-1039.
- [19] Fudenberg, Drew, David Levine, and Wolfgang Pesendorfer. “When are Nonanonymous Players Negligible?” *Journal of Economic Theory* 79.1 (1998): 46-71.
- [20] Georgiadis, George, and Balazs Szentes. “Optimal Monitoring Design.” *Econometrica* 88.5 (2020): 2075-2107.
- [21] Gossner, Olivier. “Simple Bounds on the Value of a Reputation.” *Econometrica* 79.5 (2011): 1627-1641.
- [22] Gossner, Olivier, Penelope Hernández, and Abraham Neyman. “Optimal Use of Communication Resources.” *Econometrica* 74.6 (2006): 1603-1636.
- [23] Green, Edward J. “Noncooperative Price Taking in Large Dynamic Markets.” *Journal of Economic Theory* 22.2 (1980): 155-182.
- [24] Green, Edward J., and Robert H. Porter. “Noncooperative Collusion under Imperfect Price Information.” *Econometrica* (1984): 87-100.
- [25] Hoffmann, Florian, Roman Inderst, and Marcus Opp. “Only Time will Tell: A Theory of Deferred Compensation.” *Review of Economic Studies* (2018).
- [26] Hörner, Johannes, and Satoru Takahashi. “How Fast do Equilibrium Payoff Sets Converge in Repeated Games?.” *Journal of Economic Theory* 165 (2016): 332-359.
- [27] Kandori, Michihiro, and Ichiro Obara. “Efficiency in Repeated Games Revisited: The Role of Private Strategies.” *Econometrica* 74.2 (2006): 499-519.
- [28] Khalil, Fahad, and Jacques Lawarree. “Input Versus Output Monitoring: Who is the Residual Claimant?.” *Journal of Economic Theory* 66.1 (1995): 139-157.
- [29] Kocherlakota, Narayana R. “Implications of Efficient Risk Sharing without Commitment.” *Review of Economic Studies* 63.4 (1996): 595-609.

- [30] Levine, David K., and Wolfgang Pesendorfer. "When are Agents Negligible?." *American Economic Review* (1995): 1160-1170.
- [31] Lewis, Tracy R., and David EM Sappington. "Using Markets to Allocate Pollution Permits and Other Scarce Resource Rights under Limited Information." *Journal of Public Economics* 57.3 (1995): 431-455.
- [32] Li, Anqi, and Ming Yang. "Optimal Incentive Contract with Endogenous Monitoring Technology." *Theoretical Economics* 15.3 (2020): 1135-1173.
- [33] Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies* 69.1 (2002): 209-244.
- [34] Mailath, George J., and Larry Samuelson. *Repeated Games and Reputations* (2006), Oxford University Press.
- [35] Maskin, Eric, and John Riley. "Input Versus Output Incentive Schemes." *Journal of Public Economics* 28.1 (1985): 1-23.
- [36] 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.
- [37] Mirrlees, James A. "The Theory of Moral Hazard and Unobservable Behaviour: Part I." *Working Paper* (published in *Review of Economic Studies* 66.1 (1999): 3-21).
- [38] Mookherjee, Dilip, and Ivan Png. "Optimal Auditing, Insurance, and Redistribution." *Quarterly Journal of Economics* 104.2 (1989): 399-415.
- [39] Neyman, Abraham, and Daijiro Okada. "Strategic Entropy and Complexity in Repeated Games." *Games and Economic Behavior* 29.1-2 (1999): 191-223.
- [40] Neyman, Abraham, and Daijiro Okada. "Repeated Games with Bounded Entropy." *Games and Economic Behavior* 30.2 (2000): 228-247.
- [41] Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
- [42] 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.
- [43] Pecorino, Paul. "The Effect of Group Size on Public Good Provision in a Repeated Game Setting." *Journal of Public Economics* 72.1 (1999): 121-134.
- [44] Rahman, David. "The Power of Communication." *American Economic Review* 104.11 (2014): 3737-51.
- [45] Reinganum, Jennifer F., and Louis L. Wilde. "Income Tax Compliance in a Principal-Agent Framework." *Journal of Public Economics* 26.1 (1985): 1-18.

