

An anti-folk theorem for finite past equilibria in repeated games with private monitoring

MARCIN PEŃSKI

Department of Economics, University of Toronto

We prove an anti-folk theorem for repeated games with private monitoring. We assume that the strategies have a *finite past* (they are measurable with respect to finite partitions of past histories), that each period players' preferences over actions are modified by smooth idiosyncratic shocks, and that the monitoring is sufficiently *connected*. In all repeated game equilibria, each period play is an equilibrium of the stage game. When the monitoring is approximately connected and equilibrium strategies have a uniformly bounded past, then each period play is an approximate equilibrium of the stage game.

KEYWORDS. Repeated games, anti-folk theorem, private monitoring.

JEL CLASSIFICATION. C73.

1. INTRODUCTION

The basic result of the repeated game literature, the folk theorem, shows that any feasible and individually rational payoff can be attained in an equilibrium when players are sufficiently patient (Rubinstein 1979, Fudenberg and Maskin 1986, Fudenberg et al. 1994). Recent results extend the folk theorem to classes of repeated games with private monitoring (Ely and Välimäki 2002, Matsushima 2004, Hörner and Olszewski 2006). The equilibrium strategies constructed in these results are very complex; the strategies often depend on minute details of past histories. It is hard to imagine that such strategies can be used in real-world interactions.

This paper argues that the folk theorem fails when the environment is sufficiently rich and the players have a limited capability of processing information. We make three assumptions.

1. We assume that the private monitoring is *infinite* (i.e., it has infinitely many signals) and *connected*. To explain the last property, observe first that given the players' strategies, each private signal leads to ex post beliefs about the realized actions and signals of the other players. We say that the monitoring is connected if each player's space of signals cannot be divided into two sets so that the beliefs induced by the signals from the first set are significantly different from the beliefs induced

Marcin Peński: marcin.peski@utoronto.ca

I am grateful to Bart Lipman and anonymous referees for comments that greatly improved the exposition. All remaining mistakes are mine. Financial support from the National Science Foundation (Grant 0841848) is gratefully acknowledged.

Copyright © 2012 Marcin Peński. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](http://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.

DOI: 10.3982/TE760

by the signals from the other set. We show that connectedness is a generic property of monitoring technologies with infinitely many signals.

2. We assume a *finite past*, i.e., each period's continuation strategies are measurable with respect to finite partitions of the past histories. The finite past assumption bites only because the connected monitoring has infinitely many signals. Because the assumption's absence requires players to handle infinite amounts of information, the assumption does not seem too restrictive.
3. Finally, we assume that, in each period, the payoffs are affected by smooth independent and identically distributed (i.i.d.) shocks. The payoff shocks correspond to idiosyncratic events that modify the stage-game preferences over actions.

The main result shows that, in any repeated game equilibrium, each period play is an equilibrium of the stage game. The result is a simple consequence of the following observations. Because the monitoring is connected, the set of beliefs induced by i 's past histories is topologically connected. Because each period's payoffs are affected by smooth i.i.d. shocks, the probability that player i plays action a_i in period t is continuous in the expected continuation payoffs and, as a consequence, in the beliefs about the opponents' private histories. Because the continuation strategies are measurable with respect to a finite partition of past histories, the probability of playing a_i is constant over each of the elements of the partition. Finally, the result follows from the fact that any continuous function that is constant over the elements of finite partition of a topologically connected set must be constant over the entire set.

Our result helps to clarify the assumptions behind the folk theorems in repeated games with private monitoring. The infinite and connected monitoring eliminates the possibility of constructions based on finite, almost-public monitoring, as in [Mailath and Morris \(2002, 2006\)](#), [Hörner and Olszewski \(2009\)](#), and [Mailath and Olszewski \(2011\)](#). The smooth payoff shocks eliminate belief-free equilibria ([Ely and Välimäki 2002](#), [Piccione 2002](#), [Ely et al. 2005](#)). The finite past condition excludes the possibility of fine tuning strategies in the belief-based constructions of [Sekiguchi \(1997\)](#) and [Bhaskar and Obara \(2002\)](#). In fact, each of the assumptions is necessary for the result in the following sense: If the monitoring is finite, or the strategies can have an infinite past, or there are no smooth shocks, there exist repeated games with nontrivial equilibria (i.e., equilibria that are not repetitions of the stage-game equilibria).

A version of the main result holds when the connectedness assumption is weakened. We measure the connectedness of monitoring ρ by the supremum $C(\rho)$ over distances between belief sets that are induced by two-element partitions of the players' signals. Then $C(\rho) \in [0, 1]$ and monitoring ρ is connected if and only if $C(\rho) = 0$. There exist finite monitoring technologies with arbitrarily small (but strictly positive) $C(\rho)$. A strategy has a K -bounded past if the size of the partitions of the past histories that make the continuation strategies measurable is uniformly bounded by K across all periods. For example, finite automata strategies ([Aumann 1981](#), [Rubinstein 1986](#)) have bounded pasts. We show that any equilibrium in strategies with a K -bounded past is a repetition of the approximate stage-game equilibria. The approximation is better when the memory size K is smaller or the number of signals and the connectedness of the monitoring

(inverse $C(\rho)$) are higher. The negative relation between the number of continuation strategies and the number of signals stands in some contrast to the main message of the literature of imperfect monitoring, where more signals usually lead to more outcomes.

Mailath and Morris (2002, 2006) raise the importance of bounded rationality, while discussing the robustness of public equilibria to small amounts of private monitoring. Harsanyi (1973) introduces smooth payoff perturbations to show that a large class of (static) mixed strategy equilibria can be *purified*, i.e., approximated by pure strategy equilibria of incomplete information games. Bhaskar et al. (2008) study the purifiability of strategies in the prisoner's dilemma, and show that the one-period memory equilibrium of Ely and Välimäki (2002) cannot be purified by the one-period memory equilibria of the perturbed game. Recently, Bhaskar et al. (2009) show that all purifiable equilibria of repeated games of perfect information in bounded recall strategies are Markovian.¹ In a companion paper, Pęski (2009) studies asynchronous repeated games with a finite past and rich monitoring. The latter assumption is stronger than connectedness, and requires that the set of induced beliefs is a connected and open subset of the space of beliefs. The main result shows that any equilibrium has 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 period of the game. Additionally, if the payoffs are subject to smooth i.i.d. shocks, all equilibria must be Markovian.

The next section presents the model and the definition of connected monitoring. Section 3 defines connected monitoring and studies its genericity. Section 4 describes the finite past assumption. Section 5 states and proves the main result. The last section discusses the model and extensions.

2. MODEL

2.1 Notation

We start with some notation. For each measurable space X , let $\mathcal{M}X$ be the space of (signed) measures on X with bounded total variation and let ΔX be the space of the probability measures. For any measure $\pi \in \Delta X$ and any integrable function $f: X \rightarrow \mathbb{R}$, let $\pi[f] = \int_X f(x) d\pi(x)$ denote the integral of f with respect to π . For any $\pi \in \mathcal{M}X$, let $\|\pi\|_X = \sum_x |\pi(x)|$ denote the total variation of measure π .

For each collection of sets $X^1, \dots, X^N$ and for each $i \leq N$, we denote $X = \prod_i X^i$ for the product of the sets and $X^{-i} = \prod_{j \neq i} X^j$ for the product of all sets except X^i . Similarly, if $\mu^1, \dots, \mu^N$ are measures on sets X^i , then we write $\mu = \prod_i \mu^i$ and $\mu^{-i} = \prod_{j \neq i} \mu^j$ for independent products of measures on, respectively, sets X and X^{-i} .

2.2 Stage game

There are N players. Each player i observes random variable $\varepsilon^i \in [0, 1]$, chooses action a^i from a finite set A^i , observes signal ω^i from finite or countably infinite set Ω^i , and

¹An early example of the use of payoff shocks to eliminate repeated game equilibria is contained in Bhaskar (1998).

receives payoffs equal to the sum of complete information payoffs and payoff shock

$$u^i(a^i, \omega^i, \varepsilon^i) = g^i(a^i, \omega^i) + \beta^i(a^i, \varepsilon^i).$$

Shocks ε^i are drawn independently across players from distribution $\lambda^i \in \Delta[0, 1]$. The profile of signals $\omega = (\omega^1, \dots, \omega^N)$ is drawn jointly from distribution $\rho(a) \in \Delta\Omega$, where $a \in A$ is the profile of actions. The function $\rho: A \rightarrow \Delta\Omega$ is called the *monitoring technology* (or monitoring). We assume that $|g^i(a^i, \omega^i)| \leq M$ for all actions a^i and signals ω^i . We assume that $\beta^i(\cdot, \cdot)$ is a measurable function such that $\sup_{\varepsilon^i, a^i} |\beta^i(a^i, \varepsilon^i)| \leq 1$.² Notice that $\beta^i(\cdot, \varepsilon^i) \in R^{A^i}$ can be treated as a random variable (as a function of random variable ε_i); we assume that the random variable $\beta^i(\cdot, \varepsilon^i) \in R^{A^i}$ has a distribution with a Lebesgue density bounded by $L^{-|A^i|}$ for some $L \geq 1$.

A (*stage game*) *strategy* of player i is a measurable mapping $\alpha^i: [0, 1] \rightarrow \Delta A^i$. We write $\alpha^i(a^i | \varepsilon^i)$ to denote the probability of action a^i after payoff shock ε^i . Let $\mathcal{A}^i$ be the space of strategies. A strategy profile $\alpha = (\alpha^1, \dots, \alpha^N) \in \mathcal{A}$ is an x (*interim*) *equilibrium* for some $x > 0$ if, for all players i , λ^i -almost all ε^i , and all actions $a^i, a^{i'} \in A^i$, such that $\alpha(a^i | \varepsilon^i) > 0$,

$$\lambda^{-i}[\alpha^{-i}(\varepsilon^{-i})[\rho(a^i, a^{-i})[u^i(a^i, \omega^i, \varepsilon^i)]]] \geq x + \lambda^{-i}[\alpha^{-i}(\varepsilon^{-i})[\rho(a^{i'}, a^{-i})[u^i(a^{i'}, \omega^i, \varepsilon^i)]]].$$

2.3 Repeated game

The stage game is repeated for infinitely many periods, and shocks ε^i and signal profiles ω are drawn independently across time (and across players in the case of shocks). In particular, in each period t , players observe shocks ε_t^i to payoffs in period t , choose actions a_t^i , and observe signals ω_t^i . Let $H_t^i = (A^i \times \Omega^i)^{t-1}$ and $J_t^i = [0, 1]^{(t-1)}$ be, respectively, the informative and noninformative histories before period t with typical elements $h_t^i = (a_1^i, \omega_1^i, \dots, a_{t-1}^i, \omega_{t-1}^i)$ and $j_t^i = (\varepsilon_1^i, \dots, \varepsilon_{t-1}^i)$. Let H_∞^i and J_∞^i be the sets of infinite histories.

A (*repeated game*) *strategy* of player i is a mapping $\sigma^i: \bigcup_t H_t^i \times J_{t+1}^i \rightarrow \Delta A^i$ with the interpretation that $\sigma_t^i(a^i | h_t^i, j_t^i, \varepsilon_t^i)$ is the probability of action a^i . Let Σ^i be the set of strategies of player i . For each strategy $\sigma^i \in \Sigma^i$, let $\sigma^i(\cdot | h_t^i, j_t^i): [0, 1] \rightarrow \Delta A^i$ denote the stage-game strategy after histories h_t^i and j_t^i . Let $\sigma^i(h_t^i, j_t^i)$ denote the *continuation strategy* at the beginning of period t after histories h_t^i and j_t^i . Each continuation strategy is an element of the strategy space Σ^i .

Players discount the future with discount factor $\delta < 1$. Let

$$G_i(h_\infty^i, j_\infty^i) = (1 - \delta) \sum_{t \geq 1} \delta^{t-1} u^i(a_t^i, \omega_t^i, \varepsilon_t^i)$$

denote the repeated game payoffs of player i given infinite histories h_∞^i and j_∞^i . A strategy profile $\sigma = (\sigma^1, \dots, \sigma^N)$ induces a distribution over histories $\pi^\sigma \in \Delta(\prod_i H_\infty^i \times J_\infty^i)$. Let $G_i(\sigma) = \pi^\sigma[G_i(h_\infty^i, j_\infty^i)]$ denote the expected payoff of player i given strategy profile σ .

²The exact value of the bound on the payoff from shocks is not important.

A strategy profile σ is an *equilibrium* if, for each player i , for each strategy s^i , $G_i(\sigma) \geq G_i(s^i, \sigma^{-i})$. We do not require that the equilibrium satisfy any notion of subgame perfection. Because of the negative character of our result, not requiring subgame perfection makes the result stronger.

The main results of this paper do not depend on the value of the discount factor. In fact, the results of this paper do not change if the discount factor is not constant or the game is played for finitely many periods. We focus on that case with constant discounting and infinitely many periods, because in that case, the comparison between our results and the folk theorem literature is most striking.

