

Federal Reserve Bank of Minneapolis
Research Department Staff Report 383

November 2006

Private Monitoring with Infinite Histories*

Christopher Phelan

Federal Reserve Bank of Minneapolis
and University of Minnesota

Andrzej Skrzypacz

Graduate School of Business
Stanford University

ABSTRACT

This paper develops new recursive methods for studying stationary sequential equilibria in games with private monitoring. We first consider games where play has occurred forever into the past and develop methods for analyzing a large class of stationary strategies, where the main restriction is that the strategy can be represented as a finite automaton. For a subset of this class, strategies which depend only on the players' signals in the last k periods, these methods allow the construction of *all* pure strategy equilibria. We then show that each sequential equilibrium in a game with infinite histories defines a correlated equilibrium for a game with a start date and derive simple necessary and sufficient conditions for determining if an *arbitrary* correlation device yields a correlated equilibrium. This allows, for games with a start date, the construction of all pure strategy *sequential* equilibria in this subclass.

*The authors thank Peter DeMarzo, Larry Jones, Narayana Kocherlakota, David Levine, George Mailath, Stephen Morris, Ichiro Obara, Larry Samuelson, seminar participants at the Federal Reserve Bank of Minneapolis, Stanford University, and the 2006 meetings of the Society for Economic Dynamics for helpful comments as well as the excellent research assistance of Roozbeh Hosseini. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

In this paper, we propose new recursive methods for studying repeated games with private monitoring. Our contribution is twofold. We present a model of repeated games with infinite histories — that is, games in which time extends infinitely backward and forward — and establish new set-based methods for verifying the equilibrium conditions for a large class of stationary strategies. For a subset of this class, strategies which depend only on the players' signals in the last k periods, these methods allow the construction of *all* pure strategy equilibria. We also establish a connection between equilibria of games with infinite histories and correlated stationary equilibria of traditional games with a start date and derive simple necessary and sufficient conditions for determining if an *arbitrary* correlation device yields a correlated equilibrium. This allows, for games with a start date, the construction of all pure strategy *sequential* equilibria in this subclass.

Games with infinite histories are interesting in their own right. Our model lets us abstract away from the inherent non-stationarity of the set of possible histories that plagues games with a start date. A long tradition in economic theory is to use the repeated game model to study the dynamic provision of incentives because in the case of perfect or public monitoring, such models offer a highly tractable stationary environment without end-of-horizon effects. However, with private monitoring, the stationarity is broken by beginning-of-horizon effects: the set of possible private histories changes through time and with it the possibilities to coordinate play among players. Our model of repeated games with infinite histories solves this problem.

Our formulation is simple. A joint strategy is a mapping from each player's privately observed infinite history to how he plays today. An equilibrium strategy is one where each

player's mapping, given what he can infer regarding what his opponents have observed, is a best response to the other players' strategies.

We pay particular attention to strategies which can be represented as finite automata, that is, where private histories are grouped into a finite number of private states and a player's action depends only on which private state he is in. For this class, we develop recursive methods on *sets* of beliefs for each player about the private state of each of his opponents. The key is that if all players' strategies are finite automata, a particular player's private history is relevant only to the extent that it gives him information regarding the private states of his opponents (a point first made by Mailath and Morris (2002)). This lets us summarize a player's history as a *belief* over a finite state space, a much smaller object (a point also made by Mailath and Morris (2002)). The advantages of working with *sets* of beliefs are two. One is that it is necessary and sufficient to check incentives only for those beliefs lying on the boundary of the set of beliefs which can be generated by infinite histories (Lemma 1). The other advantage of working with sets of beliefs is that these sets can be readily calculated using recursive methods which look similar to those developed by Abreu, Pearce, and Stacchetti (1990) but which operate on sets of beliefs rather than sets of continuation values (Lemma 4).

We also connect our results on games with infinite histories to the more traditional (and more problematic) formulation of games with a start date. In this part of the paper, we first show that any equilibrium of the game with infinite histories can be used to construct a correlated equilibrium of the corresponding game with a start date, where the correlation device sends each player a fictitious private history. We then develop methods for verifying whether an arbitrary correlation device signaling initial private states, when coupled with

finite state equilibrium strategy of the game with infinite histories, forms a correlated equilibrium. Since sequential equilibria are examples of correlated equilibria (with degenerate signaling devices), and since strategies which depend only on the signals a player has seen in the last k periods are finite state strategies, these methods deliver necessary and sufficient conditions for constructing all k -history dependent pure strategy sequential equilibria.

Finding equilibria in repeated games with private monitoring is known to be difficult. See, for example, the work of Kandori (2002) and Mailath and Samuelson (2006), Chap. 12. Among several difficulties, a central one is that with private monitoring, the recursive structure of public monitoring games is lost. The continuation of (sequential) equilibrium play in a game with private monitoring is not a sequential equilibrium, but rather a correlated equilibrium in which private histories function as the correlation device. But as Kandori (2002) notes, the correlation device becomes increasingly more complex over time. By introducing infinite histories, we make the correlation device stationary and regain some tractability. The mapping to correlated equilibria of games with a start date is then natural. In fact, using randomization or exogenous correlation in period 0 of the game to make it more stationary (and create interior beliefs about the private history of other players) has been suggested by Sekiguchi (1997), Compte (2002), Ely (2002), and Cripps, Mailath, and Samuelson (forthcoming). We present a robust way of applying this method to construct a family of equilibria.

Our results complement the existing literature on the construction of belief-free equilibria (for example, the work of Ely and Välimäki (2002), Piccione (2002), Ely, Hörner and Olszewski (2005), and Kandori and Obara (2006)), in which players use mixed strategies and their best responses are independent of their beliefs about the private histories of their opponents. In contrast to belief-free equilibria, the equilibria we construct are belief-dependent;

players' best responses do depend on their beliefs. Our work is also related to that of Mailath and Morris (2002), who show robustness of finite-history dependent strategies in games with almost-public monitoring (and in Mailath and Morris (2006) show problems with robustness of infinite-history dependent strategies). Mailath and Morris show that for a strict k -history dependent equilibrium of a game with public monitoring, if the game is slightly perturbed to almost-public monitoring, the strategies still form an equilibrium. We can apply our results to this class of games and calculate exactly how much correlation is necessary. They provide a sufficient bound on the correlation, but leave finding a necessary and sufficient cutoff an open question.

2. Games with Infinite Histories

Consider a stage game, Γ , with N players, $i = 1, \dots, N$, each able to take actions $a_i \in A_i$. Assume that with probability $P(y|a)$, a vector of private signals $y = (y_1, \dots, y_N)$ (each $y_i \in Y_i$) is observed conditional on the vector of private actions $a = (a_1, \dots, a_N)$, where for all (a, y) , $P(y|a) > 0$ (*full support*). Further assume that $A = A_1 \times \dots \times A_N$ and $Y = Y_1 \times \dots \times Y_N$ are both finite sets.

The current period payoff to player i is denoted $u_i(a_i, y_i)$. That is, player i 's payoff is a function of his own current-period action and private signal. Players put weight $1 - \beta$ on current utility and weight β on future payoffs and, as usual, care about the expected value of utility streams.

