Delayed-Response Strategies in Repeated Games with Observation Lags*

Drew Fudenberg† Yuhta Ishii‡ Scott Duke Kominers§

First Version: September 23, 2011
This Version: February 29, 2012

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.

*The authors thank Johannes Hörner, Yuichiro Kamada, Wojciech Olszewski, Takuo Sugaya, Satoru Takahashi, and Yuichi Yamamoto for helpful comments. This work was supported by NSF grant SES-0951462. Kominers gratefully acknowledges the support of an NSF Graduate Research Fellowship, a Yahoo! Key Scientific Challenges Program Fellowship, a Terence M. Considine Fellowship, and an AMS-Simons Travel Grant.

†Department of Economics, Harvard University; dfudenberg@harvard.edu.
‡Department of Economics, Harvard University; ishii2@fas.harvard.edu.
§Becker Friedman Institute for Research in Economics, University of Chicago; skominers@uchicago.edu.

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

[1]See e.g., Ahn, Ostrom, Schmidt, and Walker (2003); Gachter, Herrmann, and Thoni (2004).

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. Such observation lags seem plausible in many cases, but seem especially appropriate—or possibly even physically necessary—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)).

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, here 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.[2]

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.[3] 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. 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.

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. In the auxiliary game, the probability of seeing a signal immediately is assumed to be independent of the action profile played and of the signals seen by other players. 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

[2]The 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.

[3]In 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.

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 obtain equilibria in the game with lags by using delayed-response strategies as described above.

1.1 Past Work on Models without Observation Lags

As noted above, a substantial 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); 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 (Fudenberg, Levine, and Maskin (1994)) or private but accompanied by cheap-talk public messages (Compte (1998); Kandori and Matsushima (1998)),[4] or private and without communication (e.g., Sekiguchi (1997); Mailath and Morris (2002); Hörner and Olszewski (2006); 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.[5] With each type of signal structure, the key assumptions relate to the qualitative nature of the information that signals provide: Roughly

[4]We 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); 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); Sugaya (2011)).

[5]When 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.

speaking, in order for the folk theorem to obtain, signals must be informative enough to "identify deviations" in a statistical sense.[6]

1.2 Past Work on Models with Observation Lags

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 players 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 one. 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., Levy (2009); Yao, Xu, and Jiang (2011)); this sort of lag does not introduce private information and so is quite different from the lags we study.[7] Abreu, Milgrom, and Pearce (1990) showed that accumulation of signals over many periods can actually improve players' information. In their model, however, consecutive signals' are simply grouped together and delivered at once—unlike in our framework, the delay does not introduce private information.

[6]In addition, the folk theorem has been extended to recurrent stochastic games with perfectly or imperfectly observed actions (Dutta (1995); Fudenberg and Yamamoto (2011); Hörner, Sugaya, Takahashi, and Vieille (2011)).

[7]Levy (2009) studied undiscounted zero-sum stochastic games with a deterministic observation lag that increases over time; Yao, Xu, and Jiang (2011) studied trigger strategies in a continuous-time oligopoly model with deterministic lags of perfectly observed signals.

2 General Model

In this section, we introduce a general model which encompasses all the settings discussed in later sections. We consider a repeated game with n players $i \in I \equiv \{1, \ldots, n\}$, each of whom has a finite action space A_i. In each period $t = 0, 1, 2, \ldots$, players choose actions a_i^t; this generates a sequence of action profiles $\{a^t\}_{t=0}^{\infty}$. Each player i has a 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 $\pi(\omega_1, \ldots, \omega_n | a^t)$.

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 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, for each player i, we write $\Omega_i = \prod_{k=1}^{c_i} \Omega_{i,k}$, and suppose that player i observes the component $\omega_{i,k}^t$ of ω_i^t, at a stochastic time $t + L_{i,k}^t$, where $\{L_{i,k}^t\}_{i,k,t}$ is a collection of independently distributed random variables that take values in $\mathbb{N} \cup \{\infty\}$ and are (identically) distributed according to the probability density function $\lambda : \mathbb{N} \cup \{\infty\} \to [0, 1]$. We denote by Λ the cumulative distribution function, i.e. $\Lambda(l) = \sum_{m=0}^{l} \lambda(m)$.[8] (The case $L_{i,k}^t = \infty$ is interpreted as the event in which player i never receives any information about the k-th coordinate of the period-t private signal.)

Observation of $\omega_{i,k}^t$ takes place in period $t + L_{i,k}^t$ after the choice of that period's actions.

[8]The assumption that the collection $\{L_{i,l}^t\}_{i,l,t}$ is i.i.d. is not necessary for the results and is only assumed to simplify the exposition. For example: if the lags are independently distributed but not identical, then our results still hold. We believe that further relaxations of the i.i.d. assumption are also possible.

When player i observes $\omega_{i,k}^t$, he also observes a "timestamp" indicating that $\omega_{i,k}^t$ is associated with play in period t.[9] 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 private information $\omega_{i,k}^{t'}$ for which $t' + L_{i,k}^{t'} \leq t$, at each time $t = 0, 1, \ldots$, 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 $m = (m_1, \ldots, m_n)$.

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.

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

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

where we set

$$g_i(a) = \sum_{\omega \in \Omega} r_i(a_i, \omega_i)\pi(\omega|a).$$

We assume that player i does not observe the flow payoff $r_i(a_i^t, \omega_i^t)$ until i has observed all coordinates of the period-t private signal ω_i. If $L_{i,k}^t = \infty$ for some i, k, and t, then we assume for convenience that player i never observes the period-t flow payoff $r_i(a_i^t, \omega_i^t)$.[10]

We let a^* be a Nash equilibrium of the stage game and develop a series of Nash threat

[9] Thus player i cannot respond to the period-$(t + L_{i,k}^t)$ observation information until time $t + L_{i,k}^t + 1$.

[10] We make these assumptions about the observation of payoffs so that players cannot infer their private signals from payoffs.

folk theorems. For ease of exposition, we normalize payoffs of players so that $g_i(a^*) = 0$ for all i. We let V denote the convex hull of the feasible set of payoffs, and let V_{a^*} be the convex hull of the set consisting of $g(a^*) = 0$ and the payoff vectors Pareto-dominating $g(a^*) = 0$:

$$V_{a^*} \equiv \{v \in V \mid v \geq 0\}.$$

We assume that $\text{int}(V_{a^*})$ is non-empty. Furthermore, 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 $G(\delta, \pi) \equiv G(\delta, \pi, \mathsf{imm})$, where imm is the (degenerate) distribution which puts full weight on immediate observation, and define $E(\delta, \pi) \equiv E(\delta, \pi, \mathsf{imm})$ similarly.