3. CONNECTED MONITORING

Upon playing action a^i and observing signal ω^i , player i forms beliefs about the actions taken and signals observed by the opponents. The beliefs are based on the opponents' strategies as well as the monitoring. For each mixed action profile of the opponents' strategies σ^{-i} , each action a^i , and each signal ω^i , let

$$b^{\rho}(a^{-i}, \omega^{-i} | a^i, \omega^i; \sigma^{-i}) = \frac{\rho(\omega^i, \omega^{-i} | a^i, a^{-i}) \sigma^{-i}(a^{-i})}{\sum_{\omega'^{-i}, a'^{-i}} \rho(\omega^i, \omega'^{-i} | a^i, a'^{-i}) \sigma^{-i}(a'^{-i})}$$

be the ex post belief that the other players played a^{-i} and observed ω^{-i} . The ex post beliefs $b^{\rho}(a^i, \omega^i; \sigma^{-i}) \in \Delta(A^{-i} \times \Omega^{-i})$ are well defined if signal ω^i has a strictly positive probability for each action profile a . In particular, the ex post beliefs are well defined when monitoring ρ has *full support* if, for each player i , each action profile (a_i, a_{-i}) , and each signal ω_i , $\rho(\omega^i | a^i, a^{-i}) > 0$.

We say that two signal–action pairs (a^i, ω^i) and $(a^{i'}, \omega^{i'})$ are γ -close if the distance between the induced beliefs is at most γ uniformly over the opponents' strategies,

$$\sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^{\rho}(a^i, \omega^i; \sigma^{-i}) - b^{\rho}(a^{i'}, \omega^{i'}; \sigma^{-i})\| \leq \gamma.$$

We say that monitoring is (approximately) connected if it has full support and, for each player i , the set of on-path action–signal pairs cannot be divided into two sets such that the beliefs induced by the action–signal pairs from the first set are significantly different from the beliefs induced by the action–signal pairs from the other set. Formally, for any $\gamma > 0$, monitoring ρ is γ -connected if, for each player i , the following conditions are satisfied.

- For each action $a^i \in A^i$ and each subset $W \subsetneq \Omega^i$, there exist $\omega^i \in W$ and $\omega^{i'} \in \Omega^i \setminus W$ such that (a^i, ω^i) and $(a^i, \omega^{i'})$ are γ -close.
- For all actions $a^i, a^{i'} \in A^i$, there exist $\omega^i, \omega^{i'} \in \Omega^i$ such that (a^i, ω^i) and $(a^{i'}, \omega^{i'})$ are γ -close.

Notice that, with full support, the set of on-path action–signal pairs after histories (h_t^i, j_{t+1}^i) is equal to $E^i \times \Omega^i$, where E^i is the support of mixed action $\sigma_t^i(\cdot | h_t^i, j_{t+1}^i)$. Then, for any set $W \subsetneq E^i \times \Omega^i$, there exist $(a^i, \omega^i) \in W$ and $(a^{i'}, \omega^{i'}) \in (A^i \times \Omega^i) \setminus W$ that are

γ -close. Indeed, there are two cases: If there are $a^i, \omega^i, \omega^{i'}$ such that $(a^i, \omega^i) \in W$ and $(a^i, \omega^{i'}) \in (E^i \times \Omega^i) \setminus W$, then the claim follows from the first part of the definition; if not, the claim follows from the second part.

Let the *connectedness* of monitoring ρ , $C(\rho)$, be equal to the infimum over $\gamma > 0$ such that ρ is γ -connected. Say that the monitoring is *connected* if $C(\rho) = 0$. Any nontrivial connected monitoring must be infinite (i.e., $|\Omega| = \infty$).

3.1 Genericity of connected monitoring technologies

We say that monitoring is extremely rich if its signals approximate signals from any other monitoring. Formally, monitoring ρ is *extremely rich* if it has full support and for each action a^i , any other full support monitoring ρ' , any signal $\omega^{i'} \in \Omega^i$, and each $\gamma > 0$, there exists a signal $\omega^i \in \Omega^i$ such that the beliefs induced by $(a^i, \omega^{i'})$ under monitoring ρ' are γ -close to the beliefs induced by (a^i, ω^i) under monitoring ρ uniformly over all distributions over the opponents' actions

$$\sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^\rho(a^i, \omega^i, \sigma^{-i}, \rho) - b^{\rho'}(a^i, \omega^{i'}, \sigma^{-i})\| \leq \gamma.$$

In an extremely rich monitoring, any belief about the opponents' actions and signals is approximated by beliefs induced by some action–signal pairs. For example, there are action–signal pairs that assign arbitrarily high weight to any single action profile of the opponent. Similarly, there are action–signal pairs that assign arbitrarily high weight to the opponents' signals that assign arbitrarily high weight to any of the player's own actions or signals. It is easy to show that an extremely rich monitoring is connected (see [Lemma 1 in Appendix A](#)).

Suppose that Ω^i is countably infinite for each player i . Let $\Gamma = (\Delta\Omega)^A$ be the space of monitoring technologies. Define the norm on Γ : for any $\rho, \rho' \in \Gamma$,

$$\|\rho - \rho'\|_\Gamma = \sup_a \|\rho(a) - \rho'(a)\|_\Omega.$$

The norm makes Γ a Polish space. Recall that a II category subset of Γ contains a countable intersection of open and dense subsets. Because any Polish space is a Baire space, a II category subset is nonempty and dense in Γ .

THEOREM 1. *The set of extremely rich monitoring technologies is II category in Γ .*

Because extremely rich monitoring technologies are connected, the theorem implies that the connected monitoring technologies are II category. Because II category sets are dense, the theorem implies that any monitoring can be approximated by an extremely rich, hence connected, monitoring.

To make the idea of approximation clearer, we argue that any convex combination of a monitoring with finite support and an extremely rich monitoring is extremely rich. For any two monitoring technologies ρ and ρ' , each $\alpha \in (0, 1)$, define a convex combination, so that for all action and signal profiles a and ω ,

$$(\alpha\rho + (1 - \alpha)\rho')(\omega|a) = \alpha\rho(\omega|a) + (1 - \alpha)\rho'(\omega|a).$$

THEOREM 2. *Take any monitoring ρ_0 with finite support, i.e., such that there exists finite set $\Omega_0 \subseteq \Omega$ so that $\rho_0(\Omega_0|a) = 1$ for each $a \in A$. If ρ is an extremely rich monitoring, then, for any $\alpha \in [0, 1)$, monitoring $\alpha\rho_0 + (1 - \alpha)\rho$ is extremely rich.*

Theorem 2 implies that any monitoring ρ with finite support can be approximated by a sequence of extremely rich monitoring technologies that are obtained as convex combinations of ρ with some (fixed) extremely rich monitoring. The proofs of Theorems 1 and 2 can be found in [Appendix A](#).

Despite being generic in the sense described above, the connectedness assumption eliminates some theoretically important types of private monitoring. For example, no conditionally independent monitoring in the sense of [Matsushima \(2004\)](#) can be connected. Indeed, conditional independence requires that signals observed by a player are independent from the signals observed by the opponents, given the action profile. In particular, given any action profile of the opponents, two different actions of player i may (and typically do) lead to very different sets of beliefs about the opponent's actions, regardless of player i 's signal. This may (and typically does) lead to a violation of the second part of the definition of the connected monitoring.³

Below, we explain that public monitoring is not connected.

3.2 Comparison to public and almost-public monitoring

Monitoring ρ is *public* if the sets of signals are equal, $\Omega^1 = \dots = \Omega^N$, and all players observe the same signal with full probability: For each action profile $a \in A$, $\rho(\omega^i = \omega^j|a) = 1$ for all i and j . A nontrivial ($|\Omega^i| > 1$ for at least one player i) public monitoring ρ is not connected and $C(\rho) = 1$. Indeed, for any two signals $\omega^i \neq \omega^{i'}$ of player i , any action profile $a \in A$, and any strategy profile σ^{-i} , the ex post beliefs after signal ω^i assign probability 0 to signal $\omega^{i'}$ and

$$\|b^\rho(a^i, \omega^i; \sigma^{-i}) - b^\rho(a^{i'}, \omega^{i'}; \sigma^{-i})\| = 1.$$

Say that the monitoring ρ' is γ -*public* if there exists public monitoring ρ that is γ -close to ρ' , $\|\rho - \rho'\|_\Gamma \leq \gamma$. In other words, if monitoring ρ' is almost public, then the player expects to receive a signal that is equal to the signals received by other players. By [Theorem 1](#), any public monitoring can be approximated by almost-public and connected monitoring technologies.

The notion of closeness used in [Theorem 1](#) (i.e., the norm $\|\cdot\|_\Gamma$) is an *ex ante* notion of closeness. If the signal spaces $\Omega^1 = \dots = \Omega^N$ are finite and the public monitoring ρ has full support, then the ex ante notion implies a stronger *interim* notion. For sufficiently small $\gamma > 0$, for each signal $\omega^i \in \Omega^i$, player i assigns at least $1 - \gamma$ probability to other players observing the same signal ω^i *uniformly across signals* ω^i . In other words, players always assume that there is an approximate common knowledge of the observed signals. In such a case, an almost-public monitoring can be interpreted as a public monitoring perturbed by mistakes in which, with small probability, players observe the public signal

³I am grateful to an anonymous referee for pointing this out.

incorrectly. This interpretation is present in two recent papers that show the folk theorem with almost-public monitoring and finite automata (Hörner and Olszewski 2009, Mailath and Olszewski 2011). Notice that if γ is sufficiently small, then with finite signal spaces, almost public monitoring is not connected.

In this paper, we assume that the signal space is infinite. Then almost-public monitoring is not close to the public in the interim sense. For example, there exist signals that do not provide any information about the opponents' signals or that provide information that the opponents observed the public signal incorrectly. In a way, our notion of closeness allows for more types than the notion of closeness based on finite type spaces.

4. FINITE PAST

A player i 's strategy σ^i has a *finite past* if, in each period t , there exists a finite partition Π_t^i of t -period histories $H_t^i \times J_t^i$ such that the t -period continuation strategy $\sigma^i(h_t^i, j_t^i)$ is measurable with respect to Π_t^i . Equivalently, strategy σ^i has a finite past if, in each period t , there are finitely many different continuation strategies. The finite past bites only when the signal space and, as a consequence, the history space, is infinite.

A strategy has a *K -bounded past* if, in each period t , strategy σ^i induces finitely many continuation strategies, $|\{\sigma^i(h_t^i, j_t^i) : (h_t^i, j_t^i) \in H_t^i \times J_t^i\}| \leq K < \infty$. A strategy has a *bounded past* if it has a K -bounded past for some K . Thus, a bounded past strategy is a finite past strategy with a bound on the size of partitions of past histories that is uniform across all periods.

The finite and bounded past generalize 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 Σ_0^i such that $\sigma^i(h_t^i, j_t^i) \in \Sigma_0^i$ for each t and each (h_t^i, j_t^i) .⁴ Clearly, a finite automaton has a finite and $|\Sigma_0^i|$ -bounded past, but not all bounded past strategies are implementable by finite automata.

These assumptions have a number of interpretations. First, finite and bounded past capture 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 (or more 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 possible signals observed in any given period with a finite partition of the signal space.

Finally, we describe an important consequence of the finite past assumption. Consider an action taken by player i in period t . In general, this action depends on history

⁴Our definition of finite automata is equivalent to the standard definition with states and transition functions. One can think about states as continuation strategies. If transitions are not stochastic, then they are directly determined by the associated continuation strategies. To model stochastic transitions, a minor reinterpretation of ε_t^i shocks is required. The full generality of the model allows us to treat ε_t^i shocks as composed of two parts: one that affects the value of function $\beta(\cdot, \cdot)$, and one that does not and can be used for pure randomization.

(h_t^i, j_t^i) observed before period t and the payoff shock ε_t^i observed in the beginning of period t . In other words, it is measurable with respect to some partition of $H_t^i \times J_t^i \times [0, 1]$. Because the number of actions is finite, the partition can be chosen so that it has a finite number of elements.

The finite past assumption requires that the partition can be chosen so that it has a product representation. Indeed, if player i 's strategy has finite past, then the continuation strategy at the beginning of period t is measurable with respect to finite partition Π_t^i of the set of period t histories $H_t^i \times J_t^i$. The interpretation is that, at the end of period $t - 1$, the detailed information contained in histories (h_t^i, j_t^i) is effectively processed and replaced by coarse information contained in an element of the partition $\pi \in \Pi_t^i$. The actual action played by player i in period t may also depend on the payoff shock observed in period t . For each $\pi \in \Pi_t^i$, let $s_t^i(a; \pi) \subseteq [0, 1]$ denote the subset of period t payoff shocks such that action a is played after history $(h_t^i, j_t^i) \in \pi$ if and only if the payoff shock belongs to set $s_t^i(a; \pi)$. That implies that the period t action is measurable with respect to the product partition $\Pi_t^i \times \Sigma_t^i$, where Σ_t^i is a partition of set $[0, 1]$ generated by $\{s_t^i(a; \pi)\}_{a \in A_i, \pi \in \Pi_t^i}$. Notice that partition Σ_t^i consists of at most $|A^i| |\Pi_t^i| < \infty$ elements.

One can think about the procedure of replacing a signal from an infinite space by an element of a partition of the space to which the signal belongs as information processing. Then a finite past precludes that the information contained in history h_t^i and shock ε_t^i is processed simultaneously. In Section 6, we discuss the implications of an alternative modelling choice with a different kind of information processing constraint.

5. MAIN RESULT

The main result of the paper characterizes finite past equilibria in repeated games with connected monitoring.

THEOREM 3. *Suppose that the monitoring is connected and that σ is a (repeated game) equilibrium in finite past strategies. Then, for each player i , for each t , there exists stage-game profile α_t^i such that (i) $\alpha_t^i = \sigma^i(\cdot | h_t, j_t)$, π^σ -almost surely, and (ii) α_t^i is an equilibrium of the stage game.*

Theorem 3 has two parts. The first part says that finite past equilibria are essentially history-independent. Because any history-independent strategy has a finite past, the theorem completely characterizes finite past equilibria. The second part is a simple corollary to the first: if the past does not affect the history, then in each period players must play an equilibrium of the stage game.