- [46] Sabourian, Hamid. “Anonymous Repeated Games with a Large Number of Players and Random Outcomes.” *Journal of Economic Theory* 51.1 (1990): 92-110.
- [47] Sannikov, Yuliy, and Andrzej Skrzypacz. “Impossibility of Collusion under Imperfect Monitoring with Flexible Production.” *American Economic Review* 97.5 (2007): 1794-1823.
- [48] Sannikov, Yuliy, and Andrzej Skrzypacz. “The Role of Information in Repeated Games with Frequent Actions.” *Econometrica* 78.3 (2010): 847-882.
- [49] Sugaya, Takuo, and Alexander Wolitzky, “Communication and Community Enforcement.” *Journal of Political Economy* Forthcoming (2021).

B Online Appendix

B.1 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 $\mathbf{v} = (v, \dots, v) \in \mathbb{R}^N$, and let $\eta^* = cv(1 - c - v)/4 > 0$. We show that $\mathcal{C}_{\mathbf{v}}(\eta^*) \subseteq \mathcal{V}^*$ for all N .

To see this, fix any N . Since the game is symmetric, 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 $v + \eta^*$, and the remaining $N - n$ “disfavored” players receive payoffs no more than $v - \eta^*$. Fix such an n , and let $x = n/N$.

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

$$f(x) := x \frac{v + \eta^*}{1 - c} + (1 - x)(1) - c \frac{v + \eta^*}{1 - c},$$

while disfavored players receive payoff

$$g(x) := x \frac{v + \eta^*}{1 - c} + (1 - x)(1) - c.$$

Note that $f'(x) < 0$, so $f(x) \geq f(1) = v + \eta^*$.

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

$$f(x) \frac{v + \eta^*}{f(x)} = v + \eta^*,$$

while disfavored players receive payoff

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

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

B.2 Proof of Theorem 2.1 (Anti-Folk Theorem)

Fix a monitoring structure (\mathcal{Y}, p) . By standard arguments, the set E^S is a closed interval: $E^S = [\underline{v}^S, \bar{v}^S]$ for some $\underline{v}^S \leq \bar{v}^S$.

Lemma 11 *There exist $\alpha \in \Delta^* \mathcal{A}$ and $x : \mathcal{Y} \rightarrow \mathbb{R}$ such that*

$$\begin{aligned} \bar{v}^S &= u(\bar{\alpha}) - \mathbb{E}[x(y) | \bar{\alpha}], \\ \text{supp } \alpha &\subseteq \underset{a_0 \in \mathcal{A}_0}{\text{argmax}} u(a_0, \bar{\alpha}^{-i}) - \mathbb{E}[x(y) | a_0, \bar{\alpha}^{-i}] \text{ for all } i, \\ x(y) &\in \left[0, \frac{\delta}{1 - \delta} \bar{u}\right] \text{ for all } y. \end{aligned}$$

If the constraint $x(y) \in [0, (\delta/(1 - \delta)) \bar{u}]$ is replaced with $x(y) \in [-(\delta/(1 - \delta)) \bar{u}, 0]$, then the same statement holds with \underline{v}^S in place of \bar{v}^S

Proof. By standard arguments, E^S is self-generating: for any $v \in E^S$, there exist α and $w : \mathcal{Y} \rightarrow E^S$ such that

$$\begin{aligned} v &= (1 - \delta) u(\alpha) + \delta \mathbb{E}[w(y) | \alpha] \quad \text{and} \\ \text{supp } \alpha &\subseteq \underset{a_0 \in \mathcal{A}_0}{\text{argmax}} u(a_0, \bar{\alpha}^{-i}) + \delta \mathbb{E}[w(y) | a_0, \bar{\alpha}^{-i}] \text{ for all } i. \end{aligned}$$

Since \bar{v}^S is the greatest SSE payoff, if $v = \bar{v}^S$ then $w(y) \leq v$ for all $y \in \mathcal{Y}$. Hence, taking $v = \bar{v}^S = (1 - \delta) u(\alpha) + \delta \mathbb{E}[w(y) | \alpha]$ and defining $x(y) = (\delta/(1 - \delta)) (\bar{v}^S - w(y)) \geq 0$ for all y , we have

$$u(a) - \mathbb{E}[x(y) | a] = u(a) - \mathbb{E}\left[\frac{\delta}{1 - \delta} (\bar{v}^S - w(y)) | a\right] = (1 - \delta) u(a) + \delta \mathbb{E}[w(y) | a]$$

for all a , and $x(y) \leq (\delta/(1 - \delta)) \bar{u}$, and the result follows. Similarly, if $v = \underline{v}^S$ then $w(y) \geq v$ for all $y \in \mathcal{Y}$, and the symmetric argument applies. ■