We start by considering the case (denoted $\Gamma^{-\infty, \infty}$) in which time is assumed to extend infinitely both backward and forward. (Histories are infinite.) In Section 5, we connect our results to the more traditional case where time is assumed to start at date $t = 0$ and

extend forward only (denoted $\Gamma^{0,\infty}$). To this end, let $\mathbf{a}_i$ denote an infinite history of player i 's private actions $\mathbf{a}_i = \{a_{i,1}, a_{i,2}, \dots\}$ with the set of possible private action histories $\mathbf{a}_i$ denoted by $\mathbf{A}_i$. (All history sequences move backward in time with the time subscript referring to the number of periods before the present period.) Likewise, let player i 's private signal history be denoted $\mathbf{y}_i = \{y_{i,1}, y_{i,2}, \dots\}$, where $\mathbf{Y}_i$ denotes the set of possible private signal histories for player i . Finally, let $\mathbf{A} = \mathbf{A}_0 \times \dots \times \mathbf{A}_N$, $\mathbf{Y} = \mathbf{Y}_0 \times \dots \times \mathbf{Y}_N$, $\mathbf{Z} = \mathbf{A} \times \mathbf{Y}$, $\mathbf{Z}_i = \mathbf{A}_i \times \mathbf{Y}_i$ and $\mathbf{Z}_{-i} = \mathbf{A}_{-i} \times \mathbf{Y}_{-i}$, where the subscript $-i$ refers to the set $\{1, 2, \dots, i-1, i+1, \dots, N\}$. In words, $\mathbf{z}_i \in \mathbf{Z}_i$ is what player i has directly observed and $\mathbf{z}_{-i} \in \mathbf{Z}_{-i}$ is what player i has not directly observed, but his opponents have.

For player i , let a mixed strategy $\sigma_i(a_i|\mathbf{z}_i) : A \times \mathbf{Z}_i \rightarrow [0, 1]$ map infinite private histories to the probability of taking a given action. Denote a joint strategy as $\sigma = (\sigma_1, \dots, \sigma_N)$. This formulation implicitly restricts strategies to not depend on the calendar date (a concept we have, in fact, not introduced).

A. Stationarity:

Let $\pi : \mathcal{B}(\mathbf{Z}) \rightarrow [0, 1]$ be a probability measure over infinite histories where $\mathcal{B}(\mathbf{Z})$ denotes the Borel subsets of $\mathbf{Z}$. Let $G(\sigma, \pi) : \mathcal{B}(\mathbf{Z}) \rightarrow [0, 1]$ be the probability measure over infinite histories induced by π and the addition of one more period through the strategy σ and the function $P(y|a)$. That is, for all $\hat{S} \in \mathcal{B}(\mathbf{Z})$

$$G(\sigma, \pi)(\hat{S}) = \int_{\mathbf{Z}} \sum_a \left(\Pi_j \sigma_j(a_j|\mathbf{z}_j) \right) \sum_y P(y|a) I\left(\left((a, y), \mathbf{z}\right) \in \hat{S}\right) d\pi(\mathbf{z}),$$

where I is the indicator function. A pair (σ, π) is said to be *stationary* if $G(\sigma, \pi) = \pi$.

B. Beliefs:

With full support all infinite histories occur with zero probability. So to define conditional beliefs we consider sets of histories.

For sets of infinite histories $\hat{\mathbf{Z}}_i \in \mathcal{B}(\mathbf{Z}_i)$ such that $\pi(\hat{\mathbf{Z}}_i \times \mathbf{Z}_{-i}) > 0$, one can define player i 's conditional probability measure $\pi_i : \mathcal{B}(\mathbf{Z}_i) \times \mathcal{B}(\mathbf{Z}_{-i}) \rightarrow [0, 1]$ as

$$(1) \quad \pi_i(\hat{\mathbf{Z}}_{-i} | \hat{\mathbf{Z}}_i) = \pi(\hat{\mathbf{Z}}_i \times \hat{\mathbf{Z}}_{-i}) / \pi(\hat{\mathbf{Z}}_i \times \mathbf{Z}_{-i}).$$

This uniquely defines beliefs π -almost everywhere. However, for sets of infinite histories where player i has deviated, this formula has no implications since the measure π puts zero measure on the entire set of infinite private histories which can only be reached through player i 's own deviation.

In games with a start date and full support, player i 's beliefs can still be uniquely defined for histories where he has deviated. The full support assumption ensures that player i will never see evidence that his opponents have deviated (and thus will assume they haven't). Beliefs regarding what his opponents have seen after t dates can then be defined for *arbitrary* (on or off path) specifications of player i 's actions and signals using the strategy σ , the signal function P , and Bayes' rule. But this logic does not extend (as far as we can show) to games with infinite histories where player i has deviated an infinite number of times.

To see this, note that if player i deviated only in the most recent period, one can define beliefs conditional on player i 's off path actions as follows: Let (a_i, y_i) be an arbitrary action and signal representing player i 's most recent action and signal and let $\hat{\mathbf{Z}}_i$ be a set of infinite histories where player i has not deviated, representing his actions and signals for all earlier periods. In this case, the probability of $(a_{-i}, y_{-i}) \times \hat{\mathbf{Z}}_{-i}$ conditional on $(a_i, y_i) \times \hat{\mathbf{Z}}_i$ can be

defined as

$$\pi_i((a_{-i}, y_{-i}) \times \hat{\mathbf{Z}}_{-i} | (a_i, y_i) \times \hat{\mathbf{Z}}_i) = \int_{\hat{\mathbf{Z}}_{-i}} \frac{(\Pi_{-i} \sigma_{-i}(a_{-i} | \mathbf{z}_{-i})) P(y_i, y_{-i} | a_i, a_{-i})}{\sum_{\hat{a}_i, \hat{y}_i} (\Pi_{-i} \sigma_{-i}(\hat{a}_{-i} | \mathbf{z}_{-i})) P(y_i, \hat{y}_{-i} | a_i, \hat{a}_{-i})} d\pi_i(\mathbf{z}_{-i} | \mathbf{z}_i).$$

This logic generalizes to all histories where player i has deviated only a finite number of times, and thus there exists a date before which he has not deviated. But this logic does not generalize to infinite histories where player i has deviated an infinite number of times since one cannot pin down “initial” beliefs regarding the actions of the other players.

To handle these cases we consider trembling. That is, if $\sigma_i : A \times \mathbf{Z}_i \rightarrow (0, 1)$, all actions by player i are on path and conditional probabilities can be defined as in equation (1). Given this, we say π_i is a *valid conditional belief measure* if there exists a sequence $\{(\sigma_s, \pi_s, \pi_{i,s})\}_{s=0}^{\infty}$ such that $\sigma_{i,s} : A \times \mathbf{Z}_i \rightarrow (0, 1)$, $\sigma_s \rightarrow \sigma$ in the sup norm and $\pi_s \rightarrow \pi$ and $\pi_{i,s} \rightarrow \pi_i$ in the weak-* topology.

C. Equilibrium Definition

Let $v_i(\mathbf{z} | \sigma)$ be the function recursively defined by

$$v_i(\mathbf{z} | \sigma) = \sum_a \left(\Pi_j \sigma_j(a_j | \mathbf{z}_j) \right) \left(\sum_y P(y | a) [(1 - \beta) u_i(a_i, y_i) + \beta v_i((a, y), \mathbf{z} | \sigma)] \right)$$

and $Ev_i(\mathbf{z}_i | \sigma) = \int_{\mathbf{z}_{-i}} v_i((\mathbf{z}_i, \mathbf{z}_{-i}) | \sigma) d\pi_i(\mathbf{z}_{-i} | \mathbf{z}_i)$.

A triplet (σ, π, π_i) is a *Stationary Sequential Equilibrium* of the game with infinite histories $\Gamma^{-\infty, \infty}$ (SSE-ih) if

1. For all i , $\mathbf{z}$, and $\hat{\sigma}_i : \mathbf{z}_i \rightarrow [0, 1]$, $Ev_i(\mathbf{z}_i | \sigma) \geq Ev_i(\mathbf{z}_i | (\hat{\sigma}_i, \sigma_{-i}))$.
2. $G(\sigma, \pi) = \pi$.

