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

by

Roger B. Myerson and Philip J. Reny*

Draft notes March 2014

https://sites.google.com/site/philipjreny/

Abstract: Guided by several key examples, we formulate a definition of *sequential equilibrium* for multi-stage games with infinite type sets and infinite action sets, and we prove its existence for a class of games.

* Department of Economics, University of Chicago

Goal: formulate a definition of *sequential equilibrium* for multi-stage games with infinite type sets and infinite action sets, and prove general existence.
Sequential equilibria were defined for finite games by Kreps-Wilson 1982,
but rigorously defined extensions to infinite games have been lacking.
Various formulations of "perfect bayesian eqm" (defined for *finite* games in Fudenberg-Tirole 1991) have been used for infinite games. No general existence.
Harris-Stinchcombe-Zame 2000 explored definitions with nonstandard analysis.

It is natural to try to define sequential equilibria of an infinite game can by taking limits of sequential equilibria of finite games that "approximate" it.

But no general definition of "good finite approximation" has been found.
It is easy to define sequences of finite games that seem to be converging to the infinite game (in some sense) but have limits of equilibria that seem wrong.

We therefore take another route. We work with the infinite game itself, and consider for any epsilon and any finite partition of the players' type (information) spaces, strategies that are epsilon-optimal only at each of the finitely many partition elements. We then consider the limit as epsilon tends to zero and as the type-space partitions become arbitrarily fine, i.e., finer than *any* particular finite partition.

The definition presented here is based on intuitive judgments about what is reasonable. Others are encouraged to explore alternative definitions.

We begin with several examples. The first motivates the new formalism we use to describe the solution of a game. The others motivate why we do not define a solution as the limit of sequential equilibria of finite games that approximate the given infinite game.

Nonexistence and strategic entanglement (Harris-Reny-Robson 1995)
Example: Date 1: Player 1 chooses a_1 from [-1,1], player 2 chooses from {L,R}.
Date 2: Players 3 and 4 observe the date 1 choices and each choose from {L,R}.

For i=3,4, player i's payoff is $-a_1$ if i chooses L and a_1 if i chooses R.

Player 2's payoff depends on whether she matches 3's choice.
If 2 chooses L then she gets 1 if player 3 chooses L but -1 if 3 chooses R; and
If 2 chooses R then she gets 2 if player 3 chooses R but -2 if 3 chooses L.

Player 1's payoff is the sum of three terms:
(First term) If 2 and 3 match he gets $-|a_1|$, if they mismatch he gets $|a_1|$;
plus (second term) if 3 and 4 match he gets 0, if they mismatch he gets -10;
plus (third term) he gets $-|a_1|^2$.

There is no subgame perfect equilibrium of this game. But its "solution" seems obvious.

Approximations in which 3 and 4 can distinguish between $a_1 = +,0,-$ and in which 1's action set is $\{-1,\ldots,-2/m,-1/m,1/m,2/m,\ldots,1\}$ have a unique subgame perfect (hence sequential) equilibrium in which player 1 chooses $\pm 1/m$ with probability ½ each, player 2 chooses L and R each with probability ½ , and players i=3,4 both choose L if $a_1=-1/m$ and both choose R if $a_1=1/m$. The limit "solution" is $a_1= 0$, $a_2 =$ L or R each with prob. ½, and $(a_3,a_4) =$ (L,L) or (R,R) each with prob. ½ . (But, for general games, finite approximations are unreliable.)

Player 3's and player 4's strategies are entangled in the limit.

The solution of this game (and generally) cannot be described by independent strategies, so we introduce a new formalism to describe game solutions.

Problems of spurious signaling in naïve finite approximations

Example. Nature chooses $\theta \in \{1,2\}$, $p(\theta) = \theta/3$. Player 1 observes $t_1 = \emptyset$ and chooses $a_1 \in [0,1]$. Player 2 observes $t_2 = (a_1)^\theta$ and chooses $a_2 \in \{1,2\}$. Payoffs are as follows.

