Sequential Equilibria of Multi-Stage Games with Infinite Sets of Types and Actions[*]

Roger B. Myerson and Philip J. Reny

Department of Economics
University of Chicago

Abstract

Abstract: We consider how to extend Kreps and Wilson's 1982 definition of sequential equilibrium to multi-stage games with infinite sets of types and actions. A concept of *open sequential equilibrium* is defined by taking limits of strategy profiles that can consistently satisfy approximate sequential rationality for all players at arbitrarily large finite collections of observable open events. Existence of open sequential equilibria is shown for a broad class of *regular projective games*. Examples are considered to illustrate the properties of this solution and the difficulties of alternative approaches to the problem of extending sequential equilibrium to infinite games.

1. Introduction

We propose a definition of sequential equilibrium for multi-stage games with infinite type sets and infinite action sets, and prove its existence for a broad class of games.

Sequential equilibria were defined for finite games by Kreps and Wilson (1982), but rigorously defined extensions to infinite games have been lacking. Various formulations of "perfect Bayesian equilibrium" (defined for finite games in Fudenberg and Tirole 1991) have been used for infinite games, but no general existence theorem for infinite games is available.

Harris, Stinchcombe and Zame (2000) provided important examples that illustrate some of the difficulties that arise in infinite games and they also introduced a methodology for the analysis of infinite games by way of nonstandard analysis, an approach that they showed is equivalent to considering limits of a class of sufficiently rich sequences (nets, to be precise) of finite game approximations.

It may seem natural to try to define sequential equilibria of an infinite game by taking limits of sequential equilibria of finite games that approximate it. The difficulty is that no

---

[*] Current Version: April 27, 2015. We thank Pierre-Andre Chiappori and Asher Wolinsky for helpful comments. Reny thanks the National Science Foundation (SES-9905599, SES-0214421, SES-0617884) for financial support.

general definition of "good finite approximation" has been found. Indeed, it is easy to define sequences of finite games that seem to be converging to an infinite game (in some sense) but have limits of equilibria that seem wrong (e.g., examples 4.2 and 4.3 below).

Instead, we consider limits of strategy profiles that are approximately optimal (among *all* strategies in the game) on finite sets of events that can be observed by players in the game.

For any $\varepsilon > 0$, a strategy profile is an $(\varepsilon, \mathcal{F})$-sequential equilibrium on a set of open observable events $\mathcal{F}$ iff it gives positive probability to each event C in $\mathcal{F}$, and any player who can observe C has no strategy that could improve his conditional expected payoff by more than ε when C occurs.

An open sequential equilibrium is defined as a limit of $(\varepsilon, \mathcal{F})$-sequential equilibrium conditional distributions on outcomes as $\varepsilon \to 0$ and as the set of conditioning events $\mathcal{F}$ on which sequential rationality is imposed expands to include all finite subsets of a neighborhood basis for all players' open observable events.

The remainder of the paper is organized as follows. Section 2 introduces the multi-stage games that we study and provides the notation and concepts required for the definition of open sequential equilibrium given in Section 3. Section 4 provides a number of examples that motivate our definition and illustrate its limitations. Section 5 introduces the subset of "regular projective games" and states an open sequential equilibrium existence result for this class of games. All proofs are in Section 6.

2. Multi-Stage Games

A multi-stage game is played in a finite sequence of dates.[1] At each date, nature chooses first. Each player then simultaneously receives a private signal, called the player's "type" at that date, about the history of play. Each player then simultaneously chooses an action from his set of available actions at that date. Perfect recall is assumed.

Multi-stage games allow infinite action and type sets and can accommodate any finite extensive form game with perfect recall in which the information sets of distinct players never "cross" one another.[2]

Formally, a multi-stage game $\Gamma = (N, K, A, \Theta, T, \mathcal{M}, \tau, p, u)$ consists of the following items.

Γ.1. $i \in N = \{\text{players}\}$ is the finite set of players; $K = \{1, ..., |K|\}$ is the finite set of dates of the game. $L = \{(i, k) \in N \times K\}$ – write ik for (i, k).

---

[1] A countable infinity of dates can be accommodated with some additional notation.

[2] That is, in a multi-stage game with perfect recall, each player always knows, for any of his opponents' type sets, whether that opponent has been informed of his type from that set or not.

Γ.2. $A = \times_{ik \in L} A_{ik}$, where $A_{ik} = \{$possible actions for player i at date $k\}$; action sets are history independent.[3]

Γ.3. $T = \times_{ik \in L} T_{ik}$, where $T_{ik} = \{$possible informational types for player i at date $k\}$ has a topology of open sets $\mathcal{T}_{ik}$ with a countable basis.

Γ.4. $\Theta = \times_{k \in K} \Theta_k$, where $\Theta_k = \{$possible date k states$\}$.

Γ.5. σ-algebras (closed under countable intersections and complements) of measurable subsets are specified for each A_{ik} and Θ_k, and T_{ik} is given its Borel σ-algebra. All one-point sets are measurable. Products are given their product σ-algebras.

The subscript, $< k$, will always denote the projection onto dates before k, and $\leq k$ weakly before. e.g., $A_{<k} = \times_{i \in N, h<k} A_{ih} = \{$possible action sequences before date $k\}$ ($A_{<1} = \Theta_{<1} = \{\emptyset\}$), and for $a \in A$, $a_{<k} = (a_{ih})_{i \in N, h<k}$ is the partial sequence of actions before date k.

If X is any of the sets above or any of their products, $\mathcal{M}(X)$ denotes its set of measurable subsets. Let $\Delta(X)$ denote the set of countably additive probability measures on $\mathcal{M}(X)$.

Γ.6. The date k state is determined by a regular conditional probability p_k from $\Theta_{<k} \times A_{<k}$ to $\Delta(\Theta_k)$. i.e., for each $(\theta_{<k}, a_{<k})$, $p_k(\cdot|\theta_{<k}, a_{<k}) \in \Delta(\Theta_k)$, and for each $B \subset \mathcal{M}(\Theta_k)$, $p_k(B|\theta_{<k}, a_{<k})$ is a measurable function of $(\theta_{<k}, a_{<k})$. Nature's probability function is $p = (p_1, ..., p_{|K|})$.

Γ.7. Player i's date k information is given by a measurable type function $\tau_{ik} : \Theta_{\leq k} \times A_{<k} \to T_{ik}$. Assume perfect recall: $\forall ik \in L$, $\forall h < k$, there is a measurable function $\phi_{ikh} : T_{ik} \to T_{ih} \times A_{im}$ such that $\phi_{ikh}(\tau_{ik}(\theta_{\leq k}, a_{<k})) = (\tau_{ih}(\theta_{\leq h}, a_{<h}), a_{ih})$ $\forall \theta \in \Theta$, $\forall a \in A$. The game's type function is $\tau = (\tau_{ik})_{ik \in L}$.

Γ.8. Each player i has a bounded measurable utility function $u_i : \Theta \times A \to \mathbb{R}$, and $u = (u_i)_{i \in N}$.

So, at each date $k \in K$ starting with date $k = 1$, and given a partial history $(\theta_{<k}, a_{<k}) \in \Theta_{<k} \times A_{<k}$, nature chooses a date-$k$ state θ_k according to $p_k(\cdot|\theta_{<k}, a_{<k})$ producing the partial history $(\theta_{\leq k}, a_{<k})$. Each player i is then simultaneously informed of his private date-k type, $t_{ik} = \tau_{ik}(\theta_{\leq k}, a_{<k})$, after which each player i simultaneously chooses an action from his date-k action set A_{ik}. The game then proceeds to the next date. After $|K|$ dates of play this leads to an outcome $(\theta, a) \in \Theta \times A$ and the game ends with player payoffs $u_i(\theta, a)$, $i \in N$.

In the next three subsections, we formally introduce strategies, outcome distributions, and payoffs, as well as the collections of events on which we will impose sequential rationality.

---

[3] History-dependent action sets can always be modeled by letting A_{ik} be the union over all histories of player i's history-dependent date k action sets, and ending the game with a strictly dominated payoff for player i if he ever takes an infeasible action.

2.1. Strategies and Induced Outcome Distributions

A *strategy* for player $ik \in L$ is any regular conditional probability from T_{ik} to $\Delta(A_{ik})$ – i.e., for each $t_{ik} \in T_{ik}$, $s_{ik}(\cdot|t_{ik})$ is in $\Delta(A_{ik})$ and for each $B \in \mathcal{M}(A_{ik})$, $s_{ik}(B|t_{ik})$ is a measurable function of t_{ik}.

Let S_{ik} denote ik's set of strategies and let $S_i = \times_{k \in K} S_{ik}$ denote i's (behavioral) strategies. Perfect recall ensures that there is no loss in restricting attention to S_i for each player i. Let $S = \times_{ik \in L} S_{ik}$ denote the set of all strategy profiles.

Let $S_{i,<k} = \times_{h<k} S_{ih}$ and let $S_{<k} = \times_{i \in N} S_{i,<k}$ denote the strategy profiles before date k, and let $S_{\cdot k} = \times_{i \in N} S_{ik}$ denote the set of date-k strategy vectors with typical element $s_{\cdot k} = (s_{ik})_{i \in N}$.

Given any $s \in S$, let s_{ik} or $s_{i,<k}$ or $s_{\leq k}$ respectively denote the coordinates of s in S_{ik} or $S_{i,<k}$ or $S_{\leq k}$.

Each $s_{\cdot k} \in S_{\cdot k}$ determines a regular conditional probability Ψ_k from $\Theta_{<k} \times A_{<k}$ to $\mathcal{M}(\Theta_k)$ such that, for any measurable product set $Z = Z_0 \times (\times_{i \in N} Z_i) \subseteq \Theta_k \times A_{\cdot k}$, and any $(\theta_{<k}, a_{<k}) \in \Theta_{<k} \times A_{<k}$,

$$\Psi_k(Z|\theta_{<k}, a_{<k}, s_{\cdot k}) = \int_{\theta_k \in Z_0} [\Pi_{i \in N} s_{ik}(Z_i|\tau_{ik}(\theta_{\leq k}, a_{<k}))] p_k(d\theta_k|\theta_{<k}, a_{<k}).$$

For any measurable set $B \subseteq \Theta_{\leq k} \times A_{\leq k}$, and any $(\theta_{<k}, a_{<k}) \in \Theta_{<k} \times A_{<k}$, let $B_k(\theta_{<k}, a_{<k}) = \{(\theta_k, a_k) \in \Theta_k \times (\times_{i \in N} A_{ik}) : ((\theta_k, \theta_{<k}), (a_k, a_{<k})) \in B\}$.

For any strategy profile s, we inductively define measures $\Psi_{\leq k}(\cdot|s_{\leq k})$ on $\Theta_{\leq k} \times A_{\leq k}$ so that $\Psi_{\leq 1}(\cdot|s_{\leq 1}) = \Psi_1(\cdot|\emptyset, \emptyset, s_{\cdot 1})$ and, for any $k \in \{2, ..., |K|\}$, for any measurable set $B \subseteq \Theta_{\leq k} \times A_{\leq k}$,

$$\Psi_{\leq k}(B|s_{\leq k}) = \int_{(\theta_{<k}, a_{<k}) \in \Theta_{<k} \times A_{<k}} \Psi_k(B_k(\theta_{<k}, a_{<k})|\theta_{<k}, a_{<k}, s_{\cdot k}) \Psi_{\leq k-1}(d(\theta_{<k}, a_{<k})|s_{\leq k-1}).$$

Let $P(\cdot|s) = \Psi_{\leq |K|}(\cdot|s)$ be the distribution over outcomes in $\Theta \times A$ induced by the strategy profile $s \in S$. The dependence of $P(\cdot|s)$ on nature's probability function p will sometimes be made explicit by writing $P(\cdot|s; p)$.

2.2. Conditional Probabilities and Payoffs

For any $s \in S$, for any $ik \in L$ and for any $C \in \mathcal{M}(T_{ik})$, define

$$\langle C \rangle = \{(\theta, a) \in \Theta \times A : \tau_{ik}(\theta_{\leq k}, a_{<k}) \in C\},$$

and define

$$P_T(C|s) = P(\langle C\rangle\,|s).$$

Then $\langle C\rangle \in \mathcal{M}(\Theta \times A)$ is the set of outcomes that would yield types in $C \subseteq T_{ik}$, and $P_T(C|s)$ is the probability that i's date k type is in C under the strategy profile s. The dependence of $P_T(\cdot|s)$ on nature's probability function p will sometimes be made explicit by writing $P_T(\cdot|s;p)$.

