Federal Reserve Bank of Minneapolis
Research Department

April 3, 2008

Beliefs and Private Monitoring*

Christopher Phelan

University of Minnesota
Federal Reserve Bank of Minneapolis

Andrzej Skrzypacz

Graduate School of Business
Stanford University

ABSTRACT

This paper develops new recursive, *set based* methods for studying repeated games with private monitoring. For an important subclass of strategies, we provide readily checkable and computable necessary and sufficient conditions for equilibrium. In particular, for any given *finite state strategy*, we find sufficient conditions for whether there exists a distribution over initial states such that the strategy, together with this distribution, formulated a correlated sequential equilibrium. We show that with additional, checkable restrictions on strategies, these sufficient conditions are also necessary. Finally, for any given correlation device for determining initial states (including degenerate cases where players' initial states are common knowledge), we provide necessary and sufficient conditions for the correlation device and strategy to be a correlated sequential equilibrium, or in the case of a degenerate correlation device, for the strategy to be a sequential equilibrium.

*The authors thank Peter DeMarzo, Glenn Ellison, Larry Jones, Narayana Kocherlakota, David Levine, George Mailath, Stephen Morris, Ichiro Obara, Larry Samuelson, Itai Sher, Ofer Zeitouni, seminar participants at the Federal Reserve Bank of Minneapolis, the Harvard/MIT joint theory seminar, Stanford University, Iowa State University, Princeton University, the University of Chicago, the University of Minnesota and the 2006 meetings of the Society for Economic Dynamics and three anonymous referees for helpful comments as well as the excellent research assistance of Kenichi Fukushima of Roozbeh Hosseini. Financial assistance from National Science Foundation Grant # 0721090 is gratefully acknowledged. 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

This paper develops new methods for studying repeated games with private monitoring. In particular, we develop tools that allow us to answer when a particular behavior is consistent with equilibrium, as opposed to analyzing whether a particular payoff can be achieved in equilibrium. For an important subclass of strategies - those which can be represented as finite automata - we provide readily checkable and computable necessary and sufficient conditions for equilibrium. Even though checking equilibrium conditions in public monitoring games and perfect public equilibria is relatively simple, checking equilibrium conditions for games with private monitoring has previously been considered nearly impossible.

For instance, consider the following repeated game with private monitoring taken from Malaith and Morris (2002): Two partners, privately, either cooperate or defect, and in each period each, privately, has either a good or bad outcome. While each player can neither observe his partner's action, nor his partner's outcome, outcomes are correlated: the vector of joint outcomes is a probabilistic function of the vector of joint actions. (A player cooperating makes it more likely that both players have a good outcome.)

A question of obvious interest is which strategies, or families of strategies, constitute equilibria of this game. For instance, is tit-for-tat (cooperate this period if last period you had a good outcome, defect otherwise) an equilibrium? Up to now, for any particular specification of this game, the extensive literature on *repeated games with private monitoring* has given a simple and stark answer: **no one knows**.

The difficulty is that even for the simplest games, such as the one presented above, and even the simplest strategies, such as tit-for-tat, there are an infinite number of possible histories where incentives must be checked, and to check incentives one must calculate beliefs for all of them. Malaith and Samuelson (2006) call this a "difficult, if not impossible, task." This difficulty is not confined to the example above. Finding equilibria for all 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.) In this paper, we show that the checking of incentives in such repeated games with private monitoring is neither impossible nor, for simple strategies, particularly difficult.

The focus of our analysis is strategies which can be represented by a finite automaton (*finite state strategies*). A key point (first made by Mailath and Morris (2002)) 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. This lets us summarize a player's history as a *belief* over a finite state space, a much smaller object than the belief over the private histories of opponents (a point also made by Mailath and Morris (2002)). Moreover, unlike the set of possible private histories, the set of possible states for one's opponents does not grow over time.

While many private histories may put a player in the same state of his automaton, they will, in general, induce different beliefs regarding the state of his opponents. Given this, there are two advantages to working with *sets* of beliefs representing all possible beliefs a player can have in a given private state. One is that it is necessary and sufficient to check incentives only for extreme points of those sets instead of looking at beliefs after all histories. The other advantage is that these sets can be readily calculated using recursive methods (operators from sets to sets) that we describe and demonstrate.

Fixed points of our main set based operator represent the beliefs a player can have regarding his opponents' states "in the long run." We show that if incentives hold for extreme points of these sets, one can always use an initial correlation device to, in effect, start the game off as if it had been already running for a long time.[1] This technique alleviates a fundamental difficulty associated with games with private monitoring: 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. Using randomization or exogenous correlation in period 0 of the game to make it easier to satisfy incentives and hence support an equilibrium, 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 Correlated Sequential Equilibria.

Our main results are presented as follows. In Section 2, we present our model, a standard repeated game with private monitoring with finiteness and full support (all signals seen with positive probability) as its only restrictive assumptions. We also present the subclass of strategies we study — *finite state strategies*, or strategies which can be represented as finite automata. In Section 3, we present for a given *behavior* (a finite state strategy divorced from its starting conditions), sufficient conditions for the existence of starting conditions — a correlation device and a mapping from private signals to initial states of the automata — such that the behavior and starting conditions together form a *Correlated Sequential Equilibrium (CSE)* (Theorem 1). These conditions involve checking incentive constraints on fixed points of our set operator (based on Bayes' rule) which we describe how to compute. We then show that for a subclass of finite automata strategies, our sufficient conditions are also necessary:

[1] An earlier version of this project entitled "Private Monitoring with Infinite Histories" focused on this point.

if incentives do not hold for the extreme beliefs of the largest fixed point of our operator, then there exist no starting conditions such that the behavior, when coupled with the starting conditions, is a Correlated Sequential Equilibrium. (Theorem 2). We then present easy to verify conditions which determine whether a behavior falls into this subclass. In Section 4, we propose two other operators on sets of beliefs with readily computable fixed points which can be used to answer the following question: is a *particular* correlation device, when coupled with a particular finite state strategy, a CSE? (Theorems 3 and 4). Since we can apply these operators to arbitrary correlation devices, and in particular, to degenerate ones, we can use these operators to answer if a particular strategy profile is a sequential equilibrium. In Section 5 we present three simple examples which demonstrate the usefulness of these methods. In Section 6 we conclude.

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.

In terms of the focus on strategies instead of payoffs, our work is closest to Mailath and Morris (2002) and (2006). They consider robustness of particular classes of strategies - those that are equilibria in a public monitoring game - to a perturbation of the game from public to private, yet almost-public monitoring. They show that strict equilibria in strategies which look back only a finite number of periods (a subclass of the strategies we study), are robust to such perturbations. They also show when infinite-history dependent strategies (partly covered by our analysis) are not robust. Our methods allow one to extend their analysis beyond almost-public monitoring games (see Section 6 for a brief discussion.)

2. The Model

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

The current period payoff to player i is denoted $u_i : H_i \to R$. That is, player i's payoff is a function of his own current-period action and private outcome. If player i receives payoff stream $\{u_{i,t}\}_{t=0}^\infty$, his lifetime discounted payoff is $(1-\beta)\sum_{t=0}^\infty \beta^t u_{i,t}$ where $\beta \in (0, 1)$. As usual, players care about the expected value of lifetime discounted payoffs.

We assume that each player is sent a private signal at the beginning of the game. Player i's private signal is denoted $s_i \in S_i$, where the joint signal space $S = S_1 \times \ldots \times S_N$ is finite. Let $x(s)$ denote the probability of joint signal $s \in S$.

Let $h_{i,t} = (a_{i,t}, y_{i,t})$ denote player i's private action and outcome at date $t \in \{0, 1, \ldots\}$, and $h_i^t = (h_{i,0}, \ldots, h_{i,t-1})$ denote player i's private history up to, but not including, date t. A (behavior) *strategy* for player i, $\sigma_i = \{\sigma_{i,t}\}_{t=0}^\infty$, is then, for each date t, a mapping from player i's initial private signal, s_i, and his private history h_i^t, to his probability of taking any given action $a_i \in A_i$ in period t. Let σ denote the joint strategy $\sigma = (\sigma_1, \ldots, \sigma_N)$ and σ_{-i} denote the joint strategy of all players other than player i, or $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_N)$.

(Throughout the paper we use notation $-i$ to refer to all players but player i.)

A. Finite State Strategies

In this paper, we restrict attention to equilibria in *finite state strategies*, or strategies which can be described as finite automata. (However, we allow deviation strategies to be unrestricted). A finite state strategy is defined by four objects: 1) a private state space Ω_i (with D_i elements ω_i), 2) a function $p_i(a_i|\omega_i)$ giving the probability of each action a_i for each private state $\omega_i \in \Omega_i$, 3) a deterministic transition function $\omega_i^+ : \Omega_i \times H_i \to \Omega_i$ determining next period's private state as a function of this period's private state, player i's private action a_i, and his private outcome y_i, and 4) a deterministic mapping, $\omega_i^0 : S_i \to \Omega_i$ from player i's initial private signal, s_i, to his initial state, $\omega_{i,0}$.[2] Given this setup, $\sigma_{i,0}(s_i)(a_i) = p_i(a_i|\omega_i^0(s_i))$, $\sigma_{i,1}(s_i, a_{i,0}, y_{i,0})(a_i) = p_i(a_i|\omega_i^+(\omega_i^0(s_i), a_{i,0}, y_{i,0}))$ and so on.[3]

In this paper, we repeatedly make a distinction between a finite state strategy's *behavior* (objects 1 through 3) and object 4, the mapping from player i's initial signal to his initial state. This object, along with the signal space S and the probability distribution over initial signals $x(s)$, are referred to as *starting conditions*.

B. Beliefs

If an initial joint signal $s \in S$ occurs with probability $x(s)$, then player i, after receiving signal s_i such that $\sum_{\overline{s}_{-i}} x(s_i, \overline{s}_{-i}) > 0$, believes his opponents received signal s_{-i} with probability $x(s_i, s_{-i}) / \sum_{\overline{s}_{-i}} x(s_i, \overline{s}_{-i})$ (Bayes' rule). Since x is exogenous and S finite, we need

[2] The restriction to deterministic transitions is for notational convenience only. All of our methods and results apply to automata with non-deterministic transitions.

[3] 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).

not consider player i's beliefs after receiving a signal s_i such that $\sum_{\overline{s}_{-i}} x(s_i, \overline{s}_{-i}) = 0$. From the assumption of full support ($P(y|a) > 0$ for all (a, y)), the beliefs of player i regarding his opponents' private histories, h_{-i}^t, are also always pinned down by Bayes' rule. However, since the strategies of players $-i$ depend only on their joint state, ω_{-i}, to verify player i's incentive constraints, we need not directly consider player i's beliefs regarding s_{-i} and h_{-i}^t. Instead, we need focus only on player i's beliefs regarding his opponents' current state, ω_{-i}. This is a much smaller object, and, importantly, its dimension does not grow over time.

For a particular initial signal, s_i, and private history, h_i^t, player i's belief over ω_{-i} is simply a point in the $(D_{-i}-1)$-dimensional unit-simplex, denoted $\Delta^{D_{-i}}$. Let $\mu_i(s_i, h_i^t)$ denote player i's belief at the beginning of period t about ω_{-i} after private history h_i^t given initial signal s_i. Let $\mu_i(s_i, h_i^t)(\omega_{-i})$ denote the probability assigned to the particular state ω_{-i}.

