

Conditional equilibria of multi-stage games with infinite sets of signals and actions

by

Roger B. Myerson and Philip J. Reny

Department of Economics, University of Chicago

Abstract: We develop concepts of conditional equilibria to extend Kreps and Wilson's concept of sequential equilibrium to games where the sets of actions that players can choose and the sets of signals that players may observe are infinite. A strategy profile is a conditional ε -equilibrium if, given any signal event that a player could observe with positive probability, the player's conditionally expected utility would be within ε of the best that the player could achieve by deviating. With topologies on actions, if a conditional ε -equilibrium has full support, then the perfect-perturbation tests will not be necessary to evaluate ε -rationality of responses to a dense class of deviations. Perfect conditional ε -equilibria are defined by testing conditional ε -rationality also under nets of small perturbations that can make any finite collection of outcomes have positive probability. For a large class of projective games, we prove existence of full perfect conditional ε -equilibria.

http://home.uchicago.edu/~rmyerson/research/seqm_notes.pdf

Goal: We aimed to extend a definition of *sequential equilibrium* to multi-stage games with infinite signal sets and infinite action sets, prove existence for a large class of games.

Unsuccessful idea: try taking limits of sequential equilibria of finite games that "converge" to an infinite game; they often have limits of equilibria that seem wrong.

Instead we look at limits of strategy profiles that satisfy, for each player at each date, a condition of approximate optimality among all the player's possible continuation strategies.

Nash equilibrium does not test rationality of strategies in events with probability 0, which can be problematic if a deviation could give positive probability to these untested events.

For finite games, sequential equilibrium strategies are *weakly perfect* limits of totally mixed strategy profiles that satisfy conditional ε -rationality for each player at every information set.

We define *conditional ε -equilibria* by ε -rationality in all positive-probability events.

But with uncountably infinite signal spaces, it may be impossible to give positive probability to all observable events at once, and so we cannot test rationality in all events at once.

With topologies on players' actions, we can ask that such perturbations should not be needed to test rationality under all possible consequences of a dense set of deviations (*fullness*).

We can ask that, for any finite set of observable events, small perturbations of nature and players' strategies should be able to verify strategic rationality in these events (*perfectness*).

But even small perturbations of nature can change a game in counter-intuitive ways.

We also consider restrictions on the set of admissible perturbations of nature.

Existence of full perfect conditional ε -equilibria is shown for a wide class of games.

We emphasize ε -equilibria because $\varepsilon \rightarrow 0$ limits may violate strategic independence, as players can be coordinated by infinitesimal details of signals (*strategic entanglement*).

Multi-stage games $\Gamma = (I, T, A, S, \mathcal{M}, \sigma, p, u)$.

$i \in I = \{\textit{players}\}$, a finite set, $0 \notin I$. Let $I^* = I \cup \{0\}$, where 0 denotes nature (chance).

$t \in \{1, \dots, T\} = \{\textit{dates of play}\}$. $L = I \times \{1, \dots, T\} = \{\textit{dated players}\}$. We write "it" for (i,t).

$S_{it} = \{\text{possible } \textit{signals} \text{ for player } i \text{ at } t\} \quad \forall it \in L$.

$A_{it} = \{i\text{'s possible } \textit{actions} \text{ at } t\} \quad \forall it \in L$. $A_{0t} = \{\text{nature's possible } \textit{states} \text{ at } t\}$.

Sigma-algebras of measurable sets (closed under countable \cap , & complements) are specified for A_{0t} , A_{it} , S_{it} . All one-point sets are measurable. Products have product sigma-algebras.

$A = \times_{t \leq T} \times_{i \in I^*} A_{it}$. $A = \{\text{possible } \textit{outcomes} \text{ of the game}\}$.

The subscript, $<t$, denotes the projection onto periods before t , and $\geq t$ weakly after.

e.g., $A_{<t} = \times_{r < t} \times_{i \in I^*} A_{ir} = \{\text{possible state-}\&\text{-action sequences before period } t\}$. $A_{<1} = \{\emptyset\}$.

For any $a \in A$, $a_{<t} = (a_{ir})_{i \in I^*, r < t}$ is the partial sequence of states and actions before period t .

For any of the sets X above, $\mathcal{M}(X)$ is its set of measurable subsets.

Let $\Delta(X)$ denote the set of countably additive probability measures on $\mathcal{M}(X)$.

The state at each date- t is determined by *nature's probability function* $p_t: A_{<t} \rightarrow \Delta(A_{0t})$, a transition probability. [$\forall C \in \mathcal{M}(A_{0t})$, $p_t(C|a_{<t})$ is a measurable function of $a_{<t}$.]

Player i 's date- t information is given by a measurable onto *signal function* $\sigma_{it}: A_{<t} \rightarrow S_{it}$.

Assume *perfect recall*: $\forall it \in L$, $\forall r < t$, there is a measurable $\Psi_{itr}: S_{it} \rightarrow S_{ir}$ and a measurable $\psi_{itr}: S_{it} \rightarrow A_{ir}$ such that $\Psi_{itr}(\sigma_{it}(a_{<t})) = \sigma_{ir}(a_{<r})$, $\psi_{itr}(\sigma_{it}(a_{<t})) = a_{ir}$, $\forall a \in A$.

Each player i has a measurable and bounded *utility function* $u_i: A \rightarrow \mathbb{R}$.

Behavioral strategies and induced distributions

A *strategy* b_{it} , for any $it \in L$, is any transition probability $b_{it}: S_{it} \rightarrow \Delta(A_{it})$, i.e.: $\forall s_{it} \in S_{it}$, $b_{it}(\cdot | s_{it})$ is a countably additive probability on the measurable subsets of A_{it} ; and, $\forall C \in \mathcal{M}(A_{it})$, $b_{it}(C | s_{it})$ is a measurable function of the signal s_{it} .

We let B_{it} be the set of strategies for i at t , and let $B = \times_{it \in L} B_{it}$ be the set of *strategy profiles*.

Also we may write $B_i = \times_{t \leq T} B_{it}$, $B_{.t} = \times_{i \in I} B_{it}$, $B_{<t} = \times_{r < t} B_{.r}$, $B_{\geq t} = \times_{r \geq t} B_{.r}$.