Let $\mathcal{Y}$ denote the set $\mathcal{M}(\Theta \times Y)$ of measurable subsets Y of $\Theta \times A$. So $\mathcal{Y}$ is the set of all *outcome events.* If $P_T(C|s) > 0$, then we may define (for any $Y \in \mathcal{Y}$ and any $i \in N$): *conditional probabilities,*

$$P(Y|C,s) = P(\{(\theta,a) \in Y : \tau_{ik}(\theta_{\leq k}, a_{<k}) \in C\}|s)/P_T(C|s),$$

and *conditional expected payoffs,*

$$U_i(s|C) = \int_{\Theta\times A} u_i(\theta,a)P(d(\theta,a)|C,s).$$

2.3. Observable Open Events and Essential Types

An open set $C \subseteq T_{ik}$ is *observable* iff $\exists a \in A$ such that $P_T(C|a) > 0$.[4] In positive-probability events, players do not need to consider what others would do in any open event that is not observable, as they could not make its probability positive even by deviating.

Remark 1. *In most practical settings of interest, it would be equivalent to say that an open subset* C *of* T_{ik} *is observable iff* $\exists s \in S$ *such that* $P_T(C|s) > 0$. *Indeed, suppose that all* Θ_k, A_{ik} *are metric spaces with their Borel* σ*-algebras, and all* $\tau_{ik} : \Theta_{\leq k} \times A_{<k} \to T_{ik}$ *and all* $p_k : \Theta_{<k} \times A_{<k} \to \Delta(\Theta_k)$ *are continuous, with product topologies on all product sets and the weak* topology on* $\Delta(\Theta_k)$. *If* $C \subseteq T_{ik}$ *is open and* $P_T(C|s) > 0$, *then* $\exists a \in A$ *such that* $P_T(C|a) > 0$. *See Lemma 6.1 in Section 6.*

Let us call the set $\bar{T}_{ik} = \{t_{ik} \in T_{ik} :$ every open subset of T_{ik} containing t_{ik} is observable$\}$ the set *essential types* for ik. So if t_{ik} is not essential, then there is an open neighborhood of t_{ik} that will have probability 0 no matter what actions the players might use.

Remark 2. $\bar{T}_{ik}$ *is the closure of the union over all* $a \in A$ *of the supports of* $P_T(\cdot|a)$ *as probability distributions on* T_{ik}, *and so* $\bar{T}_{ik}$ *is the smallest closed set of types such that* $P_T(\bar{T}_{ik}|a) = 1\ \forall a \in A$.

---

[4] The $a \in A$ here is interpreted as the constant pure strategy profile $s \in S$ such $s_{ik}(a_{ik}|t_{ik}) = 1\ \forall t_{ik} \in T_{ik}, \forall ik \in L$.

Let $\mathcal{T} = \cup_{ik\in L}\mathcal{T}_{ik}$ (a disjoint union) denote the set of all open sets of types for dated players and let $\mathcal{T}^* = \{C \in \mathcal{T} : \exists a \in A \text{ such that } P_T(C|a) > 0\} = \{\text{open sets of types that are observable}\}$.

The set $\mathcal{T}^*$ of observable open sets contains all of the open sets on which sequential rationality will ever be imposed. But we will be content if sequential rationality is imposed only on any sufficiently rich subcollection of observable open sets that we now introduce.

A *neighborhood basis* for the essential types is any set $\mathcal{B} \subseteq \mathcal{T}^*$ that contains T_{ik} $\forall ik \in L$ and that satisfies: $\forall ik \in L$, $\forall t_{ik} \in \bar{T}_{ik}$, $\forall C \in \mathcal{T}_{ik}$, if $t_{ik} \in C$ then there exists some $B \in \mathcal{B}$ such that $t_{ik} \in B$ and $B \subseteq C$. Thus, for example, $\mathcal{T}^*$ itself is a neighborhood basis for the essential types.

We are now prepared to present our main definitions.

3. Open Sequential Equilibrium

Say that $r_i \in S_i$ is a *date-k continuation* of s_i, if $r_{ih} = s_{ih}$ for all dates $h < k$.

Definition 3.1. For any $\varepsilon > 0$ and for any $\mathcal{F} \subseteq \mathcal{T}^*$, say that $s \in S$ is an $(\varepsilon, \mathcal{F})$*-sequential equilibrium* of Γ iff for every $ik \in L$ and for every $C \in \mathcal{F} \cap \mathcal{T}_{ik}$ (so that C is open and observable by i at date k)

1. $P_T(C|s) > 0$, and

2. $U_i(r_i, s_{-i}|C) \leq U_i(s|C) + \varepsilon$ for every date-k continuation r_i of s_i.

Note. Changing i's choice only at dates $j \geq k$ does not change the probability of i's types at k, so $P_T(C|r_i, s_{-i}) = P_T(C|s) > 0$.

In an $(\varepsilon, \mathcal{F})$-sequential equilibrium, each open set of types C in $\mathcal{F}$ is reached with positive probability and the player whose turn it is to move there is ε-optimizing conditional on C.

We next define an "open sequential equilibrium" to be a limit of $(\varepsilon, \mathcal{F})$-sequential equilibrium conditional distributions on outcomes as $\varepsilon \to 0$ and as the set of conditioning events $\mathcal{F}$ on which sequential rationality is imposed expands to include all finite subsets of a neighborhood basis for all players' open observable events.

Definition 3.2. Say that a mapping $\mu : \mathcal{Y} \times \mathcal{B} \to [0, 1]$ is an *open sequential equilibrium* of Γ iff $\mathcal{B}$ is a neighborhood basis for the essential types, and, for every $\varepsilon > 0$, for every finite subset $\mathcal{F}$ of $\mathcal{B}$, and for every finite subset $\mathcal{G}$ of $\mathcal{Y}$, there is an $(\varepsilon, \mathcal{F})$-sequential equilibrium s such that,

$$|P(Y|C, s) - \mu(Y|C)| < \varepsilon, \textit{ for every } (Y, C) \in \mathcal{G} \times \mathcal{F}.$$

We then also say that μ is an *open sequential equilibrium* (*of* Γ) *conditioned on* $\mathcal{B}$.

Equivalently, $\mu : \mathcal{Y} \times \mathcal{B} \to [0,1]$ is an open sequential equilibrium of Γ conditioned on $\mathcal{B}$ iff there is a net $\{s^{\varepsilon,\mathcal{F},\mathcal{G}}\}$ of $(\varepsilon, \mathcal{F})$-sequential equilibria such that,

$$\lim_{\substack{\varepsilon>0,\ \mathcal{F}\subset\mathcal{B},\ \mathcal{G}\subset\mathcal{Y} \\ \mathcal{F} \text{ and } \mathcal{G} \text{ finite}}} P(Y|C, s^{\varepsilon,\mathcal{F},\mathcal{G}}) = \mu(Y|C), \text{ for every } (Y,C) \in \mathcal{Y} \times \mathcal{B}, \tag{3.1}$$

where smaller values of ε and more inclusive subsets $\mathcal{F}$ of $\mathcal{B}$ and $\mathcal{G}$ of $\mathcal{Y}$ are further along in the index set.

It is an easy consequence of Tychonoff's theorem that an open sequential equilibrium exists so long as $(\varepsilon, \mathcal{F})$-sequential equilibria always exist. The existence of $(\varepsilon, \mathcal{F})$-sequential equilibria is taken up in Section 5. We record here the simpler result (Section 6 contains the proof).

Theorem 3.3. *Let $\mathcal{B}$ be a neighborhood basis for the essential types. If for any $\varepsilon > 0$ and for any finite subset $\mathcal{F}$ of $\mathcal{B}$ there is at least one $(\varepsilon, \mathcal{F})$-sequential equilibrium, then an open sequential equilibrium conditioned on $\mathcal{B}$ exists.*

It follows immediately from (3.1) that if μ is an open sequential equilibrium conditioned on $\mathcal{B}$, then $\mu(\cdot|C)$ is a finitely additive probability measure on $\mathcal{Y}$ for each $C \in \mathcal{B}$, and $\mu(\cdot|\cdot)$ satisfies the Bayes' consistency condition,

$$\mu(\langle C\rangle |D)\mu(Y \cap \langle D\rangle |C)) = \mu(\langle D\rangle |C)\mu(Y \cap \langle C\rangle |D)\ \forall Y \in \mathcal{Y},\ \forall C, D \in \mathcal{B},$$

where, recalling from Section 2.2, $\langle C\rangle$ denotes the set of outcomes that would yield types in C, and similarly for $\langle D\rangle$.[5]

Since $P(\cdot|T_{ik}, s) = P(\cdot|s)$ for any $ik \in L$ and any $s \in S$, it also follows that $\mu(\cdot|T_{ik}) = \mu(\cdot|T_{nh})$ for any ik and any nh in L and so the unconditional finitely additive probability measure on outcomes can be defined by $\mu(Y) = \mu(Y|T_{ik})$ for all $Y \in \mathcal{Y}$. (Recall that a neighborhood basis $\mathcal{B}$ is defined to include each T_{ik}.)

If (3.1) holds, then so long as u_i is bounded and measurable (as we have assumed),

$$\lim_{\substack{\varepsilon>0,\ \mathcal{F}\subset\mathcal{B},\ \mathcal{G}\subset\mathcal{Y} \\ \mathcal{F} \text{ and } \mathcal{G} \text{ finite}}} \int_{\Theta\times A} u_i(\theta,a)P(d(\theta,a)|C, s^{\varepsilon,\mathcal{F},\mathcal{G}}) = \int_{\Theta\times A} u_i(\theta,a)\mu(d(\theta,a)|C)\ \forall C \in \mathcal{B}.$$

Since this holds in particular for $C = T_{ik}$, we may define i's *equilibrium expected payoff* (at μ) by

$$\int_{\Theta\times A} u_i(\theta,a)\mu(d(\theta,a)).$$

---

[5] For finite additivity, note that for any disjoint sets $Y, Z \in \mathcal{Y}$ and for any $C \in \mathcal{B}$, (3.1) and $\lim P(Y \cup Z|C, s^{\varepsilon,\mathcal{F},\mathcal{G}}) = \lim[P(Y|C, s^{\varepsilon,\mathcal{F},\mathcal{G}}) + P(Z|C, s^{\varepsilon,\mathcal{F},\mathcal{G}})]$ imply that $\mu(Y \cup Z|C) = \mu(Y|C) + \mu(Z|C)$. Bayes' consistency is obtained similarly.

Remark 3. *Since we have assumed that the set $\mathcal{T}$ of open sets of the players' types has a countable basis, any neighborhood basis $\mathcal{B}$ for the essential types has a countable neighborhood subbasis.*[6] *Let $\mathcal{B}'$ be any one of them. If μ is an open sequential equilibrium conditioned on $\mathcal{B}$, then the restriction of μ to $\mathcal{Y} \times \mathcal{B}'$ is an open sequential equilibrium conditioned on $\mathcal{B}'$ (since $\mathcal{B}' \subseteq \mathcal{B}$) and the unconditional probability measure $\mu(\cdot)$ on outcomes is unchanged (since each T_{ik} is in $\mathcal{B}'$). So if one is interested only in the unconditional probability measure on outcomes in any open sequential equilibrium, it is without loss of generality to restrict attention to countable neighborhood bases of the essential types.*

Sometimes the unconditional probability measure over outcomes $\mu(\cdot)$ is only finitely additive, not countably additive (Example 4.1). We next define an "open sequential equilibrium distribution" as a countably additive probability measure on the measurable sets of outcomes as follows.

Definition 3.4. Say that a countably additive probability measure ν on $\mathcal{Y}$ is an *open sequential equilibrium distribution of* Γ iff there is an open sequential equilibrium μ and a collection $\mathcal{C} \subseteq \{Y \in \mathcal{Y} : \nu(Y) = \mu(Y)\}$ that is closed under finite intersections and that generates the σ-algebra $\mathcal{Y}$.[7] Since there can be at most one such measure ν,[8] we then also say that ν *is the open sequential equilibrium distribution induced by* μ.

Remark 4. *If $\Theta \times A$ is a compact metric space with its Borel sigma algebra of measurable sets and μ is an open sequential equilibrium, then there exists an open sequential equilibrium distribution induced by μ.*[9] *Indeed, suppose that (3.1) holds and so, in particular, $P(Y|s^{\varepsilon,\mathcal{F},\mathcal{G}}) \to \mu(Y)$ for all $Y \in \mathcal{Y}$. Since $\{P(\cdot|s^{\varepsilon,\mathcal{F},\mathcal{G}})\}$ is a net of countably additive measures on the measurable subsets of the compact metric space $\Theta \times A$, there is a weak*-convergent subnet converging to a countably additive measure $\nu \in \Delta(\Theta \times A)$. By the portmanteau theorem (see, e.g., Billingsley 1968), $P(Y|s^{\varepsilon,\mathcal{F},\mathcal{G}}) \to \nu(Y)$ along the subnet holds for every $Y \in \mathcal{Y}$ whose boundary has ν-measure zero, and so $\nu(Y) = \mu(Y)$ for all such Y. Since the collection of Y's whose boundaries have ν-measure zero is closed under finite intersections and generates $\mathcal{Y}$,*[10] *ν is the open sequential equilibrium distribution induced by μ.*

---

[6] Indeed, let $\mathcal{T}^0$ be a countable basis for $\mathcal{T}$ and let $\mathcal{B}$ be a neighborhood basis for the essential types. Construct $\mathcal{B}' \subseteq \mathcal{B}$ as follows. First, for each $ik \in L$, include in $\mathcal{B}'$ the set T_{ik}. Also, for each pair of sets U, W in $\mathcal{T}^0$, include in $\mathcal{B}'$, if possible, a set V from $\mathcal{B}$ that is setwise between U and W (i.e., $U \subseteq V \subseteq W$). It is not difficult to show that $\mathcal{B}' \subseteq \mathcal{B}$ is a countable neighborhood basis for the essential types.

