Perfect Conditional ε-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions*

Roger B. Myerson* and Philip J. Reny**

*Department of Economics and Harris School of Public Policy
**Department of Economics
University of Chicago

Abstract

Abstract: We extend Kreps and Wilson's concept of sequential equilibrium to games with infinite sets of signals and actions. A strategy profile is a conditional ε-equilibrium if, for any of a player's positive probability signal events, his conditional expected utility is within ε of the best that he can achieve by deviating. With topologies on action sets, a conditional ε-equilibrium is *full* if strategies give every open set of actions positive probability. Such full conditional ε-equilibria need not be subgame perfect, so we consider a non-topological approach. *Perfect* conditional ε-equilibria are defined by testing conditional ε-rationality along nets of small perturbations of the players' strategies and of nature's probability function that, for any action and for almost any state, make this action and state eventually (in the net) always have positive probability. Every perfect conditional ε-equilibrium is a subgame perfect ε-equilibrium, and, in finite games, limits of perfect conditional ε-equilibria as $\varepsilon \to 0$ are sequential equilibrium strategy profiles. But limit strategies need not exist in infinite games so we consider instead the limit distributions over outcomes. We call such outcome distributions perfect conditional equilibrium distributions and establish their existence for a large class of regular projective games. Nature's perturbations can produce equilibria that seem unintuitive and so we augment the game with a net of permissible perturbations.

1 Introduction

We define perfect conditional ε-equilibrium and perfect conditional equilibrium distributions for multi-stage games with infinite signal sets and infinite action sets and prove their existence for a large class of games.

*Current Version: December 2, 2019. We thank the editor, Joel Sobel, and 5 referees for their detailed comments and suggestions. We also thank John Hillas, Arthur Robson, and Asher Wolinsky for helpful comments, and Scott Behmer, John van den Berghe, Wenji Xu, and Kai Hao Yang for their careful proofreading. Reny thanks the National Science Foundation (SES-1724747, SES-1227506, SES-0922535) for financial support.

Kreps and Wilson (1982), henceforth KW, defined sequential equilibrium for any finite game in which nature's states all have positive probability, henceforth *standard* finite games. But rigorously defined extensions to infinite games have been lacking. Various formulations of "perfect Bayesian equilibrium" (defined for standard finite games in Fudenberg and Tirole 1991) have been used for infinite games, but no general existence theorem for infinite games is available.[1]

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 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., Example 2.1).

Instead, we work directly with the infinite game itself. We define a strategy profile to be a conditional ε-equilibrium if, given the strategies of the other players, each player's continuation strategy is ε-optimal conditional on any positive probability set of signals.[2]

In standard finite games, it is not hard to see (although we have not seen it previously pointed out) that a strategy profile is part of a sequential equilibrium if and only if for every $\varepsilon > 0$ there is an arbitrarily close completely mixed strategy profile that is a conditional ε-equilibrium. It is this finite-game characterization of sequential equilibrium strategy profiles, without any reference to systems of beliefs, that we will extend to infinite games.

The central challenge in infinite games is how to test whether the players' behavior is rational off the equilibrium path of play. As we have just noted, in standard finite games the sequential equilibrium concept tests the players' rationality by checking whether, for any $\varepsilon > 0$, there is an arbitrarily close completely mixed strategy profile that is a conditional ε-equilibrium.

In infinite games, there are two serious difficulties with this approach to testing behavior off the equilibrium path. The first difficulty is that, with uncountably-infinite action spaces, we cannot make all actions have positive probability at the same time, no matter how we perturb the players' strategies. A possible response to this difficulty is to introduce separable topologies on action spaces and to test the players' behavior conditional on each signal event

[1] Watson (2017) proposed a new definition of perfect Bayesian equilibrium for standard finite games, and discussed how one might try to extend that definition to infinite games.

[2] See Radner (1980) for a study of ε-rationality in finitely repeated games.

that has positive probability under strategy profiles that have full support. A case for this approach is given in Section 5. But despite testing for rationality on a rich class of events, such a topological approach does not test for rationality everywhere, and so it can allow equilibria that fail to be subgame perfect (Example 5.3). So we emphasize a different approach.

The second difficulty is that, with uncountably many states of nature, nature's probability function must give all but countably many states probability zero. So even if rationality could be tested by perturbing the players' strategies, the players' resulting conditional probabilities over histories would be biased so as to explain, whenever possible, any probability zero event as being the result of a deviation by some player instead of perhaps being the result of the occurrence of a state of nature that had prior probability zero. This bias can be so severe that it rules out all but strictly dominated strategies (Example 6.1).

Our solution to these difficulties makes use of generalized sequences, i.e., nets. If b is any strategy profile and p is nature's probability function, we define a net $\{(b^\alpha, p^\alpha)\}$ of strategy-profile/nature-perturbation pairs to be admissible for (b, p) if two conditions are satisfied. First, the history-dependent probabilities on action-events specified in the net $\{b^\alpha\}$ must converge uniformly to those of b, and the history-dependent probabilities on state-events specified in the net $\{p^\alpha\}$ must converge uniformly to those of p. Second, for each history of play, any feasible action for any player given that history must receive positive probability under b^α for all large enough α, and almost-any state of nature that can occur given that history of play must receive positive probability under p^α for all large enough α.[3]

Admissible nets play the role in infinite games of convergent sequences of completely mixed strategies in finite games. Indeed, the strategies in an admissible sequence (hence, net) of any finite game are eventually all completely mixed. Importantly, admissible nets avoid the two serious difficulties described above because, first, for every feasible action, admissible nets eventually always give that action positive probability, and second, for almost any state of nature, admissible nets eventually always give that state positive probability thereby allowing zero probability events to be explained as the occurrence of a state of nature that has prior probability zero.[4]

A strategy profile b is defined to be a perfect conditional ε-equilibrium if there is a net of strategy-profile/nature-perturbation pairs that is admissible for (b, p) such that for each pair (b^α, p^α) in the net, b^α is a conditional ε-equilibrium in the game with nature's perturbed

[3] "Perturbations" are not synonymous with "mistakes." See KW, pp. 373-374.

[4] In standard finite games, all states of nature have strictly positive prior probability, which is why Kreps and Wilson (1982) did not need to perturb nature. (But note that their theory would have been unchanged even had they perturbed nature, because nature's strictly positive probabilities would swamp the infinitesimal perturbations.)

probability function p^α. A perfect conditional equilibrium distribution is defined as the limit of perfect conditional ε-equilibrium distributions on outcomes of the game as $\varepsilon \to 0$.[5]

Nets of perfect conditional ε-equilibria as $\varepsilon \to 0$ need not have convergent subnets even in very nice games, which is why we only consider their limit distributions on outcomes. As noted by Milgrom and Weber (1985), Van Damme (1987), Börgers (1991), and Harris et. al. (1995), the difficulty is that the randomized signals upon which players coordinate their actions along the sequence can, in the limit, have distributions that degenerate to a point, leaving the players without access to the necessary coordination device.

Our solution concept, perfect conditional ε-equilibrium, does not include systems of beliefs. In Section 6.4, we show that any perfect conditional ε-equilibrium generates a finitely consistent conditional belief system with respect to which it is sequentially ε-rational. These concepts extend to infinite games the concepts of consistency of beliefs and sequential rationality introduced in KW. Some difficulties with finite consistency are also discussed.

Perfect conditional ε-equilibria and perfect conditional equilibrium distributions are shown to exist for a large class of regular projective games (Theorems 9.3 and 9.5), and are shown to have other attractive properties. First, every perfect conditional ε-equilibrium strategy profile is a subgame perfect ε-equilibrium, and therefore also an ε-Nash equilibrium (Theorems 6.9 and 6.10). Second, if two players have the same information, they must behave, in any perfect conditional ε-equilibrium, as if they have the same beliefs about the history of play (Section 6.4). Third, in any standard finite game, a strategy profile is part of a sequential equilibrium if and only if it is the limit of perfect conditional ε-equilibria as $\varepsilon \to 0$ (Theorem 6.4). So in standard finite games, the perfect conditional equilibrium distributions defined here are precisely the distributions over outcomes that arise from sequential equilibria.

The remainder of the paper is organized as follows. Section 2 provides an example that motivates why we do not use finite-game approximations of the infinite game to define our solution. Section 3 introduces the multi-stage games that we study and provides some preliminary notation and concepts. Section 4 introduces our most basic equilibrium concept, conditional ε-equilibrium. Section 5 considers a topological approach to the problem of perfection in infinite games. Section 6 contains the definition of a perfect conditional ε-equilibrium strategy profile as well as the definition of a perfect conditional equilibrium dis-

[5] We use the term "perfect" to indicate that behavior is tested for rationality *everywhere* (i.e., at every event both on and off the equilibrium path).

Simon and Stinchcombe (1995) and Bajoori, Flesch, and Vermuelen (2013, 2016) use a topological full-support condition in defining, for infinite normal form games, solutions that they call "perfect." Such topological restrictions on the supports of strategies are used in Section 5 to refine perfect conditional ε-equilibrium, but we call the refined solution "full." The word "perfect" in English comes from a Latin word meaning "complete," and so it seems more appropriate for the condition of testing rationality everywhere versus testing rationality conditional only on sets that have positive probability under a full-support strategy profile, which, as already mentioned, need not even yield behavior that is subgame perfect.

tribution. This section also establishes several properties of perfect conditional ε-equilibria (e.g., that they are subgame perfect ε-equilibria), and introduces systems of beliefs and the concepts of finite consistency and sequential ε-rationality. Section 7 applies our definitions to several examples. Section 8 augments the game with a permissible net of nature perturbations to avoid unintuitive equilibria that can arise with arbitrary nature-perturbations. Section 9 introduces the class of "regular projective games" for which we can prove existence of perfect conditional ε-equilibria and perfect conditional equilibrium distributions. Section 10 provides some final remarks. The proof of our main existence result is in Section 11. All other proofs are in Myerson and Reny (2019).

2 Problems with Finite Approximations of Infinite Games

In this section, we provide an example that illustrates why we do not use finite approximating games as a basis for defining sequential equilibrium in infinite games. Despite many attempts, we have not found any method for providing "good" finite approximations of arbitrary multi-stage games. Instead, our solutions are based on strategies that are approximately conditionally optimal among *all* of the infinitely many strategies in the original game. To show just one of the ways that things can go wrong, the finite approximations used in this next example seem natural but lead to unacceptable results.

Example 2.1 *Spurious signaling in naïve finite approximations.*

- On date 1, nature chooses $\theta \in \{1, 2\}$ with $p(\theta = 1) = 1/4$, and player 1 chooses $x \in [0, 1]$.
- On date 2, player 2 observes the signal $s = x^\theta$ and chooses $y \in \{1, 2\}$.
- Payoffs (u_1, u_2) are as follows:

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

Consider subgame perfect equilibria of any finite approximate version of the game where player 1 chooses x in some finite subset of $[0, 1]$ that includes at least one interior point. We shall argue that player 1's expected payoff must be $1/4$.

Player 1 can obtain an expected payoff of at least $1/4$ by choosing the largest feasible $\bar{x} < 1$, as 2 should choose $y = 1$ when $s = \bar{x} > \bar{x}^2$ indicates $\theta = 1$. (In this finite approximation, player 2 has perfect information after the history $\theta = 1, x = \bar{x}$.)

Player 1's expected payoff cannot be more than $1/4$, as 1's choice of the smallest $0 < \underline{x} < 1$ in his equilibrium support would lead player 2 to choose $y = 2$ when $s = \underline{x}^2 < \underline{x}$ indicates $\theta = 2$.

But such a scenario cannot be even an approximate equilibrium of the original infinite game, because player 1 could get an expected payoff at least 3/4 by deviating to $\sqrt{\bar{x}}$ ($> \bar{x}$).

In fact, by reasoning analogous to that in the preceding two sentences, player 1 must receive an expected payoff of 0 in any Nash equilibrium of the infinite game, and so also in any sensibly defined "sequential equilibrium."

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

We next formally introduce the class of games that we study.

3 Multi-Stage Games

For ease of exposition, we restrict our analysis to a large class of extensive-form games called *multi-stage games*. A multi-stage game is played in a finite sequence of dates.[6] At each date t, each player receives a private "signal," about the history of play. Each player then simultaneously chooses an action from his set of available date-t actions, and nature simultaneously chooses a date-t state whose distribution can depend on the entire history of play. Perfect recall is assumed.[7]

Formally, a multi-stage game $\Gamma = (I, T, S, A, \mathcal{M}, \Phi, p, \sigma, u)$ consists of the following items.

Γ.1. I is the finite set of players, $0 \notin I$. Let $I^* = I \cup \{0\}$, where 0 denotes nature (chance). The finite set of dates of play is $\{1, ..., T\}$. Let $L = I \times \{1, ..., T\}$ denote the set of dated players, let $L^* = I^* \times \{1, ..., T\}$, and write it for (i, t).

Γ.2. $S = \times_{it \in L} S_{it}$, where S_{it} is the set of possible signals received by player i at date t; $S_{i1} = \{\emptyset\}$ for all $i \in I$.

Γ.3. For $it \in L$, A_{it} is the set of all possible *date-t actions* for player i, and A_{0t} is the set of all possible *date-t states* of nature.

Γ.4. $A \subseteq \times_{it \in L^*} A_{it}$ is the set of possible *outcomes* of the game. (Additional restrictions on A are given below.)

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

[7] Multi-stage games include Bayesian games, signaling games, principal-agent games, games with perfect information, games with almost perfect information, finitely-repeated games with and without private monitoring, and finite-horizon stochastic games. But if we define precedence by saying that one signal (information set) in the game precedes another when there is a path of play along which the one signal is generated first and the other signal is generated second, then the class of multi-stage games excludes all games in which the transitive closure of this binary relation fails to be acyclic. We restrict attention to multi-stage games only because they are notationally simple to describe. But there is no real difficulty in extending our definitions to games with perfect recall outside this class. See Myerson and Reny (2019).

The subscript, $< t$, indicates the projection onto dates before t, and $\leq t$ weakly before. For example, for any $a \in \times_{it\in L^*} A_{it}$, $a_{<t} = (a_{ir})_{i\in I^*, r<t}$. So if $a \in A$, then $a_{<t}$ is the associated *date-t history*, i.e., the partial sequence of actions and states before date t, and $A_{<t} = \cup_{a\in A}\{a_{<t}\} = \{\text{date-}t \text{ histories}\}$ ($A_{<1} = \{\emptyset\}$, $A_{<T+1} = A$). We also use $a_{<t}$ to denote a typical element of $A_{<t}$ without necessarily specifying an outcome in A whose projection onto dates before t is equal to $a_{<t}$.

Γ.5 The mapping $\mathcal{M}(\cdot)$ specifies sigma-algebras (closed under complements and countable intersections) of measurable subsets for each S_{it}, A_{it}, and A_{0t}, as well as for any of their finite products. So for example, $\mathcal{M}(S_{it})$ is the set of measurable subsets of S_{it}. All one-point sets are measurable, product sets are given their product sigma-algebras, and subsets of measurable spaces are given their relative sigma algebras. Assume that $A \in \mathcal{M}(\times_{it\in L^*} A_{it})$ and that $A_{<t} \in \mathcal{M}(\times_{i\in I^*, r<t} A_{ir})$ for each date $t \leq T$.

Γ.6. Player i's date t information is determined by a measurable and onto signal function $\sigma_{it} : A_{<t} \to S_{it}$.[8] Since, for every $i \in I$, $S_{i1} = A_{<1} = \{\emptyset\}$, we define $\sigma_{i1}(\emptyset) = \emptyset$. Assume perfect recall: $\forall it \in L$, $\forall r > t$, there are measurable functions $\bar{\Psi}_{irt} : S_{ir} \to S_{it}$ and $\bar{\psi}_{irt} : S_{ir} \to A_{it}$ such that $\bar{\Psi}_{irt}(\sigma_{ir}(a_{<r})) = \sigma_{it}(a_{<t})$ and $\bar{\psi}_{irt}(\sigma_{ir}(a_{<r})) = a_{it}$, $\forall a \in A$. The game's signal function is $\sigma = (\sigma_{it})_{it\in L}$.

Γ.7 For $it \in L$, $s_{it} \in S_{it}$, and $a_{<t} \in A_{<t}$, $\Phi_{it}(s_{it}) \in \mathcal{M}(A_{it})$ is the set of all *feasible* date-t actions for player i given the signal s_{it}, where $\Phi_{i1}(\emptyset) = A_{i1}$, (so on date $t = 1$ every action in A_{i1} is feasible for player i). Assume that for any $a_{it} \in A_{it}$, the set $\{s_{it} \in S_{it} : a_{it} \in \Phi_{it}(s_{it})\}$ is measurable. For any date $t \leq T$ and for any $a \in \times_{ir\in L^*} A_{ir}$, assume that $a_{<t+1} \in A_{<t+1}$ iff for every player $i \in I$ and for every date $r \leq t$, $a_{ir} \in \Phi_{ir}(\sigma_{ir}(a_{<r}))$. So the set A of outcomes of the game is the set of all paths along which the players' actions are feasible given any history.

Let $\Delta(X)$ denote the set of countably additive probability measures on the measurable subsets of X. For any two measurable spaces X and Y, a mapping $\zeta : Y \to \Delta(X)$ is a *transition probability* iff for every measurable $C \subseteq X$, $\zeta(C|y)$ is a measurable real-valued function of y on Y.

Γ.8. $p = (p_1, ..., p_T)$ is nature's probability function where, for each date t, $p_t : A_{<t} \to \Delta(A_{0t})$ is a transition probability.

Γ.9. Each player i has a bounded measurable utility function $u_i : A \to \mathbb{R}$, and $u = (u_i)_{i\in I}$.