Lemma 11 implies that \bar{v}^S is bounded by the solution to the program

$$\begin{aligned} &\max_{\alpha, x} u(\bar{\alpha}) - \mathbb{E}[x(y) | \bar{\alpha}] \quad \text{s.t.} \\ u(a_0, \bar{\alpha}^{-i}) - u(\bar{\alpha}) &\leq \mathbb{E}[x(y) | a_0, \bar{\alpha}^{-i}] - \mathbb{E}[x(y) | \bar{\alpha}] \text{ for all } a_0, \\ x(y) &\in \left[0, \frac{\delta}{1 - \delta} \bar{u}\right] \text{ for all } y, \\ |\mathbb{E}[x(y) | \bar{\alpha}]| &\leq \bar{u}. \end{aligned}$$

We show that, for every monitoring structure (\mathcal{Y}, p) , every α , every $x : \mathcal{Y} \rightarrow \mathbb{R}$ satisfying the constraints of this program, and every $\eta > 0$, there exists \bar{l} such that, whenever $l \geq \bar{l}$, we have $\mathbb{E}[x(y) | a_0, \bar{\alpha}^{-i}] - \mathbb{E}[x(y) | \bar{\alpha}] \leq \eta$ for all a_0 . This implies that $u(a_0, \bar{\alpha}^{-i}) - u(\bar{\alpha}) \leq \eta$ for all a_0 , and hence that \bar{v}^S is no greater than $\max_{\alpha \in \mathcal{A}_0(\eta)} u(\bar{\alpha})$, which completes the proof. (The argument for $\underline{v}^S \geq \min_{\alpha \in \mathcal{A}_0(\eta)} u(\bar{\alpha})$ is symmetric.)

To show this, fix α and a_0 , and consider the problem

$$\begin{aligned} \max_{\mathcal{Y}, p, x} \mathbb{E}[x(y) | a_0, \bar{\alpha}^{-i}] - \mathbb{E}[x(y) | \bar{\alpha}] \quad \text{s.t.} \\ x(y) \in \left[0, \frac{\delta}{1-\delta} \bar{u}\right] \text{ for all } y, \end{aligned} \quad (69)$$

$$|\mathbb{E}[x(y) | \bar{\alpha}]| \leq \bar{u}. \quad (70)$$

Let d denote the value of this problem. We show that if $(1-\delta) \exp(N^{1-\rho}) \rightarrow \infty$ for some $\rho > 0$ then $d \rightarrow 0$, which completes the proof.

Lemma 12 *There exists a solution (\mathcal{Y}, p, x) to the above problem that takes the form of a tail test: $Y = \{0, 1\}$, $x(0) = 0$, $x(1) = (\delta/(1-\delta))\bar{u}$, and there exists a number $n^* \in \{0, 1, \dots, N\}$ such that*

$$p(1|\hat{a}) \begin{cases} = 1 & \text{if } |\{i : \hat{a}_i = \hat{a}_0\}| > n^* \\ \in [0, 1] & \text{if } |\{i : \hat{a}_i = \hat{a}_0\}| = n^* \\ = 0 & \text{if } |\{i : \hat{a}_i = \hat{a}_0\}| < n^* \end{cases} . \quad (71)$$