[7] That is, $\mathcal{Y}$ is the smallest collection of measurable subsets of $\Theta \times A$ that is closed under countable unions and complements and that contains all sets in $\mathcal{C}$.

[8] See, e.g., Cohn (1980) Corollary 1.6.3.

[9] This conclusion can be shown to hold under the weaker conditions that for each date k: (i) A_{ik} is compact metric and Θ_k is Polish, and (ii) either Θ_k is compact or $p_k(\cdot|\theta_{<k}, a_{<k})$ is weak* continuous in $(\theta_{<k}, a_{<k})$.

[10] The set generates $\mathcal{Y}$, the Borel sigma algebra on $\theta \times A$, because for any outcome (θ, a) it contains all but perhaps countably many of the open balls centered at (θ, a). Hence, it contains a basis for the open sets.

Remark 5. *Continuing with the previous remark, because* ν *is obtained as a weak* limit of* $P(\cdot|s^{\varepsilon,\mathcal{F},\mathcal{G}})$, *player* i*'s equilibrium expected payoff (at* μ*), namely* $\int_{\Theta\times A} u_i(\theta,a)\mu(d(\theta,a))$, *will be equal to* $\int_{\Theta\times A} u_i(\theta,a)\nu(d(\theta,a))$ *so long as* u_i *is a continuous function.*

Remark 6. *It can be shown that if* $\Theta \times A$ *is a compact metric space with its Borel sigma algebra of measurable sets, then* ν *is an open sequential equilibrium distribution iff there is a countable neighborhood basis* $\mathcal{B}$ *for the essential types and a sequence* $\{s^n\}$ *of* $(\varepsilon_n,\mathcal{F}_n)$*-sequential equilibria such that* $\varepsilon_n \to 0$, $\mathcal{B} = \cup_n \mathcal{F}_n$ *and* $P(\cdot|s^n)$ *weak* converges to* ν *as* $n \to \infty$.[11] *So, in many practical settings, one can obtain all the open sequential equilibrium distributions as weak* limits of sequences of* $(\varepsilon,\mathcal{F})$*-sequential equilibrium outcome distributions.*

In any finite multi-stage game (finite A_{ik} and T_{ik}), when $\mathcal{F}$ is fixed and includes every type as a discrete open set, any $(\varepsilon,\mathcal{F})$-sequential equilibrium s^ε satisfies ε sequential rationality with positive probability at each type, and s^ε converges to a Kreps-Wilson sequential equilibrium strategy profile as $\varepsilon \to 0$ (and conversely). Consequently, when $\mathcal{B} = \mathcal{F}$, μ is an open sequential equilibrium conditioned on $\mathcal{B}$ iff a Kreps-Wilson sequential equilibrium assessment (i.e., a consistent and sequentially rational system of beliefs and strategy profile) can be recovered from μ.

4. Examples

Let us consider some examples.

Our first example illustrates a phenomenon that we may call "strategic entanglement," where a sequence of strategy profiles yields a path of randomized play that includes histories with fine details used by later players to correlate their independent actions. When these fine details are lost in the limit because the limit path does not include them, there may be no strategy profile that produces the limit distribution over outcomes.[12] This motivates our choice to base our solution not on strategy profiles – since these are insufficient to capture limit behavior – but on limits of conditional distributions over outcomes.

Example 4.1. *Strategic entanglement in limits of approximate equilibria (Harris-Reny-Robson 1995).*

- On date 1, player 1 chooses $a_1 \in [-1,1]$ and player 2 chooses $a_2 \in \{L,R\}$.

---

[11] This result can also be shown to hold under the weaker conditions given in footnote 9.

[12] Milgrom and Weber (1985) provided the first example of this kind. The example given below has the stronger property that strategic entanglement is unavoidable: it occurs along any sequence of subgame perfect ε-equilibria (i.e., ε-Nash in every subgame) as ε tends to zero.

- On date 2, players 3 and 4 observe the date 1 choices and each choose from $\{L, R\}$.
- For $i \in \{3,4\}$, player i's payoff is $-a_1$ if i chooses L and a_1 if i chooses R.
- If player 2 chooses $a_2 = L$ then player 2 gets $+1$ if $a_3 = L$ but gets -1 if $a_3 = R$;

  if player 2 chooses $a_2 = R$ then player 2 gets -2 if $a_3 = L$ but gets $+2$ if $a_3 = R$.
- Player 1's payoff is the sum of three terms:

  (first term) if 3 and 4 match he gets 0, if they mismatch he gets -10;

  plus (second term) if 2 and 3 match he gets $-|a_1|$, if they mismatch he gets $|a_1|$;

  plus (third term) he gets $-|a_1|^2$.

There is no subgame-perfect equilibrium of this game, but it has an obvious solution which is the limit of strategy profiles where everyone's strategy is arbitrarily close to optimal.

For any $\varepsilon > 0$ and $\alpha > 0$, when players 3 and 4 ε-optimize on $\{a_1 < -\alpha\}$ and on $\{a_1 > \alpha\}$, they must each, with at least probability $1 - \varepsilon/(2\alpha)$, choose L on $\{a_1 < -\alpha\}$ and choose R on $\{a_1 > \alpha\}$.

To prevent player 2 from matching player 3, player 1 should lead 3 to randomize, which 1 can do optimally by randomizing over small positive and negative a_1.

Any setwise-limit distribution over outcomes is only finitely additive, as, for any $\varepsilon > 0$, the events that player 1's action is in $\{a_1 : -\varepsilon < a_1 < 0\}$ or in $\{a_1 : 0 < a_1 < \varepsilon\}$ must each have limiting probability 1/2.

The weak*-limit distribution over outcomes is $a_1 = 0$ and $a_i = 0.5[L] + 0.5[R]$ $\forall i \in \{2,3,4\}$. But in this limit, 3's and 4's actions are perfectly correlated independently of 1's and 2's. So no strategy profile can produce this distribution and we may say that players 3 and 4 are strategically entangled in the limit.[13]

Example 4.2. *Problems of spurious signaling in naïve finite approximations.*

This example illustrates a difficulty that can arise when one tries to approximate a game by restricting players to finite subsets of their action spaces. It can happen that no such "approximation" yields sensible equilibria because new signaling opportunities necessarily arise.

- Nature chooses $\theta \in \{1,2\}$ with $p(\theta) = \theta/3$.

---

[13] Instead of considering limit distributions, a different fix might be to add an appropriate correlation device between periods as in Harris et. al. (1995). But this approach, which is not at all worked out for general multi-stage games, can add equilibria that are not close to any ε-equilibria of the real game (e.g. it enlarges the set of Nash equilibria to the set of correlated equilibria in simultaneous games).

- Player 1 observes $t_1 = \emptyset$ and chooses $a_1 \in [0, 1]$.
- Player 2 observes $t_2 = (a_1)^\theta$ and chooses $a_2 \in \{1, 2\}$.
- Payoffs (u_1, u_2) are as follows:

	$a_2 = 1$	$a_2 = 2$
$\theta = 1$	$(1, 1)$	$(0, 0)$
$\theta = 2$	$(1, 0)$	$(0, 1)$

Consider subgame perfect equilibria of any finite approximate version of the game where player 1 chooses a_1 in some $\hat{A}_1$ that is a finite subset of $[0, 1]$ including at least one $0 < a_1 < 1$. We shall argue that player 1's expected payoff must be 1/3.

Player 1 can obtain an expected payoff of at least 1/3 by choosing the largest feasible $\bar{a}_1 < 1$, as 2 should choose $a_2 = 1$ when $t_2 = \bar{a}_1 > (\bar{a}_1)^2$ indicates $\theta = 1$ (in this finite approximation, player 2 has perfect information after the history $\theta = 1, a_1 = \bar{a}_1$).

Hence, player 1's equilibrium support is contained in $(0, 1)$ since an equilibrium action of 0 or 1 would be uninformative and would lead player 2 to choose $a_2 = 2$ giving player 1 a payoff of 0, contradicting the previous paragraph.

Player 1's expected payoff cannot be more than 1/3, as 1's choice of the smallest $0 < \underline{a}_1 < 1$ in his equilibrium support would lead player 2 to choose $a_2 = 2$ when $t_2 = (\underline{a}_1)^2 < \underline{a}_1$ indicates $\theta = 2$.

But such a scenario cannot be even an approximate equilibrium of the real game, because player 1 could get an expected payoff at least 2/3 by deviating to $\sqrt{\bar{a}_1}$ ($> \bar{a}_1$).

In fact, by reasoning analogous to that in the preceding two sentences, player 1 must receive an expected payoff of 0 in any subgame perfect equilibrium of the infinite game, and so also in any sensibly defined "sequential equilibrium." (It can be shown that player 1's expected payoff is zero in any open sequential equilibrium distribution.)

Hence, approximating this infinite game by restricting player 1 to any large but finite subset of his actions, produces subgame perfect equilibria (and hence also sequential equilibria) that are all far from any sensible equilibrium of the real game.

Example 4.3. *More spurious signaling in finite approximating games (Bargaining for Akerlof's lemons).*

Instead of finitely approximating the players' action sets, one might consider using finite subsets of the players' strategy sets. This example makes use of Akerlof's bargaining game to illustrate a difficulty with this approach.

- First nature chooses θ uniformly from $[0,1]$.
- Player 1 observes $t_1 = \theta$ and chooses $a_1 \in [0,2]$.
- Player 2 observes a_1 and chooses $a_2 \in \{0,1\}$.
- Payoffs are $u_1(a_1, a_2, \theta) = a_2(a_1 - \theta)$, $u_2(a_1, a_2, \theta) = a_2(1.5\theta - a_1)$.

Consider any finite approximate game where player 1 has a given finite set of pure strategies and player 2 observes a given finite partition of $[0,2]$ before choosing a_2 (and so player 2 is restricted to the finite set of strategies that are measurable with respect to this partition).

For any $\delta > 0$, we can construct a function $f : [0,1] \to [0,1.5]$ such that: $f(y) = 0$ $\forall y \in [0,\delta)$, $f(\cdot)$ takes finitely many values on $[\delta,1]$ and, for every $x \in [\delta,1]$, it is the case that $x < f(x) < 1.5x$ and $f(x)$ has probability 0 under each strategy in 1's given finite set.

Then there is a larger finite game (a "better" approximation) where we add the single strategy f for player 1 and give player 2 the ability to recognize each a_1 in the finite range of f. This larger finite game has a perfect equilibrium where player 2 accepts $f(x)$ for any x.

But in the real game this is not an equilibrium because, when 2 would accept $f(x)$ for any x, player 1 could do strictly better by the strategy of choosing $a_1 = \max_{x \in [0,1]} f(x)$ for all θ.

Thus, restricting players to finite subsets of their strategy spaces can fail to deliver approximate equilibrium because important strategies may be left out. We eliminate such false equilibria by requiring approximate optimality among *all* strategies in the original game.

Example 4.4. *Problems of requiring sequential rationality tests with positive probability in all events.*

This example shows that requiring all events to have positive probability for reasons of "consistency" may rule out too many equilibria.

- Player 1 chooses $a_{11} \in \{L, R\}$.
- If $a_{11} = L$, then nature chooses $\theta \in [0,1]$ uniformly; if $a_{11} = R$, then player 1 chooses $a_{12} \in [0,1]$.
- Player 2 then observes $t_2 = \theta$ if $a_{11} = L$, observes $t_2 = a_{12}$ if $a_{11} = R$, and chooses $a_2 \in \{L, R\}$.

- Payoffs (battle of the sexes) are as follows:

	$a_2 = L$	$a_2 = R$
$a_{11} = L$	$(1, 2)$	$(0, 0)$
$a_{11} = R$	$(0, 0)$	$(2, 1)$

All BoS equilibria are reasonable since the choice, θ or a_{12}, from $[0, 1]$ is payoff irrelevant. However, if all events that can have positive probability under some strategies must eventually receive positive probability along a sequence (or net) for "consistency," then the only possible equilibrium payoff is (2,1).