We explain the intuition behind the theorem using a repeated game with two periods $t = 1, 2$ and two actions for player i , $A^i = \{a, b\}$.⁵ Fix the strategies of the other players σ^{-i} and suppose that the monitoring is connected. We argue that if σ^i is the

⁵In such a game, the players act and receive payoffs only for two periods. As we commented at the end of Section 2.3, the main results remain unchanged when the repeated game is played for finitely many (instead of infinitely many) periods.

best response strategy of player i with a finite past, then player i 's period 2 action does not depend on the signals and shocks observed in period 1.

Denote the expected difference between player i 's period 2 complete information payoff (i.e., absent the payoff shock) from action a and action b conditional on period 1 history h_1^i :

$$\Delta_2^i(h_1^i) = E(g(a, \cdot) | h_1^i) - E(g(b, \cdot) | h_1^i).$$

Then $\Delta_2^i(h_1^i)$ depends on the beliefs about the private histories of the other players after observing history h_1^i .

Similarly, define the difference between payoff shocks associated with actions a and b ,

$$\beta^i(\varepsilon_2^i) = \beta^i(a, \varepsilon_2^i) - \beta^i(b, \varepsilon_2^i). \quad (1)$$

The assumptions on the payoff shocks imply that $\beta^i(\varepsilon_2^i)$ is chosen from a distribution with a Lebesgue density. If strategy σ^i is the best response against the strategies of the opponents, then it should prescribe action a if

$$\Delta_2^i(h_1^i) + \beta^i(\varepsilon_2^i) > 0 \quad (2)$$

and action b if the inequality has the opposite sign.

Contrary to our claim, suppose that player i 's action in period 2 depends nontrivially on period 1's history. Based on the discussion from the previous section, there exists a positive probability set $S \subseteq [0, 1]$ of period 2 payoff shocks and partition of $H(a) \cup H(b) = H_1^i$ such that for each shock $\varepsilon_2^i \in S$, action x is played after period 1 histories in $H(x)$. Because set S has positive probability and because of the assumptions on function β^i , there exist $\varepsilon_a, \varepsilon_b \in S$ such that $\beta^i(\varepsilon_a) > \beta^i(\varepsilon_b)$. Because $H(a)$ and $H(b)$ partition set H_1^i , and due to the connectedness of the monitoring, there exist sequences of histories $h^{n,a} \in H(a)$ and $h^{n,b} \in H(b)$ such that $\lim_n \Delta_2^i(h^{n,a}) = \lim_n \Delta_2^i(h^{n,b})$. But then, for sufficiently high n , either

$$\Delta_2^i(h^{n,b}) + \beta^i(\varepsilon_a^i) > 0 \quad \text{or} \quad \Delta_2^i(h^{n,a}) + \beta^i(\varepsilon_b^i) < 0.$$

This leads to a contradiction to (2).

All three assumptions—infinite and connected monitoring, finite past, and smooth payoff shocks—are important for the theorem. The result may fail if the monitoring is infinite but not connected. For example, if the monitoring is public and it satisfies sufficient identifiability conditions, then the folk theorem in finite past strategies holds. This claim follows from an appropriately modified standard folk theorem with finite public monitoring (Fudenberg et al. 1994). Any public monitoring with infinitely many signals can be reduced to a finite monitoring with subsets of signals treated as a single signal. (Some additional care is required to deal with smooth i.i.d. payoff shocks.) Indeed, when the monitoring is finite, all strategies have a finite past and the finite past assumption does not bite.

The result may fail if the monitoring is private and finite. In [Appendix D](#), we show that if players have sufficiently many actions and signals, and the finite monitoring satisfies a certain generic property, then there exist repeated games with nontrivial equilibria.⁶ More specifically, we construct payoffs and equilibria with the following property: In odd periods, the actions do not depend on the past history and they form a strict stage-game Nash equilibrium. In even periods, the actions nontrivially depend on the signals observed in the preceding period and they form a correlated equilibrium of the stage game with all best responses being strict. The constructed strategies remain a repeated game equilibrium even when payoffs are perturbed by sufficiently small shocks.

The theorem may fail if the monitoring is infinitely connected, but the strategies are not required to have a finite past. We use the construction from [Appendix D](#) to show that for any monitoring (possibly infinite and connected) that is appropriately close to a finite monitoring with a certain generic property, there are nontrivial repeated game equilibria in strategies without finite past. The idea is to take the construction of equilibrium strategies from the game with finite monitoring and show that the construction extends to sufficiently close infinite monitoring. Note that [Theorem 1](#) implies that each finite monitoring can be approximated arbitrarily closely with connected (hence, infinite) monitoring technologies.

Finally, when there are no smooth payoff shocks, [Ely and Välimäki \(2002\)](#) show that it is possible to approximate full cooperation in the repeated prisoner's dilemma with almost perfect monitoring. Because almost perfect monitoring technologies may be connected, this indicates that the smooth payoff shocks are important for the result.

5.1 Approximately connected monitoring

Define distance on the space of stage-game strategies $\mathcal{A}^i$: for any $\alpha, \alpha' \in \mathcal{A}^i$, let

$$\|\alpha - \alpha'\|_{\mathcal{A}^i} = \int \|\alpha(\varepsilon) - \alpha'(\varepsilon)\|_{\mathcal{A}^i} d\lambda^i(\varepsilon).$$

Also, define constant

$$B = \frac{1 - \delta}{100ML|\mathcal{A}|^2}.$$

THEOREM 4. *Suppose that σ is a (repeated game) equilibrium in K -bounded past strategies. If $C(\rho) < B$, then, for each player i , for each t , there exists stage-game profile α_t^i such that (i) $\|\alpha_t^i - \sigma^i(\cdot|h_t, j_t)\|_{\mathcal{A}} < B^{-1}KC(\rho)$, π^σ -almost surely, and (ii) α_t^i is a $(2MNKB^{-1}/(1 - \delta))C(\rho)$ equilibrium of the stage game, π^σ -almost surely.*

By [Theorem 4](#), any equilibrium in bounded-past strategies is approximately history-independent and consists of a series of approximate stage-game equilibria. The quality

⁶A recent paper, [Sugaya \(2011\)](#), claims that the folk theorem holds for repeated games with finite private generic monitorings. Sugaya's construction relies heavily on belief-free techniques; for this reason, it is not clear whether the result extends to games with finite monitoring and smooth payoff shocks.

of approximation improves with the connectedness of the monitoring (i.e., it decreases with $C(\rho)$) and decreases with the size of memory K .

The theorems remain true under various modifications of the basic model: For example, the discount factors may differ across players or time, or the payoffs or the distribution of shocks may depend on time. [Section 6](#) discusses how [Theorem 3](#) changes under an alternative specification of the model timeline.

5.2 Weaker equilibrium notion

Recall that strategy profile $(\sigma^1, \dots, \sigma^I)$ is an equilibrium if there exists no other strategy $\sigma^{i'}$ that is a profitable deviation. If finite past is interpreted as a constraint on memory, one may argue that profitable deviations also have a finite past. Formally, let $\Sigma^{i, < \infty}$ be the set of strategies with a finite past. Say that finite past strategy profile $(\sigma^1, \dots, \sigma^I) \in \times_i \Sigma^{i, < \infty}$ is a *finite past equilibrium* if for all players i and finite past strategies $\sigma^{i'} \in \Sigma^{i, < \infty}$, $G_i(\sigma^{i'}, \sigma^{-i}) \leq G_i(\sigma^i, \sigma^{-i})$. Similarly, let $\Sigma^{i, K}$ denote the set of strategies with K -bounded past and define *K -bounded past equilibrium* as the profile of K -bounded past strategies such that there exists no profitable K -bounded past deviation.

It is easy to notice that [Theorem 3](#) implies that any finite past equilibrium in games with connected monitoring is a sequence of stage-game Nash equilibria. Indeed, this follows from the fact that any (not necessarily finite past) strategy can be appropriately approximated by finite past strategies.⁷ No similar result is known for K -bounded past equilibria.

6. ALTERNATIVE TIMELINE AND ONE-PERIOD MEMORY

As we discuss in [Section 5](#), finite past stops players from simultaneously processing information contained in the private history and the payoff shock. To analyze this interpretation more deeply, we discuss a modification of the model timeline that leads to a different type of constraint.

So far, we have assumed that the shock to t -period payoffs ε_t^i is observed in the beginning of period t . In this section, suppose that the shock ε_t^i to t -period payoffs is observed in the end of period $t - 1$ (and the first shock ε_1^i is observed immediately before period 1). This change implies that ε_t^i is part of a noninformative history observed before period t and that $J_t^i = [0, 1]^t$ (instead of $J_t^i = [0, 1]^{t-1}$). Let $\sigma(h_t^i, j_t^i)$, where $j_t^i = (\varepsilon_1^i, \dots, \varepsilon_t^i)$, be the continuation strategies at the beginning of period t . The definition of strategy σ^i with a finite past remains the same: A player i 's strategy σ^i has a *finite past* if, in each period t , there exists a finite partition Π_t^i of t -period histories $H_t^i \times J_t^i$ such that the t -period continuation strategy $\sigma^i(h_t^i, j_t^i)$ is measurable with respect to Π_t^i ; alternatively, strategy σ^i has a finite past if, in each period t , there are finitely many different continuation strategies.

As in the original model, the finite past assumption makes it impossible for players to simultaneously process information that arrives in different periods. With the

⁷The notion of approximation is an appropriate version of L^1 closeness, and the above claim is analogous to the fact that any measurable function can be approximated by step functions with finitely many steps.

alternative timeline, this means that information prior to period $t - 1$ is processed separately from information $(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i)$ received in period $t - 1$. More precisely, one can show that player i action in period t must be measurable with respect to the product of finite partitions $\Pi_{t-1}^i \times \Sigma_{t-1}^i$, where Π_{t-1}^i is a partition of histories observed before period $t - 1$ and Σ_{t-1}^i is a (finite) partition of the space of observations in period $t - 1$, $A_i \times \Omega_i \times [0, 1]$. The difference between the original model and the alternative timeline is that in the latter case, information about the period t payoff shock ε_t^i can be processed (i.e., replaced by a finite partition) simultaneously with the action and signal observed in period $t - 1$, $(a_{t-1}^i, \omega_{t-1}^i)$.

We describe an important class of strategies with a finite past under the alternative timeline. Say strategy σ^i has *one-period memory* if, for each t , there exist measurable functions $\alpha_t^i: A^i \times \Omega^i \times [0, 1] \rightarrow \Delta A^i$ such that

$$\alpha_t^i(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i) = \sigma^i(\cdot | h_{t-1}^i, j_{t-1}^i, a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i).$$

Thus, period t action depends only on a partition of the space of observations in period $t - 1$, $A^i \times \Omega^i \times [0, 1]$. Because there are finitely many actions, such a partition Σ_{t-1}^i can be chosen to be finite. The actions in periods $t' > t$ do not depend on $t - 1$ or any earlier information. In particular, the continuation strategy in the beginning of period t is measurable with respect to finite partition $\{H_{t-1}^i \times J_{t-1}^i\} \times \Sigma_{t-1}^i$, where $\{H_{t-1}^i \times J_{t-1}^i\}$ is the trivial partition of histories observed before period $t - 1$.

THEOREM 5. *Consider the alternative timeline. Suppose that the monitoring is extremely rich and that σ is a (repeated game) equilibrium in finite past strategies. Then, for each player i , for each t , there exists a measurable function $\alpha_t^i: A^i \times \Omega^i \times [0, 1] \rightarrow \Delta A^i$ such that*

$$\alpha_t^i(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i) = \sigma^i(\cdot | h_{t-1}^i, j_{t-1}^i, a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i),$$

π^σ -almost surely. Moreover, for almost all realizations of ε_t^i , $\alpha_t^i(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i)$ is a degenerate probability distribution concentrated only on one action.

With the alternative timeline and extremely rich monitoring, all finite past equilibria have a one-period memory. The result provides a foundation for one-period memory strategies. This contributes to the literature that analyzes the properties of such strategies (for example, see [Ely and Välimäki 2002](#) and [Bhaskar et al. 2008](#)).

One can easily show that the strategies $\alpha_t^i(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i)$ must take values in pure actions for almost all realizations of payoff shocks ε_t^i . Thus, with the exception of zero-probability events, the strategies in each period depend on partition of the last period strategies, where the size of the partition is finite and uniformly bounded across periods.

The intuition behind [Theorem 5](#) is similar to the intuition behind [Theorem 3](#). Suppose that player i has only two actions, $A^i = \{a, b\}$, and that there are $t < \infty$ periods. As in the original model, the actions are continuous in beliefs and payoff shocks, and the belief space is connected, so the same argument implies that only information received in the same period as information about the recent payoff shock may affect period t actions. In the same time, period t actions might depend on action a_{t-1}^i and signal ω_{t-1}^i , because the action–signal pair $(a_{t-1}^i, \omega_{t-1}^i)$ is processed jointly with payoff shock ε_t^i .