Beliefs $\mu_i(s_i, h_i^t)$ can be defined recursively using Bayes' rule. First, let player i's belief over the initial state of his opponents, $\omega_{-i,0}$, as a function of his initial signal s_i be defined as

$$m_i^0(s_i)(\omega_{-i,0}) = \sum_{s_{-i} \text{ s.t. } \omega_{-i}^0(s_{-i})=\omega_{-i,0}} x(s_i, s_{-i}) / \sum_{\overline{s}_{-i}} x(s_i, \overline{s}_{-i}).$$

Next, let $B_i(m_i, h_i|\sigma_{-i}) \in \Delta^{D_{-i}}$ denote the belief of player i over the state of his opponents at the beginning of period t, if his beliefs over his opponents' state at period $t-1$ were $m_i \in \Delta^{D_{-i}}$ and he subsequently observed $h_i = (a_i, y_i)$. Here,

$$B_i(m_i, h_i|\sigma_{-i})(\omega'_{-i}) \quad = \quad \frac{\sum_{\omega_{-i}} m_i(\omega_{-i}) H_i(\omega_{-i}, \omega'_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} m_i(\omega_{-i}) F_i(\omega_{-i}, h_i|\sigma_{-i})}$$

where

$$\begin{aligned} F_i(\omega_{-i}, h_i|\sigma_{-i}) &= \sum_{(a_{-i}, y_{-i})} p_{-i}(a_{-i}|\omega_{-i}) P(y_i, y_{-i}|a_i, a_{-i}), \\ H_i(\omega_{-i}, \omega'_{-i}, h_i|\sigma_{-i}) &= \sum_{h_{-i} \in G_{-i}(\omega_{-i}, \omega'_{-i}|\sigma_{-i})} p_{-i}(a_{-i}|\omega_{-i}) P(y_i, y_{-i}|a_i, a_{-i}) \end{aligned}$$

and

$$G_{-i}(\omega_{-i}, \omega'_{-i}|\sigma_{-i}) = \{h_{-i} = (a_{-i}, y_{-i}) | \omega^+_{-i}(\omega_{-i}, a_{-i}, y_{-i}) = \omega'_{-i}\}$$

or G_{-i} is the set of (a_{-i}, y_{-i}) pairs which cause players $-i$ to transit from state ω_{-i} to state ω'_{-i}.

Let $B_i^1(m_i, h_i|\sigma_{-i}) = B_i(m_i, h_i|\sigma_{-i})$ and for $s \geq 2$, let $B_i^s(m_i, h_i^s|\sigma_{-i}) = B_i(B_i^{s-1}(m_i, h_i^{s-1}|\sigma_{-i}), h_{i,s-1}|\sigma_{-i})$. Then, $\mu_i(s_i, h_i^t) = B_i^t(m_i^0(s_i), h_i^t|\sigma_{-i})$. Note that the function $B_i(m_i, h_i|\sigma_{-i})$ is defined by the *behavioral* aspects of σ_{-i} (the state space Ω_{-i} and the functions p_{-i} and ω^+_{-i}) and not by starting conditions. Further note that $B_i(m_i, h_i|\sigma_{-i})$ and $m_i^0(s_i)$, and thus $\mu_i(s_i, h_i^t)$, do not depend on σ_i at all, and thus are the same regardless of whether player i is playing a finite state strategy or not.

C. Equilibrium

Consider player i following an arbitrary strategy σ_i while players $-i$ follow a finite state strategy σ_{-i}. That is, players $-i$ are restricted to finite state strategies, but player i is not. Let $V_{i,t}(s_i, h_i^t, \omega_{-i}|\sigma_i, \sigma_{-i})$ denote the lifetime expected discounted payoff to player i conditional on initial private signal s_i, private history h_i^t, and players $-i$ being in state ω_{-i}.

Thus,

$$V_{i,t}(s_i, h_i^t, \omega_{-i}|\sigma_i, \sigma_{-i}) = \sum_{a=(a_i,a_{-i})} (\sigma_{i,t}(s_i, h_i^t)(a_i) p_{-i}(a_{-i}|\omega_{-i})) \Big(\sum_y P(y|a)[(1-\beta)u_i(a_i, y_i) + \beta V_{i,t+1}(s_i, (h_i^t, (a_i, y_i)), \omega_{-i}^+(\omega_{-i}, a_{-i}, y_{-i})|\sigma_i, \sigma_{-i})] \Big).$$

For arbitrary beliefs $m_i \in \Delta^{D_{-i}}$, let

$$EV_{i,t}(s_i, h_i^t, m_i|\sigma_i, \sigma_{-i}) = \sum_{\omega_{-i}} m_i(\omega_{-i}) V_{i,t}(s_i, h_i^t, \omega_{-i}|\sigma_i, \sigma_{-i}).$$

Player i's expected payoff given correct beliefs $\mu_i(s_i, h_i^t)$ is then $EV_{i,t}(s_i, h_i^t, \mu_i(s_i, h_i^t)|\sigma_i, \sigma_{-i})$.

If σ_i is a finite state strategy, let $\omega_i(s_i, h_i^t)$ denote the private state for player i at date t implied by initial state $\omega_i^0(s_i)$, transition rule $\omega_i^+(\omega_i, a_i, y_i)$, and history $h_i^t = ((a_{i,0}, y_{i,0}), \ldots, (a_{i,t-1}, y_{i,t-1}))$. Then, for all $(s_i, h_i^t, \hat{s}_i, \hat{h}_i^t)$ such that $\omega_i(s_i, h_i^t) = \omega_i(\hat{s}_i, \hat{h}_i^t)$, $V_{i,t}(s_i, h_i^t, \omega_{-i}|\sigma_i, \sigma_{-i}) = V_{i,t}(\hat{s}_i, \hat{h}_i^t, \omega_{-i}|\sigma_i, \sigma_{-i})$. Given this, we can write player i's lifetime payoff, conditional on ω_{-i}, as a function of his current private state ω_i as opposed to depending directly on his private history, (s_i, h_i^t). Thus we define $v_i(\omega_i, \omega_{-i}|\sigma_i, \sigma_{-i}) \equiv V_{i,t}(s_i, h_i^t, \omega_{-i}|\sigma_i, \sigma_{-i})$ for any (s_i, h_i^t) such that $\omega_i = \omega_i(s_i, h_i^t)$. Then we denote player i's expected payoff, now a function of his current state, ω_i, and his beliefs over his opponents' state, ω_{-i}, as

$$Ev_i(\omega_i, m_i|\sigma_i, \sigma_{-i}) = \sum_{\omega_{-i}} m_i(\omega_{-i}) v_i(\omega_i, \omega_{-i}|\sigma_i, \sigma_{-i}).$$

Definition 1. *A correlation device x and finite state strategy σ (paired with its uniquely defined beliefs μ) form a Correlated Sequential Equilibrium (CSE) of Γ^∞ if for all i, t, s_i, h_i^t,* *and arbitrary* $\hat{\sigma}_i$, $Ev_i(\omega_i(s_i,h_i^t),\mu_i(s_i,h_i^t)|\sigma_i,\sigma_{-i}) \geq EV_{i,t}(s_i,h_i^t,\mu_i(s_i,h_i^t)|\hat{\sigma}_i,\sigma_{-i})$.

There are two difficulties in verifying whether a proposed strategy is a CSE. First, there are infinitely many deviation strategies. Second, to verify the IC constraints we need to know the beliefs players have on and off-equilibrium after each element of the infinite set of possible private histories. The first difficulty is shared by all repeated game models and, as usual, it is solved by using the one-shot deviation principle. The resolution of the second difficulty is the main focus of this paper.

Lemma 1. *(One-shot Deviation Principle) Suppose a correlation device* x *and a finite state strategy* σ *satisfy for all* i, s_i, h_i^t *and* $\hat{a}_i$

$$Ev_i(\omega_i(s_i,h_i^t),\mu_i(s_i,h_i^t)|\sigma_i,\sigma_{-i}) \geq \sum_{\omega_{-i}}\mu_i(s_i,h_i^t)(\omega_{-i})[\sum_{a_{-i}}p_{-i}(a_{-i}|\omega_{-i})\sum_{y}P(y|\hat{a}_i,a_{-i})$$
$$[(1-\beta)u_i(\hat{a}_i,y_i)+\beta v_i(\omega_i^+(\omega_i(s_i,h_i^t),\hat{a}_i,y_i),\omega_{-i}^+(\omega_{-i},a_{-i},y_{-i})|\sigma_i,\sigma_{-i})]].$$

Then, (x,σ) *form a CSE. That is, it is sufficient to check that player* i *does not wish to deviate once and then revert to playing his finite state strategy* σ_i.

Proof. Mailath and Samuelson (2006), page 397. ∎

3. When Is a Behavior Consistent with Equilibrium?

We now turn to the main methodological contribution of the paper: set based methods delivering necessary and sufficient conditions for finite state behaviors to be consistent with equilibrium. That is, when do there exist starting conditions that when coupled with a behavior constitute an equilibrium? (In Section 4, we consider methods for directly verifying whether particular starting conditions, when coupled with a behavior, form an equilibrium).

Rather than considering separately the beliefs $m_i \in \Delta^{D_{-i}}$ that a player will have after some private history, it is useful to consider *sets* of beliefs. In particular, let $M_i(\omega_i) \subset \Delta^{D_{-i}}$ denote a closed, convex, set of beliefs, and M_i be a collection of D_i sets $M_i(\omega_i)$, one for each ω_i. Let $\mathcal{M}$ denote the space of such collections of sets M_i. To define the distance between two elements M_i and $M_i' \in \mathcal{M}$, first let the distance between two beliefs m_i and $m_i' \in \Delta^{D_{-i}}$ be defined by the sup norm (or Chebyshev distance) denoted $|m_i, m_i'| = \max_{\omega_{-i}} |m_i(\omega_{-i}) - m_i'(\omega_{-i})|$. Next, for a belief m_i and a non-empty closed set $A \subset \Delta^{D_{-i}}$ the distance between them (the Hausdorff distance) be defined as $|m_i, A| = \min_{m_i' \in A} |m_i, m_i'|$. For two non-empty, closed sets $(A, A') \subset \Delta^{D_{-i}}$, the Hausdorff distance between them is defined as $|A, A'| = \max\left\{\max_{m_i \in A} |m_i, A'|, \max_{m_i' \in A'} |m_i', A|\right\}$. If A is non-empty let $|A, \emptyset| = 1$ and $|\emptyset, A| = 1$. Finally, let $|\emptyset, \emptyset| = 0$. (Note that for non-empty A and A', $|A, A'| \leq 1$.) Then the distance between two collections of belief sets M_i, $M_i' \in \mathcal{M}$ is defined as $|M_i, M_i'| = \max_{\omega_i} |M_i(\omega_i), M_i'(\omega_i)|$.

We begin by constructing an operator from $\mathcal{M}$ to $\mathcal{M}$ where fixed points of this operator will be a focus of our main results. Let the one-step operator $T(M_i)$ be defined as

$$T(M_i) = \{T(M_i)(\omega_i') | \omega_i' \in \Omega_i\}$$

where

$$\begin{aligned} T(M_i)(\omega_i') \quad = \quad & \text{co}(\{m_i' | \text{ there exists } \omega_i \in \Omega_i, m_i \in M_i(\omega_i) \text{ and } (a_i, y_i) \in G_i(\omega_i, \omega_i' | \sigma_i) \\ & \text{such that } m_i' = B_i(m_i, a_i, y_i | \sigma_{-i})\}), \end{aligned}$$

where co() denotes the convex hull and recalling $G_i(\omega_i, \omega_i' | \sigma_i)$ as the set of (a_i, y_i) such that

$\omega_i^+(\omega_i, a_i, y_i) = \omega_i'$. The T operator works as follows: Suppose one takes as given the sets of beliefs of player i over the private state of the other players, ω_{-i}, last period. Bayesian updating 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, ω_i, the beliefs of player i over the private states of his opponents last period consistent with $m_i \in M_i(\omega_i)$, and a new realization of (a_i, y_i) such that Bayesian updating delivers beliefs m_i', then $m_i' \in T(M_i)(\omega_i^+(\omega_i, a_i, y_i))$. 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_{-i} (as well as all convex combinations of such beliefs). Note that since B_i and G_i depend only on the behavioral aspects of σ, as opposed to starting conditions, the T operator retains the property as well.

We note here that the T operator is relatively easy to operationalize. In particular, the following lemma implies that the extreme points of the collection of sets $T(M_i)$ can be calculated using only the extreme points of the collection of sets M_i.

