ASYNCHRONOUS REPEATED GAMES WITH RICH PRIVATE MONITORING AND FINITE PAST

MARCIN PESKI

Abstract. This paper analyzes asynchronous repeated games with private and rich monitoring. We assume that strategies have *finite past*, i.e., in each period, continuation strategies must be measurable with respect to finite partitions of past histories. This class includes finite automata and bounded recall strategies. Additionally, we assume that the monitoring has an infinite number of signals. We show that any equilibrium with finite past and generic infinite monitoring has to satisfy a version of the belief-free property: in each period t, the set of best responses does not depend on the information received before period t, with a possible exception of the information received in the first periods of the game. Under an additional payoff smoothness assumption, the equilibrium strategy are essentially past-independent: each period's action depends only on the information received immediately prior to the choice of the action.

1. Introduction

We study asynchronous repeated games with private monitoring. There are two players who take actions in alternating periods. Before taking an action, one of the players privately observes a signal that conveys information about current payoffs, as well as the signals of the opponent in the previous and subsequent periods. We focus on strategies that have *finite past*: the continuation strategies in each period depend only on finite partitions of histories that occurred before that period. Second, we assume that the monitoring is private and that it has an infinite number signals. The assumptions captures some features of real-world interactions. On the one hand, it is unlikely that players are able to process infinite amounts of information in order to implement their strategies. On the other hand, players must be aware of the fact that there is always a possibility of receiving a signal that carries a slightly different information than any other signal. Even if the differences are small and have a small probability of occuring, in the private monitoring case, players cannot be stopped from using such signals in tuning up their best responses. (If the strategies depend on a finite number public signals, the best responses are public and depend on exactly the same set of signals.)

Department of Economics, University of Texas at Austin, 1 University Station #C3100, Austin, Texas 78712. I am grateful to Tom Wiseman and Max Stinchombe for helpful comments. Financial support from the National Science Foundation (grant 0841848) is gratefully acknowledged.

The finite past assumption bites only because rich monitoring has an infinite number of signals.

We distinguish two properties of infinite private monitorings. The monitoring is *rich*, if the set of all plausible beliefs that may occur in the game is an *open and connected* subset of the space of belief-payoff tuples. As a special case, the monitoring is *extremely rich* if any feasible (i.e., consistent with the players' strategies) belief is induced by some signal. In particular, extremely rich monitoring does not impose any restriction on beliefs except for the fact that they need to be consistent with the players' strategies. It turns out that rich and extremely rich monitorings are generic among all private monitorings with an infinite number of signals (Theorem 1).

The main result, Theorem 2, says that any finite past equilibrium in a game with a rich monitoring has to satisfy a version of the belief-free property (Piccone (2002), Ely and Valimaki (2002)): in each period t, the set of best responses does not depend on the information received before period t, with the possible exception of the information received in the second period of the game. Under the additional payoff smoothness assumption, the main result has a simple corollary: the equilibrium strategies are essentially past-independent: each period's action depends only on the current information.

The argument relies on the double nature of signals in asynchronous games. We show that if the set of best responses in period t non-trivially depends on information from period $s < t$, then it must also depend on information in the periods immediately before and/or after period s. At the same time, we show that best responses may locally depend only on the information from one period. The result follows from a conflict between the two observations.

The results provide theoretical foundations for two classes of equilibria of asynchronous games. In the belief-free equilibria, the optimality of players' behavior is robust to misspecification of information, lack of common prior, or non-Bayesian preferences over uncertainty (Bergemann and Morris (2007), Horner and Lovo (2009)). In the repeated games with private monitoring, it is relatively easy to analyze the equilibrium properties of strategies if they are belief-free. Sometimes, but not always, the payoffs in belief-free equilibria exhaust the full range of the folk theorem payoffs Ely, Hörner, and Olszewski (2005). Past-independent strategies are motivated by complexity costs (see Bhaskar and Vega-Redondo (2002) and references therein). Due to their simplicity, they turn out to be particularly helpful in the applied and computational analysis (for example, Ericson and Pakes (1995)).

Although this paper is concerned with the properties of equilibrium strategies rather than payoffs, it is not difficult to construct examples in which the belief-free characterization of finite past payoffs reduces the range of possible equilibrium payoffs. This contrasts with

Yoon (2001) who shows the folk theorem for asynchronously repeated games with perfect monitoring. In addition, two recent papers (Horner and Olszewski (2009) and Mailath and Olszewski (2008)) show folk theorems in simultaneous repeated games with almost perfect monitoring and bounded memory (see also (Mailath and Morris (2002) and Mailath and Morris (2006))). We believe that the main difference between the asynchronous and the simultaneous models does not lie in the protocol of moves, but rather in the nature of monitoring. We discuss a modification of the standard simultaneous-move model that could lead to the belief-free property along the same lines as the current result.

A recent paper (Bhaskar, Mailath, and Morris (2009)) has considered games with perfect information (a class that includes the asynchronous games). It is shown that all equilibria that have bounded recall and that can be purified are essentially past-independent. Because our payoff-smoothness assumption is closely related to purification, the main result of Bhaskar, Mailath, and Morris (2009) is closely related to our corollary. On the one hand, the bounded recall assumption is a considerably stronger than finite past, especially in the context of asynchronous repeated games. On the other, their result is formulated for games with perfect monitoring, and, as we explain in Section 6.2, their notion of repeated game purification is weaker than the one that is needed to interpret our results.

In a companion paper (Peski (2009)), we show that any finite past equilibrium of the simultaneous move repeated game in finite past strategies is essentially a series of stage-game equilibria. For that result, we need the payoff smoothness assumption, but the monitoring has to satisfy weaker assumptions than richness (the set of beliefs has to be connected, but not necessarily open).

The next section presents the model. Section 3 defines rich and extremely rich monitoring. Section 5 contains the main results. Section 6 discusses related issues: the existence and denseness of rich monitoring, the relation of the current result to the literature on games with lexicographic costs of memory, and the possibility of extending the results to simultaneous move games.

2. Repeated game

There are 2 players $i = 1, 2$. For each player i, let $T_i = \{2k + i : k = 0, ...\}$ and, for each t, $T_i(t) = \{s \in T_i : s < t\}$. The timeline of the game is illustrated on Figure 2: In each period $t \in T_i$, player i

- privately observes a signal $\omega_t \in \Omega_i$ drawn from a distribution $\rho_i(\omega_{t-1}, a_{t-1}) \in \Delta\Omega_i$ that depends on the previous period's action and the signal of the opponent. We assume that Ω_i is equipped with a sigma-algebra of measurable sets. In order to fix arrow

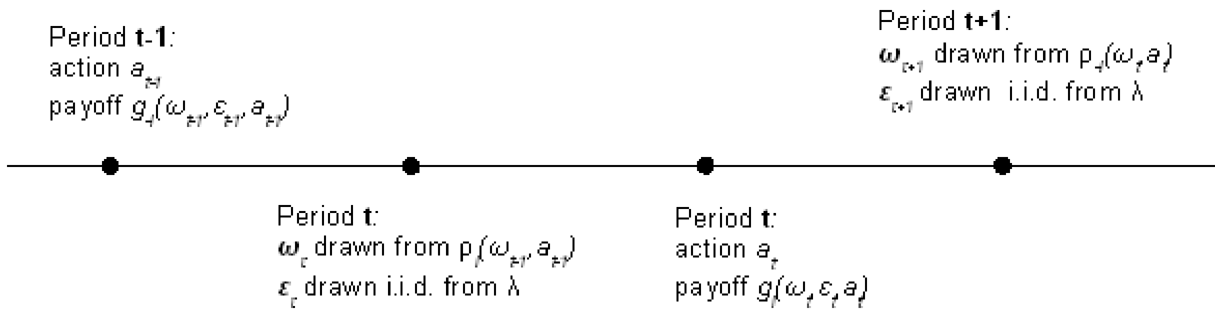

1.png

the initial conditions, we assume that signal ω_1 of player 1 is drawn from distribution $\mu_1 \in \Delta\Omega_1$,

- privately observes ε_t and γ_t, independently drawn from the Lebesgue measure $\lambda \in \Delta[0,1]$. We will interpret ε_t as an idiosyncratic shock to the payoffs and γ_t as a private randomization device,
- chooses action a_t from a finite set A_i,
- receives payoff $g_i(\omega_t, \varepsilon_t, a_t)$, where $\sup_{\omega,\varepsilon,a} |g_i(\omega, \varepsilon, a)| < \infty$. We assume that the payoff depends on signal ω_t, idiosyncratic shock ε_t, and action a_t. Because signal ω_t is affected by the opponent's action from the previous period, the current payoff is indirectly affected by the opponent's behavior.

For each $t \in T_i$, let $H_t = \left(\Omega_i \times [0,1]^2 \times A_i\right)^{T_i(t)} \times \Omega_i \times [0,1]^2$ be the history observed by player i before taking period t action. A typical element of H_t is denoted as $h_t = (...\omega_{t-2}, \varepsilon_{t-2}, \gamma_{t-2}, a_{t-2}, \omega_t, \varepsilon_t, \gamma_t)$. Let $H_{i,\infty}$ be the set of infinite histories.

Players discount the future with equal discount factors $\delta < 1$.[1] The repeated game payoffs of player i given infinite private history $h_{i,\infty} = (\omega_t, a_t, t \in T_i)$ are equal to

$$G_i(h_{i,\infty}) = (1-\delta) \sum_{t \in T_i} \delta^t g_i(\omega_t, \varepsilon_t, a_t).$$

[1] The results won't change if the discount factors differ across periods and players.

A *pure strategy* of player i is a mapping $\sigma^i : \bigcup_{t\in T_i} H_t \rightarrow A_i$ with the interpretation that period $t \in T_i$ action of player i is equal to $\sigma^i_t(h_t)$. Let Σ_i be the set of pure strategies of player i. In the subsequent analysis we assume without loss of generality that the repeated game strategies are always pure.

The two random shocks ε_t and γ_t play two different roles. The payoff shock ε_t is used in Corollaries 1 and 2 to purify the repeated game strategies in the way similar to Bhaskar, Mailath, and Morris (2008) and Bhaskar, Mailath, and Morris (2009). The payoff shock plays no role in the proof of main result, Theorem 2. Because Corollaries 1 and 2 rely on Theorem 2, we need to state the result in the general model in which the payoff shocks are present.

Although the private randomization device γ_t does not affect payoffs, it can be used to for the randmoization of pure strategies. Due the presence of γ_t, the restriction to private strategies is without the loss of generality.

Together with the initial distribution μ_1, the strategy profile $\sigma = (\sigma^1, \sigma^2)$ induces a distribution over joint histories $\pi^{\sigma,\mu_1} \in \Delta(\times_i H_{i,\infty})$. The expected payoff of player i given strategy profile $\sigma = (\sigma^1, \sigma^2)$ is equal to $G_i(\sigma;\mu_1) = \pi^{\sigma,\mu_1}[G_i(h_{i,\infty})]$. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is an *equilibrium*, if for each player i, for each strategy σ'_i, $G_i(\sigma;\mu_1) \geq G_i(\sigma'_i, \sigma_{-i};\mu_1)$.

A *continuation strategy* in period t is a pair of action and a strategy $(a, s) \in \Sigma^*_i = A_i \times \Sigma_i$. For example, for each strategy σ^i of player i, for each $t \in T_i$, private history h_t, $\bar{\sigma}^i(h_t) := (\sigma^i_t(h_t), \sigma^i(h_t, \sigma^i_t(h_t))) \in \Sigma^*_i$ is a continuation strategy after history h_t. For each history h_t, let $\pi^{\sigma,\mu_1}_{h_t} = \text{marg}_{H_{t-1}\times A_{-i}} \pi^{\sigma,\mu_1}(.|h_t) \in \Delta(H_{t-1} \times A_{-i})$, where $\pi^{\sigma,\mu_1}(.|h_t)$ is (a version of) the conditional distribution π^{σ,μ_1} given h_t. Denote the expected conditional continuation payoff of player i from continuation strategy $(a, s) \in \Sigma^*_i$, given history h_t as

$$G_i(a_t, s|h_t, \sigma) = (1-\delta)\, g(\omega_t, \varepsilon_t, a_t) + \delta \pi^{\sigma,\mu_1}_{h_t}\left[G_i\left(s, \sigma^{-i}(h_{t-1}, a_{t-1}); \rho_{-i}(\omega_t, a_t)\right)\right].$$

Here, $\pi^{\sigma,\mu_1}_{h_t}[.]$ is an expectation over histories and actions taken with respect to probability measure $\pi^{\sigma,\mu_1}_{h_t}$. Continuation strategy $b \in \Sigma^*_i$ is the *best response* after history h_t if, for each $(a', s') \in \Sigma^*_i$, $G_i(a, s, \sigma^{-i}|h_t) \geq G_i(a', s', \sigma^{-i}|h_t)$. If profile (σ^1, σ^2) is an equilibrium, then for π^{σ,μ_1}-almost all histories h_t, $\bar{\sigma}^i(h_t)$ is the best response after h_t.

3. Monitoring

For each probability space (X, μ), let $L^2(X, \mu)$ be the set of all measurable, μ-square integrable functions with L^2-norm $\|.\|_2$. Then, for each $f, g \in L^2(X, \mu)$, f, g, and fg are μ-integrable. Let $L^*(X, \mu) \subseteq L^2(X, \mu)$ consist of all functions f such that $\mu[f] := \int_X f d\mu = 1$ and that $f \geq 0$, μ-almost surely. Space $L^*(X, \mu)$ inherits its topology from $L^2(X, \mu)$.

