Supplement to '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: This supplement contains a number of proofs not included in the published paper, a corollary that is referenced in the published paper, three examples, and demonstrates how the equilibrium concepts from the published paper can be defined in games with perfect recall outside the class of multi-stage games.

1 Supplementary Material

Throughout this supplement, Myerson and Reny (2019) is referred to as MR.

1.1 Omitted Proofs and Corollary to Theorem 6.3

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 $P(C|b) > 0$.*

Proof. Let $b \in B$ be any strategy profile with full support. For any $it \in L$, and for any $\delta > 0$, let $\mathcal{V}_{it}$ be any partition of A_{it} into measurable sets that each have nonempty interior and are of diameter δ or less.[1] For any $V \in \mathcal{V}_{it}$, for any $a_{it} \in V$, for any $s_{it} \in S_{it}$, and for any Borel

*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.

[1] For example, take $\mathcal{V}_{it}$ to consist of the nonempty elements of the sequence of disjoint sets C_1, $C_2\backslash C_1, C_3\backslash(C_1 \cup C_2), ...$, where $C_1, C_2, C_3, ...$ is any sequence of open balls of diameter δ that covers the separable metric space A_{it}.

subset C of A_{it}, define $\varphi_{it}(C|a_{it}, s_{it}) = b_{it}(C \cap V|s_{it})/b_{it}(V|s_{it})$ (the denominator is positive because V has nonempty interior and b_i has full support). Then, $\varphi_{it} : A_{it} \times S_{it} \to \Delta(A_{it})$ is a transition probability and, for any Borel subset C of A_{it}, and by the definition of $b_{it} * \varphi_{it}$ given in MR Section 5,

$$\begin{aligned}
[b_{it} * \varphi_{it}](C|s_{it}) &= \int \varphi_{it}(C|a_{it}, s_{it}) b_{it}(da_{it}|s_{it}) \\
&= \sum_{V \in \mathcal{V}} \int_V b_{it}(C \cap V|s_{it})/b_{it}(V|s_{it}) b_{it}(da_{it}|s_{it}) \\
&= \sum_{V \in \mathcal{V}} b_{it}(C \cap V|s_{it}) \\
&= b_{it}(C|s_{it}),
\end{aligned}$$

and so $b_{it} * \varphi_{it} = b_{it}$. Since this holds for every $it \in L$, $b * \varphi = (b_{it} * \varphi_{it})_{it \in L} = (b_{it})_{it \in L} = b$.

Also, for any $it \in L$, for any $s_{it} \in S_{it}$, and for any Borel subset C of A_{it}, if $b_{it}(C|s_{it}) = 0$, then, by the definition of φ_{it}, $\varphi_{it}(C|a_{it}, s_{it}) = 0$ for every $a_{it} \in A_{it}$. So for any $\hat{b}_{it} \in B_{it}$, if $b_{it}(C|s_{it}) = 0$ then $[\hat{b}_{it} * \varphi_{it}](C|s_{it}) = \int \varphi_{it}(C|a_{it}, s_{it}) \hat{b}_{it}(da_{it}|s_{it}) = 0$. Hence, for any $it \in L$ and for any $s_{it} \in S_{it}$, $[\hat{b}_{it} * \varphi_{it}](\cdot|s_{it})$ is absolutely continuous with respect to $b_{it}(\cdot|s_{it})$.

Let $\hat{b}$ be any strategy profile in B. We will show by induction on the dates in the game that $P(\cdot|\hat{b} * \varphi)$ is absolutely continuous with respect to $P(\cdot|b)$.

By (3.1) and (3.2) in MR Section 3.1, $P_{<2}(\cdot|\hat{b} * \varphi)$ is equal to the product measure $p_1(\cdot|\emptyset) \times \left(\times_{i \in I}[\hat{b}_{i1} * \varphi_{i1}](\cdot|\emptyset)\right)$, and $P_{<2}(\cdot|b)$ is equal to the product measure $p_1(\cdot|\emptyset) \times (\times_{i \in I} b_{i1}(\cdot|\emptyset))$, Since, for each $i \in I$, $[\hat{b}_{i1} * \varphi_{i1}](\cdot|\emptyset)$ is absolutely continuous with respect to $b_{i1}(\cdot|\emptyset)$, we have that $p_1(\cdot|\emptyset) \times \left(\times_{i \in I}[\hat{b}_{i1} * \varphi_{i1}](\cdot|\emptyset)\right)$ is absolutely continuous with respect to $p_1(\cdot|\emptyset) \times (\times_{i \in I} b_{i1}(\cdot|\emptyset))$. (This follows, for example, by applying the Radon-Nikodym theorem to each of the absolute continuity relations.) Hence, $P_{<2}(\cdot|\hat{b} * \varphi)$ is absolutely continuous with respect to $P_{<2}(\cdot|b)$.

As an induction hypothesis, suppose that for some date $t - 1 < T$, it is the case that $P_{<t}(\cdot|\hat{b} * \varphi)$ is absolutely continuous with respect to $P_{<t}(\cdot|b)$. (We have just shown that this statement is true for $t - 1 = 1$.) To complete the induction we must show that $P_{<t+1}(\cdot|\hat{b} * \varphi)$ is absolutely continuous with respect to $P_{<t+1}(\cdot|b)$. By (3.1) and (3.2) in MR, for every $C \in \mathcal{M}(A_{<t+1})$,

$$P_{<t+1}(C|\hat{b} * \varphi) = \int_C p_t(da_{0t}|a_{<t}) \times \left(\times_{i \in I}[\hat{b}_{it} * \varphi_{it}](da_{it}|\sigma_{it}(a_{<t}))\right) P_{<t}(da_{<t}|\hat{b} * \varphi), \quad (1.1)$$

and

$$P_{<t+1}(C|b) = \int_C p_t(da_{0t}|a_{<t}) \times (\times_{i \in I} b_{it}(da_{it}|\sigma_{it}(a_{<t}))) P_{<t}(da_{<t}|b) \quad (1.2)$$

For each $it \in L$ and for each $a_{<t} \in A_{<t}$, $[\hat{b}_{it} * \varphi_{it}](\cdot|\sigma_{it}(a_{<t}))$ is absolutely continuous with respect to $b_{it}(\cdot|\sigma_{it}(a_{<t}))$. By the induction hypothesis, $P_{<t}(\cdot|\hat{b} * \varphi)$ is absolutely continuous with respect to $P_{<t}(\cdot|b)$. Consequently, the measure defined by the right-hand side of (1.1) is absolutely continuous with respect to the measure defined by the right-hand side of (1.2). (This follows, for example, by applying the Radon-Nikodym theorem to each of the absolute continuity relations.) Hence, $P_{<t+1}(\cdot|\hat{b}*\varphi)$ is absolutely continuous with respect to $P_{<t+1}(\cdot|b)$. This completes the induction and so we may conclude that $P(\cdot|\hat{b} * \varphi)$ $(= P_{<T+1}(\cdot|\hat{b} * \varphi))$ is absolutely continuous with respect to $P(\cdot|b)$ $(= P_{<T+1}(\cdot|b))$, as desired. Q.E.D.