Given $b \in B$, we let b_{it} , b_i , $b_{.t}$, $b_{<t}$, or $b_{\geq t}$ denote the components of b in B_{it} , B_i , $B_{.t}$, $B_{<t}$, or $B_{\geq t}$.

For any $a_{<t}$ in $A_{<t}$ and $b_{.t}$ in $B_{.t}$, we let $P_t(\cdot | a_{<t}, b_{.t})$ be the measure in $\Delta(A_{.t})$ such that,

for any measurable product set $H = \times_{i \in I^*} H_{it}$, $P_t(H | a_{<t}, b_{.t}) = p_t(H_{0t} | a_{<t}) \prod_{i \in I} b_{it}(H_{it} | \sigma_{it}(a_{<t}))$.

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,

$\forall t \in \{1, \dots, T\}$, $\forall H \in \mathcal{M}(A_{<t+1})$, $P_{<t+1}(H | b) = \int P_t(\{a_{.t}: (a_{<t}, a_{.t}) \in H\} | a_{<t}, b_{.t}) P_{<t}(da_{<t} | b)$.

(We could write $P_{<t}(\cdot | b) = P_{<t}(\cdot | b_{<t})$, as it depends only on strategies b_{ir} for $r < t$.)

Let $P(\cdot | b) = P_{<T+1}(\cdot | b)$ be the probability distribution on outcomes in A induced by b .

We may also write $P(\cdot | b; p) = P(\cdot | b)$ to indicate the dependence on the probability function p .

For any b , we inductively define transition probabilities from $A_{<t}$ to $\Delta(A_{\geq t})$ so that

$P_{\geq T}(\cdot | a_{<T}, b) = P_T(\cdot | a_{<T}, b_{.T})$; and for any date $t < T$, and any measurable $H \subseteq A_{\geq t}$,

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

(We could write $P_{\geq t}(\cdot | a_{<t}, b) = P_{\geq t}(\cdot | a_{<t}, b_{\geq t})$, as it depends only on strategies b_{ir} for $r \geq t$.)

At any time 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, $U_i(b | a_{<t})$ depends only on $b_{\geq t}$.)

Player i 's ex-ante expected utility is $U_i(b) = \int u_i(a) P(da | b) = \int U_i(b | a_{<t}) P_{<t}(da_{<t} | b)$.

Problems of spurious signaling in naïve finite approximations

Example 1: First, nature chooses $\tilde{a}_{01} = \theta \in \{1, 2\}$, with $p(1) = 1/4$ and $p(2) = 3/4$, while player 1 observes $s_{11} = \emptyset$ and chooses $a_{11} \in [0, 1]$.

Then player 2 observes $s_{22} = (a_{11})^\theta$, and then 2 chooses $a_{22} \in \{1, 2\}$.

Payoffs (u_1, u_2) are as follows: (a_{11} is payoff irrelevant)

	$a_{22} = 1$	$a_{22} = 2$
$\theta = 1$	(1,1)	(0,0)
$\theta = 2$	(1,0)	(0,1)

Consider subgame perfect equilibria of any finite approximate version of the game where player 1 chooses a_{11} in some \hat{A}_{11} that is a finite subset of $[0, 1]$ with at least one $0 < a_{11} < 1$. Then player 1 can get expected payoff at least $1/4$ by choosing the largest feasible $\bar{a}_{11} < 1$, as player 2 should choose $a_{22} = 1$ when $s_{22} = \bar{a}_{11} > (\bar{a}_{11})^2$ indicates $\theta = 1$.

Player 1's expected payoff cannot be more than $1/4$, as the smallest $\underline{a}_{11} > 0$ in his equilibrium support would lead player 2 to choose $a_{22} = 2$ when $s_{22} = (\underline{a}_{11})^2 < \underline{a}_{11}$ indicates $\theta = 2$.

But such a scenario cannot be even an approximate equilibrium of the real game, because player 1 could get an expected payoff at least $3/4$ by deviating to $(\bar{a}_{11})^{0.5}$.

Player 1 should get expected payoff 0 in "sequential" equilibria of the infinite game.

A full perfect conditional ε -equilibrium has $b_{22}(\{2\} | x) = 1 - \varepsilon \forall x \in [0, 1]$, $b_{11}(\cdot) = \text{Unif}[0, 1]$; perturbations give positive probability to countably many x with $\hat{b}_{11}(\{x^{0.5}\}) / \hat{b}_{11}(\{x\}) \geq 1/3$.

This example shows the importance of evaluating the optimality of each player's strategy relative to his entire set of strategies in the actual game (not just some approximating game).

Conditional ε -equilibria

$\forall s_{it} \in S_{it}$, let $\sigma^{-1}(s_{it}) = \{a_{<t} \in A_{<t} : \sigma_{it}(a_{<t}) = s_{it}\}$. $\forall Z \subseteq S_{it}$, let $\sigma^{-1}(Z) = \{a_{<t} \in A_{<t} : \sigma_{it}(a_{<t}) \in Z\}$.

For any strategy profile b in B , any player i , any date t , and any $Z \in \mathcal{M}(S_{it})$,

we let $P_{it}(Z|b) = P_{<t}(\sigma^{-1}(Z) | b) = P(\{a \in A : \sigma_{it}(a_{<t}) \in Z\} | b)$.

When $P_{it}(Z|b) > 0$, we may define the following conditional probabilities:

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

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

We can also define conditional expected utility, $\forall b \in B$: $U_i(b|Z) = \int_A u_i(a) P(da|Z, b)$.

For any player i at any date t , $c_i \in B_i$ is a *date- t continuation* of $b_i \in B_i$ iff $c_{ir} = b_{ir} \forall r < t$.

Let $\bar{B}_{i, \geq t}(b_i)$ denote the set of date- t continuations of b_i .

Let (b_{-i}, c_i) denote the strategy profile that differs from b in that player i deviates to c_i .

A strategy profile b is a *conditional ε -equilibrium* iff $\forall it \in L$, $\forall c_i \in \bar{B}_{i, \geq t}(b_i)$, $\forall Z \in \mathcal{M}(S_{it})$, if $P_{it}(Z|b) > 0$ then $U_i(c_i, b_{-i}|Z) \leq U_i(b|Z) + \varepsilon$.

