

Delayed-Response Strategies in Repeated Games with Observation Lags

Drew Fudenberg

Department of Economics, Harvard University; dfudenberg@harvard.edu

Yuhta Ishii

Department of Economics, Harvard University; ishii2@fas.harvard.edu.

Scott Duke Kominers

Society of Fellows, Department of Economics, and Program for Evolutionary Dynamics, Harvard University, and Becker Friedman Institute for Research in Economics, University of Chicago; kominers@fas.harvard.edu.

Abstract

We extend the folk theorem of repeated games to two settings in which players' information about others' play arrives with stochastic lags. In our first model, signals are almost-perfect if and when they do arrive, that is, each player either observes an almost-perfect signal of period- t play with some lag or else never sees a signal of period- t play. In the second model, the information structure corresponds to a lagged form of imperfect public monitoring, and players are allowed to communicate via cheap-talk messages at the end of each period. In each case, we construct equilibria in “delayed-response strategies,” which ensure that players wait long enough to respond to signals that with high probability all relevant signals are received before players respond. To do so, we extend past work on private monitoring to obtain folk theorems despite the small residual amount of private information.

Keywords: Repeated games; folk theorem; private monitoring; observation lag.

1. Introduction

Understanding when and why individuals cooperate in social dilemmas is a key issue not just for economics but for all of the social sciences,¹ and the theory of repeated games is the workhorse model of how and when concern for the future can lead to cooperation even if all agents care only about their own payoffs. The clearest expression of this idea comes as players become arbitrarily patient; here various folk theorems provide conditions under which approximately efficient payoffs can be supported by equilibrium strategies. Because of the influence of these results, it is important to understand which of their assumptions are critical and which are merely convenient simplifications; a large literature (discussed below) has extended the folk theorems under successively weaker assumptions about the “monitoring structures” that govern the signals players receive about one another’s actions.

Here we relax an assumption which is maintained throughout most of the prior repeated games literature: the assumption that signals of the actions taken in each period (simultaneously) arrive immediately after players’ actions in that period. Instead, we consider repeated games in which the players’ signals about other player’s actions arrive with stochastic and privately observed lags. Our folk theorems for settings with lagged signals show that the assumption that signals are observed immediately is not necessary for repeated play to support cooperation.

To prove these folk theorems, we use the idea of “delayed-response” strategies, under which players wait to respond to signals of a given period’s play for long enough that it is likely (although not certain) that every player has observed the relevant signals by the time players respond to signal information. Although the observation lags generate a form of imperfect private monitoring, the private information here has a special form that allows delayed-response strategies to construct the same set of limit equilibrium payoffs as if the lags were not present.

More specifically, we suppose that players act simultaneously each period, and that players’ actions jointly determine a probability distribution over signals, but that players

- do not observe signals immediately and
- might observe signals asynchronously.

The times at which observation occurs are private information and may be infinite, that is, a particular signal may never arrive. Some sort of observation lags seem plausible in many cases; for example there may be a small probability that a player is momentarily inattentive and temporarily does not see their partner’s actions; more strongly, in some cases a player

¹See e.g., Ahn et al. (2003); Gächter et al. (2004).

may never learn just what happened during moments of inattention. Moreover, information lags of multiple periods seem especially appropriate in settings for which the time period under consideration is extremely short (Fudenberg and Levine (2007a, 2009); Sannikov and Skrzypacz (2010)), and in continuous-time models, where the “period length” is effectively 0 (Bergin and MacLeod (1993); Sannikov (2007); Sannikov and Skrzypacz (2007); Faingold and Sannikov (2011)).²

To prove our folk theorems, we construct *delayed-response strategies*, in which the repeated game is divided into a finite number of “threads,” with play in each thread independent of play in the other threads. Section 3 examines the simplest application of this idea, which is to the case of bounded lags, where there is a K such that every signal arrives within K periods of play. Then, using strategies that have $K + 1$ threads, we can ensure that each thread is equivalent to an instance of the original game (with the original game’s underlying monitoring structure), a smaller discount factor, and *no lag*. Hence if the folk theorem holds in a given repeated game (with any sort of contemporaneous monitoring), the associated strategies can be used to establish a folk theorem—in delayed-response strategies—in the corresponding game with bounded observation lags.³

The rest of the paper allows the lag distribution to have unbounded support, and also allows for a small probability that some signals never arrive at all (corresponding to an infinite observation lag). In these cases the use of delay strategies reduces but does not eliminate the impact of lags, and the game played in each thread has some additional decision-relevant private information. Section 4 considers the case where signals are almost-perfect if and when they do arrive—that is, each player either observes an almost-perfect signal of period- t play with some lag, or else never sees a signal of period- t play.⁴ In our second model, presented in Section 5, players are allowed to communicate (via cheap talk) each period, and the underlying information structure is one of imperfect public monitoring.⁵ In each case, players do not know whether and when other players observe the signals associated with each period’s play, so there is a special but natural form of private information.

²Indeed, physics suggests that the speed of light is a constraint on the speed with which signals can travel.

³The Ellison (1994) study of contagion equilibria uses threads for a rather different purpose: to substitute for public randomization as a way to weaken the effect of a grim-trigger punishment as the discount factor tends to 1. In Sections 3 and 4, we use threads only as a way for the players to wait for lagged signals to arrive; in Section 5, we also use threading in order to weaken the effect of grim-trigger punishments.

⁴In the case of lagged almost-perfect monitoring, we consider only games with two players, so that we may invoke results of Hörner and Olszewski (2006). We do not know whether the folk theorem extends to the analogous setting with n players.

⁵Our analysis assumes that messages are received the instant they are sent, but the results extend to cases in which messages are received much more quickly than observations. We attribute the difference in speeds to the fact that messages are crafted to be easily processed, while processing and interpreting signals can take longer.

For both of our main results, we use a similar proof technique: First, we consider an auxiliary game with “rare” lags in which each player sees a private signal immediately with probability close to (but not equal to) 1. After proving a folk theorem for the auxiliary game with rare lags, we relate the perturbed game with rare lags to the game with possibly long lags by identifying the event in which the signal does not arrive immediately with the event that the signal arrives after some large time T . We then construct equilibria in the game with lags by using delayed-response strategies as described above. For the first main result we prove the folk theorem for the auxiliary game by extending the bloc-strategy construction of Hörner and Olszewski (2006) (henceforth HO2006) to treat as “erroneous” any history of the auxiliary game in which some player observes another’s action with a strictly positive lag; this corresponds to a “real lag” that is longer than the number of threads. The HO2006 construction does not directly apply here, as signals about past play may arrive outside of the relevant block, but we construct equilibria that are belief free for the past T periods provided that the probability of lagged observation is sufficiently small.

To prove the second main result, we first consider a game with private monitoring, communication, and no observation lag. In this game, each player either observes the true action profile or a null signal. We relate this to a game with a public signal that is observed by all players, but where the game ends each period with a fixed small probability, corresponding to strategies in the original game that will use reversion to static Nash equilibrium whenever the reported signals disagree. We prove a sort of folk theorem here using the techniques of Fudenberg et al. (1994) (henceforth FLM) and then again use threads and delayed responses to extend this to a proof for the original game.

1.1. *Related Work*

The repeated games literature has explored successively weaker assumptions on players’ monitoring structures, while maintaining the assumption that signals arrive immediately after play. The first wave of repeated-games models established folk theorems under the assumption that players observe each others’ actions without error at the end of each round of play (Aumann and Shapley (1976), Friedman (1971), Rubinstein (1994), and Fudenberg and Maskin (1986)). Subsequent work extended the folk theorem to cases where agents receive imperfect signals of other agents’ actions, where these signals can either be public (FLM) or private but accompanied by cheap-talk public messages (Compte (1998), Kandori and Matsushima (1998), and Obara (2009)),⁶ or private and without communication (e.g.,

⁶We allow public messages in Section 5. The role of such messages has been studied in a number of subsequent papers, including Ben-Porath and Kahneman (2003), Fudenberg and Levine (2007b), and Escobar and Toikka (2011). Public communication has also been used as a stepping stone to results for games where communication is not allowed (Hörner and Olszewski (2006), Hörner and Olszewski (2009), and Sugaya (2011)).

Sekiguchi (1997), Mailath and Morris (2002), Hörner and Olszewski (2006), and Hörner and Olszewski (2009)). As one step in our argument for the case of lagged almost-perfect monitoring (Section 4), we extend the Hörner and Olszewski (2006) construction to almost-perfect monitoring with rare lags.⁷ With each type of signal structure, the key assumptions relate to the qualitative nature of the information that signals provide: Roughly speaking, in order for the folk theorem to obtain, signals must be informative enough to “identify deviations” in a statistical sense.⁸

The papers of Fudenberg and Olszewski (2011) and Bhaskar and Obara (2011) are the closest to the present work, as in each, the time at which signals arrive is private information. Fudenberg and Olszewski (2011) studied the effect of short privately-known lags in observing the position of a state variable that evolves in continuous time, so that a player observing the state variable at slightly different times would get different readings. Bhaskar and Obara (2011) studied lags that were either deterministic or stochastic with length at most 1. Both papers considered “short lags” and also restricted to the case of a single long-run player facing a sequence of short-run opponents; this paper allows fairly general stochastic lags and considers the case of all long-run players.

Several papers in the stochastic games literature studied deterministic lags of perfect signals (e.g., Lagziel and Lehrer (2012), Levy (2009), and Yao et al. (2011)); this sort of lag does not introduce private information and so is quite different from the lags we study. In Abreu et al. (1991) consecutive signals are grouped together and delivered at once, so the delay does not introduce private information.

2. General Model

This section introduces a general model that encompasses all the settings discussed subsequently. We consider a repeated game with n players $i \in I \equiv \{1, \dots, n\}$, each of whom has a finite action space A_i . In each period $t = 0, 1, 2, \dots$, each player i chooses a possibly mixed action α_i^t ; this generates a sequence of pure action profiles $\{a^t\}_{t=0}^\infty$. Each player i has a finite signal space Ω_i , and there is a private signal structure π over $\Omega \equiv \prod_{i \in I} \Omega_i$; at each time t , a private signal profile is generated by π according to the conditional probability $\pi(\omega_1, \dots, \omega_n \mid a^t)$.

Thus far, the repeated game has the structure of a standard repeated game with private monitoring. We now relax the assumption that players receive signals of period- t play immediately after period t by replacing it with the assumption that the monitoring structure

⁷When the unlagged signals are imperfect, the signals in our auxiliary games are not almost common knowledge in the sense of Mailath and Morris (2002), so the Hörner and Olszewski (2009) construction does not apply.

⁸In addition, the folk theorem has been extended to recurrent stochastic games with perfectly or imperfectly observed actions (Dutta (1995), Fudenberg and Yamamoto (2011), and Hörner et al. (2011)).

is *private with stochastic lags*. As in the usual model, upon the choice of a period- t action profile a^t , a private signal profile ω^t is generated according to the conditional distribution $\pi(\omega^t | a^t)$. However, the players need not immediately observe their components of the signal profile. Instead, we assume that each player i observes his private signal of period- t play, ω_i^t , at a stochastic time $t + L_i^t$, where $\{L^t \equiv (L_1^t, \dots, L_n^t)\}_t$ is a collection of random variables that take values in $(\mathbb{N} \cup \{\infty\})^n$. We assume that the vectors $\{L^t\}$ are distributed identically and independently across t , with probability density function $\lambda : (\mathbb{N} \cup \{\infty\})^n \rightarrow [0, 1]$. We denote by λ_i the density of the marginal distribution of observation lags of player i , L_i . (The case $L_i^t = \infty$ is interpreted as the event in which player i never receives any information about the period- t private signal.) We let Λ_i denote the cumulative marginal distribution function of player i 's observation lags, i.e. $\Lambda_i(\ell) = \sum_{m=0}^{\ell} \lambda_i(m)$, for $\ell \in (\mathbb{N} \cup \{\infty\})$.

Observation of ω_i^t takes place in period $t + L_i^t$ after the choice of that period's actions.⁹ When player i observes ω_i^t , he also observes a “timestamp” indicating that ω_i^t is associated with play in period t . That is, for example, when a player observes that player j played “C” in a prisoner’s dilemma, she is informed about the period to which the observation applies, rather than just getting a signal that “player j played C sometime in the past.”¹⁰

As one concrete example, consider a repeated public goods game in which every period two friends must decide whether or not to exert effort to provide benefits (or gifts) for each other. The friends live far apart, so the benefits must be shared via postal mail. This induces a lag in observation of the realized signals of the friend’s action. Furthermore the postmark dates serve as natural timestamps.

Alternatively suppose that n coauthors who write numerous papers together and suppose that the quality of the paper is determined by the sum of the authors’ efforts. Each period they complete a paper and submit it to a journal. The editor then makes a decision and mails a letter to each of the authors. Here the decision reveals the project’s quality and so provides evidence about partners’ efforts; in a two-player game if the effort \rightarrow quality \rightarrow editor’s letter map is deterministic and monotone, the letter perfectly reveals partners’ efforts, but more typically letters have a stochastic component. Here the project itself serves as a natural timestamp.