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 $\alpha \geq \bar{\alpha}$.*

Proof. Let $\{(b^\alpha, p^\alpha)\}$ is admissible for (b, p). By the admissibility of $\{p^\alpha\}$ for p, 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 $a_{0t} \in C^t_{a_{<t}}$ (recall that $C^t_{a_{<t}} = \{a'_{0t} \in A_{0t} : (a'_{0t}, a_{<t}) \in C^t\}$) 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 measurable since each N^t is measurable. Let us show that N is negligible.

Let b be any strategy profile in B. To show that N is negligible, it suffices to show that $P(N^t|b) = 0$ for each $t \leq T$. So consider any $t \leq T$. Since by the definitions in MR Section 3.1, the marginal of $P(\cdot|b)$ on $A_{<t+1}$ is $P_{<t+1}(\cdot|b)$, it suffices to show that $P_{<t+1}(N^t_{<t+1}|b) = 0$. Since $N^t_{<t+1} = \{a \in A_{<t+1} : a_{0t} \notin C^t_{a_{<t}}\}$, the definitions (specifically, equations (3.1) and (3.2)) in MR Section 3.1 yield,

$$\begin{aligned} P_{<t+1}(N^t_{<t+1}|b) &= \int p_t(A_{0t}\backslash C^t_{a_{<t}}|a_{<t})\Pi_{i\in I} b_{it}(A_{it}|\sigma_{it}(a_{<t}))P_{<t}(da_{<t}|b) \\ &= 0, \end{aligned}$$

where the second equality follows because $p_t(A_{0t}\backslash C^t_{a_{<t}}|a_{<t}) = 0$ for every $a_{<t} \in A_{<t}$. Hence, N is negligible.

Next, let $\tilde{a}$ be any element of $A\backslash N$. It remains only to show that there is an index $\bar{\alpha}$ such that $P(\{\tilde{a}\}|b^\alpha; p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. We proceed by induction on the dates in the game.

First, for $t = 1$, $P_1(\{(\tilde{a}_{i1})_{i\in I^*}\}|b^\alpha; p^\alpha) = p^\alpha_1(\{\tilde{a}_{01}\}|\emptyset)\Pi_{i\in I} b^\alpha_{it}(\{\tilde{a}_{i1}\}|\emptyset)$, and so, by the admissibility of $\{(b^\alpha, p^\alpha)\}$ for (b, p) and because $\tilde{a}_{01} \in C^1_{a_{<1}}$ $(= C^1_\emptyset)$, there is α_1 such that $P_1(\{(\tilde{a}_{i1})_{i\in I^*}\}|b^\alpha; p^\alpha) > 0$ for every $\alpha \geq \alpha_1$.

As an induction hypothesis, suppose that for some date $t - 1 < T$, there is an index α_{t-1} such that $P_{<t}(\{\tilde{a}_{<t}\}|b^\alpha; p^\alpha) > 0$ for every $\alpha \geq \alpha_{t-1}$. (We have just shown that this statement is true for $t-1=1$.) By the definition of $P_{<t+1}(\cdot|b^\alpha;p^\alpha)$,

$$P_{<t+1}(\{\tilde{a}_{<t+1}\}|b^\alpha;p^\alpha) = p_t^\alpha(\{\tilde{a}_{0t}\}|a_{<t})\Pi_{i\in I}b_{it}^\alpha(\{\tilde{a}_{it}\}|\sigma_{it}(\tilde{a}_{<t}))P_{<t}(\{\tilde{a}_{<t}\}|b^\alpha;p^\alpha).$$

By the induction hypothesis, there is α_{t-1} such that $P_{<t}(\{\tilde{a}_{<t}\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \alpha_{t-1}$, and, by admissibility and because $\tilde{a}_{0t} \in C^t_{\tilde{a}_{<t}}$, there is an index α' such that $p_t^\alpha(\{\tilde{a}_{0t}\}|a_{<t}) > 0$ and $b_{it}^\alpha(\{\tilde{a}_{it}\}|\sigma_{it}(\tilde{a}_{<t})) > 0$ for every $\alpha \geq \alpha'$ and for every $i \in I$. So letting $\alpha_t = \max(\alpha', \alpha_{t-1})$, we have $P_{<t+1}(\{\tilde{a}_{<t+1}\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \alpha_t$. This completes the induction step and so we may conclude that there is an index $\bar{\alpha}$ such that $P(\{\tilde{a}\}|b^\alpha;p^\alpha) = P_{<T+1}(\{\tilde{a}\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. Q.E.D.

Corollary to Theorem 6.3. Let b be any strategy profile in B and supose that $\{(b^\alpha, p^\alpha)\}$ is admissible for (b,p). Then, for any $it \in L$ and for any observable $Z \in \mathcal{M}(S_{it})$, there is $\bar{\alpha}$ such that $P_{it}(Z|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$.