Proof. We first show that there exists a solution satisfying $Y = \{0, 1\}$, $x(0) = 0$, and $x(1) = (\delta/(1-\delta))\bar{u}$. Fix any solution, and suppose there exists $y \in Y$ such that $x(y) \in (0, (\delta/(1-\delta))\bar{u})$. If we replace this y with two signals (y^-, y^+) and specify that

$$\begin{aligned} p(y^-|\hat{a}) &= \frac{(1-\delta)x(y)}{\delta\bar{u}} p(y|\hat{a}) \text{ for every } \hat{a}, \\ p(y^+|\hat{a}) &= \left(1 - \frac{(1-\delta)x(y)}{\delta\bar{u}}\right) p(y|\hat{a}) \text{ for every } \hat{a}, \\ x(y^-) &= \frac{\delta}{1-\delta}\bar{u}, \text{ and } x(y^+) = 0, \end{aligned}$$

then the resulting triple (\mathcal{Y}, p, x) also satisfies (45) and (70), and yields the same value for $\mathbb{E}[x(y) | (a_0, \bar{\alpha}^{-i})] - \mathbb{E}[x(y) | \bar{\alpha}]$. Repeatedly applying this variation for each y yields a solution satisfying $x(y) \in \{0, \bar{u}/(1-\delta)\}$ for all y . We may then identify all signals y such that $x(y) = (\delta/(1-\delta))\bar{u}$ with a single signal $y = 1$, and identify all signals y such that $x(y) = 0$ with a single signal $y = 0$.

We next show that there exists such a solution that further satisfies $p(1|\hat{a}) = p(1|\hat{a}')$ for all \hat{a}, \hat{a}' such that $|\{i : \hat{a}_i = a_0\}| = |\{i : \hat{a}'_i = a_0\}|$. Fix any solution satisfying $Y = \{0, 1\}$, $x(0) = 0$, and $x(1) = (\delta/(1-\delta))\bar{u}$, and fix any n . Let $\tilde{p}(1|\hat{a}) = p(1|\hat{a})$ for all \hat{a} such that $|\{i : \hat{a}_i = a_0\}| \neq n$, and let

$$\tilde{p}(1|\hat{a}) = \frac{\sum_{\hat{a}': |\{i: \hat{a}'_i = a_0\}| = n} p(1|\hat{a}') \Pr(\hat{a}'|\bar{\alpha})}{\sum_{\hat{a}': |\{i: \hat{a}'_i = a_0\}| = n} \Pr(\hat{a}'|\bar{\alpha})}.$$

for all \hat{a} such that $|\{i : \hat{a}_i = a_0\}| = n$. For any a_0 and any $a^* \in \mathcal{A}^N$ such that $|\{i : a^*_i = a_0\}| = n$,

$$\begin{aligned} & \sum_{\hat{a}} p(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i}) \\ &= \sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| \neq n} p(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i}) + \sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| = n} p(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i}) \\ &= \sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| \neq n} \tilde{p}(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i}) + \left(\sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| = n} \Pr(\hat{a}|\bar{\alpha}) \right) \frac{\sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| = n} p(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i})}{\sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| = n} \Pr(\hat{a}|\bar{\alpha})} \\ &= \sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| \neq n} \tilde{p}(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i}) + \left(\sum_{\hat{a}: |\{i: \hat{a}_i = a_0\}| = n} \Pr(\hat{a}'|\bar{\alpha}) \right) \tilde{p}(1|a^*) \\ &= \sum_{\hat{a}} \tilde{p}(1|\hat{a}) \Pr(\hat{a}|a_0, \bar{\alpha}^{-i}). \end{aligned}$$

Hence, \tilde{p} and p yield the same value in the above program. Repeatedly applying this variation for each n yields a solution satisfying $p(1|\hat{a}) = p(1|\hat{a}')$ for all \hat{a}, \hat{a}' such that $|\{i : \hat{a}_i = a_0\}| = |\{i : \hat{a}'_i = a_0\}|$.

Finally, we show that there exists such a solution that further satisfies (71) for some $n^* \in \mathbb{N}$. Denote the probability that a player's realized action under α differs from a_0 by $\chi = 1 - \varepsilon/(M_0 - 1) - \alpha(a_0)(1 - M_0\varepsilon/(M_0 - 1)) \in (\varepsilon, 1 - \varepsilon/(M_0 - 1))$. For $n \in \{0, \dots, N\}$ and any \hat{a} such that $|\{i : \hat{a}_i = a_0\}| = n$, let $p_n = p(1|\hat{a})$, and let $P_n = \Pr(|\{i \neq 1 : \hat{a}_i = a_0\}| = n)$,