Players have perfect recall and receive no further information.

In one part of the paper we allow for communication in every period. Thus, we include message spaces M_i in the general model; when we want to rule out communication we set $M_i = \emptyset$ for each i . After the realization of private signal profile ω^t and after the observation of all

⁹Thus player i cannot respond to the period- $(t + L_i^t)$ observation information until time $t + L_i^t + 1$.

¹⁰The assumption of timestamps renders our model a smaller departure from the usual repeated game monitoring structure than a model in which players observe only an aggregate measure of the frequencies with which opponents took various actions. Note that it is not clear how players would interpret signals received without timestamps when the expected path of play is not constant over time.

private information $\omega_i^{t'}$ for which $t' + L_i^{t'} \leq t$, at each time $t = 0, 1, \dots$, each player i reports a message m_i chosen from the message space M_i . After all of these reports are (simultaneously) submitted, all players immediately observe the message profile $m = (m_1, \dots, m_n)$.

We let H^t denote the set of t -period histories. For a given $h^t \in H^t$ and any $t' \leq t$, we denote by $h^{t,t'}$ the profile of information about the t' -period signal that has been observed by each player. If player i has not yet observed the k -th component of his private signal, $\omega_{i,k}$ in time t' then we specify that $h_{i,k}^{t,t'} = \infty$.

Finally we describe the payoff structure. A sequence of action profiles $\{a^t\}$ chosen by the players generates a total payoff

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t g_i(a^t).$$

In Section 5, we prove a Nash threat folk theorem rather than a full folk theorem. To facilitate this, we fix a Nash equilibrium α^* of the stage game and normalize payoffs of players so that $g_i(\alpha^*) = 0$ for all i . We let V denote the convex hull of the feasible set of payoffs, and let V_{α^*} be the convex hull of the set consisting of $g(\alpha^*) = 0$ and the payoff vectors Pareto-dominating $g(\alpha^*) = 0$: $V_{\alpha^*} \equiv \{v \in V \mid v \geq 0\}$. We assume that $\text{int}(V_{\alpha^*})$ is non-empty.

In contrast, the theorems of Section 3 and 4 concern full folk theorems; thus, we define V^* to be the set of individually rational payoffs of V . With this notation, we are ready to discuss our folk theorems.

We let $G(\delta, \pi, \lambda)$ be the repeated game with discount factor δ , lag distribution λ , and monitoring structure π , and let $E(\delta, \pi, \lambda)$ denote the set of sequential equilibrium payoffs of $G(\delta, \pi, \lambda)$. We let $G(\delta, \pi) \equiv G(\delta, \pi, \text{imm})$, where imm is the (degenerate) distribution which puts full weight on immediate observation, and define $E(\delta, \pi) \equiv E(\delta, \pi, \text{imm})$ similarly.

Finally we introduce the concept of delayed-response strategies, which are used throughout the remainder of the paper to prove our folk theorems. We call σ a *delayed-response strategy profile* in the repeated game if there exists some K such that the repeated game can be divided into K “threads”, with the ℓ -th thread consisting of periods $\ell, K + \ell, 2K + \ell, \dots$, so that at any period t players condition their strategies only on messages and signal information generated within the thread containing period t .

3. Bounded Lags

We first present a simple analysis of a repeated game with observation lags in which the lag is certain to be no more than some finite bound.

Assumption 3.1. There exists some $K < \infty$ such that $\Pr(\max_i L_i \leq K) = 1$.

With this assumption, it is common knowledge that all players will have seen the signal generated in period t by period $t + K$. This restriction allows us to show that every equilibrium payoff attainable for sufficiently large discount factors in the repeated game *without observation lags* with any private monitoring structure π can also be attained in the associated repeated game *with observation lags* for sufficiently patient players. We show this using delayed-response strategies. Note that the following result does not impose any restrictions on π ; we use such conditions for our folk theorems later but they are not needed here.

Theorem 3.2. *Suppose Assumption 3.1 holds. Furthermore suppose that $v \in E(\delta, \pi)$ for all $\delta \in (\underline{\delta}, 1)$ where $0 < \underline{\delta} < 1$. Then there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta \in (\delta^*, 1)$.*

Proof. We divide the periods of the repeated game into $K + 1$ threads, with the ℓ -th thread consisting of periods ℓ , $(K + 1) + \ell$, $2(K + 1) + \ell$, \dots . Now, we suppose that $v \in E(\delta, \pi)$ is generated by the strategy profile σ in the game without lags.

As the information lag has an upper bound of K , the signals generated in periods ℓ , $(K + 1) + \ell$, \dots , $(j - 1)(K + 1) + \ell$ are observed by all players by period $j(K + 1) + \ell$. Thus, we may define a delayed-response strategy profile σ^K by specifying that in period $t = j(K + 1) + \ell$ ($0 \leq \ell \leq K$), players play according to $\sigma(h^{t,\ell}, h^{t,(K+1)+\ell}, \dots, h^{t,(j-1)(K+1)+\ell})$.

It is clear that the delayed-response strategy profile σ^K generates a payoff profile of v . Moreover, it is an equilibrium for discount factor δ^{K+1} . Thus, taking $\delta^* = \delta^{K+1}$ gives the result. \square

The proof of Theorem 3.2 relies heavily on Assumption 3.1. For example, if the support of λ were concentrated on $(0, \dots, 0)$, $(1, \dots, 1)$, \dots , (K, \dots, K) , and (∞, \dots, ∞) , then the proof above would not work, since each of the threads that it constructs would be a repeated game with a private monitoring structure $\tilde{\pi}$ that is different from π . More problematically, if $\lambda_i(k) > 0$ for all i and $k \in \mathbb{N}$, so that all lag lengths have positive probability for all players, then no matter how far apart the threads are spaced, there is always a positive probability that a realized lag will be longer than this chosen spacing, and the threads considered in the proof above cannot be identified with a private monitoring game at all. In the next two sections, we study and demonstrate how these issues can be resolved when additional assumptions are placed on the monitoring structure π . Therefore for the remainder of the paper, we dispense with Assumption 3.1 and allow λ to be an arbitrary probability distribution on $(\mathbb{N} \cup \{\infty\})^n$.

4. Lagged Almost-Perfect Monitoring with Two Players

In this section, we extend an approach of HO2006 to obtain a folk theorem for two-player games with lagged almost-perfect monitoring. We focus on the two-player case since the

techniques of HO2006 extend naturally to this setting.¹¹

4.1. Model

We restrict the general monitoring structure introduced above. First, we assume that there are only two players. We assume the monitoring structure to be that of *lagged ε -perfect monitoring*: We allow a general lag structure here, but restrict the private signal space of each player i to be $\Omega_i = A_j$ and furthermore assume that π is ε -perfect in the sense of HO2006. We assume that the private signal space of $\Omega_i = A_j$ so that we may extend the techniques of HO2006.¹²

Definition 4.1. A private monitoring structure π is *ε -perfect* if for every action profile $a \in A$, $\pi(a_2, a_1 \mid a_1, a_2) > 1 - \varepsilon$.

4.1.1. The Folk Theorem

We now prove the following folk theorem.

Theorem 4.2. *Suppose that $v \in \text{int}(V^*)$. Then there exists some $\bar{\varepsilon} \in (0, 1)$ such that for all lag distributions λ for which $\lambda_i(\infty) < \bar{\varepsilon}$ (for $i = 1, 2$), there exists some $\bar{\delta}$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \bar{\delta}$, and all private monitoring structures π that are $\bar{\varepsilon}$ -perfect.*¹³

To prove Theorem 4.2, we first analyze an auxiliary repeated game with rare observation lags, in which the probability of instantaneous observation of the private signal is very close to 1. We show that the HO2006 approach to repeated games with almost-perfect monitoring can be extended to lagged repeated games with almost-perfect monitoring, so long as positive lags are sufficiently rare, and use this to obtain a folk theorem in the auxiliary game. We then convert the associated auxiliary-game strategies to delayed-response strategies by multithreading the game with lags. A positive lag in a particular thread corresponds to a lag that exceeds the number of threads, so by taking the delay long enough we can shrink the probability of a positive lag close to 0. We thus obtain a folk theorem in the game with stochastic lags.

¹¹We do not know whether our folk theorem extends to games with n players; we discuss related issues in Section 6.

¹²Note that the work of HO2006 does contain a section that extends the analysis to more general private signal spaces where $\Omega_i \neq A_i$. However, as Yuichi Yamamoto pointed out to us, that argument contains an error so we cannot use it here.

¹³We thank Yuichi Yamamoto for pointing out a problem with our earlier proof of this result and then suggesting the approach we use now.

4.2. Auxiliary Repeated Game with Rare Observation Lags

This subsection establishes the following folk theorem for the game where $\lambda_i(0)$ is close to 1 for all players i :

Theorem 4.3. *Let $v \in \text{int}(V^*)$. Then there exist $\bar{\varepsilon}, \underline{\delta} \in (0, 1)$ such that if $\lambda_i(0) > 1 - \bar{\varepsilon}$ (for $i = 1, 2$), π is $\bar{\varepsilon}$ -perfect, and $\delta > \underline{\delta}$, then $v \in E(\delta, \pi, \lambda)$.*

Our proof of this theorem adapts a technique of HO2006 to the environment with small observation lags. The HO2006 construction for the case of almost-perfect monitoring uses the same strategies as in the perfect-monitoring construction at histories that are on the equilibrium path of that equilibrium (the “regular histories”), and then uses standard full-rank arguments to show there are continuation payoffs (at the end of the review phase) that preserve the belief-freeness property at the “erroneous” histories—those which are off the path of play under perfect monitoring. When the monitoring is close to perfect, the additional variation introduced in these continuation payoffs converges to zero. We use a similar argument, grouping histories together by treating a delayed observation as one that never arrives, and classifying as “erroneous” any history in which *some* player observes the opponent’s action with a strictly positive lag. We then construct continuation payoffs associated to these histories by applying full-rank arguments to the “immediate observation structure” defined below.

Note first that because the information lag is not bounded, it is possible that information about some past event arrives very late in the repeated game. Such possibilities cannot be ignored—even though they happen with very low probability—since they may potentially affect a player’s beliefs about his opponent’s continuation play. Our extension deals with this problem by constructing equilibria that are belief-free every T periods for the repeated game with the probability of lagged observation sufficiently small.¹⁴ This means that only information about the past T periods is relevant for computing best replies. Thus, we can ensure that effects on beliefs due to observation lags lasting more than T periods are unimportant.

Note next that lags of length less than T do affect players’ on-path beliefs, so the HO2006 arguments do not directly apply. We extend them to lags with $\lambda(0)$ close to 1 by adding the histories where observations arrive with a positive lag to the set of “erroneous” histories.

¹⁴A strategy is *belief-free at time t* if the continuation strategy at time t , $s_i | h_i^{t-1}$, is a best response against $s_{-i} | h_{-i}^{t-1}$ for all pairs of histories $(h_i^{t-1}, h_{-i}^{t-1})$. (Here, as we define formally below, “|” indicates the restriction of a strategy to a given history set.)

4.2.1. Preliminaries

We let H_i^t be the set of t -period histories in the repeated game *with observation lags*, with elements denoted in the form

$$h_i^t = (a_i^0, a_i^1, \dots, a_i^{t-1}, h_i^{1,o}, h_i^{2,o}, \dots, h_i^{t,o}).$$

Here, $h_i^{t,o}$ denotes all of the *new information* about the past play of player $-i$ that player i receives in period t . Furthermore denote by S_i^T the set of strategies in the T -times repeated game with information lags. Let \tilde{H}_i^t be the set of t -period histories in the repeated game *without observation lags and with perfect monitoring* with a typical element of \tilde{H}_i^t denoted by \tilde{h}_i^t . Also denote the set of strategies in the T -times repeated game with perfect monitoring and no observation lags by \tilde{S}_i^T .