Lemma 2. *If $M_i(\omega_i)$ is closed and convex for all ω_i, then $T(M_i)(\omega_i)$ is closed and convex for all ω_i. Further, if m_i is an extreme point of $T(M_i)(\omega_i)$, then there exists $\hat{m}_i$, $\hat{\omega}_i$, h_i such that $m_i = B_i(\hat{m}_i, h_i|\sigma_{-i})$, $h_i \in G_i(\hat{\omega}_i, \omega_i|\sigma_i)$ and $\hat{m}_i$ is an extreme point of $M_i(\hat{\omega}_i)$.*

Proof. See Appendix. ∎

A. Fixed Points of T

Our results rely on properties of the fixed points of T. We write $M_i^0 \subset M_i^1$ if $M_i^0(\omega_i) \subset M_i^1(\omega_i)$ for all ω_i. Furthermore, we write M_i is non-empty if there exists a private state ω_i such that $M_i(\omega_i)$ is non-empty. Let $\overline{\Delta}_i$ denote the collection of D_i, D_{-i}-dimensional unit simplexes and $\overline{\emptyset}$ denote the collection of D_i empty sets. Given this, the set inclusion relationship, $\subset$, defines a complete lattice on the space of D_i closed subsets of $\Delta^{D_{-i}}$. (For all (M_i^0, M_i^1), $M_i^0 \subset \overline{\Delta}_i$, $M_i^1 \subset \overline{\Delta}_i$, $\overline{\emptyset} \subset M_i^0$, $\overline{\emptyset} \subset M_i^1$.) Since T is a monotone operator, (if $M_i^0 \subset M_i^1$ then $T(M_i^0) \subset T(M_i^1)$), Tarski's fixed point theorem implies T has a unique greatest fixed point, which we denote $\overline{M}_i$, with the property that if $M_i \subset T(M_i)$, then $T(M_i) \subset \overline{M}_i$. (Since $T(\overline{\emptyset}) = \overline{\emptyset}$, the least fixed point of T is $\overline{\emptyset}$.)

Let $T^1(M_i) \equiv T(M_i)$ and for $n \geq 2$, $T^n(M_i) \equiv T(T^{n-1}(M_i))$. Since $T(\overline{\Delta}_i) \subset \overline{\Delta}_i$, and $T(T(\overline{\Delta}_i)) \subset T(\overline{\Delta}_i)$ (from monotonicity), the sequence $\{\overline{\Delta}_i, T(\overline{\Delta}_i), T(T(\overline{\Delta}_i)), \ldots\}$ must converge. That B_i is continuous implies T is continuous and thus this limit is a fixed point of T and thus equal to $\overline{M}_i$ since $\overline{M}_i$ is a subset of each element of the sequence, again from the monotonicity of T.

To this point, we have not shown that $\overline{M}_i$ is non-empty. However, if there exists any non-empty M_i such that $M_i \subset T(M_i)$, that T is monotone implies $T(M_i) \subset T(T(M_i))$ and thus $\lim_{n\to\infty} T^n(M_i)$ exists. The continuity of T then implies this limit is a fixed point of T which implies that the largest fixed point of T, $\overline{M}_i$, is non-empty since it contains all fixed points of T.

Our candidate for non-empty belief sets M_i such that $M_i \subset T(M_i)$ is the collection of single point belief sets implied by drawing initial states from an invariant distribution. That is, for a finite state strategy σ with *behavior* defined by the state space $\Omega = \Omega_1 \times \ldots \Omega_N$, action probabilities $\{p_i(a_i|\omega_i)\}_{i=1}^N$ and transition functions $\{\omega_i^+(\omega_i, a_i, y_i)\}_{i=1}^N$, let $\tau(\omega, \omega'|\Omega, p, \omega^+)$

denote the Markov transition matrix on the joint state $\omega \in \Omega$ defined by

$$\tau(\omega, \omega'|\Omega, p, \omega^+) = \sum_{(a,y) \text{ s.t. } (a_i,y_i)\in G_i(\omega_i,\omega'_i|\sigma_i) \text{ for all } i} P(y|a)\Pi_i p_i(a_i|\omega_i).$$

Since τ defines a finite state Markov chain, it has at least one invariant distribution, $\pi \in \Delta^D$, (where D is the number of joint states $\omega \in \Omega$).

Next, consider beliefs induced by letting the signal space $S = \Omega$, the probability of initial signals $x = \pi$ and for each player i, letting $\omega_i^0(s_i = \omega_i) = \omega_i$. Let $M_{\pi,i}(\omega_i) = \{m_i^0(s_i = \omega_i)\}$, for all ω_i such that $\sum_{\overline{\omega}_{-i}} \pi(\omega_i, \overline{\omega}_{-i}) > 0$. Otherwise, let $M_{\pi,i}(\omega_i) = \emptyset$. That is, for all ω_i, if ω_i occurs with positive probability under distribution π, $M_{\pi,i}(\omega_i)$ is the single point belief set consisting of what player i believes about ω_{-i} when his initial signal $s_i = \omega_i$.

Lemma 3. *For all i, $M_{\pi,i} \subset T(M_{\pi,i})$.*

Proof. See Appendix. ∎

Lemma 3 then implies that $M^*_{\pi,i} \equiv \lim_{n\to\infty} T^n(M_{\pi,i})$ exists and is a fixed point of T. That $\overline{M}_i$ is non-empty immediately follows.

B. Sufficient Conditions

We now move to establishing sufficient conditions for a finite state behavior (Ω, p, ω^+) to be compatible with equilibrium.

The following theorem establishes that to check incentives, 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 extreme points of the largest fixed point of T, $\overline{M}_i$, or are extreme points of a (weakly) smaller fixed point of T, $M^*_{\pi,i}$, derived by iterating on the beliefs associated with an invariant distribution π of $\tau(\omega, \omega'|\Omega, p, \omega^+)$. That is, it is not necessary to check incentives for every history.

Theorem 1. *Consider behaviors for players $i = 1, \ldots, N$ described by a state space $\Omega = \Omega_1 \times \ldots \Omega_N$, action probabilities $p = \{p_i(a_i|\omega_i)\}_{i=1}^N$, and transition functions $\omega^+ = \{\omega_i^+(\omega_i, a_i, y_i)\}_{i=1}^N$. If*

$$
\begin{aligned}
Ev_i(\omega_i, m_i|\sigma_i, \sigma_{-i}) \quad \geq \quad & \sum_{\omega_{-i}} m_i(\omega_{-i})[\sum_{a_{-i}} p_{-i}(a_{-i}|\omega_{-i}) \sum_y P(y|\hat{a}_i, a_{-i}) \\
& [(1-\beta)u_i(\hat{a}_i, y_i) + \beta v_i(\omega_i^+(\omega_i, \hat{a}_i, y_i), \omega_{-i}^+(\omega_{-i}, a_{-i}, y_{-i})|\sigma_i, \sigma_{-i})]]
\end{aligned} \tag{1}
$$

for all i, $\hat{a}_i$, ω_i and m_i such that

a) *m_i is an extreme point of a set $M_i(\omega_i)$ such that $\overline{M}_i(\omega_i) \subset M_i(\omega_i)$, or*

b) *m_i is an extreme point of $M_{\pi,i}^*(\omega_i)$, where $M_{\pi,i}^* \equiv \lim_{n\to\infty} T^n(M_{\pi,i})$, π is an invariant distribution of $\tau(\omega, \omega'|\Omega, p, \omega^+)$, and for all (i, ω_i) such that $\sum_{\overline{\omega}_{-i}} \pi(\omega_i, \overline{\omega}_{-i}) > 0$, $M_{\pi,i}(\omega_i)(\omega_{-i}) = \{\pi(\omega_i, \omega_{-i}) / \sum_{\overline{\omega}_{-i}} \pi(\omega_i, \overline{\omega}_{-i})\}$.*

Then there exist starting conditions (a signal space S, probability distribution over initial signals $x(s)$, and functions $\omega^0 = \{\omega_i^0(s_i)\}_{i=1}^N$) such that x and the joint strategy σ defined by $(\Omega, p, \omega^+, \omega^0)$ form a CSE.

Proof. Let the signal space S be the state space Ω, let π be an invariant distribution of the one-stage Markov process, τ, on state space Ω defined by (p, ω^+), let $x(s = \omega) = \pi(\omega)$, and let $\omega_i^0(s_i = \omega_i) = \omega_i$. That is, let the initial signal be a recommended initial state for each player where joint recommendations are drawn from an invariant distribution π. From Lemma 3 (and the monotonicity of T), the beliefs of each player i regarding the initial state of his opponents are elements of $M^*_{\pi,i}(\omega_{i,0})$ for each $\omega_{i,0}$ drawn with positive probability. Moreover, the subsequent beliefs for each player i are elements of $M^*_{\pi,i}(\omega_{i,t})$ for each date t and private history h_i^t, where $\omega_{i,t}$ is player i's state at date t after private history h_i^t.

Suppose condition (1) holds, for all i, $\hat{a}_i$, ω_i, and extreme points of $M^*_{\pi,i}(\omega_i)$, where m_i and $\hat{m}_i$ are two such points. Then since (1) is linear in these beliefs, for all $\alpha \in [0,1]$, condition (1) holds for beliefs $\alpha m_i + (1-\alpha)\hat{m}_i$, again for all i, $\hat{a}_i$, and ω_i. Thus incentives hold for all dates t and private histories h_i^t.

Finally, suppose condition (1) holds, for all i, $\hat{a}_i$, ω_i, and extreme points of a set $M_i(\omega_i)$ such that $\overline{M}_i(\omega_i) \subset M_i(\omega_i)$. Then, since $M^*_{\pi,i} \subset \overline{M}_i$ for all π and all players, incentives hold for all players at all dates and all histories when initial states are drawn from any invariant distribution π. ∎

The first sufficient condition in Theorem 1 implies that to show that a behavior is consistent with equilibrium, one does not need to iterate until convergence to prove it. Instead, one can start with $\overline{\Delta}_i$ and iterate only until the incentives hold at the extreme points of $T^n(\overline{\Delta}_i)$.

C. Necessary Conditions

The main theorem of this section shows that if, for all i and M_i, $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$, then the condition that incentives hold for all extreme points of $\overline{M}_i$ is not only sufficient, but also necessary, for the existence of starting conditions which make a behavior compatible with equilibrium.

Theorem 2. *Consider behaviors for players $i = 1, \ldots, N$ described by a state space $\Omega = \Omega_1 \times \ldots \Omega_N$, action probabilities $p = \{p_i(a_i|\omega_i)\}_{i=1}^N$, and transition functions $\omega^+ = \{\omega_i^+(\omega_i, a_i, y_i)\}_{i=1}^N$,* *and suppose for all i and non-empty $M_i \in \mathcal{M}$, $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$. Then there exist starting conditions (a signal space S, probability distribution over initial signals $x(s)$, and functions $\omega^0 = \{\omega_i^0(s_i)\}_{i=1}^N$) such that x and the joint strategy σ defined by $(\Omega, p, \omega^+, \omega_0)$ form a CSE* **only if**

$$Ev_i(\omega_i, m_i|\sigma_i, \sigma_{-i}) \geq \sum_{\omega_{-i}} m_i(\omega_{-i})[\sum_{a_{-i}} p_{-i}(a_{-i}|\omega_{-i}) \sum_y P(y|\hat{a}_i, a_{-i})$$
$$[(1-\beta)u_i(\hat{a}_i, y_i) + \beta v_i(\omega_i^+(\omega_i, \hat{a}_i, y_i), \omega_{-i}^+(\omega_{-i}, a_{-i}, y_{-i})|\sigma_i, \sigma_{-i})]] \quad (2)$$

for all i, $\hat{a}_i$, ω_i and m_i such that m_i is an extreme point of $\overline{M}_i(\omega_i)$.

Proof. Suppose correlation device x and finite state strategy σ form a finite state CSE, but there exists an i, ω_i, belief m_i such that m_i is an extreme point of $\overline{M}_i(\omega_i)$, and an action $\hat{a}_i$ such that condition (2) does not hold. This implies there exists $\epsilon > 0$ such that (2) does not hold for all $\hat{m}_i$ such that $|m_i, \hat{m}_i| < \epsilon$.

Given initial state function $\omega_i^0(s_i)$ (a component of σ_i), let M_i^0 denote the convex hull of the set of initial beliefs player i can have at date 0. That is, $M_i^0(\omega_i) = \text{co}(\{m_i | m_i = m_i^0(s_i)$ for some s_i such that $\omega_i = \omega_i^0(s_i)\})$. (Recall, $m_i^0(s_i)(\omega_{-i}) = \sum_{s_{-i} \text{ s.t. } \omega_{-i}^0(s_{-i}) = \omega_{-i}} x(s_i, s_{-i}) / \sum_{\overline{s}_{-i}} x(s_i, \overline{s}_{-i})$.) Note that M_i^0 is non-empty by construction. (That is, there exists ω_i such that $M_i^0(\omega_i)$ is non-empty.)

