

How Robust is the Folk Theorem with Imperfect Public Monitoring?*

Johannes Hörner and Wojciech Olszewski[†]

April 2008

Abstract

In this paper, we prove that, under full rank, every perfect public equilibrium payoff (above some Nash equilibrium in the two-player case, and above the minmax payoff with more players) can be achieved with strategies that have bounded memory. This result is then used to prove that feasible, individually rational payoffs (above some strict Nash equilibrium in the two-player case, and above the pure-strategy minmax payoff with more players) can be achieved under private but almost-public monitoring.

1 Introduction

The concepts and techniques developed by Abreu, Pearce and Stacchetti (1990) have significantly affected our understanding of repeated games. Building on their work, Fudenberg, Levine and

*We would like to thank Olivier Compte, David Levine, George Mailath and Rakesh Vohra, as well as the seminar participants at Université de Montréal, UCLA, University of Michigan, University of Pennsylvania, and the Conference on Economic Theory and Experimental Economics at Washington University for helpful comments and conversations. Wojciech Olszewski thanks the National Science Foundation for research support (CAREER award SES-0644930). All remaining errors are ours.

[†]Department of Managerial Economics and Decision Sciences, Kellogg School of Management and Department of Economics, Weinberg College of Arts and Sciences, respectively, Northwestern University, 2001 Sheridan Road Evanston IL 60208

Maskin (1994) have shown that, when players are sufficiently patient, imperfect monitoring imposes virtually no restriction on the set of equilibrium payoffs, provided that some (generally perceived as mild) identifiability conditions are satisfied.

Recent progress in the study of games with private monitoring has raised some issues regarding the robustness of these concepts and techniques. When monitoring is private, players can no longer perfectly coordinate their continuation play. This inability is compounded over time, to the extent that the construction of equilibria from Abreu et al. may unravel under private monitoring no matter how close it is from being public.

Mailath and Morris (2002) show that uniformly strict perfect public equilibria are robust to sufficiently small perturbations in the monitoring structure provided that the strategy profile has finite memory, i.e. that it only depends on the last M observations, for some integer M . However, the strategies used in the literature to establish the folk theorem fail to have finite memory. In addition, Cole and Kocherlakota (2005) provide an example of a game with imperfect public monitoring in which the assumption of finite memory reduces the set of equilibrium payoffs. In their example, the set of equilibrium payoff vectors that can be achieved under finite memory, no matter how large the bound may be, is bounded away from the set of payoff vectors attainable with unbounded memory. However, the analysis is restricted to strongly symmetric strategies, and the monitoring structure that is considered violates the identifiability conditions imposed by Fudenberg et al.

The purpose of this paper is to show that the restriction to finite memory perfect public equilibria is to a large extent inconsequential, and that the folk theorem with imperfect public monitoring is robust to perturbations in the private monitoring. We establish a version of the folk theorem with imperfect public monitoring that only involves public strategies with finite memory. However, we strengthen the identifiability conditions of Fudenberg et al., and we further assume that a public randomization device is available. In the case of two players, our result is also weaker, in the sense that we show that every payoff vector that dominates a (static) Nash equilibrium payoff vector (as opposed to the minmax payoff vector) can be achieved in some equilibrium by sufficiently patient players.

We build on this result to prove that the folk theorem remains valid under almost-public private monitoring. This result is not a simple application of the results of Mailath and Morris,

because the strategy profiles we consider fail to be uniformly strict. Indeed, our construction relies on the indifference across actions of some players after specific histories. Nevertheless, our proof is largely inspired by their insights and results. As in the case of imperfect public monitoring, we rely on a strong version of the identifiability conditions, on the existence of a public randomization device, and we only show that every payoff vector that dominates a *strict* Nash equilibrium payoff vector is achievable under private almost-public monitoring in the case of two players.

These results prove that the example of Cole and Kocherlakota either relies on the restriction to strongly symmetric strategies, or the failure of the identifiability conditions (or a combination of both).¹ More importantly, our construction relies on the key ideas from Abreu et al. and Fudenberg et al., so that our result does not only establish the robustness of (a version of) their folk theorem, but also shows that their methods and concepts, such as self-generation, remain useful when one insists on robustness.

The present folk theorem with imperfect private monitoring differs from those from the existing literature. Mailath and Morris (2002) assume that monitoring is not only almost public but also almost perfect. More importantly, they incorrectly claim that the profile described in the proof of their Theorem 6.1 has bounded recall.² Hörner and Olszewski (2005) also assume almost perfect monitoring; moreover, they use substantially different strategies than those from Abreu et al. and Fudenberg et al. On the other hand, almost public signals are typically very correlated (even conditionally on action profiles), while Hörner and Olszewski (2005) allow for any pattern of correlation conditional on action profiles. Matsushima (2004) and Yamamoto (2006) do not assume almost perfect monitoring either, but they obtain the folk theorem only for the Prisoner's Dilemma and games with very similar structure; also, they assume conditionally independent signals (or signals that differ from public monitoring by conditionally independent components). Finally, other constructions of equilibria under private not almost perfect monitoring either do not yield the folk theorem (Ely et al., 2005), or allow players to communicate (Compte, 1998, Kandori and Matsushima, 1998, Aoyagi, 2002, Fudenberg and Levine, 2004, Obara, 2006), or allow players to obtain, at a cost, perfect signals (Miyagawa et al., 2005).

¹Indeed, within the framework of their example but for different probability distributions over signals than those they consider, the closure of the set of strongly symmetric payoff vectors is the same with infinite or finite (but arbitrarily long) memory.

²See Section 13.6 in Mailath and Samuelson (2006) for a detailed discussion of the flaw in their proof.

2 The Model

In the present paper, we study both public and private (but almost-public) monitoring. We first introduce the notation for public monitoring, used in Section 3, and then slightly modify the notation to accommodate private monitoring, examined in Section 4.

2.1 Public Monitoring

We follow most of the notation from Fudenberg et al. (1994). In the stage game, player $i = 1, \dots, n$ chooses action a_i from a finite set A_i . We denote by m_i the number of elements of A_i , and we call a vector $a \in A = \prod_{i=1}^n A_i$ a profile of actions. Profile a induces a probability distribution over the possible public outcomes $y \in Y$, where Y is a finite set with m elements. Player i 's realized payoff $r_i(a_i, y)$ depends on the action a_i and the public outcome y . We denote by $\pi(y | a)$ the probability of y given a . Player i 's expected payoff from action profile a is

$$g_i(a) = \sum_{y \in Y} \pi(y | a) r_i(a_i, y).$$

A mixed action α_i for each player i is a randomization over A_i , i.e. an element of ΔA_i . We denote by $\alpha_i(a_i)$ the probability that α_i assigns to a_i . We define $r_i(\alpha_i, y)$, $\pi(y | \alpha)$, $g_i(\alpha)$ in the standard manner. We often denote the profile in which player i plays α_i and all other players play a profile α_{-i} by (α_i, α_{-i}) .

Definition: A profile $a \in \prod_{i=1}^n A_i$ has $2n$ -wise full rank if for every player $i = 1, \dots, n$ and every action profile b such that $b_i \neq a_i$ and $b_{-i} = a_{-i}$, the $2(\sum_{i=1}^n m_i) \times m$ matrix $\Pi_i(a, b)$ with entries

$$\{\pi(y | c_j, a_{-j}), c_j \in A_j, j = 1, \dots, n\} \text{ and } \{\pi(y | c_j, b_{-j}), c_j \in A_j, j = 1, \dots, n\}$$

has rank $m_i + 2[\sum_{j \neq i} m_j - (n - 1)]$.

Note that $m_i + 2[\sum_{j \neq i} m_j - (n - 1)]$ is the maximal rank of $\Pi_i(a, b)$. Throughout the paper, we maintain the following assumption.

Assumption: Every action profile a has $2n$ -wise full rank.

Note that this ($2n$ -wise) full-rank condition implies that every action profile α has both *individual and pairwise full rank* in the sense of Fudenberg et al. (1994). The full-rank condition also implies that there are sufficiently many public outcomes, i.e. $m \geq m_i + 2[\sum_{j \neq i} m_j - (n - 1)]$. Thus, the full-rank condition is stronger than the conditions typically imposed in the studies of repeated games with imperfect public monitoring. However, if $m \geq m_i + 2[\sum_{j \neq i} m_j - (n - 1)]$, it is generically satisfied.

As will be clear from the proof, we can weaken this assumption somewhat: it is enough that the $2n$ -wise full rank condition is satisfied for one pure-action profile, and all other action profiles satisfy the combination of individual and pairwise full rank conditions described in Fudenberg et al. (1994). We do not know whether the $2n$ -wise full-rank condition can be entirely relaxed, and whether our results can be proven under the combination of individual and pairwise full rank conditions only.

For each i , the *minmax payoff* \underline{v}_i of player i (in mixed strategies) is defined as

$$\underline{v}_i := \min_{\alpha_{-i} \in \Delta A_{-i}} \max_{a_i \in A_i} g_i(a_i, \alpha_{-i}).$$

Pick $\underline{\alpha}_{-i} \in \Delta A_{-i} := \prod_{j \neq i} \Delta A_j$ so that

$$\underline{v}_i = \max_{a_i \in A_i} g_i(a_i, \underline{\alpha}_{-i}).$$

The payoff \underline{v}_i is the smallest payoff that the other players can keep player i below in the static game, and the profile $\underline{\alpha}_{-i}$ is an action profile that keeps player i below this payoff.

Let

$$U := \{(v_1, \dots, v_n) \mid \exists a \in A, \forall i, g_i(a) = v_i\},$$

$$V := \text{Convex Hull of } U,$$

and

$$\underline{V} := \text{Interior of } \{(v_1, \dots, v_n) \in V \mid \forall i, v_i > \underline{v}_i\}.$$

The set V consists of the feasible payoffs, and \underline{V} is the set of payoffs in the interior of V that strictly Pareto-dominate the minmax point $\underline{v} := (\underline{v}_1, \dots, \underline{v}_n)$. We assume throughout that \underline{V} is non-empty, a condition first introduced by Fudenberg et al.

Given a stage-game (pure or mixed) equilibrium payoff vector $v^* := (v_1^*, \dots, v_n^*)$, we also study the set

$$V^* := \text{Interior of } \{(v_1, \dots, v_n) \in V \mid \forall i, v_i > v_i^*\},$$

and we then assume that V^* is also non-empty.