We are going to assume that all monitoring technologies are absolutely continuous with respect to some fixed non-atomic probability measure on the signal space. For each player i, suppose that (Ω_i, μ_i) is a non-atomic probability space. Let $\mu_i^A \in \Delta A_i$ be the uniform distribution on player i's action set. Let λ be the Lesbegue measure on the interval $[0,1]$. For each i, let $\Gamma_i \subseteq L^* \left(\Omega_i \times \Omega_{-i} \times A_{-i}, \mu_i \times \mu_{-i} \times \mu_{-i}^A\right)$ be the space of functions ρ_i such that $\mu_i \left[\rho\left(.|\omega_{-i}, a_{-i}\right)\right] = 1$, $\mu_{-i} \times \mu_{-i}^A$-almost surely.

A *monitoring* is a pair $\rho = (\rho_1, \rho_2)$ of measurable functions $\rho_i \in \Gamma_i$ for each player i. Abusing notation, we refer to $\rho_i(\omega_{-i}, a_{-i})$ as a probability distribution over the space of player i's signals, Ω_i with μ_i-density $\rho_i(.|\omega_{-i}, a_{-i})$. Let $\Gamma = \Gamma_1 \times \Gamma_2$ be the space of the monitorings. We assume that Γ_i has the inherited norm topology and Γ has the product topology. Then, Γ is a Polish, and hence a Baire space.

Monitoring ρ has *full support* if for each player i, $\rho_i > 0$, almost surely. Take any monitoring ρ with full support. Each signal ω_i conveys three types of information:

- the likelihood of past signals and actions: Let $P_i^* = \left(L^*\left(\Omega_{-i}, \mu_{-i}\right)\right)^{A_{-i}}$ and define mapping $\upsilon_p^\rho : \Omega_i \to P_i^*$ as
$$\upsilon_p^\rho(\omega_i)(\omega_{-i}, a_{-i}) = \frac{\rho_i(\omega_i|\omega_{-i}, a_{-i})}{\left(\mu_{-i} \times \mu_{-i}^A\right)\left[\rho_i(\omega_i|., .)\right]}.$$
Then, $\upsilon_p^\rho(\omega_i)$ is well-defined for μ_i-almost all signals ω_i. Given the strategies of the opponents, prior beliefs held before period t, and signal ω_i, the likelihood $\upsilon_p^\rho(\omega_i)(\omega_{-i}, a_{-i})$ determines the ex post beliefs in period t about actions and signals of the opponents.
- information about current payoffs: Define $\upsilon_g(\omega_i) = g_i(\omega_i, .) \in L^*\left(A_i \times [0,1], \mu_{A_i} \times \lambda\right)$, and let $G_i^* = \upsilon_g^\rho(\Omega_i)$. We assume that G_i^* is finite.[2]
- the density function of future signals: Let $F_i^* = L^*\left(\Omega_{-i}, \mu_{-i}\right)$ and define mapping $\upsilon_f^\rho : \Omega_i \to \left(F_i^*\right)^{A_i}$:
$$\upsilon_f^\rho(\omega_i)(\omega_{-i}, a_i) = \rho_{-i}(\omega_{-i}|\omega_i, a_i).$$

Space $\Omega_i^* = P_i^* \times G_i^* \times \left(F_i^*\right)^{A_i}$ is assumed to have the product topology. Define
$$\upsilon^\rho(\omega_i) := \left(\upsilon_p^\rho(\omega_i), \upsilon_g(\omega_i), \upsilon_f^\rho(\omega_i)\right) \in \Omega_i^*.$$

We say that monitoring ρ is *extremely rich*, if for each player i, the inverse image of any open set $W \subseteq \Omega_i^*$ has positive probability, $\mu_i\left(\omega_i : \upsilon^\rho(\omega_i) \in W\right) > 0$.

Recall that a subset of a topological space is non-meagre if it contains a countable intersection of open and dense sets. Non-meagre subsets of Baire spaces are dense, and they are often used as a measure of genericity in infinitely dimensional spaces.

[2] All results hold if G_i^* is a finite union of open subsets of affine subspaces of $L\left(A_i \times [0,1], \mu_{A_i} \times \lambda\right)$. Minor proof modifications are required.

Theorem 1. *The set of extremely rich monitorings is a non-meagre subset of* Γ.

The proof can be found in Appendix C.

The main result, Theorem 2, requires a slightly weaker property, which we state next. Say that full support monitoring ρ is *rich*, if for each player i, there exists an open set $W_i^\rho \subseteq \Omega_i^*$ such that two conditions are satisfied:

- open support: $\upsilon^\rho(\Omega_i) \subseteq W_i^\rho$ and for each open $W \subseteq W_i^\rho$, $\mu_i(\omega_i : \upsilon^\rho(\omega_i) \in W) > 0$,
- connectedness: set $\{(p, f(a_i)) : (p, f) \in W_i^\rho, a_i \in A_i\} \subseteq P_i^* \times F_i^*$ is connected (i.e., it is not a union of two disjoint open sets).

Of course, if ρ is extremely rich, then $W_i^\rho = \Omega_i^*$, open support and connectedness are trivially satisfied, and ρ is rich.

4. Finite past

A player i's strategy σ^i has *finite past*, if in each period t, in each period t, there exists a finite partition Π_t of t-period histories H_t such that the t-period continuation strategy $\bar{\sigma}^i(h_t)$ is measurable with respect to Π_t. Equivalently, strategy σ^i has a finite past if in each period t, there are finitely many different continuation strategies, $|\{\bar{\sigma}(h_t) : h_t \in H_t\}| < \infty$.

Finite past generalizes an assumption that is often used in the repeated game literature. Say that a strategy σ^i is *implementable by a finite automaton*, if there exists a finite set of continuation strategies $\Sigma_i^0 \subseteq \Sigma_i^*$ such that $\bar{\sigma}(h_t) \in \Sigma_i^0$ for each $t \in T_i$ and each h_t. Clearly, a finite automaton has finite past, but not all finite past strategies can be implemented by finite automata.

Notice that finite past bites only because we assume that the space of signals is infinite. With finitely many signals, there are finitely many histories in each period, and, trivially, finitely many continuation strategies.

The assumption has a number of interpretations. First, finite past captures a notion of complexity of repeated game strategies: Complex strategies depend on infinitely many details of past histories, whereas simple strategies depend only on finite representation of the past.

Second, one can think about the finite past as an assumption about memory. In general, the implementation of a strategy may require players to remember an infinite amount of information (precisely, which of the infinitely many feasible histories took place.) If the latter is impossible, players are forced to use finite past strategies: they must replace infinitely many signals observed in any given period by a finite partition of the signal space.

Finally, we describe an important consequence of a finite past assumption. Take amny finite past strategy σ^i. Then, its its continuation stratregies in period $t \in T_i$ must be measurable with respect to the 'product partition' of observations coming from different periods: for each period $s \in T_i$,$s \leq t$, there are finite partitions Ψ_s of signal-action tuples $A_i \times \Omega_i \times [0,1]^2$ with the following property: $\bar{\sigma}^i(h'_t) = \bar{\sigma}^i(h''_t)$ for any two histories h'_t and h''_t such that for each s, $\left(a'_{s-2}, \omega'_s, \varepsilon'_s, \gamma'_s\right)$ and $\left(a''_{s-2}, \omega''_s, \varepsilon''_s, \gamma''_s\right)$ belong to the same element of partition Ψ_s.

Indeed, the argument comes from an induction on t. Suppose that continuation strategy in period $t-2 \in T_i$ is measurable with respect to the product of finite partitions Ψ_s for $s \in T_i(t)$. Let $\psi = (\psi_s) \in \times_{s \in T_i(t)} \Psi_s$ be an element of such a partition. Let $\xi_t = (a_{t-2}, \omega_t, \varepsilon_t, \gamma_t)$ be the information received before period t. Because of the measurability restriction on the continuation strategies in period 2, $\bar{\sigma}(h_{t-2}, \xi_t) = \bar{\sigma}\left(h'_{t-2}, \xi_t\right)$ for any two histories h_{t-2}, h'_{t-2}. Because of finite past, there are finitely many continuation strategies following histories $h_{t-2} \in \pi$, $\left|\{\bar{\sigma}(h_{t-2}, \xi_t) : h_{t-2} \in \pi\}\right| < \infty$. In particular, there exists a finite partition $\Psi_t(\psi)$ of signal-action tuples $A_i \times \Omega_i \times [0,1]^2$ with the following property: $\bar{\sigma}^i\left(h'_{t-2}, \xi'_t\right) = \bar{\sigma}^i\left(h''_{t-2}, \xi''_t\right)$ for any two histories $h'_t, h''_t \in \psi$ and observations ξ'_t and ξ''_t belong to the same elemnt of partition of $\Psi_t(\psi)$. Finally, take Ψ_t to be the finite partition generated by the union of finite partitions $\Psi_t(\psi)$ for all $\psi \in \times_{s \in T_i(t)} \Psi_s$.

In Appendix A, we show that there always exist equilibria with finite past.

5. Main result

The main result of this paper characterizes finite past equilibria when the monitoring is rich.

Theorem 2. *Suppose that the monitoring is rich. If σ is an equilibrium profile with finite past, then for each player i and each $t \in T_i$, if continuation strategy $b \in \Sigma_i^*$ is player i's best response after history h_t, then it is a best response after each history h'_t such that $\omega_t = \omega'_t, \varepsilon_t = \varepsilon'_t$, $\omega_2 = \omega'_2$, and $a_2 = a'_2$ (the last two equalities apply only when $i = 2$).*

Theorem 2 shows that finite past equilibria under rich monitoring have to satisfy a version of the belief-free property: period t best responses to the opponent's strategies do not depend on information received in periods $s \neq 2, t$. In particular, the equilibrium behavior of a player who misstated or lost information from period $s \neq 2, t$ would be a best response to the original information. The result does not depend on the discount factor, payoffs, nor on the specific form of the monitoring, as long as it is rich.

We explain the intuition behind Theorem 2. Suppose that strategy σ^i has finite past. As we explain above, the continuation strategies in period t are emasurable with respect to the product of finite partitions of information coming from all periods $s \in T_i, s \leq t$. Figure 5 presents an example of such a partition. The horizontal axis contains information coming from period s and the vertical axis contains information from all periods but s. We focus on two continuation strategies in period t, $\bar{\sigma}_A$ and $\bar{\sigma}_B$. Let A be the set of histories followed by continuation strategy $\bar{\sigma}_A$, and let B be the set of histories followed by strategy $\bar{\sigma}_B \neq \bar{\sigma}_A$. In

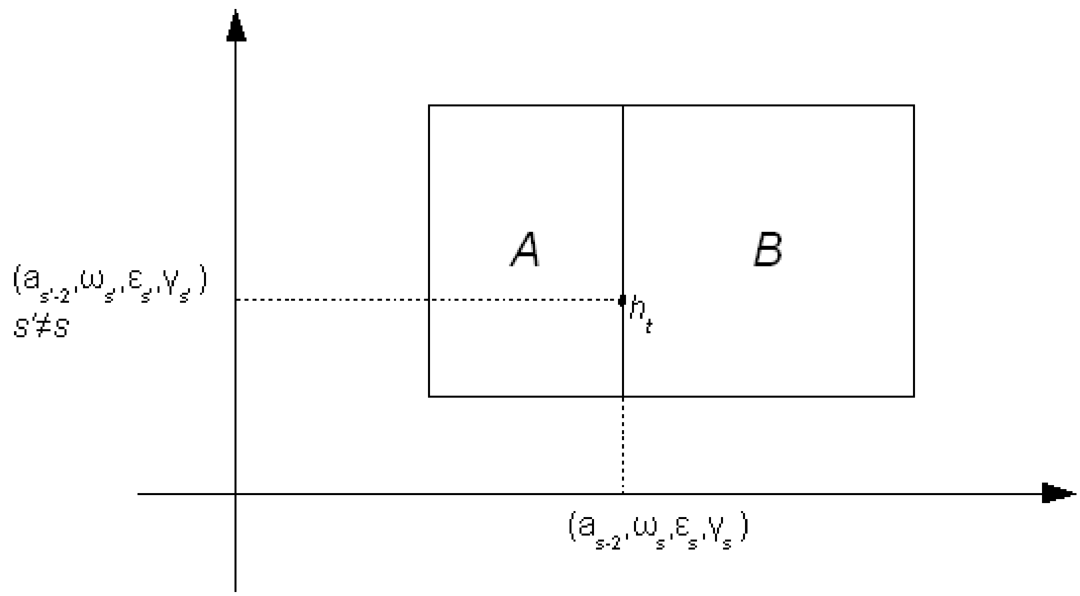

particular, for histories that belong to sets A and B, the choice between strategies $\bar{\sigma}_A$ and $\bar{\sigma}_B$ depends on action a_{s-2} and the information received at the beginning of period s.

Suppose that strategies $\bar{\sigma}_A$ and $\bar{\sigma}_B$ are continuation best responses after histories in sets, respectively, A and B. Additionally, suppose that the best responses in period $t > s$ depend non-trivially on the information in period s, and, in particular, player i strictly prefers $\bar{\sigma}_A$ to $\bar{\sigma}_B$ after some histor in set A. We argue below that because of rich monitoring, if sets A and B are not too small, then strategy $\bar{\sigma}_A$ must be strictly preferred to $\bar{\sigma}_B$ after all histories in A and, similarly, $\bar{\sigma}_B$ must be strictly preferred to $\bar{\sigma}_A$ on B with the exception of the histories on the shared boundary between A and B. In particular, locally around history h_t in Figure **??**, the preference between $\bar{\sigma}_A$ and $\bar{\sigma}_B$ depends only on the signal observed in period s, but *not* on any other signals.

We show that the situation illustrated on Figure **??** cannot happen when the monitoring is rich. Notice that s-period signal ω_s affects the best responses in period $t > s$ to the extent that it provides information about the action and signal of the opponent in period $s-1$ (through updating distribution $\rho_i(\omega_s|.)$) and/or about the signal in period $s+1$ (through distribution $\rho_{-i}(.|\omega_s, a_s)$). However, by the same observation, some information about the signal in period $s-1$ is contained in signal ω_{s-2}; similarly, some information about the signal and action in period $s+1$ is contained in signal ω_{s+2}. Because of the richness assumption, each minor variation in the signal ω_s can be replicated by an appropriate variation of signals ω_{s-2} and ω_{s+2}. Thus, if the best responses depend on s-period signal ω_s, they must depend on signals from other periods. This leads to a contradiction with the finite past property. The argument also explains the exceptional role of player 2: the period 2 signal provides information about player 1's signal ω_1 and action a_1 and this information is not replicated by any other signal of player 2.