Since $\{T^n(M_i^0)\}_{n=0}^\infty$ converges to $\overline{M}_i$, there exists $\overline{t}$ such that there exists an extreme point $\hat{m}_i \in T^{\overline{t}}(M_i^0)(\omega_i)$ such that $|m_i, \hat{m}_i| < \epsilon$ and thus (2) does not hold for belief $\hat{m}_i$. From Lemma 2, there exists a sequence $\{\overline{\omega}_{i,t}, \overline{h}_{i,t}, \overline{m}_{i,t}\}_{t=0}^{\overline{t}}$ such that $\overline{\omega}_{i,\overline{t}} = \omega_i$, $\overline{m}_{i,\overline{t}} = \hat{m}_i$, and for all $0 \leq t \leq \overline{t} - 1$, $\overline{m}_{i,t}$ is an extreme point of $T^t(M_i^0)(\overline{\omega}_{i,t})$, $\overline{\omega}_{i,t+1} = \omega_i^+(\overline{\omega}_{i,t}, \overline{h}_{i,t})$, and $\overline{m}_{i,t+1} = B_i(\overline{m}_{i,t}, \overline{h}_{i,t}|\sigma_{-i})$. Thus after initial signal s_i such that $\omega_i^0(s_i) = \overline{\omega}_{i,0}$ and history

$\{\overline{h}_{i,t}\}_{t=0}^{\bar{t}-1}$, incentives do not hold, contradicting that (x, σ) are an equilibrium. ∎

D. Conditions Ensuring a Unique Limit of T^n.

We now turn to providing conditions ensuring $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$ for all non-empty $M_i \in \mathcal{M}$. That is, when can we guarantee that regardless of the starting sets M_i, successive iteration using the T operator will converge to a unique fixed point?

Method 1:

First, let $\overline{\Omega}_i$ be the set of ω_i such that $\overline{M}_i(\omega_i)$ is non-empty, and let $\mathcal{M}(\overline{\Omega}_i)$ denote the subspace of $\mathcal{M}$ such that for all $M_i \in \mathcal{M}(\overline{\Omega}_i)$, $M_i(\omega_i)$ is non-empty for $\omega_i \in \overline{\Omega}_i$ and empty otherwise. Given our metric, $\mathcal{M}(\overline{\Omega}_i)$ is a complete metric space. Under Assumption 1 below, we show that for all non-empty $M_i \in \mathcal{M}$, there exists a K^* such that $T^K(M_i) \in \mathcal{M}(\overline{\Omega}_i)$ for all $K \geq K^*$. This then implies that if T is a contraction on $\mathcal{M}(\overline{\Omega}_i)$, then $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$ for all non-empty $M_i \in \mathcal{M}$.

Assumption 1. *(Communication) There exists L such that for all $\omega_i^0, \omega_i^1 \in \overline{\Omega}_i$, there exist $(h_{i,0}, \ldots, h_{i,L})$, $(\omega_{i,0}, \ldots, \omega_{i,L})$ such that $\omega_{i,0} = \omega_i^0$, $\omega_{i,L} = \omega_i^1$, and for all $0 \leq t \leq L-1$, $\omega_{i,t+1} = \omega_i^+(\omega_{i,t}, h_{i,t})$.*

In words, Assumption 1 requires that player i's strategy is such that there exists a way to transit in exactly L steps from any state in $\overline{\Omega}_i$ to any other state in $\overline{\Omega}_i$. Thus, for instance, there cannot be two absorbing states in $\overline{\Omega}_i$.

Lemma 4. *The operator T maps $\mathcal{M}(\overline{\Omega}_i)$ to itself. Further, given Assumption 1, there exists K^* such that for all $K \geq K^*$, $T^K(M_i) \in \mathcal{M}(\overline{\Omega}_i)$ for all non-empty $M_i \in \mathcal{M}$.*

Proof. See Appendix. ∎

Next we show that if $B_i(., h_i|\sigma_{-i})$ is a contraction for all h_i, then T is a contraction as well. We say that B_i is a contraction with modulus $\gamma < 1$, if for all h_i, m_i and m_i': $|B_i(m_i, h_i|\sigma_{-i}), B_i(m_i', h_i|\sigma_{-i})| \leq \gamma |m_i, m_i'|$. Likewise, we say that T is a contraction with modulus $\gamma < 1$, (in the complete metric space $\mathcal{M}(\overline{\Omega}_i)$) if for all M_i and $M_i' \in \mathcal{M}(\overline{\Omega}_i)$: $|T(M_i), T(M_i')| \leq \gamma |M_i, M_i'|$.

Lemma 5. *If B_i is a contraction with modulus $\gamma < 1$, then T is a contraction on $\mathcal{M}(\overline{\Omega}_i)$ with modulus γ.*

Proof. See Appendix. ∎

Lemmas 4 and 5 and the contraction mapping theorem then imply that if $B_i(., h_i|\sigma_{-i})$ is a contraction for all h_i and Assumption 1 holds, then for all non-empty $M_i \in \mathcal{M}$, $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$.

In two player games where players follow two-state strategies (such as tit-for-tat or grim trigger) it is straightforward to check whether $B_i(., h_i|\sigma_{-i})$ is a contraction for all h_i. In this case, a belief is a scalar and $B_i(., h_i|\sigma_{-i})$ maps the unit interval to itself. Thus $B_i(., h_i|\sigma_{-i})$ (for a given h_i) is a contraction if and only if the absolute value of its slope can be bounded strictly below one. For more complicated strategies, or a larger number of players, that B_i is a contraction may be harder to verify.

Method 2:

The next lemma shows an alternative, easy to check condition that guarantees that $\overline{M}_i$ is the unique limit of T^n if one starts iterating on any non-empty $M_i \in \mathcal{M}$. For it, we need an additional assumption related to Assumption 1.

Assumption 2. *(On Path Communication) There exists L such that for all $\omega_{-i}^0, \omega_{-i}^1 \in$*

$\overline{\Omega}_{-i}$, *there exist* $(h_{-i,0}, \ldots, h_{-i,L}) = ((a_{-i,0}, y_{-i,0}), \ldots, (a_{-i,L}, y_{-i,L}))$, $(\omega_{-i,0}, \ldots, \omega_{-i,L})$ *such that* $\omega_{-i,0} = \omega^0_{-i}$, $\omega_{-i,L} = \omega^1_{-i}$, *and for all* $0 \leq t \leq L-1$, $\omega_{-i,t+1} = \omega^+_{-i}(\omega_{-i,t}, h_{-i,t})$ *and* $p_{-i}(a_{-i,t}|\omega_{-i,t}) > 0$.

Assumption 1 requires that there exist an L such that it is possible, either on or off path, for player i to get from any state in $\overline{\Omega}_i$ to any other state in $\overline{\Omega}_i$ in exactly L steps. Assumption 2 requires that there exist an L such that it is possible, *on path*, for his opponents, players $-i$, to get from any state in $\overline{\Omega}_{-i}$ to any other state in $\overline{\Omega}_{-i}$ in exactly L steps. Thus if Assumption 2 holds for all players, Assumption 1 is implied.

Lemma 6. *Suppose Assumptions 1 and 2 hold for player* i. *Then* $\overline{M}_i$ *is the unique non-empty fixed point of* T *and for all non-empty* $M_i \in \mathcal{M}$, $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$.

Proof. See Appendix. ∎

4. When are Particular Starting Conditions Consistent with Equilibrium?

In the previous section we developed methods for checking, for a given behavior, whether there exist starting conditions such that the behavior, when coupled with the starting conditions, forms a CSE. In this section we develop methods for checking whether a behavior, when coupled with *particular* starting conditions, forms a CSE. Clearly, if a behavior violates the necessary conditions identified in the previous section, then no starting conditions cause that behavior to be consistent with equilibrium. However, if a behavior satisfies the sufficient conditions, there can exist many ways to start the game that are consistent with equilibrium. Moreover, it may be possible to use a degenerate correlation device (so that each player's initial state is deterministic and thus known to the other players) which would yield a *sequential* equilibrium.

We now describe two related set based methods that provide necessary and sufficient conditions for a behavior, coupled with particular starting conditions, to form a CSE. It is important to note that neither of these methods require either Assumptions 1 or 2 to hold.

A. Verifying One Starting Condition at a Time

Let the operator $T^U(M_i)$ (^{U} for union) be:

$$T^U(M_i) \quad = \quad \{T^U(M_i)(\omega_i)|\omega_i \in \Omega_i\} \text{ where } T^U(M_i)(\omega_i) = \text{co}(T(M_i)(\omega_i) \cup M_i(\omega_i)).$$

In words, the T^U operator calculates for every state ω_i, the convex hull of the union of the prior beliefs player i could hold last period, $M_i(\omega_i)$, and all the posterior beliefs he can hold in that same state, $T(M_i)(\omega_i)$.

The T^U operator has the following properties. First, Lemma 2 applies to T^U as well: since we have proven that the extreme points of $T(M_i)(\omega_i)$ can be calculated using only extreme points of $M_i(\omega_i)$, the same property holds for T^U. Moreover, T^U maps (collections of) closed convex sets to closed convex sets. Second, T^U is monotone by construction, and for any M_i, $M_i \subset T^U(M_i)$. Third, the increasing (and bounded by $\overline{\Delta}$) sequence $\{M_i, T^U(M_i), T^U\left(T^U(M_i)\right) ...\}$ converges for any M_i. We let $M^{*U}(M_i)$ denote the limit of that sequence.

Given a correlation device x and $\omega_i^0(s_i)$ (mapping private signals to initial states) define

$$M_i^0(x)(\omega_i) = co(\{m_i^0(s_i)\,|\omega_i^0(s_i) = \omega_i\})$$

In words, $M_i^0(x)(\omega_i)$ is the convex hull of the beliefs player i can have after observing any private initial signal that puts him in the initial state ω_i. Recall that the sets of possible signals, S_i, are finite, so that $M_i^0(x)(\omega_i)$ is closed and its extreme points correspond to some realizations of the private signals.

To check if the pair (x, σ) is a CSE, it is necessary and sufficient to check incentives at the extreme points of the appropriate limit of T^U :

Theorem 3. *A correlation device x and a finite state strategy σ form a Correlated Sequential Equilibrium if and only if for all i incentives (i.e. condition (1)) hold for beliefs that are extreme points of $M^{*U}(M_i^0(x))$.*

Proof. If: Since incentive compatibility conditions (1) are linear in beliefs, then if they hold for the extreme beliefs of $M^{*U}(M_i^0(x))$, they hold for all beliefs in these sets. By monotonicity, $(T^U)^t(M_i^0(x)) \subset M^{*U}(M_i^0(x))$ for all $t \geq 0$, so incentives hold in the first period for all initial signals, and in all subsequent periods for all possible continuation histories.

Only if: Suppose that incentive compatibility conditions (1) are violated for some state ω_i and extreme belief $m_i \in M^{*U}(M_i^0(x))(\omega_i)$. Since the incentive conditions (1) are continuous in beliefs and are weak inequalities, there exists an $\varepsilon > 0$ such that for all beliefs m_i' such that $|m_i', m_i| < \varepsilon$, incentives are violated in state ω_i with beliefs m_i'.

Now, by definition of T^U, for every t and ω_i, every extreme point of $(T^U)^t(M_i^0(x))(\omega_i)$ is either an extreme point of $(T^U)^{t-1}(M_i^0(x))(\omega_i)$ or an extreme point of $T\left((T^U)^{t-1}(M_i^0(x))\right)(\omega_i)$. Moreover, all extreme points of $M_i^0(x)(\omega_i)$ are generated by some initial signal s_i. Therefore, we can find a private signal s_i, as well as a private history h_i^t such that player i after (s_i, h_i^t) is in state ω_i and his beliefs $\mu_i(s_i, h_i^t)$ satisfy $|\mu_i(s_i, h_i^t), m_i| < \varepsilon$ (using that

$(T^U)^n \left(M_i^0\left(x\right)\right) \to M^{*U}\left(M_i^0\left(x\right)\right)$). Thus (x, σ) are not a CSE. ∎

This result can be related to our previous observations as follows. Suppose that Assumptions 1 and 2 hold so that the T operator has a unique limit if we start iterations with any non-empty set of beliefs. That $M_i \subset T^U(M_i)$ and $T(M_i) \subset T^U(M_i)$ for all M_i then implies, $M^{*U}\left(M_i^0\left(x\right)\right) \supset \overline{M}_i$ for any correlation device, x. This implies 1) it is necessary to satisfy incentives in the long-run (at all beliefs within $\overline{M}_i$), 2) if $M_i^0(x) \subset \overline{M}_i$, then $M^{*U}\left(M_i^0\left(x\right)\right) = \overline{M}_i$, thus it is also sufficient that incentives are satisfied in the long run, and 3) since for arbitrary correlation devices, it will not be the case that $M_i^0(x) \subset \overline{M}_i$, it is more difficult to satisfy incentives for arbitrary correlation devices rather than those constructed specifically to ensure initial beliefs are within $\overline{M}_i$.