We assume that players have access to a public randomization device. We henceforth suppose that in each period, all players observe a public signal $x \in [0, 1]$; the public signals are *i.i.d.* draws from the uniform distribution. We do not know whether the folk theorem presented in this paper is valid without a public randomization device. Such a device is necessary for Abreu, Pearce and Stacchetti (1990)'s 'bang-bang' result (Theorem 3) to be valid for monitoring structures with finitely many signals, a result on which our proof relies.

We now turn to the repeated game. At the beginning of each period $t = 1, 2, \dots$, players observe a public signal x^t . Then the stage game is played, resulting in a public outcome y^t . The public history at the beginning of period t is $h^t = (x^1, y^1, \dots, x^{t-1}, y^{t-1}, x^t)$; player i 's private history is $h_i^t = (a_i^1, \dots, a_i^{t-1})$. A (behavior) strategy σ_i for player i is a sequence of functions $(\sigma_i^t)_{t=1}^\infty$ where σ_i^t maps each pair (h^t, h_i^t) to a probability distribution over A_i .

Players share a common discount factor $\delta < 1$. All stage game payoffs are discounted and normalized by a factor $1 - \delta$. Thus, if $(g_i^t)_{t=1}^\infty$ is player i 's sequence of stage-game payoffs, the repeated game payoff is

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} g_i^t.$$

A strategy σ_i is *public* if in every period t it depends only on the public history h^t and not on player i 's private history h_i^t . We focus on equilibria in which all players' strategies are public, called perfect public equilibria (PPE). Given a discount factor δ , we denote by $E(\delta)$ the set of repeated game payoff vectors that correspond to PPE's when the discount factor is δ .

A strategy σ_i has *finite memory of length T* if it depends only on public outcomes and (randomization device) signals in the last T periods, i.e. on $(x^{t-T}, y^{t-T}, \dots, x^{t-1}, y^{t-1}, x^t)$ and $(a_i^{t-T}, \dots, a_i^{t-1})$. We denote by $E(\delta)$ the set of repeated game payoff vectors that can be achieved by a PPE, and by $E^T(\delta)$ the set of repeated game payoff vectors achievable by some PPE in which all players' strategies have finite memory of length T .

2.2 Private (Almost-Public) Monitoring

Whenever we consider monitoring structures that are private, we maintain the assumption that there is a public randomization device, which we take to be a uniform draw from the unit interval. Further, we make the assumption that the public monitoring structure has full support (as in Mailath and Morris, 2002).

Assumption: $\pi(y|a) > 0$ for all $y \in Y$ and $a \in A$.

For each i , the *minmax payoff* \underline{v}_i^P in pure strategies of player i is defined as

$$\underline{v}_i^P := \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} g_i(a_i, a_{-i}).$$

We further define the set \underline{V}^P by

$$\underline{V}^P := \text{Interior of } \{(v_1, \dots, v_n) \in V \mid \forall i, v_i > \underline{v}_i^P\}.$$

In Section 5, we assume that v^* corresponds to the equilibrium payoff vector of a *strict* Nash equilibrium.

A private monitoring structure is a collection of distributions $\{\rho(\cdot|a) : a \in A\}$ over Y^n . The interpretation is that each player i receives a signal $y_i \in Y$, and each signal profile (y_1, \dots, y_n) obtains with probability $\rho(y_1, \dots, y_n|a)$. Player i 's realized payoff $r_i(a_i, y_i)$ depends on the action a_i and the private signal y_i . Player i 's expected payoff from action profile a is therefore

$$g_i(a) = \sum_{(y_1, \dots, y_n) \in Y^n} \rho(y_1, \dots, y_n|a) r_i(a_i, y_i).$$

The private monitoring structure ρ is ϵ -close to the public monitoring structure π if $|\rho(y_1, \dots, y_n|a) - \pi(y|a)| < \epsilon$ for all y and a . Let $\rho_i(y_{-i}|a, y_i)$ denote the implied conditional probability of $y_{-i} := (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n)$, given a and y_i .

3 The Finite-Memory Folk Theorem under Public Monitoring

3.1 Result

In the appendix, we prove the following theorem.

Theorem 1 (i) ($n = 2$) For any payoff $v \in V^*$, there exists $\bar{\delta} < 1$, for all $\delta \in (\bar{\delta}, 1)$, there exists $T < \infty$ such that $v \in E^T(\delta)$. (ii) ($n > 2$) For any payoff $v \in \underline{V}$, there exists $\bar{\delta} < 1$, for all $\delta \in (\bar{\delta}, 1)$, there exists $T < \infty$ such that $v \in E^T(\delta)$

Observe that the result is stronger with three players or more. Also, observe the order of quantifiers.³ In particular, T is not independent of δ . In the proof presented below, the length of memory T is such that

$$\lim_{\delta \rightarrow 1} T = \infty.$$

3.2 Sketch of the Proof

The proofs of the cases (i) and (ii) are very similar. In both cases, we modify the proof of the folk theorem from Fudenberg et al. (1994), hereafter FLM. In particular, we take a compact set W that contains the vector v and show that $W \subset E^T(\delta)$. Our proof starts from the observation that matters would be relatively simple if players could publicly communicate, and communication was (i) cheap (i.e., not directly payoff-relevant), (ii) perfect (i.e., players observe precisely the opponents' message) and (iii) rich (i.e., players' message space is as large as the set of possible continuation payoff vectors). In this case, players could regularly communicate each others' continuation payoff, play accordingly, and ignore the history that led to those continuation payoffs.

However, even when communication is cheap, rich and perfect, there is the issue of incentive compatibility. As in mechanism design, this is problematic with two players: since no player i has an incentive to report truthfully the continuation payoff vector if it is the worst possible in the set W from his point of view, each player must have a message that would force his opponent's continuation payoff to be as low as possible in W . In turn, this implies that the set W includes a payoff vector \underline{w} that minimizes both players' payoff within W simultaneously. Therefore, the boundary of W must have a 'kink' at this point, precluding the smoothness of W assumed in the construction of equilibria from FLM. Observe, however, that a 'kink' of that kind is not a problem when it only occurs at the 'bottom left' of the set W , if this vector Pareto-dominates a Nash-equilibrium payoff vector of the stage game, since in that case the continuation payoff vector

³It might be possible to obtain the stronger conclusion that T only depends on v , rather than on both v and δ by considering more complicated sets W (see Section 4), but we have not explored this avenue.

in W can be taken along the line connecting the equilibrium payoff vector and \underline{w} , independently of the realized public signal. This is the reason why, in the two-player case, we only obtain a Nash-threat folk theorem.

We proceed in steps. We first show that for sets W satisfying certain properties, if communication is cheap, perfect and rich, then there exist protocols in which players truthfully reveal the continuation payoff vector in equilibrium. Moreover, we consider a protocol in which players take turns in sending messages, as this turns out to be convenient once we extend it to the case of imperfect communication. Of course, there is no cheap, perfect and rich communication in the repeated game. Players can use the public signals as messages, but signals satisfy neither of the three properties. Therefore, even if players wish to truthfully reveal continuation payoff vectors, they can only reveal a finite number of those; they can only do so imperfectly, and they cannot do so without affecting flow payoffs. We will show that each of these difficulties can be addressed by a suitable modification of the construction from FLM.

The strategy profile that we specify alternates phases of play signaling future continuation payoffs, with phases of play generating those payoffs:

Regular Phase (of length $M(\delta)$): in the periods of this phase, play proceeds as in FLM, given a continuation payoff vector;

Communication Phase (of length $N(\delta)$): in the periods of this phase, players use their actions to signal a continuation payoff vector.

As, the public histories from a communication phase contain information regarding the continuation payoff vector for the following regular phase, the strategies need only have memory $2N(\delta) + M(\delta)$. To guarantee that play during the communication phase hardly affects overall payoffs, it suffices to make sure that:

$$\lim_{\delta \rightarrow 1} \delta^{N(\delta)} = 1 \text{ and } \lim_{\delta \rightarrow 1} \delta^{M(\delta)} = 0. \quad (1)$$

Indeed, disregarding the flow payoffs and incentives during communication phases, one can slightly modify the construction from FLM to prove an analogue of their folk theorem for repeated games in which players' discount factor changes over time: it is δ for $M(\delta) - 1$ periods, and then

it drops, for one period, to $\delta^{N(\delta)}$. Alternatively, let $\delta' := \delta^{N(\delta)}$; taking any set W that is self-generated both for discount factors δ and δ' , we can mimic the proof of Theorem 1 in Abreu et al. (1990) to show that this implies that any payoff in W is an equilibrium payoff of our repeated game with the discount factor changing over time. Condition (1) also guarantees that flow payoffs during a communication phase hardly affect the equilibrium payoff vector as the discount factor tends to one.

As discussed, one of the issues related to communication via public signals is that this only allows for the transmission of coarse information. Observe, however, that it is enough to communicate finitely many continuation payoff vectors. Indeed, given any discount factor δ , any payoff vector from the boundary of W can be represented as the weighted average of a vector of flow payoffs in the current period and a continuation payoff from the interior of W . The set of all those continuation payoff vectors from the interior of W is a compact subset of W . Therefore there exists a (convex) polyhedron $P(\delta) \subseteq W$ that strictly contains all those vectors. Thus, one can communicate any such continuation payoff vector from the interior of W by randomizing over the vertices of $P(\delta)$ with appropriate probabilities (using the outcome of the randomization device during the first period of a communication period).⁴

However, as the discount factor δ increases, the number of vertices of $P(\delta)$ increases as well, and so does the length of the communication phase. So, calculations must be performed to ensure that condition (1) is satisfied.

To deal with the problem that communication is not perfect, notice that any desired vertex of $P(\delta)$ can be communicated with high probability by repeating for a number of times the message (playing for a number of times the action profile) that corresponds to this vertex. Suppose that players communicate in a similar manner vertices of another polyhedron $Q(\delta)$ that contains $P(\delta)$ in its interior. When the number of repetitions of the message that corresponds to any desired vertex of $Q(\delta)$ increases, the expected value of the distribution over the vertices of $Q(\delta)$ induced by this sequence of messages converges to the desired vertex. Ultimately, the polyhedron spanned by those expected values contains $P(\delta)$, and so the continuation payoff vectors for all payoff vectors from the boundary of W . A difficulty is that some players may have incentives to engage

⁴As mentioned, this is a version of Theorem 3 of Abreu, Pearce and Stacchetti (1990), which can be applied here because of the public randomization device.

in ‘signal-jamming’ when it is not their turn to ‘communicate’ in the protocol. This is where the assumption of n -wise full rank is useful.