We describe some difficulties that must be dealt with in order to formalize the above argument. In order to formally discuss minor variations of signals, we define a topology on the signal space Ω_i as the coarsest topology that makes mapping υ^{ρ} continuous. Then, the belief mapping that associates private histories with beliefs about the opponent's histories is continuous. Additionally to the notion of closeness, the choice of the topology equips the space of signals with affine structure. In particular, the belief mapping and the payoff from a continuation strategy is multilinear in the past $\upsilon_p^{\rho}(\omega_s)$ and future $\upsilon_f^{\rho}(\omega_s)$ parts of signal ω_s. The choice of the topology is helpful in the following steps:

First, we show that the finite partitions Π_s can be chosen in such a way that each element $U \in \Pi_s$ of the partition is either *almost open* (any open set V that has a non-empty intersection with U has a positive measure), or it is *almost nowhere-dense* (the measure of the intersection of any open V is equal to 0). Because the "variational" technique requires sufficient room, we use it only when A and B are products of almost open sets (the almost nowhere-dense sets are dealt with using a continuity argument). Almost open elements exist due to the richness of the monitoring (specifically, due to the openness of the support W_i^{ρ}).

Second, we can find almost open sets A and B whose boundaries have a non-empty intersection because of the connectedness of the rich monitoring.

Third, take A from the product partition that is a product of almost open sets and assume w.l.o.g. that period t payoff function $\upsilon_g^{\rho}(\omega_t)$ for all histories $h_t \in A$ is equal to g^* for some $g^* \in G_i^*$ (we can always choose the finite partitions in period t so that such a g_i^* exists). We use multilinearity and openness to show that if a continuation strategy $\bar{\sigma}_B$ is the best response after *some* history $h_t \in A$, then $\bar{\sigma}_B$ is the best response after *all* histories $h_t \in A$; if $\bar{\sigma}_B$ is *not* the best response after some history in A, it is not the best response after all histories in A.

Finally, notice that there is asymmetry between the amount of information about signal ω_s that is contained in signals ω_{s-2} and ω_{s+2}. The latter not only fully replicates the information about signal ω_{s+1} that is contained in the "future" part of signal ω_s, but ω_{s+2} additionally contains information about action a_{s+1}. On the other hand, because signal ω_{s-2} does not provide any information about action a_{s-1}, it cannot fully replicate the information contained in the "past" part of signal ω_s. Thus, the "variational" technique cannot be applied if the boundary between sets A and B at history h_t (Figure **??**) depends only on the belief about action a_{s-1}. However, because the distribution over period s-signal depends on signal ω_{s-1} *and* action a_{s-1} at the same time and the distribution over signals is non-atomic (and no signal ω_{s-1} has positive probability), we can always find a signal ω'_s such that the interpretation of information about action a_{s-1} depends on the (non-atomic) beliefs about signal ω_{s-1}. If sets A and B are equal to products of almost open sets, we can find history h'_t with such signal ω'_s that is close to h_t and that lies on the boundary between A and B. The "variational" technique can be applied to signal ω_{s-1}.

6. Extensions

6.1. **Payoff-smoothness.** Theorem 2 implies that in the finite past equilibrium, two different actions may be played after period t histories that differ only by information obtained in periods $s < t, s \neq 2$ only if the player is indifferent with respect to the two actions after each of the two histories. A stronger result can be obtained under the following assumption.

Assumption 1. *For each player i, each $g \in G_i^*$, each signal ω, any two actions a, a' of player i,*

$$\lambda \{g_i(\omega, \varepsilon, a) - g_i(\omega, \varepsilon, a') = c\} = 0 \text{ for each } c \in R.$$

Note that player i is indifferent between actions a and a' after history h_t, if and only if the difference between instantaneous payoffs from the two actions is equal to certain history dependent constant c. Assumption 1 ensures that the probability of such an event is equal to 0. Thus, the Assumption guarantees that in each period, player i has a strict best response for almost all realizations of private shock ε. This observation leads to a straightforward corollary to the main result. Say that player i's strategy σ^i is *essentially past-independent*, if for all periods $t \in T_i$, λ-almost all realizations $\varepsilon \in [0,1]$, π^{σ,μ^1}-almost all histories h_t and h'_t such that $\omega_t = \omega'_t$ and $\varepsilon_t = \varepsilon'_t = \varepsilon$,

$$\bar{\sigma}^i(h_t) = \bar{\sigma}^i(h'_t).$$

Corollary 1. *Suppose that the monitoring is rich and that Assumption 1 holds. If σ is an equilibrium profile with finite past, then for each player i, strategy σ^i is essentially past-independent.*

Proof. In case of player $i = 1$, the claim follows from Theorem 2 and the above remark. In case of player $i = 2$, the claim follows from Theorem 2, the above remark, the result about player $-i = 1$, and the fact that if the behavior of player 1 does not depend on information in periods 1 and 3, then player 2's best responses do not depend on information in period 2. □

Because any past-independent strategy has finite past, the Corollary provides a complete characterization of finite past equilibria.

6.2. **Purification.** Bhaskar, Mailath, and Morris (2008) and Bhaskar, Mailath, and Morris (2009) discuss the idea of Harsanyi's purification in the context of perfect dynamic games with perfect recall (a wider class than asynchronously repeated games) and a finite number of states. An equilibrium profile σ can be purified if there exists a sequence of games with smooth payoff shocks and a sequence of equilibria of these games such that the games and the equilibria converge to the original game and equilibrium σ. Bhaskar, Mailath, and Morris (2009) show that every equilibrium that can be purified by strategies with bounded recall must be Markovian.

We compare their result to ours. Our notion of purification is analogous to the one used in Bhaskar, Mailath, and Morris (2009). Say that a sequence of games g^n converges to game g, if

$$\lim_{n\to\infty} \sup_{i,\omega,\varepsilon,a} |g_i^n(\omega,\varepsilon,a) - g_i(\omega,\varepsilon,a)| = 0.$$

For each history , let $h_t^I = (..., \omega_{t-2}, a_{t-2}, \omega_t)$ denote the part of history h_t that consists only of own actions and informative signals (i.e., without idiosyncratic shocks and the outcomes of private randomization). For each strategy σ, and each history h_t, let $\sigma^I\left(h_t^I\right)$ denote the distribution over period actions induced by the (random) sequences of past shocks and randomization outcomes. Say that a sequence of strategies σ^n converges to strategy σ, if

$$\lim_{n\to\infty} \sup_{i,h_t^I} \left|\sigma_i^{I,n}\left(h_t^I\right) - \sigma_i^I\left(h_t^I\right)\right| = 0.$$

In other words, strategies converge, if they induce the same distributions over actions.

Say that equilibrium strategy profile σ in game g is *finite past purifiable* if there exists a sequence of games that satisfy Assumption 1 and a sequence of finite past equilibria σ^n of such games, such that games g^n converge to g and σ^n converges to σ. Corollary 1 leads immediately to the following result.

Corollary 2. *Any equilibrium that is finite past purifiable is essentially past-independent.*

Notice that finite memory assumption is significantly weaker than bounded recall.

6.3. **Lexicographic cost of memory.** Following Abreu and Rubinstein (1988) and Rubinstein (1986), we can consider players with lexicographic preferences over cost of memory. Two histories, h_{t-2} and h'_{t-2}, are *t-equivalent* with respect to strategy σ^i if, for all histories $h_t = (h_{t-2}, a_{t-2}, \omega_t, \varepsilon_t)$ and $h'_t = \left(h'_{t-2}, a_{t-2}, \omega_t, \varepsilon_t\right)$, the continuation strategies are the same, $\bar{\sigma}^i(h_t) = \bar{\sigma}^i(h'_t)$. Let $C_t(\sigma^i)$ be the number of classes of t-equivalent histories. Then, $C_t(\sigma^i)$ is a measure of the amount of information that needs to be carried over from period $t-2$

in order to implement t-period continuation of strategy σ^i. Strategy σ^i (*memory*) *dominates strategy* σ'^i, if $C_t(\sigma^i) \leq C_t(\sigma'^i)$ for each t with at least one inequality strict. Strategy σ^i is *undominated*, if it is not memory dominated by any other strategy.

Corollary 3. *Suppose that the monitoring is rich. If $\sigma = (\sigma^1, \sigma^2)$ is an equilibrium profile with finite past and such that each player's strategy is undominated, for each player i, strategy σ^i is essentially past-independent.*

Proof. In case of player $i = 1$, the result follows directly from the definitions and Theorem 2. In case of player $i = 2$, the result follows from the definitions, Theorem 2, the result for player 1, and the fact that if the behavior of player 1 does not depend on information in periods 1 and 3, then player 2's best responses do not depend on information in period 2. □

Under different assumptions, Corollaries 1 and 3 provide foundations for past-independent equilibria. Corollary 3 is very similar to a result in Bhaskar and Vega-Redondo (2002) about the asynchronous games with perfect monitoring. Bhaskar and Vega-Redondo (2002) assumes that the players have bounded recall and lexicographic preference over the size of the recall. The assumption of finite past is significantly weaker than the bounded recall. In fact, because Bhaskar and Vega-Redondo (2002) have perfect monitoring (with finitely many signals corresponding to finitely many actions), finite past does not bite in that framework.

6.4. **Belief-free equilibria in simultaneous move games.** The main result of this paper relies on the fact that in the asynchronous repeated games, each signal provides information simultaneously about the past and future signals of the opponent. This property is not present by the standard model of simultaneous repeated games with private monitoring. However, it is present in a simple modification of the standard model.

Consider a simultaneous repeated game. In each period, players choose actions $a_i \in A_i$. Given the action profile in period t, a pair of signals $(\omega_{i,t}, \omega'_{i,t})$ for each player is drawn from the joint distribution $\rho(a_1, a_2) \in \Delta(\times_{i=1,2}(\Omega_i \times \Omega'_i))$, where Ω'_i is a disjoint copy of Ω_i. In each period, after choosing the actions, player i observes signal $\omega_{i,t}$ and signal $\omega'_{i,t-1}$. Thus, in each period, each player observes a signal about actions and signals in period t, and the actions and signals in period $t-1$.

Although the above setup is not standard, it is not unrealistic. The arrival of information may be distributed across time, some parts of information may arrive immediately, and the rest may be delayed.

We conjecture that under an appropriate richness assumption, the fact that the best responses in finite past equilibria depend only on the most recent information extends to simultaneous games.

References

Abreu, D., and A. Rubinstein (1988): "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," *Econometrica*, 56, 1259–1281.

Bergemann, D., and S. Morris (2007): "Belief Free Incomplete Information Games," .

Bhaskar, V., G. Mailath, and S. Morris (2008): "Purification in the Infinitely Repeated Prisoner's Dilemma," *Review of Economic Dynamics*, 11, 515–528.

——— (2009): "A Foundation for Markov Equilibria in Infinite Horizon Perfect Information Games," .

Bhaskar, V., and F. Vega-Redondo (2002): "Asynchronous Choice and Markov Equilibria," *Journal of Economic Theory*, 103, 334–350.

Ely, J., J. Hörner, and W. Olszewski (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73, 377–415.

Ely, J., and J. Valimaki (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102, 84–105.

Ericson, R., and A. Pakes (1995): "Markov Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53–82.

Horner, J., and S. Lovo (2009): "Belief-Free Equilibria in Games with Incomplete Information," *Econometrica*.

Horner, J., and W. Olszewski (2009): "Robust is the Folk Theorem with Imperfect Public Monitoring?," *Quarterly Journal of Economics*, p. forthcoming.

Mailath, G., and S. Morris (2006): "Coordination Failure in Repeated Games with Almost-Public Monitoring," *Theoretical Economics*, 1, 311–340.

Mailath, G. J., and S. Morris (2002): "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory*, 102, 189–228.

Mailath, G. J., and W. Olszewski (2008): "Folk Theorems with Bounded Recall under (Almost) Perfect Monitoring," PIER Working Paper 08-019.

Peski, M. (2009): "Simple Anti-Folk Theorem for Finite Past Equilibria in Games with Private Monitoring," .

Piccone, M. (2002): "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102.

Rubinstein, A. (1986): "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 39, 83–96.

Yoon, K. (2001): "A Folk Theorem for Asynchronously Repeated Games," *Econometrica*, 69, 191–200.

Appendix A. Existence of finite past equilibria

In this Appendix, we argue that finite past equilibria exist. Player i's strategy σ^i is *stationary*, if there exists a measurable function $\theta^i : \Omega_i \times [0,1] \to \Delta A_i$ such that for each history h_{t-2}, actions a_{t-2}, a_t, signal ω_t, and idiosyncratic shock ε_t:

$$\lambda \left\{\gamma : \sigma^i \left(h_{t-2}, a_{t-2}, \omega_t, \varepsilon_t, \gamma\right) = a\right\} = \theta^i \left(a_t | \omega_t, \varepsilon_t\right).$$

Of course, any stationary equilibrium has finite past.

Theorem 3. *There exists a stationary equilibrium.*

In the interest of space, and because the techniques are not novel, we present a sketch of the main steps rather than a complete argument.

(1) Define the topology on the space of continuation strategies. Let Σ_i^* be the set of continuation strategies of player i. Let $\mu_t^i = \left(\mu^i \times \lambda \times \lambda \times \mu_A^i\right)^{T_i(t)} \times \left(\mu^i \times \lambda \times \lambda\right) \in \Delta H_t$ be the independent product of measures. Assume that the topology on Σ_i^* is generated by the norm $\|.\|_i^*$: for each of two strategies $(a, s), (a', s') \in \Sigma^i$, let

$$\|(a, s) - (a', s')\|_i^* = \mathbf{1}\{a = a'\} + \sum_{t \in T_i} \delta^t \int_{H_t} \sum\nolimits_{a_i} \mathbf{1}\{s(h_t) \neq s'(h_t)\} \, d\mu_t^i(h_t).$$

Then, $\mathbf{\Sigma}_i^*$ is compact.