Proof. By Theorem 6.3, there is a negligible set $N \in \mathcal{M}(A)$ such that, for any $a \in A\backslash N$, there is $\bar{\alpha}$ such that $P(\{a\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. Since Z is observable, there is $\tilde{b} \in B$ such that $P_{it}(Z|\tilde{b}) > 0$. By the definition of $P_{it}(\cdot|\tilde{b})$, $P_{it}(Z|\tilde{b}) = P(\{a \in A : \sigma_{it}(a_{<t}) \in Z\}|\tilde{b})$. Consequently, $P(\{a \in A : \sigma_{it}(a_{<t}) \in Z\}|\tilde{b}) > 0$ and so $\{a \in A : \sigma_{it}(a_{<t}) \in Z\}$ is not contained in the negligible set N (since $P(N|b') = 0$ for every $b' \in B$ and so, in particular, $P(N|\tilde{b}) = 0$). Pick any $\bar{a} \in \{a \in A : \sigma_{it}(a_{<t}) \in Z\}\backslash N$. Then $\bar{a} \in A\backslash N$ and so there is $\bar{\alpha}$ such that $P(\{\bar{a}\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. Hence, $P_{it}(Z|b^\alpha;p^\alpha) \geq P_{it}(\{\sigma_{it}(\bar{a}_{<t})\}|b^\alpha;p^\alpha) \geq P(\{\bar{a}\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. Q.E.D.

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.*

Proof. We begin by showing that $(a) \Rightarrow (b)$. Fix any $\varepsilon > 0$. If b is a sequential equilibrium strategy profile, then there is a sequence of completely mixed strategies $b^n \to b$ (e.g., any sequence that generates the system of beliefs μ associated with b in the sequential equilibrium (b,μ) of which b is a part) such that, for all n large enough, b^n is an ε-best reply for its owner conditional on any information set in the game $\Gamma(p)$. Since nature's probability function p in a standard finite game always gives every state positive probability, the sequence $\{(b^n, p)\}$ is admissible for (b,p). Hence, b is a perfect conditional ε-equilibrium.

The implication $(b) \Rightarrow (c)$ is trivial since, if (b) holds then the constant sequence $b, b, ...$ shows that b is the limit as $\varepsilon \to 0$ of a sequence of perfect conditional ε-equilibria.

It remains only to show that $(c) \Rightarrow (a)$. Let $\{b^{\varepsilon}\}$ be a sequence of perfect conditional ε-equilibria that converges to b as $\varepsilon \to 0$. Consequently, for every ε there is an ε-test net $\{(b^{\varepsilon,\alpha}, p^{\varepsilon,\alpha})\}$ for (b^{ε}, p). In particular, $b^{\varepsilon,\alpha}$ is a conditional ε-equilibrium of $\Gamma(p^{\varepsilon,\alpha})$ for every ε and α. For any ε, because nature's finitely many states all have positive probability under p, there is for every $\eta > 0$, an index α_{η} such that $b^{\varepsilon,\alpha}$ is strictly mixed, is within η of b (i.e., $||b^{\varepsilon,\alpha} - b|| < \eta$), and is a conditional $(\varepsilon + \eta)$-equilibrium in the original game $\Gamma(p)$. Consequently, for every $\varepsilon' > 0$, if we choose ε and η so that $\varepsilon + \eta < \varepsilon'$, then $b' = b^{\varepsilon,\alpha_{\eta}}$ is strictly mixed, is within η of b, and is a conditional ε'-equilibrium of $\Gamma(p)$. Letting μ' be the system of beliefs induced by the conditionals of b', and letting μ be the limit (along a subsequence if necessary) of μ' as $\varepsilon' \to 0$, we conclude that (b, μ) is a sequential equilibrium of $\Gamma(p)$. Q.E.D.

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

Proof. By hypothesis, we may choose for each $\varepsilon > 0$ a perfect conditional ε-equilibrium b^{ε}. For any $\varepsilon > 0$, $P(C|b^{\varepsilon})$ is defined for every $C \in \mathcal{M}(A)$. Consequently, $\{P(\cdot|b^{\varepsilon})\}$ is a net in $[0,1]^{\mathcal{M}(A)}$, with smaller positive numbers ε being further out in the (directed) index set. By Tychonoff's theorem, $[0,1]^{\mathcal{M}(A)}$ is compact in the product topology and so there exists $\mu \in [0,1]^{\mathcal{M}(A)}$ and a subnet $\{P(\cdot|b^{\varepsilon_{\alpha}})\}$ that converges (in the product topology) to μ, meaning precisely that (6.1) in MR holds for the subnet and so μ is a perfect conditional equilibrium distribution. Q.E.D.

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

Proof. Suppose that b is a perfect conditional ε-equilibrium. Then there is an ε-test net $\{(b^{\alpha}, p^{\alpha})\}$ for (b, p).

Consider any $it \in L$ and any $Z \in \mathcal{M}(S_{it})$ such that $P_{it}(Z|b;p) > 0$. Because $\lim_{\alpha} \|b^{\alpha} - b\| = \lim_{\alpha} \|p^{\alpha} - p\| = 0$, there is an index $\bar{\alpha}$ such that $P_{it}(Z|b^{\alpha}; p^{\alpha}) > 0$ for every $\alpha \geq \bar{\alpha}$. Since b^{α} is a conditional ε-equilibrium of $\Gamma(p^{\alpha})$,

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

for every date-t continuation c_i of b^{α}_i and for every $\alpha \geq \bar{\alpha}$.

Since any date-t continuation c_i of b^{α}_i can be written as $(b^{\alpha}_{i,<t}, c_{i,\geq t})$, we have that for

every $\alpha \geq \bar{\alpha}$, and for every $c_{i,\geq t} \in B_{i,\geq t}$,

$$U_i((b^\alpha_{i,<t}, c_{i,\geq t}), b^\alpha_{-i}|Z; p^\alpha) \leq U_i(b^\alpha|Z; p^\alpha) + \varepsilon. \tag{1.3}$$

Because $\lim_\alpha \|b^\alpha - b\| = \lim_\alpha \|p^\alpha - p\| = 0$, we have $\lim_\alpha P_{it}(Z|b^\alpha; p^\alpha) = P_{it}(Z|b; p)$, and we have, for every $C \in \mathcal{M}(A)$, $\lim_\alpha P(C|b^\alpha; p^\alpha) = P(C|b; p)$. Consequently, because $P_{it}(Z|b; p) > 0$, the definition of conditional probabilities in Section 3.2 in MR implies that, for every $C \in \mathcal{M}(A)$,

$$\begin{aligned} \lim_\alpha P(C|Z, b^\alpha; p^\alpha) &= \lim_\alpha \frac{P(\{a \in C : \sigma_{it}(a_{<t}) \in Z|b^\alpha)}{P_{it}(Z|b^\alpha; p^\alpha)} \\ &= \frac{P(\{a \in C : \sigma_{it}(a_{<t}) \in Z|b)}{P_{it}(Z|b; p)} \\ &= P(C|Z, b; p). \end{aligned}$$