Indeed, for any $x \in [0, 1]$, the event $\{t_2 = x\}$ can have positive probability, but only if positive probability is given to the history ($a_{11} = R$, $a_{12} = x$), because the event $\{\theta = x\}$ has probability 0. So, in any scenario where $P(\{t_2 = x\}) > 0$, player 2 should choose $a_2 = R$ when she observes $t_2 = x$ since the conditional probability of the history ($a_{11} = R$, $a_{12} = x$) is one. But then player 1 can obtain a payoff of 2 with the strategy ($a_{11} = R$, $a_{12} = x$) and so the unique sequential equilibrium payoffs would have to be $(2, 1)$![14]

To allow other equilibria, $(\varepsilon, \mathcal{F})$-sequential equilibrium avoids sequential rationality tests on individual points. With $a_{11} = L$, all open subsets of $T_2 = [0, 1]$ have positive probability and $a_2 = L$ is sequentially rational.

Example 4.5. *Problems from allowing perturbations of nature.*

A different solution to the problem illustrated in the previous example might be to allow perturbations of nature. This example illustrates a difficulty with such an approach.

- Nature chooses $\theta = (\omega_1, \omega_2)$ with ω_1 and ω_2 each drawn independently and uniformly from $[-1, 3]$.
- Player 1 observes $t_1 = \omega_1$ and chooses $a_1 \in \{-1, 1\}$.
- Player 2 observes $t_2 = a_1$ and chooses $a_2 \in \{-1, 1\}$.
- Payoffs are: $u_1(\omega_1, \omega_2, a_1, a_2) = a_1 a_2$; $u_2(\omega_1, \omega_2, a_1, a_2) = \omega_2 a_2$

Since no player receives any information about ω_2, and $E(\omega_2) > 0$, player 2 should choose $a_2 = 1$ regardless of the action of player 1 that she observes. But then player 1 should also choose $a_1 = 1$ regardless of the value of ω_1 that he observes. Hence, the only

---

[14] As in Kreps-Wilson (1982), "consistency" is imposed here by perturbing only the players' strategies, but not nature's probability function. Perturbing also nature's probability function may be worth exploring even though in other examples (e.g., Example 4.5 below) it can have dramatic and seemingly problematic effects on equilibrium play.

sensible equilibrium expected payoff vector is $(u_1, u_2) = (1, 1)$ and this is indeed the only open sequential euqilibrium payoff vector.

But consider the pure strategy profile (s_1, s_2) where $s_1(\omega_1) = 1$ iff $\omega_1 \neq -1$, and $s_2(a_1) = -a_1$.[15]

This profile yields the expected payoff vector $(u_1, u_2) = (-1, 1)$, and can be supported by a perturbation of nature that puts small positive probability on the event $\{(\omega_1, \omega_2) = (-1, -1)\}$. With this perturbation of nature it is sequentially rational for player 2 to choose $a_2 = -1$ when she observes $a_1 = 1$ because she attributes this observation to (ω_1, ω_2) being a mass point on $(-1, -1)$ and therefore expects the value of ω_2 to be -1.

Example 4.6. *Open sequential equilibria may not be subgame perfect if payoffs are discontinuous.*

- Player 1 chooses $a_1 \in [0, 1]$.
- Player 2 observes $t_2 = a_1$ and chooses $a_2 \in [0, 1]$.
- Payoffs are $u_1(a_1, a_2) = u_2(a_1, a_2) = a_2$ if $(a_1, a_2) \neq (1/2, 1/2)$, but $u_1(1/2, 1/2) = u_2(1/2, 1/2) = 2$.

The unique subgame-perfect equilibrium, (s_1, s_2), is pure and has $a_1 = 1/2$, $s_2(1/2) = 1/2$, and $s_2(a_1) = 1$ if $a_1 \neq 1/2$, with the result that payoffs are $u_1 = u_2 = 2$.

But there is an open sequential equilibrium distribution in which player 1 chooses a_1 randomly according to a uniform distribution on $[0, 1]$, and player 2 always chooses $a_2 = 1$, employing the pure strategy $s_2(a_1) = 1 \ \forall a_1 \in [0, 1]$, and so payoffs are $u_1 = u_2 = 1$.

When a_1 has a uniform distribution on $[0, 1]$, the observation that a_1 is in any open neighborhood around $1/2$ would still imply a probability 0 of the event $a_1 = 1/2$, and so player 2 could not increase her conditionally expected utility by deviating from $s_2(a_1) = 1$. And when player 2 always chooses $a_2 = 1$, player 1 has no reason not to randomize.

This failure of subgame perfection occurs because sequential rationality is not being applied at the exact event of $\{a_1 = 1/2\}$, where 2's payoff function is discontinuous. With sequential rationality applied only to open sets, player 2's behavior at $\{a_1 = 1/2\}$ is being justified by the possibility that a_1 was not exactly 1/2 but just very close to it, where she would prefer $a_2 = 1$.

The problem here is caused by the payoff discontinuity at $(a_1, a_2) = (1/2, 1/2)$, which could be endogenous in an enlarged game with continuous payoffs where a subsequent player

---

[15] We abuse our notation here and in the next two examples by denoting a pure strategy for player $ik \in L$ by a measurable function $s_{ik} : T_{ik} \to A_{ik}$. With this notation, for any $t_{ik} \in T_{ik}$, player ik chooses the action $s_{ik}(t_{ik})$ with probability 1.

reacts discontinuously there. To guarantee subgame perfection, even in continuous games, we would need a stronger solution concept, requiring sequential rationality at more than just open sets. (Theorem 5.4 in Section 5.1 shows that, in a large class of games, open sequential equilibrium is compatible with subgame perfection.)

Example 4.7. *Discontinuous responses may admit a possibility of other equilibria (Harris-Stinchcombe-Zame 2000).*

Even when players' payoff and type functions are continuous, discontinuities in strategies can arise in equilibrium. This can allow open sequential equilibrium – which disciplines behavior only on open sets of types, but not at every type – to include outcome distributions that may seem counterintuitive.

- Nature chooses $\theta = (\kappa, \omega) \in \{-1, 1\} \times [0, 1]$. The coordinates κ and θ are independent and uniform.
- Player 1 observes $t_1 = \omega$ and chooses $a_1 \in [0, 1]$.
- Player 2 observes $t_2 = \kappa|a_1 - \omega|$ and chooses $a_2 \in \{-1, 0, 1\}$.
- Payoffs are $u_1(\kappa, \omega, a_1, a_2) = -|a_2|$, $u_2(\kappa, \omega, a_1, a_2) = -(a_2 - \kappa)^2$.

Thus, player 2 should choose the action a_2 that is closest to her expected value of κ, and so player 1 wants to hide information about κ from 2.

In any neighborhood of any $t_2 \neq 0$, player 2 knows $\kappa = 1$ if $t_2 > 0$, and she knows $\kappa = -1$ if $t_2 < 0$, so sequential rationality requires that player 2 use the pure strategy $s_2(t_2) = 1$ if $t_2 > 0$, $s_2(t_2) = -1$ if $t_2 < 0$.

For any $\varepsilon > 0$ and for any finite collection $\mathcal{F}$ of open subsets of player 2's type space $T_2 = [-1, 1]$, there is an $(\varepsilon, \mathcal{F})$-sequential equilibrium in which player 1 hides information about ω with the pure strategy $s_1(\omega) = \omega$, and player 2 plays $s_2(0) = 0$, but $s_2(t_2) = -1$ if $t_2 < 0$, and $s_2(t_2) = 1$ if $t_2 > 0$.[16] This equilibrium seems reasonable, even though 2's behavior is discontinuous at 0.

However, there is another $(\varepsilon, \mathcal{F})$-sequential equilibrium with 2's strategy again discontinuous at $t_2 = 0$, namely: $s_1(\omega) = 1\ \forall \omega$; $s_2(t_2) = 1$ if $t_2 > 0$, $s_2(t_2) = -1$ if $t_2 \leq 0$. This equilibrium may seem less reasonable since justifying (informally) 2's choice here of $a_2 = -1$ when she observes the probability zero event $t_2 = 0$ – i.e., the event $a_1 = \omega$ – requires her to believe that it is more likely that $\kappa = -1$ than that $\kappa = +1$, even though nature's choice of κ was independent of nature's choice of ω and 1's choice of a_1.

But our doubts about this second equilibrium may be due to a presentation effect.[17] If we

---

[16] Since this strategy profile is independent of $(\varepsilon, \mathcal{F})$, the induced distribution over outcomes is an open sequential equilibrium distribution.

[17] We thank Pierre-Andre Chiappori for this observation.

had instead modeled nature with the one-dimensional random variable θ chosen uniformly from $[-2,-1]\cup[1,2]$ and had defined player 1's action set to be $A_1=[1,2]$, the types to be $t_1=|\theta|$, $t_2=(\text{sgn}\theta)|(a_1-|\theta|)|$, and 2's utility to be $u_2=-(a_2-\text{sgn}\theta)^2$, the strategic essence of the game would be unchanged. But now the independence argument is unavailable and so it might not be unreasonable for player 2 to assign more weight to the event $\theta<0$ than to $\theta>0$ (or vice versa) after observing the probability zero event $t_2=0$. So our second equilibrium may not be entirely unreasonable.

Example 4.8. *A Bayesian game where ε-sequential rationality for all types is not possible (Hellman 2014).*

Our final example illustrates why, in $(\varepsilon,\mathcal{F})$-sequential equilibrium, we apply sequential rationality only at finitely many sets of types at a time. It can be impossible to obtain sequential rationality (even ε-sequential rationality) for every type simultaneously.

- There are two players $i\in\{1,2\}$ and one period.
- Nature chooses $\theta=(\kappa,\omega_1,\omega_2)\in\{1,2\}\times[0,1]\times[0,1]$.
- κ is equally likely to be 1 or 2 and it names the player who is "on".
- When $\kappa=i$, ω_i is Uniform [0,1] and $\omega_{-i}=\begin{cases}2\omega_i, & \text{if } \omega_i<1/2\\ 2\omega_i-1, & \text{if } \omega_i\geq 1/2\end{cases}$.

  (This implies ω_{-i} is also Uniform $[0,1]$ when $\kappa=i$.)
- Player types are $t_1=\omega_1$ and $t_2=\omega_2$.
- Action sets are $A_1=A_2=\{L,R\}$.
- Payoffs: When $\kappa=i$, the other player $-i$ just gets $u_{-i}=0$, and u_i is determined by:

if $t_i<1/2$

	$a_{-i}=L$	$a_{-i}=R$
$a_i=L$	0	7
$a_i=R$	3	0

if $t_i\geq 1/2$

	$a_{-i}=L$	$a_{-i}=R$
$a_i=L$	7	0
$a_i=R$	0	3

So $t_i\geq 1/2$ wants to match $-i$ when i is "on" and prefers L if $-i$'s probability of R is less than 0.7; $t_i<1/2$ wants to mismatch $-i$ when i is "on" and prefers L if $-i$'s probability of R is greater than 0.3.

This game has no Bayesian-Nash equilibrium in which the strategic functions $s_i(R|t_i)$ are measurable functions of $t_i \in [0,1]$, by arguments of Simon (2003) and Hellman (2014).[18] Indeed, as shown in Hellman (2014), for any $\varepsilon > 0$ sufficiently small, there are no (measurable) strategies for which almost all types of the two players are ε-optimizing.

But we can construct $(\varepsilon, \mathcal{F})$-sequential equilibria for any $\varepsilon > 0$ and any finite collection $\mathcal{F}$ of open sets of types for 1 and 2. Indeed, choose an integer $m \geq 1$ such that $P(\{t_1 < 2^{-m}\}|C) < \varepsilon \ \forall C \in \mathcal{F} \cap T_1$.

First, let us arbitrarily specify that $s_1(R|t_1) = 0$ for each type t_1 of player 1 such that $t_1 < 2^{-m}$. Then for each type t_i of a player i such that $s_i(R|t_i)$ has just been specified, the types of the other player $-i$ that want to respond to t_i are $t_{-i} = t_i/2$ and $\hat{t}_{-i} = (t_i+1)/2$, and for these types let us specify $s_{-i}(R|t_{-i}) = 1 - s_i(R|t_i)$, $s_{-i}(R|\hat{t}_{-i}) = s_i(R|t_i)$, which is $-i$'s best response there. Continue repeating this step, switching i each time.

This procedure determines $s_i(R|t_i) \in \{0,1\}$ for all t_i that have a binary expansion with m consecutive 0's starting at some odd position for $i = 1$, or at some even position for $i = 2$. Wherever this first happens, if the number of prior 0's is odd then $s_i(R|t_i) = 1$, otherwise $s_i(R|t_i) = 0$. Since the remaining types t_i have probability 0, we can arbitrarily specify $s_i(R|t_i) = 0$ for all these types.[19]

5. Existence

We now introduce a reasonably large class of games within which we are able to establish the existence of both an open sequential equilibrium and an open sequential equilibrium distribution.