B. Verifying All Starting Conditions at Once

The T^U operator requires calculating limits separately for all starting conditions. We now define another operator, T^I (^{I} for incentives) that requires computing only one limit to evaluate all starting conditions.

Define $M_i^I\left(\omega_i\right)$ to be the set of beliefs such that incentives hold in the current period for all beliefs $m_i \in M_i^I\left(\omega_i\right)$ if player i is in state ω_i and plans to follow the strategy in the future. Clearly, a necessary condition for (x, σ) to be a CSE is that $M_i^0\left(x\right) \subset M_i^I$ since otherwise incentives would be violated in the first period. We need to insure, however, that incentives are satisfied not only for a particular belief generated by the correlation device, but also that incentives are satisfied for all possible successors of that belief, and successors of those beliefs, and so on.

Define the operator $T^I(M_i)$ as

$$\begin{aligned} T^I(M_i) &= \{T^I(M_i)(\omega_i)|\omega_i \in \Omega_i\} \text{ where} \\ T^I(M_i)(\omega_i) &= \text{co}(\{m_i|m_i \in M_i(\omega_i) \text{ and for all } (a_i, y_i), \\ &\quad B_i(m_i, a_i, y_i|\sigma_{-i}) \in M_i\left(\omega_i^+(\omega_i, a_i, y_i)\right)\}). \end{aligned}$$

In words, T^I eliminates an element of $M_i(\omega_i)$ if there exists a private history (a_i, y_i) and a successor belief which is not in $M_i(\omega_i^+(\omega_i, a_i, y_i))$. Clearly, T^I is montone and $T^I(M_i) \subset M_i$ for any M_i. Thus the sequence $\{(T^I)^n(M_i^I)\}_{n=0}^{\infty}$ (starting with the set of beliefs such that incentives hold in the first period), represents a sequence of (weakly) ever smaller collection of sets, guaranteeing that the limit, denoted M_i^{*I}, exists. Importantly, M_i^{*I} can be computed independently of x, allowing us to then evaluate all correlation devices to this benchmark:

Theorem 4. *A correlation device x and a finite state strategy σ form a Correlated Sequential Equilibrium if and only if for all i, $M_i^0(x) \subset M_i^{*I}$.*

Proof. If: Since $M_i^0(x) \subset M_i^{*I} \subset (T^I)^t(M_i^I)$ for all t, incentives hold in the first period ($t = 0$), for all posteriors after all possible histories in period 1 ($t = 1$) and so on. So they hold after every history.

Only if: Suppose not. That is, despite (x, σ) being a CSE, there exists a state ω_i, a private signal s_i and a belief $m_{i,0}(s_i) \in M_i^0(x)(\omega_i)$ where $m_{i,0}(s_i) \notin M_i^{*I}(\omega_i)$. Since $M_i^{*I}(\omega_i)$ is a compact set, there exists $\varepsilon > 0$ such that $m_i' \notin M_i^{*I}(\omega_i)$ for all beliefs m_i' such that $|m_i', m_{i,0}(s_i)| < \varepsilon$. However, since $(T^I)^n(M_i^I)(\omega_i) \to M_i^{*I}(\omega_i)$, we can find a finite history h_i^t such that the posterior after (s_i, h_i^t) is m_i', $|m_i', m_{i,0}(s_i)| < \varepsilon$ and the player is in state ω_i. But that implies that player i has incentives to deviate after that history, a contradiction. ∎

Note that for belief-free equilibria (such as those in Ely and Välimäki (2002)), Theorem 4 holds because $M_i^{*I} = \overline{\Delta}_i$, or that incentives hold, by construction, for all beliefs.

5. Three Examples

In this section we construct three simple examples. The first two are 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). The $P(y|a)$ function is such that 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 realize 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 β and for each player i the vector $\{u_i(C,G), u_i(C,B), u_i(D,G), u_i(D,B)\}$.

A. Tit-for-Tat

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 $i \in \{1,2\}$,

$p_i(C|R) = 1$, $p_i(D|P) = 1$, $\omega_i^+(\omega_i, a_i, G) = R$, $\omega_i^+(\omega_i, a_i, B) = P$ for $\omega_i \in \{R, P\}$ and $a_i \in \{C, D\}$. Since $D_{-i} = 2$, the set $\overline{M}_i(\omega_i)$ is simply a closed interval specifying the range of probabilities that player $-i$ is in state R, 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. Further, when the behavior is tit-for-tat, it can be analytically verified that $B_i(m_i, h_i|\sigma_{-i})$ is a contraction. Thus starting with any non-empty initial intervals and iterating delivers the unique limit $\overline{M}_i(R)$ and $\overline{M}_i(P)$.

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, starting from either both players in state R or both players in state P. For $\epsilon > 0$, beliefs matter and to check equilibrium conditions one must construct the intervals $\overline{M}_i(\omega_i)$. The procedure of iterating the T mapping is relatively easily implemented on a computer. For $\epsilon = 0.025$ the procedure converges (in less than a second) to these intervals: $\overline{M}_i(R) = [0.923, 0.972]$, and $\overline{M}_i(P) = [0.036, 0.189]$ (see Figure 1).

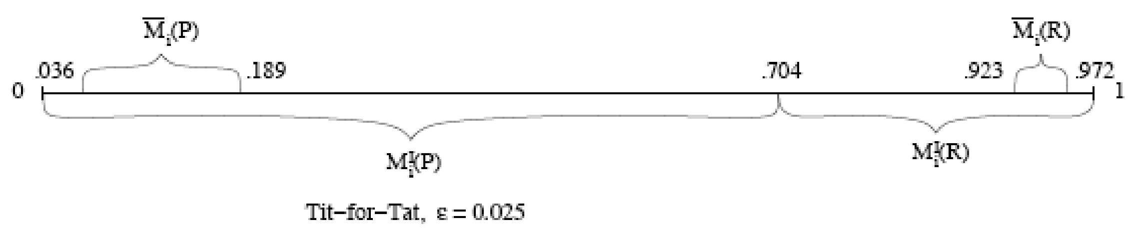

Figure 1

Since B_i is a contraction for this game and strategy, Theorems 1 and 2 imply there exist starting conditions such that tit-for-tat is an equilibrium if and only if incentives hold (equation 1) at each extreme point of $\overline{M}_i(R)$ and $\overline{M}_i(P)$. That is, one needs only to check if player 1 (player 2's incentives are identical from symmetry) indeed wishes to play C when he believes player 2 is in state R with either probability 0.923 or 0.972, and indeed wishes to play D when he believes player 2 is in state R with either probability 0.036 or 0.189 (assuming a reversion to path play after a deviation). Since equation 1 indeed holds for all four beliefs, there exist starting conditions such that tit-for-tat is an equilibrium.

In particular, Theorem 1 delivers one such starting condition. If both players follow the equilibrium, the transition matrix τ between joint state $\omega \in \Omega = \{RR, RP, PR, PP\}$ and $\omega' \in \Omega$ implies a unique invariant distribution $\pi = (0.659, 0.038, 0.038, 0.264)$. If one chooses the correlation device $x = \pi$, then if player $i \in \{1, 2\}$ receives initial signal R, he believes his opponent received signal R with probability $0.945 = 0.659/(0.659 + 0.038)$ and if he receives initial signal P, he believes his opponent received signal R with probability $0.127 = 0.038/(0.038 + 0.264)$. Note that Lemma 3 implies the belief of player i after signal R, $m_i^0(R) = 0.945 \in \overline{M}_i(R)$ and likewise, $m_i^0(P) = 0.127 \in \overline{M}_i(P)$. Thus the correlation device $x = \pi$ and tit-for-tat form a CSE.

Are there any other starting conditions for which tit-for-tat is an equilibrium? Using the T^I operator from Section 4, one can also readily calculate the sets M_i^{*I} for players $i \in \{1, 2\}$. In this example, $M_i^{*I}(R) = [0.704, 1]$ and $M_i^{*I}(P) = [0, 0.704]$. Theorem 4 then implies any correlation device x which delivers conditional beliefs $m_i^0(R) \in [0.704, 1]$ and $m_i^0(P) \in [0, 0.074]$, together with tit-for-tat, forms a CSE. Thus starting each player off in state $\omega_i = R$ with certainty (or x puts all mass on $\omega = RR$) and following tit-for-tat is a *sequential* equilibrium since $M_i^0(x)(R) = \{1\} \subset M_i^{*I}(R)$ and $M_i^0(x)(P) = \emptyset \subset M_i^{*I}(P)$. Likewise, starting each player off in state P (x puts all weight on $\omega = PP$) is also a sequential equilibrium since $M_i^0(x)(R) = \emptyset \subset M_i^{*I}(R)$ and $M_i^0(x)(P) = \{0\} \subset M_i^{*I}(P)$. Finally, letting x be such that one player starts off in state R and his opponent starts off in state P (with certainty) is *not* a sequential equilibrium since $M_i^0(x)(R) = \{0\} \not\subset M_i^{*I}(R)$. Note by calculating $\overline{M}_i$ and M_i^{*I}, we have evaluated *all* deterministic starting conditions and thus all potential sequential equilibria associated with tit-for-tat. From our assumption of finite state strategies, this holds generally.

If ϵ is increased to $\epsilon = 0.04$, then the intervals $\overline{M}_i(\omega_i)$ shift toward the middle and widen: $\overline{M}_i(R) = [0.883, 0.955]$ and $\overline{M}_i(P) = [0.057, 0.262]$. Further, we can calculate the sets $M_i^{*I}(R) = [0.918, 1]$ and $M_i^{*I}(P) = [0, 0.918]$. Now, if $\omega_i = R$ and player i believes that his opponent is in state R with probability 0.883, he wishes to deviate and play D rather than C. Thus, with $\epsilon = 0.04$, tit-for-tat is not an equilibrium for any starting conditions. Simply put, being only 88% sure your opponent saw the same good outcome as you (and thus will cooperate along with you) is an insufficient inducement for cooperation in this repeated prisoner's dilemma. Further, from all starting conditions, there exist histories where a player is supposed to cooperate, but is arbitrarily close to being only 88% sure that the other player is also cooperating. (See Figure 2).

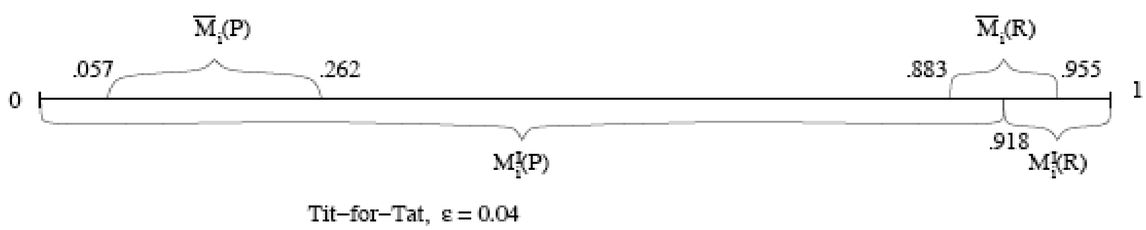

Figure 2

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 ϵ's 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 outcomes are nearly public.