The length of the communication phases increases with the number of repetitions of any message. Thus, additional calculations must be performed to ensure that condition (1) is met.

Finally, to deal with the problem that communication is not cheap, we prescribe strategies such that the play in a communication phase affects the continuation payoff vector differently in two complementary events, determined by public randomization. In one event, players will proceed as in FLM, given the continuation payoff vector that just has been communicated; in the other event, the continuation payoff vector will be contingent on the public outcomes observed in the communication phase, so that it makes every player indifferent (with respect to flow payoffs of the communication phase) across all actions (in every period of the communication phase), given any action profile of his opponents. The probability of this latter event is only of the order of $1 - \delta^{N(\delta)}$. Thus, this event affects only marginally the equilibrium payoff vector in the limit as $\delta \rightarrow 1$, and its presence can be treated as an additional discount factor (that also tends to 1 with δ) in the period that follows the communication phase.

4 The Folk Theorem under Private (Almost-Public) Monitoring

4.1 Result

In this section, we use the folk theorem under imperfect public monitoring and bounded memory to prove that a folk theorem remains valid under private almost-public monitoring. The proof follows the line of reasoning of Mailath and Morris (2002), although their result is not directly applicable, as the equilibrium with bounded memory used in the previous section is not uniformly strict. Indeed, our construction critically relies on the indifference of some players during the communication phase.

We may now state the main result of this section.

Theorem 2 *Fix a public monitoring structure π satisfying $(2n-)$ wise full rank and full support.*
(i) ($n = 2$) Suppose that v^ corresponds to the equilibrium payoff vector of a strict Nash equilibrium. For any payoff $v \in V^*$, there exists $\bar{\delta} < 1$, for all $\delta \in (\bar{\delta}, 1)$, there exists $\epsilon > 0$ such that*

$v \in E(\delta)$ for all private monitoring structures that are ϵ -close to π . (ii) ($n > 2$) For any payoff $v \in \underline{V}^P$, there exists $\bar{\delta} < 1$, for all $\delta \in (\bar{\delta}, 1)$, there exists $\epsilon > 0$ such that $v \in E(\delta)$ for all private monitoring structures that are ϵ -close to π .

4.2 A Modification of the Public Monitoring Strategies

To apply the previous results to the case of private monitoring, some modifications to the construction described above, and formally defined in the appendix, are necessary. First, it becomes important to use pure action profiles during the communication (observe that the action profiles α^* and β^* defined in appendix A.3.4. are indeed pure). More importantly, we must modify the construction that ensures that flow payoffs do not affect incentives during the communication phase (formally described in appendix A.3.4.). Recall that, with some small probability $\phi > 0$, right after the communication phase (say in period t), the realization of the public randomization device, x_t , is such that the continuation payoff vector is not equal to one signaled in the communication phase. Rather, as before, continuation play in this event is determined so as to guarantee that all players are indifferent over all actions during all periods of the communication phase that has just finished. Now, we will require the value x_t also to determine a period of the communication phase and a player. The interpretation is that continuation play in the event that period τ and player i have been determined is supposed to guarantee that player i is indifferent across all his actions in period τ .

Observe that there exists at most one player, say player j , whose action in period τ depends on the continuation payoff vector to be communicated (namely, the player who is sending a binary message in period τ). Pick some player k as follows (in the interpretation, player k controls i 's payoff in a way that ensures that player i is indifferent across all his actions in period τ of the communication phase). If $i \neq j$, let $k = j$, and if $i = j$, let k be some player different than i . It is an important property that player k always knows all equilibrium actions that were taken in period τ , except possibly the action of player i (when $j = i$).

Fix some pure action profiles $a, a' \in A$, with $a_{-k} = a'_{-k}$, $a_k \neq a'_k$; these will be action profiles to be taken in period t , by means of which player k controls that player i is indifferent across all his actions in period τ . By full-rank, we can find continuation payoff vectors $\gamma(y) \in \mathbb{R}^n$, for all $y \in Y$, such that:

(i) for all players $j \neq k$, a_j strictly maximizes

$$(1 - \delta) u_j(\tilde{a}_j, a_{-j}) + \delta \sum_y \pi(y|\tilde{a}_j, a_{-j}) \gamma_j(y),$$

and

$$(1 - \delta) u_j(\tilde{a}_j, a'_{-j}) + \delta \sum_y \pi(y|\tilde{a}_j, a'_{-j}) \gamma_j(y);$$

(ii) player k 's payoff

$$(1 - \delta) u_k(\tilde{a}_k, a_{-k}) + \delta \sum_y \pi(y|\tilde{a}_k, a_{-k}) \gamma_k(y)$$

is maximized both by a_k and a'_k ;

(iii) player i 's payoff under a exceeds that under a' :

$$(1 - \delta) u_i(a) + \delta \sum_y \pi(y|a) \gamma_i(y) - (1 - \delta) u_i(a') + \delta \sum_y \pi(y|a') \gamma_i(y) > K,$$

for some constant $K > 0$ to be specified.

That is, the action a_j is players $j \neq k$'s best-reply to $a_{-j} = a'_{-j}$ in the game in which public-outcome dependent continuation payoffs are given by γ ; player k is indifferent between both actions a_k and a'_k , given $a_{-k} = a'_{-k}$, and player i strictly prefers action profile a to a' .

The constant K is assumed not to be too small for a given δ , but any positive K is sufficient for our purposes and sufficiently large discount factors. More precisely, we require that there exists $f_k : A_k \times Y \rightarrow (0, 1)$, mapping the action played by player k in period τ and the signal player k received in period τ into the probability of him taking action a_k in period t with the following property: Player i is indifferent in period τ across all his actions, given that in the event that x_t picks player i and period τ in period t , player k chooses action a_k (rather than a'_k) with the probability given by f_k (and given that, in this event, continuation payoff vectors are given from period $t + 1$ on by γ). Notice that, by full-rank, f_k can be chosen so that the actions (in period τ) of players other than i and k do not affect the expected value of f_k .

Thus, in the event that x_t has picked player i and period τ , player k randomizes in period t between actions a_k and a'_k so as to ensure that player i is indifferent across all his actions in period τ . Meanwhile, players $j \neq k$ play their strict best-reply $a_j = a'_j$. Continuation payoffs

are then chosen according to γ , which can, without loss of generality, always be picked inside the polyhedron $P(\delta)$. Play then proceeds as before, given the public signal y observed in period t , players use the outcome x_{t+1} of the public randomization device to coordinate on some vertex of $P(\delta)$ giving $\gamma(y)$ as an expected value (where the expectation is taken over x_{t+1}), and we may as well assume that each vertex is assigned strictly positive probability.

4.3 Proof

We now return to private monitoring. The proof begins with the observation that the strategies in regular phases can be chosen uniformly strict, in the sense that there exists an $\nu > 0$ such that each player prefers his equilibrium action by at least ν to any other action. Indeed, Theorem 6.4 of FLM (and the remark that follows that theorem) establish that for any smooth set $W \subseteq V^*$, or $W \subseteq V^P$, there exists a minimal discount factor above which all elements of W correspond to strict perfect public equilibria. (A PPE is strict if each player strictly prefers his equilibrium action to any other action). Because in our construction, continuation payoffs can be drawn from a finite set of vertices (of the polyhedron $P(\delta)$), it follows that all elements of W correspond to uniformly strict perfect public equilibria, i.e. there exists $\nu > 0$ such that for any history of the regular phase for which the continuation value is not given by \underline{w} (in the case of $n = 2$), each player prefers his equilibrium action by at least ν to any other action.⁵ The assumption that v^* corresponds to the equilibrium payoff vector of a strict Nash equilibrium implies that some $\nu > 0$ also applies to the histories for which the continuation value is given by \underline{w} .

The following lemma is borrowed from Mailath and Morris (2002):

Lemma 3. (Mailath and Morris) Given η , there exists $\varepsilon > 0$ such that if ρ is ε -close to π , then for all $a \in A$ and $y \in Y$,

$$\rho_i(y, \dots, y|a, y) > 1 - \eta.$$

This lemma states that each player assigns probability arbitrarily close to one to all other players having observed the same signal as he has, independently of the action profile that has

⁵McLean, Obara and Postlewaite (2005) provides an alternative proof that the equilibrium strategies used in the construction of FLM can be chosen to be uniformly strict.

been played and the signal he has received, provided that private monitoring is sufficiently close to public monitoring. [Note that this result relies on the full support assumption.] For each integer T , this lemma immediately carries over to finite sequences of action profiles and signals of length no more than T ; it implies that, if all players use strategies of finite memory at most T , each player assigns probability arbitrarily close to one to all other players following the continuation strategy corresponding to the same sequence of signals he has himself observed (see Theorem 4.3 in Mailath and Morris).

This implies that we can pick $\varepsilon > 0$ so that player i 's continuation value to any of his strategies, given a private history h_i^t , is within $\nu/3$ of the value obtained by following the equilibrium strategy under public monitoring; we identify here the private history h_i^t under imperfect private monitoring with the corresponding history under public monitoring. (This is Lemma 3 of Mailath and Morris, 2002). This implies that, for this or lower $\varepsilon > 0$, by the one-shot deviation principle, players' actions in the regular phase under public monitoring remain optimal under imperfect private monitoring, since incentives were uniformly strict under public monitoring by the constant of ν .

It remains to prove that we can preserve players' incentives to play their public monitoring strategies during the communication phase under private monitoring, as well as in the period t right after the communication phase if the realization of public randomization x_t is such that continuation play will be making players indifferent over all actions during all periods of the communication phase that has just finished.

Consider first the communication phase. We claim that player k can still pick his action, a_k or a'_k in period t , according to some function $\tilde{f}_k : A_k \times Y \rightarrow (0, 1)$ (presumably different from f_k under public monitoring), so as to ensure that player i is indifferent between all his actions in period τ , and so that the actions of players other than i and k do not affect the expected value of \tilde{f}_k . Indeed, player k knows precisely the actions of all players $j \neq i$ in period τ , and the event that the realization of x_t selects player i and period τ is commonly known in period t , so that the expected continuation payoff vector in this event is independent of private histories up to period τ , given the prescribed strategies.

Consider now a realization of x_t in period t such that player k makes player i indifferent over all actions in period τ . First, observe that in period t all players' incentives were strict under public

monitoring, except for player k . So, the same action a_j remains optimal for all players $j \neq k$ under private monitoring (even if vertices in the next period are chosen according to distributions that are only close to those used under public monitoring). Recall that, under public monitoring, the signal y observed in period t determines a vertex v of the polyhedron according to a distribution $q(v|y)$ on which all players coordinate by means of the public randomization x_{t+1} .