[8] It is without loss of generality to assume, for every $r < t$, that σ_{it} does not depend on the date-r signal of any player since earlier signals depend on even earlier states and actions.

At each date $t \in \{1, ..., T\}$ starting with date $t = 1$, and after any date-t history $a_{<t} \in A_{<t}$, each player i is privately informed of his date-t signal, $s_{it} = \sigma_{it}(a_{<t})$, after which each player i simultaneously chooses an action from his set of feasible date-t actions $\Phi_{it}(s_{it}) \subseteq A_{it}$, and nature chooses a date-t state $a_{0t} \in A_{0t}$ according to $p_t(\cdot|a_{<t})$. The game then proceeds to the next date. After T dates of play this leads to an outcome $a \in A$ and the game ends with player payoffs $u_i(a)$, $i \in I$.

In the next two subsections, we formally introduce strategies, outcome distributions, payoffs, and conditional payoffs.

3.1 Strategies and Induced Outcome Distributions

A *strategy* for dated player $it \in L$ is any transition probability $b_{it} : S_{it} \to \Delta(A_{it})$ that satisfies $b_{it}(\Phi_{it}(s_{it})|s_{it}) = 1$ for every $s_{it} \in S_{it}$.

Let B_{it} denote it's set of strategies and let $B_i = \times_{t \in T} B_{it}$ denote i's (behavior) strategies. Perfect recall ensures that there is no loss in restricting attention to B_i for each player i. Let $B = \times_{it \in L} B_{it}$ denote the set of all strategy profiles.

For any date t, let $B_{\cdot t} = \times_{i \in I} B_{it}$ denote the set of date-t strategy vectors with typical element $b_{\cdot t} = (b_{it})_{i \in I}$. Let $A_{\cdot t} = \times_{i \in I^*} A_{it}$. Each $b_{\cdot t} \in B_{\cdot t}$ determines a transition probability P_t from $A_{<t}$ to $\mathcal{M}(A_{\cdot t})$ such that, for any measurable product set $C = \times_{i \in I^*} C_{it} \subseteq \times_{i \in I^*} A_{it}$ and for any $a_{<t} \in A_{<t}$,

$$P_t(C|a_{<t}, b_{\cdot t}) = p_t(C_{0t}|a_{<t}) \Pi_{i \in I} b_{it}(C_{it}|\sigma_{it}(a_{<t})). \tag{3.1}$$

For any $b \in B$, we inductively define measures $P_{<t}(\cdot|b)$ in $\Delta(A_{<t})$ so that $P_{<1}(\{\emptyset\}|b) = 1$ and, for all $t \in \{1, ..., T\}$ and for all measurable $C \subseteq A_{<t+1}$,

$$P_{<t+1}(C|b) = \int P_t(\{a_{\cdot t} : (a_{<t}, a_{\cdot t}) \in C\}|a_{<t}, b_{\cdot t}) P_{<t}(da_{<t}|b). \tag{3.2}$$

(Notice that $P_{<t}(\cdot|b)$ depends only on $b_{<t}$.)

Let $P(\cdot|b) = P_{<T+1}(\cdot|b)$ be the probability measure on outcome events in $\mathcal{M}(A)$ that is induced by b. The dependence of $P(\cdot|b)$ on nature's probability function p will sometimes be made explicit by writing $P(\cdot|b; p)$.

For any $b \in B$, we inductively define transition probabilities from $A_{<t}$ to $\Delta(A_{\geq t})$ so that $P_{\geq T}(\cdot|a_{<T}, b) = P_T(\cdot|a_{<T}, b_{\cdot T})$, and for any date $t < T$ and any measurable $C \subseteq A_{\geq t}$,

$$P_{\geq t}(C|a_{<t}, b) = \int P_{\geq t+1}(\{a_{\geq t+1} : (a_{\cdot t}, a_{\geq t+1}) \in C\}|a_{<t+1}, b) P_t(da_{\cdot t}|a_{<t}, b_{\cdot t}).$$

(Notice that $P_{\geq t}(\cdot|a_{<t}, b)$ does not depend on $b_{<t}$.)

At any date t, the conditional expected utility for player i with strategies b given history $a_{<t}$ is,

$$U_i(b|a_{<t}) = \int u_i(a_{<t}, a_{\geq t}) P_{\geq t}(da_{\geq t}|a_{<t}, b),$$

(notice that $U_i(b|a_{<t})$ does not depend on $b_{<t}$), and, player i's ex-ante expected utility is

$$U_i(b) = \int u_i(a) P(da|b) = \int U_i(b|a_{<T+1}) P_{<T+1}(da_{<T+1}|b).$$

3.2 Conditional Probabilities

For any $b \in B$, for any $it \in L$ and for any $Z \in \mathcal{M}(S_{it})$, define

$$P_{it}(Z|b) = P_{<t}(\sigma_{it}^{-1}(Z)|b) = P_{<t}(\{a_{<t} : \sigma_{it}(a_{<t}) \in Z\}|b).$$

Then $P_{it}(Z|b)$ is the probability that i's date t signal is in Z under the strategy profile b. The dependence of $P_{it}(\cdot|b)$ on nature's probability function p will sometimes be made explicit by writing $P_{it}(\cdot|b; p)$.

For any $it \in L$ and for any measurable $Z \subseteq S_{it}$, if $P_{it}(Z|b) > 0$, we may define: *conditional probabilities,*

$$P_{<t}(C|Z, b) = P_{<t}(C \cap \sigma_{it}^{-1}(Z)|b)/P_{it}(Z|b), \ \forall C \in \mathcal{M}(A_{<t}),$$

and,

$$P(C|Z, b) = P(\{a \in C : \sigma_{it}(a_{<t}) \in Z|b)/P_{it}(Z|b), \ \forall C \in \mathcal{M}(A),$$

and *conditional expected payoffs,*

$$U_i(b|Z) = \int_A u_i(a) P(da|Z, b).$$

(Notice that $P_{<t}(\cdot|Z, b)$ is the marginal of $P(\cdot|Z, b)$ on $A_{<t}$.) The dependence of $P(\cdot|Z, b)$ and $U_i(\cdot|Z)$ on nature's probability function p will sometimes be made explicit by writing $P(\cdot|Z, b; p)$ and $U_i(\cdot|Z; p)$.

The concepts that we will define in the rest of the paper are all based on the idea that players must choose strategies that are *approximately* optimal among all of their feasible strategies in the game Γ. One might think that strategies that are fully optimal can be obtained by taking limits of approximately optimal strategies, but this is not the case. The difficulty with *exact* optimality arises through a phenomenon that we call "strategic entanglement," where a sequence of strategy profiles yields 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, there may be no strategy profile that produces the limit outcome distribution.[9] In fact, Harris et. al. (1995) give an example in which this problem is so severe that it precludes the existence of a subgame perfect equilibrium in a two-stage game with compact action sets and continuous payoff functions.[10] Consequently, for much of what follows we consider strategies in which all players are ε-optimizing. But see Section 6.2 where we consider the limits, as $\varepsilon \to 0$, of the outcome distributions produced by such ε-optimal strategies.

We next introduce a basic solution concept that, like Nash equilibrium, only disciplines behavior in positive probability events.

4 Conditional ε-Equilibrium

For any $it \in L$, and for any $b_i \in B_i$, say that $c_i \in B_i$ is a *date-t continuation of* b_i if $c_{ir} = b_{ir}$ for all $r < t$.

Definition 4.1 *For any $\varepsilon \geq 0$, a strategy profile $b \in B$ is a conditional ε-equilibrium iff for every $it \in L$, for every measurable $Z \subseteq S_{it}$ satisfying $P_{it}(Z|b) > 0$, and for every date-t continuation c_i of b_i,*

$$U_i(c_i, b_{-i}|Z) \leq U_i(b|Z) + \varepsilon. \tag{4.1}$$

Every conditional ε-equilibrium is an ε-Nash equilibrium, which only requires inequality (4.1) to be satisfied when $Z = S_{it}$. But the converse can fail because, in an ε-Nash equilibrium a player may be able to improve his conditional payoff in some observable event by more than ε if the conditioning event occurs with sufficiently small probability in equilibrium.

Conditional ε-equilibrium ensures that no player could expect significant gains by unilaterally deviating from the equilibrium after any event that has positive probability in the equilibrium, and so predicted behavior will satisfy approximate rationality in all such positive-probability events. One might hope that we could ignore any possibility of irrational behavior in events that have zero probability in equilibrium, since they are unlikely to occur! But for any event that would be observable by some player j, if this event could get positive probability when some player i deviated from the equilibrium, then j's predicted behavior in this event may be used in the calculation of i's expected payoffs from this deviation. So

[9] Milgrom and Weber (1985) provided the first example of this kind. See also Van Damme (1987) and Börgers (1991).

[10] The nonexistence of a strategy supporting the limit outcome distribution can sometimes be remedied by adding an appropriate correlation device between periods as in Harris et. al. (1995). Manelli (1996) considers the problem of strategic entanglement in signaling games and restores existence there by adding cheap talk to the sender's message. Both of these remedies can add equilibria that are not ε-equilibria of the original game.

if j's predicted behavior in this event would not be even approximately rational, then the calculation of i's incentive to deviate in inequality (4.1) could be flawed. Thus we need to strengthen conditional ε-equilibrium so as to verify the rationality of players' behavior in any observable event that could get positive probability if players deviated from equilibrium, even if the event has zero probability in the equilibrium itself.

For a finite extensive-form game, the problem can be avoided by considering conditional ε-equilibria in which each player at each information set assigns at least some small positive probability to every feasible action. Such a completely mixed strategy would give positive probability to every event that could get positive probability after any strategic deviations by the players in the finite game. So for finite games, Kreps and Wilson could define sequential equilibria as limits (as $\varepsilon \to 0$) of completely mixed conditional ε-equilibria.

But in infinite games where players have uncountably infinite sets of actions, any behavioral strategy profile must leave many actions with zero probability. Then a deviation to such zero-probability actions could lead to events where our definition of conditional ε-equilibrium does not test the rationality of players' behavior. Thus, we will need to consider perturbations of the conditional ε-equilibrium, to test rationality in these events, unless we can find some reasonable way to restrict the set of strategic deviations that must be considered. This latter possibility is explored in the next section, where we restrict consideration to a dense set of deviations, using a topology on action spaces.

5 Full Conditional ε-Equilibrium

Although it may be impossible to give positive probability to all actions for a player in an infinite game, the player may have behavioral strategies that assign positive probability to every neighborhood of every action, under some suitable topology on A_{it}. So for a multistage game Γ as in Section 3, let us suppose now that, for each $it \in L$, the action set A_{it} is a separable metric space, and the measurable sets $\mathcal{M}(A_{it})$ are the Borel sets. For simplicity, in this section we assume that each player's set of feasible actions is history independent, so that $\Phi_{it}(s_{it}) = A_{it}$, for all $s_{it} \in S_{it}$ and for all $it \in L$.

We say that a strategy profile b has *full support* iff, for all $it \in L$, and for all $s_{it} \in S_{it}$, we have $b_{it}(C|s_{it}) > 0$ for every C that is a nonempty open subset of $\Phi_{it}(s_{it}) = A_{it}$.[11] Full-support strategies exist, by the assumption that the topology on each A_{it} is separable, as each dated player has a countable dense set of actions that could all be given positive

[11] Such strategies have been defined for various games with history-independent action sets. For signaling games, Mansuwe et. al. (1997) call such strategies pointwise completely mixed. For Bayesian games, Bajoori, Flesh, and Vermuelen (2016) call them completely mixed behavior strategies. For a class of extensive games, Jung (2018) calls them fully mixed. In normal form games, Simon and Stinchcombe (1995), and Bajoori, Flesch, and Vermuelen (2013) have used mixed strategies with full support to refine Nash equilibrium.

probability. Furthermore, any strategy profile b can be closely approximated by full support strategy profiles, because $(1-\lambda)b + \lambda\hat{b}$ has full support whenever $0 < \lambda < 1$ and $\hat{b}$ has full support.

Definition 5.1 *Say that b is a full conditional ε-equilibrium iff b is a conditional ε-equilibrium that has full support.*

With full-support strategies, any feasible action for any player has arbitrarily small neighborhoods that will get positive probability under the player's strategy after any possible signal. Using this property, we can now construct a dense set of deviations under which the problem of zero-probability events does not arise for a full conditional ε-equilibrium.

Let us define a *tremble profile* to be any $\varphi = (\varphi_{it})_{it\in L}$ such that each $\varphi_{it} : A_{it} \times S_{it} \to \Delta(A_{it})$ is a transition probability that satisfies $\varphi_{it}(\Phi_{it}(s_{it})|a_{it}, s_{it}) = 1$, for all $a_{it} \in A_{it}$ and for all $s_{it} \in S_{it}$. For any tremble profile φ and any strategy profile $b \in B$, let $b * \varphi$ denote the strategy profile $(b_{it} * \varphi_{it})_{it\in L} \in B$ where, for each $it \in L$, $b_{it} * \varphi_{it} \in B_{it}$ is defined by,

$$[b_{it} * \varphi_{it}](C|s_{it}) = \int \varphi_{it}(C|a_{it}, s_{it})b_{it}(da_{it}|s_{it}),\ \forall C \in \mathcal{M}(A_{it}),\ \forall s_{it} \in S_{it},\ \forall it \in L.$$

The tremble profile φ is *δ-local* iff $\varphi_{it}(\mathbb{B}_\delta(a_{it})|a_{it}, s_{it}) = 1$, for all $a_{it} \in A_{it}$, for all $s_{it} \in S_{it}$, and for all $it \in L$, where $\mathbb{B}_\delta(a_{it})$ is the ball of radius δ around a_{it}.

So a δ-local tremble profile φ describes a model in which, when any player i intends to choose some action a_{it} after observing some signal s_{it}, the player would tremble slightly and would really choose some nearby action, within distance δ from a_{it}, according to the probability distribution $\varphi_{it}(\cdot|a_{it}, s_{it})$. If the players' intended actions were generated by the strategy profile b, then their realized actions would depend on their signals according to the strategy profile $b * \varphi$. By taking δ to 0, we can guarantee that each player's realized actions with a δ-local tremble will always be arbitrarily close to his intended actions.

The following theorem tells us that, for any full conditional ε-equilibrium b, we can construct arbitrarily small local trembles that do not change b and are such that, for any intended deviations by any players, the corresponding deviations with trembles do not lead to any positive-probability events in which the rationality of the full conditional ε-equilibrium has not already been tested. A proof is in Myerson and Reny (2019).[12]

Theorem 5.2 *Suppose that b is a strategy profile with full support. Then for any $\delta > 0$, there is a δ-local tremble profile φ such that $b * \varphi = b$ and, for every $\hat{b} \in B$ and every $C \in \mathcal{M}(A)$, if $P(C|\hat{b} * \varphi) > 0$ then $P(C|b) > 0$.*

[12] The idea is to partition each action space A_{it} into measurable sets C, each of which is contained in a ball of radius $\delta/2$ and has a nonempty interior, so that it gets positive probability under b_{it} with full support. Then for any intended action a_{it} in the partition element C, we can let $\varphi_{it}(D|a_{it}, s_{it}) = b_{it}(D \cap C|s_{it})/b_{it}(C|s_{it})$ for all $D \in \mathcal{M}(A_{it})$. So $\varphi_{it}(\cdot|a_{it}, s_{it})$ imitates $b_{it}(\cdot|s_{it})$ within a small set of actions that are δ-close to a_{it}.

Full conditional ε-equilibria exist in a large class of (regular projective) games (Theorem 9.3 and Remark 9.4), but can fail to be subgame perfect, as the next example shows.

Example 5.3 *Failure of subgame perfection for a full conditional ε-equilibrium.*
• On date 1, nature chooses θ uniformly from $[0, 1]$.
• On date 2, player 1 observes the signal $s_1 = \theta$ and chooses $x \in [0, 1]$.
• On date 3, player 2 observes the signal $s_2 = (\theta, x)$ and chooses $y \in [0, 1]$
• Payoffs are $u_1 = u_2 = 1$ if $\theta = x = y$, and $u_1 = u_2 = 0$ otherwise.[13]
• Full-support strategies are defined with the usual topology on $[0, 1]$ as a subset of the real number line.

In this game, each player wants both players to match nature's choice of θ. Since both players observe the past history when it is their turn to move, this game has perfect information. In any subgame-perfect equilibrium, player 2 must be expected to choose $y = \theta$ whenever $x = \theta$, and so player 1 should choose $x = \theta$, so that both players get a payoff of 1. With any $\varepsilon > 0$, there exist full conditional ε-equilibria in which this outcome event $\{x = y = \theta\}$ has arbitrarily high probability (with each player having a small probability of choosing an action from a full-support distribution on $[0, 1]$).

However, we can also find full conditional ε-equilibria in which the players' expected payoffs are 0. For example, consider strategies where each player's action would be chosen from a uniform distribution on $[0, 1]$ independently of the observed history. Player 2 would strictly prefer to choose $y = \theta$ in the event $\{x = \theta\}$, but conditional ε-equilibrium does not require rationality of 2's response in this event because it has probability 0 when player 1 chooses x independently of θ. ♦

If in the above zero-payoff imperfect equilibrium player 1 understood that player 2 would rationally respond to $x = \theta$ by choosing $y = \theta$, then player 1 would certainly prefer to choose $x = \theta$. But this argument depends on the implicit assumption that player 1 can choose x exactly equal to θ, without any small local tremble. If player 1's intended choice of $x = \theta$ would lead to the realized x actually being drawn from a uniform distribution over the interval of $[\theta - \delta, \theta + \delta]$, for some small $\delta > 0$, then player 1 could not force the exact match $\{x = \theta\}$ to have positive probability even if he tried, and so the failure of subgame perfection in this event would not actually matter. It is in this sense that Theorem 5.2 tells us that any failures of sequential rationality in a full conditional ε-equilibrium could become irrelevant if players' choices are subject to arbitrarily small local trembles.