given by $P_n = \binom{N-1}{n} (1-\chi)^n \chi^{N-1-n}$. Note that

$$\begin{aligned} \Pr(y = 1|\bar{\alpha}) &= (1-\chi) \sum_{n=0}^{N-1} P_n p_{n+1} + \chi \sum_{n=0}^{N-1} P_n p_n \quad \text{and} \\ \Pr(y = 1|a_0, \bar{\alpha}^{-i}) &= (1-\varepsilon) \sum_{n=0}^{N-1} P_n p_{n+1} + \varepsilon \sum_{n=0}^{N-1} P_n p_n, \quad \text{and so} \\ \Pr(y = 1|a_0, \bar{\alpha}^{-i}) - \Pr(1|\bar{\alpha}) &= (\chi - \varepsilon) \sum_{n=0}^{N-1} P_n (p_{n+1} - p_n). \end{aligned}$$

The problem becomes

$$\begin{aligned} \max_{(p_n)_{n \in [0,1]^{N+1}}} (\chi - \varepsilon) \sum_{n=0}^{N-1} P_n (p_{n+1} - p_n) \quad \text{s.t.} \\ \sum_{n=0}^{N-1} P_n ((1-\chi)p_{n+1} + \chi p_n) \leq \frac{1-\delta}{\delta}. \end{aligned} \quad (72)$$

Letting $\lambda \geq 0$ be the multiplier on (72), the Lagrangian is

$$\max_{(p_n)_{n \in [0,1]^{N+1}}} \sum_{n=0}^{N-1} P_n ((\chi - \varepsilon - \lambda(1-\chi))p_{n+1} - (\chi - \varepsilon + \lambda\chi)p_n),$$

with solution

$$p_n \begin{cases} = 1 & \text{if } \frac{P_n}{P_{n+1}} > \frac{\chi - \varepsilon + \lambda\chi}{\chi - \varepsilon - \lambda(1-\chi)} \\ \in [0, 1] & \text{if } \frac{P_n}{P_{n+1}} = \frac{\chi - \varepsilon + \lambda\chi}{\chi - \varepsilon - \lambda(1-\chi)} \\ = 0 & \text{if } \frac{P_n}{P_{n+1}} < \frac{\chi - \varepsilon + \lambda\chi}{\chi - \varepsilon - \lambda(1-\chi)} \end{cases}.$$

Since the binomial distribution is log-concave, P_n/P_{n+1} is increasing, so (71) holds for some $n^* \in \{0, \dots, N\}$. ■

We now show that, for any $\eta > 0$, there exists \bar{l} such that, for every $l \geq \bar{l}$, we have $d \leq \eta$. Let $n = |\{i : \hat{a}_i = a_0\}|$ and let $n^{-i} = |\{j \neq i : \hat{a}_j = a_0\}|$. Note that, for any n^*

$$\begin{aligned} \Pr(n = n^*|a_0, \bar{\alpha}^{-i}) &= (1-\varepsilon) \Pr(n^{-i} = n^* - 1|\bar{\alpha}^{-i}) + \varepsilon \Pr(n^{-i} = n^*|\bar{\alpha}^{-i}), \quad \text{and} \\ \Pr(n = n^*|\bar{\alpha}) &= (1-\chi) \Pr(n^{-i} = n^* - 1|\bar{\alpha}^{-i}) + \chi \Pr(n^{-i} = n^*|\bar{\alpha}^{-i}), \end{aligned}$$

and hence

$$\Pr(n \geq n^*|a_0, \bar{\alpha}^{-i}) - \Pr(n \geq n^*|\bar{\alpha}^{-i}) = (\chi - \varepsilon) \Pr(n^{-i} = n^* - 1|\bar{\alpha}^{-i}).$$

Therefore, by Lemma 12, d is given by

$$\max_{n^* \in \{0, 1, \dots, N\}, \beta \in [0, 1]} \frac{\delta}{1 - \delta} \bar{u}(\chi - \varepsilon) \left(\beta \Pr(n^{-i} = n^* - 1 | \bar{\alpha}^{-i}) + (1 - \beta) \Pr(n^{-i} = n^* | \bar{\alpha}^{-i}) \right) \quad (73)$$

$$\text{s.t.} \quad \beta \Pr(n = n^* | \bar{\alpha}) + \Pr(n \geq n^* + 1 | \bar{\alpha}) \leq \frac{1 - \delta}{\delta}, \quad (74)$$

where

$$\begin{aligned} \Pr(n^{-i} = n^* | \bar{\alpha}^{-i}) &= \binom{N-1}{n^*} (1 - \chi)^{n^*} \chi^{N-1-n^*} \quad \text{and} \\ \Pr(n = n^* | \bar{\alpha}) &= \binom{N}{n^*} (1 - \chi)^{n^*} \chi^{N-n^*}. \end{aligned}$$