Unfortunately, this condition does not test the rationality of player's strategic responses to signal events that have zero probability in equilibrium, and changing behavior in such zero-probability events could affect the optimality of strategies in positive-probability events.

Full conditional ε -equilibria

Suppose that, for each $i \in I$, the action set A_{it} has a given separable metric topology, and the measurable sets $\mathcal{M}(A_{it})$ are the Borel sets of this topology. With such topologies, we may be able to verify rationality of responses to a dense set of deviations without perturbing nature.

We say that strategy profile b has *full support* iff: $\forall i \in I, \forall s_{it} \in S_{it}, b_{it}(C|s_{it}) > 0$ for every set C that is a nonempty open subset of A_{it} .

Full-support strategies exist, by separability (\exists countable dense set).

If $\hat{b} \in B$ has full support then, $\forall b \in B, (1-\delta)b + \delta\hat{b}$ also has full support when $0 < \delta < 1$.

A conditional ε -equilibrium b is *full* iff it has full support (with the given topologies).

For any player i , given any strategy $c_i \in B_i$ and transition probabilities $\varphi_{it}: A_{it} \times S_{it} \rightarrow \Delta(A_{it})$, we let $c_i * \varphi_i$ denote the (" φ_i -trembling") strategy such that, for each time period t ,

$$\forall s_{it} \in S_{it}, \forall C \in \mathcal{M}(A_{it}), (c_i * \varphi_i)_t(C|s_{it}) = \int_{A_{it}} \varphi_{it}(C|a_{it}, s_{it}) c_{it}(da_{it}|s_{it}).$$

φ_i is a *δ -local tremble* iff $\varphi_{it}(\mathbb{B}_\delta(a_{it})|a_{it}, s_{it}) = 1, \forall a_{it} \in A_{it}, \forall s_{it} \in S_{it}, \forall t$. ($\mathbb{B}_\delta(a_{it}) = \delta$ -ball at a_{it} .)

Fact. If b has full support then, $\forall i \in I, \forall \delta > 0$, there exists a δ -local tremble φ_i such that $b_i * \varphi_i = b_i$, and $\forall c_i \in B_i, \{H \in \mathcal{M}(A): P(H|c_i * \varphi_i, b_{-i}) > 0\} \subseteq \{H \in \mathcal{M}(A): P(H|b) > 0\}$.

Furthermore, if b is a full conditional ε -equilibrium, and $P_{jt}(\{s_{jt}: \tilde{b}_{jt}(s_{jt}) \neq b_{jt}(s_{jt})\} | b) = 0 \forall jt \in I$, then $\forall it, \forall Z \in \mathcal{M}(S_{it})$ s.t. $P_{it}(Z|\tilde{b}) > 0, \forall c_i \in \bar{B}_{i, \geq t}(b_i), U_i(c_i * \varphi_i, \tilde{b}_{-i}|Z) \leq U_i(\tilde{b}|Z) + \varepsilon$.

(*Proof:* Here $\varphi_{it}(\cdot|a_{it}, s_{it})$ imitates $b_{it}(\cdot|s_{it})$ in a small neighborhood of a_{it} contained in $\mathbb{B}_\delta(a_{it})$.)

So for any \tilde{b} that differs from a full conditional ε -equilibrium only in 0-probability events, each player i has arbitrarily small local trembles φ_i such that i could not expect to gain more than ε from any φ_i -trembling deviation from \tilde{b} in any positive-probability event.

Construction of the above δ -local trembles φ_i from i 's full-support strategy b_i .

For each $t \in \{1, \dots, T\}$, A_{it} is assumed to be a separable metric space.

So we have a countable sequence of open balls with radius of radius $\delta/2$ that cover A_{it} .

Removing redundant balls if necessary, we can assume that each ball in the sequence is not contained in the finite union of the closures of the ball's predecessors in the sequence.

Now we construct a measurable partition $\{C_{itj}\}$ of A_{it} as follows.

The first set C_{it1} is the closure of the first ball in our sequence.

So $b_{it}(C_{it1}|s_{it}) > 0 \forall s_{it} \in S_{it}$, as C_{it1} is a nonempty open subset of A_{it} and b_{it} has full support.

For each integer $j > 1$ we inductively define C_{itj} as portion of the closure of the j 'th ball that is not contained in the closure of the balls that preceded the j 'th ball in the sequence.

This C_{itj} includes the nonempty open subset of the j 'th ball that is not included in the closed union of the j 'th ball's predecessors, which has positive probability under the full-support b_{it} .

Thus $b_{it}(C_{itj}|s_{it}) > 0, \forall s_{it} \in S_{it}, \forall j \geq 1$; and $\{C_{itj}: j \in \{1, 2, \dots\}\}$ form a measurable partition of A_{it} .

For each $j, \forall a_{it} \in C_{itj}, \forall D \in \mathcal{M}(A_{it}), \forall s_{it} \in S_{it}$, let $\varphi_{it}(D|a_{it}, s_{it}) = b_{it}(D \cap C_{itj}|s_{it}) / b_{it}(C_{itj}|s_{it})$.

That is, φ_{it} maps any action a_{it} in C_{itj} to the conditional probability distribution over C_{itj} generated by $b_{it}(\bullet|s_{it})$ given the (positive-probability) event of the chosen action being in C_{itj} .

Furthermore, $C_{itj} \subseteq \mathbb{B}_\delta(a_{it})$ when $a_{it} \in C_{itj}$, because C_{itj} is a subset of a closed $\delta/2$ -radius ball that contains a_{it} , and so C_{itj} cannot contain any points that are farther than δ from a_{it} .

We get $b_i * \varphi_i = b_i$ because, for any partition element C_{itj} , for any measurable $D \subseteq C_{itj}$,

$$(b_i * \varphi_i)_t(D|s_{it}) = \int_{C_{itj}} (b_{it}(D|s_{it}) / b_{it}(C_{itj}|s_{it})) b_{it}(da_{it}|s_{it}) = (b_{it}(D|s_{it}) / b_{it}(C_{itj}|s_{it})) b_{it}(C_{itj}|s_{it}) = b_{it}(D|s_{it}).$$

If c_i is a date- t continuation of b_i , then the φ_i -trembling strategy $c_i * \varphi_i$ is also a date- t continuation of b_i because, for all $r < t$, $c_{ir} = b_{ir}$ implies $(c_i * \varphi_i)_r = (b_i * \varphi_i)_r = b_{ir}$ for all $r < t$.