3. For all i , π_i is a valid conditional belief measure.

3. Finite State Strategies

This section develops methods for analyzing a particular (but large) class of strategies. First, we assume that each player's strategy can be described as a finite automaton, and second, that along the path of play *information depreciates* (an assumption formally made below). We show how to verify whether or not a particular strategy profile of this type is an SSE-ih. Given this, we show how to calculate *all* pure strategy equilibria of an important subset of these strategies: those which depend only on the last k periods of a player's history.

Player i 's strategy σ_i can be described as a finite automaton if and only if his set of infinite private histories $\mathbf{Z}_i$ can be divided into a finite partition Ω_i such that σ_i does not distinguish between histories $\mathbf{z}_i$ and $\hat{\mathbf{z}}_i$ if both histories are in the same cell $\omega_i \in \Omega_i$. (That is, if $\mathbf{z}_i \in \omega_i$ and $\hat{\mathbf{z}}_i \in \omega_i$, then $\sigma_i(a_i | (z_{i,1}, \dots, z_{i,s}), \mathbf{z}_i) = \sigma_i(a_i | (z_{i,1}, \dots, z_{i,s}), \hat{\mathbf{z}}_i)$ for all a_i and finite histories $(z_{i,1}, \dots, z_{i,s})$.) Let D_i denote the number of elements of Ω_i and D_{-i} denote the number of elements of Ω_{-i} , the set of private states for the other players. Instead of explicitly writing strategies and payoffs as functions of histories, we (equivalently) express them as functions of this induced state space.¹

Since a player does not know the aggregate state ω , but instead knows only his own part of it, ω_i , he will use all the information available to him (his private *infinite* history $\mathbf{z}_i$) to form beliefs regarding $\omega_{-i} \in \Omega_{-i}$. For a particular infinite private history, a player's beliefs over the aggregate state of the game are simply a point in the $(D_{-i} - 1)$ -dimensional unit-simplex, denoted $\Delta^{D_{-i}}$. Let $\mu_i(\mathbf{z}_i) : \mathbf{Z}_i \rightarrow \Delta^{D_{-i}}$ denote player i 's beliefs about ω_{-i} after

¹ For a useful discussion of the validity of representing strategies as finite state automata in the context of games with private monitoring, see Mailath and Morris (2002) and Mailath and Samuelson (2006).

private history $\mathbf{z}_i$. Let $\mu_i(\omega_{-i}|\mathbf{z}_i)$ denote the probability assigned to the particular state ω_{-i} , where $\mu_i(\omega_{-i}|\mathbf{z}_i) = \int_{\mathbf{z}_{-i} \in \omega_{-i}} d\pi_i(\mathbf{z}_{-i}|\mathbf{z}_i)$.

For a strategy σ and corresponding beliefs π_i , *information is said to depreciate* if for all $\epsilon > 0$, i , ω_{-i} , and infinite private histories $\mathbf{z}_i$, $\mathbf{z}_i^0$, and $\mathbf{z}_i^1$, there exists an s such that $|\mu_i(\omega_{-i}|(z_{i,1}, \dots, z_{i,s}, \mathbf{z}_i^0)) - \mu_i(\omega_{-i}|(z_{i,1}, \dots, z_{i,s}, \mathbf{z}_i^1))| < \epsilon$ (uniform convergence of beliefs). This condition can be verified as follows: Let $X(z_i)$ specify the probability of player i 's opponents transiting from state $\bar{\omega}_{-i}$ last period to state ω_{-i} this period conditional on player i observing z_i today and his opponents following strategy σ . If there exists an s such that for all $(z_{i,1}, \dots, z_{i,s})$, the matrix $X(z_{i,1}) \dots X(z_{i,s})$ contains no zeros, then information depreciates.²

This condition is immediately satisfied if player i 's strategy σ_i (for all i) depends only on the last k periods of y_i (as opposed to depending also on his action realizations). An example of a strategy where information does not depreciate would be if player i plays whatever action he played in the previous period.

Given a strategy σ and its corresponding state space Ω , we can define expected payoffs in aggregate state ω as:

$$V_i(\omega|\sigma) = \sum_a \left(\Pi_j \sigma_j(a_j|\omega_j) \right) \left(\sum_y P(y|a) [(1 - \beta)u_i(a_i, y_i) + \beta V_i(\omega^+(a, y, \omega) | \sigma)] \right)$$

where $\omega^+(a, y, \omega)$ is the next-period state given current-period state ω and new (a, y) . That is, $V_i(\omega|\sigma)$ is the expected utility of player i if he knows the aggregate state (including the private states of the other players) and he and all the other players mechanistically play the

²If player i 's beliefs were μ_i or $\hat{\mu}_i$ s periods ago, beliefs today after $(z_{i,1}, \dots, z_{i,s})$ are, respectively, $\mu_i X(z_{i,s}) \dots X(z_{i,1})$ and $\hat{\mu}_i X(z_{i,s}) \dots X(z_{i,1})$. If the differing beliefs are pushed t periods back in time and an extra t periods of data are added for dates $s + 1$ to $s + t$, player i 's beliefs become, respectively, $\mu_i (X(z_{i,s+t}) \dots X(z_{i,s+1})) X(z_{i,s}) \dots X(z_{i,1})$ and $\hat{\mu}_i (X(z_{i,s+t}) \dots X(z_{i,s+1})) X(z_{i,s}) \dots X(z_{i,1})$. Premultiplying the transition matrix $X(z_{i,s}) \dots X(z_{i,1})$ by a matrix with no zeros uniformly decreases the distance between the rows.

strategy σ .

For arbitrary beliefs $m_i \in \Delta^{D-i}$ and a strategy σ , let

$$EV_i(\omega_i, m_i | \sigma) = \sum_{\omega_{-i}} m_i(\omega_{-i}) V_i(\omega_i, \omega_{-i} | \sigma).$$

We have now defined expected payoffs as functions of private states ω_i and beliefs over the private states of the other players m_i , instead of expected payoffs being functions of a player's infinite private history $\mathbf{z}_i$.

Rather than considering separately the beliefs $m_i \in \Delta^{D-i}$ that a player will have after some infinite history, it is useful to consider subsets of beliefs. Let $M_i^*(\omega_i) = \text{co}(\{m | m = \mu_i(\mathbf{z}_i) \text{ for some } \mathbf{z}_i \in \omega_i\})$ (where $\text{co}()$ denotes the closure of the convex hull). Here, $M_i^*(\omega_i)$ is the closure of the convex hull of the set of possible beliefs player i can have about ω_{-i} given that he's seen $\mathbf{z}_i \in \omega_i$, where a belief is possible if there is an infinite private history which induces this belief. Let M_i^* denote the collection of D_i of these subsets, one for each ω_i .

The following lemma establishes that to check the incentives of a finite state strategy, one need only check that for each player i and private state ω_i , the player does not wish to deviate when his beliefs about the other players' private states ω_{-i} are on the frontier of the convex hull of all possible beliefs $M_i^*(\omega_i)$. That is, it is not necessary to check incentives for every infinite history.

LEMMA 1. Consider a finite state strategy σ and measures (π, π_i) satisfying conditions 2 and 3 of the definition of an SSE-ih and the implied sets $M_i^*(\omega_i)$. Then (σ, π, π_i) is an SSE-ih if and only if $EV_i(\omega_i(\mathbf{z}_i), m_i | \sigma) \geq \sum_{\omega_{-i}} m_i(\omega_{-i}) v_i(\mathbf{z}_i, \mathbf{z}_{-i} \in \omega_{-i} | \hat{\sigma}_i, \sigma_{-i})$ for all i , $\mathbf{z}_i$, $\hat{\sigma}_i : A \times \mathbf{Z}_i \rightarrow [0, 1]$ and all m_i such that m_i is on the frontier of $M_i^*(\omega_i)$.