Definition 5.1. *Let* $\Gamma = (N, K, A, \Theta, T, \mathcal{M}, \tau, p, u)$ *be a multi-stage game. Then* Γ *is a* regular projective game *iff there is a finite index set* J *and sets* Θ_{kj}, A_{ikj} *such that, for every* $ik \in L$

R.1. $\Theta_k = \times_{j \in J} \Theta_{kj}$ *and* $A_{ik} = \times_{j \in J} A_{ikj}$,

R.2. *there exist sets* $M_{0ik} \subset \{1, \ldots, k\} \times J$ *and* $M_{1ik} \subset N \times \{1, \ldots, k-1\} \times J$, *such that* $T_{ik} = ((\times_{hj \in M_{0ik}} \Theta_{hj}) \times (\times_{nhj \in M_{1ik}} A_{nhj}))$ *and* $\tau_{ik}(\theta_{\leq k}, a_{<k}) = ((\theta_{hj})_{hj \in M_{0ik}}, (a_{nhj})_{nhj \in M_{1ik}})$ $\forall (\theta_{\leq k}, a_{<k})$ *is a projection map, that is,* i*'s type at date* k *is just a list of state coordinates and action coordinates from dates up to* k*,*[20]

---

[18] Nature's probability function does not satisfy the information diffusness assumption of Migrom and Weber (1985) so their existence theorem does not apply.

[19] The resulting strategies are measurable because, by construction, they are constant on each of the countably many intervals of types involved in the iterative construction as well as on the complementary (hence measurable) remainder set of types of measure zero.

[20] Perfect recall implies that for all players $i \in N$, for all dates $h < k$, and for all $j \in J$, $M_{0ih} \subseteq M_{0ik}$, $M_{1ih} \subseteq M_{1ik}$, and $ihj \in M_{1ik}$.

R.3. *Θ_{kj} and A_{ikj} are nonempty compact metric spaces $\forall j \in J$ (with all spaces, including products, given their Borel sigma-algebras),*

R.4. *$u_i : \Theta \times A \to \mathbb{R}$ is continuous,*

R.5. *there is a continuous nonnegative density function $f_k : \Theta_{\leq k} \times A_{<k} \to [0, \infty)$ and for each j in J, there is a probability measure ρ_{kj} on $\mathcal{M}(\Theta_{kj})$ such that $p_k(B|\theta_{<k}, a_{<k}) = \int_B f_k(\theta_k|\theta_{<k}, a_{<k})\rho_k(d\theta_k)$ $\forall B \in \mathcal{M}(\Theta_k)$, $\forall(\theta_{<k}, a_{<k}) \in \Theta_{<k} \times A_{<k}$, where $\rho_k = \times_{j\in J}\rho_{kj}$ is a product measure.*

If Γ satisfies R.2, we may say that Γ is a projective game *or a* game with projected types.

Remark 7. *(1) One can always reduce the cardinality of J to $(K+1)|N|$ or less by grouping, for any $ik \in L$, the variables $\{a_{ikj}\}_{j\in J}$ and $\{\theta_{kj}\}_{j\in J}$ according to the $|N|$-vector of dates at which each player observes them, if ever.*
(2) Regular projective multi-stage games can include all finite multi-stage games (simply by letting each player's type be a coordinate of the state).
(3) Since distinct players can observe the same θ_{kj}, nature's probability function in a regular projective multi-stage game need not satisfy the information diffuseness assumption of Milgrom-Weber (1985).
(4) Under the continuous utility function assumption R.4, our convention of history-independent action sets is no longer without loss of generality (see footnote 3).

In a regular projective game, define $\mathcal{B}^* \subseteq \mathcal{T}^*$ so that $B \in \mathcal{B}^* \cap \mathcal{T}_{ik}$ iff: $\exists a \in A$ such that $P_T(B|a) > 0$, and $B = (\times_{(h,j)\in M_{0ik}} B_{0hj}) \times (\times_{(n,h,j)\in M_{1ik}} B_{nhj})$, where each B_{0hj} is an open subset of Θ_{hj} and each B_{nhj} is an open subset of A_{nhj}. Then $\mathcal{B}^*$ is a neighborhood basis for the essential types in the game and we may call $\mathcal{B}^*$ the *product basis.*

A *product partition* of $\Theta \times A$ is a partition in which every element is a product of Borel subsets of the Θ_{kj} and A_{ikj} sets.

For any $ik \in L$, for any $C \subseteq T_{ik}$, recall from Section 3 that $\langle C\rangle = \{(\theta, a) \in \Theta \times A : \tau_{ik}(\theta_{\leq k}, a_{<k}) \in C\}$ is the set of outcomes that would yield types in C.

Remark 8. *For any $\mathcal{F}$ that is a finite subset of $\mathcal{B}^*$, there exists a finite product partition Q of $\Theta \times A$ such that for any $C \in \mathcal{F}$, $\langle C\rangle$ is a union of elements of Q.*

Theorem 5.2. *Let Γ be a regular projective game and let Q be any finite product partition of $\Theta \times A$. Let $\mathcal{F}$ be a finite subset of $\mathcal{T}^*$ such that for any $C \in \mathcal{F}$, $\langle C\rangle$ is a union of elements of Q. Then for any $\varepsilon > 0$, Γ has an $(\varepsilon, \mathcal{F})$-sequential equilibrium.*

Theorem 5.3. *Every regular projective game* Γ *has an open sequential equilibrium* μ *conditioned on* $\mathcal{B}^*$*. Moreover, every open sequential equilibrium* μ *of* Γ *induces an open sequential equilibrium distribution* ν*, and so* Γ *also has an open sequential equilibrium distribution.*

5.1. Subgame Perfection

In light of Example 4.6, we provide here a result showing that open sequential equilibrium and subgame perfection are mutually compatible for a large class of games.

Let Γ be any multi-stage game. For any $s \in S$ and for any date-k history $(\theta^0_{\leq k}, a^0_{<k}) \in \Theta_{\leq k} \times A_{<k}$, define player i's expected utility of s conditional on $(\theta^0_{\leq k}, a^0_{<k})$ by

$$U_i(s|\theta^0_{\leq k}, a^0_{<k}) = \int_{\Theta \times A} u_i(\theta, a) P(d(\theta, a)|(a^0_{<k}, s_{\geq k}); (\theta^0_{\leq k}, p_{>k})),$$

where $(a^0_{<k}, s_{\geq k}) \in S$ denotes the strategy profile in which the profile of actions $a^0_{<k} \in A_{<k}$ is chosen with probability 1 by all types of the players from dates 1 to $k-1$ and the strategy profile employed in dates after k is $s_{\geq k}$, and where $(\theta^0_{\leq k}, p_{>k})$ denotes the probability function of nature in which the state vector up to date k, $\theta^0_{\leq k} \in \Theta_{\leq k}$, is chosen with probability 1 independently of the history and then $p_{>k}$ is employed to choose the state vector for dates after k. Note that $U_i(s|\theta^0_{\leq k}, a^0_{<k})$ is well-defined whether or not the history $(\theta^0_{\leq k}, a^0_{<k})$ occurs with positive probability under s.

Say that a date-k history $(\theta_{\leq k}, a_{<k}) \in \Theta_{\leq k} \times A_{<k}$ is a *subgame* of Γ iff $\tau^{-1}_{ik}(\tau_{ik}(\theta_{\leq k}, a_{<k})) = \{(\theta_{\leq k}, a_{<k})\}$ $\forall i \in N$. For any $\varepsilon > 0$, a strategy $s \in S$ is an *ε-subgame perfect equilibrium* of Γ iff for every $ik \in L$ and for every subgame $(\theta_{\leq k}, a_{<k}) \in \Theta_{\leq k} \times A_{<k}$, $U_i(r_i, s_{-i}|\theta_{\leq k}, a_{<k}) \leq U_i(s|\theta_{\leq k}, a_{<k}) + \varepsilon$ for every date-k continuation $r_i \in S_i$ of s_i.

Say that a mapping $\mu : \mathcal{Y} \times \mathcal{B} \to [0, 1]$ is a *subgame perfect open sequential equilibrium* (conditioned on $\mathcal{B}$) iff $\mathcal{B}$ is a neighborhood basis for the essential types, and, for every $\varepsilon > 0$, for every finite subset $\mathcal{F}$ of $\mathcal{B}$, and for every finite subset $\mathcal{G}$ of $\mathcal{Y}$, there is an $(\varepsilon, \mathcal{F})$-sequential equilibrium s such that s is ε-subgame perfect and,

$$|P(Y|C, s) - \mu(Y|C)| < \varepsilon, \text{ for every } (Y, C) \in \mathcal{G} \times \mathcal{F}.$$

In a projective game, if some history up to date k, $(\theta_{\leq k}, a_{<k})$, is a subgame, then all players at date k always observe all of nature's states from dates 1 to k and always observe all players' actions from dates 1 to $k-1$. Hence, all histories up to date k are subgames and so we may say that date k is a *subgame date.* We will make use of this insight in the proof of Theorem 5.4.

Theorem 5.4. *Every regular projective game* Γ *has a subgame perfect open sequential equilibrium* μ *conditioned on* $\mathcal{B}^*$.

6. Proofs

Proof of Theorem 3.3. For each $\varepsilon > 0$ and for each finite subset $\mathcal{F}$ of $\mathcal{B}$, by hypothesis (and the axiom of choice) we may choose an $(\varepsilon, \mathcal{F})$-sequential equilibrium $s^{\varepsilon,\mathcal{F}} \in S$. For any $(\varepsilon, \mathcal{F})$ and for any $Y \in \mathcal{Y}$, $P(Y|C, s^{\varepsilon,\mathcal{F}})$ is defined for every $C \in \mathcal{F}$. Extend $P(Y|\cdot, s^{\varepsilon,\mathcal{F}})$ to all of $\mathcal{B}$ by defining

$$\bar{P}(Y|C, s^{\varepsilon,\mathcal{F}}) = \begin{cases} P(Y|C, s^{\varepsilon,\mathcal{F}}), & \text{if } C \in \mathcal{F} \\ 0, & \text{if } C \in \mathcal{B}\backslash\mathcal{F}. \end{cases}$$

Then, $\{\bar{P}(\cdot|\cdot, s^{\varepsilon,\mathcal{F}})\}$ is a net in $[0,1]^{\mathcal{Y}\times\mathcal{B}}$, with smaller positive numbers ε and larger finite subsets $\mathcal{F}$ of $\mathcal{B}$ being further out in the (directed) index set. By Tychonoff's theorem, $[0,1]^{\mathcal{Y}\times\mathcal{B}}$ is compact and so there exists $\mu \in [0,1]^{\mathcal{Y}\times\mathcal{B}}$ and a subnet $\{\bar{P}(\cdot|\cdot, s^{\varepsilon_\alpha,\mathcal{F}_\alpha})\}$ that converges to μ.

The convergence (under the product topology) to μ of the subnet implies that for every $\varepsilon > 0$, for every finite subset $\mathcal{F}$ of $\mathcal{B}$ and for every finite subset $\mathcal{G}$ of $\mathcal{Y}$, there exists α such that $\varepsilon_\alpha < \varepsilon$, $\mathcal{F}_\alpha \supseteq \mathcal{F}$, and

$$\left|\bar{P}(Y|C, s^{\varepsilon_\alpha,\mathcal{F}_\alpha}) - \mu(Y|C)\right| < \varepsilon, \text{ for every } (Y, C) \in \mathcal{G} \times \mathcal{F}.$$

Since $\bar{P}(Y|C, s^{\varepsilon_\alpha,\mathcal{F}_\alpha}) = P(Y|C, s^{\varepsilon_\alpha,\mathcal{F}_\alpha})$ when $C \in \mathcal{F}_\alpha$, and since $s^{\varepsilon_\alpha,\mathcal{F}_\alpha}$, being an $(\varepsilon_\alpha, \mathcal{F}_\alpha)$-sequential equilibrium is, a fortiori, an $(\varepsilon, \mathcal{F})$-sequential equilibrium, we conclude that μ is an open sequential equilibrium. Q.E.D.

Lemma 6.1. . *Suppose that all* Θ_k, A_{ik} *are metric spaces with their Borel* σ*-algebras, and all* $\tau_{ik} : \Theta_{\leq k} \times A_{<k} \to T_{ik}$ *and all* $p_k : \Theta_{<k} \times A_{<k} \to \Delta(\Theta_k)$ *are continuous, with product topologies on all product sets and the weak* topology on* $\Delta(\Theta_k)$. *If* $C \subseteq T_{ik}$ *is open and* $P_T(C|s) > 0$ *for some* $s \in S$, *then there exists* $a \in A$ *such that* $P_T(C|a) > 0$.

Proof of Lemma 6.1. Consider any $ik \in L$ and any open subset C of T_{ik} and suppose there exists $s \in S$ such that $P_T(C|s) > 0$. We wish to show that there exists $\hat{a} \in A$ such that $P_T(C|\hat{a}) > 0$. For this, it suffices to find a nonnegative function $g : \Theta \times A \to [0, \infty)$ that is positive only on those outcomes that yield types in C and that satisfies $\int g(\theta, a)P(d(\theta, a)|\hat{a}) > 0$.