[13] These payoff functions are discontinuous. A similar example with continuous payoffs can be obtained by adding more stages. Discontinuities for early players then arise because the behavior of later players is not continuous in the actions of the early players.

This interpretation of the zero-payoff imperfect full conditional ε-equilibrium relies on the possibility that players might be unable to even approximately optimize since local trembles must preclude at least one of the players from matching nature's choice of θ even when the other player matches θ. Next, we develop an approach to the problem of perfection in which all players are assumed always to approximately optimize over their entire set of feasible strategies.

6 Perfect Conditional ε-Equilibrium

From Example 5.3, we see that subgame perfection cannot be guaranteed without testing rationality of players' responses to all possible deviations (not just some dense set of deviations). Thus, we now develop our concept of perfect conditional ε-equilibrium by considering nets of perturbations of the players' strategies and nets of perturbed probability functions for nature that eventually (in the net) give all player actions and almost all states of nature positive probability, and along which the players ε-optimize conditional on all positive probability events.

Our next example motivates why we must perturb both nature's probability function and the players' strategies when we require that rationality be tested with positive probability at each signal. The example shows that such a requirement can be incompatible with the existence of equilibrium if we perturb only the players' strategies.

Example 6.1 *Nonexistence of equilibrium when only strategies are perturbed in rationality tests.*

- On date 1, Nature chooses θ uniformly from $[0, 1]$ and player 1 chooses $x \in \{-1\} \cup [0, 1]$.
- On date 2, player 2 observes the signal s, where $s = x$ if $x \in [0, 1]$ and $s = \theta$ if $x = -1$, and then chooses $y \in \{-1, 1\}$.
- Payoffs are $u_1 = -x$, $u_2 = (x + 1/2)y$.

In this game, player 2 observes a number $s \in [0, 1]$ but she does not know whether the number she observes was chosen by player 1 (which will be the case when $x \in [0, 1]$) or was chosen by nature (which will be the case when $x = -1$).

The strategy $x = -1$ is strictly dominant for player 1 and player 2 wants to choose $y = -1$ if and only if $x = -1$. So this game has an essentially unique Nash equilibrium in which player 1 chooses $x = -1$ and player 2 chooses $y = -1$ for Lebesgue almost every signal $s \in [0, 1]$ that she observes.

However, if we required that, for any signal $s \in [0, 1]$, player 2's equilibrium behavior should pass a conditional rationality test in slightly perturbed strategies that give this signal positive probability (so that conditional payoffs can be computed), then there would be no equilibrium at all. Indeed, for any $\alpha \in [0,1]$, the event $\{s = \alpha\}$ can have positive probability, but only if positive probability is given to $x = \alpha$, because the event $\{\theta = \alpha\}$ has probability 0. So in any scenario where $\{s = \alpha\}$ has positive probability, conditional rationality would require player 2 to choose $y = 1$ when she observes $s = \alpha$ since the resulting conditional probability of the event $\{x \in [0,1]\}$ is one. Applying this same argument to every signal $\alpha \in [0,1]$ would imply that player 2 must choose $y = 1$ after every signal. But, for $\varepsilon > 0$ small enough, this strategy is not even an ε-best reply for player 2 against player 1's strictly dominant choice of $x = -1$. ♦

To see the problem another way, consider any possible value of 2's signal $\hat{s} \in [0,1]$. We could try to estimate what player 2 should believe is the conditional probability of player 1 having chosen $x = -1$ given that 2 has observed $s = \hat{s}$ by taking the limit of what this Bayesian belief probability would be for strategies in a net of strategies which converge to 1's unique equilibrium strategy and which (eventually) give positive probability to the event of 2 observing $\{s = \hat{s}\}$ (so that we can apply Bayes rule). But these Bayesian belief probabilities must all be 0, because the event of 2 observing $s = \hat{s}$ can have positive probability only when player 1 gives some small positive probability to the event $\{x = \hat{s} \geq 0\}$, since the event $\{\theta = \hat{s}\}$ must have probability 0 as long as we do not perturb nature's behavior. Now this argument can be applied for every $\hat{s}$ in $[0,1]$. Thus, when we try to compute conditional belief probabilities from a net of perturbations of 1's equilibrium strategy, we find that player 2 must assign belief probability 0 to the event $\{x = -1\}$ conditional on every individual signal in $[0,1]$. But before observing this signal, knowing only that $s \in [0,1]$, player 2 must understand that the event $\{x = -1\}$ has probability 1 in equilibrium.

This problem arises here because, when only the players' strategies are perturbed, the positive probability rationality test biases player 2's conditional beliefs toward explaining prior probability-zero events as always being the result of a deviation by player 1 instead of perhaps being the result of the occurrence of a probability-zero state of nature.

To avoid such biased beliefs, and to steer clear of the problem encountered here, we perturb both the players' strategies and nature's probability function in our tests for rational behavior.

We next introduce our main solution concept which, unlike both conditional ε-equilibrium and full conditional ε-equilibrium, tests for rational behavior even at events that have probability zero in equilibrium.

6.1 Perfect Conditional ε-Equilibrium

Let $\mathcal{T}$ denote the set of $\tau = (\tau_1, ..., \tau_T)$ such that each $\tau_t : A_{<t} \to \Delta(A_{0t})$ is a transition probability. Thus $\mathcal{T}$ is the set of alternative probability functions for nature in the game Γ. Notice that nature's probability function p is in $\mathcal{T}$. For any $\tau \in \mathcal{T}$, let $\Gamma(\tau)$ denote the perturbed game in which nature's probability function is τ instead of p.

For any $\tau, \tau' \in \mathcal{T}$, define $\|\tau' - \tau\| = \sup |\tau'_t(C|a_{<t}) - \tau_t(C|a_{<t})|$, where the supremum is over all $t \leq T$, $a_{<t} \in A_{<t}$, and $C \in \mathcal{M}(A_{0t})$. For any $b', b \in B$ define $\|b' - b\| = \sup |b'_{it}(C|s_{it}) - b_{it}(C|s_{it})|$, where the supremum is over all $it \in L$, $s_{it} \in S_{it}$, and $C \in \mathcal{M}(A_{it})$.

Sequences of completely mixed strategies play a critical role in defining sequential equilibrium in finite games but are unavailable in infinite games when any player has a continuum of actions. So we extend to infinite games the concept of a sequence of completely mixed strategies by using instead nets of strategies whose tails give every action positive probability.[14]

For any $b \in B$, say that the net $\{b^\alpha\}$ of strategy profiles is *admissible for* b iff $\lim_\alpha \|b^\alpha - b\| = 0$,[15] and, for every $it \in L$, for every $s_{it} \in S_{it}$, and for every $a_{it} \in \Phi_{it}(s_{it})$, there is an index $\bar{\alpha}$ in the directed index set such that $b^\alpha_{it}(\{a_{it}\}|s_{it}) > 0$ for every $\alpha \geq \bar{\alpha}$.

Notice that in any finite game, if a sequence (and therefore a net) of strategy profiles is admissible for some strategy profile, then the sequence of strategies converges to that strategy profile and, far enough out in the sequence, all strategies always give all available actions positive probability. So admissible sequences of strategies in finite games correspond exactly to the kinds of sequences that are required to define sequential equilibria there.

For any $b \in B$, it is easy to construct a net that is admissible for b as follows. Let $A_I = \times_{it \in L} A_{it}$ denote the set of action profiles. The index set for our net will be the set, Ω, of all ordered pairs (n, F) such that n is any positive integer and F is any nonempty finite subset of A_I. This index set is a directed set when we partially order its elements by saying that (n', F') is at least as large as (n, F) iff $n' \geq n$ and $F' \supseteq F$. For any $(n, F) \in \Omega$, let F_{it} be the projection of F onto A_{it}.

For any $(n, F) \in \Omega$, for any $it \in L$, and for any $s_{it} \in S_{it}$, define $\tilde{b}^{n,F}_{it}(\cdot|s_{it})$ to be uniform on $F_{it} \cap \Phi_{it}(s_{it})$ if this intersection is nonempty, and define $\tilde{b}^{n,F}_{it}(\cdot|s_{it}) = b_{it}(\cdot|s_{it})$ otherwise.[16] Define $b^{n,F}_{it}(\cdot|s_{it}) = \left(1 - \frac{1}{n}\right) b_{it}(\cdot|s_{it}) + \frac{1}{n}\tilde{b}^{n,F}_{it}(\cdot|s_{it})$. Then, the action-probabilities assigned by $b^{n,F}_{it}$ are always within $\frac{1}{n}$ of those assigned by b_{it}, and $b^{n,F}_{it}(\cdot|s_{it})$ gives positive probability to

[14] All nets will be indexed by superscripts, e.g., $\{b^\alpha\}$. It will always be implicit that the net's set of indices comes equipped with a partial order that makes the index set a directed set, i.e., for every pair of indices, there is another index that is weakly greater than both.

[15] This limit means that, for every $\varepsilon > 0$ there exists an index $\bar{\alpha}$ in the net's directed index set such that $\|b^\alpha - b\| < \varepsilon$ for every $\alpha \geq \bar{\alpha}$.

[16] The multi-stage game measurability condtion on Φ_{it} specified in Γ.7 ensures that $\tilde{b}^{n,F}_{it}$ has the measurability property required of a transition probability.

every action in $F_{it} \cap \Phi_{it}(s_{it})$. In particular, for any $a_{it} \in \Phi_{it}(s_{it})$, $b_{it}^{n,F}(\{a_{it}\}|s_{it}) > 0$ whenever F_{it} contains a_{it}. These properties of each $b_{it}^{n,F}$ imply that the net $\{b^{n,F}\}_{(n,F)\in\Omega}$ is admissible for b.

We next define admissible nets of perturbations of nature's probability function p. For these nets, we require only that almost every state of nature (as opposed to every state) receive positive tail-probability, as formalized below. This allows one to capture the idea that after some out of equilibrium history, it is common knowledge among the players that some states that could explain that history are nevertheless impossible – e.g., states that are outside the support of nature's distribution. The simplest way to ensure that players always consider a state (with prior probability zero) to be impossible is to give it probability zero in every element of a net of perturbations.[17]

For any date t, for any $C \subseteq A_{0t} \times A_{<t}$, and for any $a_{<t} \in A_{<t}$ let $C_{a_{<t}} = \{a_{0t} \in A_{0t} : (a_{0t}, a_{<t}) \in C\}$ be the slice of C through $a_{<t}$.

Given nature's probability function, p, say that a net $\{p^\alpha\}$ of nature perturbations is *admissible for* p iff $\lim_\alpha \|p^\alpha - p\| = 0$, and, for any date t there is a measurable subset C of $A_{0t} \times A_{<t}$ such that for any $a_{<t} \in A_{<t}$, $p_t(C_{a_{<t}}|a_{<t}) = 1$ and, for any $a_{0t} \in C_{a_{<t}}$, there is an index $\bar{\alpha}$ in the directed index set such that $p_t^\alpha(\{a_{0t}\}|a_{<t}) > 0$ for every $\alpha \geq \bar{\alpha}$.[18]

For any $b \in B$, say that a net $\{(b^\alpha, p^\alpha)\}$ of strategy profiles and nature perturbations is *admissible* for (b,p) iff $\{b^\alpha\}$ is admissible for b and $\{p^\alpha\}$ is admissible for p.[19]

We can now state one of our central definitions.

Definition 6.2 *For any $\varepsilon > 0$, a strategy profile $b \in B$ is a perfect conditional ε-equilibrium iff there is a net $\{(b^\alpha, p^\alpha)\}$ of player strategies and nature perturbations that is admissible for (b,p) such that for every α, b^α is a conditional ε-equilibrium of the game $\Gamma(p^\alpha)$. The net $\{(b^\alpha, p^\alpha)\}$ is then called an ε-test net (for (b,p)).*

In a perfect conditional ε-equilibrium, behavior is ε-rational in all events given positive probability in the tail of an admissible net, which should be interpreted to mean in all events outside a "strategically irrelevant" set. The next definition makes this precise.

Say that a measurable subset N of A is *negligible* iff $P(N|b) = 0$ for every $b \in B$. So a negligible set is strategically irrelevant because, in positive probability events, it cannot be given positive probability by any strategy profile.

[17] For example, the "canonical" nets of nature perturbations defined in Section 8.1 eventually (in their tail) give probability zero to states outside the support of p. See footnote 42.

[18] Like admissible nets of strategies, admissible nets of nature perturbations are easily constructed.

[19] Notice that if $\{b^\gamma\}$ is admissible for b and $\{p^\delta\}$ is admissible for p, then defining $(b^{(\gamma,\delta)}, p^{(\gamma,\delta)}) = (b^\gamma, p^\delta)$ for each γ and δ, and partially ordering (γ,δ) pairs coordinatewise, i.e., $(\gamma',\delta') \geq (\gamma,\delta)$ iff $\gamma' \geq \gamma$ and $\delta' \geq \delta$, we obtain that the net $\{(b^{(\gamma,\delta)}, p^{(\gamma,\delta)})\}$ is admissible for (b,p).

We can now state the following result, which says that every outcome in the game outside a negligible set receives positive probability in the tail of any admissible net of strategies and nature-perturbations. A proof is in Myerson and Reny (2019).[20]

Theorem 6.3 *If $\{(b^\alpha, p^\alpha)\}$ is admissible for (b, p), then there is a negligible set of outcomes $N \subseteq A$ such that, for every $a \in A \backslash N$, there is an index $\bar{\alpha}$ such that $P(\{a\}|b^\alpha; p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$.*

In standard finite multi-stage games, we can relate perfect conditional ε-equilibria to strategy profiles that are part of a sequential equilibrium, henceforth sequential equilibrium strategy profiles. A proof is in Myerson and Reny (2019).

Theorem 6.4 *In any standard finite multi-stage game, the following conditions are equivalent.*

(a) *$b \in B$ is a sequential equilibrium strategy profile.*

(b) *$b \in B$ is a perfect conditional ε-equilibrium for every $\varepsilon > 0$, and*

(c) *$b \in B$ is the limit as $\varepsilon \to 0$ of a sequence of perfect conditional ε-equilibria.*

Given this result, it would be natural to extend the definition of sequential equilibrium to infinite games by defining $b \in B$ to be a "perfect conditional equilibrium" if and only if it is a perfect conditional ε-equilibrium for every $\varepsilon > 0$, or, if and only if it is the limit as $\varepsilon \to 0$ of a sequence of perfect conditional ε-equilibria. But such strategy profiles need not exist, even in very well-behaved infinite games.[21] So in the next section, we instead consider sequences (nets) of perfect conditional ε-equilibria and the limits of their outcome distributions as $\varepsilon \to 0$.

6.2 Perfect Conditional Equilibrium Distributions

We now define a "perfect conditional equilibrium distribution" as a limit of perfect conditional ε-equilibrium distributions on outcomes as $\varepsilon \to 0$.

[20] The idea of the proof is as follows. By admissibility, for any date t, there is $C^t \in \mathcal{M}(A_{0t} \times A_{<t})$ such that for every $a_{<t} \in A_{<t}$, $p_t(C^t_{a_{<t}}|a_{<t}) = 1$ and, for every $a_{0t} \in C^t_{a_{<t}}$ there is an index $\bar{\alpha}$ such that $p^\alpha_t(\{a_{0t}\}|a_{<t}) > 0$ for every $\alpha \geq \bar{\alpha}$. Let N be the union of $N_1, ..., N_T$, where $N_t = \{a \in A : (a_{0t}, a_{<t}) \notin C^t\}$. Then N is negligible because each N_t is negligible (by the definitions in Section 3.1).

[21] See, e.g., Example 2 in Milgrom and Weber (1985), Van Damme (1987), Börgers (1991), and Section 2 in Harris et. al. (1995).

Definition 6.5 *A mapping $\mu : \mathcal{M}(A) \to [0,1]$ is a perfect conditional equilibrium distribution iff there is a net $\{b^\alpha\}$ of perfect conditional ε_α-equilibria such that $\lim_\alpha \varepsilon_\alpha = 0$ and,*

$$\mu(C) = \lim_\alpha P(C|b^\alpha), \text{ for every } C \in \mathcal{M}(A).\text{[22]} \tag{6.1}$$

It follows immediately from (6.1) that if μ is a perfect conditional equilibrium distribution, then μ is a finitely additive probability measure on $\mathcal{M}(A)$.[23] The next result is an immediate consequence of the equivalence of (a) and (c) in Theorem 6.4.

Theorem 6.6 *In any standard finite multi-stage game, the set of perfect conditional equilibrium distributions is precisely the set of distributions over outcomes induced by the set of sequential equilibria.*

The existence of perfect conditional ε-equilibria is taken up in Section 9.1. We record here the simpler result, based on Tychonoff's theorem, that a perfect conditional equilibrium distribution exists so long as perfect conditional ε-equilibria always exist. A proof is in Myerson and Reny (2019).

Theorem 6.7 *If for each $\varepsilon > 0$ there is at least one perfect conditional ε-equilibrium, then a perfect conditional equilibrium distribution exists.*

If (6.1) holds, then so long as u_i is bounded and measurable (as we have assumed),

$$\lim_\alpha \int_A u_i(a) P(da|b^\alpha) = \int_A u_i(a)\mu(da), \tag{6.2}$$

and so we define i's *equilibrium expected payoff* (at μ) by

$$\int_A u_i(a)\mu(da).$$

Sometimes μ is only finitely additive, not countably additive (e.g., the leading example in Harris et. al., 1995). Even so, in many practical settings there is a natural countably additive probability measure over outcomes that is induced by μ.