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

It should *not* be possible for player 2 to always match the state because, if it were possible, then

- for any $a_1 \in [0,1]$ that might be chosen by player 1, $(a_1)^{1/2}$ must never be chosen, implying that choosing $(a_1)^{1/2}$ instead of a_1 would be a profitable deviation for player 1.

But if player 1 is restricted to *any finite subset*, *F*, of his action space [0,1], then player 2 *can* always match the state as follows:

- player 1 chooses the largest action in *F* that is less than 1;
- player 2 chooses $a_2 = 1$ iff $(t_2)^{1/2} \notin F$.

Thus, we must not simply discretize the players' action spaces to obtain a finite game.

To obtain a finite game while avoiding such spurious signaling, we could limit player 2's ability to distinguish her types and allow player 1 any finite subset of his actions. That is, we could coarsen the players' information by finitely partitioning each of their type spaces, T_{ik}, and then pass to the limit by expanding the players' finite action sets to fill in the action spaces A_{ik} faster than the type-space partitions become arbitrarily fine.

But this too leads to difficulties, as the next example shows.

Limitations of step-strategy approximations

Example. Player 1 chooses $a_1 \in [0,1]$. Player 2 observes $t_2 = a_1$ and chooses $a_2 \in [0,1]$. The game is zero-sum. Player 1 receives 1 if their choices do not match and -1 otherwise. (This discontinuous game is the reduced form of a two-stage continuous game.)

There should be no equilibrium in which player 2 fails to match player 1's choice.

However, for any fixed partition of player 2's type space $T_2 = [0,1]$, player 1 can mix uniformly over a large enough finite set of actions within a single element of 2's partition so that player 2 has an arbitrarily small chance of matching 1's choice.

Player 2 would like to use the strategy "choose $a_2 = t_2$," but this strategy can be approximated only by step functions when player 2 has finitely many feasible actions.
Step functions close to this strategy yield very different expected payoffs because 2's utility function is discontinuous (but such discontinuities can arise in the reduced form of a continuous game with one additional stage).

If we gave player 2 the *strategy* $s_2(t_1) = t_1$, she would use it!

This example suggests that, because strategies can be particularly important for players, we should include them (even if not measurable w.r.t. the players' finite partitions), and not merely actions, in our finite approximations.

But, as the next example shows, this brings us back again to the problem of spurious signaling.

Spurious signaling returns

Example. Nature chooses $\theta \in \{1,2\}$, $p(\theta) = \theta/3$. Player 1 observes $t_1 = \varnothing$ and chooses $a_1 \in [0,1]$. Player 2 observes $t_2 = (a_1)^\theta$ and chooses $a_2 \in \{1,2\}$. Payoffs are as follows.

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

As before, it should *not* be possible for player 2 to always match the state.

But fix any finite partition of 2's types, and consider filling in the players' pure strategy sets. For any finite set F of pure strategies (actions) for player 1, we can always add one more pure strategy $a_1 = x$ such that $(x)^{1/2}$ is not in F. And we can add to player 2's finite set of pure strategies the pure strategy s_2 such that $s_2(t_2) = 1$ iff $(t_2)^{1/2} \notin F \cup \{x\}$.

Then, player 2 always matches the state if player 1 chooses x and player 2 chooses s_2 .

Taken together, the examples indicate that attempting to approximate the true game with a finite game is unreliable because it can produce unwanted equilibria.

(Harris-Stinchcombe-Zame 2000 provide many excellent additional examples of this kind.)

We therefore always permit the players to employ ALL of their infinitely many strategies.

Multi-stage games $\Gamma=(\Theta,N,K,A,T,p,\tau,u)$

$i \in N$ = {*players*}, finite set;

$k \in \{1,...,K\}$ *periods* of the game.

Let $L = \{(i,k)|\ i\in N,\ k\in\{1,...,K\}\}$ = {*dated players*}. We write ik for (i,k).

A_{ik} = {possible *actions* for player i at date k}; history independent. $A_k = \times_{i\in N} A_{ik}$.

T_{ik} = {possible informational *types* for player i at date k}, disjoint sets. $T_k = \times_{i\in N} T_{ik}$.

Θ_k = {possible *date k states*}.