B. Grim Trigger

In this same partnership game consider an alternative strategy: Grim Trigger. The automaton representation is the same as Tit-for-Tat ($\Omega_i = \{R, P\}$) except that state P is now absorbing. (That is, $\omega_i^+(P, a_i, G) = P$ for each a_i under Grim Trigger, and $\omega_i^+(P, a_i, G) = R$ for each a_i under Tit-for-Tat, but otherwise, Tit-for-Tat and Grim Trigger are identical). For $\epsilon = 0.025$, $\overline{M}_i(P) = \overline{M}_i(R) = [0, 0.927]$. At $m_i = 0.927$, player i strictly prefers to play C (assuming reversion to grim trigger) and at $m_i = 0$, player i strictly prefers to play D. Thus incentives do not hold for two extreme points of $\overline{M}_i$ (the leftmost point of $\overline{M}_i(R)$ and the rightmost point of $\overline{M}_i(P)$.) Thus the first of our sufficient conditions from Theorem 1 (that incentives hold at the extreme points of $\overline{M}_i$) does not hold. The second sufficient condition from Theorem 1 is that incentives hold for the extreme points of $M^*_{\pi,i} = \lim_{n\to\infty} T^n(M_{\pi,i})$ where π is an invariant distribution on Ω and $M_{\pi,i}(\omega_i) = \{\pi(\omega_i, \omega_{-i}) / \sum_{\overline{\omega}_{-i}} \pi(\omega_i, \overline{\omega}_{-i})\}$ if ω_i is drawn with positive probability under π and the empty set otherwise. Under Grim Trigger, the only invariant distribution is $\pi(PP) = 1$ since PP is absorbing. Thus $M_{\pi,i}(P) = \{0\}$ and $M_{\pi,i}(R) = \emptyset$. Since $T(M_{\pi,i}) = M_{\pi,i}$, we have $M_{\pi,i} = M^*_{\pi,i}$, and indeed incentives hold at the single point in $M^*_{\pi,i}(P)$. That is, a player in state P who is certain his opponent is in state P indeed does wish to play D.

One can use Theorem 4 to show, for these parameters, that starting both players off in state P with certainty are the only starting conditions where Grim Trigger is an equilibrium. (We suspect that Theorem 4 can be used to provide a simple proof of Mailath and Morris's

(2002) result that Grim Trigger can never allow for cooperation regardless of the parameters). To see this, for these parameters, the intervals where incentives hold in the first period are $M_i^I(P) = [0, 0.268]$ and $M_i^I(R) = [0.268, 1]$. However, starting with these intervals and iterating using the T^I operator gives a fixed point of $M_i^{*I}(P) = \{0\}$ and $M_i^{*I}(R) = \emptyset$. (The T^I operator eventually eliminates all points in $M_i^I(P)$ other than 0 from the fact that $B_i(m_i, C, G|\sigma_{-i}) > m_i$ for all $m_i \in (0, 0.927)$. Given this, the fact that $B_i(m_i, a_i, B|\sigma_{-i}) \in (0, 1)$ for all $m_i \geq 0.268$ implies T^I eventually eliminates all points from $M_i^I(R)$.) Thus for (x, σ) to be a CSE, x must put all weight on $\omega = PP$.

C. A Coordination Game

The following is an example of a game and strategy where equilibrium depends on information *not* being almost public, and thus the ability to analyze general private monitoring environments is crucial. 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 $PlayBallet$. (For this game and strategy, the function $B_i(m_i, h_i|\sigma_{-i})$ can again be shown to be a contraction.) For these parameters, the intervals are $\overline{M}_i(PlayBallet) = [0.900, 0.988]$ and $\overline{M}_i(PlayHockey) = [0.012, 0.110]$, and incentives hold on the boundaries of these two intervals. But note they hold precisely because this is *not* a game with almost public outcomes. 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.

6. Concluding Remarks

Beyond using our methods directly to compute equilibria, one can extend and apply these methods in several ways.

First, as shown in a recent paper by Kandori and Obara (2007) one can use set based methods similar to ours to study strategies that can be represented by finite automata on the equilibrium path but can be much more complicated off the equilibrium path. For example, they allow the strategy off the equilibrium path to be a function of beliefs over other players' states, which implies an infinite number of the automaton states (since players believe that others are always on the equilibrium path, the beliefs are still manageable).

Second, one can prove that if incentives hold strictly (uniformly bounded) for all extreme beliefs of the fixed point operator T^U, then this CSE is robust to small perturbations of the stage game payoffs or the discount factor. The reasoning is as follows: first, the T^U operator and the initial belief sets $M_i^0(x)$ are independent of the payoffs. Hence the fixed point is independent. Second, the incentive constraints are continuous in the stage-game payoffs and the discount factor. Hence, if for the given game the incentives hold strictly for all extreme beliefs of the fixed point of the T^U operator, they also hold weakly for small perturbations of the payoffs or the discount factor. Then, Theorem 3 implies that for the perturbed game the same (x, σ) are a CSE. Similar arguments can be used for perturbations of the monitoring technology (the $P(y|a)$ function) to study robustness to changes in monitoring. In this way, we expect that one can extend the results of Mailath and Morris (2002) beyond games with almost-public monitoring.

References

[1] Compte, O. (2002): "On Failing to Cooperate when Monitoring Is Private," *Journal of Economic Theory*, 102(1, Jan.), pp. 151–188.

[2] Cripps, M., G. Mailath and L. Samuelson (forthcoming): "Disappearing Private Reputations in Long-Run Relationships," *Journal of Economic Theory.*

[3] Ely, J. C. (2002): "Correlated Equilibrium and Trigger Strategies with Private Monitor-

ing,"manuscript, Northwestern University.

[4] Ely, J. C., J. Hörner, and W. Olszewski (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73(2, Mar.), pp. 377-415.

[5] 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.

[6] Kandori, M. (2002): "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory*, 102(1, Jan.), pp. 1-15.

[7] Kandori, M., and I. Obara (2006): "Efficiency in Repeated Games Revisited: The Role of Private Strategies,"*Econometrica,* 74(2, Mar.), pp. 499-519.

[8] Kandori, M. and I. Obara (2007): "Finite State Equilibria in Dynamic Games," manuscript.

[9] Mailath, G. J., and S. Morris (2002): "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory*, 102(1, Jan.), pp. 189-228.

[10] Mailath, G. J., and S. Morris (2006): "Coordination Failure in Repeated Games with Almost-Public Monitoring," *Theoretical Economics*, 1(3, Sept.), pp. 311-340.

[11] Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*, New York: Oxford University Press.

[12] Piccione, M. (2002): "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102(1, Jan.), pp. 70–83.

[13] Sekiguchi, T. (1997): "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring," *Journal of Economic Theory*, 76(2,Oct.), pp. 345-361.

Appendix

Proof of Lemma 2

Proof. First, recall that $T(M_i)(\omega_i)$ is convex from the definition of T. Next, from its definition, we can express $T(M_i)(\omega_i')$ as

$$T(M_i)(\omega_i') = \text{co}(\cup_{\omega_i, h_i \in G_i(\omega_i, \omega_i'|\sigma_{-i})} T(M_i)(\omega_i, h_i)(\omega_i')),$$

where $T(M_i)(\omega_i, h_i)(\omega_i') = \{m_i' |$ there exists $m_i \in M_i(\omega_i)$ such that $m_i' = B_i(m_i, h_i|\sigma_{-i})\}$. Next, note that $B_i(m_i, h_i)(\omega_i')$ is continuous in m_i on the whole domain $m_i \in \Delta^{D_{-i}}$ and $M_i(\omega_i)$ is closed (and bounded). Since $T(M_i)(\omega_i, h_i)(\omega_i')$ is an image of a closed and bounded set under a continous mapping, it is closed (and bounded) as well. As a finite union of closed sets, $T(M_i)(\omega_i')$ is closed as well.

For the second part of the lemma, we use an important property of the non-linear function $B_i(m_i, h_i|\sigma_{-i})(\omega_{-i})$. For all ω_{-i}', m_i^1, m_i^2, h_i and $\alpha \in (0, 1)$,

$$B_i(\alpha m_i^1 + (1-\alpha) m_i^2, h_i|\sigma_{-i})(\omega_{-i}') = \alpha' B_i(m_i^1, h_i|\sigma_{-i})(\omega_{-i}') + (1-\alpha') B_i(m_i^2, h_i|\sigma_{-i})(\omega_{-i}')$$

for some $\alpha' \in (0, 1)$. That is, the posterior of a convex combination of beliefs m_i^1 and m_i^2 is a convex combination of their posteriors, albeit with different weights. To see this, algebraic manipulation delivers

$$\begin{aligned}
&B_i(\alpha m_i^1 + (1-\alpha) m_i^2, h_i|\sigma_{-i})(\omega_{-i}') = \\
&\frac{\alpha \sum_{\omega_{-i}} m_i^1(\omega_{-i}) F_i(\omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} (\alpha m_i^1(\omega_{-i}) + (1-\alpha) m_i^2(\omega_{-i})) F_i(\omega_{-i}, h_i|\sigma_{-i})} B_i(m_i^1, h_i|\sigma_{-i})(\omega_{-i}') + \\
&\frac{(1-\alpha) \sum_{\omega_{-i}} m_i^2(\omega_{-i}) F_i(\omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} (\alpha m_i^1(\omega_{-i}) + (1-\alpha) m_i^2(\omega_{-i})) F_i(\omega_{-i}, h_i|\sigma_{-i})} B_i(m_i^2, h_i|\sigma_{-i})(\omega_{-i}').
\end{aligned}$$

Note

$$\frac{\alpha \sum_{\omega_{-i}} m_i^1(\omega_{-i}) F_i(\omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} (\alpha m_i^1(\omega_{-i}) + (1-\alpha) m_i^2(\omega_{-i})) F_i(\omega_{-i}, h_i|\sigma_{-i})} + \frac{(1-\alpha) \sum_{\omega_{-i}} m_i^2(\omega_{-i}) F_i(\omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} (\alpha m_i^1(\omega_{-i}) + (1-\alpha) m_i^2(\omega_{-i})) F_i(\omega_{-i}, h_i|\sigma_{-i})} = 1.$$

Further, examination of the first quotient has the numerator strictly positive and strictly less than the denominator. So indeed

$$\alpha'\left(\alpha, m_i^1, m_i^2\right) = \frac{\alpha \sum_{\omega_{-i}} m_i^1(\omega_{-i}) F_i(\omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} (\alpha m_i^1(\omega_{-i}) + (1-\alpha) m_i^2(\omega_{-i})) F_i(\omega_{-i}, h_i|\sigma_{-i})} \in (0,1).$$

Now, take any m_i which is an extreme point of $T(M_i)(\omega_i)$ and suppose for any collection (m_i', ω_i', h_i') such that $m_i = B_i(m_i', h_i'|\sigma_{-i})$ and $h_i' \in G_i(\omega_i', \omega_i|\sigma_i)$, the belief m_i' is not an extreme point of $M_i(\omega_i')$. But that implies that there exist two priors $(m_i^0, m_i^1) \in M_i(\omega_i')$ such that m_i' is a strict convex combination of them. But then the posteriors $B_i(m_i^0, h_i'|\sigma_{-i})$ and $B_i(m_i^1, h_i'|\sigma_{-i})$ are both elements of $T(M_i)(\omega_i)$ and m_i is a strict convex combination of them, contradicting that m_i was an extreme point. ∎

Proof of Lemma 3

Proof. For ω_i such that $\sum_{\overline{\omega}_{-i}} \pi(\omega_i, \overline{\omega}_{-i}) > 0$, let $m_i^0(\omega_i)(\omega_{-i}) = \frac{\pi(\omega_i, \omega_{-i})}{\sum_{\overline{\omega}_{-i}} \pi(\omega_i, \overline{\omega}_{-i})}$. That is, $m_i^0(\omega_i)$ is the single point in the set $M_{\pi,i}(\omega_i)$. Since π is an invariant distribution, for all $\omega = (\omega_i, \omega_{-i})$

$$\begin{aligned} m_i^0(\omega_i)(\omega_{-i}) &= \frac{\sum_{\omega^0} \pi(\omega^0) \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} \sum_{h_{-i} \in G_i(\omega_{-i}^0, \omega_{-i}|\sigma_{-i})} p_i(a_i|\omega_i^0) p_{-i}(a_{-i}|\omega_{-i}^0) P(y|a)}{\sum_{\omega^0} \pi(\omega^0) \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} \sum_{h_{-i}} p_i(a_i|\omega_i^0) p_{-i}(a_{-i}|\omega_{-i}^0) P(y|a)} \\ &= \frac{\sum_{\omega_i^0} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} p_i(a_i|\omega_i^0) \sum_{\omega_{-i}^0} \pi(\omega^0) H_i(\omega_{-i}^0, \omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_i^0} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} p_i(a_i|\omega_i^0) \sum_{\omega_{-i}^0} \pi(\omega^0) F_i(\omega_{-i}^0, h_i|\sigma_{-i})}. \end{aligned}$$