Now we partition the set of private histories in the T -times-repeated stage game into H_i^R and H_i^E , the *regular* and *erroneous* histories. To do this we first define restricted strategy sets $\tilde{\mathcal{S}}_i$ and $\tilde{\mathcal{S}}_i^\rho$ for $i = 1, 2$ in the T -times repeated game with *perfect monitoring*. Partition the set A_i into two subsets, denoted \mathbf{G} and \mathbf{B} . We call an instance of the T -times repeated game with perfect monitoring a *block*, and say that a player i *sends message* $\mathbf{M} \in \{\mathbf{G}, \mathbf{B}\}$ if he picks an action in \mathbf{M} in the first period of a block. As in HO2006, we fix a payoff vector v to be achieved in equilibrium and pick four action profiles $a^{\mathbf{X}, \mathbf{Y}}$ for $(\mathbf{X}, \mathbf{Y}) \in \{\mathbf{G}, \mathbf{B}\}^2$ with $w_i^{\mathbf{X}, \mathbf{Y}} = g_i(a^{\mathbf{X}, \mathbf{Y}})$, $\mathbf{X}, \mathbf{Y} \in \{\mathbf{G}, \mathbf{B}\}$, where $w_i^{\mathbf{G}, \mathbf{G}} > v_i > w_i^{\mathbf{B}, \mathbf{B}}$, and

$$w_1^{\mathbf{G}, \mathbf{B}} > v_1 > w_1^{\mathbf{B}, \mathbf{G}}, \quad w_2^{\mathbf{B}, \mathbf{G}} > v_2 > w_2^{\mathbf{G}, \mathbf{B}}.^{15}$$

Choose $\underline{v}_i < \bar{v}_i$ with $v_i^* < \underline{v}_i < v_i < \bar{v}_i$ —where v_i^* is player i 's minmax payoff—such that

$$[\underline{v}_1, \bar{v}_1] \times [\underline{v}_2, \bar{v}_2] \subset \text{int}(\text{co}\{w^{\mathbf{G}, \mathbf{G}}, w^{\mathbf{B}, \mathbf{B}}, w^{\mathbf{B}, \mathbf{G}}, w^{\mathbf{G}, \mathbf{B}}\}).$$

We let \tilde{S}_i^T be the set of *block strategies* for player i , i.e. the set of strategies for the T -period perfect monitoring repeated game. We let $\tilde{\mathcal{S}}_i$ be the set of strategies $\tilde{s}_i \in \tilde{S}_i^T$ such that $\tilde{s}_i[\tilde{h}_i^t] = a_i^{\mathbf{M}_2, \mathbf{M}_1}$ for all $\tilde{h}_i^t = (a, (a_i^{\mathbf{M}_2, \mathbf{M}_1}, a_{-i}^{\mathbf{M}_2, \mathbf{M}_1}), \dots, (a_i^{\mathbf{M}_2, \mathbf{M}_1}, a_{-i}^{\mathbf{M}_2, \mathbf{M}_1}))$ with $a \in \{\mathbf{M}_i\} \times \mathbf{G}$ ($t \geq 1$). We then let

$$\begin{aligned} \tilde{\mathcal{A}}_i(\tilde{h}_i^t) &\equiv \{a_i \in A_i : \exists \tilde{s}_i \in \tilde{\mathcal{S}}_i \text{ such that } \tilde{s}_i[\tilde{h}_i^t](a_i) > 0\}, \\ \tilde{\mathcal{S}}_i^\rho &\equiv \{\tilde{s}_i \in \tilde{\mathcal{S}}_i : \tilde{s}_i[\tilde{h}_i^t](a_i) > \rho \text{ for all } \tilde{h}_i^t \text{ and } a_i \in \tilde{\mathcal{A}}_i(\tilde{h}_i^t)\}.^{16} \end{aligned}$$

¹⁵These action profiles can be assumed to be pure, either with the use of a public randomization device or by picking a quadruple of sequences of action profiles such that the average payoff of each of the sequences satisfy the above properties.

Define $\tilde{H}_i^{R,t}$ to be the set of period- t private histories of player i in the T -times-repeated game with perfect monitoring that are on the equilibrium path for some (and therefore, every) strategy profile in $\tilde{\mathcal{S}}_1^\rho \times \tilde{\mathcal{S}}_2^\rho$. Then we identify each $\tilde{h}_i^t \in \tilde{H}_i^t$ with the unique element of $h_i^t \in H_i^t$ such that h_i^t and \tilde{h}_i^t report exactly the same observations about the play of player $-i$ at all times and h_i^t contains no observations with a positive lag (all observations are observed instantaneously). Define $H_i^{R,t}$ as the image of $\tilde{H}_i^{R,t}$ under this identification, and denote this identification by $\tilde{h}_i^t \simeq h_i^t$ for $\tilde{h}_i^t \in \tilde{H}_i^t$ and $h_i^t \in H_i^t$. Also define the set of erroneous histories to be $H_i^{E,t} = H_i^t \setminus H_i^{R,t}$. This means that $H_i^{E,t}$ includes any private histories in which player i did not immediately observe the period- t' play of player $-i$ for some $t' < t$.

Additionally define the set of strategies $\mathcal{S}_i \subseteq S_i^T$ in the repeated game *with observation lags* as the set

$$\mathcal{S}_i \equiv \{s_i \in S_i^T : \exists \tilde{s}_i \in \tilde{\mathcal{S}}_i \text{ such that } \tilde{s}_i[\tilde{h}_i^t] = s_i[h_i^t] \text{ for all } \tilde{h}_i^t \in \tilde{H}_i^t \text{ where } \tilde{h}_i^t \simeq h_i^t\}.$$

Additionally, define

$$\begin{aligned} \mathcal{A}_i(h_i^t) &\equiv \{a_i \in A_i : \exists s_i \in \mathcal{S}_i \text{ such that } s_i[h_i^t](a_i) > 0\}, \\ \mathcal{S}_i^\rho &\equiv \{s_i \in \mathcal{S}_i : s_i[h_i^t](a_i) > \rho \text{ for all } h_i^t \in H_i^t \text{ and } a_i \in \mathcal{A}_i(h_i^t)\}.^{17} \end{aligned}$$

Finally we define strategies $s_i^B, s_i^G \in \mathcal{S}_i^\rho$ by mapping the strategies \tilde{s}_i^B and \tilde{s}_i^G defined by HO2006 in a perfect monitoring repeated game to strategies in our environment with private monitoring and observation lags in a natural way. (The details of this definition are included in Appendix Appendix A.)

4.2.2. Proof of Theorem 4.3

The proof of Theorem 4.3 follows from three key lemmata; once these lemmata have been established, the remainder of the proof follows exactly as in HO2006. The first lemma adapts Lemma 1 of HO2006 to our setting of repeated games with information lags. Because the proof requires some nontrivial modifications, we include the argument here. As we show in the Appendix, analogous modifications can be made to the proofs of Lemmata 2 and 3 of HO2006; Theorem 4.3 then follows.

We write $s_i | H_i$ for the restriction of strategy s_i to history set H_i . We let \tilde{U}_i^T be the payoff of player i in the T -times repeated game *with perfect monitoring and no observation lags*.

¹⁶As in the HO2006 constructions, given any history h_i^t , the set $\tilde{\mathcal{S}}_i$ imposes either no restrictions on $s_i[h_i^t]$ or restricts $s_i[h_i^t]$ to a single action. In particular any strategy $\tilde{s}_i \in \tilde{\mathcal{S}}_i^\rho$ puts positive weight on all actions after any erroneous history.

¹⁷Just as in the case of $\tilde{\mathcal{S}}_i^\rho$ in Footnote 16, $s_i^\rho \in \mathcal{S}_i^\rho$ puts positive weight on all actions after any erroneous history.

Analogously define U_i^T to be the ex-ante payoff of player i in the T -times repeated game *with private monitoring structure π and observation lags*. We consider a version of the T -times repeated game (with observation lags) which is augmented with a *transfer* $\xi_{-i} : H_i^T \rightarrow \mathbb{R}$ at the end of the T -th period. In this *auxiliary scenario*, the payoff of i under strategy profile s is taken to be

$$U_i^A(s, \xi_i) \equiv U_i^T(s) + (1 - \delta)\delta^T \mathbb{E}(\xi_i | s).$$

The set of best responses of player i in the auxiliary scenario with opponent's strategy s_{-i} and own transfer ξ_i is denoted $B_i(s_{-i}, \xi_i)$.

With these notations, we have the following lemma that defines the transfer ξ_i^B received after “bad” messages.

Lemma 4.4. *For every strategy profile $\bar{s} | H^E$, there exists $\bar{\varepsilon} > 0$ such that whenever $\lambda_i(0) > 1 - \bar{\varepsilon}$ for $i = 1, 2$ and π is $\bar{\varepsilon}$ -perfect, then there exists a nonnegative transfer $\xi_i^B : H_{-i}^T \rightarrow \mathbb{R}_+$ such that $S_i^T = B_i(\bar{s}_{-i}^B, \xi_i^B)$ where $\bar{s}_{-i}^B | H_{-i}^R = s_{-i}^B | H_{-i}^R$ and $\bar{s}_{-i}^B | H_{-i}^E = \bar{s}_{-i} | H_{-i}^E$, and for every $s_i \in B_i(\bar{s}_{-i}^B, \xi_i^B)$,*

$$\lim_{\varepsilon \rightarrow 0} U_i^A(s_i, \bar{s}_{-i}^B, \xi_i^B) = \max_{\tilde{s}_i \in \tilde{S}_i^T} \tilde{U}_i^T(\tilde{s}_i, \bar{s}_{-i}^B).$$

This generalizes Lemma 1 of HO2006 to a repeated game in which information does not arrive instantaneously. To do so, we must contend with the fact that H_{-i}^T contains many more histories than in their private monitoring environment because information may arrive with lag, so that it is not immediately clear how to construct the ξ_i^B . We handle this issue by partitioning the set of histories into sets which are *past-observation equivalent*, in the sense that for any two time- t histories h^t and $h^{t'}$ in the same set, the $(t - 1)$ -period truncations of h^t and $h^{t'}$ are equal. We then identify each of the elements of this partition with a particular history in a private monitoring repeated game with the *immediate monitoring structure* μ induced by λ and π defined over the space of signal profiles $(\Omega_1 \cup \{\infty\}) \times (\Omega_2 \cup \{\infty\})$:

$$\mu(\hat{\omega}_1, \hat{\omega}_2 | a_1, a_2) = \begin{cases} \sum_{\tau_1=1}^{\infty} \sum_{\tau_2=1}^{\infty} \lambda(\tau_1, \tau_2) & \hat{\omega}_1, \hat{\omega}_2 = \infty \\ (\sum_{\tau=1}^{\infty} \lambda(\tau, 0)) \left(\sum_{\omega'_1 \in A_2} \pi(\omega'_1, \hat{\omega}_2 | a_1, a_2) \right) & \hat{\omega}_1 = \infty \text{ and } \hat{\omega}_2 \neq \infty \\ (\sum_{\tau=1}^{\infty} \lambda(0, \tau)) \left(\sum_{\omega'_2 \in A_1} \pi(\hat{\omega}_1, \omega'_2 | a_1, a_2) \right) & \hat{\omega}_1 \neq \infty \text{ and } \hat{\omega}_2 = \infty \\ \lambda(0, 0) \pi(\hat{\omega}_1, \hat{\omega}_2 | a_1, a_2) & \hat{\omega}_1, \hat{\omega}_2 \neq \infty. \end{cases}$$

This monitoring structure represents the information about the period- t action of the opponent that is available in period- $t + 1$, treating positive lags as the null signal. With this identification, we can extend the arguments of HO2006 to arrive at our desired conclusion.

We construct the transfers ξ_i^G received after a “good” message in the repeated game with rare lags in a fashion closely similar to those specified in Lemma 4.4.

Lemma 4.5. *For every strategy profile $\bar{s} \mid H^E$, there exists $\bar{\varepsilon} > 0$ such that, whenever $\Pr(L > 0) < \bar{\varepsilon}$ and π is $\bar{\varepsilon}$ -perfect, there exists a nonpositive transfer $\xi_i^G : H_{-i}^T \rightarrow \mathbb{R}_-$ such that*

$$\{s_i \in s_i^T : s_i \mid H_i^R = \hat{s}_i \mid H_i^R \text{ for some } \hat{s}_i \in \mathcal{S}_i \text{ and } s_i \mid H_i^E = \bar{s}_i \mid H_i^E\} \subseteq B_i(\bar{s}_{-i}^G, \xi_i^G \mid \bar{s}_i)$$

where $\bar{s}_{-i}^G \mid H_{-i}^R = s_{-i}^G \mid H_{-i}^R$ and $\bar{s}_{-i}^G \mid H_{-i}^E = \bar{s}_{-i} \mid H_{-i}^E$. Furthermore $\xi_i^G : H_{-i}^T \rightarrow \mathbb{R}_-$ can be chosen so that, for every $s_i \in B_i(\bar{s}_{-i}^G, \xi_i^G \mid \bar{s}_i)$, we have

$$\lim_{\varepsilon \rightarrow 0} U_i^A(s_i, \bar{s}_{-i}^G, \xi_i^G) = \min_{\tilde{s}_i \in \tilde{\mathcal{S}}_i} \tilde{U}_i^T(\tilde{s}_i, \bar{s}_{-i}^G),$$

ξ_i^G depends continuously on \bar{s} , and ξ_i^G is bounded away from $-\infty$.

We relegate the proof to the Appendix. The remainder of the proof of Theorem 4.3 follows along the same lines as in HO2006, defining $\bar{s}_{-i}^B \mid H_{-i}^E$ and $\bar{s}_{-i}^G \mid H_{-i}^E$, ξ_i^G and ξ_i^B as the fixed point of the relevant correspondence. The construction works because of Lemma 4.5 and the fact that play at periods $T, 2T, \dots$, is belief free (by Lemma 4.4). Thus for example if player i receives information about the play of player $-i$ in period $T - m$ at some time $T + l$, this does not have any effect on his best response calculation since player i 's strategy only depends on the history of information about the events occurring after period T .