Finally we introduce the concept of delayed-response strategies, which will be used throughout the remainder of the paper to prove our folk theorems.

Definition 2.1. Let σ be a strategy profile. Then σ is a *delayed-response strategy profile* if there exists some K such that for all $t = 0, 1, \ldots, K$, $\sigma(h^t) = \sigma(\hat{h}^t)$ for all $h^t, \hat{h}^t \in H^t$ and for all $t = \ell K + t'$ where $t' \in \{0, \ldots, K\}$ and $\ell \geq 1$, $\sigma(h^{t-1}) = \sigma(\hat{h}^{t-1})$ whenever $h^{t-1, \ell' K + t'} = \hat{h}^{t-1, \ell' K + t'}$ for all $\ell' = 0, 1, \ldots, \ell - 1$.

In words, this means that σ is a delayed-response strategy profile if there exists some K such that players at time t condition their play only on information regarding a signal that was generated in period $t - \ell K$ for some $\ell \in \mathbb{N}$.

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 $\Lambda(K) = 1$.

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 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.

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. Divide the periods of the repeated game into $K + 1$ threads, with the ℓ-th thread consisting of periods ℓ, $K + 1 + \ell$, $2K + 2 + \ell$,...

For any equilibrium strategy profile σ of the game without lags, construct the associated delayed-response strategy profile σ^K by specifying that if $t = j(K+1) + \ell$ where $0 \leq \ell \leq K$, players treat history h^t as if it were period j of the game without lags. Thus if σ is the equilibrium that generates $v \in E(\delta, \pi)$, the associaed delayed-response version σ^K is an equilibrium for $\delta' = \delta^{K+1}$. ∎

The proof of Theorem 3.2 relies heavily on Assumption 3.1. For example, if the support of λ were concentrated on $0, 1, \ldots, K < \infty$ and ∞, then the proof above would not work since each of threads that it constructs would be a repeated game with a private monitoring structure $\tilde{\pi}$ that is different from π. More problematically, if $\lambda(k) > 0$ for all $k \in \mathbb{N}$ so all lag lengths have positive probability, then no matter how far apart the threads are spaced, there is always a positive probability that the 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 we place additional assumptions on the monitoring structure π. Therefore for the remainder of the paper, we dispense with Assumption 3.1 and allow λ to be any arbitrary probability distribution on $\mathbb{N} \cup \{\infty\}$.

4 Lagged Almost-Perfect Monitoring in the Case of Two Players

In this section, we extend an approach of Hörner and Olszewski (2006) (henceforth, HO2006), in order 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*

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

monitoring: We allow a general lag structure here, but restrict the private signal space of each player i to be $\Omega_i = A_j$ (so $c_i = 1$) and furthermore assume that π is ε-perfect in the sense of HO2006.

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 exist some $\bar{\varepsilon}, \bar{\delta} \in (0,1)$ such that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \bar{\delta}$, all lag distributions λ for which $\lambda(\infty) < \bar{\varepsilon}$, and all private monitoring structures π that are $\bar{\varepsilon}$-perfect.*[12]

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

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

we can shrink the probability of a positive lag close to 0. We thus obtain a folk theorem in the game with stochastic lags.

4.2 Auxiliary Repeated Game with Rare Observation Lags

We first analyze the case of repeated games with "rare" observation lags, i.e. lagged repeated games for which $\lambda(0)$ is close to one. For such games, we obtain the following folk theorem.

Theorem 4.3. *Let* $v \in \operatorname{int}(V^*)$*. Then there exist* $\bar{\varepsilon}, \underline{\delta} \in (0,1)$ *such that if* $\lambda(0) > 1 - \bar{\varepsilon}$*,* 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, and then implements the resulting strategies in delayed-response equilibria.

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. The technique of HO2006 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.[13] 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 by adding the histories where observations

[13]A strategy is *belief-free at time* t if the continuation strategy at time t, $s_i \mid h_i^{t-1}$, is a best response against $s_{-i} \mid h_{-i}^{t-1}$ 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.

arrive with a positive lag to the set of "erroneous" histories. This allows us to extend the methods of HO2006 to lagged repeated games for which $\lambda(0)$ is sufficiently close to 1.

4.2.1 Preliminaries

Before we begin this discussion, we review and extend some definitions and notations of HO2006.

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

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

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.

Analogously we define $\tilde{H}_i^t$ to be the set of t-period histories in the repeated game *without observation lags and with perfect monitoring*. We denote a typical element of $\tilde{H}_i^t$ by $\tilde{h}_i^t$. Also let us define the set of strategies in the T-times repeated game with perfect monitoring and no observation lags by $\tilde{S}_i^T$.

As in the approach of HO2006, we partition the set of private histories in the T-times-repeated stage game into two sets, H_i^R and H_i^E, which we respectively call the *regular* and *erroneous* histories.

To define these sets of histories 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* as in HO2006. We include the definitions for completeness, but since these definitions are identical to those introduced in

HO2006, we omit any discussion. We partition the set A_i into two subsets, denoted G and B. We call an instance of the T-times repeated game with perfect monitoring a *block*, and say that a player i *sends message* $\mathsf{M} \in \{\mathsf{G}, \mathsf{B}\}$ if he picks an action in 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 $a^{\mathsf{X},\mathsf{Y}}$ for $(\mathsf{X}, \mathsf{Y}) \in \{\mathsf{G}, \mathsf{B}\}^2$ with

$$w_i^{\mathsf{X},\mathsf{Y}} = g_i(a^{\mathsf{X},\mathsf{Y}}) \ i = 1, 2, \ \mathsf{X}, \mathsf{Y} \in \{\mathsf{G}, \mathsf{B}\},$$

where $w_i^{\mathsf{G},\mathsf{G}} > v_i > w_i^{\mathsf{B},\mathsf{B}}$ for $i = 1, 2$, and

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

As in HO2006, these 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.