There are two steps to the proof. The first step obtains a nonnegative function $g : \Theta \times A \to [0, \infty)$ that is positive only on outcomes (θ, a) that yield types in C, i.e., only on $\langle C \rangle$, and that satisfies $\int g(\theta, a) P(d(\theta, a)|s) > 0$. The second step establishes inductively that for each date $k \in \{2, ..., |K|\}$: If there exists $\hat{a}_{>k} \in A_{>k}$ such that

$$\int g(\theta, a) P(d(\theta, a)|(s_{\leq k}, \hat{a}_{>k})) > 0,$$

then there exists $\hat{a}_{k-1} \in A_{k-1}$ such that

$$\int g(\theta, a) P(d(\theta, a)|(s_{\leq k-1}, \hat{a}_{>k-1})) > 0. \tag{6.1}$$

These two steps suffice because if $\int g(\theta, a) P(d(\theta, a)|s) > 0$ is true, then the hypothesis in the induction step (6.1) is trivially true for $k = |K|$ and so we may apply (6.1) iteratively $|K|$ times to obtain $\hat{a} \in A$ such that $\int g(\theta, a) P(d(\theta, a)|\hat{a}) > 0$.

First Step. Let $Z = \{(\theta, a) : \tau_{ik}(\theta_{\leq k}, a_{<k}) \in C\}$, i.e., $Z = \langle C \rangle$. Then $P(Z|s) = P_T(C|s) > 0$ and Z is an open subset of $\Theta \times A$ because τ_{ik} is continuous. Choose (θ_0, a_0) in the intersection of Z and the support of $P(\cdot|s)$. Since $\Theta \times A$ is a metric space, we may define define $g(\theta, a)$ to be the distance from (θ, a) to the closed set $(\Theta \times A) \backslash Z$. Then $\int g(\theta, a) P(d(\theta, a)|s) > 0$ because the nonnegative continuous function g is positive at the point (θ_0, a_0) that is in the support of $P(\cdot|s)$. Moreover, g is positive only on Z.

Second Step. For any date $k < |K|$, for any $\bar{a} \in A$ and for any $\bar{\theta}_{\leq k} \in \Theta_{\leq k}$, let $\bar{p}_{>k}(\cdot|\bar{\theta}_{\leq k}, \bar{a})$ denote the probability measure on $\Theta_{>k}$ that is determined by $(p_{k+1}, ..., p_K)$, i.e., for any $B = B_{k+1} \times ... \times B_{|K|} \in \mathcal{M}(\Theta_{>k})$, define $\bar{p}_{>k}(B|\bar{\theta}_{\leq k}, \bar{a})$ to be equal to:

$$\int_B p_{|K|}(d\theta_K|\theta_{>k}, \bar{\theta}_{\leq k}, \bar{a}_{<K}) p_{|K|-1}(d\theta_{|K|-1}|(\theta_j)_{k<j<|K|}, \bar{\theta}_{\leq k}, \bar{a}_{<|K|-1}) ... p_{k+1}(d\theta_{k+1}|\bar{\theta}_{\leq k}, \bar{a}_{\leq k}).$$

The assumed weak* continuity of each of nature's functions $p_1, ..., p_K$ implies that $\bar{p}_{>k}(\cdot|\theta_{\leq k}, a)$ is weak* continuous in $(\theta_{\leq k}, a)$.

Suppose that there exists $\hat{a}_{>k}$ such that $\int g(\theta, a) P(d(\theta, a)|(s_{\leq k}, \hat{a}_{>k})) > 0$. We must show that there exists $\hat{a}_{k-1} \in A_{k-1}$ such that

$$\int g(\theta, a) P(d(\theta, a)|(s_{\leq k-1}, \hat{a}_{>k-1})) > 0. \tag{6.2}$$

The positive integral $\int g(\theta, a) P(d(\theta, a)|(s_{\leq k}, \hat{a}_{>k}))$ can be rewritten as,

$$\int h(\theta_{\leq k}, a_{\leq k}) s_{\cdot k}(da_k|\theta_{\leq k}, a_{<k}) \Phi_{\leq k}(d(\theta_{\leq k}, a_{<k})|s_{\leq k-1}) > 0, \tag{6.3}$$

where $h(\theta_{\leq k}, a_{\leq k}) = \int g(\theta, a_{\leq k}, \hat{a}_{>k}) \bar{p}_{>k}(d\theta_{>k}|\theta_{\leq k}, a_{\leq k}, \hat{a}_{>k})$ is continuous (by the weak* continuity of $\bar{p}_{>k}(\cdot|\theta_{\leq k}, a_{\leq k}, \hat{a}_{>k}))$ and nonnegative, and where $\Phi_{\leq k}(\cdot|s_{\leq k-1})$ is the marginal of $P(\cdot|s)$ on $\Theta_{\leq k} \times A_{<k}$.

We claim that there exists $\hat{a}_k \in A_k$ such that,

$$\int h(\theta_{\leq k}, a_{<k}, \hat{a}_k)\Phi_{\leq k}(d(\theta_{\leq k}, a_{<k})|s_{\leq k-1}) > 0. \tag{6.4}$$

Indeed, if there is no such $\hat{a}_k$, then because h is continuous and nonnegative, h must be identically zero on (support of $\Phi_{\leq k}$) $\times A_k$. But this contradicts (6.3), proving the claim.

The proof is complete by noting that the left-hand side of (6.4) is equal to left-hand side of (6.2). Q.E.D.

Proof of Theorem 5.2. For any $(\theta, a) \in \Theta \times A$, let

$$f(\theta, a) = \Pi_{k\in K} f_k(\theta_k|\theta_{<k}, a_{<k}),$$

where we define $f_1(\theta_1|\theta_{<1}, a_{<1}) = f_1(\theta_1)$.

Let ε be any strictly positive real number, let Q be any finite product partition of $\Theta \times A$, and let $\mathcal{F}$ be any finite subset of $\mathcal{T}^*$ such that, for any $C \in \mathcal{F}$, $\langle C\rangle$ is a union of elements of Q. We must show that Γ has an $(\varepsilon, \mathcal{F})$-sequential equilibrium.

For each of the finitely many events C in $\mathcal{F}$ choose an action $a \in A$ such that $P_T(C|a) > 0$ and let A^0 denote the finite set of all of these actions. Hence, $\max_{a\in A^0} P_T(C|a) > 0, \forall C \in \mathcal{F}$, and so we may define $\gamma > 0$ by $\gamma = \min_{C\in\mathcal{F}} \max_{a\in A^0} P_T(C|a)$. Since adding actions to A^0 can only increase γ, we may assume without loss of generality that A^0 is a product, i.e., that $A^0 = \times_{ik\in L, j\in J} A^0_{ikj}$ where each A^0_{ikj} is a finite subset of A_{ikj}. Hence,

$$\max_{a\in A^0} P_T(C|a) \geq \gamma > 0, \forall C \in \mathcal{F}. \tag{6.5}$$

Since payoffs are bounded, we may choose a number v so that,

$$v > \max_{i\in N,(\theta,a),(\theta',a')\in\Theta\times A} (u_i(\theta, a) - u_i(\theta', a')). \tag{6.6}$$

Because the number of periods of the game, $|K|$, is finite,[21] we may choose $\beta \in (0, 1)$ so that,

$$(1 - (1 - \beta)^{|K|})v < \varepsilon/2. \tag{6.7}$$

---

[21] Games with a countable infinity of periods can be handled by including the assumption that for any $\varepsilon > 0$ there is a positive integer n such that the history of play over the first n periods determines each player's payoff within ε (e.g., games with discounting).

Let $m = \max_{ik \in L} |A^0_{ik}|$ and choose $\eta > 0$ so that,

$$\eta < (\beta/m)^{|L|}\gamma\varepsilon/2. \tag{6.8}$$

Since Q is a finite product partition of $\Theta \times A$, it can be written as $Q = (\times_{k\in K, j\in J} Q_{\Theta_{kj}}) \times (\times_{ik\in L, j\in J} Q_{A_{ikj}})$, for some finite Borel measurable partitions $Q_{\Theta_{kj}}$ of Θ_{kj} and $Q_{A_{ikj}}$ of A_{ikj} $\forall (ik, j) \in L \times J$.

By the continuity of each player's utility function on the compact set $\Theta \times A$ and of f on the compact set $\Theta \times A$, there are sufficiently fine finite refinements $Q^1_{\Theta_{kj}}$ of $Q_{\Theta_{kj}}$ and $Q^1_{A_{ikj}}$ of $Q_{A_{ikj}}$ $\forall (i, k, j) \in N \times K \times J$, such that each element of $Q^1_{A_{ikj}}$ contains at most one action in A^0_{ikj}, and such that for any (θ, a) and (θ', a') in the same element of the partition $(\times_{k\in K, j\in J} Q^1_{\Theta_{kj}}) \times (\times_{ik\in L, j\in J} Q^1_{A_{ikj}})$ of $\Theta \times A$,

$$|u_i(\theta, a)f(\theta, a) - u_i(\theta', a')f(\theta', a')| \le \eta, \ \forall i \in N. \tag{6.9}$$

Let $Q^1 = (\times_{k\in K, j\in J} Q^1_{\Theta_{kj}}) \times (\times_{ik\in L, j\in J} Q^1_{A_{ikj}})$. Then Q^1 is a product partition of $\Theta \times A$ and a refinement of Q.

For each $(ik, j) \in L \times J$ and from each element of the partition $Q^1_{A_{ikj}}$ of A_{ikj}, choose precisely one action, where the action that is chosen is from the set A^0_{ikj} whenever possible. Letting A^1_{ikj} denote the set of all of the chosen actions, we have $A^1_{ikj} \supseteq A^0_{ikj}$ because each $Q^1_{A_{ikj}}$ contains at most one action in A^0_{ikj}.

Consider the finite extensive form game with perfect recall that results when for every $ik \in L$, player i's set of date-k strategies is restricted to those in S_{ik} that are measurable with respect to Q^1 and that give positive probability only to actions in the finite set A^1_{ik}. Call this finite game with perfect recall Γ_0. Now restrict the strategies in Γ_0 further so that for any $ik \in L$ and regardless of ik's type, player i's date k strategy must choose a uniform distribtion over $A^0_{ik} \subseteq A^1_{ik}$ with probability at least β. This more restricted game, which we will call Γ_β, is finite and has perfect recall.

The Q^1-measurability condition means that for any action-coordinate value a_{ikj} that a player observes in the (regular projective) infinite game, he observes (can condition on) in Γ_0 and Γ_β only the partition element in $Q^1_{A_{ikj}}$ that contains it, and for any state-coordinate value θ_{kj} that a player observes in the infinite game, he observes (can condition on) in Γ_0 and Γ_β only the partition element in $Q^1_{\Theta_{kj}}$ that contains it. Hence in both Γ_0 and Γ_β, a type w_{ik} of player ik is any $(\times_{hj\in M_{0ik}} q^1_{hj}) \times (\times_{nhj\in M_{1ik}} q^1_{nhj})$, where $q^1_{hj} \in Q^1_{\Theta_{hj}}$ $\forall hj \in M_{0ik}$ and $q^1_{nhj} \in Q^1_{A_{ihj}}$ $\forall nhj \in M_{1ik}$. Let W_{ik} denote the common finite set of ik's types in the finite games Γ_0 and Γ_β. Then W_{ik} is a finite partition of T_{ik}.

Let $s^* \in S$ be a Nash equilibrium of the agent normal form of Γ_β. Then for any $ik \in L$ and for any $w_{ik} \in W_{ik}$ such that $P_T(w_{ik}|s^*) > 0$ and for any $t_{ik} \in w_{ik}$, $s^*_{ik}(\cdot|t_{ik}) \in \Delta(A^1_{ik})$ places no more than total probability β on actions that are suboptimal among all actions in A^1_{ik}. Therefore, (6.7) implies that for any $ik \in L$ and for any $w_{ik} \in W_{ik}$ such that $P_T(w_{ik}|s^*) > 0$;

$$\text{In } \Gamma_0,\ s^*_i \text{ is } \varepsilon/2\text{-optimal for player } i \text{ against } s^*_{-i} \text{ given } w_{ik}. \tag{6.10}$$

We will show that s^* is an $(\varepsilon, \mathcal{F})$-sequential equilibrium of Γ. That is, we will show that for every $ik \in L$ and every $C \in \mathcal{F} \cap \mathcal{T}_{ik}$,

(a) $P_T(C|s^*) > 0$, and

(b) $U_i(r_i, s^*_{-i}|C) \leq U_i(s^*|C) + \varepsilon$ for every date-k continuation $r_i \in S_i$ of s^*_i.

Consider any $ik \in L$ and any $C \in \mathcal{F} \cap \mathcal{T}_{ik}$. Since each s^*_{ik} places probability at least β/m on each element of A^0_{ik}, s^* places probability at least $(\beta/m)^{|L|}$ on each $a \in A^0$. Hence, (6.5) implies that,

$$P_T(C|s^*) \geq (\beta/m)^{|L|}\gamma > 0,\ \forall C \in \mathcal{F} \cap \mathcal{T}_{ik},\ \forall ik \in L. \tag{6.11}$$

This proves (a). We now turn to (b).