4.3. The Repeated Game with Frequent Observation Lags

In the previous section, we required that the probability of a positive lag be small. In this section, we show that even if the lags are frequent and possibly very long, the folk theorem still obtains when $\lambda_i(\infty)$ is sufficiently small for $i = 1, 2$.

The following lemma employs a technique similar to that used in the proof of Theorem 3.2, using delayed-response strategies to relate the equilibrium payoffs in the game with rare observation lags to those with possibly long lags.

Lemma 4.6. *Suppose $v \in E(\delta, \pi, \hat{\lambda})$ for all lag distributions $\hat{\lambda}$ such that $\hat{\lambda}(0) > 1 - \bar{\varepsilon}$ and all $\delta \in (\underline{\delta}, 1)$. Then for all lag distributions λ such that $\lambda_i(\infty) < \bar{\varepsilon}/2$ for $i = 1, 2$, there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$.*

Proof. Choose $K \in \mathbb{N}$ such that $(1 - \Lambda_i(K - 1)) < \bar{\varepsilon}$ for $i = 1, 2$. and set $\delta^* = \underline{\delta}^{\frac{1}{K}}$. Then there exists a positive integer $K^* \geq K + 1$, such that $\delta^{K^*} \in (\underline{\delta}, 1)$ for every $\delta > \delta^*$.

Now divide the repeated game $G(\delta, \pi, \lambda)$ into K^* distinct repeated game ‘‘threads,’’ the ℓ -th ($1 \leq \ell \leq K^*$) of which is played in periods $\ell, K^* + \ell, 2K^* + \ell, \dots$. Because $K^* \geq K + 1$, each of these separate repeated games is equivalent to $G(\delta^{K^*}, \pi, \hat{\lambda})$ for some $\hat{\lambda}$ such that $\hat{\lambda}(0) > 1 - \bar{\varepsilon}$, and each repeated game thread can be treated independently, as players never condition their play in the ℓ -th thread on information received about play in the ℓ' -th

repeated games ($\ell' \neq \ell$). Because $v \in E(\delta^{K^*}, \pi, \hat{\lambda})$, it is then clear that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$. \square

Theorem 4.2 follows directly from Lemma 4.6 and Theorem 4.3.

Remark. Ellison (1994) used threading primarily to lower the discount factor. By contrast, we use threading to ensure that the probability of lags being longer than the thread length remains low, so that players (with high probability) observe signals of play within a thread before choosing new actions within that thread. Thus the number of threads required in our proof is independent of the discount factor, while Ellison (1994) required the number of threads to become arbitrarily large as the discount factor approaches 1.

5. Lagged Public Monitoring

5.1. Model

In this section, we consider an n -player repeated game in which the monitoring structure of the repeated game is *public with stochastic lags*: There is a set of public signals, denoted Y , and we set $\Omega_i = Y$ for all players $i \in I$. Furthermore we assume that π is supported on the set

$$\{(y_1, \dots, y_n) \in Y^n : y_1 = y_2 = \dots = y_n\}.$$

That is, the monitoring structure of the underlying repeated game without lags is public. With a slight abuse of notation, we then write $\pi(y | a)$ as shorthand for $\pi((y, \dots, y) | a)$.

We place a mild restriction on the support of the monitoring structure π .

Assumption 5.1. For every pure action profile $a \in A$, there exist $y, y' \in Y$ with $y \neq y'$ such that $\pi(y | a), \pi(y' | a) > 0$.

Note that the argument used for the case of lagged perfect monitoring does not work here because the analogous auxiliary game does not have almost-perfect monitoring. Moreover, an extension of the Hörner and Olszewski (2009) construction to repeated games with rare observation lags is not possible, because that construction assumes that each player assigns high probability to the event that all players observe the same signal as in the setting of Mailath and Morris (2002); this condition is possibly violated when a player observes the low-probability “null” signal.¹⁸ The possibility of receiving an uninformative signal also prevents the application of the folk theorem of Sugaya (2011), because the necessary full rank condition

¹⁸We believe that threading combined with Hörner and Olszewski (2009) yields a folk theorem when lags are sufficiently positively correlated because the auxiliary repeated game corresponding to a thread can be treated as an almost-public monitoring game with the possibility of an uninformative null signal. The techniques we develop in this section are more novel.

fails. Thus, instead of invoking or adapting existing results for general private monitoring games, we allow for the possibility of communication that is perfectly and publicly observed at the end of every period, i.e. $M_i \neq \emptyset$. We assume that, unlike signals, messages are observed without delay. In the context of our “joint coauthorship” example of Section 2, the authors can quickly reach each other by phone or email after the reports arrive. We show that as long as $|M_i| \geq |Y| + 1$ for all i , a folk theorem can be established.

5.2. Structure of the Observation Lags

In this section, we allow for the possibility that lags may be correlated (across agents). Define:¹⁹

$$\Lambda(T) \equiv \Pr \left(\min_i \{L_i\} \leq T \right) \quad \text{and set} \quad \gamma_i \equiv \lim_{T \rightarrow \infty} \left[\frac{1 - \Lambda(T)}{1 - \Lambda_i(T)} \right].$$

The quantity γ_i represents the limiting conditional probability that player i assigns to the event that players $j \neq i$ have not received signals about period- t play within T periods, when he himself has also not received any signal about period- t play within T periods. Note that if $\Pr(\{L : L_i = \infty\}) > 0$, then $\gamma_i = \frac{\lambda((\infty, \dots, \infty))}{\lambda_i(\infty)}$.

For our results in the section, we assume that $\lambda_i(\infty)$ and γ_i are both small for all i . It is easy to see what kind of lags satisfy the first condition. The second condition is a bit more subtle and so we illustrate it through some concrete examples in the Appendix.

5.3. The Folk Theorem

We begin our analysis with the simple observation that the repeated play of α^* is an equilibrium of the game with observation lags.²⁰ We use this fact along with techniques from Abreu et al. (1990) and FLM to construct equilibria that generate any payoff profile $v \in \text{int}(V_{\alpha^*})$.

To use the techniques of FLM, we need to impose some additional assumptions on the public monitoring structure π . Recall the following definition from FLM.

Definition 5.2. Let π be a public monitoring structure. Then a mixed action profile α has pairwise full rank for a pair $i, j \in I$ if the $((|A_i| + |A_j|) \times |Y|)$ matrix

$$\begin{pmatrix} (\pi(\cdot \mid a_i, \alpha_{-i}))_{a_i \in A_i} \\ (\pi(\cdot \mid a_j, \alpha_{-j}))_{a_j \in A_j} \end{pmatrix}$$

¹⁹Although we assume that $\Lambda_i(T) < 1$ for all $T \in \mathbb{N}$, our results extend to the case in which there exists some player i and some T^* such that $\Lambda_i(T^*) = 1$. In that case, we can simply take the number of threads to be larger than T^* , so that player i 's signal structure in the auxiliary game need not contain ∞ as one of its elements.

²⁰Note that for such play, the communication strategies are irrelevant, so we need not specify them.

has rank $|A_i| + |A_j| - 1$.

We will maintain the following restriction on π throughout the rest of this section.

Assumption 5.3. For all pairs i, j , there exists a profile α that has pairwise full rank for that pair.

We can now state our folk theorem for repeated games with public monitoring and stochastic lags with communication.

Theorem 5.4. *Let $v \in \text{int}(V_{\alpha^*})$ and suppose that π satisfies Assumptions 5.1 and 5.3. Furthermore suppose that $|M_i| \geq |Y| + 1$ for all i . Then there exist some $\varepsilon^* \in (0, 1)$ such that for every λ such that $\gamma_i < \varepsilon^*$ and $\lambda_i(\infty) < \varepsilon^*$ for all i , there exists $\delta^* \in (0, 1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$.*

As a preview of our proof, it is important that both γ_i and $\lambda_i(\infty)$ are small for all i . The need of the latter condition should be intuitive. For example, in a two-player game, if it is likely that one player never observes any information, the other player would have an incentive to play myopically. As we will see, the former condition is important for establishing truthful communication of signals. The remainder of the section proves Theorem 5.4.

5.4. Private Monitoring Game with Communication

5.4.1. Incentives for Truthful Communication

We first consider a private monitoring game with communication (in every period) and *no observation lags* for which each player's message space is $M_i = \tilde{Y} \equiv Y \cup \{\infty\}$. The results of this section are of stand-alone interest: the case where players might sometimes not see the signal seems plausible and it leads to a form of private monitoring that does not appear to be covered by past results. Let us first define some notation. For a vector $\tilde{y} \in \tilde{Y}^n$, define $\mathcal{I}(\tilde{y}) = \{i : \tilde{y}_i \neq \infty\}$ and $|\tilde{y}| \equiv |\mathcal{I}(\tilde{y})|$. Define the following set

$$\mathcal{Y} \equiv \{(\tilde{y}_1, \dots, \tilde{y}_n) \in \tilde{Y}^n : |(\tilde{y}_1, \dots, \tilde{y}_n)| > 0 \text{ and } \tilde{y}_j = \tilde{y}_k \forall j, k \text{ such that } \tilde{y}_j, \tilde{y}_k \neq \infty\}.$$

The monitoring structure is then supported on the set $\mathcal{Y} \cup \{(\infty, \dots, \infty)\}$. For any $\tilde{y} \in \mathcal{Y}$, we define $\vec{y} \in Y$ to be the $y \in Y$ such that $\tilde{y}_j = y$ for all j such that $\tilde{y}_j \neq \infty$.

Now consider a private monitoring structure π^{pr} that is supported on the set $\mathcal{Y} \cup \{(\infty, \dots, \infty)\}$ with the following additional features.

Assumption 5.5. $\pi^{\text{pr}}((\infty, \dots, \infty) | a)$ is constant across all $a \in A$.

Assumption 5.6. $\sum_{\vec{y}=y} \pi^{\text{pr}}(\tilde{y} | a) = (1 - \pi^{\text{pr}}((\infty, \dots, \infty) | a))\pi(y | a)$ for all $y \in Y$.

The reasons for these restrictions become clear when we relate this game to the repeated game with observation lags. We say that this monitoring structure is ε -close to π if π^{pr} is such that

$$\pi^{\text{pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid a, \tilde{y}_i) < \varepsilon \quad (1)$$

for all $a \in A$ and all $\tilde{y}_i \in A \cup \{\infty\}$. Note that this definition of ε -closeness to a public monitoring structure is quite different from the one used by Hörner and Olszewski (2009). The key difference is in the conditional probability $\pi^{\text{pr}}((\infty, \dots, \infty) \mid a, \tilde{y}_i = \infty)$. Hörner and Olszewski (2009) assumed this conditional probability to be close to 1. Here, we assume that it is very small.

We denote by $G^{\text{pr}}(\delta, \pi^{\text{pr}})$ the private monitoring game with discount factor δ and private monitoring structure π^{pr} (and communication) and let $E^{\text{pr}}(\delta, \pi^{\text{pr}})$ be the set of sequential equilibrium payoffs of $G^{\text{pr}}(\delta, \pi^{\text{pr}})$. We now show the following.

Theorem 5.7. *Let $v \in \text{int}(V_{a^*})$. Then there exist $\underline{\delta}, \bar{\delta} \in (0, 1)$ with $\underline{\delta} < \bar{\delta}$ and $\bar{\varepsilon} \in (0, 1)$ such that $v \in E^{\text{pr}}(\delta, \pi^{\text{pr}})$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all private monitoring structures π^{pr} that are $\bar{\varepsilon}$ -close to π .*

To prove this theorem we construct strategies that generate a payoff profile of v , and are public perfect in the sense of Kandori and Matsushima (1998): strategies in the non-communication stages of the game depend only on the sequence of message profiles reported in the history. These strategies use a form of grim-trigger reversion to static Nash equilibrium when the messages disagree, in order to provide incentives for truthful reporting. We prove the theorem in two parts. We first prove a lemma demonstrating that truth-telling is *incentive compatible* (i.e. that each player i should report message $m_i = y$ upon seeing signal $y \in Y$) when ε is sufficiently small given strategies with this grim-trigger property.