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 $\tilde{s}_i[\tilde{h}_i^t] = a_i^{\mathsf{M}_2,\mathsf{M}_1}$ for all

$$\tilde{h}_i^t = (a, (a_i^{\mathsf{M}_2,\mathsf{M}_1}, a_{-i}^{\mathsf{M}_2,\mathsf{M}_1}), \ldots, (a_i^{\mathsf{M}_2,\mathsf{M}_1}, a_{-i}^{\mathsf{M}_2,\mathsf{M}_1}))$$

with $a \in \mathsf{M}_1 \times \mathsf{G}$ $(t \geq 1)$. We then let

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

We now 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). We then define $H_i^{R,t}$ as the image of $\tilde{H}_i^{R,t}$ under this identification. We then 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 a player i did not immediately observe the period-t' play of player $-i$ for some $t' < t$.*

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

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

Finally define

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

Despite the fact that the game being studied is not a private monitoring game, it is useful to define the analogue of a private monitoring structure, the *immediate monitoring structure μ induced by λ and π*:

$$\mu(\omega_1, \omega_2|a_1, a_2) = \begin{cases} (1-\lambda(0))^2 & \omega_1, \omega_2 = \infty \\ \lambda(0)(1-\lambda(0)) \sum_{\omega_1' \in A_2} \pi(\omega_1', \omega_2 \mid a_1, a_2) & \omega_1 = \infty \text{ and } \omega_2 \neq \infty \\ \lambda(0)(1-\lambda(0)) \sum_{\omega_2' \in A_1} \pi(\omega_1, \omega_2' \mid a_1, a_2) & \omega_1 \neq \infty \text{ and } \omega_2 = \infty \\ \lambda(0)^2 \pi(\omega_1, \omega_2 \mid a_1, a_2) & \omega_1, \omega_2 \neq \infty. \end{cases}$$

This monitoring structure μ omits details about the probability distribution over signals in the future. However, it is useful for our 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 $s_i \mid H_i$ 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.* Analagously define U_i^T to be the ex-ante payoff player i in the T-times repeated game *with private monitoring structure π and observation lags.* We consider a version of the the T-times repeated game (with observation lags) which is augmented with a *transfer* $\xi_{-i} : H_i^T \to \mathbb{R}$ at the end of the T-th period. In this *auxilary scenario*, the payoff of i under strategy profile s

is taken to be

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

The set of best responses of player i in the auxilary 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.

Lemma 4.4. *For every strategy profile $\bar{s} \mid H^E$, there exists $\bar{\varepsilon} > 0$ such that whenever $\lambda(0) > 1-\bar{\varepsilon}$ and π is $\bar{\varepsilon}$-perfect, then there exists a nonnegative transfer $\xi_i^{\mathsf{B}} : H_{-i}^T \to \mathbb{R}_+$ such that*

$$S_i^T = B_i(\bar{s}_{-i}^{\mathsf{B}}, \xi_i^{\mathsf{B}})$$

where $\bar{s}_{-i}^{\mathsf{B}} \mid H_{-i}^R = s_{-i}^{\mathsf{B}} \mid H_{-i}^R$ and $\bar{s}_{-i}^{\mathsf{B}} \mid H_{-i}^E = \bar{s}_{-i} \mid H_{-i}^E$, and for every $s_i \in B_i(\bar{s}_{-i}^{\mathsf{B}}, \xi_i^{\mathsf{B}})$,

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

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 appropriate sets and identifying each of the elements of the partition with a particular history in the private monitoring repeated game.

Proof of Lemma 4.4. We wish to specify transfers $\xi_i^{\mathsf{B}} : H_{-i}^T \to \mathbb{R}_-$ 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:

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

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, $a_i^{t_1}, \ldots, a_i^{t_m}$ 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:

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

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.[14]

[14] According 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.

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

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

Note that the matrix above has full row rank when $\lambda(0)$ is sufficiently close to one and π is sufficiently close to perfect monitoring. Therefore the sub-matrix obtained by deleting the column corresponding to the "∞" signal is invertible.[15] We then set $\theta_t(h_{-i}^{T-1}, \infty) = 0$ and solve the system of equations defined by

$$\mu_{-i}(\cdot|a_i, \bar{s}_{-i}^{\mathsf{B}}(h_{-i}^{T-1})) \cdot \theta_T(h_{-i}^{T-1}, \cdot) = g_i(a_i^*, \bar{s}_{-i}^{\mathsf{B}}(h_{-i}^{T-1})) - g_i(a_i, \bar{s}_{-i}^{\mathsf{B}}(h_{-i}^{T-1})), \tag{1}$$

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

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

We proceed in a similar manner using arguments about μ_{-i} to define θ_t for $t < T$; details of this construction are presented in the Appendix. ∎

[15] In fact, this sub-matrix approaches the identity matrix as $\lambda(0) \to 1$ and π approaches perfect monitoring.

4.3 The Repeated Game with Possibly Long 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 may be very long, the folk theorem still obtains when $\lambda(\infty)$ is sufficiently small. Intuitively, this is possible because players may be taken to be arbitrarily patient; the length of the lag therefore should not affect the deterrence of deviations as long as the probability that a signal never arrives is not too large.

We first state the following lemma that 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 4.5. *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 there exists some $\delta^* \in (0,1)$ and $\varepsilon^* \in (0,1)$ such that for all lag distributions λ such that $\lambda(\infty) < \varepsilon^*$, $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$.*

Proof. First, we set $\varepsilon^* = \bar{\varepsilon}/2$. We consider any λ such that $\lambda(\infty) < \varepsilon^*$. There exists $K \in \mathbb{N}$ be such that $(1 - \Lambda(K-1)) < 2\varepsilon^* = \bar{\varepsilon}$. We choose $\delta^* = \underline{\delta}^{\frac{1}{K}}$. Then for every $\delta > \delta^*$, there exists a positive integer, $N(\delta) \geq K$, such that $\delta^{N(\delta)} \in (\underline{\delta}, 1)$.

Now we divide the repeated game $G(\delta, \pi, \lambda)$ into $N(\delta)$ distinct repeated game "threads," the ℓ-th $(1 \leq \ell \leq N(\delta))$ of which is played in periods

$$\ell, \quad N(\delta) + \ell, \quad 2N(\delta) + \ell, \ldots,$$

noting that as $N(\delta) \geq K + 1$, each of these separate repeated games can be regarded as equivalent to $G(\delta^{N(\delta)}, \pi, \hat{\lambda})$ for some $\hat{\lambda}$ such that $\hat{\lambda}(0) > 1 - 2\varepsilon^*$. Now, 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 $v \in E(\delta^{N(\delta)}, \pi, \hat{\lambda})$, it is then clear that $v \in E(\delta, \pi, \lambda)$ for all $\delta > \delta^*$. ∎

Theorem 4.2 follows directly from Lemma 4.5 and Theorem 4.3.