For each history h_t^i , let $\Delta(h_t^i)$ denote the difference between the expected complete information payoff (i.e., absent the payoff shock) from actions a and b in period t after observing history h_t^i . As in [Section 5](#), we can show that if player i 's action depends nontrivially on histories before period $t - 1$, then there exist a positive probability set $S \subseteq A^i \times \Omega^i \times [-1, 1]^{A^i}$ and two sequences of histories $h^{n,a}, h^{n,b} \in H_{t-1}^i$ such that

$$\lim_{n \rightarrow \infty} |\Delta(h^{n,a}, s) - \Delta(h^{n,b}, s)| = 0 \quad \text{for each } s \in S$$

and such that player i plays action a after histories $(h^{n,a}, s)$ and plays action b after histories $(h^{n,b}, s)$ for all $s \in S$. Due to the assumptions on the shock and because set S has a positive probability, there are action and signal pairs (a^i, ω^i) and shocks $\varepsilon_a, \varepsilon_b$ such that $(a^i, \omega^i, \varepsilon_a), (a^i, \omega^i, \varepsilon_b) \in S$ and $\beta^i(\varepsilon_a) > \beta^i(\varepsilon_b)$ (recall that $\beta_i(\varepsilon)$ is defined in (1) as the difference between shock payoffs from actions a and b). But then either

$$\limsup_{n \rightarrow \infty} \Delta(h^{n,b}, a^i, \omega^i, \varepsilon_a) > 0 \quad \text{or} \quad \liminf_{n \rightarrow \infty} \Delta(h^{n,a}, a^i, \omega^i, \varepsilon_b) < 0$$

and at least one of the actions a or b is not a best response after some histories. The proof of the theorem can be found in [Appendix C](#).

APPENDIX A: GENERICITY

Let $\Gamma^+ \subseteq \Gamma$ be the class of monitoring technologies with full support. Let $\Gamma^* \subseteq \Gamma^+$ be the set of extremely rich monitoring technologies.

A.1 Extremely rich monitoring is connected

LEMMA 1. *Each extremely rich monitoring is connected.*

PROOF. Each action–signal (a^i, ω^i) and monitoring ρ can be represented by a belief mapping $b^\rho(a^i, \omega^i): \Delta A^{-i} \rightarrow \Delta(A^{-i} \times \Omega^{-i})$, where

$$b^\rho(a^i, \omega^i)(\sigma^{-i}) = b^\rho(a^i, \omega^i; \sigma^{-i}).$$

Let B be the space of all continuous mappings $b: \Delta A^{-i} \rightarrow \Delta(A^{-i} \times \Omega^{-i})$ and let

$$B^* = \{b^\rho(a^i, \omega^i): (a^i, \omega^i) \in A^i \times \Omega^i, \rho \in \Gamma^+\}.$$

Then B^* is the set of belief mappings. One checks that, with the sup norm, B^* is a convex (hence, connected) subset of B . Moreover, if monitoring ρ is extremely rich, then for each action a^i , set

$$B^\rho(a^i) = \{b^\rho(a^i, \omega^i): \omega^i \in \Omega^i\}$$

is a dense subset of B^* .

Suppose that monitoring ρ is extremely rich. We check that ρ satisfies the first condition of the definition of a connected monitoring. Take any action $a^i \in A^i$ and subset

$W \subsetneq \Omega^i$. Let

$$B_0 = \{b^\rho(a^i, \omega^i) : \omega^i \in W\}$$

$$B_1 = \{b^\rho(a^i, \omega^i) : \omega^i \in \Omega^i \setminus W\}.$$

Then $B_0 \cup B_1 = B^\rho(a^i)$. Because ρ is extremely rich, $\text{cl } B_0 \cup \text{cl } B_1 = B^*$. Because B^* is connected, $\text{cl } B_0 \cap \text{cl } B_1 \neq \emptyset$ and, for each $\gamma > 0$, there exist belief mappings $b_0 \in B_0$ and $b_1 \in B_1$ such that

$$\sup_{\sigma^{-i}} \|b_0(\sigma^{-i}) - b_1(\sigma^{-i})\| \leq \gamma.$$

The signals associated with these belief mappings are γ -close.

We check the second condition. Take any two actions $a^i, a^{i'} \in A^i$. Because $\text{cl } B^\rho(a^i) = B^\rho(a^{i'})$, for each $\gamma > 0$, there exist $b \in B^\rho(a^i)$ and $b' \in B^\rho(a^{i'})$ such that

$$\sup_{\sigma^{-i}} \|b(\sigma^{-i}) - b'(\sigma^{-i})\| \leq \gamma.$$

The signals associated with these belief mappings are γ -close. □

A.2 Proof of *Theorem 1*

LEMMA 2. *For each $\varepsilon > 0$, full support monitoring ρ' , player i , action a^i , and signal ω' , there exists an open and dense subset $U \subseteq \Gamma$ such that for each $\rho \in U$, there exists signal $\omega \in \Sigma^i$ so that*

$$\sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^\rho(a^i, \omega; \sigma^{-i}) - b^{\rho'}(a^i, \omega'; \sigma^{-i})\| \leq \varepsilon.$$

PROOF. Fix player i , monitoring ρ' , action a^i , and signal ω' . We show that for each $\eta > 0$ and each monitoring ρ^* , there exist monitoring ρ_0 and signal ω such that $\|\rho^* - \rho_0\| \leq \eta$, and for each $\sigma^{-i} \in \Delta A^{-i}$,

$$b^{\rho_0}(a^i, \omega; \sigma^{-i}) = b^{\rho'}(a^i, \omega'; \sigma^{-i}). \quad (3)$$

Indeed, assume without loss of generality that $1/(1 - \frac{1}{100}\eta) \leq 2$. Find a signal $\omega \in \Omega^i$ such that $\sup_{a^{-i}} \rho^*(\omega | a^i, a^{-i}) \leq \frac{1}{100}\eta$. Let ρ^{**} be a monitoring such that for each action profile a , $\rho^{**}(a)$ is equal to $\rho^*(\cdot | a)$ conditionally on the fact that the signal of player i is not equal to ω :

$$\rho^{**}(a) = \rho^*(\cdot | \omega^i \neq \omega | a).$$

Let ρ be a monitoring obtained by a convex combination between ρ' and ρ^{**} :

$$\rho_0(a) = \frac{1}{100}\gamma\rho'(a) + (1 - \frac{1}{100}\gamma)\rho^{**}(a).$$

Because monitoring ρ' has full support, $\rho_0(\omega | a) > 0$ for each action profile a and $b^{\rho_0}(a^i, \omega; \sigma^{-i})$ is well defined. Because

$$\rho_0(\omega, \omega^{-i} | a) = \frac{1}{100}\gamma\rho'(\omega', \omega^{-i} | a)$$

for each a and ω^{-i} , (3) holds for any distribution σ^{-i} . Additionally,

$$\|\rho^* - \rho_0\| \leq \|\rho^* - \rho^{**}\| + \|\rho^{**} - \rho_0\| \leq \frac{2}{100}\eta + \frac{1}{100}\eta \leq \eta.$$

Due to the continuity of conditional beliefs given positive probability signals, there exists $\eta' > 0$ such that $b^{\rho_0}(a^i, \omega; \sigma^{-i})$ is well defined for each ρ so that $\|\rho - \rho_0\| \leq \eta'$ and

$$\sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^{\rho}(a^i, \omega; \sigma^{-i}) - b^{\rho'}(a^i, \omega'; \sigma^{-i})\| \leq 2\eta.$$

The lemma follows. $\square$

LEMMA 3. *For each player i and each signal ω^i , there exists an open and dense subset $U \subseteq \Gamma$ such that for each $\rho \in U$, $\rho(\omega^i|a) > 0$ for each profile $a \in A$.*

PROOF. Fix player i and signal ω^i . We show that for each $\eta > 0$ and each monitoring ρ^* , there exists monitoring ρ such that $\|\rho - \rho^*\| \leq \eta$ and $\rho(\omega^i|a) > 0$ for each profile $a \in A$. Indeed, assume without loss of generality that $\eta < 1$.

Define monitoring ρ . For each profile a such that $\rho^*(\omega^i|a) > 0$, let $\rho(\omega|a) = \rho^*(\omega|a)$ for each signal profile $\omega \in \Omega$. If $\rho^*(\omega^i|a) = 0$, then take any distribution $\mu \in \Delta \Omega^{-i}$, and let

$$\rho(\omega^i, \omega^{-i}|a) = \eta\mu(\omega^{-i})$$

and

$$\rho(\omega^{i'}, \omega^{-i}|a) = (1 - \eta)\rho^*(\omega^{i'}, \omega^{-i}|a) \quad \text{for each signal } \omega^{i'} \in \Omega^i \setminus \{\omega^i\}.$$

Then $\|\rho - \rho^*\| \leq \eta$ and $\rho(\omega^i|a) > 0$ for each $a \in A$.

The continuity of measure implies that there exists $\eta' > 0$ such that for each monitoring ρ' such that $\|\rho - \rho'\| \leq \eta'$, $\rho(\omega^i|a) > 0$ for each profile $a \in A$. The lemma follows. $\square$

PROOF OF THEOREM 1. Because space Γ is separable, Γ^+ is separable and there exists a countable dense subset $\Gamma_0^+ \subseteq \Gamma^+$. Due to the continuity of conditional beliefs given positive probability signals, for each $\rho' \in \Gamma^+$, action a^i , and $\omega' \in \Omega^i$, there exists a sequence $\rho'_n \in \Gamma_0^+$ such that

$$\lim_{n \rightarrow \infty} \sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^{\rho'_n}(a^i, \omega'; \sigma^{-i}) - b^{\rho'}(a^i, \omega'; \sigma^{-i})\| = 0.$$

For each $q = (a^i, \omega', \rho', m) \in A^i \times \Omega^i \times \Gamma_0^+ \times N$, use Lemma 2 to find open and dense set $U_q^i \subseteq \Gamma$ such that for each $\rho \in U_q^i$, there exists signal ω so that

$$\sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^{\rho}(a^i, \omega; \sigma^{-i}) - b^{\rho'}(a^i, \omega'; \sigma^{-i})\| \leq \frac{1}{m}.$$

For each player i and signal ω^i , use Lemma 3 to find open and dense set $U_{\omega^i}^i \subseteq \Gamma$ such that for each $\rho \in U_{\omega^i}^i$, $\rho(\omega^i|a) > 0$ for each profile $a \in A$. Define the set

$$U^* = \left(\bigcap_i \bigcap_{q \in Q^i} U_q^i \right) \cap \left(\bigcap_i \bigcap_{\omega^i \in \Omega^i} U_{\omega^i}^i \right).$$

Then $U^* \subseteq \Gamma^*$ and U^* is a II category subset of Γ . □

A.3 Proof of Theorem 2

Let $\rho' = \alpha\rho_0 + (1 - \alpha)\rho$. Because $\alpha < 1$, ρ' has full support. There are finite sets of signals $\Omega_0^i \subseteq \Omega^i$ such that $\rho_0(\times_i \Omega_0^i | a) = 1$ for each $a \in A$. Signals $\omega^i \notin \Omega_0^i$ do not occur with positive probability under monitoring ρ_0 . Hence, for each action–signal pair (a^i, ω^i) such that $\omega^i \notin \Omega_0^i$, the associated ρ' ex post beliefs given a^i and ω^i are equal to the ρ ex post beliefs given the same pair and the same distribution over opponents' actions,

$$b^\rho(a^i, \omega^i; \sigma^{-i}) = b^{\rho'}(a^i, \omega^i; \sigma^{-i}).$$

Because ρ is extremely rich, Ω_0^i is finite, and for each action a^i , and any other monitoring ρ'' , any signal $\omega'' \in \Omega^i$, each $\gamma > 0$, there exists a signal $\omega^i \in \Omega^i \setminus \Omega_0^i$ such that the beliefs induced by (a^i, ω'') under monitoring ρ'' are γ -close to the beliefs induced by (a^i, ω^i) under monitoring ρ uniformly over all distributions over opponents' actions,

$$\sup_{\sigma^{-i} \in \Delta A^{-i}} \|b^\rho(a^i, \omega^i, \sigma^{-i}, \rho) - b^{\rho'}(a^i, \omega^i; \sigma^{-i})\| \leq \gamma.$$

It follows that monitoring ρ' is extremely rich.

APPENDIX B: PROOFS OF RESULTS FROM SECTION 5

B.1 Proof of Theorems 3 and 4

The theorems follow from the following three lemmas. The proofs of the lemmas can be found in Sections B.5–B.7.

Fix equilibrium σ . The first result establishes an equilibrium continuity of actions and ε -shock payoffs with respect to the beliefs. For each player i and each history $h_{i,t}$ that occurs with positive π^σ probability, let $\pi^{i,\sigma}(h_{i,t}) \in \Delta(H_t^{-i} \times J_t^{-i})$ denote the equilibrium beliefs of player i about the private histories of the opponents. For all histories h_t^i and j_t^i , denote the expected utility from the payoff shock after history h_t^i as

$$\beta(h_t^i, j_t^i) = \lambda^i \left[\sum_{a^i} \beta(a^i, \varepsilon^i) \sigma^i(a^i | h_t^i, j_t^i, \varepsilon^i) \right].$$

LEMMA 4. *For each player i , period t , and positive π^σ -probability histories (h_t^i, j_t^i) and $(h_t^{i'}, j_t^{i'})$,*

$$\begin{aligned} \|\sigma^i(\cdot | h_t^i, j_t^i) - \sigma^i(\cdot | h_t^{i'}, j_t^{i'})\|_{\mathcal{A}^i} &\leq x \\ |\beta^i(h_t^i, j_t^i) - \beta^i(h_t^{i'}, j_t^{i'})| &\leq x, \end{aligned} \tag{4}$$

where

$$x = \frac{1}{50} B^{-1} \|\pi^{i,\sigma}(h_t^i) - \pi^{i,\sigma}(h_t^{i'})\|_{H_t^{-i} \times J_t^{-i}}.$$

If $\|\pi^{i,\sigma}(h_t^i) - \pi^{i,\sigma}(h_t^{i'})\|_{H_t^{-i} \times J_t^{-i}} < B$, then there exists an action $a^i \in A^i$ that is played with positive probability after histories (h_t^i, j_t^i) and $(h_t^{i'}, j_t^{i'})$.