Therefore, from the definition of conditional expected payoffs in Section 3.2 in MR,

$$\begin{aligned} U_i(b^\alpha|Z; p^\alpha) &= \int u_i(a) P(da|Z, b^\alpha; p^\alpha) \\ &\to_\alpha \int u_i(a) P(da|Z, b; p) \\ &= U_i(b|Z; p). \end{aligned}$$

Similarly, for any $c_{i,\geq t} \in B_{i,\geq t}$, because $\lim_\alpha \left\|((b^\alpha_{i,<t}, c_{i,\geq t}), b^\alpha_{-i}) - ((b_{i,<t}, c_{i,\geq t}), b_{-i})\right\| = 0$, we have $\lim_\alpha U_i((b^\alpha_{i,<t}, c_{i,\geq t}), b^\alpha_{-i}|Z; p^\alpha) = U_i((b_{i,<t}, c_{i,\geq t}), b_{-i}|Z; p)$. Consequently, by (1.3), we may conclude that, for every $c_{i,\geq t} \in B_{i,\geq t}$,

$$U_i((b_{i,<t}, c_{i,\geq t}), b_{-i}|Z; p) \leq U_i(b|Z; p) + \varepsilon.$$

Since Z is an arbitrary signal event that has positive probability under b in $\Gamma(p)$, we may conclude that b is a conditional ε-equilibrium.

That b is also an ε-Nash equilibrium follows from the fact that the null signal, $\emptyset$, which all players observe at date 1, always has positive probability. Thus, by the property of a conditional ε-equilibrium, no player, given the others' strategies under b, can improve his payoff by more than ε conditional on the null history. Q.E.D.

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

Proof. Suppose that $b \in B$ is a perfect conditional ε-equilibrium and that $\{(b^\alpha, p^\alpha)\}$ is an ε-test net for (b, p). Then, $\{(b^\alpha, p^\alpha)\}$ is admissible for (b, p) and so by Theorem 6.3 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 $\alpha \geq \bar{\alpha}$.

By the definition of a subgame perfect ε-equilibrium, it suffices to show that for any $a \in A \backslash N$ and for any date t, if $a_{<t} \in A_{<t}$ is a subgame, then

$$\sup_{c_i \in B_i} U_i(c_i, b_{-i}|a_{<t}) \leq U_i(b|a_{<t}) + \varepsilon, \text{ for every player } i \in I. \tag{1.4}$$

So suppose that a is in $A \backslash N$ and that $a_{<t} \in A_{<t}$ is a date-t subgame.

By the properties of the set $A \backslash N$, there is an index $\bar{\alpha}$ such that $P(\{a\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$. In particular, for every player $i \in I$, $P_{it}(\{\sigma_{it}(a_{<t})\}|b^\alpha;p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$.

Because $\sigma_{it}^{-1}(\sigma_{it}(a_{<t})) = \{a_{<t}\}$ and because each b^α is a conditional ε-equilibrium of $\Gamma(p^\alpha)$, we may conclude that for every $\alpha \geq \bar{\alpha}$, for every $i \in I$, and for every $c_i \in B_i$ (the specification of c_i on dates before t is irrelevant because $a_{<t}$ is given),

$$U_i(c_i, b_{-i}^\alpha|a_{<t};p^\alpha) \leq U_i(b^\alpha|a_{<t};p^\alpha) + \varepsilon.$$

Taking the limit of this inequality with respect to α gives (1.4) because $\lim_\alpha \|b^\alpha - b\| = \lim_\alpha \|p^\alpha - p\| = 0$ imply,

$$\lim_\alpha U_i(c_i, b_{-i}^\alpha|a_{<t};p^\alpha) = U_i(c_i, b_{-i}|a_{<t};p),$$

and

$$\lim_\alpha U_i(b^\alpha|a_{<t};p^\alpha) = U_i(b|a_{<t};p). \text{ Q.E.D.}$$

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.*

Proof. Suppose that $b \in B$ is a perfect conditional ε-equilibrium and that $\{(b^\alpha, p^\alpha)\}$ is an ε-test net for (b, p). Then, $\{(P_{<t}(C|Z, b^\alpha;p^\alpha))_{it \in L, \text{ observable } Z \in \mathcal{M}(S_{it}), C \in \mathcal{M}(A_{<t})}\}_\alpha$ is a net taking values in a space that is an infinite product of the compact set $[0,1]$. (The Corollary to Theorem 6.3 above implies that each $P_{<t}(C|Z, b^\alpha;p^\alpha)$ in this net is well-defined for all large enough indices α.) By Tychonoff's theorem this space is compact and so we may assume, without loss of generality, that the net $\{(P_{<t}(C|Z, b^\alpha;p^\alpha))_{it \in L, Z \in \mathcal{M}(S_{it}), C \in \mathcal{M}(A_{<t})}\}_\alpha$ converges. Therefore, we may define beliefs β as in MR equation (6.4). That is, for every $it \in L$, for every observable $Z \in \mathcal{M}(S_{it})$, and for every $C \in \mathcal{M}(A_{<t})$, we may define

$$\beta_{it}(C|Z) = \lim_\alpha P_{<t}(C|Z, b^\alpha;p^\alpha). \tag{1.5}$$

Consequently, β is finitely consistent for b. It remains only to show that (b, β) is sequentially

ε-rational.