Remark. Note that our approach requires the availability of timestamps, as otherwise players cannot play the equilibria constructed above in which players delay responding to information received early. We believe that it is difficult to imagine a repeated interaction model where timestamps are not observed (unless the payoff structure is defined in an obscure manner).

5 Public Monitoring with Arbitrary Observation Lags and Communication

5.1 Model

In this section, we again consider an n-player repeated stage game with a finite action space A_i for each player. Now, however, we assume that the monitoring structure of the repeated game is *public with stochastic lags*: There is a set of public signals, denoted Y, and we set $\Omega_i = Y$ and $c_i = 1$ for all players $i \in I$. Furthermore we assume that π is supported on the set

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

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 $\pi((y, \ldots, y)|a)$.

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

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

Assumption 5.1 is not strictly necessary for the folk theorem result. We explain how to modify the construction after we specify the strategies used in the proof.

Before presenting our results for this model, we 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 Mailath and Morris (2002); this condition is violated when a player observes the low-probability "null" signal. The possibility of receiving an uninformative signal also prohibits the application of the folk theorem of Sugaya (2011) for general private monitoring games, because the necessary full rank condition 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 show that as long as $|M_i| \geq |Y| + 1$ for all i, a folk theorem can be established.

5.2 The Folk Theorem

We begin our analysis with the simple observation that the repeated play of a^* is an equilibrium of the game with observation lags.[16] We use this fact along with techniques from Abreu, Pearce, and Stacchetti (1990) and Fudenberg, Levine, and Maskin (1994) (hereafter

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

referred to as FLM) to construct equilibria that generate any payoff profile $v \in \text{int}(V_{a^*})$.

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

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

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 $v \in \text{int}(V_{a^*})$ 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 $\delta^*, \varepsilon^* \in (0, 1)$ such that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$ and all lag distributions λ such that $\lambda(\infty) < \varepsilon^*$.*

The remainder of the section proves this result.

5.3 Step 1: Private Monitoring Game with Communication

5.3.1 Step 1a: 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 $M_i = \tilde{Y} \equiv Y \cup \{\infty\}$. Let us first define some notation. For a vector $\tilde{y} \in \tilde{Y}^n$, define $|\tilde{y}| \equiv |\{i : \tilde{y}_i \neq \infty\}|$. Define the following set

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

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

Given an $\varepsilon \in (0, 1)$, the private monitoring structure is given by

$$\pi^p(\tilde{y}|a) \equiv \begin{cases} (1-\varepsilon)^{|\tilde{y}|} \varepsilon^{n-|\tilde{y}|} \pi\big(\vec{\tilde{y}}|a\big) & \tilde{y} \in \mathcal{Y} \setminus (\infty, \ldots, \infty) \\ \varepsilon^n & \tilde{y} = (\infty, \ldots, \infty) \\ 0 & \text{otherwise.} \end{cases}$$

We denote by $G^{\text{pr}}(\delta, \varepsilon)$ the private monitoring game with parameters δ and ε (and communication) and let $E^{\text{pr}}(\delta, \varepsilon)$ be the set of sequential equilibrium payoffs of $G^{\text{pr}}(\delta, \varepsilon)$. We now show

Theorem 5.5. *Let $v \in \text{int}(V_{a^*})$. Then there exist $\underline{\delta}, \overline{\delta} \in (0, 1)$ with $\underline{\delta} < \overline{\delta}$ and $\bar{\varepsilon} \in (0, 1)$ such that $v \in E^{\text{pr}}(\delta, \varepsilon)$ for all $\delta \in [\underline{\delta}, \overline{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$.*

To prove this theorem we construct strategies that generate a payoff profile of v, and are public perfect in the sense that 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 5.6. *Let W be a convex, compact set that is a subset of $\mathrm{int}(V_{a^*})$. Consider a collection of public perfect strategy profiles $\{\sigma^{\delta,\varepsilon}\}$ indexed by δ and ε for all $\delta \in [\underline{\delta}, \overline{\delta}]$ and $\varepsilon < \overline{\varepsilon}$ with 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 in the history, there exists some t such that $m^t \notin \mathcal{Y}$, then all players i play a_i^*.*
3. *Expected continuation values are always contained in W for play of $\sigma^{\delta,\varepsilon}$ in the game $G^{\mathrm{pr}}(\delta, \varepsilon)$ whenever the message history contains only elements in the set $\mathcal{Y}$.*

Then there exists some $\varepsilon^* \leq \overline{\varepsilon}$ such that for all $\varepsilon < \varepsilon^*$ and all $\delta \in [\underline{\delta}, \overline{\delta}]$, truthful communication is incentive compatible at any private history in $G^{\mathrm{pr}}(\delta, \varepsilon)$ given continuation play determined by $\sigma^{\delta,\varepsilon}$ and truthful communication by all other players.

Proof. For sufficiently small ε, we check that there are no profitable one-stage deviations in which a player misreports once and then follows the continuation strategy prescribed by $\sigma_i^{\delta,\varepsilon}$. 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 a_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

$$(1-\varepsilon^{n-1})\sum_{y\in Y}\pi(y|\alpha)w_i(y)$$

for some $\alpha \in \prod_{i=1}^n \Delta(A_i)$ and some expected continuation value function $w : Y \to W$.[17]

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

$$(1-\varepsilon^{n-1})\pi(y'|\alpha)w_i(y') + \varepsilon^{n-1}w_i(y').$$

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

$$(1-\varepsilon^{n-1})\sum_{y\in Y}\pi(y|\alpha)w_i(y) > (1-\varepsilon^{n-1})\pi(y'|\alpha)w_i(y') + \varepsilon^{n-1}w_i(y') \tag{2}$$

for all $y' \in Y$, all $\alpha \in \prod_{i=1}^n \Delta(A_i)$, all $w : Y \to W$, $i = 1, \ldots, n$, and all $\varepsilon < \varepsilon^*$. Note that for a fixed map, $w : Y \to W$ and $\alpha \in \prod_{i=1}^n \Delta(A_i)$, we have (2) for all $i = 1, \ldots, n$ and all

[17] Note that in any sequential equilibrium if a player i observes signal ∞, then he still believes with probability 1 that all players $-i$ played according to their prescribed actions, i.e. that there have been no "unexpected" events. Thus the α in the above expression is indeed the prescribed mixed action in the current period that generated the null signal. However this is immaterial to our proof since we allow α to be any mixed action.