(2) Using the fact that payoffs are discounted, and the monitoring is absolutely continuous with respect to measures μ_i, one shows that the continuation payoffs $G_i(b|h_t, \sigma)$ are continuous in the continuation strategies. Hence, because Σ_i^* is compact, the set of continuation best responses following history h_t is a compact subset of Σ_i^*.

(3) If the opponent's strategy is stationary, we can show that continuation best response payoffs (i.e., $\max_{b \in \Sigma_i^*} G_i(b|h_t, \sigma)$) depend only on the most recent signals $(\omega_t, \varepsilon_t)$. As a consequence, one shows that for each stationary strategy, there exists a compact set of stationary best responses.

(4) There is an alternative way of defining the stationary strategies. Let Θ_i be the space of measurable functions $\theta^i : \Omega_i \times [0,1] \to \Delta A_i$. Assume that the topology on Θ_i is generated by the norm: for each $\theta, \theta' \in \Theta_i$, let

$$\|\theta - \theta'\|_i^\Theta = \int_{\Omega_i \times [0,1]} \sum\nolimits_{a_i} \left|\theta(a|\omega, \varepsilon) - \theta'(a|\omega, \varepsilon)\right| d\mu^i(\omega) \, d\lambda(\varepsilon).$$

Then, Θ_i is a compact and convex subset of a Polish space. Each element of Θ_i induces a class of equivalent stationary strategies, where two strategies are equivalent if they induce the same distribution over outcomes given any strategy of the opponent. Using a similar argument as in point 4, one shows that the set of stationary best responses is convex.

(5) The existence of stationary equilibria follows from an application of the Kakutani Fixed Point Theorem.

Appendix B. Proof of Theorem 2

This part of the Appendix is devoted to the proof of Theorem 2. We assume throughout that (σ^1, σ^2) is an equilibrium profile of finite past strategies, monitoring ρ is rich, and that for each player i, W_i^ρ is the set from the definition of the rich monitoring.

B.1. **Preliminaries.**

We start with a series of technical results.

B.1.1. *Topological measurable spaces.* Suppose that $(X, \mathcal{X}, \mu)$ is a topological measurable space, i.e., a topological space with Borel σ-field $\mathcal{X}$.

For each measurable $V \subseteq X$, let μ-closure of V, $\operatorname{cl}_\mu V$, as the smallest closed set such that $\mu\left(\left(X \backslash \operatorname{cl}_\mu V\right) \cap V\right) = 0$. In other words, $\operatorname{cl}_\mu V$ is the support of measure μ restricted to set V.

A measurable set is *(μ-)almost open*, if $V \subseteq \operatorname{int} \operatorname{cl}_\mu V$, or, in other words, for each open V', if $V' \cap V \neq \varnothing$, then $\mu\left(V \cap V'\right) > 0$. If V is almost open, then $\operatorname{cl} V = \operatorname{cl}_\mu V$ (note that the left-hand side denotes the topological closure of V, and the right-hand side corresponds to the μ-closure.) A measurable set $V \subseteq X$ is *(μ-)almost nowhere-dense*, if $\operatorname{int} \operatorname{cl}_\mu V = \varnothing$.

A *(finite) partition* of A is a finite collection $\Pi = \{V_1, ..., V_n\}$ of measurable disjoint sets $V_m \subseteq \hat{\Omega}_i$ such that $(\mu^i \times \lambda)\left(\bigcup_m V_m\right) = 1$. A *refinement* of partition Π is a finite partition Π' such that for each $V' \in \Pi'$, there is $V \in \Pi$ so that $V' \subseteq V$. Partition Π is *regular* if for each $V \in \Pi$, either V is almost open or V is almost nowhere-dense. It follows from Lemma 1 that each partition has a regular refinement.

Lemma 1. *Any measurable subset $V \subseteq X$ is a disjoint union of almost open and almost nowhere-dense sets.*

Proof. For each measurable $V \subseteq X$, let $V_0 = V \cap \operatorname{int} \operatorname{cl}_\mu V$ and $V_1 = V \backslash \operatorname{int} \operatorname{cl}_\mu V$. Then, V_0 is almost open and V_1 is almost nowhere-dense. □

Lemma 2. *Suppose that measure μ has full support, i.e., $\operatorname{cl}_\mu X = X$. For any finite sequence $V_1, ..., V_m$ of almost nowhere-dense sets, $\operatorname{cl}_\mu \left(X \backslash \bigcup_m V_m\right) = X$.*

Proof. On the contrary, suppose that there is open $V \subseteq X$ such that $\mu\left(V \backslash \bigcup_m V_m\right) = 0$. Let $V'_m = \operatorname{cl}_\mu V_m$. Then, V'_m is nowhere-dense, $\mu\left(V_m \backslash V'_m\right) = 0$, and there is measurable

$$V_0 = \left(V \backslash \bigcup_m V_m\right) \cup \left(\bigcup_m V_m \backslash V'_m\right)$$

such that $\mu(V_0) = 0$ and $V = V_0 \cup \bigcup_m V'_m$. Because V'_m is nowhere-dense and closed, $X \backslash V'_m$ is open and dense, and $\bigcap_m X \backslash V'_m$ is open and dense, as a finite intersection of open and dense sets. Thus,

$$V_0 \supseteq V \cap \bigcap_m X \backslash V'_m.$$

However, $V \cap \bigcap_m X \backslash V'_m$ is open, non-empty, and has positive measure. But this contradicts the fact that $\mu(V_0) = 0$. □

B.1.2. *Non-atomic probability spaces.* Suppose that $(X, \mathcal{X}, \mu)$ is a probability space. Probability space $(X, \mathcal{X}, \mu)$ is *non-atomic*, if for each $C \in \mathcal{X}$ such that $\mu(C) > 0$, there exists $C' \in \mathcal{X}$ such that $C' \subseteq C$ and $0 < \mu(C') < \mu(C)$.

Lemma 3. *Take any two non-atomic probability spaces $(X, \mathcal{C}, \mu)$ and $(X, \mathcal{D}, \mu)$ such that $\mathcal{D} \subseteq \mathcal{C}$. Then, for each set $C \in \mathcal{C}$ such that $\mu(C) > 0$, there exists $D \in \mathcal{D}$ such that $0 < \mu(D \cap C) < \mu(C)$.*

Proof. Suppose that for all $D \in \mathcal{D}$, either $\mu(D \cap C) = 0$, or $\mu(C \backslash D) = 0$. Let $\mathcal{D}(C) = \{D \in \mathcal{D} : \mu(C \backslash D) = 0\}$. For each countable sequence $D_1, D_2, ... \in \mathcal{D}(C)$, $\mu\left(C \backslash \bigcap_n D_n\right) = 0$, and $\mu\left(\bigcap_n D_n\right) \geq \mu(C)$. Let

$$m^* = \inf_{D_1, D_2, ... \in \mathcal{D}(C)} \mu\left(\bigcap_n D_n\right),$$

where the infimum is taken over all countable sequences of elements of $\mathcal{D}(C)$. Then, $m^* \geq \mu(C)$, and we can find a sequence $D_1^*, D_2^*, ...$ such that for all $D \in \mathcal{D}(C)$, $\mu\left(D \cap \bigcap_n D_n^*\right) = \mu\left(\bigcap_n D_n^*\right) = m^*$.

We show that $\bigcap_n D_n^*$ is an atom. Indeed, suppose that there is $D \notin \mathcal{D}(C)$ such that $0 < \mu\left(D \cap \bigcap_n D_n^*\right) < \mu\left(\bigcap_n D_n^*\right)$. Because $D \notin \mathcal{D}(C)$, it must be that $\mu(D \cap C) = 0$. But then, $D_n^* \backslash D \in \mathcal{D}(C)$ for each n, and

$$\mu\left(\bigcap_n (D_n^* \backslash D)\right) \leq m^* - \mu\left(D \cap \bigcap_n D_n^*\right) < m^*.$$

This leads to a contradiction with the choice of the sequence (D_n^*). □

B.1.3. *Functional spaces.* Suppose that $(X, \mathcal{X}, \mu)$ is a probability space. Let $L^2(X, \mathcal{X}, \mu)$ be the space of all $\mathcal{X}$-measurable and μ-square integrable functions with L^2-norm $\|.\|$. For each $f \in L^2(X, \mathcal{X}, \mu)$, we use $\mu[f]$ to denote the integral of f. Let $L^{**}(X, \mathcal{X}, \mu) \subseteq L^2(X, \mathcal{X}, \mu)$ be the space of μ-a.s bounded $\mathcal{X}$-measurable, and μ-square integrable functions $f : X \to R$ such that $\mu[f] := \int_X f d\mu = 1$, and that $f > 0$. We assume that $L^{**}(X, \mathcal{X}, \mu)$ inherits the normed topology of $L^2(X,\mathcal{X},\mu)$. (Note that space $L^*(X,\mu)$ defined in Section 3 contains functions f such that $f \geq 0$. In particular, $L^{**}(X,\mathcal{X},\mu) \subseteq L^*(X,\mu)$ if the probability space (X,μ) has σ-field $\mathcal{X}$.)

Lemma 4. *Suppose that $(X,\mathcal{X},\mu)$ is a probability space and $U \subseteq L^{**}(X,\mathcal{X},\mu)$ is open. For any $q \in L^2(X,\mathcal{X},\mu)$, if $\sup_{f\in U}\mu[fq] \leq 0$, and there exists $f_0 \in U$ such that $\mu[f_0 q] = 0$, then $q = 0$ μ-almost surely.*

Proof. We first show that for each $f \in L^{**}(X,\mathcal{X},\mu)$, there exists $\alpha^* > 0$ so that for each $\alpha \in (-\alpha^*,\alpha^*)$, there exists $f_\alpha \in L^{**}(X,\mathcal{X},\mu)$ such that $\alpha f_\alpha + (1-\alpha) f_0 \in U$ and $\lim_{\alpha\to 0} f_\alpha = f$. For all $\alpha \geq 0$, take $f_\alpha = f'$. For each $\alpha < 0$, find

$$D_\alpha = \{x \in X : \alpha f(x) + (1-\alpha) f_0(x) \geq 0\}.$$

Let $f_\alpha = (1-\alpha)\frac{1}{\mu[\mathbf{1}_{D_\alpha} f']}\mathbf{1}_{D_\alpha} f + \alpha$. Then, $\lim_{\alpha\to 0} f_\alpha = f$ and, because U is open, for sufficiently small α, $\alpha f_\alpha + (1-\alpha) f_0 \in U$.

Then, for each $\alpha \in (-\alpha^*,\alpha^*)$,

$$0 \geq \mu[(\alpha f_\alpha + (1-\alpha) f_0) q] = \alpha\mu[f_\alpha q].$$

In particular,

$$\mu[f_\alpha q] \leq 0 \text{ for } \alpha > 0, \text{ and } \mu[f_\alpha q] \geq 0 \text{ for } \alpha < 0.$$

Because $\lim_\alpha f_\alpha = f$, and because the expectation operator $\mu[.]$ is continuous, it must be that $\mu[fq] = 0$. Because the latter is true for each $f \in L^{**}(X,\mathcal{X},\mu)$, it must be that $q = 0$ μ-almost surely. □

B.1.4. *Variational technique.* Here, we present a result that provides a foundation for the "variational" technique described in the discussion of the main result in Section 5. Throughout this section, we assume that probability spaces $(X,\mathcal{C},\mu)$ and $(X,\mathcal{D},\mu)$ are non-atomic and such that $\mathcal{D} \subseteq \mathcal{C}$.

To motivate the following result, consider a decision problem, where the decision maker chooses one of $n \geq 2$ actions. The payoffs depend on the realization of the state of the world $x \in X$ and they are described by a vector of functions $q_1,...,q_n \in L(X,\mathcal{C},\mu)$. To avoid trivialities, we assume that $q_m \neq q_{m'}$ for each $m \neq m'$.

The decision maker's beliefs have μ-density $f \in L^{**}(X,\mathcal{C},\mu)$. Her expected payoffs from action m are equal to $\mu[fq_m]$. Let $M(f)$ denote the set of payoff-maximizing actions

$$M(f) = \arg\max_m \mu[fq_m].$$

Set $M(f)$ may contain one or more elements.

We study the behavior of the correspondence M with respect to smooth changes of beliefs. More precisely, we look at the behavior of correspondence $M(pf)$, where $p, f \in L^{**}(X,\mathcal{C},\mu)$

are two densities. One thinks about f as the prior density modified by additional information with density p. The following two results show that set $M(pf)$ contains exactly one element for generic modifications of p and f.

Lemma 5. *Let $f_0 \in L^{**}(X, \mathcal{D}, \mu)$ be an element of the functional space and $P \subseteq L^{**}(X, \mathcal{C}, \mu)$ be a connected open set. Assume that there are $p_1, p_2 \in P$ so that $M(p_1 f_0) \neq M(p_2 f_0)$. Then, for each neighborhood $T \ni f_0$, $T \subseteq L^{**}(X, \mathcal{D}, \mu)$, there is $p \in P$, and $f, f' \in T$ such that $M(pf) \cap M(pf') = \varnothing$.*

Lemma 6. *Let $p_0 \in L^{**}(X, \mathcal{C}, \mu)$ be an element of the functional space and $T \subseteq L^{**}(X, \mathcal{D}, \mu)$ be a connected open set such that there are $f_1, f_2 \in T$ so that $M(p_0 f_1) \neq M(p_0 f_2)$. Then, for each neighborhood $P \ni p_0$, $P \subseteq L^{**}(X, \mathcal{C}, \mu)$, there is $f \in T$, and $p, p' \in P$ such that $M(pf) \cap M(p'f) = \varnothing$.*

The proof of Lemma 5 is presented below. The proof of Lemma 6 follows the same steps as the proof of Lemma 5; the details are omitted.