We explain the first bound in (4). The bound says that the probability of ε_t^i for which the action played after history $(h_t^i, j_t^i, \varepsilon_t^i)$ is different from the action played $(h_t^{i'}, j_t^{i'}, \varepsilon_t^i)$ can be bounded by the difference between the beliefs induced by the respective histories. To see why, fix an action a_t^i . The payoff consequences of playing a_t^i can be divided into the current complete information game payoffs, the current payoff from shock, and the future continuation payoffs. The sum of the current complete information game payoffs and future continuation payoffs depends, and is continuous in, the beliefs about the private histories of the opponents. Trivially, the current payoff from the shock depends only on the realization of ε_t^i . Thus, for any ε_t^i , if action a_t^i is the best response following history $(h_t^i, j_t^i, \varepsilon_t^i)$, and the beliefs induced by histories (h_t^i, j_t^i) and $(h_t^{i'}, j_t^{i'})$ are sufficiently close, then a_t^i is an almost best response after history $(h_t^{i'}, j_t^{i'}, \varepsilon_t^i)$. If, instead, shock ε_t^i is replaced by a nearby shock $\varepsilon_t^{i'}$ that is slightly more favorable toward a_t^i , then a_t^i becomes the best response after $(h_t^i, j_t^i, \varepsilon_t^{i'})$ and $(h_t^{i'}, j_t^{i'}, \varepsilon_t^{i'})$. Similarly, if a_t^i is not the best response after $(h_t^i, j_t^i, \varepsilon_t^i)$, then it is not the best response after $(h_t^i, j_t^i, \varepsilon_t^{i'})$ and $(h_t^{i'}, j_t^{i'}, \varepsilon_t^{i'})$ for nearby shocks $\varepsilon_t^{i'}$ that are slightly less favorable toward a_t^i . Thus, for the majority of shocks ε_t^i , the actions played after histories $(h_t^i, j_t^i, \varepsilon_t^i)$ and $(h_t^{i'}, j_t^{i'}, \varepsilon_t^i)$ are similar.

The second bound follows from the first, and the last part of the lemma is a consequence of the fact that if stage-game strategies $\sigma^i(\cdot|h_t^i, j_t^i)$ and $\sigma^{i'}(\cdot|h_t^{i'}, j_t^{i'})$ have disjoint support, then their $\|\cdot\|_{\mathcal{A}^i}$ distance is equal to 1.

The second result shows that if the monitoring is approximately connected, then for any strategy profile and any division of the positive probability histories into two sets, there are histories on both sides of the division with close beliefs.

LEMMA 5. *Take any set $F \subsetneq H_t^i$ of histories such that $0 < \pi^\sigma(F) < 1$. If $C(\rho) < B$, then there are π^σ -positive probability histories $h_t^i \in F$ and $h_t^{i'} \in H_t^i \setminus F$ such that*

$$\|\pi^{i,\sigma}(h_t^i) - \pi^{i,\sigma}(h_t^{i'})\|_{H_t^{-i} \times J_t^{-i}} \leq 50C(\rho). \quad (5)$$

Suppose that the strategies in equilibrium profile σ have a finite past. Fix player i . For each $t < \infty$, let $H_t^{i,\sigma} \subseteq H_t^i$ be the set of all π^σ -positive probability histories. Let $K_t = |\{\sigma^i(h_t^i) : h_t^i \in H_t^{i,\sigma}\}| < \infty$ be the number of continuation strategies that are played by player i starting from period t . If σ^i has a K -bounded past, then $K_t \leq K$ for each t . Let $F_1, \dots, F_{K_t} \subseteq H_t^{i,\sigma} \times J_t^i$ be a partition of $H_t^{i,\sigma} \times J_t^i$ into disjoint sets such that, for each $k \leq K_t$, continuation strategies after histories $(h_t^i, j_t^i), (h_t^{i'}, j_t^{i'}) \in F_k$ are equal.

Consider a graph Γ with K_t nodes such that there is an edge between nodes k and k' if and only if there exist histories $(h_{k,t}^i, j_{k,t}^i) \in F_k$ and $(h_{k',t}^{i'}, j_{k',t}^{i'}) \in F_{k'}$ such that $\|\pi^{i,\sigma}(h_{k,t}^i) - \pi^{i,\sigma}(h_{k',t}^{i'})\|_{H_t^{-i} \times J_t^{-i}} \leq 50C(\rho)$. By Lemma 4, and because of the choice of sets F_k , if there is an edge between nodes k and k' , then inequalities (4) hold for all $(h_t^i, j_t^i) \in F_k, (h_t^{i'}, j_t^{i'}) \in F_{k'}$, and $x = B^{-1}C(\rho)$.

By Lemma 5, if $C(\rho) < B$, then graph Γ is connected, i.e., there is a path of links between each pair of nodes. Because the minimum length of such a path is bounded by K_t , it must be that, for all π^σ -positive probability histories, inequalities (4) hold for π^σ -almost all histories (h_t^i, j_t^i) and $(h_t^{i'}, j_t^{i'})$ with $x = B^{-1}K_t C(\rho)$ if $C(\rho) > 0$ and any $x > 0$ if the monitoring is connected.

The last lemma shows that inequalities (4) for sufficiently small x imply an approximate equilibrium of the stage game.

LEMMA 6. *Suppose that σ is an equilibrium profile such that for each player i , period t , inequalities (4) hold for some $x > 0$ and π^σ -almost all histories (h_t^i, j_t^i) and $(h_t^{i'}, j_t^{i'})$. Then, for all players i and π^σ -almost all histories (h_t^i, j_t^i) , (mixed) action profile $\sigma^i(\cdot|h_t^i, j_t^i)$ is $(2MN/(1 - \delta))x$ equilibrium of the stage game.*

B.2 Preliminary results

Suppose that V is a normed vector space with norm $\|\cdot\|_*$.

LEMMA 7. *For all $f, f' \in V$, $\|(1/\|f\|_*)f - (1/\|f'\|_*)f'\|_* \leq 2\|f - f'\|_*/\|f\|_*$.*

PROOF. The triangle inequality implies that $|\|f'\|_* - \|f\|_*| \leq \|f - f'\|_*$ and

$$\begin{aligned} \left\| \frac{1}{\|f\|_*}f - \frac{1}{\|f'\|_*}f' \right\|_* &\leq \left\| \frac{1}{\|f\|_*}f - \frac{1}{\|f\|_*}f' + \frac{1}{\|f\|_*}f' - \frac{1}{\|f'\|_*}f' \right\|_* \\ &\leq \frac{1}{\|f\|_*} \left(\|f - f'\|_* + \frac{|\|f'\|_* - \|f\|_*|}{\|f'\|_*} \|f'\|_* \right) \leq 2 \frac{\|f - f'\|_*}{\|f\|_*}. \quad \square \end{aligned}$$

Take any countable sets X and Y . For each $\pi \in \Delta X$ and each function $f: X \rightarrow \Delta Y$, define probability distribution $\pi * f \in \Delta(X \times Y)$,

$$(\pi * f)(x, y) = \pi(x)f(y|x).$$

For each $\pi \in \Delta(X \times Y)$, for each $y \in Y$ such that $\pi(y) > 0$, let $\pi(\cdot|y) \in \Delta X$ denote the conditional distribution given y .

LEMMA 8. *For any countable sets X , Y , and Z , any two measures $\pi, \pi' \in \Delta X$, and any function $f: X \rightarrow \Delta(Y \times Z)$, there exists $y_0 \in Y$ such that $(\pi * f)(y_0), (\pi' * f)(y_0) > 0$ and*

$$\|(\pi * f)(\cdot|y_0) - (\pi' * f)(\cdot|y_0)\|_{X \times Z} \leq 2\|\pi - \pi'\|_X.$$

PROOF. First we show that there exists positive probability y_0 such that $(\pi * f)(y_0) > 0$ and

$$\sum_{x,z} |\pi(x) - \pi'(x)|f(y_0, z|x) \leq \|\pi - \pi'\|_X \sum_{x,z} \pi(x)f(y_0, z|x). \quad (6)$$

If not, then

$$\begin{aligned} \|\pi - \pi'\|_X &= \|\pi - \pi'\|_X \sum_{x,y,z} \pi(x)f(y, z|x) \\ &< \sum_{x,y,z} |\pi(x) - \pi'(x)|f(y, z|x) \\ &= \sum_x |\pi(x) - \pi'(x)| = \|\pi - \pi'\|_X, \end{aligned}$$

which yields a contradiction. Inequality (6) implies that

$$(\pi' * f)(y_0) \geq (1 - \|\pi - \pi'\|_X)(\pi * f)(y_0) > 0.$$

The lemma follows from Lemma 7. □

LEMMA 9. For all countable sets $X^1, \dots, X^N$, and for $X = \times_i X^i$ and probability measures $\pi_i, \pi'_i \in \Delta X^i$, $i = 1, \dots, N$, if $\pi = \times_i \pi_i$ and $\pi' = \times_i \pi'_i \in \Delta X$ are independent product measures, then

$$\|\pi - \pi'\|_X \leq \sum_i \|\pi_i - \pi'_i\|_{X^i}.$$

PROOF. It is enough to show the claim for $n = 2$. Then

$$\begin{aligned} \|\pi - \pi'\|_{X^1 \times X^2} &= \sum_{x_1 \in X^1, x_2 \in X^2} |\pi_1(x_1)\pi_2(x_2) - \pi'_1(x_1)\pi'_2(x_2)| \\ &\leq \sum_{x_1 \in X^1, x_2 \in X^2} \pi_1(x_1)|\pi_2(x_2) - \pi'_2(x_2)| + \sum_{x_1 \in X^1, x_2 \in X^2} \pi'_2(x_2)|\pi_1(x_1) - \pi'_1(x_1)| \\ &= \|\pi_2 - \pi'_2\| + \|\pi_1 - \pi'_1\|. \end{aligned} \quad \square$$

B.3 Close signals

LEMMA 10. For each monitoring with full support ρ , for any two γ -close action-signal pairs (a^i, ω^i) and $(a^{i'}, \omega^{i'})$, there exists a constant c_* such that, for all a^{-i} ,

$$\|\rho(\omega^i, \cdot | a^i, a^{-i}) - c_* \rho(\omega^{i'}, \cdot | a^{i'}, a^{-i})\|_{\Omega^{-i}} \leq 11\gamma \rho(\omega^i | a^i, a^{-i}).$$

PROOF. By the definition of γ -close pair, for each a^{-i} ,

$$\left\| \frac{\rho(\omega^i, \cdot | a^i, a^{-i})}{\rho(\omega^i | a^i, a^{-i})} - \frac{\rho(\omega^{i'}, \cdot | a^{i'}, a^{-i})}{\rho(\omega^{i'} | a^i, a^{-i})} \right\| \leq \gamma.$$

Additionally, we show that for all $\gamma \leq \frac{1}{10}$, all a^{-i} and a_*^{-i} ,

$$\frac{\rho(\omega^{i'} | a^{i'}, a^{-i})}{\rho(\omega^{i'} | a^{i'}, a_*^{-i})} \leq (1 + 10\gamma) \frac{\rho(\omega^i | a^i, a^{-i})}{\rho(\omega^i | a^i, a_*^{-i})} \leq (1 + 10\gamma)^2 \frac{\rho(\omega^{i'} | a^{i'}, a^{-i})}{\rho(\omega^{i'} | a^{i'}, a_*^{-i})}.$$

It is enough to show that the first inequality holds for all a^{-i} and a_*^{-i} . On the contrary, suppose that the first inequality does not hold for some $\gamma \leq \frac{1}{10}$. Let

$$\beta = \frac{\rho(\omega^i | a^i, a_*^{-i})}{\rho(\omega^i | a^i, a^{-i})}, \quad \beta' = \frac{\rho(\omega^{i'} | a^{i'}, a^{-i})}{\rho(\omega^{i'} | a^{i'}, a_*^{-i})}, \quad \text{and} \quad s = \frac{1}{1 + \beta}.$$

Then $s = (1 - s)\beta$. Consider a distribution σ^{-i} such that $\sigma^{-i}(a^{-i}) = s$ and $\sigma^{-i}(a_*^{-i}) = 1 - s$. Then

$$\begin{aligned} \|b^\rho(a^i, \omega^i; \sigma^{-i}) - b^\rho(a^{i'}, \omega^{i'}; \sigma^{-i})\| &\geq |b(a^{-i}|a^i, \omega^i; \sigma^{-i}) - b^\rho(a^{-i}|a^{i'}, \omega^{i'}; \sigma^{-i})| \\ &= \frac{s}{s + (1-s)\beta} - \frac{s}{s + (1-s)\beta'} \\ &\geq \frac{s(1-s)(10\gamma)\beta}{(s + (1-s)(1+10\gamma)\beta)} \geq \frac{s^2}{(3s)^2}\gamma \geq \frac{10}{9}\gamma > \gamma, \end{aligned}$$

which contradicts the fact that action-signal pairs (a^i, ω^i) and $(a^{i'}, \omega^{i'})$ are γ -close. Thus,

$$\left| 1 - \frac{\rho(\omega^i|a^i, a_*^{-i})\rho(\omega^{i'}|a^{i'}, a^{-i})}{\rho(\omega^i|a^i, a^{-i})\rho(\omega^{i'}|a^{i'}, a_*^{-i})} \right| \leq 10\gamma.$$