Definition 6.8 *Suppose that A is a normal topological space and $\mathcal{M}(A)$ is its Borel sigma-algebra.*[24] *We say that $\dot{\mu}$ is the regular countably additive distribution induced by μ iff*

[22] The directed index set can always be chosen so that each element is of the form $\alpha = (\varepsilon, \mathcal{F})$, where ε is any positive real number, $\mathcal{F}$ is any finite collection of measurable subsets of A, and smaller values of ε and more inclusive finite collections $\mathcal{F}$ correspond to larger indices.

[23] For any disjoint sets $C, D \in \mathcal{M}(A)$, (6.1) and $\lim_\alpha P(C \cup D|b^\alpha) = \lim_\alpha [P(C|b^\alpha) + P(D|b^\alpha)]$ imply that $\mu(C \cup DH) = \mu(C) + \mu(D)$.

[24] Recall that a topological space is normal if any pair of disjoint closed sets can be separated by disjoint open sets.

$\dot{\mu}$ is a regular countably additive probability measure on $\mathcal{M}(A)$ such that $\int f(a)\dot{\mu}(da) = \int f(a)\mu(da)$ for all bounded continuous $f : A \to \mathbb{R}$.[25]

In most applications, e.g., whenever A is a compact Hausdorff space with its Borel sigma algebra of measurable sets, μ induces a regular countably additive distribution $\dot{\mu}$.[26] In this case, player i's equilibrium expected payoff (at μ), namely $\int_A u_i(a)\mu(da)$, is equal to $\int_A u_i(a)\dot{\mu}(da)$ whenever $u_i : A \to \mathbb{R}$ is a continuous function.[27]

6.3 Other Properties

The next result states that every perfect conditional ε-equilibrium is a conditional ε-equilibrium, and therefore also an ε-Nash equilibrium. A proof is in Myerson and Reny (2019).[28] The proof uses the fact that signal-event-probabilities in a perfect conditional ε-equilibrium are well-approximated by the ε-test net.

Theorem 6.9 *Every perfect conditional ε-equilibrium is a conditional ε-equilibrium and therefore, a fortiori, an ε-Nash equilibrium.*

Given perfect recall, we may say that a date-t history $a_{<t} \in A_{<t}$ is a *subgame* of Γ iff $\sigma_{it}^{-1}(\sigma_{it}(a_{<t})) = \{a_{<t}\}$, for all $i \in I$.

For any $\varepsilon > 0$, say that a strategy profile $b \in B$ is a *subgame perfect ε-equilibrium of* Γ iff there is a negligible subset N of A such that for every $a \in A \backslash N$ and for every date t, if $a_{<t}$ is a subgame then,

$$\sup_{c_i \in B_i} U_i(c_i, b_{-i}|a_{<t}) \leq U_i(b|a_{<t}) + \varepsilon, \text{ for every } i \in I. \tag{6.3}$$

Our next result states that perfect conditional ε-equilibria induce ε-Nash equilibria in all subgames outside a strategically irrelevant set. A proof is in Myerson and Reny (2019). The result is a consequence of the fact that, in an ε-test net for a perfect conditional ε-equilibrium, every outcome, and so also every subgame, outside a negligible set eventually has positive probability. So conditional on all such subgames, play must be ε-optimal.

[25] There can be at most one such Borel measure $\dot{\mu}$ since, by Theorem IV.6.2 in Dunford and Schwartz (1988), any two such measures must agree on all closed sets. Then, by Corollary 1.6.2 in Cohn (1980), the two measures must agree on all Borel sets since the set of closed sets is closed under finite intersections and generates the Borel sigma algebra.

[26] This follows from the Riesz representation theorem, an observation for which we are grateful to a referee.

[27] When the outcome distribution $\dot{\mu}$ can not be supported by any strategy profile, it can sometimes be supported by a correlated strategy (as happens in the examples listed in footnote 21). Whether this is true in general is not known.

[28] It is also straightforward to show that the set of perfect conditional ε-equilibria is closed under the $||\cdot||$-norm on B, and the set of perfect conditional equilibrium distributions is compact in the product topology on $[0, 1]^{\mathcal{M}(A)}$. We omit the straightforward proofs.

Theorem 6.10 *Every perfect conditional ε-equilibrium is a subgame perfect ε-equilibrium.*

Our perfect conditional ε-equilibrium concept does not specify beliefs for the players. Instead, the players' beliefs are implicitly specified through a net of perturbations that tests for ε-optimal behavior. Next, we provide one way to define systems of beliefs so that KW's consistency condition for standard finite games extends to infinite games.

6.4 Conditional Belief Systems and Sequential ε-Rationality

For any $it \in L$ and for any $Z \in \mathcal{M}(S_{it})$, say that Z is *observable* iff there is $b \in B$ such that $P_{it}(Z|b) > 0$. A player's behavior conditional on any signal event that is not observable is irrelevant since, in positive probability events, no behavior can make an unobservable event have positive probability.

Definition 6.11 *A conditional belief system β specifies, for every $it \in L$, and for every observable $Z \in \mathcal{M}(S_{it})$, a finitely additive probability measure $\beta_{it}(\cdot|Z)$ on the measurable subsets of $A_{<t}$ such that $\beta_{it}(\sigma_{it}^{-1}(Z)|Z) = 1$.*[29]

So a conditional belief system specifies, for any observable set of signals and for any dated player, a finitely additive probability measure over histories that gives probability one to the set of all histories that generate signals in the given set.

Definition 6.12 *For any $b \in B$ and for any conditional belief system β, say that (b, β) is Bayes consistent iff for all $it \in L$, for all $C \in \mathcal{M}(A_{<t})$, and for all measurable $Z \subseteq S_{it}$ such that $P_{it}(Z|b) > 0$, $\beta_{it}(C|Z) = P_{<t}(C|Z, b)$.*

So Bayes' consistency disciplines beliefs only on signal events that have positive probability under the given strategy profile. If (b, β) is Bayes consistent then we also say that β is Bayes consistent (with b).

One way to extend to infinite games KW's definition of a belief system that is consistent with a given strategy profile is the following.

Definition 6.13 *For any $b \in B$ and for any conditional belief system β, say that (b, β) is finitely consistent iff there is a net $\{(b^\alpha, p^\alpha)\}$ in $B \times \mathcal{T}$ that is admissible for (b, p) such that for every $it \in L$ and for every observable $Z \in \mathcal{M}(S_{it})$,*

$$\beta_{it}(C|Z) = \lim_{\alpha} P_{<t}(C|Z, b^\alpha; p^\alpha), \ \forall C \in \mathcal{M}(A_{<t}).^{30} \tag{6.4}$$

[29] The set $\sigma_{it}^{-1}(Z)$ is nonempty because, in a multi-stage game, each signal function $\sigma_{it} : A_{<t} \to S_{it}$ is onto, i.e., its range is S_{it}.

If (b, β) is finitely consistent then we can also say that β is finitely consistent with b.

Importantly, for any $b \in B$, there is a conditional belief system β that is finitely consistent with b.[31] Moreover, if (b, β) is finitely consistent then it is Bayes' consistent and β exhibits many additional consistency properties.[32]

We next extend KW's definition of sequential rationality to infinite games.

Definition 6.14 *For any $\varepsilon \geq 0$, for any $b \in B$ and for any conditional belief system β, say that (b, β) is sequentially ε-rational iff for every $it \in L$ and for every observable $Z \in \mathcal{M}(S_{it})$,*

$$\int U_i(c_i, b_{-i}|a_{<t})\beta_{it}(da_{<t}|Z) \leq \int U_i(b|a_{<t})\beta_{it}(da_{<t}|Z) + \varepsilon, \text{ for every } c_i \in B_i$$

It is easy to verify that if (b, β) is Bayes consistent and sequentially ε-rational, then b is a conditional ε-equilibrium. We also have the following result, whose proof is in Myerson and Reny (2019). The proof uses an ε-test net for b to construct beliefs β as in (6.4). Sequential ε-rationality then follows by continuity given that each element of the test-net is a conditional ε-equilibrium.

Theorem 6.15 *If $b \in B$ is a perfect conditional ε-equilibrium, then there is a belief system β such that (b, β) is finitely consistent and sequentially ε-rational.*

But the converse fails. That is, (b, β) can be finitely-consistent and sequentially ε-rational even though b is not a perfect conditional ε-equilibrium (see Example 1 in Myerson and Reny 2019).

A well-known property of consistent beliefs in standard finite games is that players with the same information must have the same beliefs about the history of play. This property extends to infinite games and finitely consistent beliefs. Indeed, and even more generally, suppose that at any date t, two players can each distinguish between the measurable set of histories $C_{<t}$ and its complement, i.e., for each player, no date-t history outside $C_{<t}$ generates the same signal as any history in $C_{<t}$. Suppose also that $C_{<t}$ can have positive probability

[30] An implication of Theorem 6.3 is that, because Z is observable and $\{(b^\alpha, p^\alpha)\}$ is admissible for (b, p), there is $\bar{\alpha}$ such that $P_{it}(Z|b^\alpha; p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. See Myerson and Reny (2019), Corollary to Theorem 6.3. So the conditional probability on the right-hand side of (6.4) is well-defined.

[31] For example, let $\{(b^\alpha, p^\alpha)\}$ be any net that is admissible for (b, p). Then (see footnote 30), $\{(P_{<t}(C|Z, b^\alpha; p^\alpha))_{it \in L, Z \in \mathcal{M}(S_{it}), C \in \mathcal{M}(A_{<t})}\}_\alpha$ is a net taking values in a space that is an infinite product of the compact set $[0, 1]$. By Tychonoff's theorem this space is compact and so a convergent subnet can be extracted to define beliefs as in (6.4).

[32] For example, for all $it \in L$, for all $W, Z \in \mathcal{M}(S_{it})$ and for all $C \in \mathcal{M}(A_{<t})$,

$$\beta_{it}(\sigma_{it}^{-1}(W)|Z)\beta_{it}(C \cap \sigma_{it}^{-1}(Z)|W) = \beta_{it}(\sigma_{it}^{-1}(Z)|W)\beta_{it}(C \cap \sigma_{it}^{-1}(W)|Z).$$

under some strategy profile. Then, for any finitely consistent beliefs, the two players must have the same beliefs over $C_{<t}$ conditional on each of their signal sets that is generated by $C_{<t}$.[33] Because this holds in particular when, for each player, $C_{<t}$ generates a single signal, we may conclude in addition that with finitely consistent beliefs, whenever any two players have the same information about the history of play, they must have the same beliefs.

Since for any perfect conditional ε-equilibrium b there are conditional beliefs β such that (b, β) is finitely consistent and sequentially ε-rational, the discussion in the previous paragraph implies that, in any perfect conditional ε-equilibrium, any two players with the same information about any observable event behave as if they have the same beliefs.

Despite having some good properties, finitely consistent beliefs can sometimes seem paradoxical. Indeed, returning to Example 6.1, consider any $|| \cdot ||$-convergent net $\{(b^\alpha, p^\alpha)\}$ in $B \times T$ such that $p^\alpha = p$ is constant and equal to nature's probability function and such that for each $\delta > 0$ and for each action of player 1, player 1's net of strategies eventually always gives that action positive probability and eventually always gives the strictly dominant action $x = -1$ probability at least $1 - \delta$. Any such net defines a net of belief systems that has a limit point (by Tychonoff's theorem). Moreover, each limit point is finitely consistent. But, as we have seen, these finitely consistent beliefs would put probability 0 on the event that player 1's action is $x = -1$ conditional on each of player 2's signals, even though this event should get probability 1 conditional on player 2's entire set of signals.

Thus, with only finite consistency, beliefs on one-point signal events may not be sufficient to determine beliefs more generally. In particular, the probability assigned to any set of histories conditional on any given signal event need not be a convex combination of the probabilities assigned to that set of histories conditional on each element of an arbitrary partition of that signal event. However, that probability can always be obtained as a convex combination of the conditional probabilities given each element of any *finite* partition of that event.

Much applied work on signaling games (for example) has relied on an implicit assumption that beliefs conditional on one-point signal events should be sufficient to characterize beliefs for all larger observable events. When beliefs computed pointwise are not sufficient to evaluate the sequential rationality of a strategy, this beliefs-based approach can become more difficult and so perhaps less useful.

[33] This is because their signal sets are observable events (since $C_{<t}$ can have positive probability), and because, by equation (6.4), if signal sets Z_{it} and Z_{jt} satisfy $C_{<t} = \sigma_{it}^{-1}(Z_{it}) = \sigma_{jt}^{-1}(Z_{jt})$, then for every $H \in \mathcal{M}(A_{<t})$, $\beta_{it}(H|Z_{it}) = \beta_{jt}(H|Z_{jt}) = \lim_\alpha P_{<t}(H \cap C_{<t}|b^\alpha; p^\alpha)/P_{<t}(C_{<t}|b^\alpha; p^\alpha)$.

7 Illustrative Examples

In this section we present two examples showing that perturbations of nature can sometimes lead to perfect conditional ε-equilibria that may seem unintuitive.[34]

Example 7.1 *Unintuitive consequences of non-independent perturbations of independent states of nature.*

- On date 1, nature chooses $\theta = (\theta_1, \theta_2)$ uniformly from the square $[-1, 3] \times [-1, 3]$.
- On date 2, player 1 observes θ_1 and chooses $x \in \{-1, 1\}$.
- On date 3, player 2 observes x and chooses $y \in \{-1, 1\}$.
- Payoffs are: $u_1 = xy$ and $u_2 = \theta_2 y$.

Since no player receives any information about θ_2, and $\mathbb{E}(\theta_2) > 0$, player 2 should choose $y = 1$ regardless of the action of player 1 that she observes. But then player 1 should also choose $x = 1$ regardless of the value of θ_1 that he observes. Hence, the intuitively natural equilibrium expected payoff vector is $(u_1, u_2) = (1, 1)$.

But consider the pure strategy profile (b_{12}, b_{23}) where $b_{12}(\theta_1) = [-1]$ if $\theta_1 > -1$, $b_{12}(-1) = [1]$, and $b_{23}(x) = [-x]$.[35]

This strategy profile yields the expected payoff vector $(u_1, u_2) = (-1, 1)$, but it is nonetheless a perfect conditional ε-equilibrium for any $\varepsilon > 0$ because it can be supported by a perturbation of nature that puts small positive probability on the event $\{\theta_1 = \theta_2 = -1\}$. With this perturbation of nature it would be sequentially rational for player 2 to choose $y = -1$ when she observes $x = 1$ because she would attribute this observation to the occurrence of the positive probability event $\{\theta_1 = \theta_2 = -1\}$ and therefore would expect the value of θ_2 to be -1.[36] ♦

This perfect conditional ε-equilibrium may seem unintuitive because θ_2 is observed by no one and is independent of everything in the game, yet, nature's supporting perturbation leads both players to believe that observing $\theta_1 = -1$ informs them that $\theta_2 = -1$ (player 1 observes $\theta_1 = -1$ directly; player 2 infers from 1's perturbed strategy that $\theta_1 = -1$ when she observes $x = 1$). Thus, the perturbations of nature that support perfect conditional ε-equilibria can

[34] The two examples considered here are infinite games. Perturbations of nature can also have dramatic effects in non-standard finite games, i.e., finite games in which some of nature's states have prior probability zero. This is because, when states with prior-probability zero receive positive probability in some perturbation, they suddenly become "possible" and therefore can explain events that could otherwise be explained only through a deviation by some player. See the discussion following Example 6.1 as well as Section 4.8 in Myerson (1991).

[35] Here and in the next example, the notation $[c]$ denotes the probability measure that puts probability 1 on the action c.

[36] In contrast, because full conditional ε-equilibrium does not require nature-perturbations, the only full conditional ε-equilibrium outcome in the limit as $\varepsilon \to 0$, is $(1, 1)$.

influence the informational content of nature's states in important, but perhaps unintended, ways.

When the game specifies that $\theta = (\theta_1, \theta_2)$ is uniform on $[-1, 3] \times [-1, 3]$, the modeler might intend for this to mean that neither one of nature's two coordinates, θ_1 and θ_2, can ever be informative about the other, even in zero-probability events. But, formally, the joint distribution only determines the distribution of, say θ_2 conditional on θ_1, for almost every value of θ_1. In particular, the distribution of θ_2 conditional on $\theta_1 = -1$ can be defined to assign all mass to $\theta_2 = -1$, as in the perturbation of nature in the present example.[37] The perturbations of nature in perfect conditional ε-equilibria fill in these indeterminacies that are present in, but are irrelevant for, standard probability theory. But because the way these indeterminacies are filled in can be crucial in a game-theoretic context, we may wish to better control how they are resolved. For example, the unintuitive equilibrium above can be eliminated if θ_1 and θ_2 are perturbed independently. See Section 8 for a general class of such restricted nature perturbations. Alternatively, we could apply the concept of full conditional ε-equilibrium from Section 5, which would exclude the perverse equilibrium for this example and the next one.

Example 7.2 *Unintuitive consequences of large perturbations of nature even with small probability.*

- On date 1, nature chooses $\theta = (\theta_1, \theta_2) \in [0, 1]^2$. With probability 1/2, the coordinates are independent and uniform on $[0, 1]$, and with probability 1/2 the coordinates are equal and uniform on $[0, 1]$.
- On date 2, player 1 observes $s_{12} = \theta_1$ and chooses $x \in \{-1, 1\}$.
- On date 3, player 2 observes $s_{23} = x$ and chooses $y \in \{-1, 1\}$.
- Payoffs are: $u_1 = xy$, and $u_2 = y(1/3 + \theta_2 - \theta_1)$.

Thus, player 2 should choose $y = 1$ if she expects $\theta_2 - \theta_1$ to be greater than $-1/3$ and she should choose $y = -1$ otherwise. Player 1 wants to choose an action that player 2 will match.