Next, note that

$$B_i(m_i^0(\omega_i^0), h_i|\sigma_{-i})(\omega_{-i}) = \frac{\sum_{\omega_{-i}^0} \pi(\omega_i^0, \omega_{-i}^0) H_i(\omega_{-i}^0, \omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}^0} \pi(\omega_i^0, \omega_{-i}^0) F_i(\omega_{-i}^0, h_i|\sigma_{-i})}.$$

We wish to show for all ω_i, $m_i^0(\omega_i)$ is a convex combination of $B_i(m_i^0, h_i|\sigma_{-i})$ over all (ω_i^0, h_i) such that $h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)$. For all (ω_i^0, h_i) such that $h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)$, let

$$\alpha(\omega_i^0, h_i|\omega_i) = \frac{p_i(a_i|\omega_i^0) \sum_{\omega_{-i}^0} \pi(\omega_i^0, \omega_{-i}^0) F_i(\omega_{-i}^0, h_i|\sigma_{-i})}{\sum_{\overline{\omega}_i^0} \sum_{\overline{h}_i \in G_i(\overline{\omega}_i^0, \omega_i|\sigma_i)} p_i(\overline{a}_i|\overline{\omega}_i^0) \sum_{\omega_{-i}^0} \pi(\overline{\omega}_i^0, \omega_{-i}^0) F_i(\omega_{-i}^0, \overline{h}_i|\sigma_{-i})}.$$

Since the denominator of $\alpha(\omega_i^0, h_i|\omega_i)$ is the sum of the numerators over all (ω_i^0, h_i) such that $h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)$, it is clear that $\sum_{\omega_i} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} \alpha(\omega_i^0, h_i|\omega_i) = 1$.

Next, for a given ω_i and ω_{-i}, consider

$$\begin{aligned}
& \sum_{\omega_i^0} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} \alpha(\omega_i^0, h_i|\omega_i) B_i(m_i^0(\omega_i), h_i|\sigma_{-i})(\omega_{-i}) \\
= & \sum_{\omega_i^0} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} \frac{p_i(a_i|\omega_i^0) \sum_{\omega_{-i}^0} \pi(\omega_i^0, \omega_{-i}^0) F_i(\omega_{-i}^0, h_i|\sigma_{-i}) B_i(m_i^0(\omega_i), h_i|\sigma_{-i})(\omega_{-i})}{\sum_{\overline{\omega}_i^0} \sum_{\overline{h}_i \in G_i(\overline{\omega}_i^0, \omega_i|\sigma_i)} p_i(\overline{a}_i|\overline{\omega}_i^0) \sum_{\omega_{-i}^0} \pi(\overline{\omega}_i^0, \omega_{-i}^0) F_i(\omega_{-i}^0, \overline{h}_i|\sigma_{-i})} \\
= & \frac{\sum_{\omega_i^0} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} p_i(a_i|\omega_i^0) \sum_{\omega_{-i}^0} \pi(\omega^0) H_i(\omega_{-i}^0, \omega_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_i^0} \sum_{h_i \in G_i(\omega_i^0, \omega_i|\sigma_i)} p_i(a_i|\omega_i^0) \sum_{\omega_{-i}^0} \pi(\omega^0) F_i(\omega_{-i}^0, h_i|\sigma_{-i})} \\
= & m_i^0(\omega_i)(\omega_{-i}).
\end{aligned}$$

∎

Proof of Lemma 4

Proof. For all M_i, $T(M_i)(\omega_i)$ is non-empty if and only if there exists (ω_i^0, h_i^0) such that $M_i(\omega_i^0)$ is non-empty and $\omega_i^+(\omega_i^0, h_i^0) = \omega_i$. That is, the set of ω_i such that $T(M_i)(\omega_i)$ is non-empty is determined only by the set of ω_i such that $M_i(\omega_i)$ is non-empty. Thus if the set of ω_i such that $T(M_i)(\omega_i)$ is non-empty is identical to the set of ω_i such that $M_i(\omega_i)$ is non-empty, then the set of ω_i such $T^n(M_i)(\omega_i)$ is non-empty is identical for all n. That $\overline{M}_i = T(\overline{M}_i)$ then implies T maps $\mathcal{M}(\overline{\Omega}_i)$ to itself.

Next suppose $M_i(\omega_i)$ is non-empty for some $\omega_i \in \overline{\Omega}_i$. Assumption 1 then implies $T^K(M_i)(\omega_i)$ is non-empty for all $\omega_i \in \overline{\Omega}_i$ and all $K \geq L$.

Finally, suppose $M_i(\omega_i)$ is non-empty for some $\omega_i \notin \overline{\Omega}_i$. Consider the sequence $(\overline{\Delta}_i, T(\overline{\Delta}_i), \ldots, T^{D_i}(\overline{\Delta}_i))$. From monotonicity, for all $0 \leq n \leq D_i$, $T^n(M_i) \subset T^n(\overline{\Delta}_i)$ and $T^n(\overline{\Delta}_i) \subset T^{n-1}(\overline{\Delta}_i)$. If $T^n(\overline{\Delta}_i)$ has the same non-empty sets as $T^{n-1}(\overline{\Delta}_i)$, then from the definition of T and $\overline{\Omega}_i$, $T^n(\overline{\Delta}_i)(\omega_i)$ is non-empty if and only if $\omega_i \in \overline{\Omega}_i$. If, on the other hand, $T^n(\overline{\Delta}_i)$ has fewer non-empty sets as $T^{n-1}(\overline{\Delta}_i)$, then it must have at least one fewer non-empty set $T^n(\overline{\Delta}_i)(\omega_i)$. That is, for each iteration of the T operator on $\overline{\Delta}_i$, $T^n(\overline{\Delta}_i)$ either forever stops creating empty sets or creates at least one. Since $\overline{M}_i(\omega_i)$ is empty for at most $D_i - 1$ states ω_i, $T^{D_i}(\overline{\Delta}_i)(\omega_i)$ is empty for all $\omega_i \notin \overline{\Omega}_i$. That $T^{D_i}(M_i) \subset T^{D_i}(\overline{\Delta}_i)$ then implies $T^{D_i}(M_i)(\omega_i)$ is empty for all $\omega_i \notin \overline{\Omega}_i$. Choosing $K = L + D_i$ then delivers the claim. ∎

Proof of Lemma 5

Proof. Take any two collections of closed belief sets, M and $M' \in \mathcal{M}(\overline{\Omega}_i)$. Denote the distance between them by $|M_i, M_i'| = c$. Suppose to the contrary that

$$|T(M_i), T(M_i')| = c' > \gamma c.$$

That implies that we can find a belief m_i and a state ω_i such that $m_i \in T(M_i)(\omega_i)$ or $T(M_i')(\omega_i)$ such that the distance from m_i to the other set is c'. Without loss of generality, there exists ω_i and $m_i \in T(M_i)(\omega_i)$ such that $|m_i, T(M_i')(\omega_i)| > \gamma c$.

By definition of T, there exists a finite J and a collection $\{m_{i,0}^j, \omega_{i,0}^j, h_i^j\}_{j=1}^J$ such that $h_i^j \in G_i(\omega_{i,0}^j, \omega_i|\sigma_i)$ and weights $\alpha_j \geq 0$ such that $\sum_{j=1}^J \alpha_j = 1$, $m_{i,0}^j \in M_i\left(\omega_{i,0}^j\right)$ and

$m_i = \sum_{j=1}^{j} \alpha_j B_i \left(m_{i,0}^j, h_i^j | \sigma_{-i} \right)$. Then, since the distance between M_i and M_i' is c, for each j there must exist $m_{i,0}^{j\prime} \in M_i' \left(\omega_{i,0}^j \right)$ such that $\left| m_{i,0}^j, m_{i,0}^{j\prime} \right| \leq c$. But since B_i is a contraction with modulus γ, that implies that $\left| B_i \left(m_{i,0}^j, h_i^j | \sigma_{-i} \right), B_i \left(m_{i,0}^{j\prime}, h_i^j | \sigma_{-i} \right) \right| \leq \gamma c$ for all j, and thus $| \sum_{j=1}^{J} \alpha_j B_i (m_{i,0}^j, h_i^j | \sigma_{-i}), \sum_{j=1}^{J} \alpha_j B_i (m_{i,0}^{j\prime}, h_i^j | \sigma_{-i}) | \leq \gamma c$. Since, by the definition of T, $\sum_{j=1}^{J} \alpha_j B_i \left(m_{i,0}^{j\prime}, h_i^j | \sigma_{-i} \right) \in T \left(M_i' \right) (\omega_i)$, we have a contradiction. ∎

Proof of Lemma 6

Proof. From Lemma 4, for all non-empty $M_i \notin \mathcal{M}(\overline{\Omega}_i)$, $T^K(M_i) \in \mathcal{M}(\overline{\Omega}_i)$ for some finite K. Thus if $T^n(M_i)$ converges, it must converge to an element of $\mathcal{M}(\overline{\Omega}_i)$. Thus, without loss, we restrict attention to the case where $\Omega_i = \overline{\Omega}_i$.

Next, let $\mathcal{H}(h_i)$ denote the $D_{-i} \times D_{-i}$ matrix $H_i \left(\omega_{-i}, \omega'_{-i}, h_i | \sigma_{-i} \right)$ where rows correspond to ω_{-i} and the columns to ω'_{-i}. We note that the matrix $\mathcal{H}(h_i)$ has all entries between 0 and 1 and that the rows add up to at most 1, so that if some element is positive, all other elements are strictly bounded away from 1.

From Assumption 2, there exists an L such that for all $h_{i,1}...h_{i,L}$, all elements of the matrix $\mathcal{H}(h_{i,L}) * ... * \mathcal{H}(h_{i,1})$ contain no zeros. Let $\varepsilon > 0$ be the lower bound on them (it exists since the set of h_i is finite and L is finite).

The rest of the proof has two steps. Let beliefs m_i^{E0} and m_i^{E1} be such that $m_i^{E0}(\omega_{-i}^0) = 1$ and $m_i^{E1}(\omega_{-i}^1) = 1$. That is, m_i^{E0} puts all probability on state ω_{-i}^0 and m_i^{E1} puts all weight on state ω_{-i}^1. First, we show that for all $\{h_{i,n}\}_{n=0}^{\infty}$, $\lim_{n\to\infty} |B_i^n(m_i^{E0}, h_i^n), B_i^n(m_i^{E1}, h_i^n)| = 0$. Next we show this implies $\lim_{n\to\infty} T^n(M_i) = \overline{M}_i$ for all $M_i \in \mathcal{M}$.

Step 1:

Recall from Lemma 1 that

$$B_i(m_i, h_i|\sigma_{-i})(\omega'_{-i}) = \frac{\sum_{\omega_{-i}} m_i(\omega_{-i}) H_i(\omega_{-i}, \omega'_{-i}, h_i|\sigma_{-i})}{\sum_{\omega_{-i}} m_i(\omega_{-i}) F_i(\omega_{-i}, h_i)|\sigma_{-i})}$$

Let $B_i(m_i, h_i|\sigma_{-i})$ denote the vector $B_i(m_i, h_i|\sigma_i)(\omega'_{-i})$ and $F_i(h_i|\sigma_{-i})$ denote the vector $F_i(\omega_{-i}, h_i|\sigma_{-i})$. We can then re-write Bayes' rule in the matrix form as:

$$\text{(A1)} \quad B_i(m_i, h_i|\sigma_{-i}) = \underbrace{\frac{1}{m_i \cdot F_i(h_i|\sigma_{-i})}}_{\text{scalar}} m_i \mathcal{H}(h_i)$$

where m_i is a row vector with elements $m_i(\omega_{-i})$.