Fix a strategy profile $a_*^{-i} \in A^{-i}$ and define constant

$$c_* = \frac{\rho(\omega^i|a^i, a_*^{-i})}{\rho(\omega^{i'}|a^{i'}, a_*^{-i})}.$$

Then, for all a^{-i} ,

$$\begin{aligned} &\|\rho(\omega^i, \cdot|a^i, a^{-i}) - c_*\rho(\omega^{i'}, \cdot|a^{i'}, a^{-i})\|_{\Omega^{-i}} \\ &\leq \left\| \rho(\omega^i, \cdot|a^i, a^{-i}) - \frac{\rho(\omega^i|a^i, a^{-i})}{\rho(\omega^{i'}|a^{i'}, a^{-i})}\rho(\omega^{i'}, \cdot|a^{i'}, a^{-i}) \right\|_{\Omega^{-i}} \\ &\quad + \left\| \frac{\rho(\omega^i|a^i, a^{-i})}{\rho(\omega^{i'}|a^{i'}, a^{-i})}\rho(\omega^{i'}, \cdot|a^{i'}, a^{-i}) - \frac{\rho(\omega^i|a^i, a_*^{-i})}{\rho(\omega^{i'}|a^{i'}, a_*^{-i})}\rho(\omega^{i'}, \cdot|a^{i'}, a^{-i}) \right\|_{\Omega^{-i}} \\ &\leq \left\| \frac{\rho(\omega^i, \cdot|a^i, a^{-i})}{\rho(\omega^i|a^i, a^{-i})} - \frac{\rho(\omega^{i'}, \cdot|a^{i'}, a^{-i})}{\rho(\omega^{i'}|a^{i'}, a^{-i})} \right\| \rho(\omega^i|a^i, a^{-i}) \\ &\quad + \left| 1 - \frac{\rho(\omega^i|a^i, a_*^{-i})\rho(\omega^{i'}|a^{i'}, a^{-i})}{\rho(\omega^i|a^i, a^{-i})\rho(\omega^{i'}|a^{i'}, a_*^{-i})} \right| \rho(\omega^i|a^i, a^{-i}) \\ &\leq 11\gamma\rho(\omega^i|a^i, a^{-i}). \end{aligned} \quad \square$$

B.4 Connected monitoring

For each player i , periods $s < t$, say that history $h_t^i = (h_s^i, a_s^i, \dots, \omega_{t-1}^i)$ is a *continuation* of history h_s^i and write $h_t \geq h_s$. Similarly, define continuation of uninformative histories j_t^i . Let $\pi^J \in \Delta(\times_i J_\infty^i)$ denote the (strategy-independent) distribution over uninformative histories.

For each player i and strategy $\sigma^i \in \Sigma^i$, say that informative history h_t^i is σ^i -consistent with strategy $\sigma^i \in \Sigma^i$ if there exists a π^J -positive probability set of uninformative histories $J \subseteq J_{t+1}^i$ such that for each $j_{t+1}^i \in J$, $s < t$, and histories $h_s^i \leq h_t^i$, $j_{s+1}^i \leq j_{t+1}^i$, the action taken in period s is chosen with positive probability by strategy σ^i , $\sigma^i(a_s^i|h_s^i, j_{s+1}^i) > 0$.

Because of full support, for any profile σ , an informative history has positive π^σ probability if and only if it is σ^i -consistent. In particular, if profile σ is an equilibrium, then the continuation strategy after any consistent informative history and almost all uninformative histories is the best response. (Note that because we do not assume subgame perfection, the continuation strategies do not need to be best responses after nonconsistent histories.)

Let $H_t^{i,\sigma^i} \subseteq H_t^i$ denote the set of σ^i -consistent histories. For each $s \leq t$ and each σ^i -consistent history h_s , let $H_t^i(h_s) := \{h_t \in H_t^{i,\sigma^i} : h_t \geq h_s\}$ be the set of σ^i -consistent t -period continuation histories of h_s .

LEMMA 11. *Fix equilibrium σ and player i . For each σ^i -consistent history h_{t-1}^i , and any action-signal pairs $w, w' \in A^i \times \Omega^i$ that are γ -close for some γ , and such that histories $h_t^i = (h_{t-1}^i, w)$ and $h_t^{i'} = (h_{t-1}^i, w')$ are σ^i -consistent,*

$$\|\pi^{i,\sigma}(h_t^i) - \pi^{i,\sigma}(h_t^{i'})\| \leq 22\gamma.$$

PROOF. Let $h_{t-1}^i = (a_1^i, \omega_1^i, \dots, a_{t-2}^i, \omega_{t-2}^i)$, $w = (a^i, \omega^i)$, and $w' = (a^{i'}, \omega^{i'})$. Find constant c_* from Lemma 10. Then

$$\begin{aligned} & \|\pi^\sigma(h_t^i, \cdot) - c_*\pi^\sigma(h_t^{i'}, \cdot)\|_{H_t^{-i} \times J_t^{-i}} \\ &= \sup_{E \subseteq H_t^{-i} \times J_t^{-i}, E \text{ measurable}} |\pi^\sigma(h_t^i, E) - c_*\pi^{i,\sigma}(h_t^{i'}, E)| \\ &\leq \sum_{h_t^{-i}} \int_{J_t^{-i}} \varphi(h_{t-1}^{-i}, a_{t-1}^{-i}, j_t^{-i}) |\rho(\omega^i, \omega_t^{-i} | a^i, a_t^{-i}) - c_*\rho(\omega^{i'}, \omega_t^{-i} | a^i, a_t^{-i})| d\pi^J(j_t^{-i}) \\ &\leq 2\gamma \sum_{h_t^{-i}} \int_{J_{t+1}^{-i}} \varphi(h_{t-1}^{-i}, j_t^{-i}) \sigma^{-i}(a_t^{-i} | h_t^{-i}, j_{t+1}^{-i}) \rho(\omega^i | a^i, a_{s-1}^{-i}) d\pi^J(j_{t+1}^{-i}) \\ &\leq 11\gamma \|\pi^\sigma(h_t^i, \cdot)\|. \end{aligned}$$

By Lemma 7,

$$\begin{aligned} \|\pi^{i,\sigma}(h_t^i) - \pi^{i,\sigma}(h_t^{i'})\| &\leq \left\| \frac{\pi^\sigma(h_t^i, \cdot)}{\|\pi^\sigma(h_t^i, \cdot)\|} - \frac{\pi^\sigma(h_t^{i'}, \cdot)}{\|\pi^\sigma(h_t^{i'}, \cdot)\|} \right\| \\ &= \left\| \frac{\pi^\sigma(h_t^i, \cdot)}{\|\pi^\sigma(h_t^i, \cdot)\|} - \frac{c_*\pi^\sigma(h_t^{i'}, \cdot)}{c_*\|\pi^\sigma(h_t^{i'}, \cdot)\|} \right\| \\ &\leq 22\gamma. \end{aligned}$$

□

B.5 Proof of Lemma 4

Fix player i . For each action a^i , let $\Sigma^i(a^i)$ be the strategies s that always play a^i in the first period: $s(\varepsilon_1^i) = a^i$ for each $\varepsilon_1^i \in [0, 1]$.

For all actions $a^i, a^{i'} \in A^i$, each threshold $x \in R$, and each belief over the opponents' histories $\pi \in \Delta(H_t^{-i} \times J_t^{-i})$, define

$$\begin{aligned} v(a^i, \pi) &= \sup_{s \in \Sigma(a^i)} \pi[G_i(s, \sigma^{-i}(h_{t-1}^{-i}, j_{t-1}^{-i}))] - (1 - \delta)\lambda^i[\beta(a^i, \varepsilon^i)] \\ E(a^i, a^{i'}, x) &= \{\varepsilon^i \in [0, 1] : \beta(a^i, \varepsilon^i) - \beta(a^{i'}, \varepsilon^i) \geq x\} \\ E(a^i, \pi) &= \bigcap_{a^{i'} \neq a^i} E\left(a^i, a^{i'}, \frac{1}{1 - \delta}(v(a^{i'}, \pi) - v(a^i, \pi))\right). \end{aligned}$$

Then, $v(a^i, \pi)$ is the expected current and future continuation payoff minus the current payoff shock from a strategy that starts with action a_i . Set $E(a^i, a^{i'}, x)$ consists of shocks ε^i such that the corresponding payoff shock to action a^i is higher than shock to $a^{i'}$ by some specified threshold. Set $E(a^i, \pi)$ consists of shocks for which action a^i is the current best response. In particular, for each history (h_t^i, j_t^i) ,

$$\sigma^i(a^i | h_t^i, j_t^i, \varepsilon_t) > 0 \quad \text{only if } \varepsilon_t^i \in E(a^i, \pi^{i, \sigma}(h_t^i)).$$

This observation can be used to show that for all histories (h_t^i, j_t^i) and $(h_t^{i'}, j_t^{i'})$,

$$\begin{aligned} &\|\sigma^i(\cdot | h_t^i, j_t^i) - \sigma^i(\cdot | h_t^{i'}, j_t^{i'})\|_{A^i} \\ &\leq \sum_{a^i \in A^i} |\lambda^i(E(a^i, \pi^{i, \sigma}(h_t^i))) - \lambda^i(E(a^i, \pi^{i, \sigma}(h_t^{i'})))| \\ &\leq \sum_{a^i, a^{i'} \in A^i} \lambda^i\left(E\left(a^i, a^{i'}, \frac{1}{1 - \delta}(v(a^i, \pi^{i, \sigma}(h_t^i)) - v(a^{i'}, \pi^{i, \sigma}(h_t^{i'})))\right)\right) \\ &\quad \left\|E\left(a^i, a^{i'}, \frac{1}{1 - \delta}(v(a^i, \pi^{i, \sigma}(h_t^i)) - v(a^{i'}, \pi^{i, \sigma}(h_t^{i'})))\right)\right\|. \end{aligned} \tag{7}$$

Recall that the λ^i distribution of $\beta^i(\cdot, \varepsilon^i) \in R^{A^i}$ has its Lebesgue density bounded by $L^{-|A^i|}$. Then, for all actions $a^i, a^{i'}$, all $x, y \in R$,

$$\lambda^i(E(a, a', y) \setminus E(a, a', x + y)) \leq \frac{1}{2}Lx.$$

Additionally, notice that function $v(\cdot, \cdot)$ is Lipschitz continuous in π with constant M . In particular, for any two informative histories $h_t^i, h_t^{i'} \in H_t^i$ such that $\|\pi^{i, \sigma}(h_t^i) - \pi^{i, \sigma}(h_t^{i'})\| \leq \eta$, for each action a^i ,

$$|v(a^i, \pi^{i, \sigma}(h_t^i)) - v(a^i, \pi^{i, \sigma}(h_t^{i'}))| \leq M\eta,$$

which, together with (7), implies that

$$\|\sigma^i(\cdot | h_t^i, j_t^i) - \sigma^i(\cdot | h_t^{i'}, j_t^{i'})\|_{A^i} \leq |A^i|^2 \frac{1}{2}L \frac{1}{1 - \delta} 2M\eta \leq \frac{1}{50}B^{-1}\eta.$$

The second bound in (4) follows from the first.

This shows the first part of the lemma. The second part of the lemma follows from the fact that if stage-game strategies $\sigma^i(\cdot | h_t, j_t)$ and $\sigma^i(\cdot | h_t', j_t')$ have disjoint support, then $\|\sigma(\cdot | h_t, j_t) - \sigma(\cdot | h_t', j_t')\|_{A^i} = 1$.

B.6 Proof of Lemma 5

There are two steps in the proof. First, we show that there exist a pair of histories that lead to $50C(\rho)$ -close beliefs. Second, we show that we can choose these histories so that they have positive π^σ probability. In the proof, we use the notation introduced in Appendix B.4.

Fix player i and period t , and set $F \subsetneq H_t^i$. Say that set F is *determined at history* h_s^i for some $s \leq t$ if either $H_t^i(h_s^i) \cap F = \emptyset$ or $H_t^i(h_s^i) \subseteq F$. Find $s \leq t$ and history $h_{s-1}^i \in H_{s-1}^i$ such that F is not determined at h_{s-1}^i , but it is determined at each $h_s^i \geq h_{s-1}^i$. Define sets of action–signal pairs

$$D^0 = \{(a^i, \omega^i) \in A^i \times \Omega^i : (h_{s-1}^i, a^i, \omega^i) \text{ is } \sigma^i\text{-consistent}\}$$

$$F^0 = \{(a^i, \omega^i) \in D^0 : H_t(h_{s-1}^i, a^i, \omega^i) \subseteq F\}.$$

By the choice of period s , F^0 is a nonempty and proper subset of D^0 .