$y' \in Y$ for sufficiently small ε since (2) holds if and only if

$$(1-\varepsilon^{n-1})\sum_{y\neq y'}\frac{\pi(y|\alpha)}{1-\pi(y'|\alpha)}w_i(y) > \varepsilon^{n-1}\frac{1}{1-\pi(y'|\alpha)}w_i(y'). \quad (3)$$

The inequality (3) clearly holds for sufficiently small ε, by 5.1 and the fact that W is a convex set bounded away from 0. Moreover $\prod_{i=1}^{n}\Delta(A_i)$ and the set of all maps $w : Y \to W$ are both compact. Therefore there exists such an ε^* that is uniform across all α and all maps $w : Y \to W$. Thus all players will report the null signal upon observation of a null signal when $\varepsilon < \varepsilon^*$.

Now suppose that player i observes $y \in Y$. By reporting truthfully, player i obtains a payoff of $w_i(y)$ for some map $w : Y \to W$. However by reporting $y' \in Y$ with $y' \neq y$, player i obtains a payoff of $\varepsilon^{n-1}w_i(y')$ while reporting ∞ yields a payoff of $(1-\varepsilon^{n-1})w_i(y)$. Clearly we can take ε^* sufficiently small so that

$$w_i(y) > \max_{y'\neq y}\{\varepsilon^{n-1}w_i(y')\}$$

for all $y \in Y$, all maps $w : Y \to W$, all $i = 1, \ldots, n$, and all $\varepsilon < \varepsilon^*$. Then all players have an incentive to report truthfully upon observing an informative signal when $\varepsilon < \varepsilon^*$ since $w_i(y) > \max_{y'\neq y}\{\varepsilon^{n-1}w_i(y')\}$ and $w_i(y) > (1-\varepsilon^{n-1})w_i(y)$ trivially. This concludes the proof. ∎

Remark. Note that ε^* does not depend on δ. The reason for this is that the set W does not depend on the discount factor δ. [18] This is important for our folk theorem as we must

[18] If instead W is allowed to depend on δ, then an ε^* that is independent of δ cannot necessarily be constructed.

establish a claim about all games with $\varepsilon < \varepsilon^*$ and all discount factors in an interval.

Remark. Assumption 5.1 is only used in the proof of this lemma and is not strictly required there. Instead, if the strategy construction at some histories calls for the play of an α that only generates one public signal y, we can specify that players report y even after seeing ∞. We would then introduce a public randomization device so that after the message report continuation play could still generate payoff vector $w(y)$ with probability $1 - \varepsilon^n$ (and otherwise trigger reversion to the static Nash a^*), rather than probability 1.

5.3.2 Step 1b: Non-Communication Stages

Lemma 5.6 establishes incentives for truth-telling when $\sigma^{\delta,\varepsilon}$ satisfies certain characteristics and ε is sufficiently small. We now show that given truthful communication by all players at all histories, we can construct a collection of strategies $\{\sigma^{\delta,\varepsilon}\}$ that satisfy the necessary properties of Lemma 5.6 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 $\sigma^{\delta,\varepsilon}$, we first specify that players play a^* whenever in the history there exists some t such that $m^t \notin \mathcal{Y}$. Then it is trivial that playing a_i^* is a best response at such a history since opponents play a_{-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 and 2 of Lemma 5.6, 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 $\pi(y|a)$ every period. We then modify this repeated game so that at the beginning of periods $1, 2, \ldots$, the game ends with probability ε^n and each player receives flow payoffs of $0 = g_i(a^*)$ thereafter. This corresponds exactly to the event in which all players report the null signal, triggering all players to play according to a^* forever.[19]

In the modified game, payoffs are given by

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t(1-\varepsilon^n)^t g_i(a^t). \tag{4}$$

We denote by $G^{\mathrm{pu}}(\delta,\varepsilon)$ the public monitoring game with parameters δ and ε and let $E^{\mathrm{pu}}(\delta,\varepsilon)$ be the set of sequential equilibrium payoffs of $G^{\mathrm{pu}}(\delta,\varepsilon)$. Note that in this game the feasible payoff set is not constant in δ and ε, and in particular for any fixed ε, as $\delta \to 1$, the feasible payoff set converges to $\{0\}$, just as the payoffs to grim trigger strategies converge to those of static Nash equilibrium as $\delta \to 1$ in a repeated game with imperfect public monitoring. However for any fixed δ, as $\varepsilon \to 0$, 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 $(1-\delta(1-\varepsilon^n))/(1-\delta)$ to get payoff structure

$$(1-\delta(1-\varepsilon^n))\sum_{t=0}^{\infty}\delta^t(1-\varepsilon^n)^t g_i(a^t). \tag{5}$$

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

Now, our modified game corresponds to a repeated game with discount factor given by $\delta(1-\varepsilon^n)$, 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, Pearce, and Stacchetti (1990)).

Definition 5.7. For $W \subset \mathbb{R}^n$, define $B(W, \delta, \varepsilon)$ to be the set of $v \in \mathbb{R}^n$ such that there exists some mixed action profile α and a map $w : Y \to W$ such that

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

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 $w : Y \to W$ such that

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

for all i. We say that W is *self-generating in the repeated game with payoff structure* (4) *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 (5) with discount factor δ and absorption probability ε* if $W \subseteq \hat{B}(W, \delta, \varepsilon)$.

Because the public monitoring game $G^{\text{pr}}(\delta, \varepsilon)$ 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 5.8. *Suppose W is compact and that $W \subseteq B(W, \delta, \varepsilon)$. Then $W \subseteq E^{\mathrm{pu}}(\delta, \varepsilon)$.*

The proof of Lemma 5.8 is completely standard; we present it in the web appendix.

We now use Lemma 5.8 to prove a folk theorem for our public monitoring game $G^{\mathrm{pr}}(\delta, \varepsilon)$: First, we note that we can apply FLM to the repeated game with discount factor $\delta(1-\varepsilon^n)$, to show the following lemma.