Proof. If: Suppose incentives hold for beliefs m_i and $\hat{m}_i$ each in Δ^{D-i} . (That is, $EV_i(\omega_i, m_i|\sigma) \geq \sum_{\omega_{-i}} m_i(\omega_{-i})v_i(\mathbf{z}_i, \mathbf{z}_{-i} \in \omega_{-i}|\hat{\sigma}_i, \sigma_{-i})$ and $EV_i(\omega_i, \hat{m}_i|\sigma) \geq \sum_{\omega_{-i}} \hat{m}_i(\omega_{-i})v_i(\mathbf{z}_i, \mathbf{z}_{-i} \in \omega_{-i}|\hat{\sigma}_i, \sigma_{-i})$ for all i , ω_i and $\hat{\sigma}_i : A \times \mathbf{Z}_i \rightarrow [0, 1]$.) Then since expected utility for any $\hat{\sigma}_i$ (including the on-path strategy σ_i) is linear in these beliefs, for all $\alpha \in (0, 1)$, incentives hold for beliefs $\alpha m_i + (1 - \alpha)\hat{m}_i$. Next, every infinite history generates beliefs within $M_i^*(\omega_i)$ by construction. Further, holding ω_i constant, player i 's history is relevant to him only to the extent that it affects his beliefs regarding the other players' continuation play which is determined by ω_{-i} . Thus if incentives hold for the beliefs generated by infinite history $\mathbf{z}_i$, incentives hold for infinite history $\mathbf{z}_i$. Finally, the beliefs generated by any infinite history can be constructed as the convex combination of points on the frontier of $M_i^*(\omega_i)$.

Only if: Suppose (σ, π, π_i) satisfy condition 1 of the definition of an SSE-ih, but there exists a belief m_i on the frontier of $M_i^*(\omega_i)$ and a deviation strategy $\hat{\sigma}_i : A \times \mathbf{Z}_i \rightarrow [0, 1]$ such that $EV_i(\omega_i, m_i|\sigma) < \sum_{\omega_{-i}} m_i(\omega_{-i})v_i(\mathbf{z}_i, \mathbf{z}_{-i} \in \omega_{-i}|\hat{\sigma}_i, \sigma_{-i})$. If m_i is the linear combination of beliefs generated by private histories $\mathbf{z}_i^0 \in \omega_i$ and $\mathbf{z}_i^1 \in \omega_i$, then this is a contradiction since, as shown above, if incentives hold for a set of beliefs, they hold for all linear combinations of those beliefs. But since $M_i^*(\omega_i)$ is the closure of the convex hull of beliefs generated by private histories, the only remaining possibility is that m_i is within the closure of the convex hull of beliefs generated by histories, but not the convex hull itself. But if incentives do not hold for a given belief, then, for a given deviation strategy, the gain to deviation is strictly positive. This implies deviation is preferred for some neighborhood around this belief as well, contradicting the supposition. ■

Lemma 1 is useful for checking whether or not a particular finite state strategy is an equilibrium when the sets of beliefs generated by that strategy (for each player i) M_i^* are

known. In Lemmas 2 through 4, we construct a method for generating, for each finite state strategy, the appropriate M_i^* sets.

We begin by constructing an operator from sets of beliefs to sets of beliefs with the property that the operator's largest fixed point is M_i^* . Let M_i (without the $*$) denote an arbitrary collection of D_i closed convex subsets of Δ^{D-i} , and let the one-step operator $T(M_i)$ be defined as follows: First, let $F(\omega_i)$ be the set of private states last period consistent with the private state this period being ω_i . (That is, $F(\omega_i) = \{\bar{\omega}_i \mid \text{there exists } \mathbf{z}_i \in \bar{\omega}_i \text{ and } z_i \in A_i \times Y_i \text{ such that } (z_i, \mathbf{z}_i) \in \omega_i\}$.) Likewise, let $G(\bar{\omega}_i, \omega_i) = \{z_i = (a_i, y_i) \mid \text{if } \mathbf{z}_i \in \bar{\omega}_i, \text{ then } (z_i, \mathbf{z}_i) \in \omega_i\}$. That is, $G(\bar{\omega}_i, \omega_i)$ is the set of (a_i, y_i) pairs such that player i 's private state transits from $\bar{\omega}_i$ to ω_i . The successor of belief $m_i \in \Delta_{-i}^D$ given new data (a_i, y_i) (denoted $m'_i(m_i, a_i, y_i) \in \Delta_{-i}^D$) is determined by Bayes' rule as

$$m'_i(m_i, a_i, y_i)(\omega_{-i}) = \sum_{\bar{\omega}_{-i} \in F(\omega_{-i})} \sum_{(a_{-i}, y_{-i}) \in G(\bar{\omega}_{-i}, \omega_{-i})} m_i(\bar{\omega}_{-i}) \frac{\sigma_{-i}(a_{-i} | \bar{\omega}_{-i}) P(y_i, y_{-i} | a_i, a_{-i})}{\sum_{\hat{a}_{-i}, \hat{y}_{-i}} \sigma_{-i}(\hat{a}_{-i} | \omega_{-i}) P(y_i, \hat{y}_{-i} | a_i, \hat{a}_{-i})}.$$

In the above formula, the ratio is the probability of the event (a_{-i}, y_{-i}) conditional on (a_i, y_i) . This probability is then summed over all (a_{-i}, y_{-i}) realizations consistent with the other player(s) moving from state $\bar{\omega}_{-i}$ to ω_{-i} , averaged over the probability of the other player(s) being in state $\bar{\omega}_{-i}$ according to beliefs m_i . Then

$$\begin{aligned} T(M_i)(\omega_i) &= \{m_i \mid \text{there exists } \bar{\omega}_i \in F(\omega_i), \bar{m}_i \in M_i(\bar{\omega}_i) \text{ and } (a_i, y_i) \in G(\bar{\omega}_i, \omega_i) \\ &\quad \text{such that } m_i = m'_i(\bar{m}_i, a_i, y_i)\}. \end{aligned}$$

The T operator works as follows: Suppose one takes as given the beliefs of player i over the private state of the other players, $\bar{\omega}_{-i}$, last period. Bayes' rule then implies what player i

should believe about ω_{-i} this period for each realization of (a_i, y_i) . If there exists a way to choose player i 's state last period, the beliefs of player i over the private states of his opponents last period, and a new realization of (a_i, y_i) such that Bayes' rule delivers beliefs m_i , then $m_i \in T(M_i)(\omega_i)$. In effect, the T operator gives, for a particular collection of belief sets M_i , the belief sets associated with all possible successor beliefs generated by new data and interpreted through σ .

Lemmas 2 through 4 show that the T operator can then be used to generate the true sets of valid beliefs M_i^* . We write $M_i \subset \hat{M}_i$ if $M_i(\omega_i) \subset \hat{M}_i(\omega_i)$ for all ω_i .

LEMMA 2. If $M_i^* \subset M_i$, then $M_i^* \subset T(M_i)$.