Since for almost every θ_1, θ_2 is equally likely to be equal to θ_1 (in which case $\theta_2 - \theta_1 = 0$) as to be uniform on $[0, 1]$ (in which case $\mathbb{E}(\theta_2 - \theta_1 | \theta_1) = 1/2 - \theta_1$), player 2 should expect

[37] Notice that this would not be true if θ_2 were chosen after θ_1. Then, the distribution of θ_2 would be specified by nature's transition probability function for any possible value of θ_1. In this case, the game model could specify that θ_2 is uniform on $[-1, 3]$ for every possible θ_1, which would eliminate the problem in this example. However, even then, the same problem would arise in a modified example with two additional players, 3 and 4, who, separately from players 1 and 2, play the same game, with player 3 playing the role of player 1 and player 4 playing the role of player 2, and where the roles of θ_1 and θ_2 are reversed, i.e., player 3 observes θ_2, and player 4's payoff depends on θ_1. In this modified game, the problem cannot be eliminated by specifying the temporal order in which θ_1 and θ_2 occur because each would have to occur before the other. But the refinements introduced in Sections 8 and 5 can eliminate the problem even in this modified example.

$\theta_2 - \theta_1$ to be no smaller than $-1/4$, regardless of 1's strategy. So player 2 should choose $y = 1$.

Thus, it seems that all sensible equilibria involve strategies that give probability 1 to $(x, y) = (1, 1)$.

But consider the strategy profile (b_{12}, b_{23}) where $b_{12}(\theta_1) = [-1]$ if $\theta_1 \neq 1$, $b_{12}(1) = [1]$, and $b_{23}(x) = [-x]$.[38] This profile gives probability 1 to $(x, y) = (-1, 1)$, and is supported in a perfect conditional ε-equilibrium by the perturbation of nature that does not perturb θ_2 but that with small positive probability perturbs the distribution of θ_1 so that it is a mass point on $\theta_1 = 1$. With this perturbation of nature it is conditionally rational for player 2 to choose $y = -1$ when she observes $x = 1$ because she attributes this observation to θ_1 being a mass point on 1 and therefore expects the value of $\theta_2 - \theta_1$ to be $-1/2$.[39] ♦

Once again, we have an unintuitive equilibrium that can result because the joint distribution of nature's state coordinates determines the conditionals only almost everywhere. This unintuitive equilibrium can be eliminated if nature's states can be perturbed only to nearby states so as to approximately maintain the informativeness of each coordinate θ_1 and θ_2 about the other (see the next section), or, if we apply the concept of full conditional ε-equilibrium from Section 5.

8 Augmenting a Game with a Net of Admissible Nature-Perturbations

Unintuitive perfect conditional ε-equilibria such as in Examples 7.1 and 7.2 can be eliminated if we augment a game by including in its specification a net of admissible perturbations of nature.

If $\{p^\alpha\}$ is admissible for nature's probability function p in the multi-stage game Γ, then we say that a *perfect conditional ε-equilibrium b of Γ is compatible with* $\{p^\alpha\}$ iff there is net $\{(b^\gamma, p^\gamma)\}$ of strategy profiles and nature-perturbations that is admissible for (b, p) such that $\{p^\gamma\}$ is a subnet of $\{p^\alpha\}$ and, for each γ, b^γ is a conditional ε-equilibrium of the game $\Gamma(p^\gamma)$.

For any multi-stage game Γ, its specified net of admissible nature-perturbations $\{p^\alpha\}$ should be thought of as an additional element in the structure of the game and that expresses common knowledge aspects of how players update their beliefs about nature in zero-probability events. No additional topological structure is needed to augment a game with an admissible net of nature-perturbations. However, in most applications, the various spaces

[38] See footnote 35.

[39] In contrast, because full conditional ε-equilibrium does not require nature-perturbations, the limit as $\varepsilon \to 0$ of full conditional ε-equilibria gives probability 1 to the action profile $(x, y) = (1, 1)$.

come equipped with natural topologies. We next define canonical nets of admissible nature perturbations that can be used in a wide variety of such applications.

8.1 Canonical Nets of Admissible Nature-Perturbations

We need to add something to the structure of the game because nature's probability function p may not tell us enough about what information could be inferred from observing the state of nature to be in some set that had probability 0. As Examples 7.1 and 7.2 demonstrated, even when two random variables are independent according to the prior probability distribution p, we may need some additional structure if we want to stipulate that even the observation of a probability-0 event defined by one of these random variables would still not convey any information about the other random variable. More generally, we may want to express nature's state at any date t as being composed of several different coordinates $a_{0t} = (a_{0tj})_{j \in J}$ (for some index set J), so that we can stipulate that the inference from observing a probability-0 event that is defined by any one coordinate a_{0tj} should not go beyond the range of what could be inferred from any positive-probability events that are defined by this random variable a_{0tj}.

We might also want to specify some partition on the possible values of any coordinate a_{0tj}, so that we can stipulate that the inference from an observation of a_{0tj} taking a value in any partition element should not go beyond the range of what could be inferred from observing a_{0tj} in positive-probability subsets of this partition element. Further, with a topology on the set of possible values of a_{0tj}, we may also want to specify that the players' inferences from the observation of a probability-0 value of a_{0tj} must be a limit of what could be inferred from observing a_{0tj} to be in arbitrarily small positive-probability neighborhoods of this observed value.

So let Γ be any multi-stage game with probability function for nature $p \in \mathcal{T}$. But suppose that there is a finite index set, J with $\#J \geq 1$, such that, for any date t, nature's set of date t states is written as $A_{0t} = \times_{j \in J} A_{0tj}$, where each A_{0tj} is a separable metric space with its Borel sigma algebra of measurable sets. Suppose also that, for each $j \in J$ there is a finite or countably infinite partition Q_{0tj} of A_{0tj} into measurable sets, and denote by $Q_{0tj}(a_{0tj})$ the element of Q_{0tj} that contains $a_{0tj} \in A_{0tj}$.

With this structure, we can define a *canonical net of nature-perturbations* $\{p^\alpha\}$ *for* p as follows.[40] Let $A_0 = \times_{t \leq T} A_{0t}$ be nature's state space. The index set for our net will be the set, Ω, of all ordered pairs (n, F) such that n is any positive integer and F is any nonempty finite subset of A_0. This index set is a directed set when we partially order its elements by

[40] Different partitions, Q_{0tj}, will give different canonical nets. The possibility of controlling the perturbations in the net by choosing particular partitions Q_{0tj} can be useful, as in our proof of Theorem 9.3.

saying that (n', F') is at least as large as (n, F) iff $n' \geq n$ and $F' \supseteq F$. For any $(n, F) \in \Omega$, for any date t, and for any $j \in J$, let F_{0tj} be the projection of F onto A_{0tj}.

For any $(n, F) \in \Omega$, for any date t and for any $j \in J$, define the transition probability $\phi_{tj}^{n,F} : A_{0tj} \to \Delta(A_{0tj})$ so that, for every $a_{0tj} \in A_{0tj}$, if no point in $F_{0tj} \cap Q_{0tj}(a_{0tj})$ is within distance $\frac{1}{n}$ of a_{0tj} then $\phi_{tj}^{n,F}(\{a_{0tj}\}|a_{0tj}) = 1$. Otherwise, $\phi_{tj}^{n,F}(\{a_{0tj}\}|a_{0tj}) = 1 - \frac{1}{n}$ and $\phi_{tj}^{n,F}(\cdot|a_{0tj})$ distributes the remaining probability $\frac{1}{n}$ uniformly over the finite set of points in $F_{0tj} \cap Q_{0tj}(a_{0tj})$ that are within distance $\frac{1}{n}$ of a_{0tj}.

For any $(n, F) \in \Omega$, define the perturbation of nature, $p^{n,F}$, as follows. For every date $t \leq T$, for every $a_{<t} \in A_{<t}$, and for every $C = \times_{j\in J} C_j \in \times_{j \in J} \mathcal{M}(A_{0tj})$,

$$p_t^{n,F}(C|a_{<t}) = \int_{A_{0t}} \Pi_{j\in J} \phi_{tj}^{n,F}(C_j|a_{0t}) p_t(da_{0t}|a_{<t}). \tag{8.1}$$

The perturbation $p^{n,F}$ works as follows. At each date t, and after any history $a_{<t} \in A_{<t}$, a provisional state a_{0t} is first drawn according to nature's date-t probability measure $p_t(\cdot|a_{<t})$. Then, independently for each coordinate $j \in J$, the actual j-th coordinate of the date-t state is drawn according to the distribution $\phi_{tj}^{n,F}(\cdot|a_{0tj})$, depending only on the j-th coordinate of the provisional state.[41,42]

So for large $(n, F) \in \Omega$, each perturbation $p^{n,F}$ in a canonical net perturbs nature's coordinates independently to nearby values, and only rarely. Perturbing nature's coordinates independently and to nearby values ensures that observation of any coordinate value can only convey information that could be available from events in small neighborhoods of that value. Perturbing nature only rarely ensures that, at each date, the anticipation of future perturbations of nature will not affect future expected values in the limit as $n \to \infty$ and as F expands to include all of nature's states. We can state the following result. A proof is in Myerson and Reny (2019).[43]

Theorem 8.1 *If $\{p^{n,F}\}$ is a canonical net of nature-perturbations, then $\{p^{n,F}\}$ is admissible for p.*

[41] In particular, for each date t, $p_t^{n,F}(\cdot|\cdot)$ is a Blackwell garbling of $p_t(\cdot|\cdot)$.

[42] By the definition of the $\phi_{tj}^{n,F}$ mappings, for any a_{0t} in the support of $p_t^{n,F}(\cdot|a_{<t})$ there is a'_{0t} in the support of $p_t(\cdot|a_{<t})$ such that a_{0tj} is within distance $1/n$ of a'_{0tj} for each j. Consequently, any a_{0t} outside the support of $p_t(\cdot|a_{<t})$ is given probability zero by $p_t^{n,F}(\cdot|a_{<t})$ for all large enough n.

[43] The idea of the proof is as follows. Since the support of a measure (i.e., the smallest closed set with measure-zero complement) is well-defined in a separable metric space, for any date t we can let $C^t = \{(a_{0t}, a_{<t}) : a_{0t}$ is in the support of $p_t(\cdot \cap Q_{0t}(a_{0t})|a_{<t})\}$, where $Q_{0t}(a_{0t})$ is the element of $\times_{j\in J} Q_{0tj}$ that contains a_{0t}. Then, for any $a_{<t} \in A_{<t}$, $p_t(C^t_{a_{<t}}|a_{<t}) = 1$, where $C^t_{a_{<t}}$ is the slice of C^t through $a_{<t}$. Moreover, for any $\bar{a}_{0t} \in C^t_{a_{<t}}$, and for any $(n, F) \in \Omega$ with $\bar{a}_{0t} \in F$, (8.1) implies that $p_t^{n,F}(\{\bar{a}_{0t}\}|a_{<t}) > 0$ because there is a small enough open set U containing $\bar{a}_{0t}$ such that $\Pi_{j\in J}\phi_{tj}^{n,F}(\{\bar{a}_{0t}\}|a_{0t}) > 0$ for every $a_{0t} \in U \cap Q_{0t}(\bar{a}_{0t})$ and $p_t(U \cap Q_{0t}(\bar{a}_{0t})|a_{<t}) > 0$ (the latter since $\bar{a}_{0t} \in C^t_{a_{<t}}$ is in the support of $p_t(\cdot \cap Q_{0t}(\bar{a}_{0t})|a_{<t})$).

To eliminate the unintuitive perfect conditional equilibria in example 7.1, we should set $J = \{1, 2\}$ and let $A_{011} = A_{012} = [-1, 3]$. Then $\theta = (\theta_1, \theta_2) \in A_{011} \times A_{012}$, and in any perturbation of nature from the canonical net, the coordinates θ_1 and θ_2 of nature's state θ will be perturbed independently. Moreover, in any perturbation from the canonical net in which some state $\bar{\theta} = (\bar{\theta}_1, \bar{\theta}_2)$ receives positive probability, the conditional distribution of θ_1 given $\theta_2 = \bar{\theta}_2$ will be uniformly close to a uniform distribution on $[-1, 3]$. Consequently, with this specification of the coordinates of nature, the unintuitive equilibrium fails to be a perfect conditional ε-equilibrium that is compatible with the canonical net of nature-perturbations. For finite games with a discrete topology on chance moves, there would be a positive distance between any two alternative moves by nature, and so every canonical net of nature-perturbations is eventually (in the net) constant and equal to nature's original probability function. So compatibility with the canonical net effectively rules out any nature-perturbation at all.

9 Regular Projective Games

In this section we introduce a large class of games – regular projective games – for which we can prove the existence of a perfect conditional ε-equilibrium which has full support and is compatible with the canonical net of nature perturbations, for any $\varepsilon > 0$.

Definition 9.1 *Let $\Gamma = (I, T, S, A, \mathcal{M}, \Phi, p, \sigma, u)$ be a multi-stage game. Then Γ is a* regular projective game *iff there is a finite index set J and, for all $(n, r, j) \in I^* \times T \times J$ there are sets A_{nrj} such that, for every $it \in L$,*

R.1. *$A_{it} = \times_{j \in J} A_{itj}$, $\Phi_{it}(s_{it}) = A_{it}$ for every $s_{it} \in S_{it}$, and $A_{0t} = \times_{j \in J} A_{0tj}$,*

R.2. *if $t > 1$, then there is a nonempty set $M_{it} \subset I^* \times \{1, \ldots, t-1\} \times J$ such that $S_{it} \subseteq \times_{nrj \in M_{it}} A_{nrj}$ and $\sigma_{it}(a_{<t}) = (a_{nrj})_{nrj \in M_{it}}$ $\forall a_{<t} \in A_{<t}$ is a projection map; that is, i's signal at date $t > 1$ is just a list of state coordinates and action coordinates from dates up to t, (define $M_{i1} = \emptyset$ and recall that $S_{i1} = \{\emptyset\}$ in a multi-stage game)*

R.3. *A_{itj} and A_{0tj} are nonempty compact metric spaces $\forall j \in J$, and all product spaces are given their product topologies, all subspaces are given their relative topologies, and the measurable subsets of all spaces are their Borel subsets,*

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

R.5. *nature's date-t probability function satisfies $p_t(C|a_{<t}) = \int_C f_t(a_{0t}|a_{<t})[\times_{j \in J} \rho_{0tj}](da_{0t})$ $\forall C \in \mathcal{M}(A_{0t})$, $\forall a_{<t} \in A_{<t}$, where $\rho_{0tj} \in \Delta(A_{0tj})$ has full support $\forall j \in J$, and where*

$f_t : A_{0t} \times A_{<t} \to [0, \infty)$ is continuous and the subset of $A_{0t} \times A_{<t}$ on which f_t is strictly positive is closed.

If Γ satisfies R.1 and R.2, we may say that Γ is a projective game *or a* game with projected signals.

Remark 9.2 *(1) One can always reduce the cardinality of J to $(T+1)^{\#I}$ or less by grouping, for any $it \in L^*$, the variables $\{a_{itj}\}_{j \in J}$ according to the $\#I$-vector of dates at which the players observe them, if ever.*

(2) Since distinct players can observe the same a_{0tj}, nature's probability function in a regular projective multi-stage game need not satisfy the information diffuseness assumption of Milgrom-Weber (1985). Nevertheless, the form of p_t assumed in R.5 of Definition 9.1 is reminiscent of the Milgrom-Weber assumption, and a recent counterexample to the existence of an (ex-ante) ε-Nash equilibrium in a Bayesian game due to Simon and Tomkowicz (2017) shows that some such assumption is necessary for the existence of even a conditional ε-equilibrium.

(3) Continuity of f_t implies that the subset of $A_{0t} \times A_{<t}$ on which f_t is strictly positive is open. Hence, condition R.5 implies that points of zero density are topologically isolated from points of strictly positive density. This is a restrictive condition, but it is always true for finite games with the discrete topology and for games with each f_t strictly positive. Without a condition of this kind, likelihood ratios can become unbounded in ways that our proof technique cannot handle.

Examples of regular projective games include the following.

1. **All Finite multi-stage games.** Any finite multi-stage game (i.e., finite state, action, and signal sets endowed with their discrete topologies) can be modeled as a regular projective game simply by letting each player's signal be a coordinate of the state.

2. **Compact and Continuous Multi-Stage Games.** The following compact and continuous games (i.e., all state, action and signals sets are compact metric spaces, payoff functions are continuous, nature moves only on date 1 with a date-1 probability function that is absolutely continuous with respect to the product of its marginals and with a continuous and positive Radon-Nikodym derivative) are regular projective games.

   (i) **Bayesian games.** In an N-player Bayesian game (Harsanyi model), there are two dates. On date 1, nature chooses a state vector with N coordinates. On date 2, each player i observes only the i-th coordinate of nature's date 1 state and chooses a feasible action. Payoffs can depend on all actions and on nature's state vector.

(ii) **Finite-Horizon Multi-Stage Games with Observed Actions**. In an N-player T-stage game with observed actions, all players have perfect recall. On date 1, nature chooses a state vector with N coordinates. On date 2, each player i observes only the i-th coordinate of nature's date 1 state and chooses a feasible action. On any date $t \in \{3, ..., T\}$, player i observes the actions taken by all players on the previous date and then chooses a feasible action. Payoffs can depend on all actions taken by all players on all dates and on nature's date-1 state vector.

(iii) **Signaling Games.** In a signaling game, there are three dates. On date 1, nature chooses a state. On date 2, player 1 (the "sender") observes nature's date-1 state and chooses a feasible ("message"). On date 3, player 2 (the "receiver") observes the action chosen by player 1 and then chooses a feasible action. Payoffs can depend on the actions of both players and on nature's state.