Lemma 7. *Suppose that $|M(p_0)| \geq 2$ for some $p_0 \in L^{**}(X, \mathcal{C}, \mu)$. Then, for each neighborhood $V_0 \ni p_0$ in space $L^{**}(X, \mathcal{C}, \mu)$, there is $p \in V_0$, neighborhood $V \ni p$ in space $L^{**}(X, \mathcal{C}, \mu)$, and $m \neq m'$ such that $\{m, m'\} \subseteq M(p)$, and for each $p' \in V$, $M(p') \cap \{m, m'\} \neq \varnothing$.*

Proof. W.l.o.g. assume that $\{1, 2\} \subseteq M(p_0)$. Let $U_1 = \{c : q_1(c) > q_2(c)\}$ and $U_2 = \{c : q_1(c) < q_2(c)\}$. Because $q_1 \neq q_2$, U_1 and U_2 are non-empty open half-spaces of $L^{**}(X, \mathcal{C}, \mu)$. Then, either for each p', $M(p') \cap \{1, 2\} \neq \varnothing$, or for each neighborhood $V_0 \ni p_0$, there is $p \in V_0 \cap U_1$ such that $|M(p)| \geq 2$ and $\{2\} \notin M(p)$, or for each neighborhood $V_0 \ni p_0$, there is $p \in V_0 \cap U_2$ such that $|M(p)| \geq 2$ and $\{1\} \notin M(p)$. In the first case, the thesis of the lemma holds. In the second (or the third) case, apply the same argument to functions $q_1, q_3, .., q_n$ (or $q_2, .., q_n$) and neighborhood $V \ni p$ such that $V \subseteq V_0 \cap U_1$ (or $V \subseteq V_0 \cap U_2$). □

Lemma 8. *Suppose that there are $p_0 \in L^{**}(X, \mathcal{C}, \mu)$, $f_0 \in L^{**}(X, \mathcal{D}, \mu)$ and a neighborhood $V_0 \ni p_0$ in space $L^{**}(X, \mathcal{C}, \mu)$, such that $\{1, 2\} \subseteq M(p_0 f_0)$ and for each $p' \in V_0$, $M(p' f_0) \cap \{1, 2\} \neq \varnothing$. Then, for each neighborhood $T \ni f_0$ in space $L^{**}(X, \mathcal{D}, \mu)$, there are $p \in V_0$ and $f, f' \in T$ such that $M(pf) \cap M(pf') = \varnothing$.*

Proof. Let $q = q_1 - q_2$ and define $C^+ = \{x : q(x) > 0\}$, and $C^- = \{x : q(x) < 0\}$. Because $q \neq 0$, it must be that either $\mu(C^+) > 0$ or $\mu(C^-) > 0$. Because $\{1, 2\} \subseteq M(p_0 f_0)$, it must be that $\mu(C^+) > 0$ and $\mu(C^-) > 0$.

By Lemma 3, there exists $D \in \mathcal{D}$ such that $0 < \mu(D \cap C^+) < \mu(C^+)$.

Because $p_0, f_0 > 0$, we can find $a, b, c \geq 0$ such that if

$$p = \left(1 - a\mathbf{1}_{C^+ \setminus D} - b\mathbf{1}_{C^- \cap D} - c\mathbf{1}_{C^- \setminus D}\right) p_0,$$

then $p \in V_0$, $\mu(pq) = 0$, and $\mu[pf_0 q \mathbf{1}_D] > 0$.

For each $\alpha \in (0,1)$, let

$$f_\alpha = \frac{f_0 + \alpha f_0 \mathbf{1}_D}{\mu[f_0 + \alpha f_0 \mathbf{1}_D]} \text{ and } f'_\alpha = \frac{f_0 + \alpha f_0 \mathbf{1}_{C \setminus D}}{\mu\left[f_0 + \alpha f_0 \mathbf{1}_{C \setminus D}\right]}.$$

Then, for sufficiently small α, $f_\alpha, f'_\alpha \in T$, and

$$\mu[pf_\alpha q_1] > \mu[pf_\alpha q_2], \text{ and } \mu[pf_\alpha q_2] > \mu[pf_\alpha q_1].$$

Suppose that $m \in M(pf_\alpha) \cap M(pf'_\alpha)$. Then, $m \neq 1, 2$, and, by linearity,

$$\begin{aligned}\mu[pf_0 q_m] &= \mu[f_0 + \alpha f_0 \mathbf{1}_D]\,\mu[pf_\alpha q_m] + \mu[f_0 + \alpha f_0 \mathbf{1}_D]\,\mu[pf'_\alpha q_m] \\ &> \mu[f_0 + \alpha f_0 \mathbf{1}_D]\,\mu[pf_\alpha q_1] + \mu[f_0 + \alpha f_0 \mathbf{1}_D]\,\mu[pf'_\alpha q_1] = \mu[pf_0 q_1],\end{aligned}$$

which contradicts the fact that $1 \in M(p, f_0)$. □

The hypothesis of Lemma 5 implies that there exists $p_0 \in P$ so that $|M(p_0 f_0)| \geq 2$. The Lemma follows from Lemmas 7 and 8.

B.2. **Histories, beliefs, and best responses.** Here, we define notation used later in the proof, define topologies on the spaces of signals and histories, and show that for the purpose of finding best responses, private histories can be replaced by an auxiliary information contained in the signal and described in Section 3.

From now on, we fix strategy profile σ, player i, period $t \in T_i$, $g_t^* \in G_i^*$, and $\varepsilon_t^*, \gamma_t^* \in [0,1]$.

B.2.1. *Topologies on signal spaces.* Define

$$\begin{aligned}W_i^\alpha &:= \{(p, f(a_i)) : (p, g, f(.)) \in W_i^\rho \text{ for some } g \in G_i^* \text{ and } a \in A_i\},\\ \Omega_i^\alpha &= \Omega_i \times [0,1]^2 \times A_i,\\ \mu_i^\alpha &= \mu^i \times \lambda \times \lambda \times \mu_i^A \in \Delta\Omega_i^\alpha.\end{aligned}$$

Then, W_i^α is an open and connected subset of $P^* \times F_i^*$. Define measurable mapping $\upsilon_i^\alpha : \Omega_i^\alpha \to W_i^\alpha$,

$$\upsilon_i^\alpha(\omega_i, \varepsilon_i, \gamma_i, a_i) = \left(\upsilon_p^\rho(\omega_i), \upsilon_f^\rho(\omega_i, a_i)\right).$$

Define

$$\begin{aligned}W_i^\beta &= \{(p, f(.)) : (p, g_t^*, f(.)) \in W_i^\rho\},\\ \Omega_i^\beta &= \left\{(\omega_i, \varepsilon_t^*, \gamma_t^*) : \upsilon_g^\rho(\omega_i) = g_t^*\right\} \subseteq \Omega \times [0,1]^2, \text{ and}\\ \mu_i^\beta &= \mu_i\left(.|\upsilon_g^\rho(\omega_i) = g_t^*\right) \times \delta_{\varepsilon_t^*} \times \delta_{\gamma_t^*} \in \Delta\Omega_i^\beta.\end{aligned}$$

Then, W_i^β is an open subset of $P_i^* \times (F_i^*)^{A_i}$. Define a measurable mapping $\upsilon_i^\beta : \Omega_i^\beta \to P_i^* \times F_i^*$,

$$\upsilon_i^\beta(\omega_i, \varepsilon_t^*, \gamma_t^*) = \left(\upsilon_p^\rho(\omega_i), \upsilon_f^\rho(a_i|\omega_i)\right).$$

For each $x = \alpha, \beta$, define a collection of subsets of Ω^x :

$$\mathcal{V}^x = \left\{ (\upsilon_i^x)^{-1} (U) : U \subseteq W_i^x,\ U \text{ is open} \right\}.$$

All elements of $\mathcal{V}^x$ are measurable, and $\mathcal{V}^x$ contains unions and finite intersections of elements of $\mathcal{V}^x$. Assume that the topology on Ω_i^x is defined by $\mathcal{V}^x$ as a collection of open sets. One checks that $\mathrm{cl}_{\mu_i^x} \Omega_i^x = \Omega_i^x$.

B.2.2. *Space of histories.* Define a subset of histories

$$\begin{aligned} H_t^* &= (\Omega_i^a)^{T_i(t)} \times \Omega_i^\beta = \left\{ h_t \in H_t : \upsilon_g^\rho (\omega_t) = g_t^*, \varepsilon_t = \varepsilon_t^*, \gamma_t = \gamma_t^* \right\} \subseteq H_t, \\ \mu_t^* &= (\mu^\alpha)^{T_i(t)} \times \mu^\beta \in \Delta H_t. \end{aligned}$$

The probability space (H_t^*, μ_t^*) is equipped with the product topology.

B.2.3. *Auxiliary histories.* Additionally, we define a space of *auxiliary histories*, in which signals are replaced by their density representations. Let

$$\begin{aligned} \hat{H}_t &= (P_i^* \times F_i^*)^{T_i(t)} \times P_i^* \times (F_i^*)^{A_i}, \\ \hat{H}_t^* &= (W_i^\alpha)^{T_i(t)} \times W_i^\beta \subseteq \hat{H}_t. \end{aligned}$$

Then, $\hat{H}_t^*$ is an open subset of $\hat{H}_t$. A typical auxiliary history is denoted as $\hat{h}_t = (..., p_{t-2}, f_{t-2}, p_t, f_t(.)) \in \hat{H}_t$.

The standard histories map into the auxiliary histories. Define mapping $\upsilon^H : H_t^* \to \hat{H}_t^*$. For each $s \in T_{-i}(t)$, let

$$\begin{aligned} &\upsilon^H (..., \omega_{t-2}, \varepsilon_{t-2}, \gamma_{t-2}, a_{t-2}, \omega_t, \varepsilon_t^*, \gamma_t^*) \\ &= \left(..., \upsilon_i^\alpha (\omega_{t-2}, \varepsilon_{t-2}, \gamma_{t-2}, a_{t-2}), \upsilon_i^\beta (\omega_t, \varepsilon_t^*, \gamma_t^*) \right). \end{aligned}$$

Mapping υ^H is continuous and $\upsilon^H(H_t^*)$ is dense in $\hat{H}_t^*$.

B.2.4. *Beliefs.* Next, we show that beliefs over histories of the opponent after history h_t depend in some sense only on their auxiliary representation $\upsilon^H(h_t)$.

For each strategy σ^{-i} of player $-i$, let $H_{t-1}^\sigma \subseteq H_{t-1} \times A_{-i}$ denote the space of histories and actions that are *consistent with strategy* σ^{-i}. Precisely, let $(h_{t-1}, a_{t-1}) \in H_{t-1}^\sigma$ if and only if $\sigma^{-i}(h_s) = a_s$ for each $s \in T_{-i}(t)$ and $h_s \leq h_{t-1}$.

Let

$$\mu_{t-1}^\sigma := \left(\mu^{-i} \times \lambda \times \lambda \times \mu_A^{-i} \right)^{T_{-i}(t)} \left(.|H_{t-1}^\sigma \right) \in \Delta H_{t-1}^\sigma$$

be a distribution on H_{t-1}^σ.

Take any t-period history $h_t \in H_t^*$. Because of the full support assumption, the conditional distribution $\pi_{h_t}^{\sigma,\mu_1}$ is absolutely continuous with respect to μ_{t-1}^{σ} and its Radon-Nikodym derivative is proportional to

$$\frac{d\pi_{h_t}^{\sigma,\mu_1}}{d\mu_t^{\sigma}}(h_{t-1}) \propto \prod_{s\in T_i(t)} \frac{d\rho_{-i}}{d\mu_{-i}}(\omega_{s+1}|\omega_s, a_s) \prod_{s\in T_i(t)\backslash 1} \frac{d\rho_i}{d\mu_i}(\omega_s|\omega_{s-1}, a_{s-1}) \mathbf{1}\{a_{s+1} = \sigma_{-i}(h_{s+1})\}$$

$$\propto \mathbf{1}\left\{h_{t-1} \in H_{t-1}^{\sigma}\right\} \prod_{s\in T_i(t)} \upsilon_f^{\rho}(\omega_{s+1}, a_s|\omega_s) \prod_{s\in T_i(t)\backslash 1} \upsilon_p^{\rho}(\omega_{s-1}, a_{s-1}|\omega_s). \quad \text{(B.1)}$$

B.2.5. *Best responses.* Finally, we show that the knowledge of auxiliary histories is sufficient to find the best responses.

For any continuation strategy $b = (a, s) \in \Sigma_i^*$, define payoff function $q_b : H_{t-1}^{\sigma} \times A_{-i} \times (F_i^*)^{A_i} \to R$ as

$$q_b(h_{t-1}, a_{t-1}, f(.)) = (1-\delta) g_t^*(\varepsilon_t^*, a) + \delta G_i\left(s, \sigma^{-i}(h_{t-1}, a_{t-1}); f(.|a)\mu_{-i}\right).$$

Thus, $q_b(h_{t-1}, a_{t-1}, f(.))$ is the expected payoff from period t continuation strategy $b = (a, s)$ when the opponent's private history is h_{t-1} followed by action a_{t-1} and $f(.|a)\mu_{-i}$ is the distribution of the opponent's $t+1$-period signal.

Let $Q_t^0 = \{q_b : b \in \Sigma_i^*\}$ be the space of payoff functions induced by all continuation strategies.