Lemma 5.9. *Suppose that Assumption 5.3 holds. Let $\hat{W}$ be a smooth, compact, convex set in the* $\mathrm{int}(V_{a^*})$*. 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* (5) *with discount factor δ and absorption probability ε.*

Next, we translate the payoff set used in Lemma 5.9 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 (4):

$$W = \frac{1-\delta}{1-\delta(1-\varepsilon^n)}\hat{W}. \tag{6}$$

Of course for any fixed ε and a fixed set $\hat{W}$, as $\delta \to 1$, W shrinks (setwise) towards the point-set $\{0\}$. Thus for any choice of $v \in \mathrm{int}(V_{a^*})$, v will necessarily lie outside of W for δ close to 1. Thus it is not immediate from Lemma 5.9 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 5.10. *Let $v \in \text{int}(V_{a^*})$ and suppose that Assumption 5.3 holds. Consider the repeated game with payoffs given by (4). Then there exist $\underline{\delta}, \overline{\delta} \in (0,1)$ with $\underline{\delta} < \overline{\delta}$ and $\overline{\varepsilon} \in (0,1)$ such that $v \in E^{\text{pu}}(\delta, \varepsilon)$ for all $\varepsilon < \overline{\varepsilon}$ and all $\delta \in [\underline{\delta}, \overline{\delta}]$. Furthermore there exists some compact set $W \subseteq \text{int}(V_{a^*})$ such that the equilibrium corresponding to payoff v can be taken to have continuation values that always lie in W for all $\delta \in [\underline{\delta}, \overline{\delta}]$ and all $\varepsilon < \overline{\varepsilon}$.*

Proof. Fix some $v \in \text{int}(V_{a^*})$. Then choose a compact, smooth, convex set $\hat{W} \subseteq \text{int}(V_{a^*})$ such that $v \in \hat{W}$. Since $\hat{W}$ is bounded away from 0 and contains v, there exists some $\eta < 1$ and compact set W such that $v \in \eta' \hat{W} \subseteq W \subseteq \text{int}(V_{a^*})$ for all $\eta' \in [\eta, 1]$. By Lemma 5.9, 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 $\varepsilon < \varepsilon^*$.

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

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

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

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

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

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

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

for all i for some α and some $\hat{w} : Y \to \hat{W}$ 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 (4), we get

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

We then note that

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

for all $y \in Y$ and all i. Thus we have $v \in W_{\delta,\varepsilon} \subseteq B(W_{\delta,\varepsilon}, \delta, \varepsilon)$ and $W_{\delta,\varepsilon} \subseteq W$ for all $\delta \in [\underline{\delta}, \overline{\delta}]$ and all $\varepsilon < \bar{\varepsilon}$. Then from Lemma 5.8, we can then show that 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 $\varepsilon < \overline{\varepsilon}$ and all $\delta \in [\underline{\delta}, \overline{\delta}]$. ∎

Then lemmas 5.6 and Lemma 5.10 together prove Theorem 5.5. To close the section, we link the private monitoring game with communication, $G^{\text{pr}}(\delta, \varepsilon)$, 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 $\tilde{G}(\delta, \lambda, T) = G^{\text{pr}}(\delta, 1 - \Lambda(T))$, 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.

5.4 Step 2: The Repeated Game with Observation Lags

We can now use the results established in the first two subsections to prove Theorem 5.4. We show that in the repeated game with observation lags, it is sufficient to consider equilibria of the game with communication in which the message spaces in each period are $M_i = \tilde{Y}$. Thus 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 5.11. *Suppose that* $v \in \tilde{E}(\delta, \lambda, T)$ *for all* $\delta \in [\underline{\delta}, \overline{\delta}]$ *for some fixed* λ *and all* $T \geq T^*$, *where* $0 < \underline{\delta} < \overline{\delta} < 1$. *Then there exists some* $\delta^* \in (0, 1)$ *such that* $v \in E(\delta, \lambda)$ *for all* $\delta > \delta^*$.

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

Now we divide the repeated game $G(\delta, L)$ into $N(\delta)$ distinct repeated game threads, the ℓ-th ($1 \leq \ell \leq N(\delta)$) of which is played in periods

$$\ell, \quad N(\delta) + \ell, \quad 2N(\delta) + \ell, \ldots.$$

We have players communicate the public signal generated at the end of period $(k-1)N(\delta)+m$ at the end of period $kN(\delta) + (m - 1)$, so that 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 4.5, 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 $\tilde{G}(\delta^{N(\delta)}, L, N(\delta))$ 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 $v \in \tilde{E}(\delta^{N(\delta)}, L, N(\delta))$, so it is then clear that $v \in E(\delta, L)$ for all $\delta > \delta^*$. ∎

Remark. Note that the proof of this lemma uses delayed-response strategies in two ways: delayed-response both ensures that in each thread there is very low probability of a lag longer than the thread length, *and* maps discount factors near 1 in the game $G(\delta, \lambda)$ to intermediate discount factors in the auxiliary games. We need this latter feature here because the grim strategies used to punish misreporting are excessively strong when δ is near 1. This second use of threading is closely analogous to the use of threads in the work of Ellison (1994).

We can now finish the proof of Theorem 5.4.

Proof of Theorem 5.4. By Theorem 5.5, there exist $\underline{\delta}, \overline{\delta} \in (0,1)$ with $\underline{\delta} < \overline{\delta}$ and $\varepsilon^* \in (0,1)$ such that $v \in E^{\mathrm{pr}}(\delta, \varepsilon)$ for all $\delta \in [\underline{\delta}, \overline{\delta}]$ and all $\varepsilon < \varepsilon^*$.

Now, suppose that $\lambda(\infty) < \varepsilon^*$. Then, there exists a (finite) K^* such that $1 - \Lambda(K^*) < \varepsilon^*$. Then, we have that $v \in \tilde{E}(\delta, \lambda, T)$ for all $\delta \in [\underline{\delta}, \overline{\delta}]$ and all $T \geq K^*$. This however means—by Lemma 5.11—that there exists some $\delta^* \in (0,1)$ such that $v \in E(\delta, \lambda)$ for all $\delta > \delta^*$; this concludes the proof. ∎

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) crucially 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 chosen in Lemma 4.5 to be sufficiently large so that each thread corresponds to a private monitoring game.[20] However the conjecture in Remark 4 of HO2006 regarding the partition of signals contains an error and thus cannot be applied.[21] Instead we conjecture that the set of all belief-free equilibrium payoffs in n-player games without communication can be attained in the game with lags. Using results from Yamamoto (2009), one could then obtain a lower bound on the limit equilibrium payoff set for n-player repeated games with almost-perfect monitoring structures and observation lags.