3. **Stochastic Games.** Any finite-horizon stochastic game (which includes all finitely-repeated games) in which nature's transition probability depends on the history only through a continuous and positive conditional density function. (We can take $\#J = 1$ since players observe the entire past history on each date.)

9.1 Existence

We can now state our main existence result, whose proof is in Section 11. It states that, in regular projective games, for every $\varepsilon > 0$ there is strategy profile that is a perfect conditional ε-equilibrium, and that, in addition, has full support (and so by Remark 9.4 is a full conditional ε-equilibrium) and is compatible with a canonical net of nature perturbations.

Theorem 9.3 *Let Γ be a regular projective game. Then for any $\varepsilon > 0$, Γ has a perfect conditional ε-equilibrium that has full support and that is compatible with a canonical net of nature-perturbations.*

Remark 9.4 *By Theorem 6.10 and Theorem 9.3, every regular projective game has a subgame perfect ε-equilibrium for every $\varepsilon > 0$.*[44] *By Theorem 6.9 and Theorem 9.3, every regular projective game has a full conditional ε-equilibrium.*

An immediate consequence of Theorem 9.3 and Theorem 6.7 is the following.

Theorem 9.5 *Every regular projective game Γ has a perfect conditional equilibrium distribution μ.*

[44] See Chakrabarti (1999) for an existence result concerning a related concept, subgame perfect approximate equilibria, for a different class of games.

10 Conclusion

In order to ensure that all off-path behavior (outside a negligible set) is rational in an infinite game, we have been led to perturb not only the players' strategies (as in KW), but to perturb nature's probability function as well. Although the effects of nature's perturbations can sometimes seem unintuitive, the strategy profiles that arise as perfect conditional ε-equilibria satisfy two fundamental properties. For any finite set of outcomes in the game (outside a negligible set) (i) (*finite consistency*) all players can agree on a common perturbation of nature's probability function and on a common perturbation of the players' equilibrium strategies that together give positive probability to – and so can explain the occurrence of – any of those outcomes, and (ii) (*conditional ε-optimality*) if any player were ever to observe a signal on the path to any of those outcomes, then the common explanation of the outcomes that generate that signal would make his equilibrium continuation behavior given that signal ε-optimal.

In a topological approach to the problem of rationality in extensive form games, it is natural to consider conditional ε-equilibria that have full support (with the given topologies). This full conditional ε-equilibrium concept is attractive because it does not require any perturbations of nature. However, as we have seen (Example 5.3), to obtain properties like subgame perfection, we need to consider nets of perturbations of the players' strategies and of nature's probability function, as in our perfectness concept.

In standard finite games, the sets of conditional ε-equilibria with full support and perfect conditional ε-equilibria are essentially equivalent, and their limits yield the set of sequential equilibrium strategy profiles.[45] The fact that this coincidence of perfectness and fullness does not extend to infinite games is a basic reason why it has been so difficult to define sequential equilibria for infinite games. An uncountable infinity of outcomes cannot all get positive probability from one strategy profile, and so one must either let the strategy profile satisfy a weaker topological condition of full support, or one must consider a net of perturbations of the players' strategies and of nature that can test rationality in all events but may yield only finite additivity in the limit. We have emphasized the latter approach as a general solution, but both approaches may be worth considering in particular applications.

[45] Specifically, let Γ be any standard finite multi-stage game. Any full conditional ε-equilibrium (with the discrete topology on the finite A_{it}) of Γ is a perfect conditional ε-equilibrium. Conversely, if b is a perfect conditional ε-equilibrium of Γ, then for all $\delta > 0$ and for all $\varepsilon' > \varepsilon$, there is a full conditional ε'-equilibrium b' with $\|b' - b\| \leq \delta$.

11 Proof of Theorem 9.3.

Outline.

The proof is broken into four parts. Part 1 constructs a sufficiently fine finite partition of the space of outcomes. Part 2 uses the finite partition from part 1 to define a finite approximating game played by agents $it \in L$, and fixes one of its Nash equilibria, a full-support strategy profile $\hat{b}$ in the original infinite game. Part 3 constructs a net $\{(b^\alpha, p^\alpha)\}$ of strategy profiles and nature- perturbations that is admissible for p where the net of nature-perturbations is canonical. Part 4 shows that every strategy profile b^α in the net is a perfect conditional ε-equilibrium in perturbed game $\Gamma(p^\alpha)$. Altogether, these steps show that the full-support strategy profile $\hat{b}$ is a perfect conditional ε-equilibrium of Γ that is compatible with a canonical perturbation of nature.

Preliminaries.

Recall that in any multi-stage game, $S_{i1} = A_{<1} = \{\emptyset\}$ for every $i \in I$. So $a_{<1} = \emptyset$ for any $a \in A$.

The set of Borel subsets of any metric space X will be denoted by $\mathcal{B}(X)$.

Let Γ be a regular projective game, i.e., Γ satisfies the conditions R.1-R.5 of Definition 9.1, henceforth simply R.1-R.5.

Henceforth, we will write itj for any $(it, j) \in L^* \times J$.

For any $itj \in L^* \times J$, and for any $a_{itj} \in A_{itj}$, let $\mathbb{B}_\delta(a_{itj})$ denote the δ-ball centered at a_{itj} (A_{itj} is a metric space by R.3).

If X_{itj} is any subset of A_{itj} for each $itj \in L^* \times J$ and if K is any subset of $L^* \times J$, then let $X_K = \times_{itj \in K} X_{itj}$, and let $a_K = (a_{itj})_{itj \in K}$ denote a typical element of X_K.

Recall by R.2 that for every $it \in L$, the subset M_{it} of $I^* \times \{1, ..., t-1\} \times J$ is the set of history-coordinates that player i observes at date t. Hence, $\sigma_{it}(a_{<t}) = a_{M_{it}}$ for every $a \in A$, and $S_{it} = A_{M_{it}}$. Throughout the proof we will often denote player i's set of signals by $A_{M_{it}}$ and we will often denote a typical signal for player i at date t by $a_{M_{it}} \in A_{M_{it}}$. By convention, we define $A_\emptyset = \{\emptyset\}$, $a_\emptyset = \emptyset$, and $M_{i1} = \emptyset$ for every $i \in I$.

Let $\rho_0 = (\rho_{01}, ..., \rho_{0T})$, where, for each date t, $\rho_{0t} = \times_{j \in J} \rho_{0tj}$ is the product carrying measure for nature's date t state as specified in R.5. Then ρ_0 is an element of $\mathcal{T}$, the set of alternative probability functions for nature in the game Γ.

If $\zeta : X \to \Delta(Y)$ and $\kappa : Y \to \Delta(Y)$ are any pair of transition probabilities, then define the transition probability $\zeta * \kappa : X \to \Delta(Y)$ so that for every $x \in X$, and for every measurable subset C of Y,

$$[\zeta * \kappa](C|x) = \int_Y \kappa(C|y)\zeta(dy|x). \tag{11.1}$$

Fix any positive real number ε for the remainder of the proof. The steps below establish the existence of a strategy profile with full support that is a perfect conditional ε-equilibrium of Γ that is compatible with a canonical net of nature-perturbations..

Part 1. (construct a sufficiently fine finite partition of the space of outcomes)

We shall construct a finite partition of the space of outcomes so that, within each element of the partition, the players' utilities have sufficiently small variation and so that nature's density function has bounded relative likelihoods on each partition element that can have positive probability.

Since in a multi-stage game each u_i is bounded, we may choose $\bar{m} > 0$ so that

$$\max_{a,a' \in A} (u_i(a) - u_i(a')) \leq \bar{m}, \ \forall i \in I. \tag{11.2}$$

The set of outcomes in the regular projective game Γ is the product set $A = \times_{itj \in L^* \times J} A_{itj}$. For every $a \in A$, define

$$f(a) = \Pi_{r \leq T} f_r(a_{0r}|a_{<r}), \tag{11.3}$$

define

$$g(a) = \begin{cases} f(a), & \text{if } f(a) > 0 \\ 1, & \text{if } f(a) = 0, \end{cases} \tag{11.4}$$

define

$$h(a) = \begin{cases} 1, & \text{if } f(a) > 0 \\ 0, & \text{if } f(a) = 0. \end{cases} \tag{11.5}$$

and define

$$H(a) = \Pi_{t \leq T} \rho_{0t}(\{a'_{0t} \in A_{0t} : f_t(a'_{0t}|a_{<t}) > 0\}). \tag{11.6}$$

Consequently, for every $a \in A$,

$$f(a) = g(a)h(a). \tag{11.7}$$

Since f_i is continuous and somewhere positive on the compact set $A_{0t} \times A_{<t}$, it achieves a maximum, $\bar{m}_t > 0$ say, on $A_{0t} \times A_{<t}$. Hence, for any $a_{<t} \in A_{<t}$,

$$\begin{aligned} \rho_{0t}(\{a'_{0t} \in A_{0t} : f_t(a'_{0t}|a_{<t}) > 0\} &= \int_{\{a'_{0t} \in A_{0t} : f_t(a'_{0t}|a_{<t}) > 0\}} \rho_{0t}(da_{0t}) \\ &\geq \int \frac{f_t(a_{0t}|a_{<t})}{\bar{m}_t} \rho_{0t}(da_{0t}) \\ &= 1/\bar{m}_t \\ &> 0. \end{aligned} \tag{11.8}$$

Consequently, $H(a)$ is bounded away from zero for $a \in A$. Notice also that, being a product of probabilities, $H(a) \leq 1$ for $a \in A$.

Since, by R.5, the set of outcomes on which f is strictly positive is closed, and since, by continuity, the set of outcomes on which f is zero is closed, g is continuous on A. Since g is strictly positive, it therefore achieves a positive minimum on the compact set A. So because H is positive and bounded away from zero on A, we may choose $\lambda \in (0,1)$ and $\gamma > 0$ so that,

$$2\gamma + (1-(1-\lambda)^{T(\#J)})\bar{m} \leq \varepsilon \left(\inf_{a \in A} g(a)H(a) \right). \tag{11.9}$$

For any nonempty sets $X_1, ..., X_K$ and for any partitions $\mathcal{P}_1$ of $X_1, ..., \mathcal{P}_K$ of X_K, let $\mathcal{P}_1 \otimes ... \otimes \mathcal{P}_K$ denote the (product) partition of $X_1 \times ... \times X_K$ defined by $\mathcal{P}_1 \otimes ... \otimes \mathcal{P}_K = \{E_1 \times ... \times E_K : E_k \in \mathcal{P}_k \ \forall k\}$.

We claim that we may choose a finite product partition, $Q = \otimes_{itj \in L^* \times J} Q_{itj}$ of A, composed of Borel measurable partitions Q_{itj} of A_{itj} $\forall itj \in L^* \times J$ such that for any $a, a' \in A$ in the same element of the partition Q,

$$|u_i(a)g(a) - u_i(a')g(a')| < \gamma, \text{ for every player } i \in I, \text{ and} \tag{11.10}$$

$$f_t(a_{0t}|a_{<t}) > 0 \Leftrightarrow f_t(a'_{0t}|a'_{<t}) > 0, \text{ for every date } t. \tag{11.11}$$

Let us justify this claim. For each $i \in I$, $u_i g$ is continuous on the compact set A and so $u_i g$ is uniformly continuous on A. The compactness of the A_{itj} sets ensures that, for any positive diameter, we can partition each A_{itj} into finitely many measurable sets each with that diameter or less. If that diameter is sufficiently small, then the uniform continuity of $u_i g$ on A implies that (11.10) will be satisfied. To see that (11.11) must also be satisfied for some sufficiently small diameter, notice that otherwise there would be a date t and two sequences of points in $A_{0t} \times A_{<t}$ that approach one another such that along one of the sequences f_t is strictly positive and along the other f_t is zero. By compactness, we may assume that both sequences converge, and hence they converge to the same point. But the assumption that f_t is strictly positive on a closed set would then imply that f_t is strictly positive at the limit point, and the fact (by continuity) that f_t is zero on a closed set would imply that f_t is zero at the limit point, yielding a contradiction and establishing the claim.

For any $itj \in L^* \times J$ and for any $a_{itj} \in A_{itj}$, let $Q_{itj}(a_{itj})$ denote the element of the partition Q_{itj} that contains a_{itj}, and let $Q_{it} = \otimes_{j \in J} Q_{itj}$ be the finite partition of A_{it} that is generated by the partitions Q_{itj}, $j \in J$. We henceforth assume that $Q = \otimes_{i \in I^*, t \leq T,, j \in J} Q_{itj}$ satisfies (11.10) and (11.11).

Part 2. (define a finite approximating game played by agents $it \in L$ and fix one of its equilibria)

For each $it \in L$, and for each $j \in J$, let us choose $\nu_{itj} \in \Delta(A_{itj})$ so that for each element of the finite partition Q_{itj} of A_{itj}, ν_{itj} gives positive probability to every point in a dense subset of that partition element. Such a ν_{itj} exists since, by R.3, A_{itj} is a compact metric space and therefore every partition element has a countable dense subset, to each element of which ν_{itj} can give positive probability. In particular, ν_{itj} gives positive probability to every open subset of A_{itj}. For each $it \in L$, let ν_{it} denote the history-independent strategy in B_{it} in which player i at date t chooses from A_{it} according to the product probability $\times_{j\in J}\nu_{itj}$ regardless of the date-t signal that he observes.

Given the $\lambda > 0$ chosen in (11.9), for every $itj \in L \times J$ define the transition probability $\Lambda_{itj} : A_{itj} \to \Delta(A_{itj})$ as follows. For any $a_{itj} \in A_{itj}$ and for any $D \in \mathcal{B}(A_{itj})$,

$$\Lambda_{itj}(D|a_{itj}) = (1-\lambda)\nu_{itj}(D|Q_{itj}(a_{itj})) + \lambda\nu_{itj}(D),$$

where $\nu_{itj}(D|Q_{itj}(a_{itj})) = \nu_{itj}(D\cap Q_{itj}(a_{itj}))/\nu_{itj}(Q_{itj}(a_{itj}))$ is the conditional ν_{itj}-probability of D given $Q_{itj}(a_{itj})$. (Recall that $\nu_{itj}(Q_{itj}(a_{itj})) > 0$ for every $a_{itj} \in A_{itj}$.) So for any $a_{itj} \in A_{itj}$, $\Lambda_{itj}(\cdot|a_{itj})$ chooses an element from A_{itj} according to $\nu_{itj}(\cdot|Q_{itj}(a_{itj}))$ with probability $1-\lambda$ and according to ν_{itj} with probability λ.

Then (see (11.1)), for any $b_{it} \in B_{it}$, $b_{it} * \Lambda_{it}$ is the date-t strategy for player i that, given any signal $a_{M_{it}} \in A_{M_{it}}$, first chooses a provisional $a_{it} \in A_{it}$ according to $b_{it}(\cdot|a_{M_{it}})$, and then, independently for each coordinate j, chooses the actual coordinate-j action according to $\nu_{itj}(\cdot|Q_{itj}(a_{itj}))$ with probability $1-\lambda$ and according to ν_{itj} with probability λ. In particular, because there is positive probability that all of the coordinate-j actions are chosen according to ν_{itj}, $b_{it} * \Lambda_{it}$ gives positive probability to each element of the finite partition Q_{it} of A_{it} and gives positive probability to every open subset of A_{it}, no matter what signal player i observes at date t. This last fact implies that, for any $b \in B$, the strategy profile $(b_{it} * \Lambda_{it})_{it\in L}$ has full support in the game Γ.

Define the probability function for nature $\tilde{p} = (\tilde{p}_1, ..., \tilde{p}_T) \in \mathcal{T}$ so that for every date t, for every $a_{<t} \in A_{<t}$, and for every $D \in \mathcal{B}(A_{0t})$,

$$\tilde{p}_t(D|a_{<t}) = \frac{\rho_{0t}(D \cap \{a_{0t} : f_t(a_{0t}|a_{<t}) > 0\})}{\rho_{0t}(\{a_{0t} : f_t(a_{0t}|a_{<t}) > 0\})}, \tag{11.12}$$

where the denominator is strictly positive by (11.8).

There are two important facts to note about $\tilde{p}$. First, $\tilde{p}_t(D|\cdot)$ is measurable with respect to the product partition Q (and so, in particular, $\tilde{p}_t : A_{<t} \to \Delta(A_{0t})$ is a transition probability). Indeed, if $a', a'' \in A$ are in the same element of Q, then by (11.11), $\{a_{0t} : f_t(a_{0t}|a'_{<t}) > 0\} = \{a_{0t} : f_t(a_{0t}|a''_{<t}) > 0\}$ and so $\tilde{p}_t(D|a'_{<t}) = \tilde{p}_t(D|a''_{<t})$. (This same argument implies also that $H(a') = H(a'')$, i.e., that H is Q-measurable, a fact that we will use below.) Second, given

the probability function $\tilde{p}$, and after any date-t history $a_{<t}$, the distribution of nature's date t state conditional on any positive probability element $q_{0t} = \times_{j\in J} q_{0tj} \subseteq A_{0t}$ of the partition $Q_{0t} = \otimes_{j\in J} Q_{0tj}$ is given by the history-independent product measure $\rho_{0t}(\cdot|q_{0t}) = \times_{j\in J}\rho_{0tj}(\cdot|q_{0tj})$, where $\rho_{0t}(\cdot|q_{0t})$ denotes the conditional of ρ_{0t} given $q_{0t} = \times_{j\in J} q_{0tj}$, and similarly for $\rho_{0tj}(\cdot|q_{0tj})$. This is because, if $\tilde{p}_t(q_{0t}|a_{<t}) > 0$, then by (11.11) $q_{0t} \subseteq \{a_{0t} : f_t(a_{0t}|a_{<t}) > 0\}$ and so (11.12) implies that $\tilde{p}_t(D\cap q_{0t}|a_{<t})/\tilde{p}_t(q_{0t}|a_{<t}) = \rho_{0t}(D\cap q_{0t})/\rho_{0t}(q_{0t})$.