Fix any $it \in L$, any observable $Z \in \mathcal{M}(S_{it})$, and any $c_i \in B_i$. We must show that,

$$\int U_i(c_i, b_{-i}|a_{<t})\beta_{it}(da_{<t}|Z) \leq \int U_i(b|a_{<t})\beta_{it}(da_{<t}|Z) + \varepsilon. \tag{1.6}$$

For each index α, define $c_i^\alpha = (b_{i,<t}^\alpha, c_{i,\geq t})$. Then c_i^α is a date-t continuation of b_i^α. Therefore, because each b^α is a conditional ε-equilibrium of $\Gamma(p^\alpha)$, and because, by the Corollary to Theorem 6.3, there is $\bar{\alpha}$ such that $P_{it}(Z|b^\alpha; p^\alpha) > 0$ for every $\alpha \geq \bar{\alpha}$, we have that,

$$U_i(c_i^\alpha, b_{-i}^\alpha|Z; p^\alpha) \leq U_i(b^\alpha|Z; p^\alpha) + \varepsilon. \tag{1.7}$$

By the definitions in Sections 3.1 and 3.2 in MR,[2]

$$U_i(c_i^\alpha, b_{-i}^\alpha|Z; p^\alpha) = \int U_i(c_i^\alpha, b_{-i}^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha)$$

and

$$U_i(b^\alpha|Z; p^\alpha) = \int U_i(b^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha).$$

Consequently, (1.7) implies that for every $\alpha \geq \bar{\alpha}$,

$$\int U_i(c_i^\alpha, b_{-i}^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) \leq \int U_i(b^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) + \varepsilon.$$

Moreover, since $U_i(c_i^\alpha, b_{-i}^\alpha|a_{<t})$ does not depend on the players' strategies on dates before t (since $a_{<t}$ is given), we may replace $c_i^\alpha = (b_{i,<t}^\alpha, c_{i,\geq t})$ with c_i without changing the value of $U_i(c_i^\alpha, b_{-i}^\alpha|a_{<t})$. That is, $U_i(c_i^\alpha, b_{-i}^\alpha|a_{<t}) = U_i(c_i, b_{-i}^\alpha|a_{<t})$ for every α. Hence, for every $\alpha \geq \bar{\alpha}$,

$$\int U_i(c_i, b_{-i}^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) \leq \int U_i(b^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) + \varepsilon. \tag{1.8}$$

Let us consider the limit of the term on the left-hand side of (1.8).

$$\begin{aligned}\lim_\alpha \int U_i(c_i, b_{-i}^\alpha|a_{<t})P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) \;&=\; \lim_\alpha \int \Big[U_i(c_i, b_{-i}^\alpha|a_{<t}) - U_i(c_i, b_{-i}|a_{<t}) \\ &\qquad + U_i(c_i, b_{-i}|a_{<t})\Big]P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha)\end{aligned}$$

[2] We are using the fact that, for any $it \in L$ and for any $\tilde{b} \in B$,

$$\int u_i(a)P(da|\tilde{b}) = \int_{A_{<t}} \left(\int_{A_{\geq t}} u_i(a_{<t}, a_{\geq t})P_{\geq t}(da_{\geq t}|a_{<t}, \tilde{b}) \right) P_{<t}(da_{<t}|\tilde{b}),$$

which, from the definitions in MR Section 3.1, can be established for each date t by induction, starting from date $t = 1$. We leave this induction proof to the reader.

$$
\begin{aligned}
&= \lim_\alpha \int \left(U_i(c_i, b^\alpha_{-i}|a_{<t}) - U_i(c_i, b_{-i}|a_{<t})\right) P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) \\
&\quad + \lim_\alpha \int U_i(c_i, b_{-i}|a_{<t}) P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) \\
&= 0 + \int U_i(c_i, b_{-i}|a_{<t}) \beta_{it}(da_{<t}|Z), \qquad (1.9)
\end{aligned}
$$

where the final equality follows because (first term) $\lim_\alpha ||b^\alpha - b|| = 0$ and (second term) by (1.5), and therefore the second equality follows because the two separate limits exist.

Similarly, $\lim_\alpha \int U_i(b^\alpha|a_{<t}) P_{<t}(da_{<t}|Z, b^\alpha; p^\alpha) = \int U_i(b|a_{<t}) \beta_{it}(da_{<t}|Z)$. Combined with (1.5) and (1.8), we obtain (1.6), as desired. Q.E.D.

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

Proof. By the canonical net construction in MR Section 8.1, for any date t and for any history before date t, each coordinate $j \in J$ of the provisional state that is chosen on that date by p is independently left unperturbed by $\phi^{n,F}_{tj}$ with probability at least $1 - \frac{1}{n}$. Hence, the provisional state chosen by p on any date and after any history is left entirely unperturbed with probability at least $\left(1 - \frac{1}{n}\right)^{\#J}$. Consequently, for each index $(n, F) \in \Omega$, $\left\|p^{n,F} - p\right\| \leq 1 - \left(1 - \frac{1}{n}\right)^{\#J}$ and so $\lim_{(n,F)} \left\|p^{n,F} - p\right\| = 0$ because $\lim_{(n,F)} n = +\infty$.

Let $Q_{0t} = \{\times_{j\in J} q_j : q_j \in Q_{0tj}, \forall j \in J\}$, where each Q_{0tj} is the countable (i.e., finite or countably infinite) partition of A_{0tj} that is specified in MR Section 8.1. Then Q_{0t} is a countable partition of A_{0t} into measurable (product) sets.

Fix any date $t \leq T$ for the remainder of the proof. Since each A_{0tj} is a separable metric space, so too is A_{0t} with the coordinate-wise maximum distance. So A_{0t} has a countable basis $\mathcal{U}$ of open sets. Let D^t be the union of all sets of the form $(U \cap q) \times \{a_{<t} \in A_{<t} : p_t(U \cap q|a_{<t}) = 0\}$, where U can be any element of the countable basis $\mathcal{U}$ and q can be any element of the countable partition Q_{0t}. Then $D^t \subseteq A_{0t} \times A_{<t}$ is measurable because it is the countable union of measurable sets (since p_t is a transition probability). Let $C^t = (A_{0t} \times A_{<t}) \backslash D^t$. Then $C^t \subseteq A_{0t} \times A_{<t}$ is measurable and $C^t = \{(a_{0t}, a_{<t}) : a_{0t}$ 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}. Hence, letting $C^t_{a_{<t}} = \{a_{0t} \in A_{0t} : (a_{0t}, a_{<t}) \in C^t\}$, $C^t_{a_{<t}}$ is measurable and $p_t(C^t_{a_{<t}}|a_{<t}) = 1$.

Fix any $a_{<t} \in A_{<t}$ and any $a_{0t} \in C^t_{a_{<t}}$ for the remainder of the proof. The proof will be complete if we can show that there is an index $(\bar{n}, \bar{F}) \in \Omega$ such that $p^{n,F}_t(\{a_{0t}\}|a_{<t}) > 0$ for every $(n, F) \in \Omega$ such that $n \geq \bar{n}$ and $F \supseteq \bar{F}$.