A more substantial extension of our results would be 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.

A Details Omitted from the Proof of Theorem 4.3

A.1 Details Omitted from the Proof of Lemma 4.4

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

[20] Note that this is not the case if the lag distribution's support is not finite.

[21] We thank Yuichi Yamamoto for pointing out this error.

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

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

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 $\lambda(0)$ is sufficiently close to one and π is sufficiently close to perfect monitoring. Now consider the system of equations

$$\begin{aligned}
&(1-\delta)\delta^{t-1}\left(\mu_{-i}(\cdot|a_i, \bar{s}_{-i}^{\mathsf{B}}(h_{-i}^{t-1})) \cdot \theta_t(h_{-i}^{t-1}, \cdot) + g_i(a_i, \bar{s}_{-i}^{\mathsf{B}}(h_{-i}^{t-1}))\right) + (1-\delta)U_{t+1}(h_{-i}^{t-1}, a_i) \\
&\quad = (1-\delta)U_{t+1}(h_{-i}^{t-1}, a_i^*(h_{-i}^{t-1})) + (1-\delta)\delta^{t-1} g_i(a_i^*(h_{-i}^{t-1}), \bar{s}_{-i}^{\mathsf{B}}(h_{-i}^{t-1})) \qquad (7)
\end{aligned}$$

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 (7) 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 $\lambda(0)$ and π jointly converge to 1 and perfect monitoring respectively.

Finally we define a strategy $r_i^{\mathsf{B}} \in S_i^T$ in the following way. Let $r_i^{\mathsf{B}}(h_i^{t-1})$ 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 $r_i^{\mathsf{B}}(h_i^{t-1})$ 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 $\bar{s}_{-i}^{\mathsf{B}}$. By the definition of r_i^{B}, the payoff in the T-times-repeated game without any transfers then approaches $\max_{s_i \in S_i^T} U_i^T(s_i, \bar{s}_{-i}^{\mathsf{B}})$; this implies that

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

for all $s_i \in S_i^T$.

A.2 An Additional Lemma

We follow the notation and terminology of HO2006. The next lemma defines the transfers ξ_i^{G} for the repeated game with rare lags; this provides analogs of the results of Lemmata 2 and 3 of HO2006.

Lemma A.1. *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 $\xi_i^{\mathsf{G}} : H_{-i}^T \to \mathbb{R}_-$ such that*

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

where $\bar{s}^{\mathsf{G}}_{-i} \mid H^R_{-i} = s^{\mathsf{G}}_{-i} \mid H^R_{-i}$ *and* $\bar{s}^{\mathsf{G}}_{-i} \mid H^E_{-i} = \bar{s}_{-i} \mid H^E_{-i}$*. Furthermore* $\xi^{\mathsf{G}}_i : H^T_{-i} \to \mathbb{R}_-$ *can be chosen so that, for every* $s_i \in B_i\left(\bar{s}^{\mathsf{G}}_{-i}, \xi^{\mathsf{G}}_i \mid \bar{s}_i\right)$*, we have*

$$\lim_{\varepsilon \to 0} U^A_i\left(s_i, \bar{s}^{\mathsf{G}}_{-i}, \xi^{\mathsf{G}}_i\right) = \min_{\tilde{s}_i \in \tilde{\mathcal{S}}_i} \tilde{U}^T_i(\tilde{s}_i, \bar{s}^{\mathsf{G}}_{-i}),$$

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

Proof. Let $\varepsilon > 0$ be such that π is ε-perfect and $\Pr(L > 0) < \varepsilon$. For every $\nu > 0$, observe that there exists ε/ρ small enough such that, for any history $h^{t-1}_i \in H^{R,t-1}_i$ and conditional on observing h^{t-1}_i, player i assigns probability at least $1 - \nu$ to the event that player $-i$ observe the corresponding history h^{t-1}_{-i}. Consider for some $h^{t-1}_i \in H^{R,t-1}_i$ and any action $a_i \in A_i$, the row vector consisting of the probabilities assigned by player i, conditional on history h^{t-1}_i and on action a_i taken by player i in period t, to the different equivalence classes of histories (h^{t-1}_{-i}, 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 $h^{t-1}_i \in H^{R,t-1}_i$ 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 $\theta(h^{t-1}_{-i}, \infty) = 0$ for any $h^{t-1}_{-i} \in H^{t-1}_{-i}$. This is possible since the number of rows is exactly the same as in HO2006 and the number of columns corresponding to (h^{t-1}_{-i}, a_i) for some $a_i \neq \infty$ is also the same as in HO2006. This proves the lemma. ∎

A.3 Concluding the Argument

The remainder of the proof of Theorem 4.3 follows along the same lines as in HO2006, defining $\bar{s}^{\mathsf{B}}_{-i} \mid H^E_{-i}$ and $\bar{s}^{\mathsf{G}}_{-i} \mid H^E_{-i}$, ξ^{G}_i and ξ^{B}_i as the fixed point of the relevant correspondence.[22] The construction works because of Lemma A.1 and the fact that play at periods $T, 2T, \ldots$, 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.

B Web Appendix

B.1 Proof of Lemma 5.8

Suppose that $W \subseteq B(W, \delta, \varepsilon)$. Then take any $v \in W$. We construct a strategy in the game $G^{\mathrm{pu}}(\delta, \varepsilon)$ that achieves the payoff profile v.

Note that due to the self-generation of W, there exists functions $\Sigma : W \to \prod_{i=1}^n \Delta(A_i)$ and $\Gamma : W \times Y \to W$ such that for every $v \in W$,

$$v = (1 - \delta) g(\Sigma(v)) + \delta(1 - \varepsilon^n) \sum_{y \in Y} \pi(y|\Sigma(v))\Gamma(v, y)$$

and

$$v_i \geq (1 - \delta) g_i(a_i, \Sigma(v)_{-i}) + \delta(1 - \varepsilon^n) \sum_{y \in Y} \pi(y|a_i, \Sigma(v)_{-i})\Gamma(v, y)$$

[22]See HO2006 for details.

for all $a_i \in A_i$ and all $i = 1, \ldots, n$.