Lemma 5.8. *Let W be a convex, compact set that is a subset of $\text{int}(V_{a^*})$. Consider a collection of public perfect strategy profiles $\{\sigma^{\delta, \pi^{\text{pr}}}\}$, indexed by δ and π^{pr} , for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all private monitoring structures π^{pr} that are $\bar{\varepsilon}$ -close to π and have the following properties.*

1. *In period t , each player i (truthfully) communicates the signals $\tilde{y}_i^t \in \tilde{Y} = M_i$ he observes in period t .*
2. *If there exists some t such that $m^t \notin \mathcal{Y}$, then all players i play α_i^* .*
3. *Strategies are such that $\sigma^{\delta, \pi^{\text{pr}}}(m^0, \dots, m^t) = \sigma^{\delta, \pi^{\text{pr}}}(\bar{m}^0, \dots, \bar{m}^t)$ whenever $\bar{m}^\tau = \bar{m}^\tau$ for all $\tau = 0, \dots, t$.*
4. *Expected continuation values are always contained in W for play of $\sigma^{\delta, \pi^{\text{pr}}}$ in the game $G^{\text{pr}}(\delta, \pi^{\text{pr}})$ whenever the message history contains only elements in the set \mathcal{Y} .*

Then there exists $\varepsilon^ \leq \bar{\varepsilon}$ such that for all private monitoring structures π^{pr} that are ε^* -public except at infinity and all $\delta \in [\underline{\delta}, \bar{\delta}]$, truthful communication is incentive compatible at*

any private history in $G^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ given continuation play determined by $\sigma^{\delta, \pi^{\text{Pr}}}$ and truthful communication by all other players.

Proof. We check that there are no profitable one-stage deviations in which a player misreports once and then follows the continuation strategy prescribed by $\sigma_i^{\delta, \pi^{\text{Pr}}}$. First note that if the player is at a history in which there exists some t at which $m^t \notin \mathcal{Y}$, then all players play α_i^* forever from that point on. Since then continuation play does not depend on the message being sent, all players are indifferent to the message that they send after such a history. Thus it is incentive compatible.

So it remains to analyze incentives for truth-telling after histories in which $m^t \in \mathcal{Y}$ for all t . Suppose first that player i sees the null signal. Then by reporting ∞ , player i obtains an expected payoff of

$$\sum_{\vec{y} \in \mathcal{Y}} \pi^{\text{Pr}}(\vec{y} \mid \alpha, \tilde{y}_i = \infty) w_i(\vec{y})$$

for some $\alpha \in \prod_{i=1}^n \Delta(A_i)$ and some expected continuation value function $w : Y \rightarrow W$.²¹

If instead player i reports $y' \in Y$, he obtains a payoff of

$$\pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty) w_i(y') + \sum_{\vec{y} = y'} \pi^{\text{Pr}}(\vec{y} \mid \alpha, \tilde{y}_i = \infty) w_i(y')$$

Thus, to show that truth-telling is incentive compatible after all histories in which a player observes the null signal, it suffices to show that there exists ε^* sufficiently small so that

$$\sum_{y \in Y \setminus \{y'\}} \pi^{\text{Pr}}(\{\vec{y} : \vec{y} = y\} \mid \alpha, \tilde{y}_i = \infty) w_i(y) > \pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty) w_i(y') \quad (2)$$

for all $y' \in Y$, all $\alpha \in \prod_{i=1}^n \Delta(A_i)$, all $w : Y \rightarrow W$, $i = 1, \dots, n$, and all π^{Pr} ε^* -close to π . Assumptions 5.1 and 5.6 imply:

$$\begin{aligned} M(\pi^{\text{Pr}}, y', \alpha) &\equiv \sum_{y \in Y \setminus \{y'\}} \pi^{\text{Pr}}(\{\vec{y} : \vec{y} = y\} \mid \alpha, \tilde{y}_i = \infty) \\ &= (1 - \pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha)) \sum_{y' \in Y \setminus \{y'\}} \pi(y \mid \alpha) \\ &= (1 - \pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha))(1 - \pi(y' \mid \alpha)) \\ &> 0. \end{aligned}$$

²¹Note that in any sequential equilibrium if a player observes signal ∞ , he still believes that all other players played according to their prescribed actions, i.e. that there have been no “unexpected” events.

Note that for a fixed map, $w : Y \rightarrow W$ and $\alpha \in \prod_{i=1}^n \Delta(A_i)$, (2) holds for all $i = 1, \dots, n$ and all $y' \in Y$ if and only if

$$\sum_{y \in Y \setminus \{y'\}} \frac{\pi^{\text{Pr}}(\{\tilde{y} : \vec{\tilde{y}} = y\} \mid \alpha, \tilde{y}_i = \infty)}{M(\pi^{\text{Pr}}, y', \alpha)} w_i(y) > \frac{\pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty)}{M(\pi^{\text{Pr}}, y', \alpha)} w_i(y'). \quad (3)$$

Now let π^{Pr} be ε -close to π . As $\varepsilon \rightarrow 0$, $M(\pi^{\text{Pr}}, y') \rightarrow 1 - \pi(y \mid \alpha)$ and thus because π^{Pr} is ε -close to π ,

$$\frac{\pi^{\text{Pr}}((\infty, \dots, \infty) \mid \alpha, \tilde{y}_i = \infty)}{M(\pi^{\text{Pr}}, y', \alpha)} < \frac{\varepsilon}{M(\pi^{\text{Pr}}, y', \alpha)} \rightarrow \frac{0}{1 - \pi(y' \mid \alpha)} = 0.$$

Note that for any value of $\varepsilon > 0$, because W is convex, the left hand side of inequality (3) is an element of W .

Therefore because W is compact and contained in the interior of V_{α^*} , there is some ε^* such that inequality (3) holds for all π^{Pr} ε^* -close to π . Moreover $\prod_{i=1}^n \Delta(A_i)$ and the set of all maps $w : Y \rightarrow W$ are both compact. Therefore such an ε^* can be taken uniformly across all α and all maps $w : Y \rightarrow W$. This shows that all players will report the null signal truthfully when π^{Pr} is ε^* -close to π .

Now suppose that player i observes $y \in Y$. By reporting truthfully, player i obtains a payoff of $w_i(y)$ for some map $w : Y \rightarrow W$. However by reporting $y' \in Y$ with $y' \neq y$, player i obtains a payoff of

$$\pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y) w_i(y') \quad (4)$$

while reporting ∞ yields a payoff of

$$(1 - \pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y)) w_i(y). \quad (5)$$

Clearly $w_i(y)$ is at least the expression in (5) for any π^{Pr} since $w_i(y) \geq 0$. Furthermore we can take ε^* sufficiently small so that $w_i(y) > \max_{y' \neq y} \{\varepsilon w_i(y')\}$ for all $y \in Y$, all maps $w : Y \rightarrow W$, all $i = 1, \dots, n$, and all $\varepsilon < \varepsilon^*$. Then all players have an incentive to report truthfully upon observing an informative signal when π^{Pr} is ε^* -close to π since

$$w_i(y) > \max_{y' \neq y} \{\varepsilon w_i(y')\} \geq \max_{y' \neq y} \{\pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y) w_i(y')\}$$

and

$$w_i(y) \geq (1 - \pi^{\text{Pr}}(\{\tilde{y} : \tilde{y}_{-i} = (\infty, \dots, \infty)\} \mid \alpha, \tilde{y}_i = y)) w_i(y)$$

trivially. This concludes the proof. \square

Remark. The fact that (1) is small for all $\tilde{y}_i \in A_i \cup \{\infty\}$ is crucial. Otherwise, because a message profile of (∞, \dots, ∞) results in reversion to the static Nash equilibrium, player i upon observation of the null signal may have an incentive to deviate and report some signal $y \in Y$.

Remark. Because the set W in the lemma does not depend on delta, neither does ε^* . This is important for our folk theorem as we must establish a claim about all games with a private monitoring structure that is ε^* -close to π and all discount factors in an interval.

Remark. Players have no incentive to either communicate or respond to a signal that arrives late, since in equilibrium players do not respond to such communication. This is similar to the way in which we treat late signals in Section 4, where the belief-free property of the equilibrium construction allows us to show that players do not have an incentive to respond to late signals.

5.4.2. Non-Communication Stages

Lemma 5.8 provides a sufficient condition for truth-telling to be incentive compatible. We now show that given truthful communication by all players at all histories, we can construct a collection of strategies $\{\sigma^{\delta, \pi^{\text{pr}}}\}$ that satisfy the necessary properties of Lemma 5.8 for truthful communication and in which all players are also playing best-responses in the non-communication stages of the game.

To construct such strategies $\sigma^{\delta, \pi^{\text{pr}}}$, we first specify that players play α^* whenever in the history there exists some t such that $m^t \notin \mathcal{Y}$. Then it is trivial that playing α_i^* is a best response at such a history since opponents play α_{-i}^* forever. It remains to specify play after histories in which all messages in the history are elements of \mathcal{Y} . We do this by considering public strategies that only depend on the history of messages.

Given strategies that satisfy conditions 1, 2, and 3 of Lemma 5.8 we can simplify the analysis to that of an auxiliary public monitoring game defined in the following discussion. The auxiliary game is one of standard simultaneous moves in which public signals arise according to the conditional probability distribution π every period. We then modify this repeated game so that at the beginning of periods $1, 2, \dots$, the game ends with probability ε and each player receives flow payoffs of $0 = g_i(\alpha^*)$ thereafter. This corresponds exactly to the event in which all players report the null signal, triggering all players to play according to α^* forever.²²

²²Because all players report truthfully at all histories, message profiles $m \in \tilde{Y}^n \setminus \{\mathcal{Y} \cup (\infty, \dots, \infty)\}$ never occur on the equilibrium path. Thus the “grim phase” of playing α^* forever is only triggered in the event of message profile $m = (\infty, \dots, \infty)$; this happens with probability ε .

In the modified game, payoffs are given by

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t (1 - \varepsilon)^t g_i(a^t). \quad (6)$$

We denote this public monitoring game by $G^{\text{pu}}(\delta, \varepsilon)$ and let $E^{\text{pu}}(\delta, \varepsilon)$ be the set of sequential equilibrium payoffs of $G^{\text{pu}}(\delta, \varepsilon)$. Note that in this game the feasible payoff set is not constant in δ and ε , and in particular for any fixed $\varepsilon > 0$, as $\delta \rightarrow 1$, the feasible payoff set converges to $\{0\}$, just as the payoffs to grim trigger strategies converge to those of static Nash equilibrium as $\delta \rightarrow 1$ in a repeated game with imperfect public monitoring. However for any fixed δ , as $\varepsilon \rightarrow 0$, the feasible payoff set converges to V , the feasible payoff set of the original public monitoring game. Our analysis takes care in addressing this issue.

In order to extend the arguments of FLM to this modified repeated game, we first renormalize payoffs so that the feasible payoff set is indeed equal to V . We do this by multiplying the payoffs by a factor of $(1 - \delta(1 - \varepsilon))/(1 - \delta)$ to get payoff structure

$$(1 - \delta(1 - \varepsilon)) \sum_{t=0}^{\infty} \delta^t (1 - \varepsilon)^t g_i(a^t). \quad (7)$$

Now, our modified game corresponds to a repeated game with discount factor given by $\delta(1 - \varepsilon)$, hence all of the conclusions of FLM can be applied to this game, with the appropriate assumptions on the (original) public monitoring structure.

Before we proceed with the analysis of the game, recall the definition of self-generation (Abreu et al. (1990)).

Definition 5.9. For $W \subset \mathbb{R}^n$, define the sets $B(W, \delta, \varepsilon)$ and $\hat{B}(W, \delta, \varepsilon)$ as follows. Let $B(W, \delta, \varepsilon)$ be the set of $v \in \mathbb{R}^n$ such that there exists some mixed action profile α and a map $w : Y \rightarrow W$ such that

$$\begin{aligned} v &= (1 - \delta)g(\alpha) + \delta(1 - \varepsilon) \sum_{y \in Y} w(y)\pi(y|\alpha), \quad \text{and} \\ v_i &\geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta(1 - \varepsilon) \sum_{y \in Y} w_i(y)\pi(y | a_i, \alpha_{-i}), \end{aligned}$$

for all $a_i \in A_i$ and all i . Analogously define $\hat{B}(W, \delta, \varepsilon)$ to be the set of $v \in \mathbb{R}^n$ such that

there exists some mixed action profile α and a map $w : Y \rightarrow W$ such that

$$v = (1 - \delta(1 - \varepsilon))g(\alpha) + \delta(1 - \varepsilon) \sum_{y \in Y} w(y)\pi(y|\alpha), \quad \text{and}$$

$$v_i \geq (1 - \delta(1 - \varepsilon))g_i(a_i, \alpha_{-i}) + \delta(1 - \varepsilon) \sum_{y \in Y} w_i(y)\pi(y | a_i, \alpha_{-i})$$

for all i . We say that W is *self-generating in the repeated game with payoff structure (6) with discount factor δ and absorption probability ε* if $W \subseteq B(W, \delta, \varepsilon)$. Similarly, W is *self-generating in the repeated game with payoff structure (7) with discount factor δ and absorption probability ε* if $W \subseteq \hat{B}(W, \delta, \varepsilon)$.

Because the public monitoring game $G^{\text{pf}}(\delta, \varepsilon)$ has a slightly different structure from that of a standard public monitoring game, the consequences of self-generation are not immediate from past theorems, but the same ideas apply as shown in the next lemma.

Lemma 5.10. *Suppose W is compact and that $W \subseteq B(W, \delta, \varepsilon)$. Then $W \subseteq E^{\text{pu}}(\delta, \varepsilon)$.*

The proof of Lemma 5.10 is completely standard, so we omit it. FLM applied to the repeated game with discount factor $\delta(1 - \varepsilon)$ yields the following lemma.