σ-algebras (closed under cntbl $\cap$ and complements) of measurable subsets are specified for Θ_k, A_{ik} and T_{ik}. Product spaces are given their product σ-algebras.

$A = \times_{k\leq K}\times_{i\in N}A_{ik}$ = {possible *sequences of actions* in the whole game}.

$T = \times_{k\leq K}\times_{i\in N}T_{ik}$ = {possible *sequences of types* in the whole game}.

$\Theta = \times_{k\leq K}\Theta_k$ = {possible *states* in the whole game}.

$\Theta\times A$ = {possible *outcomes* of the game}.

The subscript, $<k$, denotes the projection onto periods before k, and $\leq k$ weakly before.

e.g., $A_{<k} = \times_{h<k}\times_{i\in N}A_{ih}$ = {possible action sequences before period k} ($A_{<1} = \{\emptyset\}$),

and for $a\in A$, $a_{<k}=\times_{h<k}\times_{i\in N}a_{ih}$ is the partial sequence of actions before period k.

The date k state is determined by a regular conditional probability p_k from $\Theta_{<k}\times A_{<k}$ to Θ_k.

i.e., for each $(\theta_{<k},a_{<k})$, $p_k(.|\theta_{<k},a_{<k})$ is a cntbly additive prob. on the msbl subsets of Θ_k, and for each msbl subset C of Θ_k, $p_k(C|\theta_{<k},a_{<k})$ is a msbl function of $(\theta_{<k},a_{<k})$.

Player i's period k information is given by a measurable *type function* $\tau_{ik}:\Theta_{\leq k} \times A_{<k} \rightarrow T_{ik}$.

Assume *perfect recall*: $\forall ik\in L$, $\forall m < k$, there is a measurable $\rho_{ikm}:T_{ik}\rightarrow T_{im}\times A_{im}$ such that $\rho_{ikm}(\tau_{ik}(\theta_{\leq k},a_{<k})) = (\tau_{im}(\theta_{\leq m},a_{<m}),a_{im})$, $\forall\theta\in\Theta$, $\forall a\in A$.

Each player i has a measurable and bounded *utility function* u_i: $\Theta\times A \rightarrow\Re$.

Strategies and induced distributions

A *strategy*, s_{ik}, for $ik \in L$ is any regular conditional probability from T_{ik} to $\Delta(A_{ik})$.

i.e., for each $t_{ik} \in T_{ik}$, $s_{ik}(\cdot|t_{ik})$ is a countably additive probability on the measurable subsets of A_{ik}, and for each measurable subset C of A_{ik}, $s_{ik}(C|t_{ik})$ is a measurable function of t_{ik}.

Let S_{ik} denote ik's set of strategies, let $S_i = \times_k S_{ik}$, denote i's set of strategies, and let $S_k = \times_{i \in N} S_{ik}$ and $S = \times_{ik \in L} S_{ik}$.

For any $s_k \in S_k$ and any $t_k = (t_{ik})_{i \in N} \in T_k$, let $s_k(\cdot|t_k)$ denote the product of the measures $s_{ik}(\cdot|t_{ik})$, for $i \in N$.

Each $s_k \in S_k$ determines a regular conditional probability h_k from $\Theta_{<k} \times A_{<k}$ to $\Theta_k \times A_k$ as follows. For any measurable subset Z of $\Theta_k \times A_k$, and any $(\theta_{<k}, a_{<k}) \in \Theta_{<k} \times A_{<k}$,

$$h_k(Z|\theta_{<k}, a_{<k}, s_k) = \int_{\Theta\_k} s_k(Z_k(\theta_k)|\tau_k(\theta_{\leq k}, a_{<k}))p_k(d\theta_k|\theta_{<k}, a_{<k}),$$

where $Z_k(\theta_k) = \{a_k : (\theta_k, a_k) \in Z\}$, and $\tau_k(\theta_{\leq k}, a_{<k}) = (\tau_{ik}(\theta_{\leq k}, a_{<k}))_{i \in N}$.