Possible failure of subgame perfection for full conditional equilibria

Example 2: First, nature chooses $\tilde{a}_{01}=\theta$ uniformly from $A_{01}=[0,1]=\{\theta: 0\leq\theta\leq 1\}$.

Then player 1 observes $s_{12}=\theta$ and chooses $a_{12}\in A_{12}=[0,1]$.

Then player 2 observes $s_{23}=(a_{12},\theta)$ and chooses $a_{23}\in A_{23}=[0,1]$

If $a_{23} = a_{12} = \theta$ then payoffs are $u_1 = u_2 = 1$, otherwise $u_1 = u_2 = 0$.

A subgame perfect equilibrium must yield the outcome $a_{23}=a_{12}=\theta$ with probability 1.

But there is also a full conditional ε -equilibrium (for all $\varepsilon\geq 0$) where both players randomize uniformly over $[0,1]$ independently of each other and θ .

Here, player 2 fails to optimize a_{23} in the event $\{a_{12}=\theta\}$ because it has probability 0, and 2's failure to coordinate then eliminates player 1's incentive to match a_{12} with θ .

Even if player 2 changed to $a_{23}=\theta$ when $a_{12}=\theta$ (a change in an event which has probability 0 in this equilibrium), arbitrarily small local trembles in player 1's choice of a_{12} could eliminate any expected gain for 1 from attempting $a_{12}=\theta$, because u_1 is not continuous there.

(This illustrates the Fact in the above page on full conditional ε -equilibria.)

Why perfect conditional-rationality tests must include perturbations of nature

Example 3: First, nature chooses $\tilde{a}_{01}=(\theta_1,\theta_2)\in\{1,2\}\times[0,1]$, with θ_1 and θ_2 independent, $P(\{\theta_1=1\})=1/4$, and θ_2 uniform on $[0,1]$. Simultaneously, player 1 chooses $a_{11}\in[0,1]$. On date 2, player 2 observes $s_{22}=a_{11}$ if $\theta_1=1$, or $s_{22}=\theta_2$ if $\theta_1=2$, and then player 2 chooses $a_{22}\in\{1,2\}$.

Payoffs (u_1,u_2) are as follows: (a_{11} is payoff irrelevant)

	$a_{22}=1$	$a_{22}=2$
$\theta_1=1$	(1,1)	(0,0)
$\theta_1=2$	(1,0)	(0,1)

If we test "perfectness" of a conditional ε -equilibrium by evaluating conditional rationality with a net of perturbations on strategies alone, without perturbing nature, we get trouble here. If we only perturb players' strategies then, for any $x\in[0,1]$, the event $\{s_{22}=x\}$ could have positive probability only when $\theta_1=1$ and 1's strategy assigns positive probability to $a_{11}=x$, and so then conditional ε -rationality for player 2 would require $\hat{b}_{22}(1|x) > 1-\varepsilon, \forall x\in[0,1]$. But such a strategy would not be ex-ante rational for player 2!

Such perturbations would also verify beliefs $\beta_{22}(\{\theta_1=1\}|\{s_{22}=x\})=1, \forall x\in[0,1]$, even though the event $\{\theta_1=1\}$ has prior probability $1/4 = \beta_{22}(\{\theta_1=1\}|\{s_{22}\in[0,1]\})$. Knowing beliefs after all individual signals might not determine beliefs in larger events!

(This game has many perfect conditional ε -equilibria, as defined in next slide.

In one eqm, player 1 chooses a_{11} uniformly over $[0,1]$ and player 2 always chooses $a_{22}=2$. In another, player 1 chooses $a_{11}=0.5$, and player 2 chooses $a_{22}=1$ if $s_{22}=0.5$, else $a_{22}=2$.)

Perfect conditional ε -equilibria

To evaluate rationality of an equilibrium in 0-probability events, we must perturb it.

An infinite net of perturbations of both players' strategies and nature's behavior may be needed to test all events that have zero probability, to verify *perfectness* of the equilibrium.

For any two strategies b and \hat{b} in B , let $\|\hat{b}-b\| = \sup |\hat{b}_{it}(C|s_{it}) - b_{it}(C|s_{it})|$ where the supremum is taken over all $it \in L$, all $s_{it} \in S_{it}$, and all $C \in \mathcal{M}(A_{it})$.

For any perturbation \hat{p} that is a perturbation of nature's probability function p , we let $\|\hat{p}-p\| = \sup |\hat{p}_t(C|a_{<t}) - p_t(C|a_{<t})|$, where the sup is over all $t \in \{1, \dots, T\}$, $C \in \mathcal{M}(A_{0t})$, $a_{<t} \in A_{<t}$.

Let $\Gamma(\hat{p})$ denote the *perturbed game* where nature's probability function is \hat{p} instead of p .

An event $N \subseteq A$ is *negligible* iff $N \in \mathcal{M}(A)$ and $P(N|b) = 0$, $\forall b \in B$.

A *perfect conditional ε -equilibrium* is any $b \in B$ such that there is a negligible set $N \subseteq A$ and

there is a net of perturbations $\{(\hat{p}^\alpha, \hat{b}^\alpha) : \alpha \in \Xi\}$ indexed on some directed set Ξ , such that:

each \hat{b}^α is a conditional ε -equilibrium in the perturbed game $\Gamma(\hat{p}^\alpha)$,

and for every $\delta > 0$ and every outcome $a \in A \setminus N$ and there exists some $\xi(\delta, a) \in \Xi$ such that

$\|\hat{p}^\alpha - p\| \leq \delta$, $\|\hat{b}^\alpha - b\| \leq \delta$, and $P(\{a\}|\hat{b}^\alpha; \hat{p}^\alpha) > 0$ for all $\alpha \geq \xi(\delta, a)$.

(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 perturbed strategies that can test rationality in all events but may yield only finite additivity in the limit.