Lemma 5.11. *Suppose that Assumption 5.3 holds. Let \hat{W} be a smooth, compact, convex set in the $\text{int}(V_{\alpha^*})$. Then there exists $\bar{\delta} \in (0, 1)$ and $\bar{\varepsilon} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ and all $\varepsilon < \bar{\varepsilon}$, $\hat{W} \subseteq \hat{B}(\hat{W}, \delta, \varepsilon)$, that is, \hat{W} is self-generating in the repeated game with payoff structure (7) with discount factor δ and absorption probability ε .*

Next, we translate the payoff set used in Lemma 5.11 back into payoffs without the renormalization. To do this, we define (for a set \hat{W}) a set W under the payoff normalization given by (6):

$$W = \frac{1 - \delta}{1 - \delta(1 - \varepsilon)} \hat{W}. \quad (8)$$

Of course for any fixed ε and a fixed set \hat{W} , as $\delta \rightarrow 1$, W shrinks (setwise) towards the point-set $\{0\}$. Thus for any choice of $v \in \text{int}(V_{\alpha^*})$, v will necessarily lie outside of W for δ close to 1, so it is not immediate from that for any discount factor δ , one can construct a self-generating set containing v according to the B operator rather than the \hat{B} operator. The next lemma shows that this can be done for a non-empty interval of discount factors.

Lemma 5.12. *Let $v \in \text{int}(V_{\alpha^*})$ and suppose that Assumption 5.3 holds. Consider the repeated game with payoffs given by (6). Then there exist $\underline{\delta}, \bar{\delta} \in (0, 1)$ with $\underline{\delta} < \bar{\delta}$ and $\bar{\varepsilon} \in (0, 1)$ such that $v \in E^{\text{pu}}(\delta, \varepsilon)$ for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$. Furthermore there exists some compact set $W \subseteq \text{int}(V_{\alpha^*})$ such that the equilibrium corresponding to payoff v can be taken to have continuation values that always lie in W for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$.*

Proof. Fix some $v \in \text{int}(V_{\alpha^*})$. Then choose a compact, smooth, convex set $\hat{W} \subseteq \text{int}(V_{\alpha^*})$ such that $v \in \text{int}(\hat{W})$. Since \hat{W} is bounded away from 0 and contains v , there exists some $\eta < 1$ and compact set W such that $v \in \eta' \hat{W} \subseteq W \subseteq \text{int}(V_{\alpha^*})$ for all $\eta' \in [\eta, 1]$. By Lemma 5.11, there exists some $\underline{\delta}$ and ε^* such that $\hat{W} \subseteq \hat{B}(\hat{W}, \delta, \varepsilon)$ for all $\delta \geq \underline{\delta}$ and all $\varepsilon < \bar{\varepsilon}$.

Now choose $\bar{\delta} \in (\underline{\delta}, 1)$ arbitrarily. Then choose

$$\bar{\varepsilon} = \min \left\{ \frac{(1-\eta)(1-\bar{\delta})}{\bar{\delta}\eta}, \varepsilon^* \right\}.$$

This then implies that for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$,

$$v \in W_{\delta, \varepsilon} \equiv \frac{1-\delta}{1-\delta(1-\varepsilon)} \hat{W} \subseteq W \subseteq \text{int}(V_{\alpha^*}).$$

Furthermore $\hat{W} \subseteq \hat{B}(\hat{W}, \delta, \varepsilon)$ for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$.

This observation allows us to establish the claims of the lemma. To see this, we note that for every $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$, every $\check{w} \in \hat{W}$ can be written in the form

$$\check{w}_i = (1-\delta(1-\varepsilon))g_i(\alpha) + \delta(1-\varepsilon) \sum_{y \in Y} \hat{w}_i(y)\pi(y|\alpha)$$

for all i for some α and some $\hat{w} : Y \rightarrow \hat{W}$ so that α_i is a best response given the expected continuation payoff \hat{w}_i and opponents' current mixed action profile α_{-i} . Translating payoffs into the original normalization under (6), yields

$$\frac{1-\delta}{1-\delta(1-\varepsilon)} \check{w}_i = (1-\delta)g_i(\alpha) + \delta(1-\varepsilon) \sum_{y \in Y} \frac{1-\delta}{1-\delta(1-\varepsilon)} \hat{w}_i(y)\pi(y|\alpha).$$

We then see that

$$\frac{1-\delta}{1-\delta(1-\varepsilon)} \hat{w}_i(y) \in W_{\delta, \varepsilon}$$

for all $y \in Y$ and all i . Thus $v \in W_{\delta, \varepsilon} \subseteq B(W_{\delta, \varepsilon}, \delta, \varepsilon)$ and $W_{\delta, \varepsilon} \subseteq W$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$. Then from Lemma 5.10, if $v \in W \subseteq B(W_{\delta, \varepsilon}, \delta, \varepsilon)$ then $v \in E^{\text{pu}}(\delta, \varepsilon)$. Therefore $v \in E^{\text{pu}}(\delta, \varepsilon)$ for all $\varepsilon < \bar{\varepsilon}$ and all $\delta \in [\underline{\delta}, \bar{\delta}]$. \square

Then we relate the auxiliary game $G^{\text{pu}}(\delta, \varepsilon)$ back to the original private monitoring game $G^{\text{pr}}(\delta, \pi^{\text{pr}})$ as follows. We let $\varepsilon = \pi^{\text{pr}}((\infty, \dots, \infty) | a)$.²³ Furthermore when constructing

²³Here we use Assumption 5.5 so that ε does not depend on $a \in A$.

strategies that satisfy condition 3 of Lemma 5.8, players play as if they are observing a public signal structure over $Y \cup \{\infty\}$ with $\pi^{\text{pu}}(\infty | a) = \varepsilon$ and $\pi^{\text{pu}}(y | a) = \sum_{\tilde{y}=y} \pi^{\text{pr}}(\tilde{y} | a) = (1 - \varepsilon)\pi(y | a)$ by Assumption 5.5. With these observations, lemmas 5.8 and 5.12 together prove Theorem 5.7.

5.5. The Repeated Game with Observation Lags

We now prove Theorem 5.4. To this end, let us first link the private monitoring game with communication, $G^{\text{pr}}(\delta, \pi^{\text{pr}})$, to the original repeated game with public monitoring and observation lags: For a given lag distribution λ and some $T \in \mathbb{N}$, we define the induced private monitoring structure π^{pr} in the following way:

$$\pi^{\text{pr}}(\tilde{y} | a) = \begin{cases} \Pr(\{L : L_i \geq T \ \forall i \in \mathcal{I}(\tilde{y}), L_i < T \ \forall i \notin \mathcal{I}(\tilde{y})\})\pi(y | a) & \text{if } \tilde{y} \in \mathcal{Y}, \vec{y} = y \\ \Pr(\{L : L_i \geq T \ \forall i\}) & \text{if } \tilde{y} = (\infty, \dots, \infty) \\ 0 & \text{otherwise.} \end{cases}$$

Note that π^{pr} satisfies Assumptions 5.5 and 5.6. Then given this monitoring structure, we define the game $\tilde{G}(\delta, \lambda, T) = G^{\text{pr}}(\delta, \pi^{\text{pr}})$, and let $\tilde{E}(\delta, \lambda, T)$ be the set of sequential equilibrium payoffs of $\tilde{G}(\delta, \lambda, T)$ for which equilibrium play depends only on the message histories.

In constructing an equilibrium for the repeated game with observation lags, we suppose that the message spaces in each period are $M_i = \tilde{Y}$. Henceforth $G(\delta, \lambda)$ and $E(\delta, \lambda)$ specifically refer to the repeated game with observation lag distribution λ , discount factor δ , and message spaces $M_i = \tilde{Y}$.

Lemma 5.13. *Suppose that $v \in \tilde{E}(\delta, \lambda, T)$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ for some fixed λ and all $T \geq T^*$, where $0 < \underline{\delta} < \bar{\delta} < 1$. Then there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$.*

As in the proof of Lemma 4.6, the proof here also divides the repeated game into threads, mapping each thread to an auxiliary game of the form described in the preceding sections. However because the lemma here additionally allows for communication, care in defining the communication strategies is necessary in order to appropriately construct the map from threads to auxiliary games.

Proof. We set $\delta^* = (\underline{\delta}/\bar{\delta})^{\frac{1}{T^*+1}}$, so that for every $\delta > \delta^*$, there exists a positive integer multiple of $T^* + 1$, $N(\delta)$, such that $\delta^{N(\delta)} \in [\underline{\delta}, \bar{\delta}]$.

Now we divide the repeated game $G(\delta, \lambda)$ into $N(\delta)$ distinct repeated game threads, the ℓ -th ($1 \leq \ell \leq N(\delta)$) of which is played in periods $\ell, N(\delta) + \ell, 2N(\delta) + \ell, \dots$. In our construction, players communicate the public signal generated at the end of period $(k - 1)N(\delta) + m$ at the end of period $kN(\delta) + (m - 1)$. If they have not yet seen the signal of that period's play they report the null signal. Then each repeated game thread is equivalent to a private monitoring game of the form described in the previous section.

As in the proof of Lemma 4.6, each repeated game can be treated independently, as players never condition their play in the ℓ -th repeated game on information received about play in the ℓ' -th repeated games ($\ell' \neq \ell$). Moreover, any equilibrium of $\tilde{G}(\delta^{N(\delta)}, \lambda, N(\delta))$ where play depends only on the message history can be embedded into an equilibrium of one of the repeated game threads. But since $N(\delta) > T^* + 1$, we have $v \in \tilde{E}(\delta^{N(\delta)}, \lambda, N(\delta))$, so it is then clear that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$. \square

We can now finish the proof of Theorem 5.4.

Proof of Theorem 5.4. By Theorem 5.7, there exist $\underline{\delta}, \bar{\delta} \in (0, 1)$ with $\underline{\delta} < \bar{\delta}$ and $\varepsilon^* \in (0, 1)$ such that $v \in E^{\text{Pr}}(\delta, \pi^{\text{Pr}})$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all π^{Pr} that is ε^* -close to π .

Then choose $\bar{\varepsilon} > 0$ such that $\frac{\varepsilon}{1-\varepsilon} < \varepsilon^*$ for all $\varepsilon \leq \bar{\varepsilon}$. Now suppose that $\lambda_i(\infty) < \bar{\varepsilon}$ and $\gamma_i < \bar{\varepsilon}$ for all i . Then there exists a (finite) K^* such that

$$\frac{\Pr(L_i \leq K, L_j > K \forall j \neq i)}{\Lambda_i(K)} \leq \frac{\Pr(L_j > K \forall j \neq i)}{\Lambda_i(K)} < \varepsilon^* \quad \text{and} \quad \frac{1 - \Lambda(K)}{1 - \Lambda_i(K)} < \varepsilon^*$$

for all i and $K \geq K^*$. Thus $v \in \tilde{E}(\delta, \lambda, T)$ for all $\delta \in [\underline{\delta}, \bar{\delta}]$ and all $T \geq K^*$ since it is easy to show that the π^{Pr} induced by λ and T is in fact ε^* -close to π for all $T \geq K^*$. This however means—by Lemma 5.13—that there exists some $\delta^* \in (0, 1)$ such that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$; this concludes the proof. \square

Remark. Note that the proof of this theorem uses delayed-response strategies in three ways: to ensure that in each thread there is very low probability of all players' lags being longer than the thread length; so that even after not observing any signal from the previous period in a thread, players believe with high probability that others have observed an informative signal; *and* to map discount factors near 1 in the game $G(\delta, \lambda)$ to intermediate discount factors in the auxiliary games. The first feature is also present in the proof of Lemma 4.6. The second and third features are specific to the proof here. The second, ensured by the assumption that γ_i is small, is key to establishing incentives for truthful communication, so that π^{Pr} in the auxiliary game can be shown to satisfy condition (1). The third feature is closely analogous to the use of threads in the work of Ellison (1994).

Remark. Note that an important part of the proof of Theorem 5.4 is that messages are instantaneously observed. As in the literature on private monitoring games with communication, this is important since the messages serve to make private information public.

However, it is straightforward to extend our argument to settings in which messages are observed with a bounded lag. To see this, suppose that lags arrive within \tilde{K} periods with probability 1. We separate the game into $\hat{K} \equiv \max\{K, \tilde{K}\} + 1$ threads, where K is as in the proof of Theorem 5.4 (page 26). Each thread is further subdivided into a pair of subthreads, respectively played in “even” and “odd” thread periods; players communicate information

observed in the even (resp. odd) subthread in periods of the odd (resp. even) subthread. Since the gap between thread periods is at least \hat{K} , all messages sent in the even (resp. odd) subthread arrive with probability 1 before the next period of the odd (resp. even) subthread. Thus messages about play in the odd (resp. even) subthread arrive in time for the next round of play in that subthread. More formally, the ℓ -th thread is separated into two subthreads so that:

1. In periods $(2k)\hat{K} + \ell$, the players send messages about the signals generated in period $(2k - 1)\hat{K} + \ell$, and in periods $(2k + 1)\hat{K} + \ell$, the players play the appropriate responses to the messages sent in periods $2\hat{K} + \ell, 4\hat{K} + \ell, \dots, (2k)\hat{K} + \ell$.
2. In periods $(2k + 1)\hat{K} + \ell$, the players send messages about the signals generated in period $(2k)\hat{K} + \ell$, and in periods $(2k + 2)\hat{K} + \ell$, the players play the appropriate responses to the messages sent in periods $\hat{K} + \ell, 3\hat{K} + \ell, \dots, (2k + 1)\hat{K} + \ell$.