Fix a sequence of games Γ (with $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$ for some $\rho > 0$) and pairs (n^*, β) indexed by l that satisfy the constraint (74). Suppose towards a contradiction that, for every \bar{l} , there is some $l \geq \bar{l}$ such that the value of the objective (73) exceeds η . Taking a subsequence and relabeling \bar{l} if necessary, this implies that there exists \bar{l} such that, for every $l \geq \bar{l}$, the value of the objective (73) exceeds η .

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

First, suppose that there exists $c > 0$ such that, for every \tilde{l} , there is some $l \geq \tilde{l}$ satisfying $|1 - \chi - (n^* - 1)/(N - 1)| > c$. By Hoeffding's inequality,

$$\begin{aligned} \beta \Pr(n^{-i} = n^* - 1 | \bar{\alpha}^{-i}) + (1 - \beta) \Pr(n^{-i} = n^* | \bar{\alpha}^{-i}) &\leq \Pr(n^{-i} \geq n^* - 1 | \bar{\alpha}^{-i}) \\ &\leq \exp\left(-2 \left(1 - \chi - \frac{n^* - 1}{N - 1}\right)^2 (N - 1)\right). \end{aligned}$$

Hence, for every \tilde{l} , there is some $l \geq \tilde{l}$ such that the value of (73) is at most

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

Since $(1 - \delta) \exp(N^{1-\rho}) \rightarrow \infty$, we have $\exp(-c^2 N)/(1 - \delta) \rightarrow 0$ for all $c > 0$, and hence (73) is less than η for sufficiently large l , a contradiction.

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

$$\left|1 - \chi - \frac{n^* - 1}{N - 1}\right| \leq c. \quad (75)$$

Note that

$$\begin{aligned}
\frac{\beta \Pr(n = n^* | \bar{\alpha}) + \Pr(n \geq n^* + 1 | \bar{\alpha})}{\Pr(n^{-i} = n^* | \bar{\alpha}^{-i})} &\geq \frac{\Pr(n \geq n^* + 1 | \bar{\alpha})}{\Pr(n^{-i} = n^* | \bar{\alpha}^{-i})} \\
&= \sum_{n=n^*+1}^N \frac{N\chi}{N-n^*} \frac{(N-n^*)!n^*!}{(N-n)!n!} \left(\frac{1-\chi}{\chi}\right)^{n-n^*} \\
&= \sum_{n=n^*+1}^N \frac{N\chi}{N-n^*} \frac{\frac{N-n^*}{\chi} \frac{N-n^*-1}{\chi} \dots \frac{N-n}{\chi}}{\frac{\frac{n}{N}}{1-\chi} \frac{\frac{n-1}{N}}{1-\chi} \dots \frac{\frac{n^*}{N}}{1-\chi}}.
\end{aligned}$$

By (75), for each k and $\gamma > 0$, there exists \tilde{l} such that, for every $l \geq \tilde{l}$, we have

$$\frac{N\chi}{N-n^*} \frac{\frac{N-n^*}{\chi} \frac{N-n^*-1}{\chi} \dots \frac{N-n}{\chi}}{\frac{\frac{n}{N}}{1-\chi} \frac{\frac{n-1}{N}}{1-\chi} \dots \frac{\frac{n^*}{N}}{1-\chi}} \geq 1 - \gamma \text{ for each } n^* + 1 \leq n \leq n^* + k,$$

and hence

$$\frac{\beta \Pr(n = n^* | \bar{\alpha}) + \Pr(n \geq n^* + 1 | \bar{\alpha})}{\Pr(n^{-i} = n^* | \bar{\alpha}^{-i})} \geq k(1 - \gamma).$$