Let $\xi : A \to A$ be measurable with respect to Q (i.e., constant on each partition element) and such that, for every $a \in A$, $\xi(a)$ is in the same element of the partition Q as a.

For every $i \in I$, and for every $a \in A$, define $v_i(a) = u_i(a)g(a)H(a)$.

Let the game $\Gamma_{v\circ\xi}(\tilde{p})$ be identical to Γ except that, for each $i \in I$, player i's payoff function is $v_i(\xi(a)) = u_i(\xi(a))g(\xi(a))H(\xi(a))$ instead of $u_i(a)$, and nature's probability function is $\tilde{p}$ instead of p.

Notice that $H(\xi(a)) = H(a)$ because, as observed in the paragraph following (11.12), H is Q-measurable. Consequently, because $0 \leq H \leq 1$, (11.10) implies that for every $a \in A$,

$$|u_i(\xi(a))g(\xi(a))H(\xi(a)) - u_i(a)g(a)H(a)| < \gamma, \text{ for every player } i \in I. \tag{11.13}$$

For each $itj \in L \times J$, select precisely one action from each element of the partition Q_{itj} of A_{itj} and let the finite set of all of the selected actions be denoted by $\bar{C}_{itj}$. Let $\bar{C}_{it} = \times_{j\in J}\bar{C}_{itj}$.

Let $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ denote the agent normal form of $\Gamma_{v\circ\xi}(\tilde{p})$ in which each dated player $it \in L$ is a separate agent and is restricted to strategies $b_{it} \in B_{it}$ of the form $b_{it} = \tilde{b}_{it} * \Lambda_{it}$ for some $\tilde{b}_{it} \in B_{it}$ that is measurable with respect to Q and that assigns probability 1 to the finite set $\bar{C}_{it}$.

For any strategy b_{it} that is feasible for an agent it in $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$, and for any $q_{it} = (q_{itj})_{j\in J} \in Q_{it}$, the conditional distribution of b_{it} given q_{it} is the product measure $\times_{j\in J}\nu_{itj}(\cdot|q_{itj})$. So the coordinates of it's actions are always chosen conditionally independently. Also, for any signal $a_{M_{it}} \in A_{M_{it}}$, the probability measure $b_{it}(\cdot|a_{M_{it}})$ chooses an action in A_{it} according to $\Pi_{j\in J}\nu_{itj}$ with probability at least $\lambda^{\#J} > 0$. Consequently, every strategy profile $b \in B$ that is feasible in $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ has full support in the original game Γ.

The Q-measurability condition means that for any $nrj \in L^* \times J$ and for any signal coordinate $a_{nrj} \in A_{nrj}$ that a player observes in the original (regular projective) infinite game, he observes (can condition on) in $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ only the partition element in Q_{nrj} that contains a_{nrj}. Hence, in $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$, for any date $t > 1$, a signal w_{it} for agent $it \in L$ is any $\times_{nrj\in M_{it}} q_{nrj}$, where $q_{nrj} \in Q_{nrj}$ $\forall nrj \in M_{it}$.[46] Let W_{it} denote the finite set of it's signals in the game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$. Then $W_{it} = \otimes_{nrj\in M_{it}} Q_{nrj}$ is a finite partition of player i's date-t signal space $A_{M_{it}}$ in the original infinite game Γ.

[46] For $t = 1$, agent it's signal in the game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ is always equal to the null signal, $\emptyset$.

Together, the measurability condition and the fact that the support of each agent's strategy in $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ is always a subset of a fixed finite set of actions, imply that $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ is a finite game.

Let $\hat{b} \in B$ be a Nash equilibrium of the finite game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ played by agents $it \in L$. Then, in particular, $\hat{b}$ is of the form $(\tilde{b}_{it} * \Lambda_{it})_{it\in L}$ for some $\tilde{b} \in B$, and $\hat{b}$ is a full-support strategy profile in the original game Γ.

The remainder of the proof will establish that the full-support strategy profile $\hat{b}$ is a perfect conditional ε-equilibrium of Γ that is compatible with a canonical net of perturbations.

Part 3. (define a net $\{(b^\alpha, p^\alpha)\}$ that is admissible for $(\hat{b}, p)$, and where $\{p^\alpha\}$ is a subnet of a canonical net for p given the partition $Q = \otimes_{t\leq T, j\in J} Q_{0tj}$ of $\times_{t\leq T, j\in J} A_{0tj}$)

The index set for our net will be the set, Ω, of all ordered pairs (n, F) such that n is any positive integer and F is any nonempty finite subset A. This index set is a directed set when we partially ordered its elements by saying that (n', F') is at least as large as (n, F) iff $n' \geq n$ and $F' \supseteq F$. For any $(n, F) \in \Omega$, and for any $itj \in L^*$, let F_{itj} be the projection of F onto A_{itj}.

For any index $(n, F) \in \Omega$, and for any $itj \in L^* \times J$, define the transition probability $\phi^{n,F}_{itj}$: $A_{itj} \to \Delta(A_{itj})$ so that, for every $a_{itj} \in A_{itj}$, if no point in $F_{itj} \cap Q_{itj}(a_{itj})$ is within distance $\frac{1}{n}$ of a_{itj} then $\phi^{n,F}_{itj}(\{a_{itj}\}|a_{itj}) = 1$. Otherwise, $\phi^{n,F}_{itj}(\{a_{itj}\}|a_{itj}) = 1 - \frac{1}{n}$ and $\phi^{n,F}_{itj}(\cdot|a_{itj})$ distributes the remaining probability $\frac{1}{n}$ uniformly over the finite set of points in $F_{itj} \cap Q_{itj}(a_{0tj})$ that are within distance $\frac{1}{n}$ of a_{0tj}.

For any index (n, F) and for any $it \in L^*$, define the transition probability $\phi^{n,F}_{it} : A_{it} \to \Delta(A_{it})$ so that for every $C = \times_{j\in J} C_j \in \times_{j\in J}\mathcal{M}(A_{itj})$, and for every $a_{it} \in A_{it}$, $\phi^{n,F}_{it}(C|a_{it}) = \Pi_{j\in J}\phi^{n,F}_{itj}(C_j|a_{itj})$.

Define a net of strategy profiles and nature perturbations $\{(b^{n,F}, p^{n,F})\}$ as follows. For every index $(n, F) \in \Omega$ and for every $it \in L$, define $b^{n,F}_{it} = \hat{b}_{it} * \phi^{n,F}_{it}$ and define $p^{n,F}_t = p_t * \phi^{n,F}_{0t}$. Then (see Section 8.1), $\{p^{n,F}\}_{(n,F)\in\Omega}$ is a subnet of a canonical net of perturbations of p and so, by Theorem 8.1, $\{p^{n,F}\}$ is admissible for p.[47] We next show that $\{b^{n,F}\}_{(n,F)\in\Omega}$ is admissible for $\hat{b}$, from which we can conclude that $\{(b^{n,F}, p^{n,F})\}$ is admissible for $(\hat{b}, p)$.

Since $b^{n,F}_{it} = \hat{b}_{it} * \phi^{n,F}_{it}$, we have $\left\|b^{n,F} - \hat{b}\right\| \leq \frac{1}{n}$ for every index (n, F) and so $\lim_{(n,F)} \left\|b^{n,F} - \hat{b}\right\| = 0$ since $\lim_{(n,F)} n = +\infty$. Fix any $it \in L$, fix any $s_{it} \in S_{it}$, and fix any $a_{it} \in \Phi_{it}(s_{it})$. To show that $\{b^{n,F}\}$ is admissible for $\hat{b}$, we must show that there is an index $(\bar{n}, \bar{F})$ such that $b^{n,F}_{it}(a_{it}|s_{it}) > 0$ for every $(n, F) \in \Omega$ such that $n \geq \bar{n}$ and $F \supseteq \bar{F}$. Choose $(\bar{n}, \bar{F})$ so that $a_{itj} \in \bar{F}_{itj}$ for every $j \in J$. and let (n, F) be any index such that $n \geq \bar{n}$ and $F \supseteq \bar{F}$. Hence, $a_{itj} \in F_{itj}$

[47] For any $(n, F) \in \Omega$, let F_0 be the projection of F onto nature's states $\times_{t\leq T} A_{0t}$. The net $\{p^{n,F}\}$ defined here is a subnet of the canonical net because $p^{n,F}$ is equal to the canonical perturbation of nature defined in Section 8.1 for the canonical index (n, F_0).

for every $j \in J$. Also, $\phi_{itj}^{n,F}(\{a_{itj}\}|a'_{itj}) > 0$ for any $a'_{itj} \in Q_{itj}(a_{itj})$ that is within distance $\frac{1}{n}$ of a_{itj}. But since the product measure $\times_{j\in J}\nu_{itj}$ is absolutely continuous with respect to $\hat{b}_{it}(\cdot|s_{it})$ (by the definition of $\hat{b}_{it}$), and since each ν_{itj} gives positive probability to each action in a dense subset of $Q_{itj}(a_{itj})$, $\hat{b}_{it}(\cdot|s_{it})$ gives positive probability to every action in a dense subset of $\times_{j\in J}Q_{itj}(a_{itj})$. In particular, $\hat{b}_{it}(\cdot|s_{it})$ gives positive probability to some $a'_{0t} \in \times_{j\in J}Q_{itj}(a_{itj})$ such that for every $j \in J$, $a'_{0tj} \in Q_{itj}(a_{0tj})$ is within distance $\frac{1}{n}$ of a_{0tj}, which implies that $\Pi_{j\in J}\phi_{itj}^{n,F}(\{a_{0tj}\}|a'_{0tj}) > 0$. Since $b_{it}^{n,F}(\{a_{0t}\}|s_{it}) \geq \Pi_{j\in J}\phi_{itj}^{n,F}(\{a_{0tj}\}|a'_{0tj})\hat{b}_{it}(\{a'_{it}\}|s_{it})$, we may conclude that $b_{it}^{n,F}(\{a_{0t}\}|s_{it}) > 0$, as desired. Hence, $\{b^{n,F}\}$ is admissible for $\hat{b}$.

Having established that $\{(b^{n,F}, p^{n,F})\}$ is admissible for $(\hat{b}, p)$, where $\{p^{n,F}\}$ is a subnet of a canonical net for p, let us note an important property of each $b^{n,F}$. For every $it \in L$, for every $a_{it} \in A_{it}$, and for every $q_{it} \in Q_{it}$, $\phi_{it}^{n,F}(q_{it}|a_{it})$ is equal to 1 if $a_{it} \in q_{it}$ and is equal to 0 otherwise. Therefore,

$$b_{it}^{n,F}(q_{it}|a_{M_{it}}) = \hat{b}_{it}(q_{it}|a_{M_{it}}). \tag{11.14}$$

So no matter what signal $a_{M_{it}}$ is observed by agent it, $b_{it}^{n,F}(\cdot|a_{M_{it}})$ generates the same distribution over the elements of the finite partition Q_{it} as does $\hat{b}_{it}(\cdot|a_{M_{it}})$.

Part 4 (show that for each index $(n, F) \in \Omega$, $b^{n,F}$ is a conditional ε-equilibrium of the perturbed game $\Gamma(p^{n,F})$)

To simplify the notation, we use α to denote a typical element (n, F) of the directed index set Ω constructed in Part 3 above. So the net $\{(b^{n,F}, p^{n,F})\}_{(n,F)\in\Omega}$ constructed in Part 3 will be denoted by $\{(b^\alpha, p^\alpha)\}_{\alpha\in\Omega}$ for the remainder of the proof.

In this part of the proof it will be useful to make explicit the dependence of the outcome distribution and of the players' expected utilities on the probability function for nature that is in effect. For example, $P(\cdot|b; p^\alpha)$ is the probability distribution over outcomes under the strategy profile b in the game $\Gamma(p^\alpha)$, i.e., in the game Γ when nature's probability function is $p^\alpha \in \mathcal{T}$ instead of p.

Fix any index $\alpha \in \Omega$, fix any $it \in L$, fix any measurable $Z \subseteq A_{M_{it}}$ such that $P_{it}(Z|b^\alpha; p^\alpha) > 0$ and fix any date-t continuation c_i of b_i^α. To complete the final step of the proof we must show that,

$$U_i(c_i, b_{-i}^\alpha|Z; p^\alpha) \leq U_i(b^\alpha|Z; p^\alpha) + \varepsilon. \tag{11.15}$$

Recall that W_{it} is a finite partition of S_{it} and the elements of W_{it} are the signals for agent it in the finite approximating game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$. Without loss of generality, we may assume that there is $w_{it} \in W_{it}$ such that $Z \subseteq w_{it}$ (since otherwise we could consider separately each $Z \cap w_{it}$ that has positive probability in $\Gamma(p^\alpha)$ under b^α, where w_{it} varies over all elements of the finite partition W_{it}).

If $\tau \in \mathcal{T}$ is any perturbation of nature, define $\tau * \phi_0^\alpha \in \mathcal{T}$ to be the perturbation of nature such that, for every date $t \leq T$,

$$[\tau * \phi_0^\alpha]_t = \tau_t * \phi_{0t}^\alpha.$$

Then, we may define $\tilde{p}^\alpha \in \mathcal{T}$ by,

$$\tilde{p}^\alpha = \tilde{p} * \phi_0^\alpha. \tag{11.16}$$

Hence, for any date t, $\tilde{p}_t^\alpha : A_{<t} \to \Delta(A_{0t})$ is a transition probability that, like $\tilde{p}_t$, is measurable with respect to Q (see the paragraph following (11.12)).

Because $f = gh$, and by the definition of $\tilde{p}$ in (11.12), we have that for every $b \in B$, and for every $C \in \mathcal{B}(A)$,

$$\begin{aligned}
P(C|b; p^\alpha) &= P(C|b; p * \phi_0^\alpha) \\
&= \int_C f(a) P(da|b; \rho_0 * \phi_0^\alpha) \\
&= \int_C g(a) h(a) P(da|b; \rho_0 * \phi_0^\alpha) \\
&= \int_C g(a) H(a) \left(\frac{h(a)}{H(a)} \right) P(da|b; \rho_0 * \phi_0^\alpha) \\
&= \int_C g(a) H(a) P(da|b; \tilde{p} * \phi_0^\alpha) \\
&= \int_C g(a) H(a) P(da|b; \tilde{p}^\alpha),
\end{aligned} \tag{11.17}$$

where the fourth equality follows because $H(a) > 0$ for every $a \in A$, and the fifth equality follows because, for every $C \in \mathcal{B}(A)$,

$$P(C|b; \tilde{p} * \phi_0^\alpha) = \int_C (h(a)/H(a))\, P(da|b; \rho_0 * \phi_0^\alpha).$$

Therefore, for any $b \in B$ and for any $C \in \mathcal{B}(A)$,

$$\int_C u_i(a) P(da|b; p^\alpha) = \int_C u_i(a) g(a) H(a) P(da|b; \tilde{p}^\alpha) \tag{11.18}$$

Because c_i and b_i^α agree on dates before t and because date-t signal event probabilities depend only on the player's strategies on dates before t, $P_{it}(Z|(c_i, b_{-i}^\alpha); p^\alpha) = P_{it}(Z|b^\alpha; p^\alpha)$. So, because $v_i(a) = u_i(a) g(a) H(a)$, (11.18) gives,[48]

[48] Notice that, by (11.17), $P(\cdot|b; p^\alpha)$ is absolutely continuous with respect to $P(\cdot|b; \tilde{p}^\alpha)$ for any $b \in B$. So in particular, $P_{it}(Z|b^\alpha; p^\alpha) > 0$ implies $P_{it}(Z|b^\alpha; \tilde{p}^\alpha) > 0$.

$$\begin{aligned}
U_i(c_i, b^\alpha_{-i}|Z; p^\alpha) &= \frac{\int_{\{a: a_{M_{it}} \in Z\}} u_i(a) P(da|(c_i, b^\alpha_{-i}); p^\alpha)}{P_{it}(Z|b^\alpha; p^\alpha)} \\
&= \frac{\int_{\{a: a_{M_{it}} \in Z\}} v_i(a) P(da|(c_i, b^\alpha_{-i}); \tilde{p}^\alpha)}{P_{it}(Z|b^\alpha; \tilde{p}^\alpha)} \frac{P_{it}(Z|b^\alpha; \tilde{p}^\alpha)}{P_{it}(Z|b^\alpha; p^\alpha)} \\
&= \left(\int v_i(a) P(da|Z, (c_i, b^\alpha_{-i}); \tilde{p}^\alpha) \right) \frac{P_{it}(Z|b^\alpha; \tilde{p}^\alpha)}{P_{it}(Z|b^\alpha; p^\alpha)}. \qquad (11.19)
\end{aligned}$$

By (11.13) and because $\{\xi(a) : a \in q\} \subseteq q$ for every $q \in Q$, we may bound the integral in parentheses on the right-hand side of (11.19) as follows:

$$\int v_i(a) P(da|Z, (c_i, b^\alpha_{-i}); \tilde{p}^\alpha) \leq \int v_i(\xi(a)) P(da|Z, (c_i, b^\alpha_{-i}); \tilde{p}^\alpha) + \gamma. \qquad (11.20)$$

Also, by (11.17), we may bound the ratio of probabilities $P_{it}(Z|b^\alpha; \tilde{p}^\alpha)/P_{it}(Z|b^\alpha; p^\alpha)$ on the right-hand side of (11.19) as follows:

$$\begin{aligned}
\frac{P_{it}(Z|b^\alpha; \tilde{p}^\alpha)}{P_{it}(Z|b^\alpha; p^\alpha)} &= \frac{\int_{\{a: a_{M_{it}} \in Z\}} P(da|b^\alpha; \tilde{p}^\alpha)}{\int_{\{a: a_{M_{it}} \in Z\}} g(a) H(a) P(da|b^\alpha; \tilde{p}^\alpha)} \\
&\leq \frac{\int_{\{a: a_{M_{it}} \in Z\}} P(da|b^\alpha; \tilde{p}^\alpha)}{\int_{\{a: a_{M_{it}} \in Z\}} (\inf_{a \in A} g(a) H(a)) P(da|b^\alpha; \tilde{p}^\alpha)} \\
&= \frac{1}{\inf_{a \in A} g(a) H(a)}. \qquad (11.21)
\end{aligned}$$

(Recall from Part 1 above that $\inf_{a \in A} g(a) H(a) > 0$.)