Under this construction, with the number of threads larger than $\hat{K} = \max\{K, \check{K}\} + 1$, messages sent in period $(2k)\hat{K} + \ell$ (resp. period $(2k + 1)\hat{K} + \ell$) are observed with probability 1 by the time at which players must act on them—period $(2k + 1)\hat{K} + \ell$ (resp. period $(2k + 2)\hat{K} + \ell$). However, we do not know whether a folk theorem would obtain if lags of message transmission are possibly unbounded.²⁴

6. Discussion and Conclusion

As we argued in the introduction, the key role of the repeated games model makes it important to understand which of its many simplifications are essential for the folk theorem. We have extended this result to two settings in which players' information about others' play arrives with stochastic lags. In both of the settings we consider, there is a special but natural form of private information, as players do not know whether and when their opponents observe signals.

Our proof in the case of almost-perfect monitoring (and no communication) depends on the methods of HO2006. Unfortunately, our proof technique does not extend to repeated games with n players. We could attempt to classify any history containing the null signal as an erroneous history and follow the approach of HO2006 for n -player games, but this approach is invalid because of the HO2006 n -player proof's requirement of *communication phases*. For repeated games with observation lags having finite support (possibly including ∞), it may seem that the discussion in Remark 4 of HO2006 regarding almost-perfect monitoring private monitoring games with general signal spaces could be useful. This is due to the fact that as long as the lag distribution has finite support, we can take the K

²⁴In any event, as motivated in footnote 5, it seems reasonable to assume that message delays are much shorter than signal lags.

chosen in Lemma 4.6 to be sufficiently large so that each thread corresponds to a private monitoring game.²⁵ However the conjecture in Remark 4 of HO2006 regarding the partition of signals contains an error and thus cannot be applied.²⁶ Instead, we conjecture that the set of all belief-free equilibrium payoffs in n -player games without communication can be attained in the analogous games with lags. Using results from Yamamoto (2009), one could then obtain a lower bound on the limit equilibrium payoff sets for n -player repeated games with almost-perfect monitoring structures and observation lags.

A more substantial extension of our results would be to the case in which the lag distribution varies with the discount factor. It seems likely that our results would extend to settings in which longer lags become somewhat more likely as players become more patient, but we do not know how rapid an increase can be accommodated.

Acknowledgments

The authors thank Attila Ambrus, Daniel Fudenberg, Johannes Hörner, Yuichiro Kamada, Wojciech Olszewski, Takuo Sugaya, Satoru Takahashi, Yuichi Yamamoto, the associate editor, and several referees for helpful comments. This work was supported by NSF grants SES-0951462 and CCF-1216095, an NSF Graduate Research Fellowship, a Yahoo! Key Scientific Challenges Program Fellowship, a Terence M. Considine Fellowship funded by the John M. Olin Center at Harvard Law School, and an AMS-Simons Travel Grant.

Appendix A. Definitions of s_i^G and s_i^B

We recall the following definitions of $\tilde{s}_i^G, \tilde{s}_i^B \in \tilde{\mathcal{S}}_i^\rho$ of HO2006.

First, define \tilde{s}_i^g as some strategy such that $\tilde{s}_i^g[\emptyset] \in \Delta G$ and

$$\text{for all } \tilde{h}_i^t = \left(a, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}), \dots, (a_i^{M_2, M_1}, a_{-i}^{M_2, M_1}) \right), a \in M_i \times M_{-i}, t \geq 1 : \\ \tilde{s}_i^g[h_i^t] = a_i^{M_2, M_1};$$

²⁵Note that this is not the case if the lag distribution's support is not finite.

²⁶Specifically, Remark 4 suggests that one can find a partition of the private signals to restore the invertibility of the appropriate information matrix so that their results go through, with the elements of the partition treated as the set of private signals. However inference about others' private histories is different across different signals within the same element of the partition so that it is not clear whether the appropriate incentive compatibility conditions would hold.

and define define \tilde{s}_i^b such that $\tilde{s}_i^b[\emptyset] \in \Delta B$ and

$$\begin{aligned} \text{for all } \tilde{h}_i^t &= \left(a, (a_i^{\text{M}_2, \text{M}_1}, a_{-i}^{\text{M}_2, \text{M}_1}), \dots, (a_i^{\text{M}_2, \text{M}_1}, a_{-i}^{\text{M}_2, \text{M}_1}) \right), a \in \text{M}_i \times \text{M}_{-i}, t \geq 1 : \\ \tilde{s}_i^b[\tilde{h}_i^t] &= a_i^{\text{M}_2, \text{M}_1}. \end{aligned}$$

Moreover define $\tilde{s}_i^b[h_i^t] = \alpha_i^m$ for every history \tilde{h}_i^t that is a continuation of a history

$$\begin{aligned} \tilde{h}_i^r &= \left(a, (a_i^{\text{M}_2, \text{M}_1}, a_{-i}^{\text{M}_2, \text{M}_1}), \dots, (a_i^{\text{M}_2, \text{M}_1}, a_{-i}^{\text{M}_2, \text{M}_1}), (a_i^{\text{M}_2, \text{M}_1}, a'_{-i}) \right), \\ a &\in B \times \text{M}_{-i}, a'_{-i} \neq a_{-i}^{\text{M}_2, \text{M}_1}, t \geq r \geq 1, \end{aligned}$$

where α_i^m is a possibly mixed minmax action against player $-i$. Then \tilde{s}_i^G and \tilde{s}_i^B are defined as small perturbations of \tilde{s}_i^g and \tilde{s}_i^b to obtain a pair of strategies \tilde{s}_i^G and \tilde{s}_i^B in $\tilde{\mathcal{S}}_i^\rho$. HO2006 showed that by choosing ρ sufficiently small, we may ensure that there exists some T such that

$$\min_{\tilde{s}_i} \tilde{U}_i^T(\tilde{s}_i, \tilde{s}_{-i}^G) > \bar{v}_i > v_i > \underline{v}_i > \max_{\tilde{s}_{-i}^T} \tilde{U}_i^T(\tilde{s}_i, \tilde{s}_{-i}^B)$$

for sufficiently patient players.

Given the above definitions, we can obtain strategies in repeated games *with observation lags and private monitoring structure* π in a natural way by identifying with each $\tilde{h}_i^t \in \tilde{H}_i^t$ the unique element of $h_i^t \in H_i^t$ such that h_i^t and \tilde{h}_i^t report exactly the same observations about the play of player $-i$ at all times and h_i^t contains no observations with a positive lag. We denote this identification by $\tilde{h}_i^t \simeq h_i^t$.

Similarly we identify $s_i \in \mathcal{S}_i^T$ to a strategy $\tilde{s}_i \in \tilde{\mathcal{S}}_i^T$ in a natural way. We say that $\tilde{s}_i \simeq s_i$ if $s_i[h_i^t] = \tilde{s}_i[\tilde{h}_i^t]$ for all $\tilde{h}_i^t \in \tilde{H}_i^t$ and all $h_i^t \in H_i^t$ such that $\tilde{h}_i^t \simeq h_i^t$. Then we simply define s_i^G and s_i^B to be strategies such that $\tilde{s}_i^G \simeq s_i^G$ and $\tilde{s}_i^B \simeq s_i^B$. It is easy to see that we can appropriately define s_i^G and s_i^B at all histories $h_i^t \in H_i^t$ that are not identified with some $\tilde{h}_i^t \in \tilde{H}_i^t$ to obtain a pair of strategies $s_i^G, s_i^B \in \mathcal{S}_i^\rho$. Moreover when the probability of observation lag is small and π is very close to perfect monitoring, it is clear that $U_i^T(s_i, s_{-i})$ is close to $\tilde{U}_i^T(\tilde{s}_i, \tilde{s}_{-i})$ where $\tilde{s}_i \simeq s_i$ and $\tilde{s}_{-i} \simeq s_{-i}$. Thus we have for sufficiently patient players and ρ sufficiently small,

$$\min_{s_i} U_i^T(s_i, s_{-i}^G) > \bar{v}_i > v_i > \underline{v}_i > \max_{s_{-i}^T} U_i^T(s_i, s_{-i}^B).$$

Appendix B. Details of the Proof of Theorem 4.3

Appendix B.1. Proof of Lemma 4.4

Proof of Lemma 4.4: We wish to specify transfers $\xi_i^B : H_{-i}^T \rightarrow \mathbb{R}_-$ in such a way that players are indifferent between all possible strategies in the T -period repeated game given auxiliary

transfers ξ_i^B . To do this, we define equivalence classes over T -period histories in the following way:

$$(h_{-i}^{T-1}, a_i^{t_1}, a_i^{t_2}, \dots, a_i^{t_m}, a_i^T, a_{-i}^T) \sim (\hat{h}_{-i}^{T-1}, \hat{a}_i^{t_1}, \hat{a}_i^{t_2}, \dots, \hat{a}_i^{t_m}, \hat{a}_i^T, \hat{a}_{-i}^T)$$

if and only if $h_{-i}^{T-1} = \hat{h}_{-i}^{T-1}$ and $a_i^T = \hat{a}_i^T$. Here, if player $-i$ does not obtain information about the play of player i in time T , then a_i^T is taken to be ∞ (representing a null signal). Also notationally, $a_i^{t_1}, \dots, a_i^{t_m}$ are the elements of $h_i^{T,o}$ that are not equal to a_i^T . We may represent this equivalence class of T period histories in the form (h_{-i}^{T-1}, a_i^T) ; note that this indicates that neither

1. the action played by player $-i$ in period T , nor
2. new information gained about past actions

matter for the determination of the equivalence class.

We define equivalence classes over t -period histories similarly, and represent such an equivalence class by (h_{-i}^{t-1}, a_i^t) . We now define a transfer function ξ_i^B as in HO2006 for some functions θ_t defined over equivalence classes of t -period histories:

$$\xi_i^B(h_{-i}^T) = \frac{1}{\delta^T} \sum_{t=1}^T \delta^{t-1} \theta_t(h_{-i}^{t-1}, a_i^t).$$

Here, h_{-i}^{t-1} is the t -period truncation of h_{-i}^T and a_i^t is the signal that player $-i$ observed of player i 's period- t action in period t . That is, $a_i^t = \infty$ if player $-i$ does not observe i 's play immediately and is otherwise equal to the actual period- t action of player i .²⁷

Given any h_{-i}^{T-1} , consider the matrix

$$\left(\mu_{-i}(\cdot \mid a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})) \right)_{a_i \in A_i}.$$

Note that the matrix above has full row rank when $\lambda_{-i}(0)$ is sufficiently close to 1 and π is sufficiently close to perfect monitoring. Therefore the sub-matrix obtained by deleting the column corresponding to the " ∞ " signal is invertible.²⁸ We then set $\theta_t(h_{-i}^{T-1}, \infty) = 0$ and solve the system of equations defined by

$$\mu_{-i}(\cdot \mid a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})) \cdot \theta_T(h_{-i}^{T-1}, \cdot) = g_i(a_i^*, \bar{s}_{-i}^B(h_{-i}^{T-1})) - g_i(a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})), \quad (\text{B.1})$$

where a_i^* is the stage game best response to $\bar{s}_{-i}^B(h_{-i}^{T-1})$. Our preceding observations show that system (B.1) has a unique solution when $\lambda_{-i}(0)$ is sufficiently large and π is sufficiently

²⁷According to this definition, if for example the play of player 1's period-1 action is not observed immediately (i.e. in period 1) by player 2, then the observation of player 1's period-1 action in a later period only has an effect on ξ_i^B through its effect on player $-i$'s play.

²⁸In fact, this sub-matrix approaches the identity matrix as $\lambda(0) \rightarrow 1$ and π approaches perfect monitoring.

close to perfect monitoring.