Let v_1, \dots, v_L denote the set of vertices. Observe that, under public monitoring, each public history determines a vertex, and the corresponding vertex determines the players' continuation play. Therefore, we can identify each player's strategy with a finite automaton, as in Mailath and Morris, and the private states can be identified with the vertices. Let $v(x_{t+1}|y_i)$ denote the function mapping the outcome of the randomization device in period $t+1$ given i 's signal in period t into the vertex determining his continuation play (the same function applies for all players). For definiteness, assume that, if $y \neq y'$,

$$\begin{aligned} & \Pr \{x_{t+1}|v(x_{t+1}|y) = v, v(x_{t+1}|y') = v'\} = \\ & = \Pr \{x_{t+1}|v(x_{t+1}|y) = v\} \Pr \{x_{t+1}|v(x_{t+1}|y') = v'\}, \end{aligned}$$

that is, conditional on two different signals, players choose their vertices independently. To see that this is possible, we can view the randomization device as consisting of m independent uniform draws from $[0, 1]$ (recall that m is the number of public outcomes), with players using one distinct draw for each distinct public outcome. (Of course, the randomization device has been assumed so far to be a unique draw x from $[0, 1]$, but to see that this is equivalent, we can consider the dyadic expansion of x and split this expansion in m infinite expansions, from which we obtain m independent and uniform draws). While this assumption is inessential, it simplifies the argument. Let $q(v|y) := \Pr \{x_{t+1}|v(x_{t+1}|y) = v\}$.

Under imperfect private monitoring, players do not necessarily observe the same signal. Therefore each player j may select different vertices (or states) $v(x_{t+1}|y_j)$ based on the same realization of x_{t+1} , since the signal y_j is private.

Our objective is to show that, if monitoring is sufficiently close to public, we can find distributions $\{q(v|y) : v, y\}$ such that, if players use these distributions in period $t+1$ (and in all corresponding periods for other communication phases), player k is indifferent between the two actions a_k or a'_k (in the event that he must make player i indifferent), and prefers those actions

to all others. We shall prove that such distributions exist, using the implicit function theorem, given that ρ is in a neighborhood of π .

Observe that player k 's payoff from, say, action profile a is:

$$(1 - \delta) g_k(a) + \delta \sum_{(y_1, \dots, y_n)} \rho(y_1, \dots, y_n | a) \mathbb{E}_{x_{t+1}} \{V_k(v(x_{t+1}|y_1), \dots, v(x_{t+1}|y_n))\}, \quad (2)$$

where $V_k(v(x_{t+1}|y_1), \dots, v(x_{t+1}|y_n))$ is player k 's expected continuation payoff from the equilibrium strategy, given that each player j is in private state $v(x_{t+1}|y_j)$. By making player k 's payoffs from all actions equal to those with public monitoring, we obtain the number of equations equal to m_k , the number of actions of player k . The variables of choice are $\{q(v|y) : v = v_1, \dots, v_{L-1}, y = y_1, \dots, y_m\}$ they will be a function of the monitoring structure $\{\rho(y_1, \dots, y_n | a) : a \in A, (y_1, \dots, y_n) \in Y^n\}$. (Observe that we omit v_L , since $q(v_L|y)$ is pinned down by the other values). Consider the Jacobi matrix whose rows correspond to the equations and whose columns correspond to the variables of choice, evaluated at the public monitoring structure $\{\pi(y|a) : y, a\}$. The entry corresponding to the equation for action \tilde{a}_k and column for $v = v_l$ and $y = y_{m'}$ is equal to:

$$\pi(y_{m'} | (\tilde{a}_k, a_{-k})) (V_{k,l} - V_{k,L}) + \lambda,$$

where $V_{k,l}$ is player k 's continuation payoff under public monitoring if vertex l is selected, and λ is the remainder involving only terms premultiplied by $\delta^{M(\delta)}$. This expression is easy to obtain due to the assumption that conditional on two different signals players choose their vertices independently.

By construction, $V_{k,l} \neq V_{k,L}$ for some l . Since $\delta^{M(\delta)} \rightarrow 0$, also $\lambda \rightarrow 0$. Thus, since π satisfies individual full rank, the rank of our Jacobi matrix is equal to m_k . It is easy to see that the expression (2) is continuously differentiable with respect to each $q(v|y)$ (it is actually analytic). That is, the assumptions of the implicit function theorem are satisfied. This guarantees the existence of distributions $\{q(v|y) : v, y\}$ such that the payoff to player k under public and almost-public monitoring are equal (for each of his actions). Finally, since the values of $q(v|y)$ under public monitoring are in $(0, 1)$, it follows that there exists a neighborhood of the public monitoring structure in which the values of $q(v|y)$ under private monitoring are in $(0, 1)$ as well.

References

- [1] Abreu, D., D. Pearce, and E. Stacchetti (1990), “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, **58**, 1041–1063.
- [2] Aoyagi M. (2002), “Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication,” *Journal of Economic Theory*, **102**, 229–248.
- [3] Böröczky, K. (2001), “About the Error Term for Best Approximation with Respect to the Hausdorff Related Metrics,” *Discrete Computational Geometry*, **25**, 293–309.
- [4] Cole, H.L. and N.R. Kocherlakota (2005), “Finite Memory and Imperfect Monitoring,” *Games and Economic Behavior*, **53**, 59–72.
- [5] Compte O. (1998), “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica* **66**, 597–626.
- [6] Ely, J., J. Hörner and W. Olszewski (2005), “Belief-free Equilibria in Repeated Games”, *Econometrica*, **73**, 377–415.
- [7] Fudenberg, D., D. Levine, and E. Maskin (1994), “The Folk Theorem with Imperfect Public Information,” *Econometrica*, **62**, 997–1040.
- [8] Fudenberg D. and D. Levine (2004), “The Nash Threats Folk Theorem with Communication and Approximate Common Knowledge in Two Player Games,” Mimeo.
- [9] Hellwig, M. and W. Leininger (1987), “On the Existence of Subgame-Perfect Equilibrium in Infinite-Action Games of Perfect Information,” *Journal of Economic Theory*, **43**, 55–75.
- [10] Hörner, J. and W. Olszewski (2006), “The Folk Theorem for Games with Private Almost-Perfect Monitoring”, *Econometrica*, **74**, 1499–1544.
- [11] Kandori, M., and H. Matsushima (1998), “Private Observation, Communication and Collusion,” *Econometrica*, **66**, 627–652.

- [12] Mailath, G.J. and S. Morris (2002), “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory*, **102**, 189-228.
- [13] Mailath, G.J. and S. Morris (2006), “Coordination failure in repeated games with almost-public monitoring,” *Theoretical Economics*, **1**, 311-340.
- [14] Mailath G.J. and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.
- [15] McLean, R., Obara, I. and A. Postlewaite (2005), “Information Smallness and Private Monitoring in Repeated Games,” Mimeo.
- [16] Matsushima, H. (2004), “Repeated Games with Private Monitoring: Two Players,” *Econometrica*, **72**, 823-852.
- [17] Miyagawa, E., Y. Miyahara, and T. Sekiguchi (2005), “The Folk Theorem for Repeated Games with Observation Costs”, *Journal of Economic Theory*, forthcoming.
- [18] Obara, I. (2006), “Folk Theorem with Communication,” *Journal of Economic Theory*, forthcoming.
- [19] Yamamoto, Y. (2006), “Efficiency Results in N Player Games with Imperfect Private Monitoring,” *Journal of Economic Theory*, forthcoming.

A Details of the Proof of Theorem 1

We shall break the proof in several steps. First, we will show that if the condition (1) is satisfied, then any payoff $v \in \underline{V}$ can be achieved in an equilibrium of the (modified) repeated game in which players’ discount factor changes over time: it is δ for $M(\delta) - 1$ periods, and then it drops, for one period, to $\delta^{N(\delta)}$.

Given some stage game, define a repeated game with time-varying discount factor $\{\delta_t\}_{t=1}^{\infty}$ ($\delta_t < \delta' < 1$ for all t and some $\delta' < 1$) such that payoffs from period $t + 1$ onward are discounted

by the discount factor δ_t from the point of view of period t . That is, given flow payoffs g_t , total average payoff from period t onward is

$$v_t = \frac{g_t + \sum_{s=t+1}^{\infty} \prod_{\tau=t+1}^s \delta_s g_t}{1 + \sum_{s=t+1}^{\infty} \prod_{\tau=t+1}^s \delta_s}.$$

Following Mailath and Samuelson (2006), define $B(W; \delta)$ as the set of equilibrium payoffs, given that W is the set of continuation equilibrium payoffs (see their Definition 7.3.3).

Lemma 1. Let $W \subseteq \mathbb{R}^n$ be a bounded and convex set with $W \subseteq B(W; \delta)$ for some $\delta < 1$. Then for all $v_0 \in W$ and $\delta' \in (\delta, 1)$, v_0 is an equilibrium payoff vector for the game with any sequence of discount factors $\{\delta_t\}_{t=1}^{\infty}$, $\delta_t \in \{\delta, \delta'\}$ all t .

Proof: Since W is convex, $W \subseteq B(W; \delta)$ implies, by Proposition 7.3.4 in Mailath and Samuelson (2006), that $W \subseteq B(W; \delta'')$ for every $\delta'' \in [\delta, 1)$. Defining a public-outcome dependent payoff vector $v_t \in W$ for $t = 0, 1, \dots$ iteratively by observing that for each t , v_t can be written as a weighted average of a stage-game payoff vector g_t and a public-outcome dependent continuation payoff vector v_{t+1} , with weight $\delta'' \in [\delta, \delta']$ on v_{t+1} , we may mimic the proof of Proposition 7.3.1 in Mailath and Samuelson (2006). ■

This implies, for the modified repeated game that we analyze, that it is enough to show that $W \subseteq B(W; \delta^{N(\delta)})$ in order to conclude that W is a set of equilibrium payoffs in the modified repeated game.