Simon-Stinchcombe 1995 and Bajoori-Flesch-Vermuelen 2013&2016 use the topological full-support condition in defining solutions that they call "perfect." But "perfect" in English comes from a Latin word meaning "complete", and so it seems more appropriate for the condition of testing rationality everywhere by a net of perturbations.)

Problems from allowing perturbations of nature

Example 4: Nature chooses $\tilde{a}_{01}=(\theta_1,\theta_2)\in[-1,3]\times[-1,3]$, independent, uniform.

Then player 1 observes $s_{12}=\theta_1$ and chooses $a_{12}\in\{-1,1\}$.

Then player 2 observes $s_{23}=a_{12}$ and chooses $a_{23}\in\{-1,1\}$.

Payoffs are $u_1=a_{12}a_{23}$, $u_2=a_{23}\theta_2$.

So player 2 wants a_{23} to match the sign of θ_2 , and player 1 wants to match a_{23} .

Player 1 has no information about θ_2 , and player 2 thinks $E\theta_2>0$.

In one perfect conditional ε -equilibrium, player 2 chooses $a_{23}=1$ for any s_{23} , and player 1 chooses $a_{12}=1$ regardless of s_{12} . With small ($<\varepsilon/2$) probabilities of each player independently trembling to the other action, we get a conditional ε -equilibrium that is perfect and full.

Now consider a perturbation of nature that puts a small positive probability δ on the event $\theta_1=\theta_2=-1$, otherwise (θ_1,θ_2) are drawn independently from the given uniform distribution.

These perturbations verify a perfect conditional ε -equilibrium in which player 1 chooses $a_{12}=-1$ if $\theta_1>-1$, but 1 chooses $a_{12}=1$ if $\theta_1=-1$, and player 2 always chooses $a_{23}=-a_{12}$.

Here 1 usually avoids $a_{12}=1$ because 2 would take this surprise as evidence of $\theta_1=\theta_2=-1$.

The perturbation of nature here admitted a possibility that player 1's signal θ_1 might convey information about θ_2 , which affects beliefs even though its probability vanishes as $\delta\rightarrow 0$.

However, if we added small-probability trembles to these equilibrium strategies, to yield a full-support strategy profile, then nature's δ -perturbations would not affect 2's limiting beliefs after any $s_{23}\in\{-1,1\}$, and so a full conditional ε -equilibrium cannot be so perverse.

Such perverse equilibria could also be avoided by perturbing θ_1 and θ_2 independently.

Augmenting a game with a specified net of admissible perturbations of nature

When nature generates pairs of independent random variables in our model, we might want to consider only perturbations of nature that preserve this independence.

So we might want to specify a net of admissible perturbations of nature $\{\tilde{p}^\alpha\}$, as an additional part of the structure of the game.

The net should satisfy $\lim_\alpha \|\tilde{p}^\alpha - p\| = 0$ and should include perturbations so that any outcome outside some negligible set can eventually in the net get positive probability with some $\hat{b} \in B$. A perfect conditional ε -equilibrium is *compatible* with a specified net if the above definition can be satisfied with perturbations $\{(\hat{p}^\alpha, \hat{b}^\alpha)\}$ such that $\{\hat{p}^\alpha\}$ is a subnet of the specified net.

We will consider models where, at each period t , nature's domain is a product of coordinates $A_{0t} = \times_{j \in J} A_{0tj}$, and each state-coordinate domain A_{0tj} is a metric space.

We may also specify a measurable partition Q_{0tj} of each A_{0tj} .

With this structure, we may specify a *canonical net of nature-perturbations* as follows:

The index set is the set of pairs (δ, \mathcal{W}) where δ is a positive number and \mathcal{W} is a finite subset of nature's state space $A_0 = \times_t A_{0t}$, partially ordered by smallness of δ and inclusiveness of \mathcal{W} .

Let $\mathcal{W}_{tj} = \{a_{0tj} : a_0 \in \mathcal{W}\}$. Let $Q_{0tj}(a_{0tj})$ denote the element in partition Q_{0jt} that contains a_{0tj} .

Define $\chi^{\delta, \mathcal{W}}_{tj} : A_{0tj} \rightarrow \Delta(A_{0tj})$ so that $\forall a_{0tj}$, if $Q_{0tj}(a_{0tj}) \cap \mathcal{W}_{tj} = \emptyset$ then $\chi^{\delta, \mathcal{W}}_{tj}(\{a_{0tj}\} | a_{0tj}) = 1$,

otherwise $\chi^{\delta, \mathcal{W}}_{tj}(\{a_{0tj}\} | a_{0tj}) = 1 - \delta$ and $\chi^{\delta, \mathcal{W}}_{tj}(\bullet | a_{0tj})$ distributes probability δ uniformly over the finitely many points in $Q_{0tj}(a_{0tj}) \cap \mathcal{W}_{tj}$ that are closest to a_{0tj} (in the given metric on A_{0tj}).

Then the (δ, \mathcal{W}) -perturbation has, for any rectangular set $\times_{j \in J} C_j$ (with $C_j \in \mathcal{M}(A_{0tj})$, $\forall j \in J$),

$$\tilde{p}^{\delta, \mathcal{W}}_t(\times_{j \in J} C_j | a_{<t}) = \int_{A_{0t}} [\prod_{j \in J} \chi^{\delta, \mathcal{W}}_{tj}(C_j | a_{0tj})] p(da_{0t} | a_{<t}).$$

Beliefs and subgame-perfectness, for a perfect conditional ε -equilibrium \mathbf{b} .

Perfectness gives us a net of perturbations indexed on $\delta > 0$ and finite sets $\mathcal{W} \subset Y$, such that, conditional probabilities $P_{<t}(H|Z, \hat{\mathbf{b}}^{\delta, \mathcal{W}}; \hat{\mathbf{p}}^{\delta, \mathcal{W}})$ are defined for all (δ, \mathcal{W}) such that $\mathcal{W} \cap Z \neq \emptyset$.

By Tychonoff's theorem, there is a subnet which yields well-defined *belief probabilities*, $\beta_{it}(H|Z) = \lim_{\delta, \mathcal{W}} P_{<t}(H|Z, \hat{\mathbf{b}}^{\delta, \mathcal{W}}; \hat{\mathbf{p}}^{\delta, \mathcal{W}})$, $\forall it \in L$, $\forall Z \in \mathcal{M}(Y_{it})$, $\forall H \in \mathcal{M}(A_{<t})$.