If player i starts with prior m_i^0 and observes $(h_{i,L}, ...h_{i,1})$ (with $h_{i,1}$ being the most recent observation), then his posterior beliefs after L periods are:

$$\begin{aligned} & B_i^L\left(m_i^0, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i}\right) \\ = & \frac{1}{B_i^{L-1}(m_i^0, h_{i,L}, \ldots, h_{i,2}|\sigma_{-i}) \cdot F_i(h_{i,1}|\sigma_{-i})} B_i^{L-1}\left(m_i^0, h_{i,L}, \ldots, h_{i,2}|\sigma_{-i}\right) \mathcal{H}(h_i) \\ = & \frac{1}{(m_i^0 \mathcal{H}(h_{i,L}) \ldots \mathcal{H}(h_{i,2})) \cdot F_i(h_{i,1}|\sigma_{-i})} m_i^0 \mathcal{H}(h_{i,L}) \ldots \mathcal{H}(h_i) \end{aligned}$$

This implies for $j \in \{0,1\}$, $B_i^L(m_i^{E_j}, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i})$ is equal to the ω_{-i}^j row of matrix

$$\frac{1}{\left(m_i^{E_j} \mathcal{H}(h_{i,L}) \ldots \mathcal{H}(h_{i,2})\right) \cdot F_i(h_{i,1}|\sigma_{-i})} \mathcal{H}(h_{i,L}) \ldots \mathcal{H}(h_{i,1})$$

For a matrix Q let $R_l^Q = \sum_k q_{lk}$ be the sum of the elements of row l of this matrix. Denote by $R(Q)$ a matrix obtained by dividing each element of matrix Q by the corresponding R_l^Q, that is if $B = R(Q)$ then $b_{lk} = \frac{q_{lk}}{R_l^Q}$. By definition the rows of $R(Q)$ add up to 1. Hence, $R(\mathcal{H}(h_{i,L}) ...\mathcal{H}(h_{i,1}))$ is a probability matrix and the posterior belief $B_i^L\left(m_i^{E0}, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i}\right)$ is equal to the ω_{-i}^0 row of $R(\mathcal{H}(h_{i,L}) ...\mathcal{H}(h_{i,1}))$.

Let $d_k(Q)$ be the difference between the largest and smallest elements of $Q's$ column k : $d_k(Q) = \max_{l,j}(q_{lk} - q_{jk}))$. Let $d(Q)$ be the vector of these differences. Then

$\max_{\omega'_{-i}} d(R(\mathcal{H}(h_{i,L})...\mathcal{H}(h_{i,1})))(\omega'_{-i})$ is the maximum distance of the posterior beliefs $B_i^L\left(m_i^{E0}, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i}\right)$ and $B_i^L\left(m_i^{E1}, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i}\right)$ over all extreme priors, m_i^{E0} and m_i^{E1}. To continue, we invoke the following technical lemma (proven below):

Technical Lemma:

Suppose that $\{Q_n\}_{n=1}^{\infty}$ is a sequence of square matrices with all elements $q_{ij} \in (\varepsilon, 1-\varepsilon)$ for some $\varepsilon > 0$. Then there exists a $\delta \in (0,1)$ such that for every n :

$$d\left(R(Q_n...Q_1)\right) \leq \delta d\left(R(Q_{n-1}...Q_1)\right) \leq \delta^{n-1} d\left(R(Q_1)\right)$$

i.e. the distance between the normalized rows of $Q_n...Q_1$contracts by a factor at least δ as we left-multiply it by another matrix from the sequence.

Now, since there exists $L \geq 1$ and $\varepsilon > 0$ such that for all $(h_{i,L}, ...h_{i,1})$ all elements of $\mathcal{H}(h_{i,L})...\mathcal{H}(h_{i,1})$ are bounded between $(\varepsilon, 1-\varepsilon)$, this technical lemma implies that there exists a $\delta \in (0,1)$ such that for any integer n :

$$d(R(\mathcal{H}(h_{i,nL})...\mathcal{H}(h_{i,1}))) \leq \delta\ d(R(\mathcal{H}\left(h_{i,(n-1)L}\right)...\mathcal{H}(h_{i,1}))) \leq \delta^{n-1}\mathbf{1}$$

where $\mathbf{1}$ is a vector of ones (of length D_{-i}). Therefore, for any ε' we can find n large enough so that for any history of length nL and any two extreme priors, m_i^{E0} and m_i^{E1}, the distance between the posteriors will be less than ε'. So, for every history h_i^n, as $n \to \infty$, the posteriors converge to the same belief for all extreme priors.

Step 2:

As we have shown in the proof of Lemma 3, beliefs $B_i(m_i^0, h_i|\sigma_{-i})$ are a convex combination of beliefs $B_i\left(m_i^E, h_i|\sigma_{-i}\right)$ of all extreme priors m_i^E. Applying this reasoning iteratively (that if prior belief m_i is a convex combination of priors m_i' and m_i'', then after applying B_i the posterior of m_i is a convex combination of the posteriors of m_i' and m_i''), we get that for any history sequence, the posteriors after all possible beliefs are convex combinations of posteriors $B_i^L\left(m_i^E, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i}\right)$. Since for any sequence $\left\{h_i^L\right\}_{L=1}^{\infty}$, for all m_i^E the posteriors $B_i^L\left(m_i^E, h_{i,L}, \ldots, h_{i,1}|\sigma_{-i}\right)$ converge, the same is true for posteriors after arbitrary priors. In other words, after long enough histories, the posteriors depend (almost) only on the history and not on the prior.

More formally, let the set $\hat{\Delta}_i$ be a collection of unit symplexes for all $\omega_i \in \overline{\Omega}_i$ and empty sets otherwise (i.e. the largest set in $\mathcal{M}(\overline{\Omega}_i)$). Clearly, $\lim_{n\to\infty} T^n(\hat{\Delta}_i) = \overline{M}_i$ (by the Tarski fixed point theorem and Lemma 4). Now, suppose there exists a set $M_i^0 \in \mathcal{M}(\overline{\Omega}_i)$ such that $\lim_{n\to\infty} T^n(M_i^0) \neq \overline{M}_i$ (either because the sequence $\{T^n(M_i^0)\}_{n=0}^{\infty}$ converges to something else or does not converge at all). First, by monotonicity of T, for all n, $T^n(M_i^0) \subset T^n(\hat{\Delta}_i)$ so that for any $\varepsilon > 0$ we can find n large enough so that for all $\omega_i \in \overline{\Omega}_i$ and all $m_i \in T^n(M_i^0)(\omega_i)$, $\left|m_i, \overline{M}_i(\omega_i)\right| < \varepsilon$.

So the only remaining possibility for $\lim_{n\to\infty} T^n(M_i^0) \neq \overline{M}_i$ is that there exists $\varepsilon > 0$ such that for all n' we can find $n \geq n'$ and a state ω_i^n such that $\max_{m_i \in \overline{M}_i(\omega_i^n)} |T^n(M_i^0)(\omega_i^n), m_i| > \varepsilon$. If so, then we can find an extreme belief $m_i^n \in \overline{M}_i(\omega_i^n)$ that satisfies $|m_i^n, T^n(M_i^0)(\omega_i^n)| > \varepsilon$. Fix n' such that the distance between $B_i^n\left(m_i^{E0}, h_i^n|\sigma_{-i}\right)$ and $B_i^n\left(m_i^{E1}, h_i^n|\sigma_{-i}\right)$ is uniformly bounded by $\varepsilon/2$ for all histories h_i^n (for all $n \geq n'$) and extreme priors m_i^{E0}, m_i^{E1}. Since $\lim_{n\to\infty} T^n(\hat{\Delta}_i) = \overline{M}_i$, we can find a history h_i^n and a prior m_i^{E0} such that $|B_i^n\left(m_i^{E0}, h_i^n|\sigma_{-i}\right), m_i^n| \leq \varepsilon/2$ and a starting state ω_i^0 such that after that history, player i is in the state ω_i^n. Now, take any prior $m_i^0 \in M_i^0(\omega_i^0)$. It is a convex combination of the priors m_i^E. Moreover, after history h_i^n, the posterior $B_i^n(m_i^0, h_i^n|\sigma_{-i}) \in T^n(M_i^0)(\omega_i^n)$ and it is a convex combination of the posteriors $B_i^n\left(m_i^E, h_i^n|\sigma_{-i}\right)$ (from $(A1)$ it easily follows that a posterior of a convex combination of priors is a convex combination of posteriors, albeit with different weights, see also proof of

Lemma 2) so that

$$\left|B_i^n\left(m_i^0, h_i^n|\sigma_{-i}\right), B_i^n\left(m_i^{E0}, h_i^n|\sigma_{-i}\right)\right| \leq \max_{m_i^{E1}, m_i^{E2}} \left|B_i^n\left(m_i^{E1}, h_i^n|\sigma_{-i}\right), B_i^n\left(m_i^{E2}, h_i^n|\sigma_{-i}\right)\right| \leq \varepsilon/2$$

Using the triangle inequality, $|B_i^n(m_i^0, h_i^n|\sigma_{-i}), m_i^n| \leq \varepsilon$, but that contradicts that $|m_i^n, T^n(M_i^0)(\omega_i^n)| > \varepsilon$. ∎

Proof of Technical Lemma.

Proof. Consider a general multiplication: $Q = Q_n...Q_1$. Let $C = Q_n$, $F = Q_{n-1}$, $B = Q_{n-2}...Q_1$. Also, let $G = FB$, so that $Q = CG = CFB$. By assumption all the elements of C and F are bounded from below by $\varepsilon > 0$, but we do not know that about B or G.

For arbitrary matrix A, let R_k^A be the sum of elements in row k of that matrix. Then:

$$R_i^Q = \sum_j q_{ij} = \sum_j \left(\sum_k c_{ik} g_{kj}\right) = \sum_k c_{ik} \sum_j g_{kj} = \sum_k c_{ik} R_k^G$$

Moreover,

$$\frac{q_{ij}}{R_i^Q} = \sum_k \Gamma_k^i \frac{g_{kj}}{R_k^G}$$

where

$$\Gamma_k^i = \frac{c_{ik} R_k^G}{\sum_l c_{il} R_l^G}$$

In words, the elements of $R(Q_n G)$ are a weighted average of elements of $R(G)$ (note that $\sum_k \Gamma_k^i = 1$).

We now bound the weights Γ_k^i uniformly away from zero for all G. To this end, bound

$$\Gamma_k^i = \frac{c_{ik} R_k^G}{\sum_l c_{il} R_l^G} > c_{ik} \frac{R_k^G}{\sum_l R_l^G}$$

Next,

$$
\begin{aligned}
\frac{R_i^G}{\sum_l R_l^G} &= \frac{\sum_k f_{ik} R_k^B}{\sum_l \sum_k f_{lk} R_k^B} = \frac{\sum_k f_{ik} R_k^B}{\sum_k \sum_l f_{lk} R_k^B} = \frac{\sum_k f_{ik} R_k^B}{\sum_k R_k^B L_k^F} \\
&= \sum_k \frac{f_{ik}}{L_k^F} \frac{L_k^F R_k^B}{\sum_k R_k^B L_k^F} = \sum_k \frac{f_{ik}}{L_k^F} \gamma_k
\end{aligned}
$$

where L_k^F is the sum of elements of column k of matrix F and

$$
\gamma_k = \frac{L_k^F R_k^B}{\sum_k R_k^B L_k^F} \in [0,1] .
$$

Note that for any matrices F and B, $\sum_k \gamma_k = 1$.

Therefore we can find a bound $\varepsilon_L \in \left(0, \frac{1}{2}\right)$ that depends only on F and C :

$$
\Gamma_k^i \geq c_{ik} \frac{R_k^G}{\sum_l R_l^G} \geq \varepsilon \min_k \frac{f_{ik}}{L_k^F} > \varepsilon_L
$$

where ε_L can be chosen independently of i and k.

To finish the proof we show how to choose δ. Consider any column k. Any element of column k of matrix $R\left(Q_n ... Q_1\right)$ is a weighted average of elements in the same column of $R\left(Q_{n-1} ... Q_1\right)$, with the weights bounded uniformly away from zero by ε_L. Suppose that the largest and smallest elements of column k of $R\left(Q_{n-1} ... Q_1\right)$ are equal to q_h and q_l respectively. Then

$$
d_k\left(R\left(Q_n ... Q_1\right)\right) \leq \left(1-\varepsilon_L\right) q_h + \varepsilon_L q_l - \left(\varepsilon_L q_h + \left(1-\varepsilon_L\right) q_l\right) = \left(1 - 2\varepsilon_L\right) d_k\left(R\left(Q_{n-1} ... Q_1\right)\right) .
$$

So we can pick $\delta = \left(1 - 2\varepsilon_L\right)$. ∎