Choose an index $(\bar{n}, \bar{F})$ such that $a_{0t} \in \bar{F}$, and let (n, F) be any index such that $n \geq \bar{n}$ and $F \supseteq \bar{F}$ (and so $a_{0t} \in F$). We must show that $p^{n,F}_t(\{a_{0t}\}|a_{<t}) > 0$, where we recall from

MR Section 8.1 that,

$$p_t^{n,F}(\{a_{0t}\}|a_{<t}) = \int \Pi_{j\in J}\phi_{tj}^{n,F}(\{a_{0tj}\}|a'_{0tj})p_t(da'_{0t}|a_{<t}). \tag{1.10}$$

Let V be an open set in $\mathcal{U}$ that contains a_{0t} such that for every $a'_{0t} \in V$ and for every $j \in J$, a'_{0tj} is within distance $1/n$ of a_{0tj}. Therefore, since $a_{0t} \in F \cap Q_{0t}(a_{0t})$, if $a'_{0t} \in V \cap Q_{0t}(a_{0t})$ then for every $j \in J$, a_{0tj} is within distance $1/n$ of a'_{0tj} and both points are in the same element, $Q_{0tj}(a_{0tj})$, of the partition Q_{0tj}. So by the definition of the $\phi_{tj}^{n,F}$ mappings in MR Section 8.1, $\Pi_{j\in J}\phi_{tj}^{n,F}(\{a_{0tj}\}|a'_{0tj}) > 0$ for every $a'_{0t} \in V \cap Q_{0t}(a_{0t})$. Also, $p_t(V \cap Q_{0t}(a_{0t})|a_{<t}) > 0$ because the open set V contains a_{0t} and a_{0t} is in the support of $p_t(\cdot \cap Q_{0t}(a_{0t})|a_{<t})$ (because $a_{0t} \in C^t_{a_{<t}}$). Hence, for this a_{0t}, the nonnegative function, $\Pi_{j\in J}\phi_{tj}^{n,F}(\{a_{0tj}\}|a'_{0tj})$ of a'_{0t}, is strictly positive for every a'_{0t} in $V \cap Q_{0t}(a_{0t})$. Therefore, since $p_t(V \cap Q_{0t}(a_{0t})|a_{<t}) > 0$, (1.10) implies that $p_t^{n,F}(\{a_{0t}\}|a_{<t}) > 0$. Q.E.D.

1.2 Examples

Our first example fulfills a promise in MR Section 6.4 to provide an example of a strategy profile b and belief system β that is finitely-consistent and sequentially ε-rational even though b is not a perfect conditional ε-equilibrium. The example shows that, even with finitely consistent beliefs, sequentially ε-rational behavior can be unintuitive because beliefs are only finitely additive.

Example 1. *Problems with sequential rationality and finite consistency.*

- On date 1, nature chooses θ uniformly from the open interval $(0, 1)$, where neither $\theta = 0$ nor $\theta = 1$ is a possible state of nature.
- On date 2, player 1 observes θ and then chooses $x \in \{0, 1\}$.
- On date 3, player 2 observes x and then chooses $y \in [0, 1)$.
- Payoffs are as follows:

  If $x = 0$ ("out"), then $u_1 = u_2 = 0$.

  If $x = 1$ ("in") and $y = 0$ ("out"), then $u_1 = -1$ and $u_2 = 0$.

In any perfect conditional ε-equilibrium of this game with $\varepsilon < 1$, player 1 must choose $x = 1$ with probability at least $1 - \varepsilon/(1 - \varepsilon)$, and player 2's strategy, conditional on the signal $x = 1$, must give the event $\{y \geq \theta\}$ probability at least $1 - \varepsilon/2$. This is because, in any perturbation that gives the signal $x = 1$ positive probability (even if that perturbation involves nature), player 2 can obtain a conditional expected payoff arbitrarily close to 1 by choosing y sufficiently close to 1 (this latter fact is a consequence of the countable additivity of the conditional probability measures defined by the perturbations). Consequently, by choosing $x = 1$, player 1 can obtain a payoff close to 1 when $\varepsilon > 0$ is close to zero.

However, consider the strategy profile b in which player 1 chooses $x = 0$ and player 2 chooses $y = 0$ no matter what signal he observes. In addition, consider the beliefs β_{23} for player 2 such that $\beta_{23}(\cdot|x = 0)$ is uniform on $(0, 1)$ and $\beta_{23}((1 - \delta, 1)|x = 1) = 1$ for every $\delta > 0$. Then $\beta_{23}(\cdot|x = 1)$ is only finitely additive. The strategy profile b gives both players a payoff of zero, contrary to every perfect conditional ε-equilibrium when $\varepsilon > 0$ is small. Nevertheless, (b, β) is finitely consistent and sequentially 0-rational.[3] To see that this (b, β) is sequentially 0-rational, note first that, given 2's behavior, it is optimal for player 1 to choose $x = 0$ no matter what θ he observes, and note second that, because for any $\delta > 0$ player 2's beliefs after observing $x = 1$ put probability 1 on the event that nature's θ is greater than $1 - \delta$, player 2's expected payoff from choosing any $y > 0$ is -1. Hence, player 2's unique optimal choice is $y = 0$ after $x = 1$. And any choice of y is optimal for player 2 after $x = 0$.

The next two examples are taken from MR. In each of them, the index set for the nets that are constructed is the set of all pairs (n, F) such that n is any positive integer and F is any nonempty finite subset of $[0, 1]$. Larger positive integers n and more inclusive finite subsets F of $[0, 1]$ correspond to larger indices.