If $P_{it}(Z|\mathbf{b}; \mathbf{p}) > 0$ then the beliefs $\beta_{it}(\bullet|Z)$ do not depend on the perturbations $(\hat{\mathbf{b}}^{\delta, \mathcal{W}}; \hat{\mathbf{p}}^{\delta, \mathcal{W}}) \rightarrow (\mathbf{b}; \mathbf{p})$.

Facts. The beliefs $\beta_{it}(\bullet|Z)$ are finitely additive probability distributions on $A_{<t}$.

With these beliefs, the perfect conditional ε -equilibrium \mathbf{b} satisfies sequential ε -rationality:

$$\int U_i(c_i, \mathbf{b}_{-i} | a_{<t}) \beta_{it}(da_{<t} | Z) \leq \int U_i(\mathbf{b} | a_{<t}) \beta_{it}(da_{<t} | Z) + \varepsilon, \quad \forall it \in L, \quad \forall c_i \in \bar{B}_{i, \geq t}(\mathbf{b}_i), \quad \forall Z \in \mathcal{M}(Y_{it}).$$

In examples, we can get $\beta_{it}(\bullet|Z)$ not countably additive and $\beta_{it}(H|Z) < \min_{s_{it} \in Z} \beta_{it}(H|\{s_{it}\})$.

With perfect recall, a date- t history $a_{<t}$ in $A_{<t}$ is a *subgame* of Γ iff $\sigma^{-1}(\sigma_{it}(a_{<t})) = \{a_{<t}\} \quad \forall i \in I$.

Given any $\varepsilon > 0$, a strategy profile \mathbf{b} is a *subgame-perfect ε -equilibrium* of Γ iff such that for every $it \in L$ and every subgame $a_{<t}$ outside some negligible set,

if $\sigma_{it}(a_{<t}) \in Y_{it}$, then $U_i(c_i, \mathbf{b}_{-i} | a_{<t}) \leq U_i(\mathbf{b} | a_{<t}) + \varepsilon \quad \forall c_i \in \bar{B}_{i, \geq t}(\mathbf{b}_i)$.

Fact. If \mathbf{b} is a perfect conditional ε -equilibrium then \mathbf{b} is a subgame-perfect ε -equilibrium.

Limits of equilibrium distributions as $\varepsilon \rightarrow 0$

A *[perfect, full] conditional-equilibrium distribution* is any μ such that, for any finite $\mathcal{F} \subset \mathcal{M}(A)$, for any $\varepsilon > 0$, there exists some $b^{\varepsilon, \mathcal{F}}$ such that $b^{\varepsilon, \mathcal{F}}$ is a [perfect, full] conditional ε -equilibrium and $|\mu(H) - P(H|b^{\varepsilon, \mathcal{F}})| < \varepsilon \forall H \in \mathcal{F}$.

A conditional-equilibrium distribution is a finitely additive probability measure on the set of outcomes A . It might be only finitely additive (as when $\mu(\{0 < a_{it} < \delta\}) = 1 \forall \delta > 0$).

Fact. Let Γ^f be any finite extensive-form game, with discrete topology on the finite A_{it} . If all alternatives have positive probability at each chance node in Γ^f (as KW assumed) then the perfect conditional-equilibrium distributions and full conditional-equilibrium distributions both coincide with the distributions over outcomes that can result from sequential equilibria. For finite games where some chance moves have zero probability, two definitions of sequential equilibrium have been used: one considering limits of perturbed chance probabilities that are all positive, the other holding chance probabilities fixed. The perfect conditional-equilibrium distributions coincide with the outcome distributions for sequential equilibria with limits of positive perturbations of chance probabilities. The full conditional-equilibrium distributions coincide with the outcome distributions for sequential equilibria without perturbations of chance. For infinite games where chance (nature) has an uncountable infinity of possible moves, the question of zero-probability chance events cannot be avoided, and so these two definitions which seemed almost the same for finite games become essentially distinct for infinite games.

Regular projective games

Let $\Gamma = (I, T, A, S, \mathcal{M}, \sigma, p, u)$ be a multi-stage game (with perfect recall).

Γ is a *regular projective game* iff there is a finite index set J and sets A_{ikj} such that, $\forall i \in I$:

(R.1) $A_{it} = \times_{j \in J} A_{itj} \quad \forall i \in I^*, \forall t \in \{1, \dots, T\}$;

(R.2) there exist sets $M_{it} \subseteq I^* \times \{1, \dots, t-1\} \times J$ such that $S_{it} = \times_{hrj \in M_{it}} A_{hrj}$ and

$\sigma_{it}(a_{<t}) = (a_{hrj})_{hrj \in M_{it}} \quad \forall a_{<t} \in A_{<t}$, that is, i 's signal at t is a list of state and action coordinates from periods before t ;

(R.3) A_{itj} are nonempty compact metric spaces $\forall j \in J$, and all products of these spaces, including the signal sets S_{it} , have their product topologies and Borel sigma-algebras,

(R.4) $u_i: A \rightarrow \mathbb{R}$ is continuous,

(R.5) for each $j \in J$ there is a probability measure ρ_{tj} on $\mathcal{M}(A_{0tj})$, with full support on A_{0tj} , and there is a continuous density function $f_t: A_{0t} \times A_{<t} \rightarrow \mathbb{R}$ such that

$$p_t(C|a_{<t}) = \int_C f_t(a_{0t}|a_{<t}) \prod_{j \in J} \rho_{tj}(da_{0tj}) \quad \forall C \in \mathcal{M}(A_{0t}), \forall a_{<t} \in A_{<t},$$

and $\{a_{\leq t}: f_t(a_{0t}|a_{<t}) > 0\}$ is a closed subset of $A_{\leq t}$.

Remarks. One can always reduce the cardinality of J to $(T+1)^{\#I}$ or less by grouping ($\forall i \in I^*, \forall t \in \{1, \dots, T\}$) the variables $\{a_{itj}\}_{j \in J}$ by the dates when each player observes them, if ever.

The final condition in (R.5) says that points of zero density are topologically isolated from points of positive density. This is always true for finite games with the discrete topology.