For each $q \in Q_t^0$, each auxiliary history $\hat{h}_t \in \hat{H}_t$, define

$$E_{\hat{h}_t} q = \int_{H_{t-1}^{\sigma}} q(h_{t-1}, a_{t-1}, f(.)) \left(\prod_{s\in T_i(t)} f_s(\omega_{s+1}) \prod_{s\in T_i(t)\backslash\{1\}} p_s(\omega_{s-1}, a_{s-1}) \right) d\mu_{t-1}^{\sigma}(h_{t-1}, a_{t-1}).$$

(In the above integration, we use the representation $h_{t-1} = (...\omega_{t-3}, \varepsilon_{t-3}, \gamma_{t-3}, a_{t-3}, \omega_{t-1}, \varepsilon_{t-1}, \gamma_{t-1})$.)

For each $q, q' \in Q_t^0$ write $q' \equiv q$ iff $E_{\hat{h}_t} q = E_{\hat{h}_t} q'$ for all auxiliary histories $\hat{h}_t \in H_t$. Let $Q_t := Q_t^0 / \equiv$ be the quotient of the space of payoff functions with respect to relation "$\equiv$". In other words, Q_t consists of equivalence classes of payoff functions with equal expectations with respect to auxiliary histories.

For each $q \in Q_t$ and each $\hat{h}_t \in H_t$, define

$$B(q) = \left\{\hat{h}_t \in \hat{H}_t : E_{\hat{h}_t} q > E_{\hat{h}_t} q' \text{ for each } q' \in Q_t \backslash \{q\}\right\},$$

$$\bar{B}(q) = \left\{\hat{h}_t \in H_t^* : E_{\hat{h}_t} q \geq E_{\hat{h}_t} q' \text{ for each } q' \in Q_t \backslash \{q\}\right\}.$$

Lemma 9. *For each history $h_t \in H_t^*$, each continuation strategy $b \in \Sigma_i^*$, b is the best response after history h_t iff $v^H(h_t) \in \bar{B}(q_b)$.*

Proof. The Lemma is an immediate consequence of (B.1). □

Lemma 10. *For each $q \in Q_t$, $\operatorname{cl} B(q) \subseteq \bar{B}(q)$.*

Proof. The Lemma follows from the continuity of expectation $E_{\hat{h}_t} q$ with respect to $\hat{h}_t$. □

B.3. **Finite past.** We state and prove a product property for finite past strategies. For each sequence of finite partitions $\Pi_s, s \in T_i(t)$ of signal space Ω_i^α and Π_t^* of signal space Ω_i^β, define a product partition $\Pi = \bigwedge_{s \in T_i(t)} \Pi_s \wedge \Pi_t^*$ on the space of histories H_t: $U \in \Pi$ if and only if $U = \times_{s \in T_i(t)} V_s \times V_t$ for some $V_s \in \Pi_s$ and $V_t \in \Pi_t$. If partitions Π_s and Π_t^* are regular, then Π is regular and U is almost open if and only if sets V_s are almost open.

Player i's t -period continuation of strategy σ is *measurable with respect to partition* Π of H_t if for each subset $U \in \Pi$, $\bar{\sigma}(h_t) = \bar{\sigma}(h_t')$ for μ_t-almost all histories $h_t, h_t' \in U$. If partitions Π_s' refine partitions Π_s, and σ is measurable with respect to the product partition $\bigwedge_{s \in T_i(t)} \Pi_s \wedge \Pi_t^*$, then it is measurable with respect to the refinement $\bigwedge_{s \in T_i(t)} \Pi_s \wedge \Pi_t^*$.

Lemma 11. *If player i's strategy $\sigma \in \Sigma_i$ has finite memory, then there exist regular partitions Π_t, of Ω_i^α and Π_t^* of Ω_i^β such that, for each $t \in T_i$, the t-period continuation of strategy σ is measurable with respect to $\bigwedge_{s \in T_i(t)} \Pi_s \wedge \Pi_t^*$.*

Proof. The existence of partitions such that the t-period continuation of strategy σ is measurable with respect to their product is immediate. The existence of a regular refinement follows from Lemma 1. □

B.4. **Variational technique.** We describe the intermediary results that together form the variational technique described in Section 5.

We start with some notation. For each auxiliary history $\hat{h}_t \in \hat{H}_t$, for each $s \in T_i(t)$, for each $p \in P_i^*$, define auxiliary history $\hat{h}_t^{s,p}$ as equal to history $\hat{h}_t$ with p_s replaced by p :

$$\hat{h}_t^{s,p} = (...., f_{s-2}, p, f_s, ..) \in \hat{H}_t.$$

In a similar way, for each auxiliary history $\hat{h}_t \in \hat{H}_t$,

- for each $s \in T_i(t)$, for each $f \in F_i^*$, define $\hat{h}_t^{s,f}$ as history $\hat{h}_t$ with f_s replaced by f;
- for each $s \in T_i(t)$, for each $(p, f) \in P_i^* \times F_i^*$, define $\hat{h}_t^{s,(p,f)}$ as history $\hat{h}_t$ with p_s and f_s replaced by, respectively, p and f;
- for each $(p, f(.)) \in P_i^* \times (F_i^*)^{A_i}$, define $\hat{h}_t^{(p,f)}$ as history $\hat{h}_t$ with p_t and $f_t(.)$ replaced by, respectively, p and $f(.)$;

Lemma 12. *Take period $s \in T_i(t)$, $q, q' \in Q_t$, and $\hat{h}_t \in \hat{H}_t$ such that $E_{\hat{h}_t} q = E_{\hat{h}_t} q'$. Then,*

1. *for each $s \in T_i(t)$, if there is an open subset $W \subseteq P_i^* \times F_i^*$ such that $(p_t, f_t) \in W$, and $E_{\hat{h}_t^{s,w}} q \geq E_{\hat{h}_t^{s,w}} q'$ for each $w \in W$, then $E_{\hat{h}_t^{s,w}} q = E_{\hat{h}_t^{s,w}} q'$ for each $w \in P_i^* \times F_i^*$;*
2. *if there exists an open $W \subseteq P_i^* \times (F^*)^{A_i}$ such that $(p_t, f_t) \in W$, and $E_{\hat{h}_t^w} q \geq E_{\hat{h}_t^w} q'$ for each $w \in W$, then $E_{\hat{h}_t^w} q = E_{\hat{h}_t^w} q'$ for each $w \in P_i^* \times (F_i^*)^{A_i}$;*

Proof. We prove part (1) (the proof of part (2) is analogous and it is omitted). Part (1) follows from the two following claims. For each $s \in T_i(t)$,

- if there is an open subset $P \subseteq P_i^*$ such that $p_t \in P$, and $E_{\hat{h}_t^{s,p}} q \geq E_{h_t^{s,p}} q'$ for each $p \in P$, then $E_{\hat{h}_t^{s,p}} q = E_{\hat{h}_t^{s,p}} q'$ for each $p \in P_i^*$, and
- if there is an open subset $F \subseteq F_i^*$ such that $f_t \in F$, and $E_{\hat{h}_t^{s,f}} q \geq E_{\hat{h}_t^{s,f}} q'$ for each $f \in F$, then $E_{\hat{h}_t^{s,f}} q = E_{\hat{h}_t^{s,f}} q'$ for each $f \in F_i^*$.

Consider the first claim. Define $q^* \in L^\infty\left(\Omega_{-i} \times A_{-i}, \mu_{-i} \times \mu_{-i}^A\right)$:

$$\begin{aligned} & q^*(\omega_{s-1}, a_{s-1}) \\ & = \int_{H_{t-1}^\sigma} (q - q')(h_{t-1}, a_{t-1}, f) \cdot \\ & \left(\prod_{s' \in T_i(t)} f_{s'}(\omega_{s'+1}) \prod_{s' \in T_i(t+1) \setminus \{s,1\}} p_{s'}(\omega_{s'-1}, a_{s'-1}) \right) d\mu_{t-1}^\sigma(h_{t-1}, a_{t-1} | \omega_{s-1}, a_{s-1}). \end{aligned}$$

Then, for each $p \in P$,

$$E_{h_t^{s,p}} q = \mu_{t-1}^\sigma [p q^*].$$

Because $p_s \in P$, by Lemma 4, $q^* = 0$ $\mu_{-i} \times \mu_{-i}^A$-almost surely, and $E_{\hat{h}_t^{s,p}} q = 0$ for all $p \in P_i^*$. The proof of the second claim is analogous. □

Lemma 13. *For each $q \in Q_t$,* $\operatorname{int} \bar{B}(q) \subseteq B(q)$.

Proof. Fix $q \in Q_t$ such that $\operatorname{int} \bar{B}(q)$ is not empty. Take any $q' \in Q_t \setminus \{q\}$. Then, $E_{\hat{h}_t} q \geq E_{\hat{h}_t} q'$ for each $\hat{h}_t \in \operatorname{int} \bar{B}(q)$ and there exists $\hat{h}_t^0 \in U$ such that $E_{\hat{h}_t^0} q = E_{\hat{h}_t^0} q'$, then $q = q'$. Because $\operatorname{int} \bar{B}(q) \subseteq \hat{H}_t^*$ is open, there are open sets $W_s \subseteq P_i^* \times F_i^*$, $s \in T_i(t)$, and $W_t \subseteq P_i^* \times (F_i^*)^{A_i}$ such that $\times_{s \in T_i(t+1)} W_s \subseteq U$. The result follows from a series of applications of Lemma 12. □

In the next Lemma, we fix period $s \in T_i(t) \setminus \{2\}$ and assume that there exists a collection of open sets $V_{s'} \subseteq W_i^\alpha$ for $s' \neq s, V_t \subseteq W_i^\beta$ and $V_s^m \subseteq W_i^\alpha$ for $m = 1, ..., M < \infty$ such that sets V_s^m and $V_s^{m'}$ are disjoint for $m \neq m'$ and $\bigcup_m V_s^m$ is a dense subset of W_i^α. Define

$$\begin{aligned} U &= V_1 \times ... \times V_{s-1} \times W_i^\alpha \times V_{s+1} \times ... \times V_t, \text{ and for each } m, \\ U^m &= V_1 \times ... \times V_{s-1} \times V_s^m \times V_{s+1} \times ... \times V_t. \end{aligned}$$

Then, U and U^m are open sets of auxiliary histories, and $\bigcup_m U^m$ is a dense subset of U.

Lemma 14. *Suppose that for each m, there exists $q_m \in Q_t$ such that $U^m \subseteq \bar{B}(q_m)$. Then, $q_m = q_{m'} =: q^*$ for all m, m' and $U \subseteq B(q)$.*

Proof. We can assume w.l.o.g. that $q_m \neq q_{m'}$ for all $m \neq m'$. If $M = 1$, then V_s^1 is a dense subset of W_i^α and U^1 is a dense subset of U. Because $U^1 \subseteq B(q_1)$, Lemma 10 implies that $U \subseteq \bar{B}(q)$. Because U is open, Lemma 13 implies that $U \subseteq B(q)$. Thus, if $M = 1$, the Lemma holds.

We show that $M > 1$ leads to a contradiction. Let $V_s = \bigcup_m V_s^m$. For each $\hat{h}_t \in U$ and each m, define

$$V^m\left(\hat{h}_t\right) = \left\{(p, f) \in W_i^\alpha : \hat{h}_t^{s,(p,f)} \in B(q_m)\right\}.$$

Then, $V^m\left(\hat{h}_t\right)$ is open. Moreover, for each $m' \neq m$,

(1) $V^m\left(\hat{h}_t\right) \subseteq \operatorname{cl} V_s^m \backslash \operatorname{cl} V_s^{m'}$. Indeed, take any $(p, f) \in V^m\left(\hat{h}_t\right)$. Then, $\hat{h}_t^{s,(p,f)} \in B(q_m)$, which implies that $\hat{h}_t^{s,(p,f)} \notin \bar{B}(q_{m'})$ and, by Lemma 10, $(p, f) \notin \operatorname{cl} V_s^{m'}$. Because V_s is dense in W_i^α, it must be that $(p, f) \in \operatorname{cl} V_s^m$;

(2) for each $\hat{h}'_t \in U$, $V^m\left(\hat{h}_t\right) \cap V^{m'}\left(\hat{h}'_t\right) = \varnothing$. Indeed, by continuity, if $(p, f) \in V^m\left(\hat{h}_t\right)$, then there exists a neighborhood $V \ni (p, f)$ such that $V \cap V_s^{m'} = \varnothing$. Because V_s is dense, each open subset $V' \subseteq V$ has a non-empty intersection with V_m. If $V^m\left(\hat{h}_t\right) \cap V^{m'}\left(\hat{h}'_t\right) \neq \varnothing$, then there exists open $V \subseteq V^m\left(\hat{h}_t\right) \cap V^{m'}\left(\hat{h}'_t\right)$ that simultaneously has empty and non-empty intersections with V^m and $V^{m'}$. The contradiction shows the claim.

For each m, define

$$\bar{V}^m = \bigcup_{\hat{h}_t \in U} V_m\left(\hat{h}_t\right).$$

Then, sets $\bar{V}^m \supseteq V^m$ are open (as unions of open sets), disjoint, $\bar{V}^m \subseteq \operatorname{cl} V^m \backslash \left(\bigcup_{m' \neq m} \operatorname{cl} V^{m'}\right)$, and for each $\hat{h}_t \in U$ such that $(p_s, f_s) \in \bar{V}^m$, $\hat{h}_m \in \operatorname{cl} B(q)$.

For each $(p, f) \in W_i^\alpha$, define $P(f; p) \subseteq P_i^*$ as the connected component of set $\{p' : (p', f) \in W_i^\alpha\}$ such that $p \in P(f; p)$. Similarly, define $F(p; f)$ as the connected component of set $\{f' : (p, f') \in W_i^\alpha\}$ such that $f \in F(p; f)$. Then, $P(f; p)$ and $F(p; f)$ are open.

We make two claims:

- for each $\hat{h}_t \in U$ such that $(p_s, f_s) \in P^* \times F^* \cap \bar{V}^m$, it must be that $P(f_s; p_s) \subseteq \operatorname{cl} V^m$;
- for each $\hat{h}_t \in U$ such that $(p_s, f_s) \in P^* \times F^* \cap \bar{V}^m$, $F(p_s; f_s) \subseteq \operatorname{cl} V^m$.