A.1 The definition of the set W

Given any compact set W , let \underline{w}^i denote the vector that minimizes the i -th coordinate over W :

$$\underline{w}^i := \arg \min \{w_i : w \in W\}.$$

Let $w_i := \underline{w}_i^i$ be its i -th coordinate. Also, let \bar{w}^i denote the vector that maximizes the i -th coordinate over W :

$$\bar{w}^i := \arg \max \{w_i : w \in W\}.$$

Further, given any $\tilde{w} = (\tilde{w}_1, \tilde{w}_2)$ in $W \subset \mathbb{R}^2$, define \tilde{w}^1 by $\tilde{w}_2^1 := \tilde{w}_2$ and

$$\tilde{w}_1^1 := \min \{w_1 : w = (w_1, w_2) \in W \text{ and } w_2 = \tilde{w}_2\}.$$

That is, $\tilde{w}^1 = (\tilde{w}_1^1, \tilde{w}_2^1)$ is the horizontal projection of \tilde{w} onto the boundary of W that minimizes the first coordinate over W . Given \tilde{w} in W , and thus \tilde{w}^1 , let $\tilde{w}_1^{12} := \tilde{w}_1^1$ and

$$\tilde{w}_2^{12} := \min \{w_2 : w = (w_1, w_2) \in W \text{ and } w_1 = \tilde{w}_1^1\}.$$

That is, $\tilde{w}^{12} = (\tilde{w}_1^{12}, \tilde{w}_2^{12})$ is the vertical projection of \tilde{w}^1 onto the boundary of W that minimizes the second coordinate over W . See the right panel of Figure 1 for an illustration.

Finally, given any $\tilde{w} = (\tilde{w}_1, \dots, \tilde{w}_n)$ in (not necessarily two-dimensional) W , define \tilde{w}^2 as the vector in W that minimizes the second coordinate over all vectors w in W such that $w_1 = \tilde{w}_1$. Formally,

$$\tilde{w}_2^2 := \min \{w_2 : w = (w_1, \dots, w_n) \in W \text{ and } w_1 = \tilde{w}_1\},$$

and \tilde{w}_j^2 for $j \neq 2$ are the coordinates of $\arg \min \{w_2 : w = (w_1, \dots, w_n) \in W \text{ and } w_1 = \tilde{w}_1\}$.

We now consider the cases $n = 2$ and $n > 2$ in turn.

- $n = 2$:

We pick W such that its boundary consists of three quarters of a circle and a quarter of a square, as depicted in the right panel of Figure 1. Formally, for some $w^0 \in \mathbb{R}^2$ and $r > 0$,

$$W = \bar{B}(w^0, r) \cup \{\tilde{w} \in \mathbb{R}^2 : \tilde{w}_i \in [w_i^0 - r, w_i^0], i = 1, 2\},$$

where $\bar{B}(w^0, r)$ denotes the closed ball of radius r and center w^0 . We choose w and r such that $v \in W \subset V^*$. Observe that this set W has the following two properties:

(i) for any vector $\tilde{w} \in W$,

$$\tilde{w}_2^{12} = \underline{w}_2,$$

and

(ii) the boundary of W is smooth except at the point $\underline{w} := (\underline{w}_1, \underline{w}_2)$.

Given the (individual and pairwise) full-rank conditions, we know from FLM that, provided δ is large enough, W is self-generating (and is therefore a set of equilibrium payoff vectors). Although W is not smooth at \underline{w} , this is guaranteed by the existence of a Nash payoff vector that is strictly lower than \underline{w} .




Figure 1

- $n > 2$:

The set W is any closed ball such that $v \in W$ and $W \subset \underline{V}$.

Observe that:

$$\forall \tilde{w} \in W, \underline{w}_2^3 \geq \tilde{w}_2^2 \quad (3)$$

for any vector $\tilde{w} \in W$. That is, identifying coordinate i with player i 's payoff, player 2 always weakly prefers the payoff vector \underline{w}^3 to the vector \tilde{w}^2 . See the right panel of Figure 2 in the case $n = 3$.

Since W is a smooth subset of \underline{V} , we know from FLM that, under full rank and provided δ is large enough, W is self-generating (and is therefore a set of equilibrium payoff vectors).




Figure 2

A.2 Polyhedra $P(\delta)$

Given any sets $W, W' \subseteq \mathbb{R}^n$, let bdW and $intW$ denote, respectively, the boundary and the interior of W , let $d(W, W')$ denote the Hausdorff distance between W and W' .

Lemma 2. Let $W \subseteq \mathbb{R}^n$ be a compact and convex set. There exists $c > 0$ such that:

$$W \subseteq B(intW; \delta') \text{ for some } \delta' < 1 \Rightarrow P(\delta) \subseteq B(P(\delta); \delta) \text{ for every } \delta \in [\delta', 1),$$

and for all convex polyhedra $P(\delta)$ such that

$$\{v \in W : d(v, bdW) > c(1 - \delta)\} \subseteq P(\delta) \subseteq W.$$

Proof: For any $\delta \in [\delta', 1)$, any payoff vector from bdW can be represented as the weighted average of a stage-game payoff vector and a convex combination of payoff vectors from $intB(W; \delta)$ with total weight δ ; moreover, as δ converges to 1, those continuation payoff vectors converge to the payoff on the boundary approximately along straight lines. (See the proof of Proposition 7.3.4 in Mailath and Samuelson, 2006.) Thus, the distance of those continuation payoff vectors to the boundary, for any single payoff vector from bdW , is of order $1 - \delta$; by standard continuity and compactness arguments, there exists a $c > 0$ such that those continuation payoff vectors, for all payoff vectors from bdW , belong to the set

$$\{v \in W : d(v, bdW) > c(1 - \delta)\}.$$

This yields the result. ■

Given any number $c > 0$ from Lemma 2, let

$$W_\delta := \{v \in W : d(v, bdW) > c(1 - \delta)\}. \quad (4)$$

We now define some specific polyhedra $P(\delta)$ with $W_\delta \subseteq P(\delta) \subseteq W$. We must make sure that those polyhedra have not “too many” vertices, as otherwise, we would not be able to make sure later that condition (1) is satisfied. In the case of $n = 2$, start from a regular polygon D whose vertices lie on the boundary of $\bar{B}(w^0, r)$; additionally, assume that D is symmetric with respect to the vertical line passing through the point w^0 . Consider the convex hull of D and the point

$$(\min \{w_1 : w \in D\}, \min \{w_2 : w \in D\}).$$

For the family of discount factors $\delta_m = 1 - 2^{-m}$, we claim that it is enough to consider polygon D with 2^m vertices, at least when m is large enough.




Figure 3: The polyhedron $P(\delta_m)$

Denote by x_m the distance between the middle of an edge and the boundary of W . See Figure 3. Notice that

$$x_m = r \left(1 - \cos \frac{\pi}{2^m} \right),$$

where r denotes the radius of W , and so

$$\frac{x_m}{2^{-m}} = r \frac{1 - \cos \frac{\pi}{2^m}}{\frac{1}{2^m}};$$

this expression converges to 0 (as $m \rightarrow \infty$), because

$$\frac{1 - \cos \pi x}{x} \rightarrow_{x \rightarrow 0} 0.$$

Thus, x_m converges to 0 faster than 2^{-m} , which yields that

$$W_{\delta_m} \subseteq P(\delta_m)$$

(at least when m is large enough). In fact, the number of vertices $P(\delta_m)$ need not be exponential in m ; there exist polyhedra $P(\delta_m)$ with the required properties, and a polynomial number of vertices. See Böröczky (2001), which also contains the argument for the general case of $n \geq 2$.⁶

A.3 Truthful Communication of Continuation Payoffs

A.3.1 Cheap, rich and perfect *sequential* communication

In this section, we define an infinite-action game of perfect information with n players. (For a rigorous definition of such games, see Hellwig and Leininger, 1987.) In this game, Nature moves at date 0 and its actions are labeled by $w \in W$. Payoff vectors are elements of W . We require that the players' action sets be independent of previous play, and that payoffs be only a function of the play from period 1 on. Our objective is to construct a game such that in some subgame-perfect equilibrium of this game, for any initial action $w \in W$, the equilibrium payoff is w . That is, in equilibrium, players' actions from period 1 on uniquely identify the initial action of Nature. Our game will be of perfect information. We also wish that no two players move simultaneously.

⁶Although we will not need this, Böröczky (2001) establishes this result also for the case in which W is not a ball. We are indebted to Rakesh Vohra for this reference.

This last property is important for our construction of equilibria from Section 4. However, it prevents us from using now standard arguments from mechanism design.⁷ Nevertheless, we show that games with the required property exist, for sets W as defined in Section A.1. Consider both cases in turn.

- $n = 2$:

The extensive form is described in the left panel of Figure 1, in which Nature's initial move has been omitted. Player 2 moves first, choosing any action $\tilde{w} \in W$. Player 1 moves second, and chooses **Yes** or **No**. If player 1 chose **Yes**, player 2 moves a last time, and chooses **Yes** or **No**. Payoff vectors are as follows. If the history is $(\tilde{w}, \mathbf{Yes}, \mathbf{Yes})$, the payoff vector is \tilde{w} . If it is $(\tilde{w}, \mathbf{Yes}, \mathbf{No})$, the payoff vector is \tilde{w}^1 . If the history is (\tilde{w}, \mathbf{No}) , the payoff vector is \tilde{w}^{12} .

Given w , the equilibrium strategies $\sigma := (\sigma_1, \sigma_2)$ are as follows:

$$\begin{aligned} \sigma_2(\emptyset) &= w, \\ \sigma_2(\tilde{w}, \mathbf{Yes}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w}, \\ \mathbf{No} & \text{otherwise,} \end{cases} \\ \sigma_1(\tilde{w}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w}, \\ \mathbf{No} & \text{otherwise.} \end{cases} \end{aligned}$$

To check that this is an equilibrium, observe that player 2 is indifferent between both outcomes if the history reaches $(\tilde{w}, \mathbf{Yes})$ by definition of \tilde{w}^1 . Therefore, his continuation strategy is optimal. Given this, after history $\tilde{w} = w$, player 1 chooses between getting \tilde{w}_1 (if he chooses **Yes**) or \tilde{w}_1^{12} (if he chooses **No**), so that choosing **Yes** is optimal; after history $\tilde{w} \neq w$, he chooses between \tilde{w}_1^1 (if he chooses **Yes**) or \tilde{w}_1^{12} (if he chooses **No**), and choosing **No** is optimal as well. Given this, player 2 faces an initial choice between getting w_2 by choosing $\tilde{w} = w$ or getting $\tilde{w}_2^{12} = \underline{w}_2$ by choosing another action. Since $w_2 \geq \underline{w}_2$, choosing $\tilde{w} = w$ is optimal.

⁷If players were allowed to move simultaneously, the construction would become significantly simpler, but it is unclear how we could then adapt it to the repeated game, given the other difficulties that arise there.