Fix, for the remainder of the proof, any $ik \in L$ and any $C \in \mathcal{F} \cap \mathcal{T}_{ik}$.

Because $C \in \mathcal{F}$, its set of outcomes $\langle C \rangle$ is a union of elements of Q. Together with condition (R.2), this implies that C is the disjoint union of sets of the form $(\times_{hj \in M_{0ik}} q_{hj}) \times (\times_{nhj \in M_{1ik}} q_{nhj})$, where each q_{hj} is an element of $Q_{\Theta_{hj}}$ and where each q_{nhj} is an element of $Q_{A_{nhj}}$. On the other hand, because Q^1 refines Q, each q_{hj} is a union of elements q^1_{hj} of $Q^1_{\Theta_{hj}}$ and each q_{nhj} is the union of elements q^1_{nhj} of $Q^1_{A_{ihj}}$. Hence, C is a union of sets of the form $(\times_{hj \in M_{0ik}} q^1_{hj}) \times (\times_{nhj \in M_{1ik}} q^1_{nhj})$, each of which is a type of player ik in the finite game. Consequently,

$$C \text{ is the disjoint union of types of player } ik \text{ in the finite game.} \tag{6.12}$$

Let r_i be any strategy for player i in the original infinite game that is a date-k continuation of s^*_i. We must show that (b) holds.

Define the date-k continuation strategy $r'_i \in S_i$ of s^*_i as follows. For any $h < k$, define $r'_{ih} = s^*_{ih}$. For any $h \geq k$, for any $w_{ih} \in W_{ih}$, for any $t_{ih} \in w_{ih}$, and for any $a^1_{ih} \in A^1_{ih}$, let $q_{ih}(a^1_{ih})$ denote the element of the partition $\times_{j \in J} Q^1_{A_{ihj}}$ of A_{ih} that contains a^1_{ih} and define $r'_{ih}(a^1_{ih}|t_{ih})$ so that,[22]

$$r'_{ih}(a^1_{ih}|t_{ih}) P_T(w_{ih}|(r_i, s^*_{-i}); \rho) = \int_{(\theta,a) \in \langle w_{ih} \rangle} r_{ih}(q_{ih}(a^1_{ih})|\tau_{ih}(\theta_{\leq h}, a_{<h})) P(d(\theta,a)|(r_i, s^*_{-i}); \rho).$$

---

[22] Recall from Section 2.1 that $P(\cdot|s;\rho)$ is the probability measure over outcomes when the strategy profile is s and nature's probability function is ρ.

This defines $r'_{ih}(\cdot|t_{ih}) \in \Delta(A^1_{ih})$ uniquely when $P_T(w_{ih}|(r_i, s^*_{-i}); \rho) > 0$ and we may define $r'_{ih}(\cdot|t_{ih})$ to be constant in t_{ih} on w_{ih} and equal to any element of $\Delta(A^1_{ih})$ when $P_T(w_{ih}|(r_i, s^*_{-i}); \rho) = 0$. Because $r'_{ih}(\cdot|t_{ih}) \in \Delta(A^1_{ih})$ is constant for $t_{ih} \in w_{ih}$, the strategy $r'_i \in S_i$ is feasible for the finite game Γ_0.

Because s^* is measurable with respect to Q^1, the definition of r'_i yields the following:

> The distribution over the elements of the finite partition Q^1
> is the same under each of the two probability measures
> $P(\cdot|(r_i, s^*_{-i}); \rho)$ and $P(\cdot|(r'_i, s^*_{-i}); \rho)$ on $\Theta \times A$.[23] (6.13)

Then,

$$\begin{aligned}
& U_i(r_i, s^*_{-i}|C) \\
= \;& \frac{\int_{\langle C\rangle} u_i(\theta,a) P(d(\theta,a)|(r_i, s^*_{-i}); p)}{P_T(C|s^*; p)} \\
= \;& \frac{\int_{\langle C\rangle} u_i(\theta,a) f(\theta,a) P(d(\theta,a)|(r_i, s^*_{-i}); \rho)}{P_T(C|s^*; p)}, \quad \text{since by R.5 } p \text{ has density } f \text{ and carrying measure } \rho \\
\le \;& \frac{\int_{\langle C\rangle} u_i(\theta,a) f(\theta,a) P(d(\theta,a)|(r'_i, s^*_{-i}); \rho)}{P_T(C|s^*; p)} + \frac{\eta}{P_T(C|s^*; p)}, \quad \text{by (6.13) and (6.9)} \\
= \;& \frac{\int_{\langle C\rangle} u_i(\theta,a) P(d(\theta,a)|(r'_i, s^*_{-i}); p)}{P_T(C|s^*; p)} + \frac{\eta}{P_T(C|s^*; p)}, \quad \text{since by R.5 } p \text{ has density } f \text{ and carrying measure } \rho \\
= \;& U_i(r'_i, s^*_{-i}|C) + \frac{\eta}{P_T(C|s^*; p)}, \qquad (6.15) \\
\le \;& U_i(s^*|C) + \frac{\varepsilon}{2} + \frac{\eta}{P_T(C|s^*; p)}, \quad \text{by (6.10) and since } C \text{ is a union of types for } ik \text{ in the finite game by (6.12)}
\end{aligned}$$

---

[23] This requires perfect recall of player i, the ρ-independence of the coordinates of θ, and the property of Q^1 that: $\forall a, a' \in A^1$, $\forall (nh, j) \in L \times J$, if a_{nhj} and a'_{nhj} are in the same element of $Q^1_{A_{nhj}}$, then $a_{nhj} = a'_{nhj}$.

$$\leq \; U_i(s^*|C) + \frac{\varepsilon}{2} + \frac{\eta}{(\beta/m)^{|L|}\gamma}, \text{ by (6.11), where } P_T(C|s^*) = P_T(C|s^*;p)$$

$$\leq \; U_i(s^*|C) + \varepsilon, \text{ given the choice of } \eta \text{ in (6.8). Q.E.D.}$$

Proof of Theorem 5.3. Since $\Theta \times A$ is a compact metric space it suffices, by Remark 4, to show that an open sequential equilibrium conditioned on $\mathcal{B}^*$ exists. But this follows from Theorem 3.3 because, by Remark 8 and Theorem 5.2, for any $\varepsilon > 0$ and for any finite subset $\mathcal{F}$ of $\mathcal{B}^*$, there exists an $(\varepsilon, \mathcal{F})$-sequential equilibrium of Γ. Q.E.D.

Proof of Theorem 5.4. The proof has two steps.
Step 1. Let ε be any strictly positive real number and let $\mathcal{F}$ be any finite subset of $\mathcal{B}^*$. In this first step, we will show that there exists $s^* \in S$ that is both an ε-subgame perfect equilibrium and an $(\varepsilon, \mathcal{F})$-sequential equilibrium.

For any $(\theta, a) \in \Theta \times A$, and for any $k \in K$, let

$$f^k(\theta, a) = \Pi_{h>k} f_h(\theta_h|\theta_{<h}, a_{<h}),$$

where we define the product over the empty set to be 1, and so $f^{|K|}(\theta, a) = 1$.

As noted in Remark 8 there exists a finite product partition $Q = (\times_{k\in K, j\in J} Q_{\Theta_{kj}}) \times (\times_{ik\in L, j\in J} Q_{A_{ikj}})$ of $\Theta \times A$ such that for any $C \in \mathcal{F}$, $\langle C \rangle$ is a union of elements of Q.

With these ε, $\mathcal{F}$, and Q, follow the proof of Theorem 5.2 up to the point where s^* is about to be defined, but replace (6.9) with stronger condition: $\forall(\theta, a), (\theta', a')$ in the same element of the partition $(\times_{k\in K, j\in J} Q^1_{\Theta_{kj}}) \times (\times_{ik\in L, j\in J} Q^1_{A_{ikj}})$,

$$\left| u_i(\theta, a) f^k(\theta, a) - u_i(\theta', a') f^k(\theta', a') \right| \leq \eta, \; \forall ik \in L, \tag{6.16}$$

and choose $\eta > 0$ so that (6.8) holds and also so that $\eta < \varepsilon/4$.

Thus, we have v satisfying (6.6), $\beta > 0$ satisfying (6.7), $\eta \in (0, \varepsilon/4)$ satisfying (6.8), a refinement $Q^1 = (\times_{k\in K, j\in J} Q^1_{\Theta_{kj}}) \times (\times_{ik\in L, j\in J} Q^1_{A_{ikj}})$ of Q satisfying (6.16), finite subsets $A^0 \subseteq A^1$ of A, a finite partition W_{ik} of T_{ik} for each $ik \in L$, a finite game Γ_0 that is obtained by restricting players to strategies that are measurable with respect to Q^1 and that give positive probability only to actions in A^1, and a finite game Γ_β that further restricts the strategies so that for any $ik \in L$ and regardless of ik's type, player i's date k strategy must choose a uniform distribtion over $A^0_{ik} \subseteq A^1_{ik}$ with probability at least β.

Since $\beta \in (0, 1)$ satisfies (6.7), the inequalities

$$\beta + \zeta < 1 \text{ and } (1 - (1 - \beta - \zeta)^{|K|})v < \varepsilon/2, \tag{6.17}$$

hold for every $\zeta > 0$ small enough.

For any date k, let Θ_k^1 be a finite subset of Θ_k that contains precisely one state from each element of the finite partition $\times_{,j\in J} Q_{\Theta_{kj}}$ of Θ_k. Choose any $\zeta > 0$ satisfying (6.17) and perturb the finite game Γ_β so that at every date k, with independent probability ζ nature chooses θ_k uniformly from Θ_k^1, and further restrict the players' strategies so that for any $ik \in L$ and regardless of player ik's type, player i's date k strategy must choose a uniform distribtion over A_{ik}^1 with probability at least ζ. This perturbed game, denoted by $\Gamma_{\beta,\zeta}$, is finite, has perfect recall, and is such that every type w_{ik} of every player ik occurs with positive probability. Henceforth, we restrict attention to values of $\zeta > 0$ that satisfy (6.17).

Let s^ζ be a Nash equilibrium of the agent normal form of $\Gamma_{\beta,\zeta}$, and let s^* be the limit of s^ζ along a convergent subsequence as $\zeta \to 0$. Then, by continuity, s^* is a Nash equilibrium of the agent normal form of the finite game Γ_β and so, exactly as in the proof of Theorem 5.2, s^* is an $(\varepsilon, \mathcal{F})$-sequential equilibrium of Γ. It remains only to show that s^* is ε-subgame perfect.

Consider any subgame $(\hat{\theta}_{\leq k}, \hat{a}_{<k})$ and any player $ik \in L$. There is a unique $\hat{w}_{ik} \in W_{ik}$ such that $(\hat{\theta}_{\leq k}, \hat{a}_{<k}) \in \hat{w}_{ik}$.[24] Since every type $w_{ik} \in W_{ik}$, including $\hat{w}_{ik} \in W_{ik}$, has positive probability under s^ζ, there is a unique probability measure μ^ζ on $\hat{w}_{ik}$ that represents ik's beliefs about the past history given $\hat{w}_{ik}$. Furthermore, because for any date h and for any type of player ih, s_{ih}^ζ places no more than probability $\beta + \zeta$ on suboptimal actions in A_{ih}^1, the second inequality in (6.17) implies that s_i^ζ is $\varepsilon/2$-optimal for player i against s_{-i}^ζ given $\hat{w}_{ik}$ in the game Γ_0.

For any ζ, μ^ζ is a convex combination of a fixed finite set of probability measures on $\hat{w}_{ik}$. Indeed, let $(\theta_{\leq k}^1, a_{<k}^1)$ be the (unique) element of $\Theta_{\leq k}^1 \times A_{<k}^1$ that is contained in $\hat{w}_{ik}$. Then $\hat{w}_{ik}$ can occur with positive probability under s^ζ only if the players choose action $a_{<k}^1$. Let $\mathcal{H}$ be the set of subsets H of $\{1, ..., k\}$ such that $\hat{w}_{ik}$ would have positive probability with $a_{<k}^1$ if H were the set of dates up to k where nature trembles to $\theta_{\leq k}^1$. We note that $\mathcal{H}$ is finite and nonempty, because it includes the set $\{1, ..., k\}$ itself. We can note also that the players' actions $a'_{\geq k}$ cannot affect probabilities of types at date k. Then, we can observe, for any fixed $a'_{\geq k}$, μ^ζ must be a convex combination of the probability measures:

$$\{P_T(\cdot|\hat{w}_{ik}, (a_{<k}^1, a'_{\geq k}); ((\theta_h^1)_{h\in H}, (p_h)_{h\in K\setminus H}))\}_{H\in\mathcal{H}},$$[25]

The weight for any H would be the conditional probability, in the ζ-pertubed game $\Gamma_{\beta,\zeta}$, of

---

[24] As observed in Section 5.1, in a projective game, if some history up to date k, $(\theta_{\leq k}, a_{<k})$, is a subgame, then all players at date k always observe all of nature's states from dates 1 to k and always observe all players' actions from dates 1 to $k-1$. So, in particular, $\hat{w}_{ik} \subseteq \Theta_{\leq k} \times A_{<k}$.