There are action–signal pairs $w \in F^0$ and $w' \notin D^0 \setminus F^0$ that are $C(\rho)$ -close. Indeed, if there are $(a^i, \omega^i) \in F^0$ and $(a^i, \omega^{i'}) \in D^0 \setminus F^0$, then the claim is implied by the first part of the definition of connected monitoring; otherwise, the claim is implied by the second part. Then histories $h_s^i = (h_{s-1}^i, w)$ and $h_s^{i'} = (h_{s-1}^i, w')$ are σ^i -consistent, $H_t^i(h_s^i) \subseteq F$, $H_t^i(h_s^{i'}) \cap F = \emptyset$, and, by Lemma 11,

$$\|\pi^{i,\sigma}(h_s^i) - \pi^{i,\sigma}(h_s^{i'})\|_{H_{s-1}^{-i} \times J_s^{-i}} \leq 22C(\rho). \quad (8)$$

We show that there exist σ^i -consistent histories h_t^i and $h_t^{i'}$ that are continuations of histories h_s^i and $h_s^{i'}$, and that induce $50C(\rho)$ -close beliefs. Let ξ^i be some s -period continuation strategy of player i . (We put more care into the choice of ξ^i below.) Let

$$X = H_{t-1}^{-i}$$

$$Y = (A^i \times \Omega^i)^{t-s}$$

$$Z = (A^{-i} \times \Omega^{-i})^{t-s}.$$

For each x , let $f(x) \in \Delta(Y \times Z)$ be the distribution over s -period continuation histories induced by strategy profile σ^{-i} (continued after private histories x) and continuation strategy ξ^i . By (8) and Lemma 8, there exists player i 's continuation history $h_{t-s}^i \in Y$ such that

$$\|\pi^{i,\sigma}(h_s^i, h_{t-s}^i) - \pi^{i,\sigma}(h_s^{i'}, h_{t-s}^{i'})\|_{H_{t-1}^{-i} \times J_t^{-i}} \leq 50C(\rho).$$

Although histories h_s and h_s' are σ^i -consistent, histories (h_s^i, h_{t-s}^i) and $(h_s^{i'}, h_{t-s}^{i'})$ do not need to be.

We show that histories (h_s^i, h_{t-s}^i) and $(h_s^{i'}, h_{t-s}^{i'})$ are σ^i -consistent given an appropriate choice of the continuation strategy ξ^i and history h_{t-s}^i . The argument follows by induction on $t = s, s+1, \dots$. Suppose that the inductive claim is proven for some $t \geq s$. Then there exist continuation strategy ξ^i and history h_{t-s}^i such that (h_s^i, h_{t-s}^i) and $(h_s^{i'}, h_{t-s}^{i'})$ are σ^i -consistent and (5) holds. By Lemma 4, there exists an action $a^i \in A^i$

that is played with positive probability after histories (h_s^i, h_{t-s}^i) and $(h_s^{i'}, h_{t-s}^{i'})$. Consider a continuation strategy $\bar{\xi}^i$ that is equal to ξ^i but such that $\bar{\xi}^i$ plays action a^i with probability 1 after histories (h_s^i, h_{t-s}^i) and $(h_s^{i'}, h_{t-s}^{i'})$. By the above argument, there exists a continuation history h_{t+1-s}^i such that

$$\|\pi^{i,\sigma}(h_s^i, h_{t+1-s}^i) - \pi^{i,\sigma}(h_s^{i'}, h_{t+1-s}^{i'})\|_{H_t^{-i} \times J_t^{-i}} \leq 50C(\rho).$$

By the choice of continuation strategy $\bar{\xi}^i$, histories (h_s^i, h_{t+1-s}^i) and $(h_s^{i'}, h_{t+1-s}^{i'})$ are σ^i -consistent.

B.7 Proof of Lemma 6

For each player i , fix a σ^i -consistent infinite history h_∞^{i*} and a history j_∞^{i*} . Define a strategy s^i : let $s^i(h_t^i, j_t^i) = \sigma^i(h_t^{i*}, j_t^{i*})$ after all informative and uninformative histories h_t^{i*} and j_t^{i*} . Let $s = (s^1, \dots, s^N)$ be the profile of so-defined strategies. Define

$$\beta^s(h_t^i, j_t^i) = \lambda^i \left[\sum_{a^i} \beta(a^i, \varepsilon^i) s^i(a^i | h_t^i, j_t^i, \varepsilon^i) \right] = \beta^s(h_t^{i*}, j_t^{i*}).$$

By Lemma 9, for all profiles of histories $(h_t, j_t) \in H_t \times J_t$ and each player i ,

$$\begin{aligned} \|\sigma^i(\cdot | h_t^i, j_t^i) - s^i(\cdot | h_t^i, j_t^i)\|_{\mathcal{A}^i} &\leq x \\ \|\sigma(\cdot | h_t, j_t) - s(\cdot | h_t, j_t)\|_{\mathcal{A}^i} &\leq Nx \\ |\beta(h_t^i, j_t^i) - \beta^s(h_t^i, j_t^i)| &\leq x. \end{aligned} \tag{9}$$

Because the strategy profile s does not depend on past histories, for some $h_{t+1} \in H_{t+1}$, let

$$V^{i,s} = \pi^{i,s}(h_{t+1}^i) [G_i(s(h_{t+1}^i), s(h_{t+1}^{-i}))]$$

be the expected $(t+1)$ -continuation payoff of player i given profile s .

For each player i , history h_t^i , and action a^i , let $V^{i,\sigma}(h_t^i, a^i)$ be the expected continuation payoff in period $t+1$ after player i chooses a^i in period t ,

$$V^{i,\sigma}(h_t^i, a^i) = \sup_{s \in \Sigma^i} \pi^{i,\sigma}(h_t^i) \left[\frac{\sigma^{-i}(a^{-i} | h_t^{-i}, j_t^{-i})}{[\rho(\cdot | a_t^i, a^{-i}) [G_i(s(h_t^i, \varepsilon_t^i, a^i, \omega^i), \sigma^{-i}(h_t^{-i}, \varepsilon_t^{-i}, a^{-i}, \omega^{-i}))]]} \right].$$

Simple computations involving bounds (9) show that

$$|V^{i,\sigma}(h_t^i, a^i) - V^{i,s}| \leq MNx. \tag{10}$$

Moreover, if $\sigma^i(a^i | h_t^i, j_t^i, \varepsilon_t^i) > 0$ for some action a^i , histories (h_t^i, j_t^i) , and payoff shock ε_t^i , then for each action $a \neq a^i$,

$$\begin{aligned} \lambda^{-i} [\alpha^{-i}(\varepsilon^{-i}) [\rho(a, a^{-i}) [u^i(a, \omega^i, \varepsilon^i)] - \rho(a^i, a^{-i}) [u^i(a^i, \omega^i, \varepsilon^i)]] \\ \leq \frac{1}{1-\delta} |V^{i,\sigma}(h_t^i, a) - V^{i,\sigma}(h_t^i, a^i)|. \end{aligned} \tag{11}$$

The lemma follows from (10) and (11).

APPENDIX C: PROOF OF THEOREM 5

Suppose that strategies in equilibrium profile σ have a finite past. Fix player i . For each $t < \infty$, let $H_t^{i,\sigma} \subseteq H_t^i$ be the set of all π^σ -positive probability histories.

LEMMA 12. *There exists a function $\hat{\sigma}^i: H_{t-1}^i \times P_i \times F_i \times [0, 1] \rightarrow \Delta A_i$ such that for each $(h_{t-1}^i, a_{t-1}^i, \omega_{t-1}^i) \in H_{t-1}^{i,\sigma}$, almost all $(j_{t-1}^i, \varepsilon_t^i) \in J_t^i$,*

$$\sigma^i(\cdot | h_{t-1}^i, j_{t-1}^i, a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i) = \hat{\sigma}^i(h_{t-1}^i, \theta^\rho(a_{t-1}^i, \omega_{t-1}^i), \varepsilon_t^i)$$

and there exists a constant such that for each $h_{t-1}^{i'} \in H_{t-1}^{i',\sigma}$, all $(p, f) \in P_i \times F_i$,

$$\begin{aligned} & \int \|\hat{\sigma}^i(\cdot | h_{t-1}^i, (p, f), \varepsilon) - \sigma^i(\cdot | h_{t-1}^{i'}, (p, f), \varepsilon)\|_{A^i} d\lambda^i(\varepsilon) \\ & \leq C \|\pi^{i,\sigma}(h_{t-1}^i) - \pi^{i,\sigma}(h_{t-1}^{i'})\|_{H_{t-1}^{-i} \times J_{t-1}^{-i}} + Cd(\theta^\rho(a_{t-1}^i, \omega_{t-1}^i), \theta^\rho(a_{t-1}^{i'}, \omega_{t-1}^{i'})). \end{aligned}$$

The proof of Lemma 12 follows from Lemma 11 and the argument from the proof of Lemma 4. We omit the details.

The proof of Lemma 5 remains unchanged. This implies that for each $\gamma > 0$, each subset of informative positive probability histories $F \subseteq H_{t-1}^{i,\sigma}$, action a_{t-1}^i , and signal ω_{t-1}^i such that $0 < \pi^\sigma(F \times \{(a_{t-1}^i, \omega_{t-1}^i)\}) < 1$, there exists π^σ -positive probability histories $h_t^i \in (F \times \{(a_{t-1}^i, \omega_{t-1}^i)\}) \cap H_t^{i,\sigma}$ and $h_t^{i'} \in H_t^{i',\sigma} \setminus F \times \{(a_{t-1}^i, \omega_{t-1}^i)\}$ so that

$$\|\pi^{i,\sigma}(h_t^i) - \pi^{i,\sigma}(h_t^{i'})\| \leq \gamma.$$

Now, let $F_1, \dots, F_{K_{t-1}}$ be a partition of $H_{t-1}^{i,\sigma} \times J_{t-1}^i$ into disjoint sets such that for each $k \leq K_{t-1}$, continuation strategies after histories $(h_{t-1}^i, j_{t-1}^i), (h_{t-1}^{i'}, j_{t-1}^{i'}) \in F_k$ are equal. Let $\sigma^{i,k}$ be a $(t-1)$ -period continuation strategy after histories in F_k . Let $\hat{\sigma}^{i,k}: P_i \times F_i \times [0, 1] \rightarrow \Delta A_i$ be the function that is associated with $\sigma^{i,k}$ through the thesis of Lemma 12.

LEMMA 13. *For all k and k' , $\hat{\sigma}^{i,k} = \hat{\sigma}^{i,k'}$.*

PROOF. Consider a graph Γ with K_{t-1} nodes such that there is an edge between nodes k and k' if and only if, for each $\gamma > 0$, there exist histories $(h_{k,t-1}^i, j_{k,t-1}^i) \in F_k$ and $(h_{k',t-1}^i, j_{k',t-1}^i) \in F_{k'}$ such that $\|\pi^{i,\sigma}(h_{k,t-1}^i) - \pi^{i,\sigma}(h_{k',t-1}^i)\|_{H_{t-1}^{-i} \times J_{t-1}^{-i}} \leq \gamma$. By the proof of Lemma 5, graph Γ is (graph-theoretically) connected, i.e., a path exists between any two vertices.

Fix $(p, f) \in P^i \times F^i$. Suppose that vertices k and k' are connected with an edge. Then there exist sequences of histories $(h_{t-1}^{i,n}, j_{t-1}^{i,n}) \in F_k$ and $(g_{t-1}^{i,n}, l_{t-1}^{i,n}) \in F_{k'}$ such that

$$\lim_{n \rightarrow \infty} \|\pi^{i,\sigma}(h_{t-1}^{i,n}) - \pi^{i,\sigma}(g_{t-1}^{i,n})\|_{H_{t-1}^{-i} \times J_{t-1}^{-i}} = 0.$$

Fix any actions $a^{i,k}$ and $a^{i,k'}$ that are played with positive probability in the first stage of the continuation strategies σ^k and $\sigma^{k'}$: $\sigma^k(a^k | \varnothing) > 0$ and $\sigma^{k'}(a^{k'} | \varnothing) > 0$. Because the monitoring is extremely rich, there exist sequences of signals $\omega^{i,n}, u^{i,n} \in \Omega^i$ such that

$$\theta^\rho(a^{i,k}, \omega^{i,n}) \rightarrow (p, f) \quad \text{and} \quad \theta^\rho(a^{i,k'}, u^{i,n}) \rightarrow (p, f).$$

Then, by Lemma 12, for almost all $\varepsilon \in [0, 1]$,

$$\begin{aligned}
& \int \|\hat{\sigma}^{i,k}(\cdot | (p, f), \varepsilon) - \hat{\sigma}^{i,k'}(\cdot | (p, f), \varepsilon)\|_{A^i} d\lambda^i(\varepsilon) \\
&= \lim_{n \rightarrow \infty} \int \|\hat{\sigma}^{i,k}(\cdot | \theta^\rho(a_{t-1}^{i,k}, \omega_{t-1}^{i,n}), \varepsilon) - \hat{\sigma}^{i,k'}(\cdot | \theta^\rho(a_{t-1}^{i,k'}, u_{t-1}^{i,n}), \varepsilon)\|_{A^i} d\lambda^i(\varepsilon) \\
&= \lim_{n \rightarrow \infty} \int \|\sigma^i(\cdot | h_{t-1}^{i,n}, j_{t-1}^{i,n}, a_{t-1}^{i,k}, \omega_{t-1}^{i,n}, \varepsilon) - \sigma^i(\cdot | g_{t-1}^{i,n}, j_{t-1}^{i,n}, a_{t-1}^{i,k'}, u_{t-1}^{i,n}, \varepsilon)\|_{A^i} d\lambda^i(\varepsilon) \\
&\leq C \lim_{n \rightarrow \infty} \|\pi^{i,\sigma}(h_{t-1}^{i,n}) - \pi^{i,\sigma}(g_{t-1}^{i,n})\|_{H_{t-1}^{-i} \times J_{t-1}^{-i}} \\
&\quad + C \lim_{n \rightarrow \infty} d(\theta^\rho(a_{t-1}^{i,k}, \omega_{t-1}^{i,n}), \theta^\rho(a_{t-1}^{i,k'}, u_{t-1}^{i,n})) \\
&= 0.
\end{aligned}$$