Regular projective games can include all finite multi-stage games simply by letting each player's signal be a coordinate of the state.

We defined canonical nets of nature-perturbations for such games (with partitions Q_{0tj}).

Existence Theorems

Theorem. Given any regular projective game, for any $\varepsilon > 0$, there exists full and perfect conditional ε -equilibrium that is compatible with a canonical net of nature-perturbations, for some measurable partitions Q_{0tj} of the nature-coordinate spaces A_{0tj} .

The proof uses a finite approximating game where each A_{itj} is partitioned into small sets on which the functions $u_i(\theta, a) \prod_{t \in \{1, \dots, T\}} f(a_{0t} | a_{<t})$ have small variation.

Theorem. Any regular projective game has a full and perfect conditional-equilibrium distribution μ .

The proof applies Tychonoff's theorem to the compact product topology on $[0, 1]^{\mathcal{M}(A)}$.

Problems of applying sequential rationality with finitely additive beliefs

Example 5: First, nature chooses $\tilde{a}_{01}=\theta$ uniformly from $A_{01}=(0,1)=\{\theta: 0<\theta<1\}$.

Then player 1 observes $s_{12}=\theta$ and chooses $a_{12}\in A_{12}=\{0,1\}$.

Then player 2 observes $s_{23}=a_{12}$ and chooses $a_{23}\in A_{23}=[0,1)=\{a_{23}: 0\leq a_{23}<1\}$

If $a_{12}=0$ ("1 quits") then payoffs are $u_1 = u_2 = 0$.

If $a_{12}=1$ but $a_{23}=0$ ("2 quits") then payoffs are $u_1 = -1$ and $u_2 = 0$.

If $a_{12}=1$ and $a_{23}>0$ then $u_1 = 1$, and $u_2 = 1$ if $a_{23} \geq \theta$, but $u_2 = -1$ if $a_{23} < \theta$.

A perfect conditional ε -equilibrium has $a_{12}=1$, $a_{23}=1-\varepsilon/2$. (No optimum for 2 in $A_{23}=[0,1)$.)

Given $a_{12}=1$, for any countably additive conditional belief on A_{01} , there would exist some feasible $a_{23}<1$ worth choosing (having probability greater than $1-\varepsilon$ of satisfying $a_{23}>\theta$).

However, consider a strategy profile b such that player 1 chooses $a_{12}=0$ with probability 1, but if player 1 chose $a_{12}=1$ then player 2 would choose $a_{23}=0$ with probability 1.

Consider beliefs for 2 when $a_{12}=1$ generated by a net of perturbations of b indexed on $\delta>0$ in which player 1 would choose $a_{12}=1$ with conditional probability ε only when $\theta>1-\delta$.

The limiting beliefs are finitely additive, with $\beta_{23}(\{a_{23}<\theta<1\}|\{a_{12}=1\}) = 1 \forall a_{23}<1$.

So with these beliefs, player 2 would strictly prefer the quit option $a_{23}=0$ over any $a_{23}>0$.

But these perturbations are not conditional ε -equilibria because, for any $\delta>0$ perturbation in the net, player 2 would strictly prefer to choose $a_{23}=1-\delta\varepsilon/2$.

An example with asymptotic strategic entanglement everywhere

Example 6 (Hellman 2014): $I=\{1,2\}$, $T=2$. $a_{01} = (\theta_0, \theta_1, \theta_2) \in A_{01} = \{1,2\} \times [0,1] \times [0,1]$.

θ_0 is equally likely to be 1 or 2; it names the player who is "on". Signals are $s_1=s_{12}=\theta_1$, $s_2=s_{22}=\theta_2$.

When $\theta_0=i$, s_i is Uniform $[0,1]$, other $-i$ has signal $s_{-i} = 2s_i$ if $s_i < 0.5$, $s_{-i} = 2s_i - 1$ if $s_i \geq 0.5$.

(This implies θ_{-i} is also Uniform $[0,1]$ when $\theta_0=i$.) Action sets are $A_{12} = A_1 = \{L,R\} = A_{22} = A_2$.

When $\theta_0=i$, the other player $-i$ just gets $u_{-i} = 0$, and u_i is determined by:

when $s_i \geq 0.5$ then i gets $u_i = 0.7$ if $a_i=L=a_{-i}$, $u_i = 0.3$ if $a_i=R=a_{-i}$, and $u_i=0$ if $a_i \neq a_{-i}$;

when $s_i < 0.5$ then i gets $u_i = 0.7$ if $a_i=L \neq a_{-i}$, $u_i = 0.3$ if $a_i=R \neq a_{-i}$, and $u_i=0$ if $a_{i2}=a_{-i2}$.

So $s_i \geq 0.5$ wants to match $-i$ when i is "on" and prefers L if $-i$'s probability of R is < 0.7 ;

$s_i < 0.5$ wants to mismatch $-i$ when i is "on" and prefers L if $-i$'s probability of R is > 0.3 .

This game has no Bayesian equilibrium in which the strategic functions $b_i(R|s_i)$ are measurable functions of $s_i \in [0,1]$, by arguments of Simon (2003) and Hellman (2014).

But for any $\varepsilon > 0$, we can construct strategy profiles that are conditionally ε -rational on an ε -sure set.

Pick an integer $m \geq 1$ such that $P(s_1 < 2^{-m}) < \varepsilon^2$.

In the binary expansion of any s_1 , find the first string of m consecutive 0's starting at an odd position, and if the number of prior 0's is odd then let $b_1(R|s_1)=1$, else $b_1(R|s_1)=0$.

In the binary expansion of any s_2 , find the first string of m consecutive 0's starting at an even position, and if the number of prior 0's is odd then let $b_2(R|s_2)=1$, else $b_2(R|s_2)=0$.