- $n > 2$:

The extensive form is described in the left panel of Figure 2, in which Nature's initial move has been omitted. While there may be more than three players, only three of them take actions. Player 3 moves first, choosing any action $\tilde{w} \in W$. Player 2 moves second, and chooses **Yes** or **No**. Player 1 then moves a last time, and chooses either **Yes** or **No**. Payoff vectors are as follows. If the history is $(\tilde{w}, \mathbf{Yes}, \mathbf{Yes})$, the payoff vector is \tilde{w} . If it is $(\tilde{w}, \mathbf{Yes}, \mathbf{No})$, the payoff vector is \tilde{w}^2 . If the history is $(\tilde{w}, \mathbf{No}, \mathbf{Yes})$, the payoff vector is \underline{w}^3 . If the history is $(\tilde{w}, \mathbf{No}, \mathbf{No})$, the payoff vector is \underline{w}^2 .

Given w , the equilibrium strategies $\sigma := (\sigma_1, \sigma_2, \sigma_3)$ are as follows:

$$\begin{aligned}\sigma_3(\emptyset) &= w, \\ \sigma_2(\tilde{w}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w}, \\ \mathbf{No} & \text{otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \mathbf{Yes}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w}, \\ \mathbf{No} & \text{otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \mathbf{No}) &= \begin{cases} \mathbf{No} & \text{if } w = \tilde{w}, \\ \mathbf{Yes} & \text{otherwise.} \end{cases}\end{aligned}$$

Again, it is straightforward to verify that this is an equilibrium. Player 1 is indifferent between his actions after all histories at which he is to move; this follows from the definition of \tilde{w}^2 , and the fact that $\underline{w}_1^3 = \underline{w}_1^2$ (recall that W is a ball). Given this, after history $\tilde{w} = w$, player 2 chooses between getting \tilde{w}_2 (if he chooses **Yes**) or \underline{w}_2^2 (if he chooses **No**), so that choosing **Yes** is optimal; after history $\tilde{w} \neq w$, he chooses between \tilde{w}_2^2 (if he chooses **Yes**) or \underline{w}_2^3 (if he chooses **No**), and choosing **No** is optimal as well given (3). Given this, player 3 faces an initial choice between getting w_3 by choosing $\tilde{w} = w$ or getting \underline{w}_3^3 by choosing another action. Choosing $\tilde{w} = w$ is optimal.

It is not difficult to see that similar constructions exist for sets W which do not have all the properties specified earlier. However, the only feature that has been assumed which turns out to be restrictive for our purposes is that $\underline{w} \in W$ in case $n = 2$, and this feature cannot be dispensed

with: if $\underline{w} \notin W$, there exists no game solving the problem described in this section. Indeed, since there must be an equilibrium in which player 1 receives his minimum over the set W , player 2 must have a strategy s_2 such that player 1's payoff is equal (at best) to this lowest payoff, independently of his own strategy; similarly, player 1 must have a strategy s_1 such that player 2's payoff is equal to his lowest payoff independently of his own strategy. The pair of strategies (s_1, s_2) must yield a payoff vector from the set W , and so it must be the vector that minimizes simultaneously the payoffs of both players. Thus, we cannot avoid a kink in the set W , and thus, using the present approach, we cannot strengthen our result in the case $n = 2$ beyond a Nash-threat folk theorem.

A.3.2 Cheap, rich and *imperfect* sequential communication

We shall now replace W by a convex polyhedron $P(\delta)$ constructed in Section A.2. In the case of two players, we can assume that for any vertex w of $P(\delta)$, the resulting vectors w^1 and w^{12} are *vertices* of $P(\delta)$ (where, in the definition of w^{12} the original set W is replaced by $P(\delta)$); by construction, this can be done by increasing the number of vertices (by no more than a half). Also, for any of its vertex w , the vector w^{12} satisfies $w_2^{12} = \underline{w}_2$ (again, in the definition of \underline{w}^i the original set W is replaced by $P(\delta)$). This guarantees that the construction of the extensive form game also applies to this polyhedron.

In the case of $n > 2$, again increasing the number of vertices, we can assume that the vectors \underline{w}^i and \bar{w}^i , $i = 1, 2$, and for every vertex w , the resulting vector w^2 are *vertices* of $P(\delta)$. The number of vertices of the so modified $P(\delta)$ is a linear function of the number of vertices of $P(\delta)$.

Since players will only need to convey which vertex of $P(\delta)$ is called for, and this polyhedron has a finite number of vertices, the (two) games of perfect information can be simplified to finite-action games, in which Nature chooses its initial action from the set of vertices of $P(\delta)$, and the first player's (player 2 if $n = 2$, player 3 if $n = 3$) action set is also this set of vertices. The construction of the extensive form game and an equilibrium of this game applies without change.

However, we now modify the extensive form to obtain the following stronger property: each player's action is optimal even if he expects all other players to tremble with fixed, exogenous probabilities not exceeding some small $\varepsilon > 0$.

Case $n = 2$: Begin with the case $n = 2$. The extensive form is defined as follows. Nature's initial move is omitted. See also Figure 4.




Figure 4

- Player 2 moves first, choosing any \tilde{w} (a vertex of the polyhedron $P(\delta)$).
- Nature (the public randomization device) moves second, choosing **Left** and **Right** with probability ξ and $1 - \xi$, respectively (to be defined).
- Player 1 moves third, choosing **Yes** or **No**.

- If Nature has chosen **Right** and Player 1 has chosen **No**, Nature chooses **Left** and **Right** with probability $\eta(\tilde{w})$ and $1 - \eta(\tilde{w})$, respectively (to be defined).
- If Nature has chosen **Right** and Player 1 has chosen **Yes**, Player 2 chooses **Yes** or **No**.

Payoff vectors: The payoff vector is defined as follows. If the history is $(\tilde{w}, \mathbf{Left}, \mathbf{Yes})$, the payoff vector is w' to be defined; if it is $(\tilde{w}, \mathbf{Left}, \mathbf{No})$, the payoff vector is w'' to be defined; if it is $(\tilde{w}, \mathbf{Right}, \mathbf{Yes}, \mathbf{Yes})$, the payoff vector is \tilde{w} ; if it is $(\tilde{w}, \mathbf{Right}, \mathbf{Yes}, \mathbf{No})$, the payoff vector is \tilde{w}^1 ; if the history is $(\tilde{w}, \mathbf{Right}, \mathbf{No}, \mathbf{Left})$, the payoff vector is \tilde{w} ; finally, if the history is $(\tilde{w}, \mathbf{Right}, \mathbf{No}, \mathbf{Right})$, it is \tilde{w}^{12} .

The vectors w' and w'' are vertices of $P(\delta)$ such that player 1 is indifferent between both, but player 2 strictly prefers w' to w'' (thus, any strategy of player 1 in the event that Nature picks **Left** will be optimal). By construction, there exist vertices w' and w'' with the required properties.

Compared to the construction from Section A.3.1, note that Nature appears in two additional instances. This allows us to deal with the trembles.

Strategies: Given w , the strategy profile $\sigma := (\sigma_1, \sigma_2)$ defined as follows is a subgame-perfect equilibrium given the trembles.

$$\begin{aligned} \sigma_2(\emptyset) &= w, \\ \sigma_2(\tilde{w}, \mathbf{Right}, \mathbf{Yes}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w}, \\ \mathbf{No} & \text{otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \mathbf{Right}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w}, \\ \mathbf{No} & \text{otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \mathbf{Left}) &= \begin{cases} \mathbf{Yes} & \text{if } w = \tilde{w} \text{ and } w \text{ minimizes player 2's payoff over } P(\delta), \\ \mathbf{No} & \text{otherwise.} \end{cases} \end{aligned}$$

Observe that, given these strategies, if the history is $(\tilde{w}, \text{Right})$, with $w \neq \tilde{w}$, player 1 faces a lottery between getting \tilde{w} with low probability and \tilde{w}^1 with high probability (by choosing **Yes**), and a lottery between \tilde{w} and \tilde{w}^{12} (by choosing **No**), with probability $\eta(\tilde{w})$ and $1 - \eta(\tilde{w})$, respectively. We can thus pick $\eta(\tilde{w})$ so as to guarantee that player 1 is indifferent between both choices in this event. Note that $\eta(\tilde{w})$ is accordingly low (it is equal to the probability of a tremble on **Yes** when player 2's planned action is **No**).

In case w minimizes player 2's payoff over $P(\delta)$, player 2's incentive for revealing truthfully $\tilde{w} = w$ comes from the event that Nature picks **Left** in the second round, since in that case truth-telling (resp. lying) gives a lottery between w' and w'' with high (resp. low) probability on w' . Since in the event that Nature picks **Right** in the second period, player 2 gets $\tilde{w}_2^{12} = \underline{w}_2^2$ with very high probability, we can ensure truthful revelation by choosing a low $\xi > 0$ (again, this probability can be set proportional to a tremble). The remaining equilibrium conditions are immediate to verify for small $\varepsilon > 0$, given the definitions of \tilde{w}^1 and \tilde{w}^{12} .

It is important to note that $\varepsilon > 0$ must be small relative to the differences in coordinates of any two vertices of $P(\delta)$, whenever these coordinates are different.

Case $n > 2$: Consider now the case $n > 2$. The extensive form is defined as follows. As before, Nature's initial move is omitted. See also Figure 4.

- Player 3 moves first, choosing any \tilde{w} (a vertex of the polyhedron $P(\delta)$).
- Nature (the public randomization device) moves second, choosing **Left** and **Right** with probability ξ and $1 - \xi$, respectively (to be defined).
- Player 2 moves third, choosing **Yes** or **No**.
- If Nature has chosen **Right** in the second stage, Nature then moves again, choosing **Left** and **Right** with probability ξ and $1 - \xi$, respectively.
- If Nature has chosen **Right** in the second stage, player 1 moves last, choosing **Yes** or **No**.

Payoff vectors: The payoff vector is defined as follows. If the history is $(\tilde{w}, \text{Left}, \text{Yes})$, the payoff vector is w' to be defined; if it is $(\tilde{w}, \text{Left}, \text{No})$, the payoff vector is w'' to be defined; if it is $(\tilde{w}, \text{Right}, \text{Yes}, \text{Right}, \text{Yes})$, the payoff vector is \tilde{w} ; if it is $(\tilde{w}, \text{Right}, \text{Yes}, \text{Right}, \text{No})$, the payoff vector is \tilde{w}^2 ; if the history is $(\tilde{w}, \text{Right}, m_2, \text{Left}, \text{Yes})$, the payoff vector is \underline{w}^2 for $m_2 = \text{Yes}, \text{No}$; if the history is $(\tilde{w}, \text{Right}, m_2, \text{Left}, \text{No})$, it is \bar{w}^2 (to be defined), ; if the history is $(\tilde{w}, \text{Right}, \text{No}, \text{Right}, \text{Yes})$, it is \underline{w}^3 ; finally, if the history is $(\tilde{w}, \text{Right}, \text{No}, \text{Right}, \text{No})$, it is \underline{w}^2 . Modifying the polyhedron $P(\delta)$ if necessary, we can assume that all payoff vectors are vertices of this polyhedron.