[25] Recall from Section 2.2 that $P_T(\cdot|(a_{<k}^1, a'_{\geq k}), ((\theta_h^1)_{h\in H}, (p_h)_{h\in K\setminus H}))$ provides the distributions over the players' types when the strategy profile is $(a_{<k}^1, a'_{\geq k})$ and nature's probability function is

nature having trembled at exactly the dates in H, when it is given that i's date-k type is $\hat{w}_{ik}$.

Taking a further subsequence of $\{\zeta\}$ along which the weights in the finite convex combination μ^{ζ} all converge as $\zeta \to 0$, we may let μ^* be the probability measure on $\hat{w}_{ik}$ that is the convex combination obtained using the limit weights. Then, since s_i^{ζ} is $\varepsilon/2$-optimal for player i against s_{-i}^{ζ} given $\hat{w}_{ik}$ in the game Γ_0, by the continuity of finite convex combinations, s_i^* is $\varepsilon/2$-optimal for player i against s_{-i}^* in Γ_0 given beliefs μ^* on $\hat{w}_{ik}$. That is, for all date-k continuations r_i^0 of s_i^* that are feasible in Γ_0,

$$\int_{\hat{w}_{ik}} U_i(r_i^0, s_{-i}^* | \tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \mu^*(d(\tilde{\theta}_{\leq k}, \tilde{a}_{<k})) \leq \int_{\hat{w}_{ik}} U_i(s^* | \tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \mu^*(d(\tilde{\theta}_{\leq k}, \tilde{a}_{<k})) + \varepsilon/2 \quad (6.18)$$

Let $r_i \in S_i$ be any strategy for player i in the original infinite game that is a date-k continuation of s_i^*. Analogous to the proof of Theorem 5.2 we wish to define a date-k continuation $r_i' \in S_i$ of s_i^* that is feasible in Γ_0 and that implies the following.

$$\begin{aligned}&\text{The distribution over the elements of the finite partition } Q^1\\&\text{is the same under each of the two probability measures}\\&P(\cdot|(\hat{a}_{<k}, r_{i,\geq k}, s^*_{-i,\geq k}); (\hat{\theta}_{\leq k}, \rho_{>k})) \text{ and}\\&P(\cdot|(\hat{a}_{<k}, r'_{i,\geq k}, s^*_{-i,\geq k}); (\hat{\theta}_{\leq k}, \rho_{>k})) \text{ on } \Theta \times A.\end{aligned} \quad (6.19)$$

Note that the two probability distributions in (6.19) each give probability 1 to the subgame $(\hat{\theta}_{\leq k}, \hat{a}_{<k})$ in Γ. Since the subgame starting at $(\hat{\theta}_{\leq k}, \hat{a}_{<k})$ is itself a multi-stage game, the existence of such a strategy for player i within this subgame has already been demonstrated by the general construction in the proof of Theorem 5.2 and so there is no need to repeat the argument. We can then extend this subgame strategy to the desired strategy r_i' for whole game Γ by defining r_i' to choose any available action with probability 1 for any types t_{ih} such that $h < k$ or such that the projection of t_{ih} onto $\Theta_{\leq k} \times A_{<k}$ is not in $\hat{w}_{ik}$. For any $h \geq k$ and for any t_{ih} whose projection onto $\Theta_{\leq k} \times A_{<k}$ is in $\hat{w}_{ik}$ we may define $r_i'(\cdot|t_{ih})$ to be equal to the already defined subgame strategy $r_i'(\cdot|\hat{t}_{ih})$, where $\hat{t}_{ih}$ is ih's type in the subgame $(\hat{\theta}_{\leq k}, \hat{a}_{<k})$ that coincides with t_{ih} except perhaps in the coordinate values $(\hat{\theta}_{\leq k}, \hat{a}_{<k})$.

Because every $(\tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \in \hat{w}_{ik}$ is in the same element of the partition $(\times_{h \leq k, j \in J} Q^1_{\Theta_{hj}}) \times (\times_{i \in N, h<k, j \in J} Q^1_{A_{ihj}})$ of $\Theta_{\leq k} \times A_{<k}$, both $r'_{i,\geq k}$ and $s^*_{-i,\geq k}$ are constant on $\hat{w}_{ik}$ (a common

---

$((\theta_h^1)_{h\in H}, (p_h)_{h \in K\setminus H}))$. So $P_T(\cdot|\hat{w}_{ik}, (a^1_{<k}, a'_{\geq k}), ((\theta^1_h)_{h\in H}, (p_h)_{h\in K\setminus H}))$ is the conditional distribution on $\hat{w}_{ik}$. Here, $(a^1_{<k}, a'_{\geq k})$ denotes the pure strategy profile that chooses the action $(a^1_{<k}, a'_{\geq k})$ with probability 1, and $(\theta^1_h)_{h\in H}$ denotes the degenerate probability function for nature for dates in H that chooses the state $(\theta^1_h)_{h\in H}$ with probability 1.

type for all players in the finite game Γ_0, by construction, since k is a subgame date). Because $\rho_{>k}$ is history independent, this means that the distribution over the elements of Q^1 under $P(\cdot|(\tilde{a}_{<k}, r'_{i,\geq k}, s^*_{-i,\geq k}); (\tilde{\theta}_{\leq k}, \rho_{>k}))$ is the same for all $(\tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \in \hat{w}_{ik}$. Therefore, since $(\hat{\theta}_{\leq k}, \hat{a}_{<k}) \in \hat{w}_{ik}$, (6.19) implies that $\forall (\tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \in \hat{w}_{ik}$,

$$\begin{aligned}&\text{The distribution over the elements of the finite partition } Q^1\\ &\text{is the same under each of the two probability measures}\\ &P(\cdot|(\hat{a}_{<k}, r_{i,\geq k}, s^*_{-i,\geq k}); (\hat{\theta}_{\leq k}, \rho_{>k})) \text{ and}\\ &P(\cdot|(\tilde{a}_{<k}, r'_{i,\geq k}, s^*_{-i,\geq k}); (\tilde{\theta}_{\leq k}, \rho_{>k})), \text{ on } \Theta \times A. \end{aligned} \tag{6.20}$$

Then,

$$\begin{aligned}
&U_i(r_i, s^*_{-i}|\hat{\theta}_{\leq k}, \hat{a}_{<k})\\
&= \textstyle\int_{\Theta\times A} u_i(\theta, a) P(d(\theta, a)|(\hat{a}_{<k}, r_{i,\geq k}, s^*_{-i,\geq k}); (\hat{\theta}_{\leq k}, p_{>k})),\\
&= \textstyle\int_{\Theta\times A} u_i(\theta, a) f^k(\theta, a) P(d(\theta, a)|(\hat{a}_{<k}, r_{i,\geq k}, s^*_{-i,\geq k}); (\hat{\theta}_{\leq k}, \rho_{>k}))\\
&\leq \textstyle\int_{\hat{w}_{ik}} \left(\int_{\Theta\times A} u_i(\theta, a) f^k(\theta, a) P(d(\theta, a)|(\tilde{a}_{<k}, r'_{i,\geq k}, s^*_{-i,\geq k}); (\tilde{\theta}_{\leq k}, \rho_{>k}))\right) \mu^*(d(\tilde{\theta}_{\leq k}, \tilde{a}_{<k})) + \eta,\\
&= \textstyle\int_{\hat{w}_{ik}} \left(\int_{\Theta\times A} u_i(\theta, a) P(d(\theta, a)|(\tilde{a}_{<k}, r'_{i,\geq k}, s^*_{-i,\geq k}); (\tilde{\theta}_{\leq k}, p_{>k}))\right) \mu^*(d(\tilde{\theta}_{\leq k}, \tilde{a}_{<k})) + \eta,\\
&= \textstyle\int_{\hat{w}_{ik}} U_i(r'_i, s^*_{-i}|\tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \mu^*(d(\tilde{\theta}_{\leq k}, \tilde{a}_{<k})) + \eta,\\
&\leq \textstyle\int_{\hat{w}_{ik}} U_i(s^*|\tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \mu^*(d(\tilde{\theta}_{\leq k}, \tilde{a}_{<k})) + \varepsilon/2 + \eta,\\
&\leq U_i(s^*|\hat{\theta}_{\leq k}, \hat{a}_{<k}) + \varepsilon/2 + 2\eta,\\
&\leq U_i(s^*|\hat{\theta}_{\leq k}, \hat{a}_{<k}) + \varepsilon,
\end{aligned}$$

where the first inequality follows from (6.20) and (6.16); the second inequality follows from (6.18); the third inequality follows from (6.16) because every $(\tilde{\theta}_{\leq k}, \tilde{a}_{<k}) \in \hat{w}_{ik}$, including $(\hat{\theta}_{\leq k}, \hat{a}_{<k})$, is in the same element of the partition $(\times_{h\leq k, j\in J} Q^1_{\Theta_{hj}}) \times (\times_{i\in N, h<k, j\in J} Q^1_{A_{ihj}})$ of $\Theta_{\leq k} \times A_{<k}$, and because $s^*_{\geq k}$ is therefore constant on $\hat{w}_{ik}$ being is measurable with respect to Q^1; and the final inequality follows because $\eta < \varepsilon/4$.

We conclude that s^* is ε-subgame perfect, which completes the first step.

Step 2. By the first step, there is a net $\{s^{\varepsilon,\mathcal{F}}\}$ such that for every $\varepsilon > 0$ and for every finite subset $\mathcal{F}$ of $\mathcal{B}^*$, $s^{\varepsilon,\mathcal{F}}$ is both ε-subgame perfect and an $(\varepsilon, \mathcal{F})$-sequential equilibrium. Then, as in the proof of Theorem 3.3, there exists $\mu : \mathcal{Y} \times \mathcal{B}^* \to [0,1]$ and a subnet $\{s^{\varepsilon_\alpha,\mathcal{F}_\alpha}\}$ such that for every $\varepsilon > 0$, for every finite subset $\mathcal{F}$ of $\mathcal{B}^*$ and for every finite subset $\mathcal{G}$ of $\mathcal{Y}$, there exists α such that $\varepsilon_\alpha < \varepsilon$, $\mathcal{F}_\alpha \supseteq \mathcal{F}$, and

$$\left|P(Y|C, s^{\varepsilon_\alpha,\mathcal{F}_\alpha}) - \mu(Y|C)\right| < \varepsilon, \text{ for every } (Y,C) \in \mathcal{G} \times \mathcal{F}.$$

Since $\varepsilon_\alpha < \varepsilon$ and $\mathcal{F}_\alpha \supseteq \mathcal{F}$ imply that $s^{\varepsilon_\alpha,\mathcal{F}_\alpha}$ is, a fortiori, ε-subgame perfect and an $(\varepsilon, \mathcal{F})$-sequential equilibrium, we may conclude that μ is a subgame perfect open sequential equilibrium conditioned on $\mathcal{B}^*$. Q.E.D.

References

Billingsley, P., (1968): *Convergence of Probability Measures*, Wiley Series in Probability and Mathematical Statistics, New York.

Cohn, D. L., (1980): *Measure Theory*, Birkhauser, Boston.

Fudenberg, Drew and Jean Tirole (1991): "Perfect Bayesian and Sequential Equilibrium," *Journal of Economic Theory,* 53, 236-260.

Harris, Christopher J. , Philip J. Reny, and Arthur J. Robson (1995): "The existence of subgame-perfect equilibrium in continuous games with almost perfect information," *Econometrica* 63, 507-544.

Harris, Christopher J., Maxwell B. Stinchcombe, and William R. Zame (2000): "The Finitistic Theory of Infinite Games," University of Texas working paper. http://www.laits.utexas.edu/~maxwell/finsee4.pdf

Harris, Christopher J., Philip J. Reny, and Arthur J. Robson (1995): "The existence of subgame-perfect equilibrium in continuous games with almost perfect information," *Econometrica* 63(3):507-544.

Hellman, Ziv (2014): "A game with no Bayesian approximate equilibria," *Journal of Economic Theory*, 153, 138-151.

Hellman, Ziv and Yehuda Levy (2013): "Bayesian games with a continuum of states," Hebrew University working paper.

Kelley, John L. (1955), *General Topology,* Springer-Verlag.

Kreps, David and Robert Wilson (1982): "Sequential Equilibria," *Econometrica,* 50, 863-894.

Milgrom, Paul and Robert Weber (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10, 619-32.

Simon, Robert Samuel (2003): “Games of incomplete information, ergodic theory, and the measurability of equilibria,” Israel Journal of Mathematics 138:73-92.