Proof. For a given ω_i choose beliefs $m_i \in M_i^*(\omega_i)$ such that m_i is the linear combination of beliefs m_i^0 and m_i^1 for which there exist infinite histories $\mathbf{z}_i^0$ and $\mathbf{z}_i^1$ which generate beliefs m_i^0 and m_i^1 . Now consider these histories except for the last period. That is, let $\hat{\mathbf{z}}_i^0 = \{z_{i,2}^0, \dots\}$ and $\hat{\mathbf{z}}_i^1 = \{z_{i,2}^1, \dots\}$. Beliefs after histories $\hat{\mathbf{z}}_i^0$ and $\hat{\mathbf{z}}_i^1$ (call them $\hat{m}_i^0$ and $\hat{m}_i^1$) satisfy $\hat{m}_i^0 \in M_i^*(\omega_i^0)$ and $\hat{m}_i^1 \in M_i^*(\omega_i^1)$, where $\hat{\mathbf{z}}_i^0 \in \omega_i^0$ and $\hat{\mathbf{z}}_i^1 \in \omega_i^1$ from the definition of M_i^* . Since $\hat{m}_i^0 \in M_i(\omega_i^0)$ and $\hat{m}_i^1 \in M_i(\omega_i^1)$ from $M_i^* \subset M_i$, $m_i^0 = m'_i(\hat{m}_i^0, a_i, y_i) \in T(M_i)(\omega_i)$ and $m_i^1 = m'_i(\hat{m}_i^1, a_i, y_i) \in T(M_i)(\omega_i)$. Since T maps closed convex sets to closed convex sets (from the linearity of our Bayes' rule operator in beliefs), $m_i \in T(M_i)(\omega_i)$. This leaves only the possibility that $m_i \in M_i^*(\omega_i)$ is not a linear combination of beliefs generated by infinite histories, which, from the definition of M_i^* implies m_i is on the frontier of $M_i^*(\omega_i)$. Suppose then $m_i \notin T(M_i)(\omega_i)$. Since the above logic can be applied to a sequence of points in $M_i^*(\omega_i)$ converging to m_i , each point in this sequence is in $T(M_i^*)(\omega_i)$, implying $T(M_i^*)(\omega_i)$ is an open set, a contradiction. ■

LEMMA 3. If $M_i \subset T(M_i)$, then $T(M_i) \subset M_i^*$.

Proof. For a given ω_i , choose beliefs $m_i \in T(M_i)(\omega_i)$. Since $m_i \in T(M_i)(\omega_i)$ there exists a private realization $z_i = (a_i, y_i)$ and beliefs $m_i^1 \in M_i(\omega_i^1)$ (where $\omega_i^1 \in F(\omega_i)$) such that $m_i = m'(m_i^1, a_i, y_i)$. That $M_i \subset T(M_i)$ ensures that $m_i^1 \in T(M_i)(\omega_i^1)$, thus this can be repeated indefinitely, generating any finite length history of private outcomes $\{z_{i,1}, \dots, z_{i,s}\}$ and beliefs $m_i^s \in M_i(\hat{\omega}_i^s)$, with the property that beliefs $m_i \in T(M_i)(\omega_i)$ are the beliefs player i would hold if he started with beliefs m_i^s and proceeded to experience private history $\{z_{i,1}, \dots, z_{i,s}\}$. That information depreciates ensures that beliefs converge to the actual probability of ω_{-i} conditional on the entire infinite history. Given this, $m_i \in M_i^*(\omega_i)$. ■

Let $T^s(M_i)$ denote the application of the T operator s times on M_i and $\bar{\Delta} = (\Delta^{D-i})^{D_i}$ (that is, $M_i = \bar{\Delta}$ implies for all ω_i , all beliefs are acceptable).

LEMMA 4. $\lim_{s \rightarrow \infty} T^s(\bar{\Delta}) = M_i^*$.

Proof. Examination of the T operator shows it to be monotonic in that if $M_i \subset \hat{M}_i$, $T(M_i) \subset T(\hat{M}_i)$. Since $T(\bar{\Delta}) \subset \bar{\Delta}$, $T^2(\bar{\Delta}) \subset T(\bar{\Delta})$ and so on. Thus $T^s(\bar{\Delta})$ represents a sequence of (weakly) ever smaller included sets, guaranteeing that the limit exists. From Lemma 2, $M_i^* \subset T(M_i^*)$. Lemma 3 then implies $T(M_i^*) \subset M_i^*$, thus $M_i^* = T(M_i^*)$. Further, since $M_i^* \subset \bar{\Delta}$ monotonicity implies $T(M_i^*) = M_i^* \subset T(\bar{\Delta})$ and so on. Thus $M_i^* \subset \lim_{s \rightarrow \infty} T^s(\bar{\Delta})$. Since $\lim_{s \rightarrow \infty} T^s(\bar{\Delta}) = T(\lim_{s \rightarrow \infty} T^s(\bar{\Delta}))$, Lemma 3 implies $\lim_{s \rightarrow \infty} T^s(\bar{\Delta}) \subset M_i^*$. ■

While Lemma 4 shows that M_i^* is the largest fixed point of T (and thus M_i^* can be computed by successively applying T to $\bar{\Delta}$) unlike the value sets calculated in Abreu, Pearce, and Stacchetti (1990) we can show that M_i^* is, in fact, the unique fixed point of our operator if the strategy depends only on the last k periods of the player's private history

and information depreciates. However, iterating on $\bar{\Delta}$ is particularly useful since at every iteration $M_i^* \subset T^s(\bar{\Delta})$. Thus, if after s iterations, incentives hold on the boundary of $T^s(\bar{\Delta})$ one need iterate no further in order to verify that σ is an equilibrium.

So far, we have focused on the finite state strategies σ and left implicit the corresponding measures on infinite histories π and π_i . For a given finite state strategy σ , one can construct π as follows: Strategy σ and the function P define a Markov transition matrix X mapping the aggregate state ω yesterday to the aggregate state ω' today. That information depreciates implies that this matrix has a unique ergodic distribution, $m^e(\omega)$. The probability of an infinite sequence ending in a particular $z_1 = (a_1, y_1)$ is $\sum_{\omega} m^e(\omega)(\Pi_i \sigma_i(a_{i,1}|\omega_i))P(y_1|a_1)$. The probability of an infinite sequence ending in (z_1, z_2) (where again the most recent realization is first) is $\sum_{\omega} m^e(\omega)(\Pi_i \sigma_i(a_{i,2}|\omega_i))P(y_2|a_2)(\Pi_i \sigma_i(a_{i,1}|\omega'_i(\omega_i, z_2)))P(y_1|a_1)$, and so on. Likewise, the conditional probability measures π_i can be calculated using Bayes' rule similarly conditioning only on the last s periods, for all s .

Finally, note that the ability to verify whether or not a finite state strategy σ such that information depreciates is an SSE-ih (when coupled with the appropriate beliefs π and π_i) allows one to calculate *all* such pure-strategy k -history dependent SSE-ih for the simple reason that there exist a finite number of candidate pure strategies.

4. Two Examples

In this section we construct two simple examples. The first is based on Mailath and Morris (2002). Consider the two player partnership game in which each player $i \in \{1, 2\}$ can take action $a_i \in \{C, D\}$ (cooperate or defect) and each can realize a private outcome $y_i \in \{G, B\}$ (good or bad). If m players cooperate, then with probability $p_m(1 - \epsilon)^2 + (1 - p_m)\epsilon^2$,

both players realize the good private outcome. With probability $(1 - \epsilon)\epsilon$, player 1 realizes the good outcome while player 2 realizes the bad. (Likewise, with this same probability, player 2 realizes the good outcome and player 1 the bad.) Finally, with probability $p_m\epsilon^2 + (1 - p_m)(1 - \epsilon)^2$, both players realize the bad outcome. Essentially, this game is akin to one in which p_m determines the probability of an unobservable common outcome and ϵ is the probability that player i 's outcome differs from the common outcome. Thus when $\epsilon = 0$, outcomes are public, and when ϵ approaches zero, outcomes are almost public. Payoffs are determined by specifying for each player i the vector $\{u_i(C, G), u_i(C, B), u_i(D, G), u_i(D, B)\}$.