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

$$\frac{\beta \Pr(n = n^* | \bar{\alpha}) + \Pr(n \geq n^* + 1 | \bar{\alpha})}{\Pr(n^{-i} = n^* - 1 | \bar{\alpha}^{-i})} \geq k(1 - \gamma).$$

Thus, for each k and $\gamma > 0$, there exists \tilde{l} such that, for every $l \geq \tilde{l}$, we have

$$\frac{\beta \Pr(n = n^* | \bar{\alpha}) + \Pr(n \geq n^* + 1 | \bar{\alpha})}{\beta \Pr(n^{-i} = n^* - 1 | \bar{\alpha}^{-i}) + (1 - \beta) \Pr(n^{-i} = n^* | \bar{\alpha}^{-i})} \geq k(1 - \gamma),$$

and therefore

$$\begin{aligned}
&\frac{\delta}{1-\delta} \bar{u}(\chi - \varepsilon) (\beta \Pr(n^{-i} = n^* - 1 | \bar{\alpha}^{-i}) + (1 - \beta) \Pr(n^{-i} = n^* | \bar{\alpha}^{-i})) \\
&\leq \bar{u}(\chi - \varepsilon) \frac{\beta \Pr(n^{-i} = n^* - 1 | \bar{\alpha}^{-i}) + (1 - \beta) \Pr(n^{-i} = n^* | \bar{\alpha}^{-i})}{\beta \Pr(n = n^* | \bar{\alpha}) + \Pr(n \geq n^* + 1 | \bar{\alpha})} \quad (\text{by (74)}) \\
&\leq \frac{\bar{u}(\chi - \varepsilon)}{k(1 - \gamma)}.
\end{aligned}$$

Taking k and γ such that $\bar{u}(\chi - \varepsilon) / (k(1 - \gamma)) < \eta$ gives a contradiction.

B.3 Proof of Theorem 2.2 (Folk Theorem)

Proof of Theorem 2.2(a). Fix \bar{a}_0 such that $u(\bar{a}_0) = \bar{v}$ and $\max_{a \in \mathcal{A}_0} u(a, \bar{a}_0^{-i}) - \bar{v} = d$, fix α^N such that $\bar{\alpha}^N$ is a static Nash equilibrium and $u(\bar{\alpha}^N) = \underline{v}^N$, and fix $\eta \in \left(\limsup \frac{(M_0-1)\varepsilon d}{M_0-1-2\varepsilon}, \liminf \frac{\bar{v}-\underline{v}^N}{2} \right)$. We construct an equilibrium that yields expected payoff $\bar{v} - \eta$. An equilibrium that yields any payoff $v \in [\underline{v}^N, \bar{v} - \eta]$ is then given by using public randomization to appropriately mix between this equilibrium and the infinite repetition of $\bar{\alpha}^N$.

Define (\mathcal{Y}, p) by letting $Y = \{0, 1\}$, $p(y = 1|\hat{a}) = 0$ if $\hat{a}_i = a_0$ for any player $i \in I$, and $p(y = 1|\hat{a}) = 1$ if $\hat{a}_i \neq a_0$ for all $i \in I$. Consider Nash reversion strategies where all players start by taking \bar{a}_0 and, in each period, permanently switch to $\bar{\alpha}^N$ with probability β if $y = 1$ (coordinated by public randomization), for some $\beta \in (0, 1)$ to be determined, and never switch if $y = 0$. We show that there exists a value for β such that these strategies yield expected payoff $\bar{v} - \eta$ and form an equilibrium.

Note that, when all players take a_0 , $\Pr(\hat{a}_i \neq a_0 \forall i) = \varepsilon^N$, and hence $\Pr(y = 1) = \varepsilon^N$. If instead one player takes an action other than a_0 , then $\Pr(\hat{a}_i \neq a_0 \forall i) = (1 - \varepsilon / (M_0 - 1)) \varepsilon^{N-1}$, and hence $\Pr(y = 1) = (1 - \varepsilon / (M_0 - 1)) \varepsilon^{N-1}$. Therefore, taking a_0 is optimal before the “switch” iff