We next adjust the deviation c_i so that it becomes measurable with respect to Q, without changing the value of the integral on the right-hand side of (11.20). This is possible because $v_i \circ \xi$, b^α_{-i} and $\tilde{p}^\alpha$ are all measurable with respect to Q and so any achievable expected value of $v_i \circ \xi$ by i, is achievable with a strategy for i that is Q-measurable.

For any $nr \in L$, recall from Part 3 above that the finite subset $\bar{C}_{nr}$ of A_{nr} contains precisely one action from each element of the finite partition Q_{nr} of A_{nr}, and that for any $a_{nr} \in A_{nr}$, $Q_{nr}(a_{nr})$ is the element of Q_{nr} that contains a_{nr}.

For any date $r > t$, recall from Section 3 the perfect recall map $\bar{\Psi}_{irt} : S_{ir} \to S_{it}$ in Γ.6 of the definition of a multi-stage game. When $r = t$, define $\bar{\Psi}_{irt} : S_{ir} \to S_{it}$ to be the identity map.

Define the strategy $\tilde{c}_i \in B_i$ as follows. For any date $r < t$, let $\tilde{c}_{ir} = c_{ir}$ $(= b^\alpha_{ir})$. For any date $r \geq t$, for any $w_{ir} \in W_{ir}$, for any $a_{M_{ir}} \in w_{ir}$, and for any $a_{ir} \in \bar{C}_{ir}$, if $P_{ir}(w_{ir} \cap$

$\bar{\Psi}_{irt}^{-1}(Z)|(c_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}) > 0$, then let

$$\tilde{c}_{ir}(\{a_{ir}\}|a_{M_{ir}}) = \int c_{ir}(Q_{ir}(a_{ir})|s_{ir})P_{ir}(ds_{ir}|w_{ir} \cap \bar{\Psi}_{irt}^{-1}(Z), (c_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}), \tag{11.22}$$

but if $P_{ir}(w_{ir} \cap \bar{\Psi}_{itr}^{-1}(Z)|(c_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}) = 0$, then let $\tilde{c}_{ir}(\{a_{ir}\}|a_{M_{ir}}) = 1/(\#\bar{C}_{ir})$.

For each date $r \geq t$, $\tilde{c}_{ir} \in B_{ir}$ is measurable with respect to Q, and, for each $w_{ir} \in W_{ir}$ and for each $a_{M_{ir}} \in w_{ir}$, $\tilde{c}_{ir}(\bar{C}_{ir}|a_{M_{ir}}) = 1$. Consequently, $\tilde{c}_{ir} * \Lambda_{ir}$ is feasible for agent ir in the finite game $\Gamma_{v \circ \xi, \Lambda}(\tilde{p})$. Moreover, (11.22) implies that for every date $r \geq t$, for every $w_{ir} \in W_{ir}$ such that $P_{ir}(w_{ir} \cap \bar{\Psi}_{irt}^{-1}(Z)|(c_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}) > 0$, and for every $q_{ir} \in Q_{ir}$,

$$\tilde{c}_{ir}(q_{ir}|w_{ir}) = \int c_{ir}(q_{ir}|s_{ir})P_{ir}(ds_{ir}|w_{ir} \cap \bar{\Psi}_{irt}^{-1}(Z), (c_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}), \tag{11.23}$$

and so $\tilde{c}_{ir}$ conditional on w_{ir} induces the same distribution over the elements of Q_{ir} as does c_{ir} conditional on w_{ir} and Z.

In the game $\Gamma_{v \circ \xi}(\tilde{p}^{\alpha})$ and under the strategy profile b^{α}, for each date r and for any $n \in I^*$ (n may be a player or nature), the distribution of the j-th coordinate of n's date-r action/state conditional on any element q_{nrj} of the finite partition Q_{nrj} of A_{nrj}, is independent of any of the other coordinates of n's date-r action and is independent of the date-r history (for $n = 0$, see the paragraph following (11.12)). Consequently, because b^{α} and $\tilde{p}^{\alpha}$ are measurable with respect to Q, the occurrence of Z and the occurrence of any $q \in Q$ are independent events conditional on w_{it}. Therefore, because $Z \subseteq w_{it}$,

$$P(q|Z, b^{\alpha}; \tilde{p}^{\alpha}) = P(q|w_{it}, b^{\alpha}; \tilde{p}^{\alpha}) \text{ for every } q \in Q. \tag{11.24}$$

In particular, $P_{<t}(q_{<t}|Z, b^{\alpha}; \tilde{p}^{\alpha}) = P_{<t}(q_{<t}|w_{it}, b^{\alpha}; \tilde{p}^{\alpha})$ for every element $q_{<t}$ of the finite partition $Q_{<t} = \otimes_{nrj \in L^* \times J : r < t} Q_{nrj}$ of $A_{<t}$. Therefore, since changing i's behavior at dates $r \geq t$ does not affect the probability of any date-t history event,

$$P_{<t}(q_{<t}|Z, (\tilde{c}_i, b^{\alpha}); \tilde{p}^{\alpha}) = P_{<t}(q_{<t}|w_{it}, (\tilde{c}_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}) \text{ for every } q_{<t} \in Q_{<t}. \tag{11.25}$$

A consequence of (11.25), of (11.23) for $r \geq t$, and of the Q-measurability of $\tilde{c}_{ir}$, b_{jr}^{α}, and $\tilde{p}_r^{\alpha}$ for all $j \neq i$ and all $r \geq t$ is that,

$$P(q|Z, (c_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}) = P(q|w_{it}, (\tilde{c}_i, b_{-i}^{\alpha}); \tilde{p}^{\alpha}), \text{ for every } q \in Q. \tag{11.26}$$

Therefore, since $v_i(\xi(a))$ is measurable with respect to Q,

$$\int v_i(\xi(a))P(da|Z,(c_i,b^\alpha_{-i});\tilde{p}^\alpha) = \int v_i(\xi(a))P(da|w_{it},(\tilde{c}_i,b^\alpha_{-i});\tilde{p}^\alpha). \tag{11.27}$$

Recall that $\hat{b} \in B$ is a Nash equilibrium of the agent normal form game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ played by agents in L. For every date $r \leq T$, $\hat{b}_{ir}$ is a feasible strategy for agent ir in the game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$.

Define $\breve{c}_i \in B_i$ as follows. For each date $r < t$, let $\breve{c}_{ir} = \hat{b}_{ir}$, and for each date $r \geq t$, let $\breve{c}_{ir} = \tilde{c}_{ir} * \Lambda_{ir}$. Then, for every date $r \leq T$, $\breve{c}_{ir}$ is feasible for agent ir in $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$.

Because changing i's behavior at dates $r \geq t$ does not affect the probability of any date-t history event, we have $P_{it}(w_{it}|(\tilde{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) = P_{it}(w_{it}|b^\alpha;\tilde{p}^\alpha)$ because $\tilde{c}_{ir} = b^\alpha_{ir}$ for $r < t$, and we have $P_{it}(w_{it}|(\breve{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) = P_{it}(w_{it}|(\hat{b}_i,b^\alpha_{-i});\tilde{p}^\alpha)$ because $\breve{c}_{ir} = \hat{b}_{ir}$ for $r < t$. Also, by (11.14), we have $P_{it}(w_{it}|b^\alpha;\tilde{p}^\alpha) = P_{it}(w_{it}|(\hat{b}_i,b^\alpha_{-i});\tilde{p}^\alpha)$ because w_{it} is a union of elements of Q, and because both b^α and $\tilde{p}^\alpha$ are Q-measurable. Hence, we may conclude that

$$P_{it}(w_{it}|(\tilde{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) = P_{it}(w_{it}|(\breve{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) \tag{11.28}$$

By the definition of $\breve{c}_i$, there is probability at least $(1-\lambda)^{(T-t+1)(\#J)}$ that for every $r \geq t$, $\breve{c}_{ir}$ gives each element of Q_{ir} the same probability as does $\tilde{c}_{ir}$ regardless of the history of play. Consequently, by (11.28), (11.2), the measurability of $v_i(\xi(a))$ with respect to Q, and because w_{it} is a union of elements of Q, we have,

$$\int v_i(\xi(a))P(da|w_{it},(\tilde{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) \leq \int v_i(\xi(a))P(da|w_{it},(\breve{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) + (1-(1-\lambda)^{T(\#J)})\bar{m}, \tag{11.29}$$

where we have used the fact that $(1-(1-\lambda)^{(T-t+1)(\#J)})\bar{m} \leq (1-(1-\lambda)^{T(\#J)})\bar{m}$.

For every date r, for every $a \in A$, and for every $q_{0r} \in Q_{0r}$, the product $\Pi_{j\in J}\phi^\alpha_{0rj}(q_{0rj}|a_{0tj})$ is equal to 1 if $a_{0t} \in q_{0t}$ and is equal to 0 otherwise. Therefore,

$$\tilde{p}^\alpha_r(q_{0r}|a_{<r}) = \tilde{p}_r(q_{0r}|a_{<r}). \tag{11.30}$$

So no matter what is the date-t history, $\tilde{p}^\alpha_t$ generates the same distribution over the elements of the finite partition Q_{0t} of A_{0t} as does $\tilde{p}_t$.

Together, (11.30) and (11.14) imply that $P_{it}(w_{it}|\hat{b};\tilde{p}) = P_{it}(w_{it}|b^\alpha;\tilde{p}^\alpha)$. Hence, $P_{it}(w_{it}|\hat{b};\tilde{p}) > 0$ because $P_{it}(w_{it}|b^\alpha;\tilde{p}^\alpha) \geq P_{it}(Z|b^\alpha;\tilde{p}^\alpha) > 0$. And since $\breve{c}_i$ agrees with $\hat{b}_i$ on dates before t, $P_{it}(w_{it}|(\breve{c},\hat{b}_{-i});\tilde{p}) = P_{it}(w_{it}|\hat{b};\tilde{p}) > 0$. Therefore, (11.30) and (11.14) together with the measurability of $\breve{c}_i$, b^α_{-i}, and $\tilde{p}^\alpha$ with respect to Q imply that,

$$\int v_i(\xi(a))P(da|w_{it},(\breve{c}_i,b^\alpha_{-i});\tilde{p}^\alpha) = \int v_i(\xi(a))P(da|w_{it},(\breve{c}_i,\hat{b}_{-i});\tilde{p}). \tag{11.31}$$

Together, (11.20), (11.27), (11.29), and (11.31) imply that,

$$\int v_i(\xi(a))P(da|Z,(c_i,b^\alpha_{-i});\tilde{p}^\alpha) \leq \int v_i(\xi(a))P(da|w_{it},(\breve{c}_i,\hat{b}_{-i});\tilde{p}) + (1-(1-\lambda)^{T(\#J)})\bar{m}.$$

Consequently,

$$\int v_i(a)P(da|Z,(c_i,b^\alpha_{-i});\tilde{p}^\alpha) \leq \int v_i(\xi(a))P(da|w_{it},(\breve{c}_i,\hat{b}_{-i});\tilde{p}) + \gamma + (1-(1-\lambda)^{T(\#J)})\bar{m}.$$

$$\begin{aligned}
&\leq \int v_i(\xi(a))P(da|w_{it},(\breve{c}_i,\hat{b}_{-i});\tilde{p}) + \gamma + (1-(1-\lambda)^{T(\#J)})\bar{m}. \\
&\leq \int v_i(\xi(a))P(da|w_{it},\hat{b};\tilde{p}) + \gamma + (1-(1-\lambda)^{T(\#J)})\bar{m}. \\
&= \int v_i(\xi(a))P(da|w_{it},b^\alpha;\tilde{p}^\alpha) + \gamma + (1-(1-\lambda)^{T(\#J)})\bar{m}. \\
&= \int v_i(\xi(a))P(da|Z,b^\alpha;\tilde{p}^\alpha) + \gamma + (1-(1-\lambda)^{T(\#J)})\bar{m}. \\
&\leq \int v_i(a)P(da|Z,b^\alpha;\tilde{p}^\alpha) + 2\gamma + (1-(1-\lambda)^{T(\#J)})\bar{m}.
\end{aligned} \tag{11.32}$$

where the second inequality follows from the one-shot deviation principle for finite games with perfect recall because $\breve{c}_{ir} = \hat{b}_{ir}$ for $r < t$ and $\hat{b}$ gives w_{it} positive probability (see the paragraph following (11.30)) and is an equilibrium of the agent normal form of the perfect recall game $\Gamma_{v\circ\xi,\Lambda}(\tilde{p})$ when played by separate agents $nr \in L$. The first equality follows because, by (11.14) and (11.30), $P(\cdot|w_{it},\hat{b};\tilde{p})$ and $P(\cdot|w_{it},b^\alpha;\tilde{p}^\alpha)$ generate the same distribution over the elements of Q. The second equality follows from (11.24), and the final inequality follows from (11.13).

Multiplying both sides of (11.32) by $P_{it}(Z|b^\alpha;\tilde{p}^\alpha)/P_{it}(Z|b^\alpha;p^\alpha)$ and using (11.19), (11.21), and (11.9) gives,

$$
\begin{aligned}
U_i(c_i, b_{-i}^{\alpha}|Z;p^{\alpha}) &\leq \left(\int v_i(a)P(da|Z,b^{\alpha};\tilde{p}^{\alpha})\right)\frac{P_{it}(Z|b^{\alpha};\tilde{p}^{\alpha})}{P_{it}(Z|b^{\alpha};p^{\alpha})}+\varepsilon \\
&= \frac{\int_{\{a:a_{M_{it}}\in Z\}} u_i(a)g(a)H(a)P(da|b^{\alpha};\tilde{p}^{\alpha})}{P_{it}(Z|b^{\alpha};\tilde{p}^{\alpha})}\frac{P_{it}(Z|b^{\alpha};\tilde{p}^{\alpha})}{P_{it}(Z|b^{\alpha};p^{\alpha})}+\varepsilon \\
&= \frac{\int_{\{a:a_{M_{it}}\in Z\}} u_i(a)g(a)H(a)P(da|b^{\alpha};\tilde{p}^{\alpha})}{P_{it}(Z|b^{\alpha};p^{\alpha})}+\varepsilon \\
&= \frac{\int_{\{a:a_{M_{it}}\in Z\}} u_i(a)P(da|b^{\alpha};p^{\alpha})}{P_{it}(Z|b^{\alpha};p^{\alpha})}+\varepsilon \\
&= U_i(b^{\alpha}|Z;p^{\alpha})+\varepsilon,
\end{aligned}
$$

where the third equality follows from (11.17), proving (11.15). Q.E.D.

References

Bajoori, E., J. Flesch, and D. Vermeulen (2013), "Perfect equilibrium in games with compact action spaces," *Games and Economic Behavior,* 82, 490-502.

Bajoori, E., J. Flesch, and D. Vermeulen (2016), "Behavioral perfect equilibrium in Bayesian games," *Games and Economic Behavior* 98, 78-109.

Börgers, T. (1991): "Upper Hemicontinuity of the Correspondence of Subgame-Perfect Equilibrium Outcomes," *Journal of Mathematical Economics*, 20, 89–106.

Chakrabarti , S. K. (1999): "Finite and Infinite Action Dynamic Games with Imperfect Information," *Journal of Mathematical Economics*, 32, 243–266.

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

Dunford, N., and J. T. Schwartz (1988) *Linear Operators Part I General Theory*, John Wiley and Sons, New York.

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

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

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

Jung, H. M. (2018): "Extension of Sequential Equilibrium to Games with Infinite Types and Actions," mimeo, The Institute of Economics, Academia Sinica, Taipei, Taiwan.

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

Manelli, A. (1996): "Cheap Talk and Sequential Equilibria in Signaling Games," *Econometrica* 64, 917-942.

Mansuwé, A. P., Jansen, M., Peters, H. (1997): "Consistency of Assessments in Infinite Signaling Games," *Journal of Mathematical Economics*, 27, 425-449.

Milgrom, P. and R. Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*, 10, 619-32.

Myerson, R. B. (1991): *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Massachusetts.

Myerson, R. B., and P. J. Reny (2019): "Supplement to 'Perfect Conditional ε-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions'," University of Chicago.

Radner, R. (1980): "Collusive Behavior in Non-Cooperative Epsilon-Equilibria of Oligopolists with Long But Finite Lives, *Journal of Economic Theory*, 22, 136–154.

Simon, Leo K., and Maxwell B. Stinchcombe (1995), "Equilibrium refinement for infinite normal-form games," *Econometrica,* 63 (6), 1421-1443.

Simon, R. S. and G. Tomkowicz (2017): "A Bayesian Game without ε-Equilibria," *Israel Journal of Mathematics* (to appear).

Van Damme, E. (1987): "Equilibria in Non-Cooperative Games," in *Surveys in Game Theory and Related Topics*, ed. by H. Peters and 0. Vrieze. Amsterdam, C.W.I. Tract 39.

Watson, J. (2017): "A General, Practicable Definition of Perfect Bayesian Equilibrium," UCSD working paper.