We prove the first claim. If $s = 1$, the claim is immediate. Assume that $s \neq 1$. For each $n \leq M$, define the measurable and bounded function $\phi_n : \Omega_{-i} \times A_{-i} \to R$,

$$\phi_n(\omega_{s-1}, a_{s-1}) = \int_{H_{t-1}^\sigma} q_n(h_{t-1}, a_{t-1}, f_t(.)) \cdot \prod_{s' \in T_i(t) \backslash \{s-2\}} f_{s'}(\omega_{s'+1}) \prod_{s' \in T_i(t+1) \backslash \{1, s\}} p_{s'}(\omega_{s'-1}, a_{s'-1})\, d\mu_{t-1}^\sigma(h_{t-1}, a_{t-1} | \omega_{s-1}, a_{s-1}).$$

Then, for each $p \in P_i^*$ and $f \in F_i^*$,

$$E_{\left(h_t^{s-2,f}\right)^{s,p}} q_n = \int\limits_{\Omega_{-i} \times A_{-i}} \phi_n(\omega_{s-1}, a_{s-1}) f(\omega_{s-1}) p(\omega_{s-1}, a_{s-1}) d\mu_{t-1}^{\sigma}(\omega_{s-1}, a_{s-1})$$
$$= \mu_{t-1}^{\sigma}[fp\phi_n].$$

Suppose that $P(f_s; p_s) \cap \operatorname{cl} V_s^{m'} \neq \varnothing$ for some $m' \neq m$. Because $P(f_s; p_s)$ and $V_s^{m'}$ are open, there is $p' \in P(f_s; p_s) \cap V_s^{m'}$ such that

$$\mu_{t-1}^{\sigma}[f_s p_s \phi_m] > \mu_{t-1}^{\sigma}[f_s p_s \phi_{m''}] \text{ for each } m'' \neq m, \text{ and}$$
$$\mu_{t-1}^{\sigma}[f_s p_s \phi_m] \leq \mu_{t-1}^{\sigma}[f_s p_s \phi_{m'}].$$

By Lemma 5, there is $p^* \in P(f_s, p_s)$, $f^1, f^2 \in F(p_{s-2}; f_{s-2})$, and m^1, m^2, $m^1 \neq m^2$ such that

$$\left(\hat{h}_t^{s-2,f^k}\right)^{s,p^*} \in U \cap B(q_{m^k}) \text{ for } k = 1, 2.$$

Then, $(p^*, f_s) \in \bar{V}^{m^k}$ for $k = 1, 2$, which contradicts the fact that $\bar{V}^{m^1}$ and $\bar{V}^{m^2}$ are disjoint.

A similar argument together with Lemma 6 establishes the second claim.

Because set W_i^{α} is connected, the two claims together show that $W_i^{\alpha} \subseteq \operatorname{cl} V_q$. This ends the proof of the Lemma. □

B.5. **Proof of Theorem 2.** In the course of the proof, we assume that strategy profile σ, player i, payoffs $g_t^* \in G_i^*$, and $\varepsilon_t^*, \gamma_t^* \in [0,1]$ are fixed, and we use the notation and results developed above.

Define an equivalence relation on the set of continuation strategies Σ_i^* : say that continuation strategies b and b' are equivalent, $b \simeq b'$, if for each history $h_t \in H_t^*$, b is the best response after h_t iff b' is the best response after h_t'. By Lemma 9, continuation strategies b and b' are equivalent if and only if the associated payoff functions are equivalent, $q_b \simeq q_{b'}$. Let $B_t = \Sigma_i^* / \simeq$ be the space of classes of equivalence. Then, B_t has one-to-one correspondence with set Q_t.

For each period $s \in T_i(t) \setminus \{2\}$, each history $h_t \in H_t^*$, define $H_s^*(h_t^*)$ to be the set of histories $h_t' \in H_t^*$ that (possibly) differ from h_t only with respect to information in period s : for each $s' \neq s, t$, $(\omega_{s'}, \varepsilon_{s'}, a_{s'}) = (\omega_{s'}', \varepsilon_{s'}', a_{s'}')$ and $(\omega_t, \varepsilon_t) = (\omega_t', \varepsilon_t')$. For each set $U \subseteq H_t^*$, define $H_s^*(U) = \bigcup_{h_t} H_s^*(h_t)$.

Let $\Pi = \bigwedge_{s \in T_i(t)} \Pi_s \wedge \Pi_t^*$ be a regular partition from Lemma 11. For each $U \in \Pi$, let $b(U) \in B_t$ be the equivalence class of the continuation strategy that is played after histories in U. Then, $b(U)$ is the best response after μ_t-almost all histories $h_t \in U$.

We show that if $U \in \Pi$ is almost open, then $b(U)$ is the unique best response after each history $h_t \in H_s^*(U)$. Let $U = \times_{s' \in T_i(t)} V_{s'} \times V_t$. Let $V_s^1, \ldots, V_s^M$ be an enumeration of all almost open elements of Π_s and w.l.o.g. assume that $V_s = V_s^1$. Let

$$\begin{aligned}
\hat{V}_{s'} &= \operatorname{int}\operatorname{cl}\upsilon_i^{\alpha}\left(V_{s'}\right) \text{ for } s' \in T_i\left(t\right)\setminus\left\{s\right\},\\
\hat{V}_s^m &= \operatorname{int}\operatorname{cl}\upsilon_i^{\alpha}\left(V_s^m\right) \text{ for each } m,\\
\hat{V}_t &= \operatorname{int}\operatorname{cl}\upsilon_i^{\beta}\left(V_t\right),\\
U^m &= \left(\times_{s'\in T_i(t)\setminus\{s\}}V_{s'}\right)\times V_s^m \times V_t \text{ for each } m,\\
\hat{U}^m &= \left(\times_{s'\in T_i(t)\setminus\{s\}}\hat{V}_{s'}\right)\times \operatorname{cl}\hat{V}_s^m \times \hat{V}_t \text{ for each } m,\\
\hat{U}^* &= \times_{s'\in T_i(t)\setminus\{s\}}\hat{V}_{s'}\times W_i^{\alpha}\times\hat{V}_t.
\end{aligned}$$

By Lemmas 9 and 10, $\upsilon^H\left(U_m\right)\subseteq\bar{B}\left(q_{b(U^m)}\right)$ and $\hat{U}^m\subseteq\operatorname{cl}\upsilon^H\left(U_m\right)$, it must be that $\hat{U}^m\subseteq\bar{B}\left(q_{b(U^m)}\right)$. Because $\hat{U}^m\subseteq\hat{H}_t^*$ is open, Lemma 13 implies that $\hat{U}^m\subseteq B\left(q_{b(U^m)}\right)$. By Lemma 2, and by the choice of topology on Ω_i^{α}, $\bigcup_m\hat{V}_s^m$ is a dense subset of W_i^{α}. By Lemma 14, $q_{b(U^m)} = q_{b\left(U^{m'}\right)} = q$ for all m and m'. Because $U = U^m$ for some m, $\bigcup_m\hat{U}^m\subseteq B\left(q_{b(U)}\right)$. Because $\bigcup_m\hat{U}^m$ is dense in $\hat{U}^*$, by Lemmas 10 and 13, $\hat{U}^*\subseteq B\left(q_{b(U)}\right)$. Because $\upsilon^H\left(H_s^*\left(U\right)\right)\subseteq\hat{U}^*$, by Lemma 9, $b\left(U\right)$ is the unique best response after each history $h_t\in H_s^*\left(U\right)$.

In order to finish the proof of the Theorem, it is enough to show that for each history h_t, there exists (an equivalence class of) strategy b such that b is the unique (up to equivalence) best response after each history $h_t'\in H_s^*\left(h_t\right)$. Take any history $h_t\in H_t^*$ and suppose that (the equivalence class of) continuation strategy $b\in B_t$ is the best response after h_t. By Lemma 2, either there is almost open $U\in\Pi$ such that $h_t\in U$, or $h_t\in\operatorname{cl}_{\mu_t}U$ for some almost open $U\in\Pi$. In the former case, the above argument shows that $b=b\left(U\right)$ is the best response after each history $h_t'\in H_s^*\left(h_t\right)$. In the later case, $h_t\in\operatorname{cl}_{\mu_t}U\backslash U$ for some almost open $U\in\Pi$. By the above argument, $b\left(U\right)$ is the unique best response after history $h_t'\in H_s^*\left(U\right)$. By continuity, $b\left(U\right)$ is the (not necessarily unique) best response after history h_t, and after each history $h_t'\in H_s^*\left(h_t\right)\subseteq\operatorname{cl}_{\mu_t}H_s^*\left(U\right)$. By part (3) of Lemma 12, $E_{\upsilon^H(h_t)}q_b = E_{\upsilon^H(h_t')}q_b$ for each $h_t'\in H_s^*\left(h_t\right)$. By Lemma 9, b is the best response after each $h_t'\in H_s^*\left(h_t\right)$.

Appendix C. Proof of Theorem 1

We start with some notation. For each player i and $g\in G_i^*$, let $\Omega_i\left(g\right)=\left\{\omega_i : g\left(\omega_i,.\right)=g\right\}$. Without loss of generality, we assume that for each $g\in G_i^*$, $\mu_i\left(\Omega_i\left(g\right)\right)>0$.

Let $\Gamma_0\subseteq\Gamma$ be the set of monitorings with full support.

For each $\rho, \rho' \in \Gamma$, let $\|\rho - \rho'\|$ be the norm that induces topology on Γ:

$$\begin{aligned}\|\rho - \rho'\| &= \sum\nolimits_i \|\rho_i - \rho_i'\|_2 \\ &= \sum\nolimits_i \sqrt{\int \left|\rho\left(\omega_i|\omega_{-i}, a_{-i}\right) - \rho_i'\left(\omega_i|\omega_{-i}, a_{-i}\right)\right|^2 d\left(\mu_i \times \mu_{-i} \times \mu_{-i}^A\right)}.\end{aligned}$$

Topology on Ω_i^* is metrizable with the following metric: for any $(p, g, f), (p', g', f') \in \Omega_i^*$, let

$$d^*\left((p, g, f), (p', g', f')\right) = \mathbf{1}\{g = g'\} + \|p - p'\|_{\Omega_{-i}\times A_{-i}} + \frac{1}{|A_i|}\sum\nolimits_{a_i} \|f(a_i) - f'(a_i)\|_{\Omega_{-i}},$$

where $\|.\|_{\Omega_{-i}\times A_{-i}}$ is the L^2-norm on the probability space $\left(\Omega_{-i} \times A_{-i}, \mu_{-i} \times \mu_{-i}^A\right)$, and $\|.\|_{\Omega_{-i}}$ is the L^2-norm on the probability space (Ω_{-i}, μ_{-i}). For each $\gamma > 0$, let $B(p, g, f, \gamma) \subseteq \Omega_i^*$ be the open ball with center at (p, g, f) and radius γ.

Lemma 15. Γ_0 *is a non-meagre subset of* Γ.

Proof. We show that for each player i, and each $\gamma > 0$, there exists an open and dense subset $E_{i,\gamma} \subseteq \Gamma$ such that for each $\rho \in E$, $\left(\mu_{-i} \times \mu_{-i}^A \times \mu_i\right)\{(\omega_{-i}, a_{-i}, \omega_i) : \rho_i(\omega_i|\omega_{-i}, a_{-i}) > 0\} \geq 1 - \gamma$. The result follows from the fact that

$$\Gamma_0 = \bigcap\nolimits_i \bigcap\nolimits_n E_{i,\frac{1}{n}}.$$

Indeed, take any monitoring ρ^0 and $\varepsilon > 0$. Let $\rho = \frac{\varepsilon}{10} + \left(1 - \frac{\varepsilon}{10}\right)\rho_0$ and let $\delta = \frac{\varepsilon}{100}$. Then, $\|\rho^0 - \rho\| < \frac{\varepsilon}{2}$, and for each monitoring ρ' such that $\|\rho - \rho'\| < \delta$, $\rho' \in E_{i,\gamma}$. □

Lemma 16. *For each monitoring* $\rho \in \Gamma_0$, *each* $\eta > 0$, *there exists constant* $C < \infty$, *and monitoring* $\rho^C \in \Gamma_0$ *such that* $\left\|\rho - \rho^C\right\| \leq \eta$, *and for each player* i, $\rho_i^{C*} \leq C$, *almost surely.*

Proof. W.l.o.g. assume that $\eta < 1$. For each player i, each ω_{-i}, a_{-i}, define sets

$$\begin{aligned}A_i^C(\omega_{-i}, a_{-i}) &= \left\{\omega_i : \rho_i(\omega_i|\omega_{-i}, a_{-i}) > \frac{C}{4}\right\}, \\ A_i^C &= \left\{(\omega_i, \omega_{-i}, a_{-i}) : \omega_i \in A_i^C(\omega_{-i}, a_{-i})\right\}.\end{aligned}$$

Let $\mathbf{1}_{A_i^C}$ be the indicator function of set A_i^C. Because ρ is square integrable,

$$\lim_{C\to\infty} \max_i \left(\left\|\mathbf{1}_{A_i^C}\right\|_2, \left\|\rho_i \mathbf{1}_{A_i^C}\right\|_2\right) = 0.$$

Find $C < \infty$ large enough, so that for each player i,

$$\left\|\mathbf{1}_{A_i^C}\right\|_2, \left\|\rho_i \mathbf{1}_{A_i^C}\right\|_2 \leq \left(\frac{\eta}{100}\right)^2.$$

For each player i, each ω_{-i}, a_{-i}, define set

$$D_i^C = \left\{(\omega_{-i}, a_{-i}) : \mu_i\left[\mathbf{1}_{A_i^C(\omega_{-i}, a_{-i})}\rho_i(.|\omega_{-i}, a_{-i})\right] > \frac{\eta}{100}\right\}.$$