Thus, $\hat{\sigma}^{i,k} = \hat{\sigma}^{i,k'}$ for all k and k' that are connected by an edge in graph Γ . The lemma follows from the fact that graph Γ is (graph-theoretically) connected. $\square$

Finally, take any σ^i -consistent histories $(h_{t-1}^i, a_{t-1}^i, \omega_{t-1}^i), (h_{t-1}^{i'}, a_{t-1}^{i'}, \omega_{t-1}^{i'}) \in H_t^{i,\sigma}$. Find k and k' such that $(h_{t-1}^i, j_{t-1}^i) \in F_k$ and $(h_{t-1}^{i'}, j_{t-1}^{i'}) \in F_{k'}$. Then, for almost any $j_{t-1}^i, j_{t-1}^{i'}$, and ε_t^i ,

$$\begin{aligned}
\sigma^i(\cdot | h_{t-1}^i, j_{t-1}^i, a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i) &= \sigma^{i,k}(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i) \\
&= \hat{\sigma}^{i,k}(\theta^\rho(a_{t-1}^i, \omega_{t-1}^i), \varepsilon_t^i) \\
&= \hat{\sigma}^{i,k'}(\theta^\rho(a_{t-1}^i, \omega_{t-1}^i), \varepsilon_t^i) \\
&= \sigma^{i,k'}(a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i) \\
&= \sigma^i(\cdot | h_{t-1}^{i'}, j_{t-1}^{i'}, a_{t-1}^i, \omega_{t-1}^i, \varepsilon_t^i).
\end{aligned}$$

APPENDIX D: NONTRIVIAL EQUILIBRIA WITH FINITE MONITORING

In this appendix, we show that if players have sufficiently many actions and signals, the monitoring is finite, and it satisfies a certain generic condition, then there are game payoffs such that the repeated game has nontrivial equilibria. First, we describe a generic property of monitoring technologies. Second, we describe a repeated game without payoff shocks and a nontrivial equilibrium of such a game. We construct payoffs and equilibria with the following properties: In odd periods, the actions do not depend on the past history and they form a strict stage-game Nash equilibrium. In even periods, the actions nontrivially depend on the signals observed in the preceding period and they form a correlated equilibrium of the stage game with all best responses being strict. Third, we argue that the constructed strategies remain a repeated game equilibrium even when payoffs are perturbed by sufficiently small shocks. Finally, we use the same construction to show that there exist nontrivial equilibria without a finite past in a game with the

same payoffs and possibly infinite monitoring technologies that are sufficiently close to the monitoring described in the finite case.

Fix finite signal spaces. Suppose that at least two players have at least five actions and at least two signals. Without any further loss of generality, assume that $|\mathcal{A}^1|, |\mathcal{A}^2| \geq 5$, and $|\Omega^1|, |\Omega^2| \geq 2$. Consider monitoring technologies ρ that satisfy three properties.

1. Full support.

2. *Convex independence* (CI): For each player i , each profile a , the marginal distribution over signals of player i given profile a , $\text{marg}_{\Omega^i} \rho(a)$, does not belong to the convex hull of marginal distributions given all other action profiles, $\{\text{marg}_{\Omega^i} \rho(a'), a' \in A \setminus \{a\}\}$.

3. *Strong identification* (SI): For all players $i \neq j$, action profile a , any two signals $\omega^i, \omega^{i'} \in \Omega^i$, and proper subset $\Omega_0^j \subsetneq \Omega^j$, if $\omega^i \neq \omega^{i'}$, then

$$\rho(\Omega_0^j | \omega^i, a) \neq \rho(\Omega_0^j | \omega^{i'}, a).$$

Here, $\rho(\omega^j | \omega^i, a)$ is the conditional probability of player j observing a signal in set Ω_0^j given action profile a and player i observing signal ω^i .

It is easy to check that when there are sufficiently many (but finitely many) signals, the above properties are satisfied by an open, dense, and full Lebesgue measure subset of all monitoring technologies.

Fix an action profile $a^* = (a^{i*})$. Fix positive probability signal $\omega^{2*} \in \Omega^2$. Find $p^1 \in (0, 1)$ and $e^1 > 0$ such that sets

$$\Omega_-^1 = \{\omega^1 \in \Omega^1 : \rho(\omega^{2*} | \omega^1, a^*) \leq p^1 - e^1\}$$

and

$$\Omega_+^1 = \{\omega^1 \in \Omega^1 : \rho(\omega^{2*} | \omega^1, a^*) \geq p^1 + e^1\}$$

are nonempty and $\Omega^1 = \Omega_-^1 \cup \Omega_+^1$. Let $p^2 = \rho(\Omega_+^1 | \omega^2, a^*)$ and find $e^2 > 0$ such that $\Omega^2 = \Omega_-^2 \cup \{\omega^{2*}\} \cup \Omega_+^2$, where

$$\Omega_-^2 = \{\omega^2 \in \Omega^2 : \rho(\Omega_+^1 | \omega^2, a^*) \leq p^2 - e^2\}$$

and

$$\Omega_+^2 = \{\omega^2 \in \Omega^2 : \rho(\Omega_+^1 | \omega^2, a^*) \geq p^2 + e^2\}.$$

Such p^i and e^i exist because of SI.

By CI, we can find payoff functions $g_i: A^i \times \Omega^i \rightarrow R$ such that the following statements hold.

- For any player $j \neq 1, 2$, action a^{*j} is strictly dominant.
- Action profile a^* is stage-game Nash equilibrium.

- Action a_-^1 is the strict best response of player 1 if player 2 plays action a_-^2 or a_+^2 with probability no more than $p^1 - e^1$ and action a_0^2 with the remaining probability.
- Action a_+^1 is the strict best response of player 1 if player 2 plays action a_-^2 or a_+^2 with probability no less than $p^1 + e^1$ and action a_0^2 with the remaining probability.
- Action a_-^2 is the strict best response of player 2 if player 1 plays action a_+^1 with probability no more than $p^2 - e^2$ and action a_-^2 with the remaining probability.
- Action a_0^2 is the strict best response of player 2 if player 1 plays action a_+^1 with probability in the interval $(p^2 - \frac{1}{2}e^2, p^2 + \frac{1}{2}e^2)$ and action a_-^2 with the remaining probability.
- Action a_+^2 is the strict best response of player 2 if player 1 plays action a_+^1 with probability no less than $p^2 + e^2$ and action a_-^2 with the remaining probability.
- The absolute value of player 1 payoffs if player 2 plays one of the actions a_-^2, a_0^2 , or a_+^2 is smaller than the lowest cost of deviation of player 1 from stage-game equilibrium a^* .
- The absolute value of player 2 payoffs if player 1 plays one of the actions a_-^1 or a_+^1 is smaller than the lowest cost of deviation of player 2 from stage-game equilibrium a^* .

We construct a nontrivial repeated game equilibrium. In odd periods $t = 1, 3, 5, \dots$, all players play action profile a^* . In even periods $t = 2, 4, 6, \dots$, the play depends on the signals observed in the previous periods. Specifically, the following statements hold.

- Assuming that player 1 chose a_-^1 in period $t - 1$, player 1 plays action a_-^1 if $\omega_{t-1}^1 \in \Omega_-^1$; otherwise, he plays a_+^1 .
- Assuming that player 2 chose a_-^2 in period $t - 1$, player 2 plays action a_-^2 if $\omega_{t-1}^2 \in \Omega_-^2$ and action a_0^2 if $\omega_{t-1}^2 = \omega_-^{2*}$; otherwise, he plays a_+^1 .
- Other players $j \neq 1, 2$ choose a^{*j} .

Because of the choice of payoffs, the play in the even periods is a stage-game correlated equilibrium, and the signals from the previous period play the role of a correlating device.

We verify that the above profile is a strict repeated game equilibrium. Players $j \neq 1, 2$ have no reason to deviate from their stage-game dominant action a^{*j} . Because the cost of deviation from the odd-period stage-game Nash equilibrium outweighs any potential gain for players 1 and 2, the two players follow the strategy in the odd periods. Because the continuation play does not depend on the signals observed in the even periods, assuming that the players followed the equilibrium strategy in the previous period, the equilibrium prescription is a best response in the even periods.

So far, we have discussed a game without payoff shocks. If the shocks are sufficiently small, they do not change the fact that the prescribed actions are strict best responses and the above strategy profile remains a repeated game equilibrium.

Finally, suppose Ω_i are infinite for all players and that ρ is a monitoring with finite support (i.e., there are finite subsets $\Omega_\rho^i \subseteq \Omega^i$ such that for all action profiles $a \in A$, $\rho(\times_i \Omega_\rho^i | a) = 1$) and such that ρ satisfies full support, CI, and SI when restricted to the support. Consider any (possibly infinite and connected) monitoring ρ' that is γ -close to monitoring ρ in the sense of norm $\|\cdot\|$ from Section 3.1. For sufficiently small $\gamma > 0$, if a player observes a signal from the support of ρ , his posterior beliefs are very close to the beliefs that he would hold under the monitoring ρ .

Consider strategies that play a^* in the odd periods and that replicate the behavior of the above constructed profile in odd periods after signals from the support of ρ . For sufficiently small $\gamma > 0$, such a behavior is a best response behavior no matter what the other players are doing on the small probability signals outside the support of ρ . Using an appropriate equilibrium existence theorem (for incomplete information dynamic games with countably many types), we can complete the strategies on the signals outside the support of ρ so that the strategies are best responses to each other after all histories. Because of Theorem 3, strategies so obtained do not have a finite past.

REFERENCES

- Aumann, Robert J. (1981), “Survey of repeated games.” In *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern* (Robert J. Aumann, John C. Harsanyi, Werner Hildenbrand, Michael Maschler, Micha A. Perles, Joachim Rosenmuller, Reinhard Selten, Martin Shubik, and Gerald L. Thompson, eds.), 11–42, Bibliographisches Institut, Mannheim. [26]
- Bhaskar, V. (1998), “Informational constraints and the overlapping generations model: Folk and anti-Folk Theorems.” *Review of Economic Studies*, 65, 135–149. [27]
- Bhaskar, V. and Ichiro Obara (2002), “Belief-based equilibria in the repeated Prisoners’ Dilemma with private monitoring.” *Journal of Economic Theory*, 102, 40–69. [26]
- Bhaskar, V., George J. Mailath, and Stephen Morris (2008), “Purification in the infinitely repeated Prisoner’s Dilemma.” *Review of Economic Dynamics*, 11, 515–528. [27, 37]
- Bhaskar, V., George J. Mailath, and Stephen Morris (2009), “A foundation for Markov equilibria in infinite horizon perfect information games.” Unpublished paper, Department of Economics, University of Pennsylvania. [27]
- Ely, Jeffrey C. and Juuso Välimäki (2002), “A robust Folk Theorem for the Prisoner’s Dilemma.” *Journal of Economic Theory*, 102, 84–105. [25, 26, 27, 35, 37]
- Ely, Jeffrey C., Johannes Hörner, and Wojciech Olszewski (2005), “Belief-free equilibria in repeated games.” *Econometrica*, 73, 377–415. [26]
- Fudenberg, Drew and Eric Maskin (1986), “The Folk Theorem in repeated games with discounting or with incomplete information.” *Econometrica*, 54, 533–554. [25]
- Fudenberg, Drew, David Levine, and Eric Maskin (1994), “The Folk Theorem with imperfect public information.” *Econometrica*, 62, 997–1039. [25, 34]

- Harsanyi, John C. (1973), “Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points.” *International Journal of Game Theory*, 2, 1–23. [27]
- Hörner, Johannes and Wojciech Olszewski (2006), “The Folk Theorem for games with private almost-perfect monitoring.” *Econometrica*, 74, 1499–1544. [25]
- Hörner, Johannes and Wojciech Olszewski (2009), “How robust is the Folk Theorem?” *Quarterly Journal of Economics*, 124, 1773–1814. [26, 32]
- Mailath, George J. and Stephen Morris (2002), “Repeated games with almost-public monitoring.” *Journal of Economic Theory*, 102, 189–228. [26, 27]
- Mailath, George J. and Stephen Morris (2006), “Coordination failure in repeated games with almost-public monitoring.” *Theoretical Economics*, 1, 311–340. [26, 27]
- Mailath, George J. and Wojciech Olszewski (2011), “Folk Theorems with bounded recall under (almost) perfect monitoring.” *Games and Economic Behavior*, 71, 174–192. [26, 32]
- Matsushima, Hitoshi (2004), “Repeated games with private monitoring: Two players.” *Econometrica*, 72, 823–852. [25, 31]
- Pęski, Marcin (2009), “Asynchronous repeated games with rich private monitoring and finite past.” Unpublished paper, Department of Economics, University of Toronto. [27]
- Piccione, Michele (2002), “The repeated Prisoner’s Dilemma with imperfect private monitoring.” *Journal of Economic Theory*, 102, 70–83. [26]
- Rubinstein, Ariel (1979), “Equilibrium in supergames with the overtaking criterion.” *Journal of Economic Theory*, 21, 1–9. [25]
- Rubinstein, Ariel (1986), “Finite automata play the repeated Prisoner’s Dilemma.” *Journal of Economic Theory*, 39, 83–96. [26]
- Sekiguchi, Tadashi (1997), “Efficiency in repeated Prisoner’s Dilemma with private monitoring.” *Journal of Economic Theory*, 76, 345–361. [26]
- Sugaya, Takuo (2011), “Folk Theorem in repeated games with private monitoring.” Unpublished paper, Department of Economics, Princeton University. [35]