The vertices w', w'' are such that both player 1 and 2 are indifferent between them and player 3 strictly prefers w' to w'' . (It may be necessary to modify the definition of $P(\delta)$ by adding a pair of vertices w' and w'' with the required property). Therefore, any strategy of player 2 when Nature chooses **Left** in the second stage will be optimal.

- *Strategies:* Given w , the strategy profile $\sigma := (\sigma_1, \sigma_2, \sigma_3)$ defined as follows is a sequential equilibrium.

$$\begin{aligned} \sigma_3(\emptyset) &= w, \\ \sigma_2(\tilde{w}, \text{Left}) &= \begin{cases} \text{Yes if } w = \tilde{w} \text{ and } w \text{ minimizes 3's payoff over } P(\delta), \\ \text{No otherwise,} \end{cases} \\ \sigma_2(\tilde{w}, \text{Right}) &= \begin{cases} \text{Yes if } w = \tilde{w}, \\ \text{No otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \text{Right}, \text{Yes}, \text{Right}) &= \begin{cases} \text{Yes if } w = \tilde{w} \\ \text{No otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \text{Right}, \text{Yes}, \text{Left}) &= \begin{cases} \text{No if } w = \tilde{w} = \underline{w}^2 \\ \text{Yes otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \text{Right}, \text{No}, \text{Right}) &= \begin{cases} \text{No if } w = \tilde{w} \\ \text{Yes otherwise,} \end{cases} \\ \sigma_1(\tilde{w}, \text{Right}, \text{No}, \text{Left}) &= \begin{cases} \text{No if } w \neq \tilde{w} \text{ and } \tilde{w} = \hat{w}^1 \text{ or } \hat{w}^2 \\ \text{Yes otherwise,} \end{cases} \end{aligned}$$

where \hat{w}^1, \hat{w}^2 are to be defined.

Observe that player 1 is indifferent between \bar{w}^2 and \underline{w}^2 . Since player 1 is also indifferent between \underline{w}^3 and \underline{w}^2 as well as between \tilde{w} and \tilde{w}^2 , this ensures that any strategy of player 1 will be optimal for every possible history.

To verify incentives of player 2 when Nature has chosen **Right**, we consider the cases $w = \tilde{w}$ and $w \neq \tilde{w}$ separately.

- $w = \tilde{w}$: By choosing **Yes**, player 2 faces a compound lottery; with probability $1 - \xi$, the compound lottery yields a lottery on $\{\tilde{w}, \tilde{w}^2\}$ with high probability on \tilde{w} ; with probability ξ , the compound lottery yields a lottery on $\{\bar{w}^2, \underline{w}^2\}$; by choosing **No**, player 2 also faces a compound lottery; with probability $1 - \xi$, it yields a lottery on $\{\underline{w}^3, \underline{w}^2\}$ with high probability on \underline{w}^2 ; with probability ξ , it yields a lottery on $\{\bar{w}^2, \underline{w}^2\}$. Therefore, it is clear that player 2 prefers **Yes**, unless $w = \underline{w}^2$. In that case, however, choosing **Yes** is optimal because it guarantees that the probability of \bar{w}^2 is high, while it is low otherwise.
- $w \neq \tilde{w}$: By choosing **No**, player 2 faces a compound lottery; with probability $1 - \xi$, it yields a lottery on $\{\underline{w}^3, \underline{w}^2\}$ with high probability on \underline{w}^3 ; with probability ξ , it yields a lottery on $\{\bar{w}^2, \underline{w}^2\}$; by choosing **Yes**, player 2 also faces a compound lottery; with probability $1 - \xi$, it yields a lottery on $\{\tilde{w}, \tilde{w}^2\}$ with high probability on \tilde{w}^2 ; with probability ξ , it yields a lottery on $\{\bar{w}^2, \underline{w}^2\}$. Therefore, player 2 prefers **No**, unless \tilde{w} is one of the two extreme points \hat{w}^1, \hat{w}^2 for which both $\tilde{w}_2^2 = \underline{w}_2^3$. In that case, however, choosing **No** is optimal because it guarantees that the probability of \bar{w}^2 is high, while it is low otherwise.

Finally, the incentives of player 3 are immediate to verify.

A.3.3 Cheap, *coarse* and imperfect sequential communication

If a player has only a fixed finite number of actions and there is only a fixed finite number of public outcomes, his ability of communicating a large set of vertices in a single period is quite limited. To accommodate for this, we shall now further modify the extensive form.

By ($2n$ -wise) full rank, each player has a pair of actions that induce distinct probability distributions over the set of public outcomes (given any fixed action profile of his opponents). Pick a mapping $f : Y \rightarrow [0, 1]$ such that the expected value of f conditional on one of the two actions is different from the expected value conditional on the other action. For any $\varepsilon > 0$, we may then find an integer T and a threshold $\mu \in [0, 1]$ such that, replacing each single period in which a given player i moves by a sequence of periods of length T in which this same player moves, the probability that the average value of f conditional on one (resp. the other) action being played in all these periods is above (resp. below) μ is less than $\varepsilon/(m+1)$, where 2^{m+1} is an upper bound for the number of vertices of the polyhedron $P(\delta)$. We introduce a set of binary *messages*, **a** or **b**, such that the sequence of actions of length T is mapped into one or the other binary message depending on whether the average value of f over these T periods is above or below μ .

In the communication game from Section A.3.2, player i 's action set was the set of vertices. Now, we replace it with sequences of $(m+1)T$ stage-game actions. We associate with any such sequence of stage-game actions the $(m+1)$ successive binary messages (there are then 2^m such sequences of binary messages), and we associate with any such sequence of binary messages a specific vertex of the polyhedron $P(\delta)$. The probability that the vertex determined by a sequence of binary messages is different from the vertex encoded in the stage-game actions is less than ε . The specific relationship between ε and $(m+1)T$ is examined in Section A.4.

It follows from Section A.3.2 that as far as there is a positive probability that the intended vertex will be designated, player i has an incentive to keep on choosing the actions corresponding to the intended vertex. However, it is possible that, during those $(m+1)T$ periods, it becomes apparent that the vertex that will be designated will differ from the vertex that was intended. In this case, there is no reason to expect that it is optimal for player i to keep on choosing the action corresponding to the intended vertex, and we leave the optimal continuation strategy unspecified, noting that it does not affect the probability of successful communication.

A.3.4 Cheap, coarse and imperfect *simultaneous* communication

Suppose now that (as in the repeated game), all players $i = 1, \dots, n$ choose actions in every period of the extensive-form game. This creates the possibility of “signal-jamming” by players

other than the player that was supposed to take an action according to the construction from Section A.3.3. We shall now assume that, as in the repeated game, players' action set in every period is A_i , all i , and that the public outcomes are drawn according to the distribution $\pi(\cdot|a)$. However, we still ignore flow payoffs.

We address next the following problem: Given some mapping $f : Y \rightarrow [0, 1]$, consider the one-shot game where each player j chooses an action $\alpha_j \in \Delta A_j$, but his payoff depends only on the message, which is equal to \mathbf{a} with probability $f(y)$, where y is the public outcome that results from the action profile that is chosen; in the repeated game, the message is equal to \mathbf{a} if the public randomization device takes (in the following period) a value $x \in [0, f(y)]$. Fix one player, say player i . Can we find f and action profiles α^* and β^* such that

- (i) for every $j \neq i$, actions of player j do not affect the probability of message \mathbf{a} , given that his opponents take the action profile α_{-j}^* or β_{-j}^* ;
- (ii) player i maximizes the probability of message \mathbf{a} by taking action α_i^* , given that his opponents take the action profile α_{-i}^* , and player i maximizes the probability of message \mathbf{b} by taking action β_i^* , given that his opponents take the action profile β_{-i}^* ;
- (iii) the probability of message \mathbf{a} is (strictly) higher under α^* than under β^* ?

We argue here that the answer is positive, under the $2n$ -wise full-rank condition that we have imposed. Note that this will turn out the only place in the proof in which we refer to the $2n$ -wise full rank. We take first an arbitrary pure action profile α^* . Then we define β^* as an arbitrary pure action profile that differs from α^* only by the action of player i . Finally, we will define the values $\{f(y) : y \in Y\}$ as a solution of a system of

$$m_i + 2 \sum_{j \neq i} (m_j - 1)$$

linear equations, which we shall now describe. In this system of equations, there is one equation that corresponds to the action profile α^* , one equation that corresponds to β^* , one equation for every pure action profile that differs from α^* (and so from β^*) only by the action of player i , and one equation for every pure action profile that differs from α^* or β^* only by the action of one of the players $j \neq i$. Every pure action profile determines a probability distribution over public signals $y \in Y$. We take the probability assigned to the public signal y as the coefficient of the

unknown $f(y)$, and the right-hand constant is equal to one of two numbers l or h , where $l < h$. It is h for α^* and any pure action profile that differs from α^* only by the action of one of the players $j \neq i$, and it is l for the remaining action profiles.

This system of equations has a solution because the vectors of coefficients of particular equations are linearly independent (by the $2n$ -wise full-rank condition). If it happens that the values $\{f(y) : y \in Y\}$ do not belong to $[0, 1]$, we may replace them with $\{f'(y) : y \in Y\}$ where

$$f'(y) = cf(y) + d,$$

for some constants $c > 0$ and d ensuring that $\{f'(y) : y \in Y\} \subseteq [0, 1]$. This changes the right-hand constants to $l' = cl + d$ and $h' = ch + d$, but the property that $l' < h'$ is preserved.

By construction, if player i takes action α_i^* the probability of message \mathbf{a} is equal to h given the other players take the action profile α_{-i}^* ; if he takes any other action, it is equal to l . On the other hand, no other player $j \neq i$ can unilaterally affect the probability of message \mathbf{a} ; it is h independently of his action. Similarly, if player i takes action β_i^* (or any action other than α_i^*) the probability of message \mathbf{a} is equal to l given that the other players take the action profile $\beta_{-i}^* = \alpha_{-i}^*$ (it is h if he takes action α_i^*), and no other player $j \neq i$ can unilaterally affect the probability of message \mathbf{a} . This yields properties (i)-(ii); the probability of message \mathbf{a} is equal to h and l contingent on α^* and β^* , respectively; thus, (iii) follows from the assumption that $l < h$.