Next consider perhaps the simplest non-trivial pure strategy: tit-for-tat. That is, let each player i play C if his private outcome was good in the previous period and D otherwise. This is a two-state strategy with $\Omega_i = \{R, P\}$, for “reward” and “punish.” For infinite histories ending in $y_i = G$, player i is in state $\omega_i = R$ (and the strategy calls for the player to play C) and for all other histories, player i is in state P where he plays D . Thus, for computation purposes, the set $M_i^*(\omega_i)$ is simply an interval specifying the range of probabilities that player $-i$ realized a good outcome last period, given that player i is in state ω_i . The mapping T from Section 3 then maps a collection of two intervals (one for each ω_i) to a collection of two intervals, and the results in that section imply that starting with the unit interval for each of these and iterating delivers the true intervals $M_i^*(R)$ and $M_i^*(P)$.

For $\beta = 0.9$, $p_0 = 0.3$, $p_1 = 0.55$, and $p_2 = 0.9$ and a payoff of 1 for receiving a good outcome and a payoff of -0.4 for cooperating, we can easily verify that the static game is a prisoner's dilemma and that tit-for-tat is an equilibrium of the public outcome ($\epsilon = 0$) game. For $\epsilon > 0$, beliefs matter and one must construct the intervals $M_i^*(\omega_i)$. The procedure of iterating the T mapping (starting with unit intervals) is relatively easily implemented on a

computer. For $\epsilon = 0.025$ the procedure converges (in less than a second) to these intervals: $M_i^*(R) = [0.923, 0.972]$, and $M_i^*(P) = [0.036, 0.199]$. For each specification of ω_i , if player i believes the other player saw a good outcome with a probability within $M_i^*(\omega_i)$, he wishes to follow the equilibrium strategy (C if $\omega_i = R$, D otherwise); thus tit-for-tat is an equilibrium.

If ϵ is increased to $\epsilon = 0.03$, then the intervals $M_i^*(\omega_i)$ shift toward the middle and widen: $M_i^*(R) = [0.908, 0.966]$ and $M_i^*(P) = [0.043, 0.229]$. Now, if $\omega_i = R$ and player i believes that his opponent is in state R with probability 0.908, he wishes to deviate and play D rather than C . Thus, with $\epsilon = 0.03$, tit-for-tat is not an equilibrium, since, by construction, there exists an infinite private history for player i ending in G such that he believes the other player saw a good outcome last period with probability 0.908 (the lower end of the interval). Simply put, being only 91 percent sure your opponent saw the same good signal as you (and thus will cooperate along with you) is an insufficient inducement for cooperation in this repeated prisoner's dilemma.

From Mailath and Morris (2002) we know that in this example, for sufficiently small ϵ , tit-for-tat is an equilibrium, and obviously for sufficiently high ϵ it is not. Our analysis of this example allows us to go further: to establish exactly for which epsilons the profile is an equilibrium. That is, our methods allow us to consider whether any proposed strategy is an equilibrium strategy, regardless of whether the signals are nearly public. In fact, one can construct equilibria which depend on the private signals *not* being nearly public.

Consider a two-player battle of the sexes game where each player $i \in \{1, 2\}$ can take action $a_i \in \{Ballet, Hockey\}$ and each can realize a private outcome $y_i \in \{G, B\}$ (good or bad). If both players take the same action, they both realize a good outcome with probability 0.9, both receive a bad outcome with probability 0.08, and player i realizes a good outcome

while player $-i$ receives a bad outcome with probability 0.01. If the players take differing actions, they both realize a good outcome with probability 0.05, both receive a bad outcome with probability 0.05, and player i realizes a good outcome while player $-i$ receives a bad outcome with probability 0.45. If player 1 realizes a bad outcome, her payoff is zero, and if she realizes a good outcome, her payoff is $\frac{3}{2}$ if she played *Ballet* and 1 if she played *Hockey*. Likewise, if player 2 realizes a bad outcome, his payoff is zero, and if he realizes a good outcome, his payoff is $\frac{3}{2}$ if he played *Hockey* and 1 if he played *Ballet*. As in the previous example, $\beta = 0.9$.

Our methods can be used to check if the following simple strategy is an equilibrium: if a player's private outcome was good, repeat last period's play regardless of whether it was on or off path. If his (or her) private outcome was bad, switch away from last period's play regardless of whether it was on or off path. This strategy is a two-state automaton $\omega_i = ([PlayBallet], [PlayHockey])$, and belief sets are intervals specifying the probability that the other player is in state *PlayBallet*. (This strategy depends on previous actions as well as signals, but nevertheless, information depreciation is easily verified.) For these parameters, the intervals are $M_i^*(PlayBallet) = [0.889, 0.988]$ and $M_i^*(PlayHockey) = [0.012, 0.111]$, and incentives hold on the boundaries of these two intervals. But note they hold precisely because this is *not* a game with almost public signals. That is, suppose player 1 is in state *PlayHockey* and deviates by playing *Ballet*, while believing (with high probability) that player 2 is in state *PlayHockey*. If she realizes a bad outcome, the function P above implies she believes player 2 most likely received a good outcome (and thus will not switch states), and thus it is in her interest to follow the equilibrium by playing *Hockey* next period. If P were such that she believed player 2 also had a bad outcome, as would be the case if outcomes were almost

public, after this deviation, player 1 would no longer be willing to follow the strategy.

5. Games with a Start Date

In this section we connect our results in Section 3 for games with infinite histories to the more traditional class of infinitely repeated games where there exists a start date, $t = 0$. We first show how to construct, for each SSE-ih, a correlated equilibrium in the corresponding game with a start date in which the correlation device sends fictitious infinite private histories to each player. Then, as we did in Section 3, we restrict ourselves to finite state strategies where information depreciates and show how, for each finite state SSE-ih, to construct a correlated equilibrium in which the correlation device signals an initial state ω_i for each player i , as opposed to signaling infinite private histories. Finally, we develop conditions for checking whether an *arbitrary* correlation device which signals starting states, when coupled with a finite state strategy σ , is a correlated equilibrium. Since sequential equilibria are examples of correlated equilibria (with degenerate signaling devices), these methods deliver simple necessary and sufficient conditions for constructing all k -history dependent pure strategy sequential equilibria of private monitoring games with a start date.

If players receive private signals $s_i \in S_i$ at the beginning of the game, a joint strategy σ is such that $\sigma_{i,t}(a_{i,t}|s_i, (z_{i,1}, \dots, z_{i,t})) : A \times S_i \times Z_i^t \rightarrow [0, 1]$, where $Z_i^t = (Z_i)^t$. Let $v_{i,t}(s, (z_{i,1}, \dots, z_{i,t})|\sigma)$ be the expected discounted payoff to player i if he knows the joint signal s , every player's private history $(z_{j,1}, \dots, z_{j,t})$, and he and the other players follow σ . Let $Ev_{i,t}(s_i, (z_{i,1}, \dots, z_{i,t})|\sigma) = \int_{s_{-i}} v_{i,t}(s, (z_{i,1}, \dots, z_{i,t})|\sigma) d\pi_{i,t}(s_{-i}, (z_{i,1}, \dots, z_{i,t}))$, where $\pi_{i,t}$ describes player i 's beliefs regarding the signals and histories of the other players given what he knows at date t . By a *correlated equilibrium* we mean a joint strategy σ , a signaling device

or probability measure $x : \mathcal{B}(S) \rightarrow [0, 1]$, (where S is the joint signal space), and conditional probability measures $\pi_{i,t} : \mathcal{B}(S_i) \times Z_i^t \times \mathcal{B}(S_{-i}) \times Z_{-i}^t \rightarrow [0, 1]$ such that

1. Strategies σ_i are mutual best responses (given beliefs) after all signals and histories.
2. Beliefs, π_i , after all signals and histories are consistent with the signaling device x , the strategy σ , and P .