Set $H_1 = h_1$, and define the probability measures $H_2, \ldots, H_K$ inductively as follows. For each k and any measurable subset W of $\Theta_{\leq k} \times A_{\leq k}$,

$$H_k(W|s) = \int_{\Theta\_<k \times A\_<k} h_k(W_k(\theta_{<k}, a_{<k})|\theta_{<k}, a_{<k}, s_k)H_{k-1}(d(\theta_{<k}, a_{<k})|s),$$

where $W_k(\theta_{<k}, a_{<k}) = \{(\theta_k, a_k) : (\theta_k, \theta_{<k}, a_k, a_{<k}) \in W\}$.

$H_K(.|s)$ is the probability measure over the outcome set $\Theta \times A$ induced by the strategy s.

Conditional probabilities, payoffs, and observable events

Consider any $s \in S$ and any $ik \in L$.

For any measurable $R_{ik} \subseteq T_{ik}$, let $\Pr(R_{ik}|s) = H_K(\{(\theta,a): \tau_{ik}(\theta_{\leq k},a_{<k}) \in R_{ik}\}|s)$.

Define the conditional probability P on $\Theta \times A$ as follows. For any measurable $Y \subseteq \Theta \times A$, if $\Pr(R_{ik}|s) > 0$, let

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

If $\Pr(R_{ik}|s) > 0$, define player i's conditional expected payoff by

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

The set of *observable events for i at k* that can have positive probability is
$\mathcal{Q}_{ik} = \{Q_{ik} \subseteq T_{ik} | \ Q_{ik} \text{ is measurable and } \exists s \in S \text{ such that } \Pr(Q_{ik}|s) > 0\}$.

Let $\mathcal{Q} = \cup_{ik \in L}\, \mathcal{Q}_{ik}$ (a disjoint union) denote the set of *observable events*, i.e., the set of all events that can be observed with positive probability by some dated player.

Type-set partitions and $(\varepsilon,\prod)$-sequential equilibria

A *type-set partition* is any $\prod = \times_{ik\in L}\prod_{ik}$ such that each $\prod_{ik}$ is a finite partition of measurable subsets of T_{ik}. (So elements of $\prod_{ik}$ are disjoint measurable sets with union T_{ik}.)

Type-set partitions are partially ordered by their fineness. Say that $\prod=\times_{ik\in L}\prod_{ik}$ is *finer* than $\prod^o=\times_{ik\in L}\prod^o_{ik}$ if, $\forall ik\in L$, each element of $\prod_{ik}$ is a subset of some element of $\prod^o_{ik}$.

Say that $c_i\in S_i$ is a *date-k continuation* of $s_i\in S_i$ if $c_{ij} = s_{ij}$ for all dates $j < k$.

For any $\varepsilon > 0$ and any type-set partition $\prod$, say that $s\in S$ is an *$(\varepsilon,\prod)$-sequential equilibrium* of Γ if $\forall ik\in L$, $\forall$observable $\pi_{ik}\in\prod_{ik}\cap\mathcal{Q}_{ik}$,

1. $\Pr(\pi_{ik}|s) > 0$, and
2. $U_i(c_i,s_{-i}|\pi_{ik}) \leq U_i(s|\pi_{ik}) + \varepsilon$, for every date-k continuation c_i of s_i.

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

Fact. For any finite game with perfect recall, if $\prod$ is the finest possible type-set partition (the partition given by the players' information sets), then a strategy profile is part of a sequential equilibrium iff it is the limit as $\varepsilon\to 0$ of a sequence of $(\varepsilon,\prod)$-sequential equilibria.

Outcome events and assessments

Recall: $\mathcal{Q} = \cup_{ik\in L}\, \mathcal{Q}_{ik}$ (a disjoint union) is the set of observable events, i.e., the set of all events that can be observed with positive probability by some dated player.

Let $\mathcal{Y}$ = {measurable subsets Y of $\Theta\times A$} be the set of all *outcome events*.

For any dated player ik and any observable event $Q_{ik}\in\mathcal{Q}_{ik}$, let
$I(Q_{ik}) = \{(\theta,a)\in\Theta\times A \mid \tau_{ik}(\theta,a_{<k})\in Q_{ik}\}$.