MR Example 2.1. *Problems of spurious signaling in naïve finite approximations.*

For this example, let us first show that player 1's expected payoff must converge to zero as $\varepsilon \to 0$ in any sequence of perfect conditional ε-equilibria. So consider any such sequence. For each state $\theta = 1, 2$, define q_θ to be the limiting probability that player 2 chooses $y = 1$ conditional on state θ (extract a convergent subsequence if necessary). Since player 1's payoff in this game is nonzero if and only if player 2 chooses $y = 1$, we must show that $q_1 = q_2 = 0$. Player 2 can obtain an expected payoff of 3/4 by always choosing $y = 2$, and so 2's limit equilibrium payoff, $(1/4)(q_1)1 + (3/4)(1 - q_2)1$, must be at least $3/4$,[4] which means that $q_1 \geq 3q_2$. Player 1's limit equilibrium payoff, $(1/4)(q_1)1 + (3/4)(q_2)1$, cannot be less than the limit of the payoffs that he would achieve if, along the sequence, he always deviated to the square root of the action that his equilibrium strategy dictated. Since such deviations along the sequence would yield player 1 a limit payoff of at least $(3/4)q_1$,[5] we must have

[3] Finite consistency can be verified by taking the limit as $\delta \to 0$ of joint perturbations that, for nature, put positive probability on $\theta = 1 - \delta$ and, for player 1, put positive probability on $x = 1$ after observing $\theta = 1 - \delta$.

[4] Otherwise, for some $\varepsilon > 0$ small enough along the sequence, player 2 would not be ε-optimizing conditional on the null state at date 1.

[5] Because, under player 1's square root deviation, when the state is $\theta = 2$ player 2 chooses $y = 1$ with limiting probability q_1, and when the state is $\theta = 1$ player 1's payoff is nonnegative no matter what action player 2 chooses.

$(1/4)(q_1)1 + (3/4)(q_2)1 \geq (3/4)q_1$, which is equivalent to $3q_2 \geq 2q_1$. But because $q_1 \geq 3q_2$, this means that $q_1 = q_2 = 0$ as desired.

Next we define a strategy profile, b^*, that is a perfect conditional ε-equilibrium for every $\varepsilon > 0$. Let player 1's strategy, b_1^*, choose action $x = 0$ with probability 1, and let player 2's strategy b_2^* choose action $y = 2$ regardless of the signal that she observes. Construct an ε-test net $\{(b^{n,F}, p^{n,F})\}$ for (b^*, p) as follows. For any nonempty finite subset F of $[0,1]$, let F' be the smallest set containing F that is closed under the taking of square roots (i.e., $F' = \cup_{x \in F}\{x, \sqrt{x}, \sqrt{\sqrt{x}}, ...\}$). For any positive integer n and for any finite subset F of $[0,1]$, define $b_1^{n,F} \in B_1$ so that $b_1^{n,F}(\{0\}) = 1 - 1/n$, $b_1^{n,F}(F') = 1/n$, and $b_1^{n,F}(x) = 3b_1^{n,F}(\sqrt{x})$ for every $x \in F'$. Define $b_2^{n,F} \in B_2$ so that for every $s \in [0,1]$, $b_2^{n,F}(\{2\}|s) = 1 - 1/n$ and $b_2^{n,F}(\{1\}|s) = 1/n$. Finally, define $p^{n,F} = p$. Then, $\lim_{n,F} \left\| b^{n,F} - b^* \right\| = \lim_{n,F} \left\| p^{n,F} - p \right\| = 0$, and every outcome in the game is, eventually in the net $\{(b^{n,F}, p^{n,F})\}$, given positive probability. Hence, $\{(b^{n,F}, p^{n,F})\}$ is admissible for (b^*, p). In the perturbed game $\Gamma(p^{n,F})$ ($= \Gamma$ since $p^{n,F} = p$), $b^{n,F}$ gives probability 1 to the set of player 2 signals $F' \cup \{0\}$. Since conditional on any positive signal in F' that is less than 1, the states $\theta = 1$ and $\theta = 2$ are equally likely under $b^{n,F}$ in $\Gamma(p^{n,F})$, it is conditionally optimal for player 2 to choose any $y \in \{1, 2\}$ after observing any such signal. And since the signals $s = 0$ and $s = 1$ are uninformative about nature's state, it is conditionally optimal for player 2 to choose $y = 2$ after observing $s = 0$ or 1. So $b_2^{n,F}$ is a $1/n$-best reply for player 2 conditional on every positive probability signal event. Since player 1's payoff in $\Gamma(p^{n,F})$ under $b^{n,F}$ is $1/n$ no matter what action he chooses, player 1 is fully optimizing. Hence, $b^{n,F}$ is a conditional $1/n$-equilibrium of $\Gamma(p^{n,F})$, and so, for every $\varepsilon > 0$, b^* is a perfect conditional ε-equilibrium of Γ.[6]

MR Example 6.1. *Why nature must be perturbed to test rational behavior with positive probability in all events.*

In this example it is obvious that player 1 must put probability at least $1 - \varepsilon$ on his strictly dominant strategy $x = -1$ in any perfect conditional ε-equilibrium. What we wish to show here is that perfect conditional ε-equilibrium solves the existence problem that was presented in the text by allowing perturbations of nature's probability function in addition to perturbations of the players' strategies. For this, it is enough to define a strategy profile, b^*, and to show that it is a perfect conditional ε-equilibrium for every $\varepsilon > 0$. Define $b^* \in B$ so that $b_1^*(\{-1\}) = 1$ and $b_2^*(\{-1\}|s) = 1$ for every $s \in [0,1]$. Define an ε-test net $\{(b^\alpha, p^\alpha)\}$ for (b^*, p) as follows. For every positive integer n and for every nonempty finite subset F of $[0,1]$, define $b_1^{n,F} \in B_1$ so that $b_1^{n,F}(\{-1\}) = 1 - (1/n)^2$ and $b_1^{n,F}(\{x\}) = 1/(n^2(\#F))$ for every $x \in F$;

[6] There are many perfect conditional ε-equilibria of this game (e.g., player 1 chooses x uniformly from $[0,1]$ and player 2 always chooses $y = 2$ regardless of the signal that she observes).