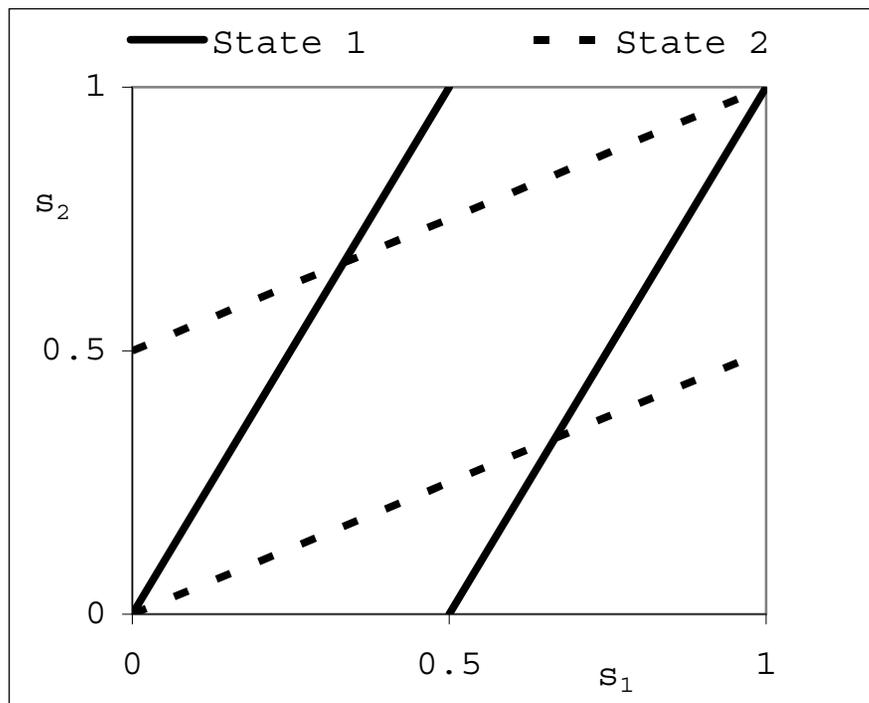
In this specification, only signals $s_1 < 2^{-m}$ are not best responding (to $s_{-i}=2s_i$ or $s_{-i}=2s_i-1$).

The signals not covered here have probability 0, for them we can let $b_i(R|s_i)=0$.

Beliefs and strategies are entangled in limit: on any small open interval, 2 believes that, when $\theta_0=2$, the unobserved a_1 is equally likely to be L or R, but $a_2 = a_1$ when $s_2 > 0.5$, $a_2 \neq a_1$ when $s_2 < 0.5$.

The players' actions are correlated by infinitesimal details of their signal information.

Example 6 (Hellman 2014) picture



$u_{-i}=0$ in state $\theta_0=i$.

u_i in state $\theta_0=i$ if $s_i \geq 0.5$:

	$a_{-i}=L$	$a_{-i}=R$
$a_i=L$	0.7	0
$a_i=R$	0	0.3

$(s_i = (1+s_{-i})/2)$

u_i in state $\theta_0=i$ if $s_i < 0.5$:

	$a_{-i}=L$	$a_{-i}=R$
$a_i=L$	0	0.7
$a_i=R$	0.3	0

$(s_i = s_{-i}/2)$

Suppose 1 does L when $s_1 < 2^{-3}$, that is $s_1 = 0.000\dots$ in binary expansion (where " \dots " can be any string of 0s & 1s).

Then rational responses for all other signal-types would imply:

2 does R when $s_2 = 0.0000\dots$, 2 does L when $s_2 = 0.1000\dots$;

1 does R when $s_1 = 0.10000\dots$,

1 does R when $s_1 = 0.01000\dots$, 1 does L when $s_1 = 0.11000\dots$;

2 does L when $s_2 = 0.010000\dots$, 2 does R when $s_2 = 0.110000\dots$,

2 does L when $s_2 = 0.001000\dots$, 2 does R when $s_2 = 0.101000\dots$,

2 does R when $s_2 = 0.011000\dots$, 2 does L when $s_2 = 0.111000\dots$;

In the binary expansion of s_1 , find the first string of 3 consecutive 0's starting at an odd position, and if the number of prior 0's is odd then 1 does R at s_1 , else 1 does L at s_1 .

In the binary expansion of s_2 , find the first string of 3 consecutive 0's starting at an even position,

and if the number of prior 0's is odd then 2 does R at s_2 , else 2 does L at s_2 .

References

- Robert J. Aumann, "Mixed and behavior strategies in infinite extensive games" (1964).
- Elnaz Bajoori, János Flesch, Dries Vermeulen, "Perfect equilibrium in games with compact action spaces," *Games and Economic Behavior* 82:490-502 (2013).
- Elnaz Bajoori, János Flesch, Dries Vermeulen, "Behavioral perfect equilibrium in Bayesian games," *Games and Economic Behavior* 98:78-109 (2016).
- Drew Fudenberg and Jean Tirole (1991): "Perfect Bayesian and Sequential Equilibrium," *Journal of Economic Theory* 53:236-260.
- Christopher J. Harris, Maxwell B. Stinchcombe, and William R. Zame (2000): "The Finitistic Theory of Infinite Games," UTexas.edu working paper. <http://www.laits.utexas.edu/~maxwell/finsee4.pdf>
- Christopher J. Harris, Philip J. Reny, and Arthur J. Robson (1995): "The existence of subgame-perfect equilibrium in continuous games with almost perfect information," *Econometrica* 63(3):507-544.
- Ziv Hellman (2014): "A game with no Bayesian approximate equilibria," *Journal of Economic Theory* 153:138-151.
- Ziv Hellman and Yehuda Levy, "Bayesian games with a continuum of states," Hebrew University working paper (2013).
- John L. Kelley, *General Topology* (Springer-Verlag, 1955).
- David Kreps and Robert Wilson (1982): "Sequential Equilibria," *Econometrica* 50:863-894.
- Paul Milgrom and Robert Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research* 10:619-32.
- Leo K. Simon and Maxwell B. Stinchcombe, "Equilibrium refinement for infinite normal-form games," *Econometrica* 63(6):1421-1443 (1995).
- Robert Samuel Simon (2003): "Games of incomplete information, ergodic theory, and the measurability of equilibria," *Israel Journal of Mathematics* 138:73-92.
- Robert Samuel Simon and Grzegorz Tomkowicz (2017), "A Bayesian game without ε -equilibria," *Israel Journal of Mathematics* (to appear).
- Joel Watson (2017), "A general practicable definition of perfect Bayesian equilibrium," UCSD working paper.