$$\begin{aligned} & (1 - \delta) \bar{v} + \delta (\varepsilon^N \beta \underline{v}^N + (1 - \varepsilon^N \beta) (\bar{v} - \eta)) \\ & \geq (1 - \delta) \max_{a \in \mathcal{A}_0} u(a, \bar{a}_0^{-i}) + \delta \left(\left(1 - \frac{\varepsilon}{M_0 - 1}\right) \varepsilon^{N-1} \beta \underline{v}^N + \left(1 - \left(1 - \frac{\varepsilon}{M_0 - 1}\right) \varepsilon^{N-1} \beta\right) (\bar{v} - \eta) \right), \text{ or} \\ \beta & \geq \frac{1}{\varepsilon^{N-1} \left(1 - \frac{2\varepsilon}{M_0-1}\right)} \frac{1 - \delta}{\delta} \frac{d}{\bar{v} - \underline{v}^N - \eta}. \end{aligned} \quad (76)$$

Since taking α^N is clearly optimal after the switch, the prescribed strategies form an equilibrium whenever (76) is satisfied. Moreover, the prescribed strategies yield expected payoff $\bar{v} - \eta$ iff

$$\begin{aligned} (1 - \delta) \bar{v} + \delta (\varepsilon^N \beta \underline{v}^N + (1 - \varepsilon^N \beta) (\bar{v} - \eta)) & = \bar{v} - \eta, \quad \text{or} \\ \beta & = \frac{1}{\varepsilon^N} \frac{1 - \delta}{\delta} \frac{\eta}{\bar{v} - \underline{v}^N - \eta}. \end{aligned} \quad (77)$$

If $(1 - \delta) \exp(N^{1+\rho}) \rightarrow 0$ for some $\rho > 0$ then $(1 - \delta) / \varepsilon^N \rightarrow 0$. Together with the assumption that $\eta < \liminf (\bar{v} - \underline{v}^N) / 2$, this implies that $\beta \in (0, 1)$ for sufficiently large l . Finally,

(76) is satisfied for this value of β iff

$$\begin{aligned} \frac{1}{\varepsilon^N} \frac{1-\delta}{\delta} \frac{\eta}{\bar{v} - \underline{v}^N - \eta} &\geq \frac{1}{\varepsilon^{N-1} \left(1 - \frac{2\varepsilon}{M_0-1}\right)} \frac{1-\delta}{\delta} \frac{d}{\bar{v} - \underline{v}^N - \eta}, & \text{or} \\ \eta &\geq \frac{(M_0-1)\varepsilon d}{M_0-1-2\varepsilon}. \end{aligned}$$

Since $\eta > \limsup \frac{(M_0-1)\varepsilon d}{M_0-1-2\varepsilon}$, this holds for sufficiently large l , which completes the proof.

Proof of Theorem 2.2(b). The proof is parallel to that of part (a), except that now the target action \bar{a}_0 satisfies $u(\bar{a}_0) = \underline{v}$ and $\max_{a \in \mathcal{A}_0} u(a, \bar{a}_0^{-i}) - \underline{v} = \underline{d}$, reversion a static Nash equilibrium yielding payoff \bar{v}^N occurs with positive probability β only if $\hat{a}_i = a_0$ for every player i . The condition required for these strategies to form an equilibrium is the same as (76), and the condition required for them to yield expected payoff $\underline{v} + \eta$ is the same as (77), except with $(1-\varepsilon)^{N-1}$ in place of ε^{N-1} in both equations. Hence, (76) (with $(1-\varepsilon)^{N-1}$ in place of ε^{N-1}) is satisfied for the required value of β iff

$$\begin{aligned} \frac{1}{(1-\varepsilon)^N} \frac{1-\delta}{\delta} \frac{\eta}{\bar{v}^N - \underline{v} - \eta} &\geq \frac{1}{(1-\varepsilon)^{N-1} \left(1 - \frac{2\varepsilon}{M_0-1}\right)} \frac{1-\delta}{\delta} \frac{d}{\bar{v}^N - \underline{v} - \eta}, & \text{or} \\ \eta &\geq \frac{(M_0-1)(1-\varepsilon)d}{M_0-1-2\varepsilon}. \end{aligned}$$

Since $\eta > \limsup \frac{(M_0-1)(1-\varepsilon)d}{M_0-1-2\varepsilon}$, this holds for sufficiently large l , which completes the proof.