[We may now combine this argument with the argument of the Section A.3.3 to ensure that, by picking mT large enough, the result of Section A.3.3 extend to this richer environment.]

Notice finally that the method of interpreting signals as messages described in this section is relatively simple due to the communication protocol that allowed only one player to “speak” at a time. If several players were allowed to speak at a time, we would have to design a method of interpreting signals as message profiles such that players can choose their own messages but cannot affect the messages of other players.

A.3.5 *Costly, coarse and imperfect simultaneous communication*

Players’ payoffs in the repeated game also depend on the flow payoffs in the communication phase, as modeled by the extensive-form game described in Section A.3.4. Those flow payoffs have been

ignored so far, yet they may affect players' incentives to communicate the vertex, which is the continuation payoff vector in the repeated game. This difficulty is easy to deal with.

Suppose that in the first period that follows any communication phase, with probability $\phi > 0$, instead of playing the strategies corresponding to the vertex that the communication phase has designated, players play strategies that yield another payoff vector $v \in W$, which also depends on the communication phase and which will be defined shortly. More precisely, if the outcome of the public randomization device in this period is less than ϕ , the payoff vector v is enforced; otherwise, the vertex designated by the communication phase is enforced. The vector v is contingent on the public outcomes observed in the communication phase, so that it makes every player indifferent (in every period of the communication phase) across all actions given any action profile of the opponents. The $2n$ -wise full-rank condition⁸ and the full-dimensionality condition guarantee the existence of v with the required property, provided that the differences in flow payoffs in the communication phase are small enough compared to the continuation payoff vector v , even if it is received only with the probability ϕ . However, it is indeed the case for every $\phi > 0$ if δ is large enough. Further, if ϕ is small and communication phases are sufficiently infrequent, the repeated game payoff vector approximates that for $\phi = 0$.

More precisely, we can assume that ϕ is of order proportional to the flow payoffs during the communication phase, i.e. it is proportional to

$$1 - \delta^{N(\delta)}.$$

A.4 The Length of Communication Phase

To establish our result, it remains to show that

$$\lim_{m \rightarrow \infty} (\delta_m)^{N(\delta_{m+1})} = 1 \tag{5}$$

where

$$\delta_m := 1 - \frac{1}{2^m}.$$

⁸Only individual full-rank, referring to the terminology from Fudenberg et al., suffices here.

Indeed, given our sequence of polyhedra $P(\delta_m)$, $m = 1, 2, \dots$, we shall define $P(\delta)$ as $P(\delta_{m+1})$ for every δ such that $\delta_m \leq \delta < \delta_{m+1}$. Then

$$(\delta_m)^{N(\delta_{m+1})} \leq (\delta)^{N(\delta)} \leq 1,$$

and the first part of (1) follows from (5). It is straightforward to guarantee that the second part of condition (1) is satisfied, as one can always take an $M(\delta)$ that is large enough.

Recall that the polyhedra $P(\delta_m)$ were chosen so that the number of their vertices is no larger than 2^{m+1} (it is even a polynomial expression of m). If players were able to communicate binary messages accurately, this would already conclude the proof of (5); indeed only $m + 1$ periods would be necessary for any vertex to be communicated, and so the communication phase would have to last only $m + 3$ (two periods have been included for **Yes** or **No** messages following the sequence of messages that determines a vertex); it follows now from

$$(1 - 2^{-m})^{m+3} \rightarrow 1$$

that discounting from one regular phase to the next is arbitrarily close to one.

Because communication is imperfect, we must now examine more carefully the relationship between ε and $(m + 1)T$ from Section A.3.3. Imperfect communication creates the following difficulty. Since players do not communicate accurately when they take the sequence of action profiles that corresponds to a vertex of the polyhedron $P(\delta_m)$, they end up only with a probability distribution over the messages that correspond to various vertices of $P(\delta_m)$, and the expected value of this probability distribution typically does not coincide with the vertex to be communicated.

To address this problem, we first slightly modify the polyhedron $P(\delta_m)$ so that the distance of any vertex of $P(\delta_m)$ to the boundary of W can be bounded by a constant times $1 - \delta_m$, preserving the property that $W'_m \subseteq P(\delta_m) \subseteq W$ where $W'_m := W_{\delta_m}$ defined by formula (4). In the case of $n = 2$, replace each vertex of $P(\delta_m)$ with a vertex that lies on the line joining the original vertex with the origin of W whose distance from the boundary of W is equal to a half of the distance between W and W'_m ; moreover, for every edge of the boundary of $P(\delta_m)$, add a vertex that lies on the line joining the origin of W and the middle of this edge whose distance from the boundary of W is also equal to a half of the distance between W and W'_m (see Figure 5). The modified polyhedron $P(\delta_m)$ contains W'_m , because the point from its edge that is closest to bdW'_m coincides

with the intersection point of this edge and an edge of the original $P(\delta_m)$. Moreover, the modified $P(\delta_m)$ has only twice as many vertices as the original $P(\delta_m)$.

In the general case of $n \geq 2$, see Böröczky (2001) for a rigorous proof of the existence of a polyhedron $P(\delta_m)$ such that $W'_m \subseteq P(\delta_m) \subseteq W$ and the distance of any vertex of $P(\delta_m)$ to the boundary of W can be bounded by a constant times $1 - \delta_m$. By a similar argument to that used in Section A.3.2, it can also be assumed that all terminal payoff vectors of the communication game described in that section are vertices of $P(\delta_m)$. This at most doubles the number of vertices of the polyhedron $P(\delta_m)$, so the number is bounded by 2^{m+2} .

Now, suppose that players communicate the vertices of another, larger and homothetic polyhedron $Q(\delta_m)$.⁹ If the communication is sufficiently accurate, the sequence of action profiles that corresponds to a vertex of $Q(\delta_m)$ induces a probability distribution over the messages that correspond to various vertices of $Q(\delta_m)$ with the expected value that is close to the desired vertex; if the desired vertex is communicated successfully at least with probability $1 - \varepsilon$, then the distance of this expected value to the desired vertex is proportional to ε (of course, this expected value belongs to the convex polyhedron $Q(\delta_m)$). Denote by $Q'(\delta_m)$ the polyhedron spanned by the expected values of the probability distributions induced by sequences of action profiles that correspond to all vertices of $Q(\delta_m)$. By taking ε small enough, but still proportional to $1 - \delta_m$, we can make sure that the polyhedron $Q'(\delta_m)$ contains $P(\delta_m)$. Thus, the vertices of $P(\delta_m)$ can be communicated in expectation in the communication phase in which players communicate the vertices of $Q(\delta_m)$, and the vertex to be communicated is determined by the public randomization device from the first period of the communication phase.

Finally, recall that we require ε to be small also for another reason. The probability of receiving the opposite binary message to one that a player is prescribed to communicate (through the sequence of action profiles α^* or β^*) must be low compared to the differences in coordinates of $Q(\delta)$ whenever they are different. Otherwise, players could not have incentives to play the prescribed strategies in the communication game.

However, $Q(\delta_m)$ can be easily constructed such that those differences are at least of order

$$x'_m = r' \left(1 - \cos \frac{\pi}{2^{m+1}} \right),$$

⁹The assumption that $Q(\delta_m)$ and $P(\delta_m)$ are homothetic guarantees all payoff vectors used the communication phase described in Section A.3.2 are vertices of the polyhedron $Q(\delta_m)$, as so they were for the polyhedron $P(\delta_m)$.

where r' stands for the “radius” of the polyhedron $Q(\delta_m)$ (in Figure 6, we depict the case in which the difference is the lowest possible if the extreme points of the circle are vertices of $Q(\delta_m)$). Since

$$\frac{x'_m}{r' \left(\frac{\pi}{2^{m+1}}\right)^2} \rightarrow_m \frac{1}{2},$$

it suffices to take an ε proportional to 2^{-2m-2} . By construction, this estimates remain unaltered when we add new vertices in order to include all terminal payoff vectors of the communication game into the set of vertices.

In the general case of $n \geq 2$, we first take $P(\delta_m)$ as any polyhedron that contains W'_m and that is contained in W , and such that the distance of any vertex of $P(\delta_m)$ to the boundary of W can be bounded by a constant times $1 - \delta_m$. We take $Q(\delta_m)$ to a larger and homothetic polyhedron, exactly as in the case $n = 2$, and then we perturb vertices of $Q(\delta_m)$ (coordinate by coordinate) to guarantee that the differences in coordinates of $Q(\delta)$ (whenever they are different) are proportional to 2^{-2m-2} .

The sufficient accuracy of communication can be achieved only if the action profiles α^* and β^* (corresponding to each binary message) are played repeatedly for a sufficiently large number of periods. We shall now estimate that number of periods.

By Hoeffding’s inequality, players need a number of periods (in which they play α^* or β^*) that is proportional to $-\log \varepsilon$ to make sure that the binary message is interpreted as **a** or **b** at least with probability $1 - \varepsilon$. Therefore, since the number of vertices of $Q(\delta_m)$ is bounded by 2^{m+2} , it suffices to have a number of periods that is proportional to

$$-(m+2) \log \left(\frac{\varepsilon}{m+2} \right)$$

to communicate successfully any vertex of $Q(\delta_m)$ at least with probability $1 - \varepsilon$; indeed, if every binary message (out of $m+2$ of them) is successful at least with probability

$$1 - \frac{\varepsilon}{(m+2)},$$

then the desired vertex is communicated successfully at least with probability

$$\left(1 - \frac{\varepsilon}{m+2} \right)^{m+2} \geq 1 - \varepsilon.$$

Since we take an ε proportional to $(2^{-m-1})^2$, one needs only a number of periods that is proportional to

$$-(m+2) \log \left(\frac{1}{(m+2)2^{2m+2}} \right) = (m+2) \log ((m+2)2^{2m+2})$$

to communicate any desired vertex of $P(\delta_m)$ in expectation, and so (5) follows from the fact that

$$\left(1 - \frac{1}{2^m}\right)^{(m+3) \log[(m+3)2^{2m+4}]} \rightarrow_m 1.$$




Figure 5: The modified polyhedron $P(\delta)$. The boundary of the original $P(\delta)$ is depicted by the straight line, and the boundary of the modified $P(\delta)$ is depicted by the kinked line.




Figure 6