We now define recursively sequences of maps $v^t : Y^t \to W$ and $\sigma^t : Y^t \to \prod_{i=1}^n \Delta(A_i)$ for all $t = 0, 1, 2, \ldots$. Begin by setting $v^0 = v$ and $\sigma^0 = \Sigma(v)$. Suppose that $v^t : Y^t \to W$ and $\sigma^t : Y^t \to W$. We define $v^{t+1} : Y^{t+1} \to W$ by

$$v^{t+1}(y^0, \ldots, y^t) = \Gamma(v^t(y^0, \ldots, y^t), y^t)$$

and

$$\sigma^{t+1}(y^0, \ldots, y^t) = \Sigma(v^{t+1}(y^0, \ldots, y^t)).$$

Then consider the strategy profile σ that is the collection $\{\sigma^t\}_{t=0}^{\infty}$. To see that this strategy generates the payoff profile v,

$$\mathbb{E}_\sigma \left[(1-\delta) \sum_{t=0}^{T} \delta^t (1-\varepsilon^n)^t g(a^t) \right] = v - \delta^{T+1}(1-\varepsilon^n)^{T+1} \mathbb{E}_\sigma \left[v^{T+1}(y^0, \ldots, y^T) \right].$$

It is clear that the left hand side converges to $\mathbb{E}_\sigma \left[(1-\delta) \sum_{t=0}^{\infty} \delta^t (1-\varepsilon^n)^t g(a^t) \right]$ while the right hand side of the equation above converges to v since v^{T+1} is a function that maps into a compact set W. Thus we have shown that

$$v = \mathbb{E}_\sigma \left[(1-\delta) \sum_{t=0}^{\infty} \delta^t (1-\varepsilon^n)^t g(a^t) \right].$$

Moreover by the same argument, the continuation payoff at any public history $(y^0, \ldots, y^t)$ induced by σ is also equal to $v^{t+1}(y^0, \ldots, y^t)$. Then the optimality of σ_i against σ_{-i} follows immediately by construction due to the one-stage deviation principle.

References

Abreu, D., P. Milgrom, and D. Pearce (1990): "Information and Timing in Repeated Partnerships," *Econometrica*, 59(6), 1713–1733.

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1064.

Ahn, T. K., E. Ostrom, D. Schmidt, and J. Walker (2003): "Trust in two-person games: Game structures and linkages," in *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, ed. by E. Ostrom, and J. Walker, pp. 323–351. Russell Sage, New York.

Aumann, R. J., and L. S. Shapley (1976): "Long-Term Competition—A Game Theoretic Analysis," in *Essays in Game Theory in Honor of Michael Maschler*, ed. by N. Megiddo, pp. 1–15. Springer-Verlag, New York.

Ben-Porath, E., and M. Kahneman (2003): "Communication in repeated games with costly monitoring," *Games and Economic Behavior*, 44(2), 227–250.

Bergin, J., and W. B. MacLeod (1993): "Continuous Time Repeated Games," *International Economic Review*, 34, 21–37.

Bhaskar, V., and I. Obara (2011): "Reputation and Information Lag," mimeo.

Compte, O. (1998): "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66(3), 597–626.

Dutta, P. (1995): "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66(1), 1–32.

Ellison, G. (1994): "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61, 567–588.

Escobar, J., and J. Toikka (2011): "A Folk Theorem with Markovian Private Information," mimeo.

Faingold, E., and Y. Sannikov (2011): "Reputation in Continuous-Time Games," *Econometrica*, 79(3), 773–876.

Friedman, J. W. (1971): "A Noncooperative Equilibrium for Supergames," *Review of Economic Studies*, 38(1), 1–12.

Fudenberg, D., and D. Levine (2007a): "Continuous time limits of repeated games with imperfect public monitoring," *Review of Economic Dynamics*, 10, 173–192.

Fudenberg, D., and D. Levine (2007b): "The Nash-threats folk theorem with communication and approximate common knowledge in two player games," *Journal of Economic Theory*, 132(1), 461–473.

Fudenberg, D., and D. Levine (2009): "Repeated Games with Frequent Signals," *Quarterly Journal of Economics*, 124(1), 233–265.

Fudenberg, D., D. K. Levine, and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5), 997–1031.

Fudenberg, D., and E. Maskin (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54(3), 533–554.

Fudenberg, D., and W. Olszewski (2011): "Repeated games with asynchronous monitoring of an imperfect signal," *Games and Economic Behaivor*, 72, 86–99.

Fudenberg, D., and Y. Yamamoto (2011): "The folk theorem for irreducible stochastic games with imperfect public monitoring," *Journal of Economic Theory*, 146, 1664–1683.

Gachter, S., B. Herrmann, and C. Thoni (2004): "Trust, voluntary cooperation, and socio-economic background: Survey and experimental evidence," *Journal of Economic Behavior & Organization*, 55, 505–531.

Hörner, J., and W. Olszewski (2006): "The Folk Theorem for Games with Private Almost-Perfect Monitoring," *Econometrica*, 74(6), 1499–1544.

Hörner, J., and W. Olszewski (2009): "How Robust is the Folk Theorem?," *Quarterly Journal of Economics*, 124(4), 1773–1814.

Hörner, J., T. Sugaya, S. Takahashi, and N. Vieille (2011): "Recursive methods in discounted stochastic games: An algorithm for δ? 1 and a folk theorem," *Econometrica*, 79, 1277–1318.

Kandori, M., and H. Matsushima (1998): "Private Observation, Communication and Collusion," *Econometrica*, 66(3), 627–652.

Levy, Y. (2009): "Stochastic Games with Information Lag," mimeo.

Mailath, G. J., and S. Morris (2002): "Repeated Games with Almost-Public Monitoring," *Journal of Economic theory*, 102(1), 189–228.

Rubinstein, A. (1994): "Equilibrium in Supergames," in *Essays in Game Theory in Honor of Michael Maschler*, ed. by N. Megiddo, pp. 17–28. Springer-Verlag, New York, 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., and A. Skrzypacz (2007): "Impossibility of collusion under imperfect monitoring with flexible production," *American Economic Review*, 97, 1794–1823.

——— (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(2), 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(2), 802–824.

Yao, T., S. H. Xu, and B. Jiang (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(3), 1197–1220.