Then,

$$(\mu_{-i} \times \mu_{-i}) \left[\mathbf{1}_{D_i^C}\right] \leq \left\|\mathbf{1}_{A_i^C}\rho_i\right\|_2 \left(\frac{\eta}{100}\right)^{-1} \leq \frac{\eta}{100}.$$

For each $(\omega_{-i}, a_{-i}) \notin D_i^C$, define constant

$$a_i^C(\omega_{-i}, a_{-i}) = \frac{\mu_i\left[\mathbf{1}_{A_i^C(\omega_{-i},a_{-i})}\rho_i(.|\omega_{-i}, a_{-i})\right] - \frac{C}{4}\mu_i\left(A_i^C(\omega_{-i}, a_{-i})\right)}{1 - \mu_i\left[\mathbf{1}_{A_i^C(\omega_{-i},a_{-i})}\rho_i(.|\omega_{-i}, a_{-i})\right]}.$$

Then, because $\frac{C}{4}\mu_i\left(A_i^C(\omega_{-i}, a_{-i})\right) \leq \left\|\mathbf{1}_{A_i^C}\rho_i\right\|_2$, it must be that,

$$\left|a_i^C(\omega_{-i}, a_{-i})\right| \leq \frac{\eta/100 + \eta/100}{1 - \eta/100} \leq \frac{\eta}{10}.$$

Construct monitoring ρ^C : For each player i, let

$$\begin{aligned}
&\rho_i^C(\omega_i|\omega_{-i}, a_{-i}) = 1, \text{ if } (\omega_{-i}, a_{-i}) \in D_i^C \\
&\rho_i^C(\omega_i|\omega_{-i}, a_{-i}) = \frac{C}{4}, \text{ if } \omega_i \in A_i^C(\omega_{-i}, a_{-i}), \text{ and } (\omega_{-i}, a_{-i}) \notin D_i^C, \\
&\rho_i^C(\omega_i|\omega_{-i}, a_{-i}) = \left(1 + a_i^C(\omega_{-i}, a_{-i})\right)\rho_i(\omega_i|\omega_{-i}, a_{-i}), \text{otherwise.}
\end{aligned}$$

Then, $\mu_i\left[\rho_i^C(.|\omega_{-i}, a_{-i})\right] = 1$, almost surely. Moreover, $\rho^C \in \Gamma_0$, $\rho^C < C$, almost surely, and

$$\begin{aligned}
&\left\|\rho - \rho^C\right\| \\
&\leq \sum\nolimits_i \left[2(\mu_{-i} \times \mu_{-i})\left[\mathbf{1}_{D_i^C}\right] + \left\|\mathbf{1}_{A_i^C}\rho_i\right\|_2 + \left|\sup_{\omega_{-i},a_{-i}} a_i^C(\omega_{-i}, a_{-i})\right| \|\rho_i\|_2\right] < \eta.
\end{aligned}$$

□

Lemma 17. *Suppose that monitoring $\rho \in \Gamma_0$ is bounded, i.e., there exists constant $C < \infty$, so that for each player i, $\rho_i \leq C$, almost surely. Then, for each $\eta > 0$, each player i, each $(p, g, f) \in \Omega_i^*$ such that $p > 0$ and $f > 0$, each $\gamma > 0$, there exists monitoring $\rho' \in \Gamma_0$ such that $\|\rho - \rho'\| \leq \eta$, and $\mu_i\left(\omega_i : \upsilon^{\rho'}(\omega_i) = (p, g, f)\right) > 0$.*

Proof. Without loss of generality, assume that $\eta < 1$. Define constants

$$\begin{aligned}
A_i &= \int p(\omega_{-i}, a_{-i})\, d\left(\mu_{-i} \times \mu_{-i}^A\right), \\
A_{-i} &= \int f(\omega_{-i}|a_i)\, d\left(\mu_{-i} \times \mu_i^A\right).
\end{aligned}$$

Let $\mathbf{1}_S$ be the indicator function of set S. Find set $S \subseteq \Omega_i(g)$ such that $\mu_i(S) > 0$, and that

$$\mu_i(S)(1 + A_1 + A_2 + C), \max_j \|\mathbf{1}_S\rho_j\|_2 \leq \frac{\eta}{100}.$$

Let

$$B_i(\omega_{-i}, a_{-i}) = \int \mathbf{1}_S \rho_i(\omega_i|\omega_{-i}, a_{-i})\, d\mu_i,$$
$$D_i(\omega_{-i}, a_{-i}) = \frac{\mu_i(S)\, p(\omega_{-i}, a_{-i}) - B_i(\omega_{-i}, a_{-i})}{1 - B_i(\omega_{-i}, a_{-i})},$$

Then,

$$\sup_{\omega_{-i}, a_{-i}} B_i(\omega_{-i}, a_{-i}) \le C\mu_i(S) \le \frac{\eta}{100},$$
$$\|D_i\|_2 \le 2\mu_i(S) A_i + 2 \sup_{\omega_{-i}, a_{-i}} B_i(\omega_{-i}, a_{-i}) \le \frac{\eta}{10}.$$

Define monitoring $\rho' = (\rho'_1, \rho'_2)$: For each $\omega_{-i} \in \Omega_{-i}$, and each a_i, a_{-i},

- for each $\omega_i \in S$, let

$$\rho'_i(\omega_i|\omega_{-i}, a_{-i}) = p(\omega_{-i}, a_{-i}),$$
$$\rho'_{-i}(\omega_{-i}|\omega_i, a_i) = f(\omega_{-i}|a_i), \text{ and}$$

- for each $\omega_i \notin S$, let

$$\rho'_i(\omega_i|\omega_{-i}, a_{-i}) = (1 - D_i(\omega_{-i}, a_{-i}))\, \rho^*_i(\omega_i|\omega_{-i}, a_{-i}),$$
$$\rho'_{-i}(\omega_{-i}|\omega_i, a_i) = \rho^*_{-i}(\omega_{-i}|\omega_i, a_i).$$

Then,

$$\mu_i[\rho'_i(.|\omega_{-i}, a_{-i})] = \mu_{-i}\left[\rho'_{-i}(.|\omega_i, a_i)\right] = 1,$$

ρ' has full support,

$$\mu_i\left(\omega_i : \upsilon^{\rho'}(\omega_i) = (p, g, f)\right) \ge \mu_i(S_i) > 0,$$

and

$$\|\rho - \rho'\|_\Gamma \le \|D_i\|_2 + \sum\nolimits_j \|\mathbf{1}_S \rho_j\|_2 + \sum\nolimits_j \mu_i(S) A_j \le \eta.$$

□

Lemma 18. *For each player i, each $(p, g, f) \in \Omega^*_i$, each $\gamma > 0$, and monitoring $\rho \in \Gamma_0$ such that $\mu_i(\omega_i : \upsilon^\rho(\omega_i) = (p, g, f)) > 0$, there exists $\eta > 0$ so that for each monitoring $\rho' \in \Gamma_0$, if $\|\rho - \rho'\| \le \eta$, then $\mu_i\left(\omega_i : \upsilon^{\rho'}(\omega_i) \in B(p, g, f, \gamma)\right) > 0$.*

Proof. Fix player i, $(p, g, f) \in \Omega^*_i$, $\gamma > 0$, and monitoring $\rho \in \Gamma_0$ such that

$$\mu_i(\omega_i : \upsilon^\rho(\omega_i) = (p, g, f)) > 0.$$

There exists

$$S \subseteq \{\omega_i : \upsilon^\rho(\omega_i) = (p, g, f)\}$$

and $r > 0$ such that $\mu_i(S) > 0$, and for each $\omega_i \in S$,

$$\left|\left(\mu_{-i} \times \mu_i^A\right)\left[\rho_i\left(\omega_i|\omega_{-i}, a_{-i}\right)\right] - r\right| \leq \frac{1}{2}r.$$

Let

$$x = \frac{\gamma(1+r)}{10r}, \text{ and } \eta = \frac{1}{2}x\mu_i(S).$$

For each monitoring ρ' such that $\|\rho - \rho'\| \leq \eta$, define sets

$$S_i(x, \rho') = \left\{\omega_i \in S : \|\rho_i(\omega_i|\omega_{-i}, a_{-i}) - \rho'_i(\omega_i|\omega_{-i}, a_{-i})\|_{\Omega_{-i}\times A_{-i}} \leq x\right\},$$
$$S_{-i}(x, \rho') = \left\{\omega_i \in S : \frac{1}{|A_i|}\sum_{a_i \in A_i}\left\|\rho_{-i}(\omega_{-i}|\omega_i, a_i) - \rho'_{-i}(\omega_{-i}|\omega_i, a_i)\right\|_{\Omega_{-i}} \leq x\right\}.$$

Then,

$$\begin{aligned}
\eta &\geq \|\rho - \rho'\| \geq \|(\rho - \rho')\mathbf{1}_S\| = \|(\rho_1 - \rho'_1)\mathbf{1}_S\| + \|(\rho_2 - \rho'_2)\mathbf{1}_S\| \\
&\geq \mu_i\left[\mathbf{1}_S\|\rho_i(\omega_i|\omega_{-i}, a_{-i}) - \rho'_i(\omega_i|\omega_{-i}, a_{-i})\|_{\Omega_{-i}\times A_{-i}}\right] \\
&+ \left(\mu_i \times \mu_i^A\right)\left[\mathbf{1}_S\left\|\left(\rho_{-i}(\omega_{-i}|\omega_i, a_i) - \rho'_{-i}(\omega_{-i}|\omega_i, a_i)\right)\right\|_{\Omega_{-i}}\right] \\
&\geq \left(\mu_i(S\backslash S_i(x, \rho')) + \mu_i(S\backslash S_{-i}(x, \rho'))\right)x \\
&\geq x\mu_i\left(S\backslash\left(S_i(x, \rho') \cap S_{-i}(x, \rho')\right)\right),
\end{aligned}$$

and

$$\mu_i\left(S_i(x, \rho') \cap S_{-i}(x, \rho')\right) \geq \mu_i(S) - \frac{\eta}{x} \geq \frac{1}{2}\mu_i(S) > 0.$$

Moreover, for each $\omega_i \in S_i(x, \rho') \cap S_{-i}(x, \rho')$,

$$\begin{aligned}
&d^*\left(\upsilon^\rho(\omega_i), \upsilon^{\rho'}(\omega_i)\right) \\
&= \left\|\frac{\rho_i(\omega_i|\omega_{-i}, a_{-i})}{\left(\mu_{-i} \times \mu_{-i}^A\right)\left[\rho_i(\omega_i|., .)\right]} - \frac{\rho'_i(\omega_i|\omega_{-i}, a_{-i})}{\left(\mu_{-i} \times \mu_{-i}^A\right)\left[\rho'_i(\omega_i|., .)\right]}\right\|_{\Omega_{-i}\times A_{-i}} \\
&+ \frac{1}{|A_i|}\sum\nolimits_{a_i}\left\|\rho_{-i}(.|\omega_i, a_i) - \rho'_{-i}(.|\omega_i, a_i)\right\|_{\Omega_{-i}} \\
&\leq \frac{x}{\frac{1}{2}r} + \frac{x\left(\frac{3}{2}r + x\right)}{\frac{1}{2}r} + x = 4x\left(\frac{1+r}{r}\right) < \gamma,
\end{aligned}$$

and $\upsilon^{\rho'}(\omega_i) \in B(p, g, f, \gamma)$. □

Lemma 19. *For each player i, each $(p, g, f) \in \Omega_i^*$ such that $p > 0$ and $f > 0$, almost surely, each $\gamma > 0$, there exists an open and dense subset $E \subseteq \Gamma_0$ (open and dense in Γ_0) such that for each $\rho \in E$, $\mu_i(\omega_i : \upsilon^\rho(\omega_i) \in B(p, g, f, \gamma)) > 0$.*

Proof. Fix player i. In steps, we show that for each monitoring $\rho^* \in \Gamma_0$, each $\eta > 0$, there exists monitoring ρ and $\eta' > 0$ such that $\|\rho - \rho^*\| \leq \eta$, and for each ρ' so that $\|\rho - \rho'\| \leq \eta'$, $\mu_i (\omega_i : \upsilon^\rho (\omega_i) \in B(p, g, f, \gamma)) > 0$.

By Lemma 16, there exists constant $C < \infty$, and monitoring ρ^{C*} such that $\left\|\rho^* - \rho^{C^*}\right\| \leq \frac{\eta}{2}$, and for each player i, $\rho_i^{C*} \leq C$, almost surely.

By Lemma 17, there exists monitoring ρ such that $\left\|\rho^{C*} - \rho\right\| \leq \frac{\eta}{2}$, and $\mu_i (\omega_i : \upsilon^\rho (\omega_i) = (p, g, f)) > 0$.

By Lemma 18, there exists $\eta' > 0$ such that for each $\rho' \in \Gamma_0$, if $\|\rho - \rho'\| \leq \eta'$, then $\mu_i \left(\omega_i : \upsilon^{\rho'} (\omega_i) \in B(p, g, f, \gamma)\right) > 0$.

The Lemma follows. □

We can finish the proof of Theorem 1. For each player i, Ω_i^* is Polish, hence separable, and there exists a countable dense subset $Q_i \subseteq \Omega_i^* \times (0, 1)$. We can assume that for each $(p, g, f, \gamma) \in Q_i$, $p > 0$ and $f > 0$, almost surely. By Lemma 19, for each $q = (p, g, f, \gamma) \in \Omega_i$, there exists open and dense subset $E_{i,q} \subseteq \Gamma_0$ such that for each $\rho \in E$, $\mu_i (\omega_i : \upsilon^\rho (\omega_i) \in B(p, g, f, \gamma)) > 0$.

Because Γ_0 is non-meagre,

$$\Gamma^* = \Gamma_0 \cap \bigcap\nolimits_i \bigcap\nolimits_{q \in Q_i} E_{i,q}$$

is a non-meagre subset of Γ. Of course, each monitoring $\rho \in \Gamma^*$ is extremely rich.