An *assessment* for Γ is a vector μ of conditional probabilities $\mu(Y|Q)\in[0,1]$ $\forall Y\in\mathcal{Y}$, $\forall Q\in\mathcal{Q}$, such that for any outcome events Y and Z and any observable events Q_{ik} and Q_{jm}:

1. $\mu(Y|Q_{ik})\in[0,1]$, $\mu(\Theta\times A|Q_{ik}) = 1$, $\mu(\varnothing|Q_{ik}) = 0$ (*probabilities*);
2. if $Y\cap Z=\varnothing$ then $\mu(Y\cup Z|Q_{ik}) = \mu(Y|Q_{ik}) + \mu(Z|Q_{ik})$ (*finite additivity*);
3. $\mu(Y|Q_{ik}) = \mu(Y\cap I(Q_{ik})|Q_{ik})$ (*conditional support*);
4. $\mu(Y\cap I(Q_{jm})|Q_{ik})\, \mu(I(Q_{ik})|Q_{jm}) = \mu(Y\cap I(Q_{ik})|Q_{jm})\, \mu(I(Q_{jm})|Q_{ik})$ (*Bayes consistency*).

So {assessments μ} is a compact (product topology) subset of $[0,1]^{\mathcal{Y}\times\mathcal{Q}}$.

Note. Bayes consistency implies that $\mu(Y|T_{ik}) = \mu(Y|T_{jm})$, for all ik and jm in L.
So the unconditional distribution on outcomes $\Theta\times A$ can be defined by $\mu(Y) = \mu(Y|T_{ik})$, $\forall$measurable $Y\subseteq\Theta\times A$, $\forall ik\in L$.

Sequential equilibrium

Recall: $\mathcal{Y}$ = {measurable $Y \subseteq \Theta\times A$}, and $\mathcal{Q} = \cup_{ik\in L}\mathcal{Q}_{ik}$ = {observable events}.

An assessment $\mu\in[0,1]^{\mathcal{Y}\times\mathcal{Q}}$ is a *sequential equilibrium* if for every $\varepsilon > 0$, for every type-space partition Π, and for every finite subset Φ of $\mathcal{Y}\times\mathcal{Q}$, there is an (ε,Π)-sequential equilibrium $\sigma[\varepsilon,\Pi,\Phi]$ of Γ, such that $P(Y|Q,\sigma[\varepsilon,\Pi,\Phi])$ is well-defined and within ε of $\mu(Y|Q)$, $\forall(Y,Q)\in\Phi$.

Remark. $\{(\varepsilon,\Pi,\Phi)\}$ is a directed set when $(\varepsilon',\Pi',\Phi')$ is considered larger than (ε,Π,Φ) iff $\varepsilon'\leq\varepsilon$, Π' is finer than Π, and Φ' contains Φ. Then, an assessment $\mu\in[0,1]^{\mathcal{Y}\times\mathcal{Q}}$ is a sequential equilibrium iff there is a mapping $\sigma[\cdot]$ such that the net of assessments $\{P(\cdot|\cdot,\sigma[\varepsilon,\Pi,\Phi])\}$ converges to μ, i.e., such that

$$\lim_{(\varepsilon,\Pi,\Phi)} P(Y|Q,\sigma[\varepsilon,\Pi,\Phi]) = \mu(Y|Q),\ \forall(Y,Q)\in\mathcal{Y}\times\mathcal{Q},$$

where each $\sigma[\varepsilon,\Pi,\Phi]$ is some (ε,Π)-sequential equilibrium of Γ.

Note. For any $ik\in L$ and any $Q\in\mathcal{Q}_{ik}$, $\Pr(Q|\sigma[\varepsilon,\Pi,\Phi]) > 0$ for all fine enough Π (e.g., whenever Q is the union of elements of Π_{ik}). So $P(Y|Q,\sigma[\varepsilon,\Pi,\Phi])$ is well-defined for all fine enough Π.

Regular multi-stage games with projected types

Let $\Gamma=(\Theta,N,K,A,T,p,\tau,u)$ be a multi-stage game (hence with perfect recall).