The next lemma demonstrates that any SSE-ih (σ, π, π_i) of $\Gamma^{-\infty, \infty}$ induces a correlated equilibrium of the game with a start date $\Gamma^{0, \infty}$, by letting the signal space S be a set of infinite fictitious histories and having the strategy for the game with a start date combine a player's fictitious history signal and his actual history up to date t in such a way that the player treats his fictitious history as if it were real. That is, if player i receives fictitious history $(\hat{z}_{i,1}, \hat{z}_{i,2}, \dots)$ (with date subscripts going backward in time as in the previous sections), and then experiences actual history $(z_{i,0}, z_{i,1}, \dots, z_{i,t})$ (with date subscripts referring to the calendar date), we let $\sigma_{i,t}(a_{i,t} | s = (\hat{z}_{i,1}, \hat{z}_{i,2}, \dots), (z_{i,0}, \dots, z_{i,t})) = \sigma_i(a_{i,t} | (z_{i,t}, \dots, z_{i,0}, \hat{z}_{i,1}, \hat{z}_{i,2}, \dots))$. In effect, through the use of fictitious histories, we have constructed fully stationary strategies in a game with a start date — a somewhat difficult task since in a game with a start date, the calendar date is automatically encoded into a player's history through the length of that history.

LEMMA 5. Take as given an SSE-ih (σ, π, π_i) of $\Gamma^{-\infty, \infty}$. For the game with a start date, $\Gamma^{0, \infty}$, let $S = \mathbf{Z}$, $x = \pi$, and $\sigma_{i,t}(a_{i,t} | s = (\hat{z}_{i,1}, \dots), (z_{i,0}, \dots, z_{i,t})) = \sigma_i(a_{i,t} | (z_{i,t}, \dots, z_{i,0}, \hat{z}_{i,1}, \dots))$. Then (σ, x, π_i) is a correlated equilibrium of $\Gamma^{0, \infty}$.

Proof. The result immediately follows from (σ, π, π_i) being an SSE-ih and that $\sigma_{i,t}$ reacts to a fictitious history signal, or an actual history with a fictitious history prepended to it, as

if it were an actual history in $\Gamma^{-\infty, \infty}$. ■

While general, Lemma 5 may not be very useful since it relies on a signal space with a continuum of signals. If σ is a finite state strategy where information depreciates, Lemmas 6 and 7 below construct correlated equilibria with the signal space S equaling the (finite) set of states Ω .

LEMMA 6. Take as given a finite state SSE-ih (σ, π, π_i) of $\Gamma^{-\infty, \infty}$ where information depreciates. For the game with a start date, $\Gamma^{0, \infty}$, there exists a correlated equilibrium with $S = \Omega$ and $\sigma_{i,t}(a_{i,t}|\omega_i, (z_{i,0}, \dots, z_{i,t-1})) = \sigma_i(a_{i,t}|\omega'_i(\omega_i, z_{i,1}, \dots, z_{i,t-1}))$, where $\omega'_i(\omega_i, z_{i,1}, \dots, z_{i,t-1})$ denotes player i 's private state if he starts in private state ω_i and realizes $(z_{i,1}, \dots, z_{i,t-1})$.

Proof. Let X denote the probability transition matrix from state ω yesterday to state ω' today defined by σ and the function P . Since information depreciates, X defines a unique ergodic distribution over states ω . If joint signals are drawn from this distribution, beliefs for each player i must lie within M_i^* since the beliefs of player i regarding ω_{-i} (conditional on ω_i) are a weighted average of player i 's beliefs regarding ω_{-i} when conditioning on his entire infinite history. By construction, if beliefs over ω_{-i} start within M_i^* , they will always lie within M_i^* . That (σ, π, π_i) is an SSE-ih then ensures that incentives hold at all signals and histories. ■

Lemma 6 ensures that every finite state SSE-ih of $\Gamma^{-\infty, \infty}$ where information depreciates defines at least one correlated equilibrium of $\Gamma^{0, \infty}$. But one still may ask if other correlation devices may work to permit, say, higher initial values than the correlation device constructed above. Thus we now turn to developing necessary and sufficient conditions for any correlation device $x : \Omega \rightarrow [0, 1]$, when coupled with a finite state SSE-ih σ , to constitute a correlated

equilibrium.

Define $\hat{M}_i$ to be set of beliefs such that incentives hold for all beliefs $m_i \in \hat{M}_i$. Lemma 1 showed that a necessary and sufficient condition for a finite state dependent strategy to be (part of) an SSE-ih is that $M_i^* \subset \hat{M}_i$ for all i . We need to ensure, however, that incentives are satisfied not only for a particular belief generated by a correlation device, but also for all possible successors of that belief, and successors of those beliefs, and so on. Recall (from Section 3) $m'_i(m_i, a_i, y_i) \in \Delta_{-i}^D$ as the successor beliefs (regarding ω_{-i}) of beliefs m_i given new data (a_i, y_i) . Further recall $\omega_i^+(a_i, y_i, \omega_i)$ as the next-period private state given current-period private state ω_i and a new (a_i, y_i) . Define the operator $\hat{T}(M_i)$ (mapping sets of beliefs to sets of beliefs) as

$$\begin{aligned} \hat{T}(M_i)(\omega_i) &= \{m_i | m_i \in M_i(\omega_i), \text{ and for all } (a_i, y_i), \\ &\quad m'_i(m_i, a_i, y_i) \in M_i(\omega_i^+(a_i, y_i, \omega_i))\} \end{aligned}$$

In words, $\hat{T}$ eliminates an element of $M_i(\omega_i)$ if there exists a successor belief which is not in $M_i(\omega_i^+)$. Clearly, $\hat{T}$ is monotonic and $\hat{T}(\hat{M}_i) \subset \hat{M}_i$. Thus $\{\hat{T}^s(\hat{M}_i)\}_{s=0}^\infty$ represents a sequence of (weakly) ever smaller included sets, guaranteeing that the limit, denoted $\overline{M}$, exists.

LEMMA 7. Take as given a finite state SSE-ih (σ, π, π_i) of $\Gamma^{-\infty, \infty}$ where information depreciates. For the game with a start date, $\Gamma^{0, \infty}$, the strategy $\sigma_{i,t}(a_{i,t} | \omega_i, (z_{i,0}, \dots, z_{i,t-1})) = \sigma_i(a_{i,t} | \omega'_i(\omega_i, z_{i,1}, \dots, z_{i,t-1}))$ constitutes a correlated equilibrium when coupled with correlation device $x : \Omega \rightarrow [0, 1]$ with implied conditional beliefs $\pi_i(\omega_{-i} | \omega_i) = x(\omega_i, \omega_{-i}) / \sum_{\bar{\omega}_{-i}} x(\omega_i, \bar{\omega}_{-i})$, if and only if $\pi_i(\omega_i) \in \overline{M}_i(\omega_i)$ for all i and ω_i .