Then in period $T - 1$, player i is indifferent between all of his actions given that player $-i$ plays according to the strategy prescribed by \bar{s}_{-i}^B at history h_{-i}^{T-1} and transfers given by $\theta_T(h_{-i}^{T-1}, \cdot)$, as playing any action a_i generates a payoff of

$$\begin{aligned} & (1 - \delta)\delta^{T-1}g_i(a_i, \bar{s}_{-i}^B(h_{-i}^{T-1})) + (1 - \delta)\delta^{T-1} \sum_{\omega_i \in A_i \cup \{\infty\}} \mu_{-i}(\omega_i | a_i, \bar{s}_{-i}^B)\theta_T(h_{-i}^{T-1}, \omega_i) \\ & = (1 - \delta)\delta^{T-1}g_i(a_i^*, \bar{s}_{-i}^B(h_{-i}^{T-1})). \end{aligned}$$

Suppose that all transfers θ_τ for $\tau \geq t$ have been defined so that player i is indifferent across all of his strategies from period $t + 1$ on. Then define $U_{t+1}(h_{-i}^{t-1}, a_i)$ to be the expected continuation payoff given the transfers at period $t + 1$, given that player $-i$'s history in period $t - 1$ is h_{-i}^{t-1} and player i played a_i in period t .

We now define θ_{t-1} in a similar manner. Again we consider any $h_{-i}^T \in H_{-i}^T$ and consider the following expression:

$$\frac{1}{\delta^T} \sum_{s=t}^T \delta^{s-1} \theta_s(h_{-i}^{s-1}, a_i^s).$$

Again define $\theta_{t-1}(h_{-i}^{t-1}, \infty) = 0$ and consider the matrix

$$\left(\mu_{-i}(\cdot | a_i, \bar{s}_{-i}^B(h_{-i}^{t-1})) \right)_{a_i \in A_i}.$$

Let us denote the sub-matrix obtained by deleting the column corresponding to the null signal “ ∞ ” by $D(h_{-i}^{t-1})$. This is again invertible when $\lambda_{-i}(0)$ is sufficiently close to 1 and π is sufficiently close to perfect monitoring. Now consider the system of equations

$$\begin{aligned} & (1 - \delta)\delta^{t-1} \left(\mu_{-i}(\cdot | a_i, \bar{s}_{-i}^B(h_{-i}^{t-1})) \cdot \theta_t(h_{-i}^{t-1}, \cdot) + g_i(a_i, \bar{s}_{-i}^B(h_{-i}^{t-1})) \right) + (1 - \delta)U_{t+1}(h_{-i}^{t-1}, a_i) \\ & = (1 - \delta)U_{t+1}(h_{-i}^{t-1}, a_i^*(h_{-i}^{t-1})) + (1 - \delta)\delta^{t-1}g_i(a_i^*(h_{-i}^{t-1}), \bar{s}_{-i}^B(h_{-i}^{t-1})) \end{aligned} \quad (\text{B.2})$$

where $a_i^*(h_{-i}^{t-1})$ is the term that maximizes the expression on the right hand side of the equation above.

Because the matrix $D(h_{-i}^{t-1})$ is invertible, the system (B.2) has a unique solution when we set $\theta_t(h_{-i}^{t-1}, \infty) = 0$. Iterating in this manner allows us to obtain the first part of the lemma.

To achieve non-negativity of transfers, we observe that as the square matrices $D(h_{-i}^{t-1})$ converge to the identity matrix, the solutions $\theta_t(h_{-i}^{t-1}, a_i)$ must be non-negative in the limit. Thus we can make all transfers $\theta_t(h_{-i}^{t-1}, a_i)$ non-negative by adding to all of them a positive constant that converges to zero as $\lambda_{-i}(0)$ and π jointly converge to 1 and perfect monitoring respectively.

Finally we define a strategy $r_i^B \in S_i^T$ in the following way. Let $r_i^B(h_i^{t-1})$ be the action $a_i^*(h_{-i}^{t-1})$ as defined above for all histories h_{-i}^{t-1} that do not contain any null signals, where h_i^{t-1} is the history that corresponds to h_{-i}^{t-1} . Define $r_i^B(h_i^{t-1})$ arbitrarily for all other histories. Then note that as monitoring becomes perfect, the expected value of ξ_i^B goes to zero if players play according to r_i^B and \bar{s}_{-i}^B . By the definition of r_i^B , the payoff in the T -times-repeated game without any transfers then approaches $\max_{s_i \in S_i^T} U_i^T(s_i, \bar{s}_{-i}^B)$; this implies that

$$\lim_{\varepsilon \rightarrow 0} U_i^A(s_i, \bar{s}_{-i}^B, \xi_i^B) = \max_{\tilde{s}_i} U_i^T(\tilde{s}_i, \bar{s}_{-i}^B)$$

for all $s_i \in S_i^T$. □

Appendix B.2. Proof of Lemma 4.5

Proof. Let $\varepsilon > 0$ be such that π is ε -perfect and $\Pr(L_{-i} > 0) < \varepsilon$. For every $\nu > 0$, observe that there exists ε/ρ small enough such that, for any history $h_i^{t-1} \in H_i^{R,t-1}$ and conditional on observing h_i^{t-1} , player i assigns probability at least $1 - \nu$ to the event that player $-i$ observed the corresponding history h_{-i}^{t-1} . Consider for some $h_i^{t-1} \in H_i^{R,t-1}$ and any action $a_i \in A_i$, the row vector consisting of the probabilities assigned by player i , conditional on history h_i^{t-1} and on action a_i taken by player i in period t , to the different equivalence classes of histories (h_{-i}^{t-1}, a_i) observed by player $-i$ in period t . As in HO2006, we construct a matrix D^{t-1} by stacking the row vectors for all regular histories $h_i^{t-1} \in H_i^{R,t-1}$ and actions $a_i \in A_i$. Note that for small enough ε/ρ , the matrix D^{t-1} has full row rank for every t .

With this we can define $\theta(\cdot, \cdot)$ by setting $\theta(h_{-i}^{t-1}, \infty) = 0$ for any $h_{-i}^{t-1} \in H_{-i}^{t-1}$. This is possible since the number of rows is exactly the same as in HO2006 and the number of columns corresponding to (h_{-i}^{t-1}, a_i) for some $a_i \neq \infty$ is also the same as in HO2006. This proves the lemma. □

Appendix C. Conditions Guaranteeing Small γ_i

When the measures L_1, \dots, L_n are independent and identically distributed, we have

$$\frac{\lambda(\infty, \dots, \infty)}{\lambda_i(\infty)} = \frac{\lambda_1(\infty) \cdots \lambda_n(\infty)}{\lambda_i(\infty)}. \quad (\text{C.1})$$

Clearly when $\lambda_1(\infty) = \dots = \lambda_n(\infty)$ and $\lambda_1(\infty)$ small, (C.1) is close to 0. Even if L_1, \dots, L_n are not identically distributed but are independent, we again have (C.1) small if $\frac{(\max_i \lambda_i(\infty))^n}{\min_i \lambda_i(\infty)}$ is sufficiently small, i.e. when no player i 's probability of never observing a signal is much smaller than some other player's probability of never observing a signal.

Consider lags for which L_i is split into two components, X and A_i : $L_i = X + A_i$. Note that X is common across all players. We assume that A_i is independent and identically

distributed across players, and that A_i is independent of X . Let ξ be the density of X and χ denote the density of A_i .

Then γ_i is small if

$$\frac{(1 - \xi(\infty))(\chi(\infty))^n + \xi(\infty)}{\xi(\infty) + (1 - \xi(\infty))(\chi(\infty))} \quad (\text{C.2})$$

is small—which is true if $\xi(\infty)$ is much smaller than $\chi(\infty)$. For example, suppose that $\xi(\infty) = 0$ but $\chi(\infty) > 0$. Then if $\chi(\infty)$ is small, (C.2) is small.²⁹

References

- Abreu, D., Milgrom, P., Pearce, D., 1991. Information and timing in repeated partnerships. *Econometrica* 59, 1713–1733.
- Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58, 1041–1064.
- Ahn, T.K., Ostrom, E., Schmidt, D., Walker, J., 2003. Trust in two-person games: Game structures and linkages, in: Ostrom, E., Walker, J. (Eds.), *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, Russell Sage, New York. pp. 323–351.
- Aumann, R.J., Shapley, L.S., 1976. Long-term competition—A game theoretic analysis, in: Megiddo, N. (Ed.), *Essays in Game Theory in Honor of Michael Maschler*. Springer-Verlag, New York, pp. 1–15.
- Ben-Porath, E., Kahneman, M., 2003. Communication in repeated games with costly monitoring. *Games and Economic Behavior* 44, 227–250.
- Bergin, J., MacLeod, W.B., 1993. Continuous time repeated games. *International Economic Review* 34, 21–37.
- Bhaskar, V., Obara, I., 2011. Reputation and information lag. Mimeo.
- Compte, O., 1998. Communication in repeated games with imperfect private monitoring. *Econometrica* 66, 597–626.
- Dutta, P., 1995. A folk theorem for stochastic games. *Journal of Economic Theory* 66, 1–32.
- Ellison, G., 1994. Cooperation in the prisoner’s dilemma with anonymous random matching. *Review of Economic Studies* 61, 567–588.
- Escobar, J., Toikka, J., 2011. A folk theorem with markovian private information. Mimeo.
- Faingold, E., Sannikov, Y., 2011. Reputation in continuous-time games. *Econometrica* 79, 773–876.
- Friedman, J.W., 1971. A noncooperative equilibrium for supergames. *Review of Economic Studies* 38, 1–12.
- Fudenberg, D., Levine, D., 2007a. Continuous time limits of repeated games with imperfect public monitoring. *Review of Economic Dynamics* 10, 173–192.
- Fudenberg, D., Levine, D., 2007b. The Nash-threats folk theorem with communication and approximate common knowledge in two player games. *Journal of Economic Theory* 132, 461–473.
- Fudenberg, D., Levine, D., 2009. Repeated games with frequent signals. *Quarterly Journal of Economics* 124, 233–265.

²⁹By contrast, consider the case in which $\chi(\infty) = \xi(\infty) = \varepsilon > 0$. Then (C.2) equals

$$\frac{(1 - \varepsilon)\varepsilon^n + \varepsilon}{\varepsilon + (1 - \varepsilon)\varepsilon},$$

which converges to 1/2 as $\varepsilon \rightarrow 0$ —so γ_i cannot be taken to be small even when ε is small.

- Fudenberg, D., Levine, D.K., Maskin, E., 1994. The folk theorem with imperfect public information. *Econometrica* 62, 997–1031.
- Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54, 533–554.
- Fudenberg, D., Olszewski, W., 2011. Repeated games with asynchronous monitoring of an imperfect signal. *Games and Economic Behavior* 72, 86–99.
- Fudenberg, D., Yamamoto, Y., 2011. The folk theorem for irreducible stochastic games with imperfect public monitoring. *Journal of Economic Theory* 146, 1664–1683.
- Gächter, S., Herrmann, B., Thoni, C., 2004. Trust, voluntary cooperation, and socio-economic background: Survey and experimental evidence. *Journal of Economic Behavior & Organization* 55, 505–531.
- Hörner, J., Olszewski, W., 2006. The folk theorem for games with private almost-perfect monitoring. *Econometrica* 74, 1499–1544.
- Hörner, J., Olszewski, W., 2009. How robust is the folk theorem? *Quarterly Journal of Economics* 124, 1773–1814.
- Hörner, J., Sugaya, T., Takahashi, S., Vieille, N., 2011. Recursive methods in discounted stochastic games: An algorithm for $\delta < 1$ and a folk theorem. *Econometrica* 79, 1277–1318.
- Kandori, M., Matsushima, H., 1998. Private observation, communication and collusion. *Econometrica* 66, 627–652.
- Lagziel, D., Lehrer, E., 2012. Regret-free strategy with delayed information. Mimeo.
- Levy, Y., 2009. Stochastic games with information lag. Mimeo.
- Mailath, G.J., Morris, S., 2002. Repeated games with almost-public monitoring. *Journal of Economic theory* 102, 189–228.
- Obara, I., 2009. Folk theorem with communication. *Journal of Economic Theory* 144, 120 – 134.
- Rubinstein, A., 1994. Equilibrium in supergames, in: Megiddo, N. (Ed.), *Essays in Game Theory in Honor of Michael Maschler*. Springer-Verlag, New York, pp. 17–28. Master’s thesis, Hebrew Univeristy of Jerusalem, 1977.
- Sannikov, Y., 2007. Games with imperfectly observed actions in continuous time. *Econometrica* 75, 1285–1329.
- Sannikov, Y., Skrzypacz, A., 2007. Impossibility of collusion under imperfect monitoring with flexible production. *American Economic Review* 97, 1794–1823.
- Sannikov, Y., Skrzypacz, A., 2010. The role of information in repeated games with frequent actions. *Econometrica* 78, 847–882.
- Sekiguchi, T., 1997. Efficiency in repeated prisoner’s dilemma with private monitoring. *Journal of Economic Theory* 76, 345–361.
- Sugaya, T., 2011. Folk theorem in repeated games with private monitoring. Mimeo.
- Yamamoto, Y., 2009. A limit characterization of belief-free equilibrium payoffs in repeated games. *Journal of Economic Theory* 144, 802–824.
- Yao, T., Xu, S.H., Jiang, B., 2011. Trigger strategy equilibriums in stochastic differential games with information time lags: an analysis of cooperative production strategies. *SIAM Journal of Control and Optimization* 49, 1197–1220.