Γ is a *regular game with projected types* if there is a finite set J such that $\forall ik\in L$, $\forall j\in J$, $\exists A_{ikj}$ and $\exists\Theta_{kj}$ such that: $\forall ik\in L$,

(i) $A_{ik} = \times_{j\in J}A_{ikj}$, $\Theta_k = \times_{j\in J}\Theta_{kj}$,

(ii) Θ_{kj} is a nonempty complete separable metric space and A_{ikj} is a nonempty compact metric space $\forall j\in J$, and all spaces, including products, are given their Borel sigma-algebras,

(iii) $\exists$a measurable $f_k:\Theta_{\leq k} \times A_{<k} \rightarrow [0,\infty)$, and $\forall j\in J$, $\exists$a finite countably-additive measure ρ_{kj} on the Borel subsets of Θ_{kj} such that $\forall(\theta_{<k},a_{<k})\in\Theta_{<k}\times A_{<k}$ and $\forall$Borel subsets $C\subseteq\Theta_k$, $p_k(C|\theta_{<k},a_{<k}) = \int_C f_k(\theta_k|\theta_{<k},a_{<k})\rho_k(d\theta_k)$, where $\rho_k = \times_{j\in J}\rho_{kj}$ is a product measure on $\Theta_k=\times_{j\in J}\Theta_{kj}$,

(iv) $\{u_i(\theta,\cdot)f_1(\theta_1|\theta_{<1},\cdot)\ldots f_K(\theta_K|\theta_{<K},\cdot)\}_{\theta\in\Theta}$ is an equicontinuous family,

(v) $\tau_{ik}:(\times_{h\leq k}\times_{j\in J}\Theta_{hj})\times(\times_{n\in N}\times_{h<k}\times_{j\in J}A_{nhj})\rightarrow T_{ik}$ is a projection onto $(\times_{hj\in H}\Theta_{hj})\times(\times_{nhj\in M}A_{nhj})$ for some $H\subseteq\{(h,j): h\leq k, j\in J\}$ and some $M\subseteq\{(n,h,j): n\in N, h<k, j\in J\}$, where H and M may depend on ik, and where the projection onto the empty set is a constant function.

Note. A family $\mathcal{G}$ of real-valued functions on a metric space X is *equicontinuous* if $\forall x\in X$, $\forall\varepsilon>0$, $\exists\delta>0$ such that if $x'\in X$ is within δ of x, then $|g(x)-g(x')| < \varepsilon$ for all $g\in\mathcal{G}$.

Remark. (a) The equicontinuity condition (iv) holds if all sets are compact and the function within curly brackets is jointly continuous in (θ,a). (b) One can always reduce the number of coordinates |J| of, say θ_k, to $(K+1)^{|N|}$ or less by grouping them according to the |N|-vector of dates at which each player observes them, if ever.

Existence

Theorem. The set of sequential equilibria is nonempty for all regular games with projected types and is equivalent to the set of Kreps-Wilson sequential equilibria in all finite games.

Remark. (a) Regular games with projected types include all finite games and allow games with perfect information, multi-stage games with observable actions, signaling games, …
(b) Since distinct players can observe the same θ_{kj}, regular games with projected types need not satisfy the Milgrom-Weber (1985) absolute continuity condition.

References

David Kreps and Robert Wilson, "Sequential Equilibria," 50:863-894 (1982).
Drew Fudenberg and Jean Tirole, "Perfect Bayesian and Sequential Equilibrium,", *Journal of Economic Theory* 53:236-260 (1991).
Christopher J. Harris, Maxwell B. Stinchcombe, and William R. Zame, "The Finitistic Theory of Infinite Games," UTexas.edu working paper. http://www.laits.utexas.edu/~maxwell/finsee4.pdf
C. Harris, P. Reny, A. Robson, "The existence of subgame-perfect equilibrium in continuous games with almost perfect information," Econometrica 63(3):507-544 (1995).
J. L. Kelley, General Topology (Springer-Verlag, 1955).
Milgrom, P. and R. Weber (1985): "Distributional Strategies for Games with Incomplete Information," Mathematics of Operations Research, 10, 619-32.