define $b_2^{n,F} \in B_2$ so that $b_2^{n,F}(\{-1\}) = 1 - 1/n$ and $b_2^{n,F}(\{1\}) = 1/n$; and define $p^{n,F}$ so that for any interval $[c, d] \subseteq [0, 1]$, $p^{n,F}([c, d]\backslash F) = (1 - 1/n)(d - c)$ and $p^{n,F}(\{x\}) = 1/(n(\#F))$ for every $x \in F$. Then, $\lim_{n,F} \left\| b^{n,F} - b^* \right\| = \lim_{n,F} \left\| p^{n,F} - p \right\| = 0$, and every outcome in the game is, eventually in the net $\{(b^{n,F}, p^{n,F})\}$, given positive probability. Hence, $\{(b^{n,F}, p^{n,F})\}$ is admissible for (b^*, p). For any $\varepsilon > 0$, when n is sufficiently large and the players use $b^{n,F}$ in the game $\Gamma(p^{n,F})$, player 1 is clearly ε-optimizing and, conditional on any signal s that has positive probability (i.e., any signal $s \in F$), player 2 is also ε-optimizing because the event $\{x = -1 \text{ and } \theta = s\}$ is $n - 1$ times as likely as the event $\{x = s\}$. Also, conditional on any positive probability signal event in $[0, 1]\backslash F$, there is probability 1 that player 1 chose $x = -1$ and so player 2 is again ε-optimizing (for n sufficiently large). Hence, for any $\varepsilon > 0$, there is a large enough $\bar{n}$ such that $b^{n,F}$ is a conditional ε-equilibrium of $\Gamma(p^{n,F})$ for all (n, F) with $n \geq \bar{n}$. Thus, b^* is a perfect conditional ε-equilibrium.

1.3 Beyond Multi-Stage Games

The class of multi-stage games encompasses the vast majority of games considered by economists. Nevertheless, we show here that the definitions of conditional ε-equilibrium, full conditional ε-equilibrium, and perfect conditional ε-equilibrium, all easily extend to extensive form games with perfect recall outside of this class.

Let Γ be any finite-player extensive form game with perfect recall. We list below only a subset of all of the items needed to fully specify such a game Γ. The subset that we specify is just enough so that our equilibrium concepts can be defined. A full specification of Γ would have to list even more items (e.g., signal maps would have to be specified, conditions ruling out cycles in the game tree would have to be specified, etc.) So assume that Γ satisfies at least the following conditions.

1. I is the finite set of players, $0 \notin I$. Let $I^* = I \cup \{0\}$, where 0 denotes nature.

2. For $i \in I$ and for any integer $t \geq 1$, S_{it} is the set of all possible signals that player i can observe in the game after having already observed precisely $t-1$ previous signals in the game;[7] A_{it} is a nonempty action set for player i; for each $s_{it} \in S_{it}$, $\Phi_{it}(s_{it}) \subseteq A_{it}$ is the set of feasible actions for player i after observing the signal s_{it}; and A_{0t} is a nonempty set of states of nature.

3. A is set of all possible outcomes (paths of play) in the game. Each $a \in A$ is an infinite

[7] So S_{i1} is the set of all possible first signals that player i can ever obersve, S_{i2} is the set of all possible second signals that player i can ever observe, etc. If the game has finite length, then there is T such that $S_{it} = \emptyset$ for every $t > T$. Perfect recall ensures that each player's set of signals can be serially partitioned in this way.

sequence, $a = (a_1, a_2, ...)$, [8] where, for each point a_r in the sequence, there is a nonempty subset $J \subseteq I^*$ such that $a_r \in \times_{i \in J} \left(\cup_{t=1}^{r} A_{it}\right)$ indicating that all of the players in J, and also nature if $0 \in J$, moved simultaneously.[9] For any $a = (a_1, a_2, ...) \in A$, and for any integer $t > 1$, $a_{<t}$ is the subsequence $(a_r)_{r<t}$; let $A_{<t} = \{a_{<t} : a \in A\}$. $(A_{<1} = \emptyset)$.

4. All of the sets A_{it}, S_{it}, $A_{<t}$, and A, are equipped with sigma algebras of measurable subsets.

5. For any player $i \in I$, for any integer $t \geq 1$, and for any measurable subset Z of S_{it}, the set of paths of play in A along which player i observes a signal in Z is a measurable subset of A and is denoted by $[Z] \subseteq A$.

6. For any integer $t \geq 1$, $p_t : A_{<t} \to \Delta(A_{0t})$ is a transition probability specifying nature's distribution over states in A_{0t} after histories in $A_{<t}$.[10] Let $p = (p_t)_{t \geq 1}$ denote nature's probability function.

7. $u_i : A \to \mathbb{R}$ is player i's bounded and measurable payoff function.

A *strategy* $b_i = (b_{it})_{t \geq 1}$ for player i specifies, for each $t \geq 1$, a transition probability $b_{it} : S_{it} \to \Delta(A_{it})$ such that $b_{it}(\Phi_{it}(s_{it})|s_{it}) = 1$ for every $s_{it} \in S_{it}$.

Given nature's probability function $p = (p_t)_{t \geq 1}$, we assume that any strategy profile $b = (b_i)_{i \in I}$, induces a unique distribution over the set, A, of outcomes of the game, and we denote this outcome distribution by $P(\cdot|b;p) \in \Delta(A)$. In fact, we assume such a distribution over A is well-defined for any strategy profile and for any probability function for nature that satisfies condition 6.

For any strategy profile b, for any player i, for any integer $t \geq 1$, and for any measurable subset Z of S_{it}, define $P_{it}(Z|b;p) = P([Z]|b;p)$, and, if $P_{it}(Z|b;p) > 0$, define $U_i(b|Z;p) = \int_{[Z]} u_i(a) P(da|b;p)/P_{it}(Z|b;p)$.

We can now define conditional ε-equilibrium precisely as in Definition 4.1 of MR. Then we can define full conditional ε-equilibrium precisely as in Section 5 of MR,[11] and we can define perfect conditional ε-equilibrium precisely as in Definition 6.2 of MR.

[8] Paths of play that are finite in length can always be represented by infinite sequences that are eventually constant.

[9] We include simultaneous moves in this way so as to include the class of multi-stage games as defined in MR, where all players and nature move simultaneously at each one of T dates of play.

[10] As in MR, for any measurable space, X, $\Delta(X)$ is the set of countably additive probability measures on the measurable subsets of X.

[11] As in Section 5, we should assume that $\Phi_{it}(s_{it}) = A_{it}$. We can relax this assumption by defining b_{it} to have *full support* iff $b_{it}(C|s_{it}) > 0$ for every open subset C of A_{it} such that $C \cap \Phi_{it}(s_{it})$ is nonempty, and for every $s_{it} \in S_{it}$.

References

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