Proof. If $\pi_i(\omega_i) \in \overline{M}_i(\omega_i)$ for all i and ω_i , then by the construction of $\overline{M}_i$, incentives hold for all histories. Only if: $\hat{M}_i$ is a collection of compact sets because it is defined by weak incentive compatibility constraints that are linear in beliefs. Second, the $\hat{T}$ operator maps compact sets into compact sets because Bayes' rule is linear in prior beliefs. Hence, $\{\hat{T}^s(\hat{M}_i)\}_{s=0}^\infty$ is a sequence of ever smaller (in the sense of set inclusion) collections of compact sets so the limit $\overline{M}_i(\omega_i)$ is also compact. Now, suppose that $\pi_i(\omega_i) \notin \overline{M}_i(\omega_i)$ for some ω_i . That implies that there exists an s such that $\pi_i(\omega_i) \notin \hat{T}^s(\hat{M}_i)$, but that means that there exists a history of the length s such that player i has a profitable deviation after that history. For example, if $s = 0$, then $\pi_i(\omega_i) \notin \hat{M}_i(\omega_i)$ for some ω_i , and then incentive compatibility does not hold at date $t = 0$ for signal ω_i . If $s = 1$, then $\pi_i(\omega_i) \in \hat{M}_i(\omega_i)$ but $\pi_i(\omega_i) \notin \hat{T}(\hat{M}_i)(\omega_i)$ for some ω_i so that incentive compatibility does not hold at date $t = 1$ for a private signal $\overline{\omega}_i$ and a private history $(a_{i,0}, y_{i,0})$ at date 0 such that $\pi_i(\omega_i) = m'(\pi_i(\overline{\omega}_i), a_{i,0}, y_{i,0})$, and so on. ■

The implications of Lemma 7 for constructing sequential equilibria can be seen by referring to our first example from Section 4. In that example, with $\epsilon = 0.025$, tit-for-tat is a $k = 1$ -history dependent SSE-ih. That is, incentives hold for the beliefs which can be generated by infinite histories, $M_i^*(R) = [0.923, 0.972]$, $M_i^*(P) = [0.036, 0.198]$. But incentives hold for wider beliefs than these intervals. For this example, $\hat{M}_i(R) = [0.704, 1]$ and $\hat{M}_i(P) = [0, 0.704]$. Further, for this example, $\hat{T}(\hat{M}) = \hat{M}$, thus $\hat{M} = \overline{M}$. A correlation device signaling states R and P according to the ergodic distribution of signals G and B has $\pi_i(\omega_{-i} = R | \omega_i = R) = 0.966$ and $\pi_i(\omega_{-i} = R | \omega_i = P) = 0.085$. Since both of these beliefs are within $\overline{M}(\omega_i)$, this signaling device can be used to create a correlated equilibrium for the tit-for-tat strategy. The point of Lemma 7, however, is that *any* correlation device which delivers $\pi_i \in \overline{M}$ can be used as well. In particular, since $\overline{M}(\omega_i = R)$ includes the

belief that $\omega_{-i} = R$ with probability 1, putting all probability on $\omega = (R, R)$ delivers a correlated equilibrium. But since this correlation device is degenerate, we have constructed a *sequential* equilibrium where both agents cooperate in the first period and play tit-for-tat after that. Likewise, since $\overline{M}(\omega_i = P)$ includes $\omega_{-i} = R$ with probability 0, we have also constructed a sequential equilibrium where both agents defect in the first period and play tit-for-tat after that. Finally, since $\overline{M}(\omega_i = R)$ does not include the belief that $\omega_{-i} = R$ with probability 0, we have demonstrated that having one player cooperate and the other defect in the first period (and playing tit-for-tat after that) is *not* a sequential equilibrium. Thus we have exhaustively determined which starting conditions, when coupled with a particular SSE-ih, are and are not sequential equilibria.

This logic generalizes. Since the number of pure-strategy k -history dependent strategies where information depreciates is finite (and we have derived methods for checking whether each is an SSE-ih) and the number of pure-strategy possibilities for determining play in the first k periods of a game with a start date is also finite (and we have derived methods for determining whether each of these starting conditions satisfy incentive compatibility), we have constructed methods for determining *all* pure-strategy k -history dependent sequential equilibria (such that information depreciates) of traditional games with a start date.

6. Concluding Remarks

The equilibria we have characterized here can be used to construct others. For example, we can construct sequential equilibria for games with a start date in which players mix in the first period and then follow pure stationary strategies conditional on the initial randomization simply by letting the players draw randomly their private fictitious histories (a construction

similar to Sekiguchi (1997)). Or, we can also design some other non-stationary strategies in the first few periods and append to it one of our equilibria. Furthermore, once we find an equilibrium for a discount factor β , we can apply Ellison's (1994) method and construct an equilibrium for discount factor $\beta^{\frac{1}{2}}$ by simply dividing the game into odd and even periods and making the players treat these two parts of the game independently. Since incentive constraints are continuous in β , the sets $\overline{M}_i$ are too, so finding a generic equilibrium for one β provides a neighborhood of equilibria around it. (Recall that the M_i^* sets are independent of β .)

Some questions remain open. First, we have assumed away any problems associated with interpreting off equilibrium behavior by assuming full support of signals. Can our methods be extended to games without full support? These questions, albeit important, seem tangential to the stationarity issues we have addressed. Some relaxation of the full support assumption is possible, for example by the introduction of public signals (that is, making some realizations of y_i perfectly correlated). But the analysis is much more complicated when some private signals indicate a deviation by another player, and this has not yet been studied much. Second, we have considered finite state strategies. It is an open question how important this restriction is in repeated games with private monitoring. For example, Cole and Kocherlakota (2005) show that for some games with public monitoring, the set of payoffs achievable with public k -history dependent strategies (an important subset of finite state strategies) is strictly smaller than the whole set of PPE payoffs (at least for strongly symmetric strategies). We don't know how rich the class of games with that property is and whether the same is true for sequential or correlated equilibria of games with private monitoring.

References

- [1] Abreu, D., D. Pearce, and E. Stacchetti (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58(5, Sept.), pp. 1041–1063.
- [2] Cole, H. L., and N. R. Kocherlakota (2005): “Finite Memory and Imperfect Monitoring,” *Games and Economic Behavior*, 53(1, Oct.), pp. 59–72.
- [3] Compte, O. (2002): “On Failing to Cooperate When Monitoring Is Private,” *Journal of Economic Theory*, 102(1, Jan.), pp. 151–188.
- [4] Cripps, M., G. Mailath, and L. Samuelson (forthcoming): “Disappearing Private Reputations in Long-Run Relationships,” *Journal of Economic Theory*.
- [5] Ellison, G. (1994): “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching,” *Review of Economic Studies*, 61(3, Jul.), pp. 567–588.
- [6] Ely, J. C. (2002): “Correlated Equilibrium and Trigger Strategies with Private Monitoring,” Manuscript, Northwestern University.
- [7] Ely, J. C., J. Hörner, and W. Olszewski (2005): “Belief-Free Equilibria in Repeated Games,” *Econometrica*, 73(2, Mar.), pp. 377–415.
- [8] Ely, J. C., and J. Välimäki (2002): “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory*, 102(1, Jan.), pp. 84–105.
- [9] Kandori, M. (2002): “Introduction to Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 102(1, Jan.), pp. 1–15.

- [10] Kandori, M., and H. Matsushima (1998): “Private Observation, Communication and Collusion,” *Econometrica*, 66(3, May), pp. 627–652.
- [11] Kandori, M., and I. Obara (2006): “Efficiency in Repeated Games Revisited: The Role of Private Strategies,” *Econometrica*, 74(2, Mar.), pp. 499–519.
- [12] Mailath, G. J., and S. Morris (2002): “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory*, 102(1, Jan.), pp. 189–228.
- [13] Mailath, G. J., and S. Morris (2006): “Coordination Failure in Repeated Games with Almost-Public Monitoring,” *Theoretical Economics*, 1(3, Sept.), pp. 311–340.
- [14] Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*, New York: Oxford University Press.
- [15] Piccione, M. (2002): “The Repeated Prisoner’s Dilemma with Imperfect Private Monitoring,” *Journal of Economic Theory*, 102(1, Jan.), pp. 70–83.
- [16] Sekiguchi, T. (1997): “Efficiency in Repeated Prisoner’s Dilemma with Private Monitoring,” *Journal of Economic Theory*, 76(2, Oct.), pp. 345–361.